Verified propagation of imprecise probabilities in non-linear ODEs

Ander Gray[a,b], Marcelo Forets[c], Christian Schilling[d], Scott Ferson[b], Luis Benet[e]

[a] *United Kingdom Atomic Energy Authority, Culham, UK*
[b] *Institute for Risk and Uncertainty, University of Liverpool, UK*
[c] *Departamento de Matemática y Aplicaciones, Centro Universitario Regional del Este, Universidad de la República, Maldonado, Uruguay*
[d] *Department of Computer Science, Aalborg University, Aalborg, Denmark*
[e] *Instituto de Ciencias Físicas, Universidad Nacional Autónoma de México (UNAM), México*

Abstract

We combine reachability analysis and probability bounds analysis, which allow for imprecisely known random variables (multivariate intervals or p-boxes) to be specified as the initial states of a dynamical system. In combination, the methods allow for the temporal evolution of p-boxes to be rigorously computed, and they give interval probabilities for *formal verification* problems, also called *failure probability* calculations in reliability analysis. The methodology places no constraints on the input probability distribution or p-box and can handle dependencies generally in the form of copulas. We also provide a consonant approximation method for multivariate p-boxes, which allows for the prediction sets of dynamical systems to be efficiently computed. The presented methodology is rigorous and automatically verified, as both the dynamics and uncertainties are represented and solved with guaranteed enclosures.

Keywords: Reachability analysis, dynamical systems, automatically verified, imprecise probabilities, p-boxes

1. Introduction

We often face the situation of needing to compute the evolution of a system under known dynamics, or at least a very good approximation of them. Knowing the initial conditions allows to determine the evolution under very general assumptions such as continuity and smoothness of the functions dictating the dynamics. Yet, the actual initial state of the system may often be poorly known, with uncertainty attached to it. This is the case, for instance, for a near-Earth object whose position and velocity is measured by (many) amateur astronomers in different parts of the Earth: atmospheric conditions, weather, telescope quality, and other factors enter the measurement process, and are quantified with an uncertainty [55]. Initial conditions of the asteroid with its associated uncertainties are then obtained as the best fit of observations to a dynamical model using standard methods, which may include unknown parameters. Another example is the problem of space debris [18, 42] – vast numbers of small (centimeter-sized) objects located in the low Earth orbit, moving at high speed – which pose collision risks with other objects. In this case, the size and quantity of these objects lead to rather large uncertainties in terms of potential impact risk with satellites.

Prescribing a unique and precise initial probability distribution requires a large amount of high-quality data, or specialist domain knowledge. In situations where information is scarce, such as the examples mentioned above, the usual practice is to assume (i.e., choose) a particular distributional model (usually normal) and dependence (usually independence). However, making these kinds of unwarranted probabilistic assumptions will lead to a serious underestimation (or at least an inaccurate estimation) of uncertainties, and subsequently risks. Imprecise probability relaxes such assumptions and generalizes probability measures from precise values to set values, bounding the contributions to risks from all individual distributions in the set. This allows robust statistical calculations to be made even when data is unreliable, conflicting, imprecise, or limited.

Balch et al. [6] and Martin [53] further argue that using any additive measure of uncertainty (e.g., precise probability) for inference will lead to *false confidence*. They introduce a false confidence theorem, stating that inferences based on additive measures will consistently assign high probabilities to incorrect propositions. Balch et al. [6] expose this in a very practical engineering situation: computing the collision risk of satellites under uncertain measurements. They show a counter-intuitive phenomenon whereby the computed risk somehow *reduces* as the input data uncertainty *increases*, leaving the analyst with a paradoxical situation where collecting less data makes their system safer. Martin [53] offers a remedy based on non-additive measures (e.g., imprecise probabilities), and along with a validity criteria required for uncertainty models to be devoid of false confidence.

For the above reasons, we believe that advancing the capabilities of dealing with imprecise probabilities in dynamical systems is worthwhile.

1.1. Problem statement

In this work, we describe a verified numerical method for propagating uncertainties through a system of (nonlinear) ordinary differential equations (ODEs) of the form

$$x'(t) = f(x(t), t, p), \qquad x(0) = x_0 \in \mathbb{R}^n, \tag{1}$$

where $x(t) \in \mathbb{R}^n$ is the state vector, t is the time variable, p is a vector of model parameters, and x_0 is the initial state. Besides the system dynamics[1] (1) including the model parameters p, we assume that the following specification is given to us. First, a probability distribution $G_{\mathcal{X}_0}$ with density $g_{\mathcal{X}_0}$ over a set of initial conditions $\mathcal{X}_0 \subseteq \mathbb{R}^n$, and second, a set of states called the *failure domain* $\mathcal{U} \subseteq \mathbb{R}^n$ associated with some time interval $\Delta_t = [t_0, t_1] \subseteq \mathbb{R}_{\geqslant 0}$ (the generalization to multiple time intervals is straightforward).

The goal in this work is to estimate the *failure probability*, which intuitively is the probability that a random trajectory emerging from $\mathcal{X}_0$ enters the set $\mathcal{U}$ within the time frame Δ_t. Let $\mathcal{Y}_0 := \{x_0 \in \mathcal{X}_0 : \exists t \in \Delta_t \text{ s.t. } \xi_{x_0}(t) \in \mathcal{U}\}$ be the set of initial states, $\mathcal{Y}_0 \subseteq \mathcal{X}_0$, that eventually enter the failure domain, where $\xi_{x_0}(t)$ is the solution of (1) for $x(0) = x_0$ at time t. Then the failure probability is

$$\mathbb{P}(\mathcal{U}) = \int_{\mathcal{Y}_0} dG_{\mathcal{X}_0} = \int_{\mathcal{Y}_0} g_{\mathcal{X}_0}(x_0) dx_0. \tag{2}$$

In the above equation, the distribution is precisely known. But as mentioned before, in reality, we often lack precise models of uncertainty. There are many imprecise probability models, with various degrees of generality and computational complexity. The simplest, and crudest, imprecise probability model is an interval. Intervals are usually interpreted as closed sets of real numbers, but they may be extended to a probabilistic interpretation by considering they bound all distributions with ranges bounded by the interval. A generalization of intervals are probability boxes (p-boxes) [28], which define a set of distributions with two bounding cumulative distribution functions. Other imprecise probability models include: Dempster-Shafer structures (also called belief functions) [19, 64], random sets [56], possibility distributions [73, 21, 37], lower previsions [69], and credal sets [47]. In this work, we study imprecise probabilities modeled with multivariate p-boxes, for the reason that p-boxes generalize both intervals and distribution functions, so that the presented methodology also applies to precise distributions, and because the multivariate extension of p-boxes in terms of copulas has been well studied [57].

[1] We assume that system (1) has a unique solution.

1.2. Verified propagation of imprecise probabilities

The most common practice to estimate the failure probability explores different behaviors via simulation, e.g., by Monte-Carlo methods [4, 2]. However, if uncertainty is present, e.g., if the distribution over the initial states is modeled with imprecise distributions, simulation-based approaches require a "double loop" to explore both different initial states and distributions. Some applications also require an exhaustive exploration of the state space. In both cases, simulation-based approaches become computationally intractable.

In a first step, our approach computes a set-based solution of the system, borrowing ideas used in the field of *reachability analysis*. This solution yields a guaranteed outer approximation of the states that can be reached in the dynamical system (1), and thus our approach provides absolute guarantees (in contrast with only statistical guarantees provided by simulations). The technique we use is known as *set propagation*, where the solution of an ODE is expressed in terms of sets of states rather than a single state.

In a second step, we compute rigorous bounds for probabilistic quantities (e.g., quantiles or moments) of interest, taking into account the multivariate p-boxes for the initial states. This step returns conservative (outer) approximations of failure probabilities, and we can control the precision of the results. When only considering sets, traditional reachability analysis can yield one of three possible probabilities: $\mathbb{P}(\mathcal{U}) = 0$ (guaranteed to be safe), $\mathbb{P}(\mathcal{U}) = 1$ (guaranteed to be unsafe), and the interval probability $\mathbb{P}(\mathcal{U}) = [0, 1]$ (unknown safety). The last of these options is a drawback of the usual method because we cannot prove or disprove the safety of the system. By introducing p-box initial conditions, the analysis can be generalized to return any $\mathbb{P}(\mathcal{U}) \subseteq [0, 1]$ for the system's safety. For this, two outer approximations for multivariate p-boxes with precise copulas are described, one based on belief functions, and a second based on consonant belief functions, i.e., possibility measures. The second of these representations allows to efficiently compute prediction sets of the dynamical system.

1.3. Related work

To our knowledge, there has been limited literature on computing *rigorous* interval bounds on probabilities in continuous dynamical systems, with most methods relying on Monte-Carlo sampling and simulation and requiring a precise definition of the initial distributions, i.e., they cannot handle imprecise probabilities. In this work we consider a distribution over the initial condition, but the system dynamics are deterministic. Below we focus on approaches with similar assumptions.

Enszer et al. [24] describe a method for p-box propagation in ODE systems. However, their method only uses stochastic independence, and they cannot compute failure probabilities. Their method also relies on "subintervalization" to propagate p-boxes. Maces [49] studies uncertainty represented as a fuzzy set and also uses similar set-propagation techniques. However, the approach does not support a fuzzy set theory from the perspective of bounding probabilities, and is thus arguably unsuitable for robust risk analysis [7].

Shmarov & Zuliani [66] describe a method to compute bounds on probabilities based on validated integration and a δ-complete decision procedure [31]. That method relies on state-space partitioning, which suffers from scalability issues. Huang et al. [39] synthesize a so-called barrier certificate, whose existence proves probabilistic safety (i.e., non-failure) in *unbounded* time, but is only applicable to simple distributions (uniform, normal, and exponential).

Sandretto [60] proposes a contractor for interval propagation, assuming that the initial interval comes with confidence levels. While generally orthogonal to our approach, the main computation also has to be applied only once and a-posteriori propagation is relatively efficient.

Tardioli et al. [68] estimate an upper bound for the impact probability of an asteroid based on unknown parametric distributions. In a similar setting, Serra et al. [63] assume that the initial positions of objects in space follow a Gaussian distribution and that the system has linear dynamics, which allows to preserve Gaussian distributions over time.

Stochastic systems have probabilities associated with the dynamics instead of the initial condition, which makes them orthogonal to the systems we study here. A common approach to compute the failure probability for such systems is called statistical model checking [46], which is essentially a Monte-Carlo estimation. Given a required confidence threshold, a number of simulations is drawn from the initial set. While conceptually simple, such approaches suffer from large numbers of required simulations, in particular for rare events, and only come with statistical guarantees, as opposed to absolute guarantees.

There are other, orthogonal classes of dynamical systems with stochastic elements. Stochastic differential equations (SDEs) [3, 59] are differential equations with a stochastic process (such as Brownian motion). Much less studied, random ODEs (RODEs) [67] are actual ODEs with a stochastic process; compared to SDEs, RODEs are sometimes the more suitable model. There are established rigorous analysis techniques for SDEs, e.g., in [25], while this problem has only recently been studied for RODEs [71].

Bouissou et al. [15] employ Dempster-Shafer structures for the analysis of programs, where they use affine arithmetic to mitigate the wrapping effect occurring with pure interval analysis.

1.4. Summary

The remainder of this article is structured as follows. In Section 2 we introduce notation and recall theoretical background. In Section 3 we present our contributions, which we summarize here.

- We show how to over-approximate a multivariate p-box into a belief function with copula dependencies (Section 3.1) in order to propagate them through functions (Section 3.2).
- We present an algorithm for estimating the failure probability, which integrates the p-box propagation and reachability techniques (Section 3.3).
- We describe an consonant approximation method for multivariate p-boxes which may be used to efficiently compute prediction sets of the dynamical system (Section 3.4).

In Section 4 we evaluate our methods in several numerical case studies. Finally, we conclude the article in Section 5.

An earlier version of this work was presented as an extended abstract in [35].

2. Preliminaries

In this work we borrow concepts from two different fields: probability theory and reachability analysis. The purpose of this section is to fix the notation and to familiarize the reader with the relevant aspects that we need later. We start with the probability part.

2.1. Probability boxes

A real-valued random variable is characterized by its cumulative distribution function (cdf) F, which is a monotonically increasing function from the real numbers onto the interval $[0, 1]$ such that the value of the function at negative infinity is zero, $F(-\infty) = 0$, and the value of the function at positive infinity is one, $F(\infty) = 1$. A probability box (p-box) consists of a pair of such functions that are used to circumscribe an imprecisely known distribution function F. The p-box, in its simplest form consisting of the pair of bounding functions, identifies a set of probability distributions (a credal set) that lie entirely within these bounds $\underline{F}(x) \leqslant F(x) \leqslant \overline{F}(x)$. Additional information about the random variable may be available, such as bounds on its mean and variance and its distribution family, which may be used to further restrict the set of distributions.

The origins of p-boxes can be related back to Kolmogorov and Marakov [50], who produced two cdf bounds on sums of random variables with unknown dependencies. Sklar, Schweizer and Frank generalized this result to other positive binary operations [30, 62]. Williamson [70] described a robust and efficient outer representation of p-boxes, describing them as *dependency bounds*, and using them to bound arithmetic operations with unknown dependencies. Their further generalization, modern definition and popularisation in risk and uncertainty analysis is often attributed to Ferson and Krenovich [26, 27, 28].

Real-valued statistics and features of distributions typically become intervals for p-boxes. The cdf of a p-box is

$$[\underline{F}(x), \overline{F}(x)].$$

A sample realization of the p-box may be drawn using the inverses of the bounding cdfs

$$[\overline{F}^{-1}(\alpha), \underline{F}^{-1}(\alpha)]$$

where $\alpha \sim U(0,1)$ is a sample from a uniform distribution. The probability measure on some interval $[a,b]$ is bounded as follows:

$$\underline{\mathbb{P}}([a,b]) = \max(0, \underline{F}(b) - \overline{F}(a))$$
$$\overline{\mathbb{P}}([a,b]) = \overline{F}(b) - \underline{F}(a),$$

where the max operator is required when $\underline{F}(b) < \overline{F}(a)$.

P-boxes generalize random variables and intervals in the following way. A random variable is modeled as a p-box with equal cdf bounds: $\underline{F}(x) = \overline{F}(x)$. In this case, all of the above probabilistic interval features of the p-box become precise. An interval is modeled as a p-box with step functions for bounds.

2.2. Belief functions

A Dempster-Shafer structure (also called a belief function) is a finite collection of intervals, called *focal elements* X, with probabilities attached to each interval (summing up to 1), called basic mass assignments m. An example of a belief function with three focal elements is: $E = \{[0,2],[0.5,1.3],[1,3]\}$ and $m = \{0.25, 0.5, 0.25\}$. Like a p-box, a Dempster-Shafer structure also bounds a set of probability distributions, using two set-based functions called *belief* and *plausibility*:

$$\begin{aligned} \text{bel}(\mathcal{U}) &= \sum_{X: E \subseteq \mathcal{U}} m(X), \\ \text{pl}(\mathcal{U}) &= \sum_{X: E \cap \mathcal{U} \neq \emptyset} m(X), \end{aligned} \tag{3}$$

where the sum is taken over the focal elements X from E that satisfy the given property (subset of resp. nonempty intersection with $\mathcal{U}$). For example, the belief and plausibility of $\mathcal{U} = [0, 2.5]$ are $\text{bel}([0, 2.5]) = 0.75$ ($\{[0,2],\ [0.5,1.3]\}$ are subsets) and $\text{pl}([0, 2.5]) = 1$ (all focal elements intersect with $\mathcal{U}$). The belief and plausibility serve as lower and upper bounds on the probability measure:

$$\text{bel}(\mathcal{U}) \leqslant \mathbb{P}(\mathcal{U}) \leqslant \text{pl}(\mathcal{U}),$$

thus for the previous example $\mathbb{P}(\mathcal{U}) = [0.75, 1]$. A p-box is a special case of a belief function, where the focal elements are ordered by $\leqslant$. Note that the focal elements do not have to be ordered in general belief functions, as was the case in the previous example: $E = \{[0,2],[0.5,1.3],[1,3]\}$.

2.3. Multivariate p-boxes

P-boxes have been extended to higher dimensions [57]. However, as is the case for precise probabilities when considering multivariate p-boxes, dependence information (like correlation) must be accounted for. Considering complex dependencies amongst p-boxes is important, as they play a key role in computations involving p-boxes. That is, the result of a function (and any computed risks) of p-boxes not only depends on the input p-boxes, but also on how they are correlated. Covariance, and subsequently the Pearson correlation coefficient, is insufficient to fully determine the dependence between two random variables. Like for the univariate moments, a single multivariate distribution cannot be prescribed for a given value of the covariance without making distributional assumptions (like normality). A more descriptive (and general) model for dependence is required to exactly specify a multivariate distribution.

Copulas [58] provide a solution to this problem. A copula $C : [0,1]^d \to [0,1]$ is a multivariate cdf with uniform marginals on the unit hypercube $[0,1]^d$, which have all marginal (univariate) information stripped away, leaving only the dependence. Any precise stochastic dependence information can therefore be modelled by a copula exactly and completely separately from its marginals. Multivariate problems in probability can often be reduced to an analysis of copulas.

A bivariate cdf F_{XY} can be constructed from two marginal cdfs F_X and F_Y and copula C using Sklar's theorem:

$$F_{XY}(x, y) = C(F_X(x), F_Y(y)),$$

where the constructed bivariate distribution will follow the dependence encoded in the copula and have the prescribed marginals. Copulas also give a convenient way to sample from F_{XY} in a generalization of inverse transform sampling, where first the copula is sampled $(u, v) \sim C$, and then the samples are transformed through the inverses of the marginals $(F_X^{-1}(u), F_Y^{-1}(v))$.

Sklar's theorem has a straightforward imprecise extension [57], where the two bivariate bounds of a p-box $[\underline{F}_{XY}(x, y), \overline{F}_{XY}(x, y)]$ can be expanded in terms of two marginal p-boxes and bounds on a copula $[\underline{C}(u, v), \overline{C}(u, v)]$:

$$\begin{aligned} \underline{F}_{XY}(x, y) &= \underline{C}(\underline{F}_X(x), \underline{F}_Y(y)), \\ \overline{F}_{XY}(x, y) &= \overline{C}(\overline{F}_X(x), \overline{F}_Y(y)). \end{aligned}$$

In this work, we use the imprecise Sklar's theorem to construct multivariate p-boxes using precise copulas.

2.4. Reachable states of dynamical systems

A *dynamical system* with d dimensions is characterized by a system of ordinary differential equations (ODEs) of the form (1), $x'(t) = f(x(t), p, t)$. Here we assume that f is analytic and typically given as the composition of standard arithmetic operations, which include addition, multiplication, exponentiation, and trigonometric functions. An *initial-value problem* (IVP) consists of a dynamical system together with an initial condition. In the simplest case, the initial condition is just a point $x(0) = x_0 \in \mathbb{R}^d$. Under the above assumption on f, an IVP has a unique solution, called a *trajectory*, which we write ξ_{x_0}, such that $\xi_{x_0}(t)$ satisfies the ODE f subject to $\xi_{x_0}(0) = x_0$.

Here we consider uncertain initial conditions coming from a set: $x(0) \in \mathcal{X}_0 \subseteq \mathbb{R}^d$, which induces a set of trajectories $\{\xi_{x_0} : x_0 \in \mathcal{X}_0\}$. Fixing a time point $t \in \mathbb{R}_{\geqslant 0}$, we define the set of *reachable states* at time t as $\mathcal{R}(\mathcal{X}_0, t) = \{\xi_{x_0}(t) : x_0 \in \mathcal{X}_0\}$. The reachable states for a time interval $[0, T]$,

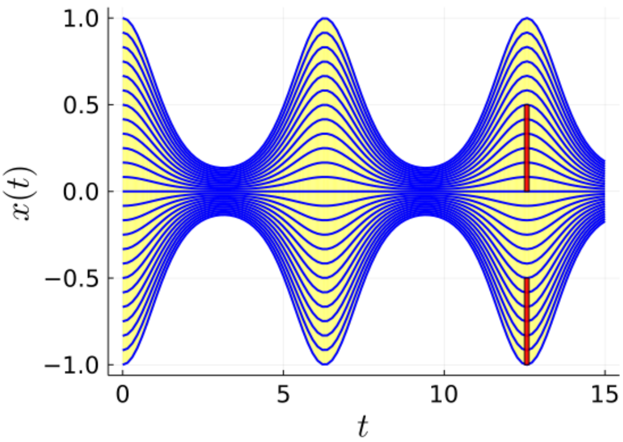

Figure 1: Flowpipe (yellow) for the univariate oscillator, together with 25 trajectories (blue) from uniformly picked initial states. In red we show the failure domain $\mathcal{U}$.

$\mathcal{R}(\mathcal{X}_0, [0, T]) = \{\xi_{x_0}(t) : x_0 \in \mathcal{X}_0, t \in [0, T]\}$, is also called a (exact) *flowpipe*. We note that the sets $\mathcal{R}$ are generally not computable [36], but over-approximations can be obtained [1].

Example 1. As a simple illustrative example we consider a nonlinear model of a univariate oscillator $x'(t) = -x(t)\sin(t)$. The analytic solution of this ODE is $x(t) = x(0)e^{\cos(t)-1}$. Hence, assuming that the initial condition is $x(0) \in \mathcal{X}_0 = [-1, 1]$, the reachable states at time t are $\mathcal{R}(\mathcal{X}_0, t) = \{x_0 e^{\cos(t)-1}, x_0 \in [-1, 1]\}$. Similarly, the flowpipe until the time horizon $T = 15$ is $\mathcal{R}(\mathcal{X}_0, [0, 15]) = \{x_0 e^{\cos(t)-1}, x_0 \in [-1, 1], t \in [0, 15]\}$, which is shown in yellow in Figure 1. As the failure domain we consider the union of two disjoint intervals at a time point $t = 4\pi$: $\mathcal{U} = [-1, -0.5] \cup [0, 0.5]$ (shown in red in Figure 1). The failure domain only partly encloses the flowpipe, and thus system safety cannot be guaranteed in this example.

2.5. Euclidean sets

We consider geometric sets $\mathcal{X} \subseteq \mathbb{R}^d$ of d-dimensional points. Common classes of sets include (1) the *interval box* (or hyperrectangle) $\{x \in \mathbb{R}^d : \forall i = 1, \ldots, d.\ \ell_i \leqslant x_i \leqslant u_i\}$, which is the Cartesian product of intervals $[\ell_i, u_i]$, (2) the *half-space* $\{x \in \mathbb{R}^d : a^T x \leqslant b\}$, which is parametric in $a \in \mathbb{R}^d$ and $b \in \mathbb{R}$ and characterizes all points on one side of the hyperplane $a^T x = b$, and (3) the *polyhedron*, which is the intersection of finitely many half-spaces. A central class of non-convex sets is described by Taylor models, which we introduce below in Section 2.6. In our application we consider pairs of sets $\mathcal{X}, \mathcal{U} \subseteq \mathbb{R}^d$ where $\mathcal{X}$ corresponds to a Taylor model and $\mathcal{U}$ is the failure

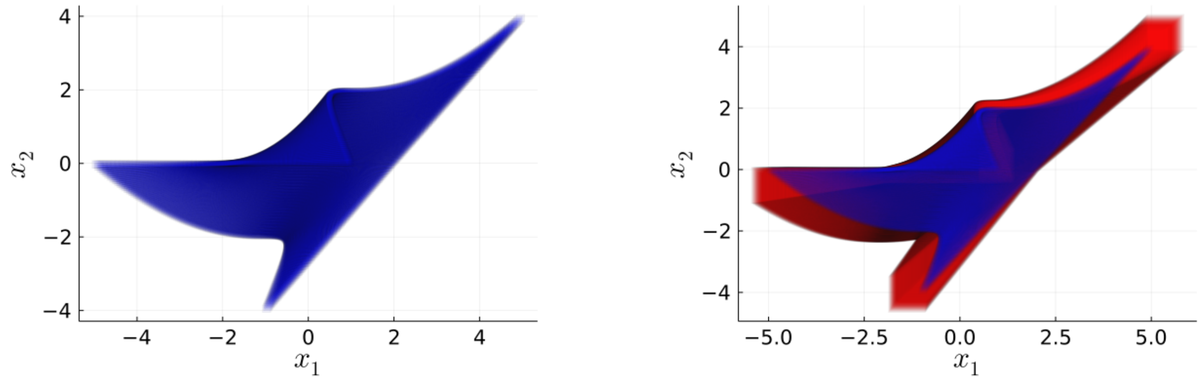

Figure 2: Example plot of a Taylor model (blue) and a Taylor-model reach set evaluated in its time domain (red).

domain, typically a polyhedron, or a union of polyhedra. In that setting, inclusion $\mathcal{X} \subseteq \mathcal{U}$ and disjointness $\mathcal{X} \cap \mathcal{U} = \emptyset$ can be checked efficiently in an approximate manner using suitable set transformations [29].

2.6. Taylor models as representation of reachable state

Taylor models provide a way to rigorously bound a function f over a domain by means of 1) a polynomial Taylor expansion around a point inside the domain and 2) an interval remainder. More formally, an d-dimensional *Taylor model* of order k is a tuple $\mathcal{T} = (p, I, x_0, \mathcal{D})$ where $p = (p_1, \ldots, p_d)^T$ is a vector of multivariate polynomials of degree at most k, $p_i : \mathcal{D} \to \mathbb{R}$ for $i = 1, \ldots, d$, obtained by Taylor-expanding a function around $x_0 \in \mathcal{D}$, $I = I_1 \times \cdots \times I_d$ is an interval box, called the remainder, containing the origin, and $\mathcal{D} \subseteq \mathbb{R}^d$ is the domain [11, 51, 52, 43]. Then, for all $x \in \mathcal{D}$ we have that $f(x) \in p(x) + I$, where the Taylor expansion is defined around x_0, i.e., $p_i(x_0) = 0$ for all i. Thus $\mathcal{T}$ represents the vector-valued function $p(x) + I$, which is an interval tube around the polynomial p. We often use the common convention that $\mathcal{D}$ is normalized to the symmetric box $[-1, 1]^d$, which can be enforced algorithmically.

Example 2. Let $d = 2$ and the polynomials $p_1(x) = 2x_1 + x_2 + 2x_1^2 x_2$, $p_2(x) = 2x_2 + 2x_1^3 x_2$. We consider a Taylor model of order 4, $\mathcal{T} = (p, I_1 \times I_2, (0, 0), [-1, 1]^2)$, with remainders $I_1 = I_2 = [-0.05, 0.05]$, which is plotted in Figure 2 (left).

A *Taylor-model reach set* (or *reach set* for short) $\mathcal{R}$ is a structure used in reachability algorithms when propagating Taylor models through an ODE in time. For an d-dimensional system, $\mathcal{R}$ is an d-vector of Taylor models. Each coordinate itself is a Taylor model in one variable representing time, of order k_t, with domain $\mathcal{D}_t$. The coefficients of these Taylor models are themselves multivariate polynomials in the d state variables, of order k and whose domain is assumed to be $\mathcal{D} = [-1, 1]^d$. (The time domain of $\mathcal{R}$ is not normalized to $[-1, 1]$.) Evaluating $\mathcal{R}$ over a time interval yields an d-dimensional Taylor model.

Example 3. We modify the previous example to consider a Taylor-model reach set in the time interval $[0, 0.2]$ with components $f_1(x, t) = p_1(x) + p_2(x)t + I_1$ and $f_2(x, t) = p_2(x) + p_1(x)t + I_2$. Each component is a Taylor model in one variable (time, t) whose coefficients are multivariate polynomials in the spatial variables (x_1 and x_2). A plot is shown in Figure 2 (right).

Taylor models can be used to approximate the reachable states. The procedure is outlined below; for a more technical presentation we refer to [43, 16]. First, we compute a non-validated integration of (1) using a Taylor model of order k_t. As said above, the coefficients of that series are polynomials of order k in the variables that denote small deviations of the initial conditions. We obtain a time step from the last two coefficients of this time series. In order to validate the integration step, we compute a second integration using intervals as coefficients of the polynomials in time, and we obtain a bound for the integration using a Lagrange-like remainder. The remainder is used to check the contraction of a Picard iteration. If the combination of the time step and the remainder do not satisfy the contraction, we iteratively enlarge the remainder or possibly shrink the time step. Finally, we evaluate the initial Taylor series with the valid remainder at the time step for which the contraction has been proven, which is also evaluated in the initial set to yield an over-approximation. The approach is (numerically) sound due to rigorous interval bounds in the Taylor approximation.

Example 4. We consider the univariate oscillator from Example 1. The Taylor-model reach set obtained for the time interval $\mathcal{D}_t = [0, 0.059472]$, domain $\mathcal{D} = [-1, 1]$, and orders $k = 2, k_t = 10$ is

$$\begin{aligned} &x_1 - 0.5x_1t^2 + 0.16666666666666666x_1t^4 - 0.043055555555555555x_1t^6 \\ &+ 0.009399801587301585x_1t^8 - 0.0018066578483245145x_1t^{10} + \underbrace{[-1.3938, 1.3938] \times 10^{-21}}_{\text{remainder}}. \end{aligned}$$

We can evaluate the Taylor-model reach set for any time interval contained in $\mathcal{D}_t$, which yields a Taylor model over x_1. For instance, choosing $[0, 0.01]$, we get the Taylor model

$$0.999975x_1 + [-2.50001, 2.50001] \times 10^{-5},$$

while choosing $[0.04, 0.05]$, we get

$$0.9989756663137614x_1 + [-2.25001, 2.25001] \times 10^{-4}.$$

Evaluating this Taylor model further in the spatial domain using interval arithmetic, i.e., for $x_1 \in [-1, 1]$, we get the intervals $[-1.00001, 1.00001]$ and $[-0.999201, 0.999201]$ respectively.[2] This is aligned with the solution shown in Figure 1, which shrinks toward the origin over time for the time interval $\mathcal{D}_t$. In our algorithm we will make use of the following key observation: we can partition the spatial domain $\mathcal{D} = [-1, 1]$ into *subsets* and evaluate the Taylor model for each of them. Indeed, evaluating over $[0.04, 0.05]$ and $\mathcal{D} = [0.6, 0.9]$ yields $[0.838937, 0.884296]$. The idea is that each subset of $\mathcal{D}$ will be uniquely associated with a focal element of a p-box.

3. Methodology

In this section we present the main contributions of this article. In the first part (Section 3.1) we outline a construction on how to over-approximate a multivariate p-box into a belief function with copula dependencies (Section 3.1), and then consider the problem of propagating multivariate p-boxes through functions under weak assumptions (Section 3.2). In the second part we develop an algorithm to estimate the failure probability of a dynamical system based on reachability analysis (Section 3.3). The algorithm combines over-approximation and propagation of a p-box from the first part with flowpipe approximation using Taylor models (see Section 2.6). We conclude the section with an extension of our algorithm to compute the prediction sets of dynamical systems under uncertain initial conditions (Section 3.4). Numerical results are presented in Section 4.

3.1. Over-approximating a multivariate p-box

Although it is well known how to convert a p-box into a belief function in one dimension, and in multiple dimensions under the assumption of independence, very few authors describe the general conversion of a p-box to a multivariate belief function with a copula dependence. We will do that next. Beginning with a set of d marginal p-boxes $[\underline{F}_i, \overline{F}_i]$ and an d-copula C_d, a multivariate p-box can be constructed using the imprecise Sklar's theorem [57].

$$\underline{F}(x_1, x_2, \ldots, x_d) = C_d(\underline{F}_1(x_1), \underline{F}_2(x_2), \ldots, \underline{F_d}(x_d)),$$
$$\overline{F}(x_1, x_2, \ldots, x_d) = C_d(\overline{F}_1(x_1), \overline{F}_2(x_2), \ldots, \overline{F_d}(x_d)).$$

[2] The bounds exceed -1 resp. 1 due to rigorous computations.

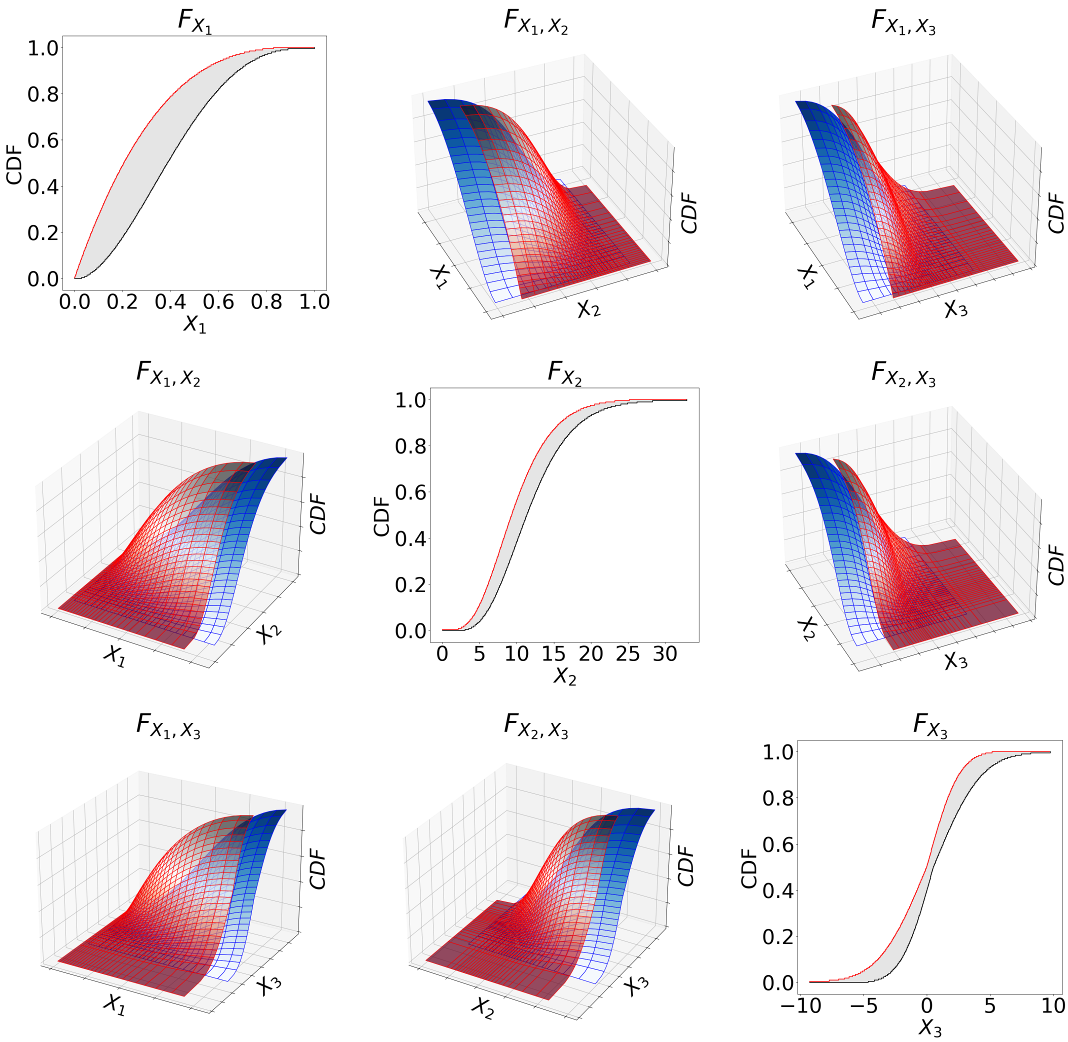

Figure 3: A visualization of a three-dimensional p-box. The diagonals show the univariate marginals, which are $X_1 \sim \text{Beta}([1,2],3)$, $X_2 \sim \text{Gamma}([5,6],2)$, and $X_3 \sim N([0,0.5],[2,3])$. The off-diagonals show the bivariate marginals.

Figure 3 is a visualization of a three-dimensional p-box constructed with Sklar's theorem, which has the marginals $X_1 \sim \text{Beta}([1,2],3)$, $X_2 \sim \text{Gamma}([5,6],2)$, and $X_3 \sim N([0,0.5],[2,3])$ (note that $N(\cdot)$ here denotes the normal distribution) shown along the diagonals, and a three-dimensional Gaussian copula with a parameter matrix

$$\begin{pmatrix} 1 & -0.8 & 0.8 \\ -0.8 & 1 & -0.8 \\ 0.8 & -0.8 & 1 \end{pmatrix}.$$

The two-dimensional cdf pairs are shown on the off-diagonals in Figure 3. The goal is to convert an d-dimensional p-box into an d-dimensional belief function with a finite number of focal elements, such that the belief function bounds at least the same (but possibly a larger) set of distributions as the p-box. The multivariate belief function will be a finite collection of focal elements, each with a basic mass assignment m_i. The focal elements of the multivariate belief function will be d-dimensional interval boxes $[a_1, b_1] \times [a_2, b_2] \times \cdots \times [a_d, b_d]$, determined by the shape of the marginals. The basic mass assignments m_i will be determined and correlated by the copula C_d.

We follow the discretization first introduced by Williamson & Downs [70], first converting the marginals into belief functions. Although these authors describe a discretization where each of the marginals has the same number of focal elements, the construction can be generalized such that the discretization of each marginal is different. Picking a particular number of desired focal elements N for a marginal, two vectors[3] u and d (for *up* and *down*) of length N are constructed by evaluating the marginal's inverse cdfs for uniformly spaced probability levels $p_j = \frac{j-1}{N}$ for $j = 1, \ldots, N+1$, e.g., $p_j = [0, 0.25, 0.5, 0.75, 1]$ for $N = 4$. The vectors u and d are defined as follows, for $j = 1, \ldots, N$

$$
\begin{aligned}
u_j &= \overline{F}^{-1}(p_j) \\
d_j &= \underline{F}^{-1}(p_{j+1})
\end{aligned}
$$

The vectors u and d serve as bounds on the inverses $\overline{F}^{-1}$ and $\underline{F}^{-1}$

$$
\begin{aligned}
u_j &\leqslant \overline{F}^{-1}(p) && p \in [p_j, p_{j+1}) \\
d_j &\geqslant \underline{F}^{-1}(p) && p \in (p_j, p_{j+1}],
\end{aligned}
$$

and give the interval bounds on the focal elements

$$
X_j = [u_j, d_j]. \tag{4}
$$

In one dimension, the basic mass assignment of each focal element is simply the difference of the probability level $m_j = p_{j+1} - p_j$, which for a uniform spacing gives equal mass assignment. Figure 4 shows converted marginal p-boxes from Figure 3, with numbers of focal elements $N_1 = 5$, $N_2 = 3$, $N_3 = 10$. The belief functions give rigorous outer approximations of the marginal p-boxes. Note that the larger the number of focal elements, the more tightly the belief functions enclose the

[3] We begin vector indexing with 1.

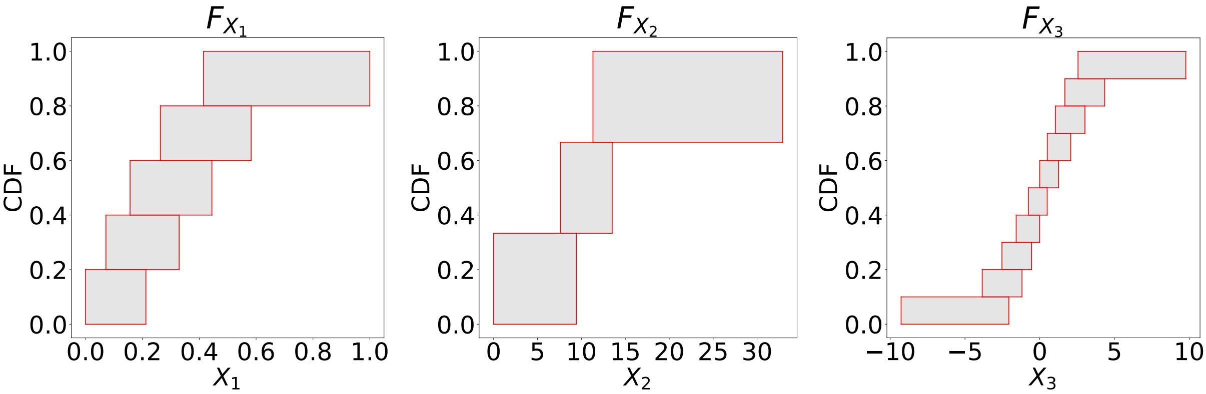

Figure 4: The over-approximated marginals of Figure 3, with $N_1 = 5$ (for x_1), $N_2 = 3$ (for x_2), and $N_3 = 10$ (for x_3), respectively.

p-boxes, and also the more tightly the credal sets of the belief functions enclose the p-boxes' credal sets. Although a uniform partition of the probability levels is shown here, the construction can be performed with any arbitrary spacing, for example if one desired to give finer resolution to the tails of the p-boxes. Note that this discretization will not necessarily produce a *partition* of the input space (the x axis on Figure 4), since the focal elements may overlap. The scheme does however perfectly partition the probability spaces (the y axis). In determining the mass assignment of the multivariate belief function, the copula will be integrated over this partition.

3.1.1. Mass assignment using copulas

The focal elements of the multivariate structure are found by taking Cartesian products between the marginal focal elements

$$K_{i_1,i_2,\ldots,i_d} = X_{1,i_1} \times X_{2,i_2} \times \cdots \times X_{d,i_d}, \tag{5}$$

which is done over all combinations of the marginal focal elements. For example, the Cartesian product between the focal elements of X_1 and X_2 from the example in Figure 4 (showing three significant digits) are:

	$X_{2,1}$	$X_{2,2}$	$X_{2,3}$
$X_{1,1}$	$[0, 0.213] \times [0, 9.43]$	$[0, 0.213] \times [7.61, 13.6]$	$[0, 0.213] \times [11.3, 33]$
$X_{1,2}$	$[0.072, 0.33] \times [0, 9.43]$	$[0.072, 0.33] \times [7.61, 13.6]$	$[0.072, 0.33] \times [11.3, 33]$
$X_{1,3}$	$[0.156, 0.445] \times [0, 9.43]$	$[0.156, 0.445] \times [7.61, 13.6]$	$[0.156, 0.445] \times [11.3, 33]$
$X_{1,4}$	$[0.263, 0.583] \times [0, 9.43]$	$[0.263, 0.583] \times [7.61, 13.6]$	$[0.263, 0.583] \times [11.3, 33]$
$X_{1,5}$	$[0.415, 1] \times [0, 9.43]$	$[0.415, 1] \times [7.61, 13.6]$	$[0.415, 1] \times [11.3, 33]$

When performing the mass assignment in higher dimensions, most authors choose random set independence, where the multivariate mass assignment is found by a simple multiplication between the marginal assignments m_{i,j_i}. In the more general case where the dependence is modeled by a copula C, the mass assignment will be determined by computing the C-*volume* on the probability partition of the marginals. The C-volume [62], denoted V_C, evaluates the probability measure from an d-copula C on some interval box $P = [\underline{p}_{x_1}, \overline{p}_{x_1}] \times [\underline{p}_{x_2}, \overline{p}_{x_2}] \times \cdots \times [\underline{p}_{x_d}, \overline{p}_{x_d}]$, which is a proper subset of the unit hypercube, $P \subseteq [0,1]^d$, as follows:

$$V_C(P) = \sum_{\mathbf{p} \in \text{vertices}(P)} \text{sign}_P(\mathbf{p}) C(\mathbf{p}), \tag{6}$$

where the sum is taken over all the vertices of P and where

$$\text{sign}_P(\mathbf{p}) = \begin{cases} 1 & \text{if } \mathbf{p} \text{ has an even number of lower bounds of } P \\ -1 & \text{if } \mathbf{p} \text{ has an odd number of lower bounds of } P. \end{cases} \tag{7}$$

As an example, the two-dimensional calculation is

$$V_C([\underline{p_x}, \overline{p_x}] \times [\underline{p_y}, \overline{p_y}]) = C(\overline{p_x}, \overline{p_y}) - C(\overline{p_x}, \underline{p_y}) - C(\underline{p_x}, \overline{p_y}) + C(\underline{p_x}, \underline{p_y}), \tag{8}$$

and the three-dimensional calculation is

$$\begin{aligned} V_C([\underline{p_x}, \overline{p_x}] \times [\underline{p_y}, \overline{p_y}] \times [\underline{p_z}, \overline{p_z}]) = {} & C(\overline{p_x}, \overline{p_y}, \overline{p_z}) - C(\overline{p_x}, \overline{p_y}, \underline{p_z}) - C(\overline{p_x}, \underline{p_y}, \overline{p_z}) \\ & - C(\underline{p_x}, \overline{p_y}, \overline{p_z}) + C(\overline{p_x}, \underline{p_y}, \underline{p_z}) + C(\underline{p_x}, \underline{p_y}, \overline{p_z}) \\ & + C(\underline{p_x}, \overline{p_y}, \underline{p_z}) - C(\underline{p_x}, \underline{p_y}, \underline{p_z}). \end{aligned}$$

The C-volume may be evaluated on the partition of the unit hypercube $[0,1]^d$ produced by the Cartesian product of all the marginal probability levels $P_i = [p_{i,j}, p_{i,j+1}]$. For example, the probability partition of X_1 and X_2 from Figure 4 is:

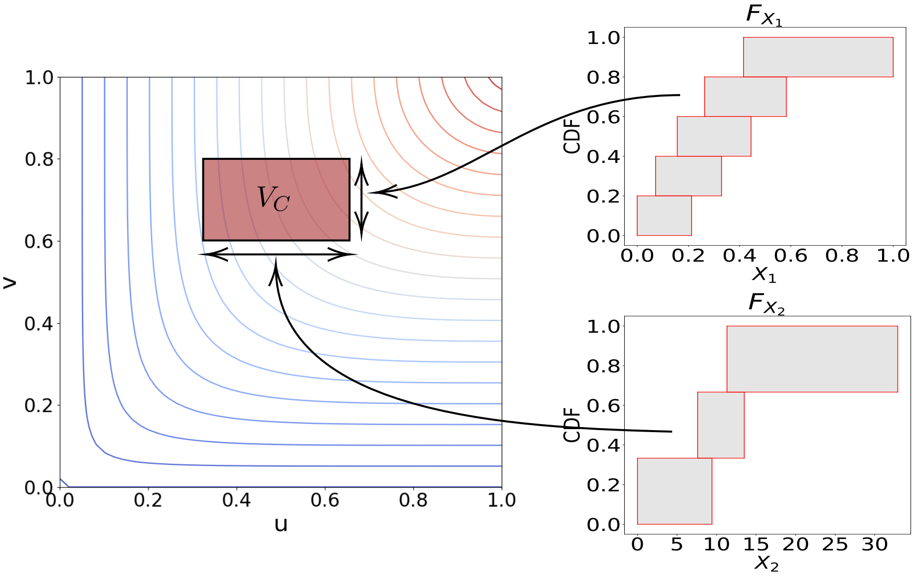

Figure 5: Computing the mass assignment of two focal elements using a copula and the C-volume V_C.

	$P_{2,1}$	$P_{2,2}$	$P_{2,3}$
$P_{1,1}$	$[0, \frac{1}{5}] \times [0, \frac{1}{3}]$	$[0, \frac{1}{5}] \times [\frac{1}{3}, \frac{2}{3}]$	$[0, \frac{1}{5}] \times [\frac{2}{3}, \frac{3}{3}]$
$P_{1,2}$	$[\frac{1}{5}, \frac{2}{5}] \times [0, \frac{1}{3}]$	$[\frac{1}{5}, \frac{2}{5}] \times [\frac{1}{3}, \frac{2}{3}]$	$[\frac{1}{5}, \frac{2}{5}] \times [\frac{2}{3}, \frac{3}{3}]$
$P_{1,3}$	$[\frac{2}{5}, \frac{3}{5}] \times [0, \frac{1}{3}]$	$[\frac{2}{5}, \frac{3}{5}] \times [\frac{1}{3}, \frac{2}{3}]$	$[\frac{2}{5}, \frac{3}{5}] \times [\frac{2}{3}, \frac{3}{3}]$
$P_{1,4}$	$[\frac{3}{5}, \frac{4}{5}] \times [0, \frac{1}{3}]$	$[\frac{3}{5}, \frac{4}{5}] \times [\frac{1}{3}, \frac{2}{3}]$	$[\frac{3}{5}, \frac{4}{5}] \times [\frac{2}{3}, \frac{3}{3}]$
$P_{1,5}$	$[\frac{4}{5}, \frac{5}{5}] \times [0, \frac{1}{3}]$	$[\frac{4}{5}, \frac{5}{5}] \times [\frac{1}{3}, \frac{2}{3}]$	$[\frac{4}{5}, \frac{5}{5}] \times [\frac{2}{3}, \frac{3}{3}]$

Figure 5 shows an example of computing the bivariate mass assignment of two marginal focal elements of X_1 and X_2 using the C-volume V_C. When computed on all pairs of focal elements, the following mass assignments are retrieved for a Gaussian copula with correlation -0.8:

$V_C(P_{1,i} \times P_{2,j})$	$j = 1$	$j = 2$	$j = 3$
$i = 1$	0.00211	0.03046	0.1674
$i = 2$	0.01638	0.08431	0.0993
$i = 3$	0.04811	0.10379	0.0481
$i = 4$	0.09931	0.08431	0.0164
$i = 5$	0.16743	0.0305	0.0021

As a partition of $[0, 1]^d$ was produced, all the mass assignments will sum to one. Figure 6 shows a visualization of the resulting three-dimensional belief function produced from the described proce-

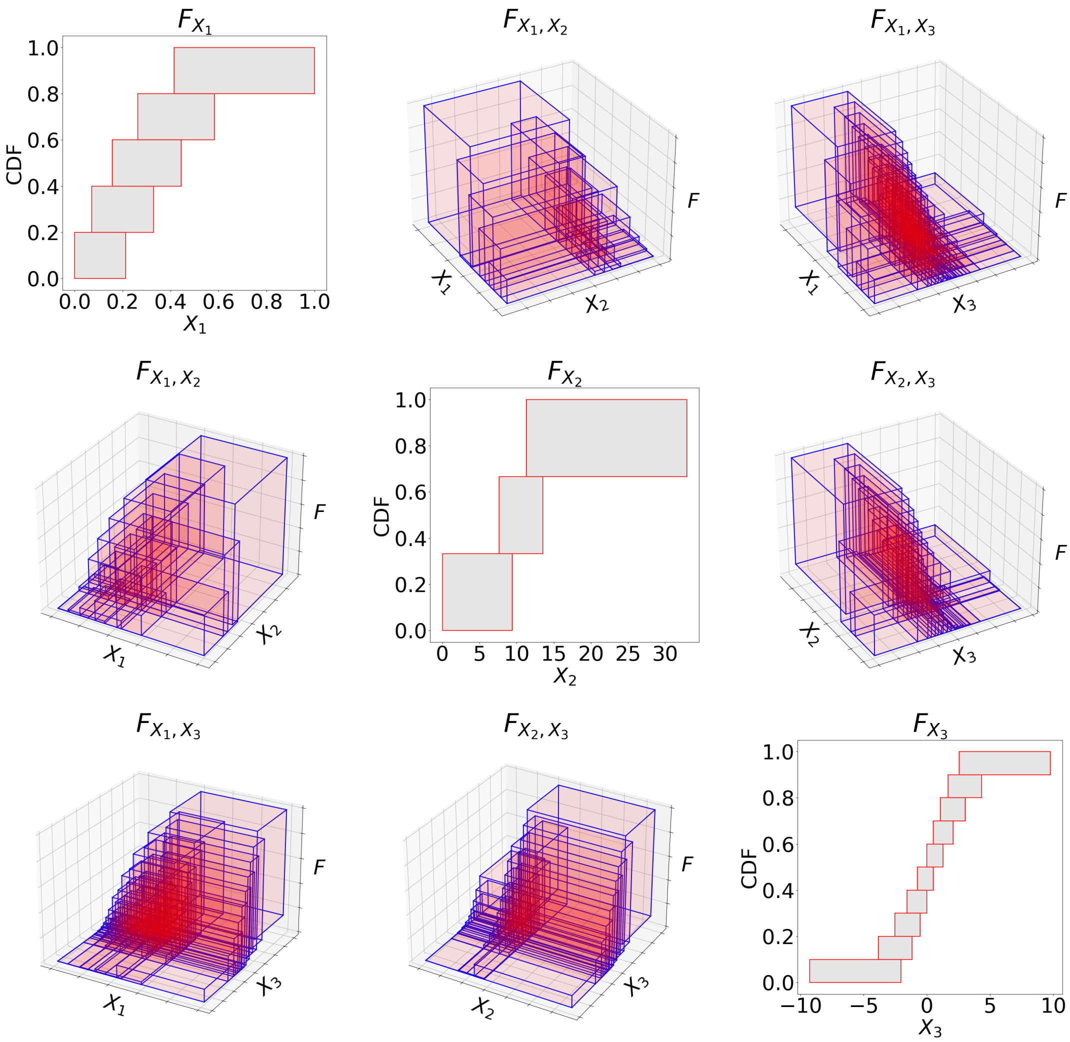

Figure 6: A visualization of a three-dimensional belief function, constructed from the multivariate p-box from Figure 3. The diagonals show the converted marginals, with the off-diagonals showing cumulative belief and plausibility of the bivariate pairs.

dure. The diagonals show the marginal belief functions are shown, while the off-diagonals showing the cumulative belief and plausibility of the bivariate pairs. The union of the shown boxes gives an outer approximation of the multivariate p-box shown in Figure 3.

3.2. Propagation of p-boxes

Once over-approximated with a belief function, the p-box can be propagated through any real function f that has an interval extension. Following Dempster–Shafer theory, the focal elements may be evaluated with f as

$$K_{out} = f(K_{i_1,i_2,\ldots,i_d}), \tag{9}$$

forming the focal elements of an output belief function. The mass assignment of the focal elements are retained by the images [72]

$$m(K_{\mathrm{out}}) = m(f(K_{i_1,i_2,\ldots,i_d})) = m(K_{i_1,i_2,\ldots,i_d}). \tag{10}$$

The resulting output belief function may then be used to bound the probability measure on any set $\mathcal{U}$ in the output space of f using the formulas (3). The belief and plausibility give interval bounds on the probability that the imprecisely known random variable falls in the set $\mathcal{U}$

$$\mathbb{P}(\mathcal{U}) = [\mathrm{bel}(\mathcal{U}), \mathrm{pl}(\mathcal{U})].$$

The belief function may also be converted back to a multivariate p-box. Some information is however lost in this process: the set of distributions is larger for the p-box, and the belief function generally provides tighter interval probabilities. However, the multivariate p-box can be useful for visualizing the imprecise random variable. Considering sets of the form $A_{y_i} = (-\infty, y_i]$, bounds on the cdf may be computed by accumulating the belief and plausibility:

$$\underline{F}_{\mathrm{out}}(y_1, y_2, \ldots, y_m) = \mathrm{bel}(A_{y_1} \times A_{y_2} \times \cdots \times A_{y_m}) \tag{11}$$

$$\overline{F}_{\mathrm{out}}(y_1, y_2, \ldots, y_m) = \mathrm{pl}(A_{y_1} \times A_{y_2} \times \cdots \times A_{y_m}) \tag{12}$$

The propagated belief function may also be used to compute prediction sets of the output, however, in Section 3.4 we provide a more efficient method to determine this.

3.3. Algorithm for estimating the failure probability

In this section we describe our algorithm to compute lower and upper bounds on the failure probability (see Eq. (2)) based on the observations in the previous sections. Assume we are given an IVP with dynamics f and initial condition $\mathcal{X}_0$, where $\mathcal{X}_0$ is associated with a p-box, and a failure domain $\mathcal{U}$ associated with a time interval Δ_t. The goal is to estimate the failure probability.

A naive approach to propagating the p-box is to repeat the reachability analysis for each focal element, i.e., for all ℓ interval boxes, constituting the multivariate p-box, compute ℓ flowpipe solutions of the dynamical system. This approach quickly becomes computationally prohibitive.

Algorithm 1: Estimation of the failure probability.

Input: $(f, \mathcal{X}_0)$: IVP; $\mathcal{U}$: failure domain; Δ_t: time interval for failure

Output: bel: belief; pl: plausibility of the failure probability

$\mathcal{R}_1, \ldots, \mathcal{R}_r \leftarrow$ solve(f, $\mathcal{X}_0$, $[0, \max(\Delta_t)]$); `// obtain` r `Taylor-model reach sets`
$(K_1, m_1), \ldots, (K_\ell, m_\ell) \leftarrow$ over-approx($\mathcal{X}_0$); `// obtain` ℓ `focal elements and masses`
bel $\leftarrow 0$; pl $\leftarrow 0$;
for $k \leftarrow 1$ **to** ℓ **do** `// loop over focal elements`
    intersects $\leftarrow$ false;
    for $j \leftarrow 1$ **to** r **do** `// loop over Taylor-model reach sets`
        if $\Delta_t \cap \mathcal{D}_t(\mathcal{R}_j) = \emptyset$ **then** continue;
        $\mathcal{X} \leftarrow$ evaluate($\mathcal{R}_j$, $\Delta_t \cap \mathcal{D}_t(\mathcal{R}_j)$, K_k); `// evaluate Taylor-model reach set for time interval and chosen focal element`
        if $\neg$ intersects $\wedge$ $\mathcal{X} \cap \mathcal{U} \neq \emptyset$ **then**
            pl $\leftarrow$ pl $+ m_k$; `// increase plausibility`
            intersects $\leftarrow$ true;
        end
        if intersects $\wedge$ $\mathcal{X} \subseteq \mathcal{U}$ **then**
            bel $\leftarrow$ bel $+ m_k$; `// increase belief`
            break;
        end
    end
end
return [bel, pl]

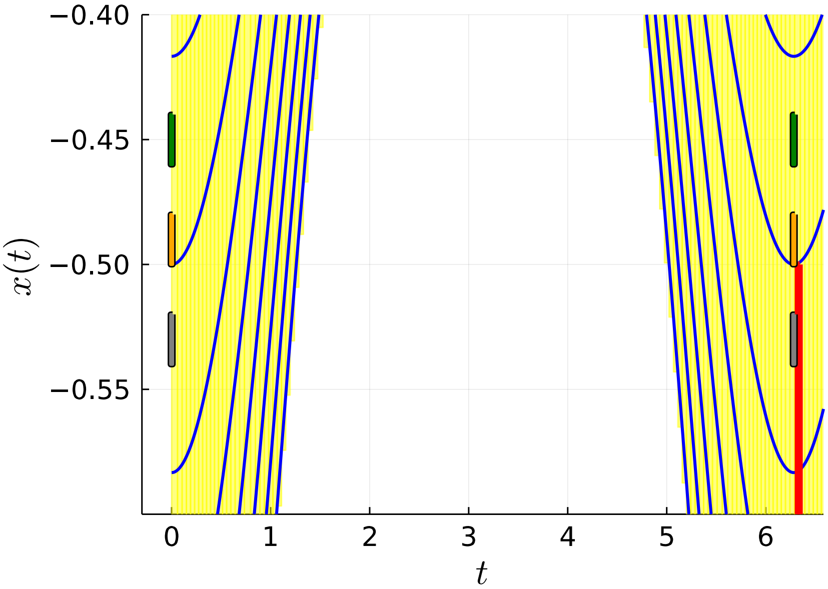

Figure 7: Propagation of three focal elements (gray, orange, green) for the univariate oscillator, from $t = 0$ to $t = 2\pi$ (see Figure 1 for a zoomed-out plot). The failure domain is drawn in red, with a small x-offset for better visibility.

Instead we compute a single flowpipe and propagate the focal elements by evaluating the Taylor models on the sub-domains, as shown in Example 4.

The procedure is summarized in Algorithm 1, which consists of two conceptual phases. In the first phase (line 1) we compute the flowpipe, for which we use an established procedure as mentioned in Section 2.6. Here we focus on the second phase, in which we estimate the failure probability using p-box propagation. First we obtain the focal elements and associated masses from the distribution over the initial states $\mathcal{X}_0$, according to a prescribed p-box discretization (line 2), as explained in Section 3.1. The purpose of the following double loop is to compute how each focal element affects the plausibility and belief. Fixing a focal element K_k, we iterate over the reach sets that affect the failure probability. More specifically, given a reach set $\mathcal{R}_j$, we obtain a new set $\mathcal{X}$ that represents the evaluation of $\mathcal{R}_j$ over the time interval $\Delta_t \cap \mathcal{D}_t(\mathcal{R})$ and over the spatial domain K_k (line 8). Once $\mathcal{X}$ has been obtained, one among three cases is considered: if $\mathcal{X}$ and $\mathcal{U}$ are disjoint, then the focal element contributes to neither the plausibility nor the belief. Otherwise, according to Eq. (3), the plausibility is incremented by the mass m_k. If, in addition, $\mathcal{X}$ is fully contained in the failure domain $\mathcal{U}$, then the belief is also incremented by the mass m_k.

In Figure 7, we illustrate the core of the algorithm on the univariate oscillator. We consider three focal elements (gray, orange, green) and compute their image at time 2π. (Since this model is 2π-periodic, these images are identical to the original intervals at $t = 0$.) We consider a slightly simplified failure domain for easier illustration (only the lower interval is kept, and the time component is shifted from 4π to 2π). The failure domain (red) fully contains the gray image, and intersects (here: at a point) the orange image, while it is disjoint from the green image. Hence the mass of the gray focal element contributes to both the plausibility and the belief, whereas the mass of the orange focal element only contributes to the plausibility, and the mass of the green focal element contributes to neither of them.

Implementation. For better scalability, our implementation preprocesses the list of reach sets for the loop in line 6 to skip useless computations for each focal element 1) if the time domain of the j-th reach set does not intersect with the time component of the failure domain (line 7), and, more crucially, 2) if the reach set evaluated over the whole spatial domain (i.e., all focal elements together) does not intersect with the failure domain. These are single checks that are independent of the discretization into focal elements, and they speed up the analysis in the common case that most of the reach sets are irrelevant for the failure property.

Our implementation is rigorous, i.e., it deals with potential imprecision due to floating-point arithmetic in a conservative way. This means that our implementation always over-estimates the failure probability. That is the case because 1) the Taylor-model reach sets we compute are provably outer approximations of the true reachable states and 2) all computations are performed with interval arithmetic and conservative rounding.

Further, our implementation supports multiple failure domains with different time components. The generalization of the algorithm is straightforward and essentially requires another loop over the failure domains (plus some clutter and bookkeeping steps, which is why we decided to present the simpler version in Algorithm 1).

Convergence. Our algorithm is effective and terminates for any precision settings (Taylor-model approximation and p-box discretization). By tweaking these parameters, we argue that the interval computed by Algorithm 1 converges to the true probability (a scalar for a precise distribution and an interval for an imprecise distribution). In the limit, the Taylor-model approximation converges to the true reachable states at a given time interval, and this time interval converges to a time

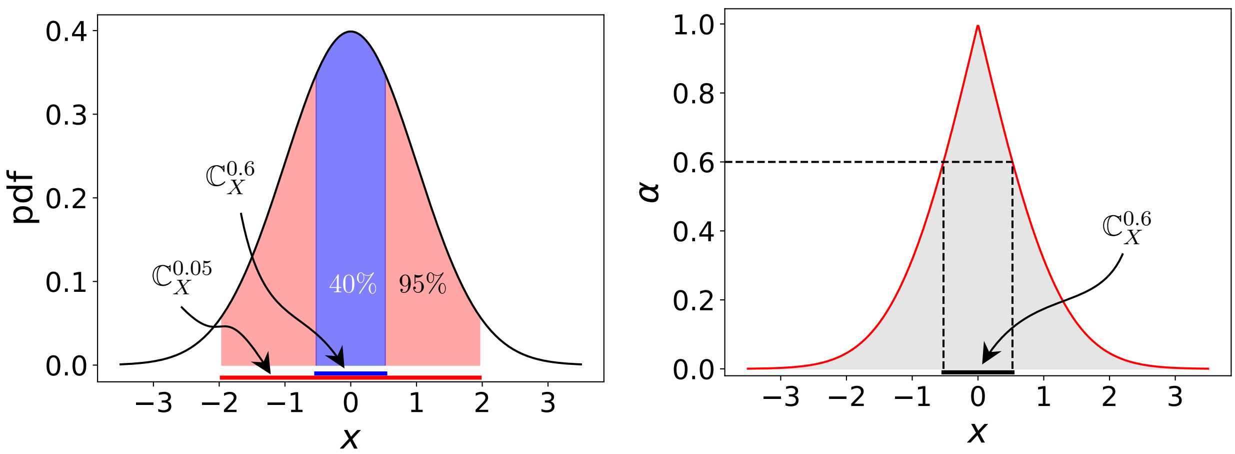

Figure 8: Left shows the construction of two-sided prediction sets $\mathbb{C}_X^{0.05}$ and $\mathbb{C}_X^{0.6}$ of a standard normal distribution, containing 95% and 40% of the probability respectively, from integrating the density. Right shows the possibility distribution obtained when all $\mathbb{C}_X^{\alpha}$ are stacked, and how $\mathbb{C}_X^{0.6}$ is drawn.

point. Similarly, with a more precise discretization of the p-box, the focal elements shrink, which shrinks the width between plausibility and belief, converging to the true failure probability.

3.4. Consonant approximation of a multivariate p-box

In this section we describe a method to transform a multivariate p-box under a precise copula dependence into a possibility distribution (or a fuzzy-set) [22], such that the possibility distribution forms an outer approximation to the original multivariate p-box. This transformation is based on the viewpoint of a possibility distribution as a consonant belief function, a belief function with nested focal elements. This transformation is useful, as it allows one to evaluate the probabilistic prediction sets of a flowpipe efficiently.

A possibility distribution can be interpreted as a collection of nested prediction sets (or percentile sets), e.g., the sets containing at least 95%, 70%, 20%, etc. of the probability measure at all levels [65]. The α prediction set (or an α-cut) $\mathbb{C}_X^{\alpha}$ of a precise random variable X can be defined as

$$\mathbb{P}_X\left(\mathbb{C}_X^{\alpha}\right) = 1 - \alpha \qquad \forall\, \alpha \in [0, 1]. \tag{13}$$

For example, the $\alpha = 0.2$ prediction set $\mathbb{C}_X^{0.2}$ contains 80% of the probability. Like confidence intervals, α-cuts can be constructed in various ways and may be one-sided or two-sided. The left of Figure 8 shows the usual construction of the prediction sets using the probability density and

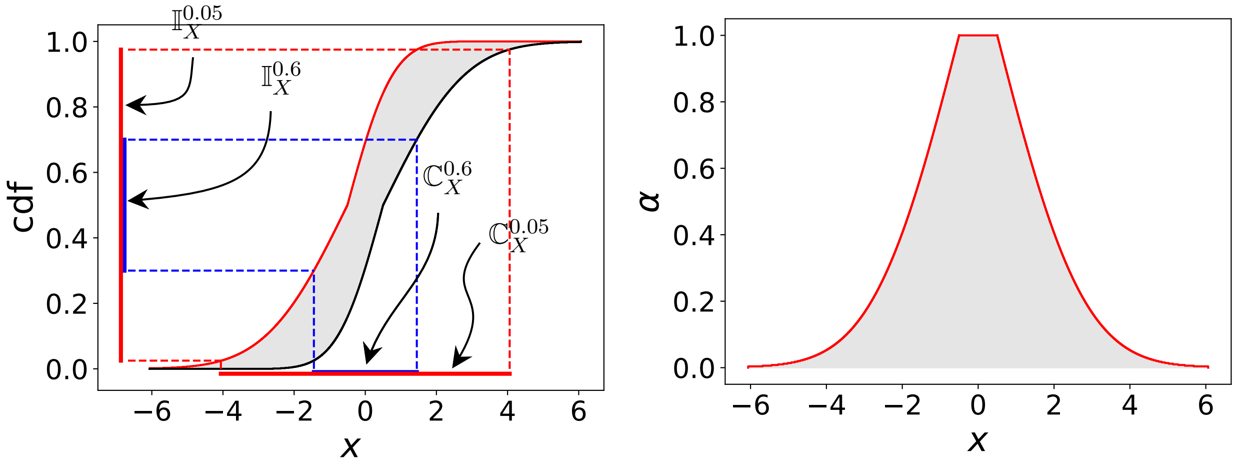

Figure 9: Visualization of constructing two-sided nested prediction sets from a p-box. Right shows the structure obtained when all $\mathbb{C}_X^\alpha$ are stacked.

integration, showing the 95% ($\mathbb{C}_X^{0.05}$) and 40% ($\mathbb{C}_X^{0.6}$) sets of a standard normal distribution. The right of Figure 8 shows a possibility distribution obtained when all $\mathbb{C}_X^\alpha$ are stacked for all α-levels, and shows how $\mathbb{C}_X^{0.6}$ is obtained from the structure.

When X is described by a p-box, the probability measure $\mathbb{P}_X$ returns intervals, and therefore the prediction sets of a p-box can be defined as the set containing *at least* $1-\alpha$ of the probability,

$$\mathbb{P}_X\left(\mathbb{C}_X^\alpha\right) \geqslant 1-\alpha \qquad \forall\, \alpha \in [0,1]. \tag{14}$$

Compared to integrating the density, it is simpler to use the cdf, where first symmetric nested sets are constructed on $[0,1]$,

$$\mathbb{I}^\alpha = \left[\frac{\alpha}{2}, 1-\frac{\alpha}{2}\right],$$

(e.g., $\mathbb{I}^{\alpha=0} = [0,1]$ is the entire probability interval and $\mathbb{I}^{\alpha=1} = 0.5$) and then evaluated in the inverse cdf:

$$\mathbb{C}_X^\alpha = [\underline{F}_X^{-1}(\underline{\mathbb{I}^\alpha}), \overline{F}_X^{-1}(\overline{\mathbb{I}^\alpha})].$$

Because the inverse distribution F_X^{-1} is an isoprobabilistic mapping from a uniform distribution $u \sim U(0,1)$ to X, we have that

$$\mathbb{P}_X(\mathbb{C}_X^\alpha) \geqslant \mathbb{P}(u \in \mathbb{I}^\alpha) = 1-\alpha, \tag{15}$$

and therefore by construction the inequality (14) holds, with equality given when $\underline{F}_X = \overline{F}_X$. The left of Figure 9 is a visualization of this construction from a normally distributed p-box, showing

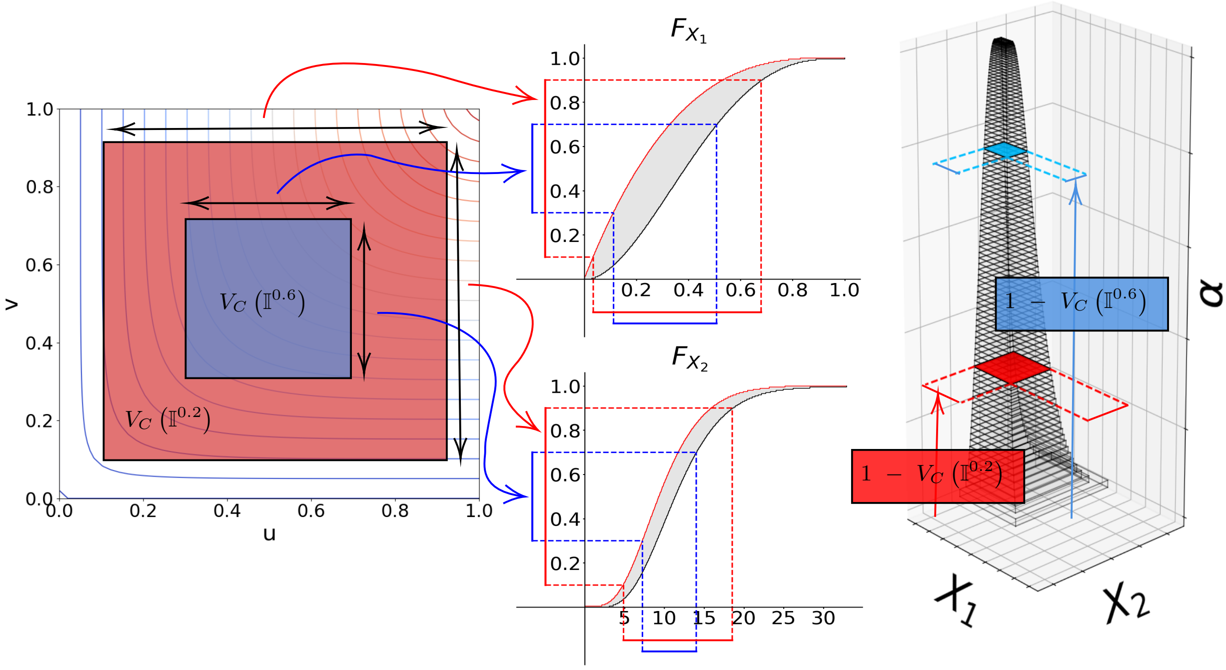

Figure 10: Visualization of the proposed consonant approximation of a multivariate p-box under copula dependence. Left shows the probability measure V_C of the copula being computed on two symmetric boxes $\mathbb{I}^{\beta=0.2}$ and $\mathbb{I}^{\beta=0.6}$. The center shows these sets being transformed through the marginal p-boxes, determining the margins of $\mathbb{C}_X^{\alpha}$. The height of $\mathbb{C}_X^{\alpha}$ is determined by $\alpha = 1 - V_C(\mathbb{I}^{\beta})$, shown on the right plot. When this is performed for all $\beta \in [0, 1]$, the gray possibility distribution on the right is obtained.

two sets $\mathbb{I}^{0.05}$ and $\mathbb{I}^{0.6}$ and their corresponding prediction sets $\mathbb{C}_X^{0.05}$ and $\mathbb{C}_X^{0.6}$. The sets $\mathbb{C}_X^{0.05}$ and $\mathbb{C}_X^{0.6}$ have a lower probability of $\underline{\mathbb{P}}_X(\mathbb{C}_X^{0.05}) = 0.95$ and $\underline{\mathbb{P}}_X(\mathbb{C}_X^{0.6}) = 0.4$ respectively. The right of Figure 9 shows the resulting possibility distribution when $\mathbb{C}_X^{\alpha}$ are stacked.

In the multivariate case, we construct a nested family of d-dimensional interval boxes $\mathbb{C}_X^{\alpha}$ which bound the probability measure of the multivariate p-box with (14). First a symmetric nested family $\mathbb{I}^{\beta}$ of boxes on $[0, 1]^d$ is constructed as

$$\mathbb{I}^{\beta} = \bigtimes_{1}^{d} \left[\frac{\beta}{2}, 1 - \frac{\beta}{2}\right], \tag{16}$$

for example in two dimensions $\mathbb{I}^{\beta=0.2} = [0.1, 0.9] \times [0.1, 0.9]$. The probability measure from the copula can then be computed using the C-volume $V_C(\mathbb{I}^{\beta})$ (6), and transforming $\mathbb{I}_X^{\beta}$ for all $\beta \in [0, 1]$

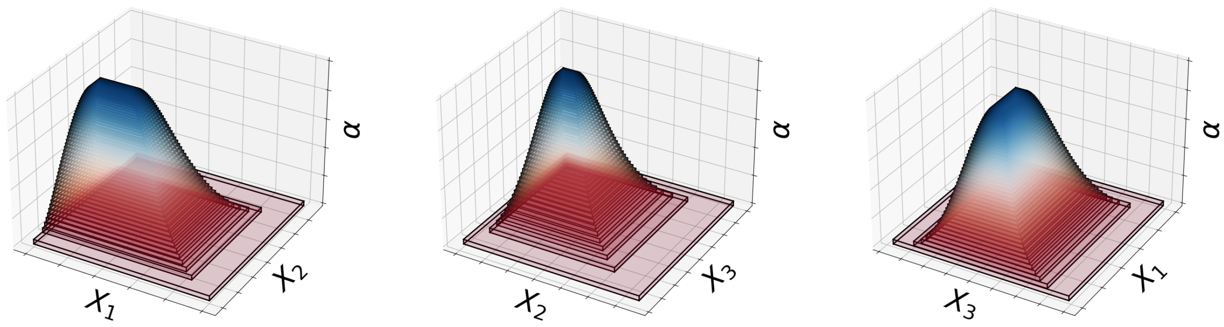

Figure 11: The two-dimensional marginals of the possibility distribution obtained from the example p-box in Figure 3.

through the inverse of the marginals yields the following nested set of interval boxes

$$\mathbb{C}_X^{1-V_C(\mathbb{I}^\beta)} = \bigtimes_{i=1}^{d} [\underline{F}_{X_i}^{-1}(\underline{\mathbb{I}}_i^\beta), \overline{F}_{X_i}^{-1}(\overline{\mathbb{I}}_i^\beta)], \tag{17}$$

where $F_{X_i}^{-1}$ are the inverse marginals. Again by construction (14) holds,

$$\mathbb{P}_X(\mathbb{C}_X^{1-V_C(\mathbb{I}^\beta)}) \geqslant V_C(\mathbb{I}^\beta),$$

or $\alpha = 1 - V_C(\mathbb{I}^\beta)$, as the inverses $F_{X_i}^{-1}$ are an isoprobabilistic mapping to X from the unit hypercube $[0,1]^d$ with a multivariate uniform distribution following the copula C. Figure 10 is a visualization of this construction in two dimensions, showing how two symmetric boxes $\mathbb{I}^{\beta=0.2}$ and $\mathbb{I}^{\beta=0.6}$ are transformed through the marginal distributions, forming the margins of α-cuts $\mathbb{C}_X^\alpha$, where the height of these α-cuts is given by $1 - V_C(\mathbb{I}^\beta)$. When this procedure is performed for all β levels, the entire possibility distribution in black on the rightmost plot is found. Figure 11 shows the two-dimensional marginal confidence structures obtained from the example multivariate p-box in Figure 3 using this construction.

The constructed possibility distribution is strongly related to Birnbaum confidence curves [13], confidence structures [5], and inferential models [54, 48], and can be interpreted as belief functions whose focal elements are nested [22, 23, 7, 37]. An important aspect of this is that possibility theory can be used to efficiently evaluate the confidence intervals of a flowpipe. Once multivariate α-cuts $\mathbb{C}_X^\alpha$ have been constructed, they can be propagated through any function $f : X \mapsto Y$ which has an interval extension, with the α-level preserved by the output [22, 37]:

$$\mathbb{C}_Y^\alpha = f(\mathbb{C}_X^\alpha) \tag{18}$$

Therefore, for any reach set, the α-level prediction set $\mathbb{C}_Y^\alpha$ can be found by a single set propagation of $\mathbb{C}_X^\alpha$. This is much more efficient than propagating the original multivariate p-box as described in Section 3.2, which has exponential complexity in the dimension.

Note that this is distinct to the 'level-wise' propagation originally proposed by Zadeh [73] in his fuzzy-set theory, where marginal fuzzy numbers are cut at equal α-levels, combined, and evaluated through the functions. This has been shown not to be consistent with probability theory by several authors [7, 37], and is argued not to be suitable for robust risk analysis. The distinction between the proposed construction and Zadeh's level-wise propagation is that here the entire multivariate fuzzy set has been derived, and not just the marginals, and as such the dependence of the original p-box has been captured. It has been argued by Hose & Hanss [37] that it is the disregard of dependence information that makes level-wise propagation unsuitable, and they have proposed several combination rules to capture different dependence assumptions. The proposed consonant-approximation method may also be seen as an application of the imprecise-probability-to-possibility transformation [38] for multivariate p-boxes with a precise copula dependence.

The constructed possibility distribution forms an outer approximation of the multivariate p-box in the following way. Like a p-box, a possibility distribution can also be viewed as a special case of a belief function, but with nested focal elements. As such, a possibility distribution can be rigorously represented by a finite number of focal elements K_i with basic mass assignment m_i, both constructed from the α-cuts. For a partition of length N of α-levels with $0 = \alpha_1 < \alpha_2 < \cdots < \alpha_{N+1} = 1$, the focal elements are the cuts of the multivariate possibility distribution at the α_i's,

$$K_i = \mathbb{C}_X^{\alpha_i}, \tag{19}$$

with basic mass assignment given by difference between the levels,

$$m_i = \alpha_{i+1} - \alpha_i. \tag{20}$$

Note that $K_i \subseteq K_{i+1}$ because the α-cuts are nested, and $\sum_i m_i = 1$ because α was partitioned. The computed interval probability using (3) from the original multivariate p-box will always be contained in the interval probability using (19) and (20). Stated another way, the credal set (set of probability distributions) defined by the original p-box is contained in the credal set of the possibility distribution.

The belief function created using the consonant approximation (17), with focal elements (19) and basic mass assignment (20), may also be used in Algorithm 1 to rigorously bound the failure probability. Again, the advantage is that the number of focal elements in (19) does not scale exponentially in the dimension, whilst it does for the original p-box (5).

However, the produced interval probability may be quite wide, depending on how well the possibility distribution approximates the original multivariate p-box. We also highlight that the above transformation (17) is not unique, and that there are multiple distinct possibility distributions that can be created from the same multivariate p-box. Indeed there are infinitely many consonant approximations of a multivariate p-box, depending on the selection of (the nested family of) the sets $\mathbb{I}^\beta \subseteq [0,1]^d$ used in (17). We proposed that symmetric boxes are used in (16), but we do not claim that this will give the tightest approximation of the p-box. Any nested family of subsets on $[0,1]^d$ could be used for $\mathbb{I}^\beta$, leading to different consonant approximations, some of which will be tighter than others. Irregardless of the selection of the family for $\mathbb{I}^\beta$, the consonant approximation will be rigorous, and we leave determining the optimal transformation for future work. The sets could potentially even be non hyper-rectangular, but computing the probability measure V_C and the inverse transformation through F_X^{-1} would have to be determined for such sets.

We also note that the above consonant approximation, and indeed all of the methods proposed in this paper, apply to precise probability distributions as well as p-boxes.

4. Numerical evaluation

We implemented our approach in the Julia language [12] in several open-source libraries. Probability bounds are computed using the package ProbabilityBoundsAnalysis.jl [34]. Flowpipe approximations are computed using the JuliaReach suite [14], the JuliaIntervals suite, TaylorSeries.jl [10], and TaylorModels.jl [9]. All experiments are executed on a laptop computer running Linux, with an Intel Core i7-8705G@3.10 GHz CPU and 16 GB RAM.

4.1. Example 1: Univariate oscillator

As our first model, we look again at the nonlinear univariate oscillator $x'(t) = -x(t)\sin(t)$. The analytic solution is $x(t) = x_0 e^{\cos(t)-1}$. We consider two initial conditions: a uniform distribution $x(0) \sim U(-1,1)$, which we coin the *precise* condition, and an *imprecise* condition where the lower and upper bounds are only known to lie in intervals $x(0) \sim U([-1,0],[0,1])$.

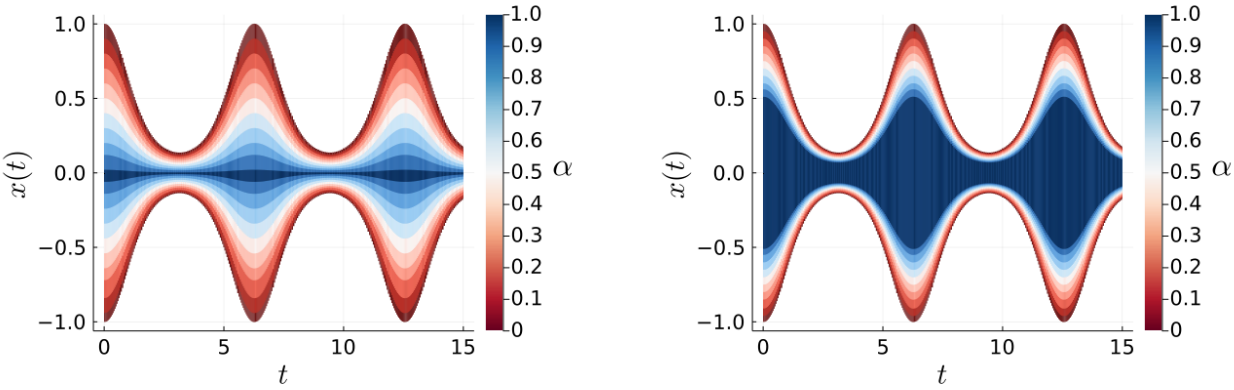

Figure 12: Two-sided prediction sets for a precise initial condition (left) and an imprecise initial condition (right).

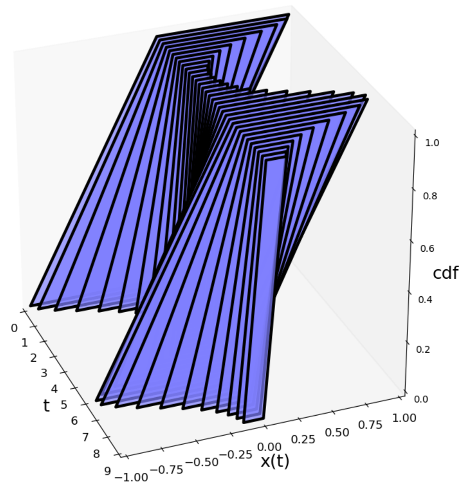

Figure 13: Propagation of the initial p-box $x(0) \sim U([-1, 0], [0, 1])$ through the univariate oscillator.

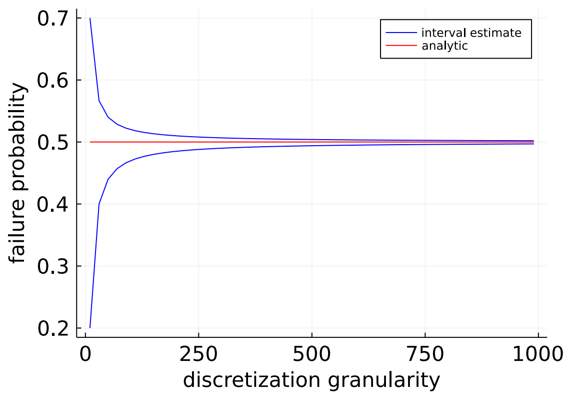

Figure 14: Lower and upper bounds (blue) on the failure probability for failure domain $\mathcal{U}_2$ and different discretization choices of the initial p-box. In red we show the analytic solution.

We consider the failure domain to be the union of two disjoint intervals at a time point $t = 4\pi$: $\mathcal{U} = [-1, -0.5] \cup [0, 0.5]$. Figure 1 shows the flowpipe and the failure domain $\mathcal{U}$. As can be seen from the outermost trajectories, the flowpipe approximation is very precise.

In Figure 12 we show flowpipes for different subsets of the initial condition, corresponding to two-sided prediction sets. The two plots respectively show the precise and imprecise initial conditions. In the latter case, we can expect the failure probability to change as well. Figure 13 further shows the computed time evolution of the p-box initial condition.

Since this system is periodic with period 2π, the probability of picking an initial state x_0 whose

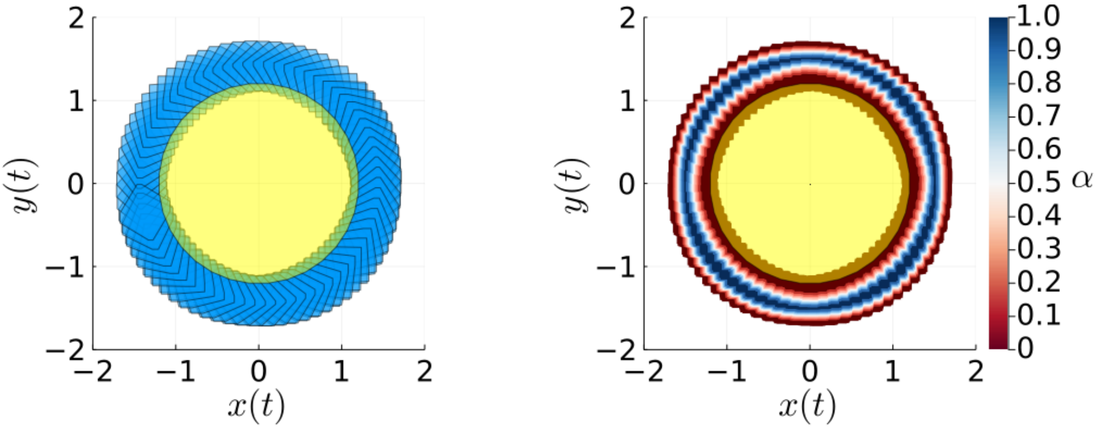

Figure 15: Left shows the flowpipe of the 2D harmonic oscillator in blue and the failure domain in yellow. Right shows the flowpipes containing α prediction sets from two initial beta distributions $X_0 \sim \text{beta}(3,3)$, $Y_0 \sim \text{beta}(8,3)$ scaled to the $\mathcal{X}_0 = [0.8, 1.2]^2$ box, and with correlation $\rho = -0.8$ between the random variables.

trajectory reaches the failure domain is identical to the probability that x_0 itself is in the failure domain (at $t = 0$). Thus we know that the failure probability is 0.5. Figure 14 shows the results with our approach. In the plot we vary the p-box discretization of the p-box (in the previous plots we have used 100 discretization levels). We see that more discretization level typically help narrowing the probability bounds.

4.2. Example 2: Harmonic oscillator

Next we consider a two-dimensional version of the oscillator model which requires a definition of stochastic dependence, as well as the marginal p-boxes. The linear harmonic oscillator is defined by the following system of differential equations:

$$\begin{cases} x'(t) = y(t) \\ y'(t) = -x(t) \end{cases} \qquad \text{subject to } (x(0), y(0)) \in \mathcal{X}_0 = [0.8, 1.2] \times [0.8, 1.2] \subset \mathbb{R}^2$$

The failure domain $\mathcal{U}$ is a sphere centered at $(0, 0)$ with radius 1.2. The temporal component of the failure domain is $\Delta_t = [8, 10]$. Figure 15 shows the flowpipe containing the entire range of the uncertainty on the left in blue, and the circular failure domain in yellow, which intersects with the flowpipe, and so this system cannot be verified to be safe in the classical reachability view.

Setting the initial condition as the distributions $X_0 \sim \text{beta}(3,3)$, $Y_0 \sim \text{beta}(8,3)$ scaled to the $\mathcal{X}_0 = [0.8, 1.2]^2$ box, and with a correlation of $\rho = -0.8$ using a Gaussian copula, we can prove

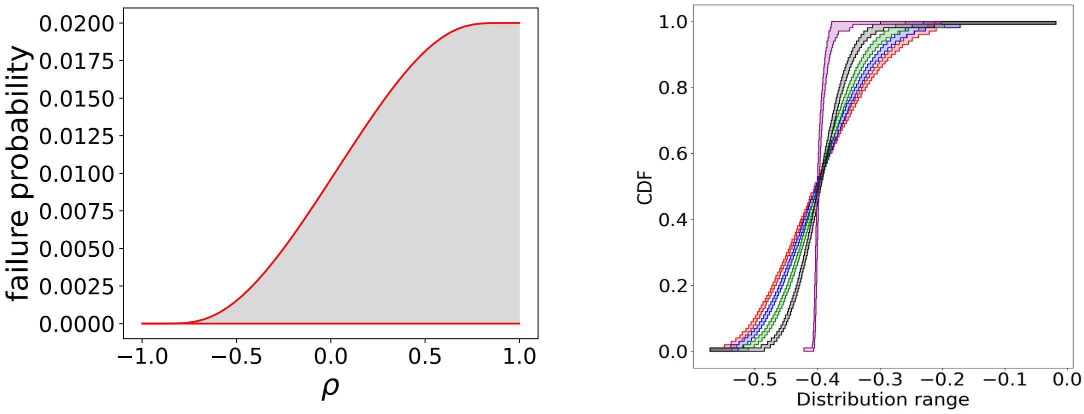

Figure 16: Left shows the change in the failure probability of the 2D harmonic oscillator as the correlation ρ is changed. Right shows the variation of the p-boxes of y at $t = 10$ as ρ is changed, with the colors corresponding to $\rho = \{-1 \text{ (red)}, -0.5 \text{ (blue)}, 0 \text{ (green)}, 0.5 \text{ (black)}, 1 \text{ (purple)}\}$

that the failure probability lies in $[0, 1.74887 \times 10^{-5}]$, using only 50 focal elements per dimension. The computation takes 25.56 seconds. The α prediction sets are shown on the right of Figure 15, where it can be seen that the majority of the probability mass is not in the failure domain.

The inclusion of correlation is important in multidimensional systems, often as important as the definition of the marginal distributions. Stochastic dependence can change the system's failure probability in an non-obvious way, and in the left of Figure 16 we show how the interval failure probability changes as ρ is varied. Setting $\rho = -1$ yields a failure probability $\mathbb{P}(\mathcal{U}) = 0$, with the upper bound increasing as correlation increases. The lower bound did not increase, and indicates that no focal elements entirely entered the failure domain, as the resolution of 50 focal elements is quite coarse. The right of Figure 16 also shows how the computed p-box for the final state of y at $t = 10\,\text{s}$ changes as the dependence varies.

4.3. Example 3: Lorenz attractor

We consider the three-dimensional nonlinear and chaotic Lorenz system, which is a simplified model for atmospheric convection. The chaotic model diverges at some point, however, in this example we show that some internal probabilistic information can be retained. The system of

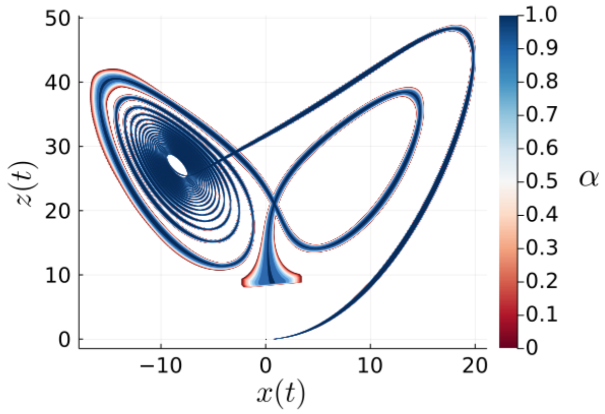

Figure 17: The α prediction sets for the Lorenz model, for uniformly distributed initial conditions. The flowpipe is shown until the time at which it begins to diverge.

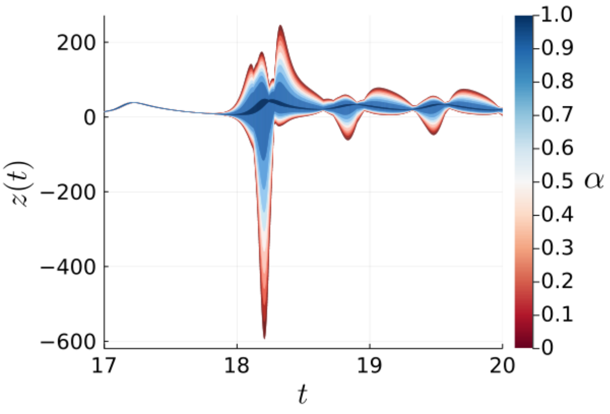

Figure 18: The α prediction sets for the z variable of the Lorenz system for time $t \in [17, 20]$ s. The range ($\alpha = 0$) of the flowpipe diverges around $t = 17.8$ s.

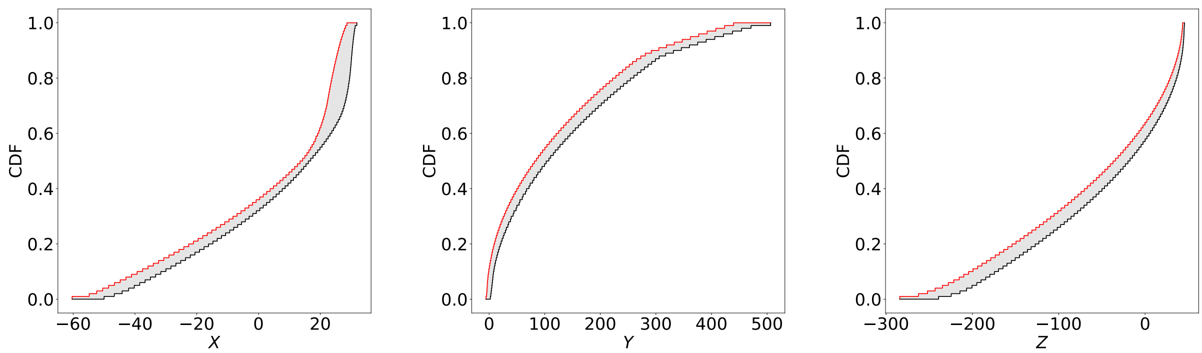

Figure 19: The p-boxes for the Lorenz attractor at $t = 18.25$ s, using 200 focal elements.

differential equations is

$$\begin{cases} x'(t) = \sigma(y - x) \\ y'(t) = x(\rho - z) - y \\ z'(t) = xy - \beta z \end{cases} \qquad \text{subject to } (x(0), y(0), z(0)) \in \mathcal{X}_0 = [0.9, 1.1] \times [-1e{-2}, 1e{-2}]^2 \subset \mathbb{R}^3,$$

where we use the parameter values $\sigma = 10, \rho = 28, \beta = 8/3$. We consider the initial states to be the uniformly distributed p-boxes $X \sim U([0.9, 1], 1.1)$, $Y \sim U(-0.01, [0, 0.01])$, and $Z \sim U(-0.01, 0.01)$. Figure 17 shows the α prediction sets in the x-z plane for time $t \in [0, 17.8]$ s, at which point the trajectories begin to diverge. Figure 18 shows the z dimension beyond this point, where the flowpipe diverges. The prediction sets show however that much of the probability of the

N	$\mathbb{P}(\mathcal{U})$	Run-time (sec)
20	$[0, 0.076297]$	125
50	$[0, 0.0338668]$	6576

Table 1: Failure probabilities for varying p-box discretization steps of the Duffing model, and the run-time (in seconds) of the corresponding computations.

probability mass does not diverge as much, and so something can be said about the distribution of trajectories within this divergent region. Figure 19 shows the p-boxes for the Lorenz system at $t = 18.25\,\text{s}$, and for example Y has a large range of $[-5.1749, 506.43]$, but the probability bound that a trajectory from the initial distribution is negative at $t = 18.25\,\text{s}$ is $\mathbb{P}_Y((-\infty, 0]) = [0, 0.095]$, that it is between 100 and 200 is $\mathbb{P}_Y([100, 200]) = [0.175, 0.255]$, and that it is greater than 300 is $\mathbb{P}_Y([300, \infty)) = [0.105, 0.125]$, which are fairly tight statements given that the system is chaotic, the calculation is verified, and we began with some imprecision in the input distribution functions. Computing the p-boxes at $t = 18.25$ took 2,023 seconds with 200^3 focal elements.

4.4. Example 4: Duffing oscillator

In the previous examples we have only considered uncertainty in the initial states. In this example we show that our framework can also model parameter uncertainty, i.e., in the coefficients defining the model, by including additional variables, one for each uncertain parameter. Consider the non-linear Duffing oscillator:

$$\begin{cases} x'(t) = y \\ y'(t) = -\alpha x - \delta y - \beta x^3 + f(t) \end{cases} \qquad \text{where } f(t) = \gamma \cos \omega t$$

We consider the case with the initial condition as the beta-distributed p-boxes $X_0 \sim \text{beta}([2,3],[3,4])$ and $Y_0 \sim \text{beta}([7,8],[2,3])$ defined by interval ranges for the traditional beta parameters, and with a correlation of $\rho = -0.8$ using a Gaussian copula. We also consider that the parameter β follows a uniform distribution $U(-0.01, 0.01)$ with constant (zero) dynamics. The numerical values of the remaining parameters are $\alpha = -1.0$, $\delta = 0.3$, $\gamma = 0.37$, and $\omega = 1.2$.

The prediction sets are shown in Figure 20 in the x-y plane, and the time horizon considered is $T = 20\pi/1.2 \approx 52.35$. We use a failure domain in the shape of a sphere centered at $(-0.3, 0.35)$ with radius 0.18 (in gray). Table 1 shows the obtained failure probability for two different choices

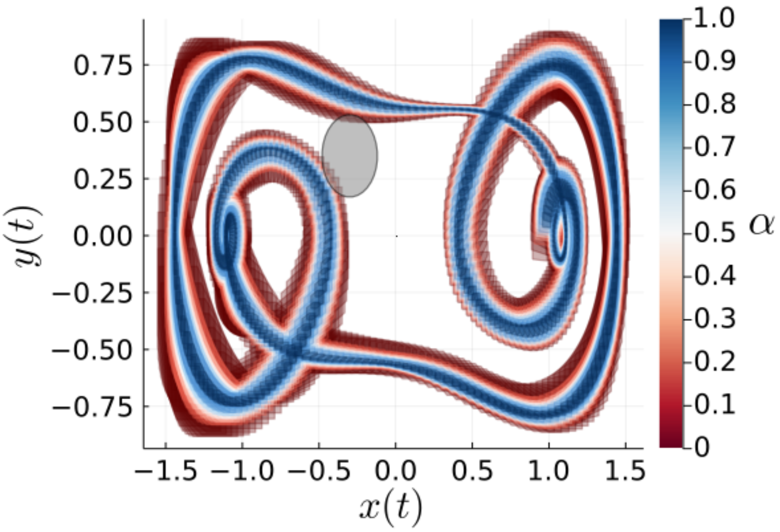

Figure 20: The α prediction sets for the Duffing oscillator. An elliptical failure domain that intersects with the flowpipe is shown in gray.

of N, the number of focal elements per dimension. It is observed from this table that the failure probability can be refined (by a factor of two) at a higher computational cost (by a factor of 50).

4.5. Example 5: Quadrotor

To investigate the scalability for higher-dimensional systems, we study the dynamics of a quadrotor as derived in [8, eq. (16) - (19)]. The variables consist of the inertial position x_1 (north) and x_2 (east), the altitude x_3, the longitudinal velocity x_4, the lateral velocity x_5, the vertical velocity x_6, the roll angle x_7, the pitch angle x_8, the yaw angle x_9, the roll rate x_{10}, the pitch rate x_{11}, and the yaw rate x_{12}. We also have the following parameters: gravity constant $g = 9.81$ [m/s^2], radius of center mass $R = 0.1$ [m], distance of rotors to center mass $l = 0.5$ [m], rotor mass $M_{rotor} = 0.1$ [kg], center mass $M = 1$ [kg], and total mass $m = M + 4M_{rotor}$. The above parameters yield the moments of inertia

$$J_x = \frac{2}{5}\, M\, R^2 + 2\, l^2\, M_{rotor}, \qquad J_y = J_x, \qquad J_z = \frac{2}{5}\, M\, R^2 + 4\, l^2\, M_{rotor}.$$

The system of differential equations is:

$$\begin{cases} x_1' = \cos(x_8)\cos(x_9)x_4 + \Big(\sin(x_7)\sin(x_8)\cos(x_9) - \cos(x_7)\sin(x_9)\Big)x_5 \\ \qquad + \Big(\cos(x_7)\sin(x_8)\cos(x_9) + \sin(x_7)\sin(x_9)\Big)x_6 \\ x_2' = \cos(x_8)\sin(x_9)x_4 + \Big(\sin(x_7)\sin(x_8)\sin(x_9) + \cos(x_7)\cos(x_9)\Big)x_5 \\ \qquad + \Big(\cos(x_7)\sin(x_8)\sin(x_9) - \sin(x_7)\cos(x_9)\Big)x_6 \\ x_3' = \sin(x_8)x_4 - \sin(x_7)\cos(x_8)x_5 - \cos(x_7)\cos(x_8)x_6 \\ x_4' = x_{12}x_5 - x_{11}x_6 - g\sin(x_8) \\ x_5' = x_{10}x_6 - x_{12}x_4 + g\cos(x_8)\sin(x_7) \\ x_6' = x_{11}x_4 - x_{10}x_5 + g\cos(x_8)\cos(x_7) - \dfrac{F}{m} \\ x_7' = x_{10} + \sin(x_7)\tan(x_8)x_{11} + \cos(x_7)\tan(x_8)x_{12} \\ x_8' = \cos(x_7)x_{11} - \sin(x_7)x_{12} \\ x_9' = \dfrac{\sin(x_7)}{\cos(x_8)}x_{11} + \dfrac{\cos(x_7)}{\cos(x_8)}x_{12} \\ x_{10}' = \dfrac{J_y - J_z}{J_x}x_{11}x_{12} + \dfrac{1}{J_x}\tau_\phi \\ x_{11}' = \dfrac{J_z - J_x}{J_y}x_{10}x_{12} + \dfrac{1}{J_y}\tau_\theta \\ x_{12}' = \dfrac{J_x - J_y}{J_z}x_{10}x_{11} + \dfrac{1}{J_z}\tau_\psi \end{cases}$$

The quadrotor is stabilized with PD controllers for height, roll, and pitch. The inputs to these controllers are the desired values u_1, u_2, and u_3, respectively. The equations of the controllers are

$$\begin{aligned} F &= m\,g - 10(x_3 - u_1) + 3x_6 && \text{(height control)}, \\ \tau_\phi &= -(x_7 - u_2) - x_{10} && \text{(roll control)}, \\ \tau_\theta &= -(x_8 - u_3) - x_{11} && \text{(pitch control)}. \end{aligned}$$

We leave the heading uncontrolled by setting $\tau_\psi = 0$.

The specification consists of three properties concerning the altitude (dimension x_3), each with an individual failure domain. The first property ($\mathcal{U}_1$) is that the quadrotor stays below 1.4 m for 5 seconds. The second property ($\mathcal{U}_2$) is that the quadrotor stays above 0.9 m after 1 second. The third property ($\mathcal{U}_3$) is that the quadrotor is lifted from the reference altitude of 0 m to an altitude of 1 m ± 0.02 m after 5 seconds.

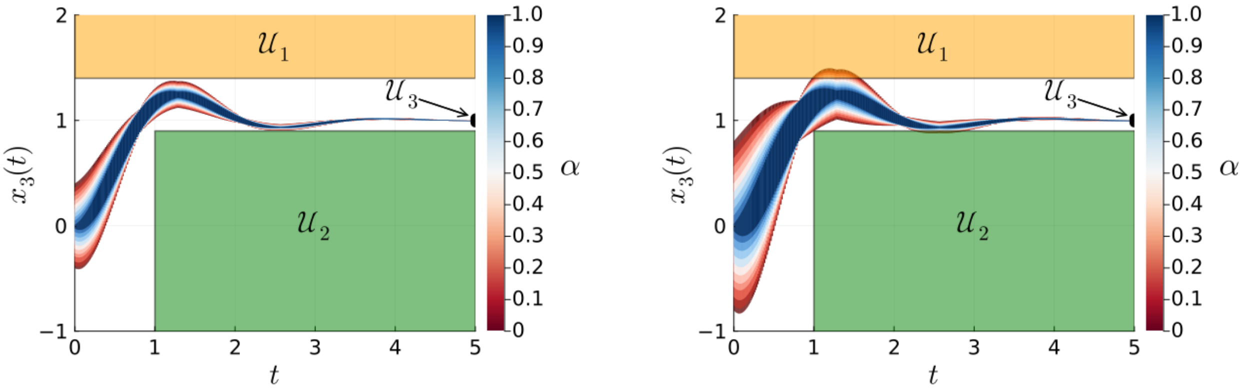

Figure 21: The α prediction sets for the two considered scenarios. Left shows the scenario with $W = 0.4$. Right shows the scenario with $W = 0.8$.

	$\mathcal{U}_1$	$\mathcal{U}_2$	$\mathcal{U}_3$	$\mathcal{U}_1 \wedge \mathcal{U}_2 \wedge \mathcal{U}_3$
	$t \in [0, 5] : x_3(t) \leqslant 1.4$	$t \in [1, 5] : x_3(t) \geqslant 0.9$	$x_3(5) \in [0.98, 1.02]$	
$W = 0.4$	$[0, 0]$	$[0, 0]$	$[0, 0]$	$[0, 0]$
$W = 0.8$	$[0.03, 0.25]$	$[0.13, 0.22]$	$[0, 0]$	$[0.13, 0.25]$

Table 2: Failure probabilities of the quadrotor model for each of the three failure domains ($\mathcal{U}_1$,$\mathcal{U}_2$,$\mathcal{U}_3$) and their combination, and for each scenario ($W = 0.4$, $W = 0.8$).

The initial value for the position and velocities (i.e., from x_1 to x_6) is unknown and given by $[-W, W]$ m, and we consider two scenarios with $W \in \{0.4, 0.8\}$ taken from a reachability competition [33]. The initial value of all other variables x_7 to x_{12} are unknown and given by $[0, 0.1]$. Here we assume a uniform distribution of the initial states in dimension x_3 (altitude) as $U(-W, W)$.

Figure 21 shows the α prediction sets for the two scenarios, which gives some intuition about the model behavior. The quadrotor approaches the target altitude of 1 m, but it overshoots this threshold and then stabilizes by oscillating around it. We compute the failure probability with our algorithm. The results are summarized in Table 2, using 100 focal elements. We can see that the first scenario ($W = 0.4$) satisfies all three properties and our failure-probability estimates are tight, as expected. On the contrary, the second scenario violates both the upper-bound and the lower-bound property, which is reflected in the obtained failure-probability estimates. We also investigate the combination of all three failure domains. Note that they each have different time

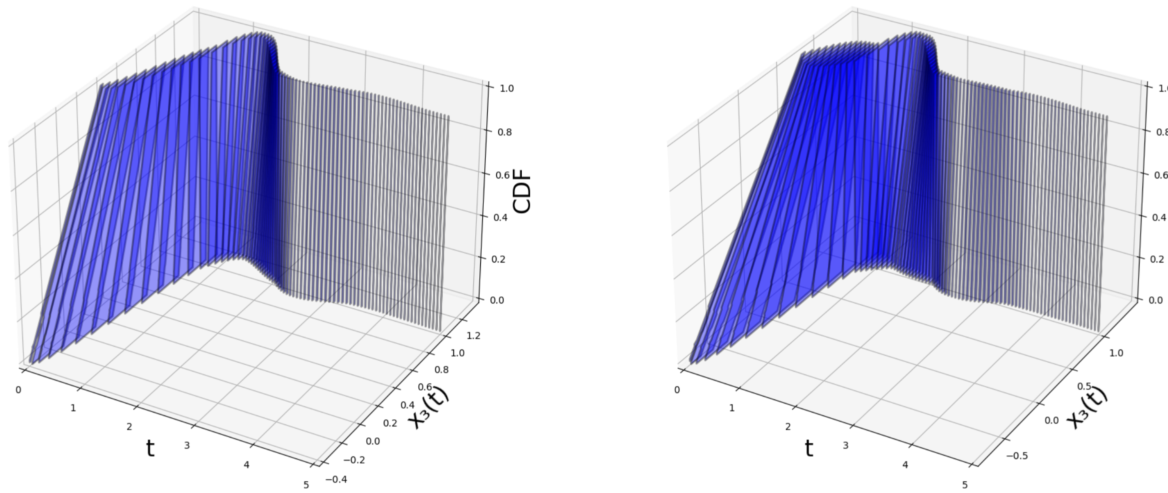

Figure 22: Propagation of the initial p-box along coordinate $x_3(t)$ (altitude) through the nonlinear dynamics of the quadrotor. Left shows the scenario with $W = 0.4$. Right shows the scenario with $W = 0.8$.

components, so Algorithm 1 does not directly apply, but the generalization is simple (essentially adding another loop over the individual failure domains). As we can see in Table 2, the failure probability of the combination in this case is just the tightest possible failure probability given the individual properties (here: for $\mathcal{U}_1$ and $\mathcal{U}_2$, which have non-zero failure probability). This is explained from the fact that the trajectories leading to failure domain $\mathcal{U}_1$, i.e., overshooting the height of 1.4 m, also lead to failure domain $\mathcal{U}_2$, i.e., undershooting the height of 0.9 m. That conclusion is also supported by Figure 21.

Finally, the propagation of the p-box projected along $x_3(t)$ is shown in Figure 22 for each scenario. The transition from a precise distribution to a p-box is observed in both scenarios, since x_3 interacts with the other variables (which are intervals); for the second scenario ($W = 0.8$), this transition is more prominent.

5. Conclusions

In this article we described a general algorithm to compute rigorous bounds on the failure probability of a multivariate, nonlinear dynamical system with a space- and time-dependent failure domain, and also extended the algorithm to percentile sets. Our algorithm makes the following mild assumptions. For the initial states, we assume a distribution represented as a p-box, which can express both precise and imprecise probabilities. For the dynamical system, we require knowledge about the system of differential equations and applicability of a reachability algorithm based on Taylor models; a common sufficient condition for the latter is Lipschitz continuity. Thus our algorithm is applicable in general and realistic settings where precise distributions are not available.

A key idea of our algorithm is to exploit that the Taylor-model representation preserves dependencies between inputs and outputs (a similar observation was also made for another set representation in [44]). Intuitively, the information of each focal element can be "pushed" along the flowpipe without the need to recompute one flowpipe per focal element, which greatly reduces the computational complexity of our algorithm.

One could alternatively use p-box arithmetic [27] by performing convolutions for each operation in the Taylor polynomials. Although this would not have exponential complexity, it would be sub-optimal due to the dependence problem experienced in p-box arithmetic. That is, like for intervals, even if dependencies are known at the beginning of an expression, as calculations are performed, dependence information is lost and must be over-approximated, leading to often overly conservative bounds. Although p-box arithmetic does allow for p-boxes to be evaluated without dependence (copula) assumptions, due to repeated instances of the same variable, the computed bounds are not the best possible [45]. This paper does not explore partially known or unknown copulas, only precise dependencies, and extending the current methodology to the imprecise case is of interest.

In summary, reachability analysis traditionally gives only qualitative results, which in many applications is unrealistic because failure cannot be completely excluded. Instead of proving absence of errors, in safety-critical applications it is required to show that the probability of failure is below a certain threshold. We thus believe that our work allows for wider application of reachability analysis.

We add one small caveat about the rigor of the presented method. While automatically verified in both reachability and probability, the method relies on the ability to evaluate the copula C exactly. Although this is possible for some copulas such as independence, perfect and opposite (in 2 dimensions) associations, and some parametric families such as the Frank and Clayton copulas, to our knowledge there is no exact or rigorous algorithm to evaluate the Gaussian copula. In this work, we used the algorithm for the multivariate Gaussian cdf by Genz [32], which, although accurate, is not rigorous or even deterministic. This therefore leads to slight variations of computed failure probability intervals when using this copula family.

For future work, we see several directions. A direct improvement of our algorithm is to use a more sophisticated, hierarchical over-approximation of the p-box. Instead of over-approximating the p-box into all focal elements directly, we could first perform computations with a coarser granularity. Only if we find that the corresponding reach set is intersecting with but not included in the failure domain, do we have to refine. This idea is related to a technique called *counterexample-guided abstraction refinement* (CEGAR) in the reachability community [17].

A challenging future direction is to extend the system dynamics to additional stochasticity (e.g., SDEs or RODEs), which may require a synthesis with orthogonal approaches [25]. Another direction is to extend our approach to control systems, where in addition to the dynamical system there is a digital controller [20]. The corresponding reachability problem has been studied successfully with Taylor-model techniques, e.g., for neural-network controllers [41, 61, 40].

Acknowledgements

The authors acknowledge discussions with David P. Sanders at initial stages of this work. The authors also thank the attendees of the virtual Fuzzy and Possibility interest group (FnPig), at the University of Liverpool, for the discussions which gave inspiration to the consonant p-box approximation in this paper, in particular Dominik Hose, Marco de Angelis, and Alexander Wimbush.

M. Forets has been partly supported by the *Agencia Nacional de Investigación e Innovación*, Uruguay. C. Schilling acknowledges support in part by DIREC - Digital Research Centre Denmark and the Villum Investigator Grant S4OS. L. Benet acknowledges funding by UNAM–PAPIIT grant number IG-101122. A. Gray has been part-funded by the EPSRC Energy Programme [grant number EP/W006839/1]. A. Gray and S. Ferson were partly funded by the Engineering & Physical Sciences Research Council (EPSRC), United Kingdom with grant no. EP/R006768/1.

References

[1] Althoff, M., Frehse, G., & Girard, A. (2020). Set propagation techniques for reachability analysis. *Annu. Rev. Control Robot. Auton. Syst.*, *4*. doi:10.1146/annurev-control-071420-081941.

[2] de Angelis, M., Patelli, E., & Beer, M. (2015). Advanced line sampling for efficient robust reliability analysis. *Structural Safety*, *52*, 170–182. doi:10.1016/j.strusafe.2014.10.002.

[3] Arnold, L. (1974). *Stochastic differential equations: Theory and applications*. Wiley Interscience.

[4] Au, S.-K., & Beck, J. L. (2001). Estimation of small failure probabilities in high dimensions by subset simulation. *Probabilistic Engineering Mechanics*, *16*, 263–277. doi:10.1016/S0266-8920(01)00019-4.

[5] Balch, M. S. (2012). Mathematical foundations for a theory of confidence structures. *Int. J. Approx. Reason.*, *53*, 1003–1019. doi:10.1016/j.ijar.2012.05.006.

[6] Balch, M. S., Martin, R., & Ferson, S. (2019). Satellite conjunction analysis and the false confidence theorem. *Proc. R. Soc. A*, *475*. doi:10.1098/rspa.2018.0565.

[7] Baudrit, C., Dubois, D., & Guyonnet, D. (2006). Joint propagation and exploitation of probabilistic and possibilistic information in risk assessment. *IEEE Trans. Fuzzy Syst.*, *14*, 593–608. doi:10.1109/TFUZZ.2006.876720.

[8] Beard, R. (2008). *Quadrotor Dynamics and Control Rev 0.1*. Technical Report Brigham Young University.

[9] Benet, L., Forets, M., Sanders, D. P., & Schilling, C. (2019). TaylorModels.jl: Taylor models in Julia and its application to validated solutions of ODEs. In *SWIM*. URL: https://swim2019.ensta-paristech.fr/.

[10] Benet, L., & Sanders, D. P. (2019). TaylorSeries.jl: Taylor expansions in one and several variables in Julia. *Journal of Open Source Software*, *4*. doi:10.21105/joss.01043.

[11] Berz, M. (1989). Differential algebraic description of beam dynamics to very high orders. *Part. Accel.*, *24*, 109–124.

[12] Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. (2017). Julia: A fresh approach to numerical computing. *SIAM Rev.*, *59*, 65–98. doi:10.1137/141000671.

[13] Birnbaum, A. (1961). Confidence curves: An omnibus technique for estimation and testing statistical hypotheses. *J. Am. Stat. Assoc.*, *56*, 246–249. doi:10.2307/2282249.

[14] Bogomolov, S., Forets, M., Frehse, G., Potomkin, K., & Schilling, C. (2019). JuliaReach: a toolbox for set-based reachability. In *HSCC* (pp. 39–44). ACM. doi:10.1145/3302504.3311804.

[15] Bouissou, O., Goubault, E., Goubault-Larrecq, J., & Putot, S. (2012). A generalization of p-boxes to affine arithmetic. *Computing*, *94*, 189–201. doi:10.1007/s00607-011-0182-8.

[16] Chen, X. (2015). *Reachability analysis of non-linear hybrid systems using Taylor models*. Ph.D. thesis RWTH Aachen University, Germany. URL: http://publications.rwth-aachen.de/record/465295.

[17] Clarke, E. M., Grumberg, O., Jha, S., Lu, Y., & Veith, H. (2000). Counterexample-guided abstraction refinement. In *CAV* (pp. 154–169). Springer volume 1855 of *LNCS*. doi:10.1007/10722167_15.

[18] Crowther, R. (2002). Orbital debris: a growing threat to space operations. *Phil. Trans. R. Soc. A.*, *361*, 157–168. doi:10.1098/rsta.2002.1118.

[19] Dempster, A. P. (2008). Upper and lower probabilities induced by a multivalued mapping. In *Classic Works of the Dempster-Shafer Theory of Belief Functions* (pp. 57–72). Springer volume 219 of *Studies in Fuzziness and Soft Computing*. doi:10.1007/978-3-540-44792-4_3.

[20] Doyle, J. C., Francis, B. A., & Tannenbaum, A. R. (2013). *Feedback control theory*. Dover Publications.

[21] Dubois, D., & Prade, H. (1988). *Possibility Theory - An Approach to Computerized Processing of Uncertainty*.

Springer.

[22] Dubois, D., & Prade, H. (1990). Consonant approximations of belief functions. *Int. J. Approx. Reason.*, *4*, 419–449. doi:10.1016/0888-613X(90)90015-T.

[23] Dubois, D., Prade, H., & Sandri, S. (1993). On possibility/probability transformations. In *Fuzzy logic* (pp. 103–112). Springer. doi:10.1007/978-94-011-2014-2_10.

[24] Enszer, J. A., Lin, Y., Ferson, S., Corliss, G. F., & Stadtherr, M. A. (2008). Propagating uncertainties in modeling nonlinear dynamic systems. In *REC* (pp. 89–105).

[25] Feng, S., Chen, M., Xue, B., Sankaranarayanan, S., & Zhan, N. (2020). Unbounded-time safety verification of stochastic differential dynamics. In *CAV* (pp. 327–348). Springer volume 12225 of *LNCS*. doi:10.1007/978-3-030-53291-8_18.

[26] Ferson, S., Ginzburg, L., & Akçakaya, R. (1996). Whereof one cannot speak: when input distributions are unknown. *Risk Analysis*, .

[27] Ferson, S., & Hajagos, J. G. (2004). Arithmetic with uncertain numbers: rigorous and (often) best possible answers. *Reliab. Eng. Syst. Saf.*, *85*, 135–152. doi:10.1016/j.ress.2004.03.008.

[28] Ferson, S., Kreinovich, V., Grinzburg, L., Myers, D., & Sentz, K. (2003). *Constructing probability boxes and Dempster-Shafer structures*. Technical Report Sandia National Lab., Albuquerque, NM (United States).

[29] Forets, M., & Schilling, C. (2021). LazySets.jl: Scalable symbolic-numeric set computations. *Proceedings of the JuliaCon Conferences*, *1*. doi:10.21105/jcon.00097.

[30] Frank, M. J., Nelsen, R. B., & Schweizer, B. (1987). Best-possible bounds for the distribution of a sum—a problem of Kolmogorov. *Probability theory and related fields*, *74*, 199–211. doi:10.1007/BF00569989.

[31] Gao, S., Avigad, J., & Clarke, E. M. (2012). Delta-decidability over the reals. In *LICS* (pp. 305–314). IEEE Computer Society. doi:10.1109/LICS.2012.41.

[32] Genz, A. (2004). Numerical computation of rectangular bivariate and trivariate normal and t probabilities. *Stat. Comput.*, *14*, 251–260. doi:10.1023/B:STCO.0000035304.20635.31.

[33] Geretti, L., Sandretto, J. A. D., Althoff, M., Benet, L., Chapoutot, A., Collins, P., Duggirala, P. S., Forets, M., Kim, E., Linares, U., Sanders, D. P., Schilling, C., & Wetzlinger, M. (2021). ARCH-COMP21 category report: Continuous and hybrid systems with nonlinear dynamics. In *ARCH* (pp. 32–54). EasyChair volume 80 of *EPiC Series in Computing*. doi:10.29007/2jw8.

[34] Gray, A., Ferson, S., & Patelli, E. (2022). ProbabilityBoundsAnalysis.jl: Arithmetic with sets of distributions. (submitted).

[35] Gray, A., Forets, M., Schilling, C., Benet, L., & Ferson, S. (2022). Rigorous time evolution of p-boxes in non-linear ODEs. In *ESREL* (pp. 154–155).

[36] Hainry, E. (2008). Reachability in linear dynamical systems. In *CiE* (pp. 241–250). Springer volume 5028 of *LNCS*. doi:10.1007/978-3-540-69407-6_28.

[37] Hose, D., & Hanss, M. (2019). Possibilistic calculus as a conservative counterpart to probabilistic calculus. *Mechanical Systems and Signal Processing*, *133*. doi:10.1016/j.ymssp.2019.106290.

[38] Hose, D., & Hanss, M. (2021). A universal approach to imprecise probabilities in possibility theory. *Int. J. Approx. Reason.*, *133*, 133–158. doi:10.1016/j.ijar.2021.03.010.

[39] Huang, C., Chen, X., Lin, W., Yang, Z., & Li, X. (2017). Probabilistic safety verification of stochastic hybrid systems using barrier certificates. *ACM Trans. Embed. Comput. Syst.*, *16*. doi:10.1145/3126508.

[40] Huang, C., Fan, J., Chen, X., Li, W., & Zhu, Q. (2022). POLAR: A polynomial arithmetic framework for verifying neural-network controlled systems. In *ATVA* (pp. 414–430). Springer volume 13505 of *LNCS*. doi:10.1007/978-3-031-19992-9_27.

[41] Ivanov, R., Carpenter, T. J., Weimer, J., Alur, R., Pappas, G. J., & Lee, I. (2021). Verisig 2.0: Verification of neural network controllers using taylor model preconditioning. In *CAV* (pp. 249–262). Springer volume 12759 of *LNCS*. doi:10.1007/978-3-030-81685-8_11.

[42] Johnson, N. L. (2010). *Orbital Debris: the Growing Threat to Space Operations*. Technical Report NASA.

[43] Joldes, M. (2011). *Rigorous Polynomial Approximations and Applications.*. Ph.D. thesis École normale supérieure de Lyon, France. URL: https://tel.archives-ouvertes.fr/tel-00657843.

[44] Kochdumper, N., Schürmann, B., & Althoff, M. (2020). Utilizing dependencies to obtain subsets of reachable sets. In *HSCC*. ACM. doi:10.1145/3365365.3382192.

[45] Kreinovich, V., & Ferson, S. (2006). Computing best-possible bounds for the distribution of a sum of several variables is NP-hard. *Int. J. Approx. Reason.*, *41*, 331–342. doi:10.1016/j.ijar.2005.06.009.

[46] Legay, A., Delahaye, B., & Bensalem, S. (2010). Statistical model checking: An overview. In *RV* (pp. 122–135). Springer volume 6418 of *LNCS*. doi:10.1007/978-3-642-16612-9_11.

[47] Levi, I. (1983). *The enterprise of knowledge: An essay on knowledge, credal probability, and chance*. MIT press.

[48] Liu, C., & Martin, R. (2020). Inferential models and possibility measures. doi:10.48550/arXiv.2008.06874.

[49] Maces, D. A. (2013). *Uncertainty propagation in models for dynamic nonlinear systems: Methods and applications*. Ph.D. thesis University of Notre Dame. doi:10.7274/8p58pc30b0t.

[50] Makarov, G. (1982). Estimates for the distribution function of a sum of two random variables when the marginal distributions are fixed. *Theory of Probability & its Applications*, *26*, 803–806. doi:10.1137/1126086.

[51] Makino, K. (1998). *Rigorous analysis of nonlinear motion in particle accelerators*. Ph.D. thesis Michigan State University. doi:10.25335/M5F18SJ1N.

[52] Makino, K., & Berz, M. (2003). Taylor models and other validated functional inclusion methods. *International Journal of Pure and Applied Mathematics*, *6*, 239–316.

[53] Martin, R. (2019). False confidence, non-additive beliefs, and valid statistical inference. *Int. J. Approx. Reason.*, *113*, 39–73. doi:10.1016/j.ijar.2019.06.005.

[54] Martin, R., & Liu, C. (2015). *Inferential Models: Reasoning with Uncertainty*. Chapman and Hall/CRC. doi:10.1201/b19269.

[55] Milani, A., & Gronchi, G. (2009). *Theory of Orbit Determination*. Cambridge University Press. doi:10.1017/cbo9781139175371.

[56] Molchanov, I. (2005). *Theory of random sets*. Springer. doi:10.1007/978-1-4471-7349-6.

[57] Montes, I., Miranda, E., Pelessoni, R., & Vicig, P. (2015). Sklar's theorem in an imprecise setting. *Fuzzy Sets Syst.*, *278*, 48–66. doi:10.1016/j.fss.2014.10.007.

[58] Nelsen, R. B. (2007). *An introduction to copulas*. Springer Science & Business Media. doi:10.1007/0-387-28678-0.

[59] Protter, P. E. (2005). Stochastic differential equations. In *Stochastic integration and differential equations* (pp. 249–361). Springer. doi:10.1007/978-3-662-10061-5_6.

[60] Sandretto, J. A. D. (2021). Confidence-based contractor, propagation and potential clouds for differential equations. *Acta Cybern.*, *25*, 49–68. doi:10.14232/actacyb.285177.

[61] Schilling, C., Forets, M., & Guadalupe, S. (2022). Verification of neural-network control systems by integrating Taylor models and zonotopes. In *AAAI* (pp. 8169–8177). AAAI Press. doi:10.1609/aaai.v36i7.20790.

[62] Schweizer, B., & Sklar, A. (2011). *Probabilistic metric spaces*. Courier Corporation.

[63] Serra, R., Arzelier, D., Joldes, M., & Rondepierre, A. (2015). Probabilistic collision avoidance for long-term space encounters via risk selection. In *Advances in Aerospace Guidance, Navigation and Control* (pp. 679–698). Springer. doi:10.1007/978-3-319-17518-8_39.

[64] Shafer, G. (1976). *A mathematical theory of evidence*. Princeton university press.

[65] Shao, J. (2003). *Mathematical statistics*. Springer. doi:10.1007/b97553.

[66] Shmarov, F., & Zuliani, P. (2015). ProbReach: verified probabilistic delta-reachability for stochastic hybrid systems. In *HSCC* (pp. 134–139). ACM. doi:10.1145/2728606.2728625.

[67] Strand, J. (1970). Random ordinary differential equations. *Journal of Differential Equations*, *7*, 538–553. doi:10.1016/0022-0396(70)90100-2.

[68] Tardioli, C., Farnocchia, D., Vasile, M., & Chesley, S. R. (2020). Impact probability under aleatory and epistemic uncertainties. *Celestial Mechanics and Dynamical Astronomy*, *132*, 1–12. doi:10.1007/s10569-020-09991-3.

[69] Troffaes, M. C., & De Cooman, G. (2014). *Lower previsions*. John Wiley & Sons.

[70] Williamson, R. C., & Downs, T. (1990). Probabilistic arithmetic. I. Numerical methods for calculating convolutions and dependency bounds. *Int. J. Approx. Reason.*, *4*, 89–158. doi:10.1016/0888-613X(90)90022-T.

[71] Xue, B., Fränzle, M., Zhan, N., Bogomolov, S., & Xia, B. (2020). Safety verification for random ordinary differential equations. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, *39*, 4090–4101. doi:10.1109/TCAD.2020.3013135.

[72] Yager, R. R. (1986). Arithmetic and other operations on Dempster-Shafer structures. *Int. J. Man Mach. Stud.*, *25*, 357–366. doi:10.1016/S0020-7373(86)80066-9.

[73] Zadeh, L. A. (1978). Fuzzy sets as a basis for a theory of possibility. *Fuzzy sets and systems*, *1*, 3–28. doi:10.1016/0165-0114(78)90029-5.