

Robust Probabilistic Model Checking with Continuous Reward Domains

Xiaotong Ji*, Hanchun Wang[†], Antonio Filieri*, Ilenia Epifani[‡]

*Department of Computing, Imperial College London

[†]Department of Mathematics, Imperial College London

[‡]Dipartimento di Matematica, Politecnico di Milano

Abstract—Probabilistic model checking traditionally verifies properties on the expected value of a measure of interest. This restriction may fail to capture the quality of service of a significant proportion of a system’s runs, especially when the probability distribution of the measure of interest is poorly represented by its expected value due to heavy-tail behaviors or multiple modalities. Recent works inspired by distributional reinforcement learning use discrete histograms to approximate integer reward distribution, but they struggle with continuous reward space and present challenges in balancing accuracy and scalability.

We propose a novel method for handling both continuous and discrete reward distributions in Discrete Time Markov Chains using moment matching with Erlang mixtures. By analytically deriving higher-order moments through Moment Generating Functions, our method approximates the reward distribution with theoretically bounded error while preserving the statistical properties of the true distribution. This detailed distributional insight enables the formulation and robust model checking of quality properties based on the entire reward distribution function, rather than restricting to its expected value. We include a theoretical foundation ensuring bounded approximation errors, along with an experimental evaluation demonstrating our method’s accuracy and scalability in practical model-checking problems.


I. INTRODUCTION

Probabilistic model checking (PMC) [29] plays a critical role in many self-adaptation methods, allowing the detection of quantitative requirements violations and supporting the planning of appropriate adaptation actions [7], [15], [32], [46].


However, most PMC methods allow to reason only about the expected value of a system property of interest, which may neglect the variability in the distribution of possible outcomes. For example, even if the average execution time of the system is less than 100ms, a large number of runs may exceed such expected value. Not accounting for the variability in a measure of interest’s distribution may provide a partial assessment of a system’s performance, in turn limiting the awareness of the adaptation logic and the robustness of its decisions.

Recent advancements in PMC [2], [12], also inspired by related results in reinforcement learning [5], [11] have incorporated distributional perspectives in the analysis of models abstracted as Discrete-Time Markov Chains (DTMCs) and Markov Decision Process (MDPs) by approximating probability or reward distributions using histograms. However, the precision of histogram approximations is constrained by the

setting of bins, with fewer bins leading to coarser representations. Furthermore, the resulting piecewise constant approximations fail to capture subtle features such as multimodality or skewness, and their applicability is restricted to integer reward scheme to ensure convergence. These methods also face scalability challenges, particularly in large or continuous reward spaces, where sparse data worsens the inaccuracy of the discretization process.



(a) UAV’s Flight Process DTMC



(b) Cumulative reward Distribution from s_0

Fig. 1. UAV flight process with states $\{s_0, s_1, s_2, s_3, s_4\}$ representing Takeoff, Stable Flight, Maneuvering, Descent, and Landing, respectively. Transition probabilities and rewards are shown on the edges in the form of p/r . The cumulative reward distribution at s_0 is illustrated using the expected value (grey), probability histogram estimation (light blue), and true simulated distribution points (green), highlighting the different levels of information provided by various methods on the same distribution.

Example. Consider a UAV’s flight process modeled as the DTMC shown in Figure 1(a), where each state corresponds to a specific flight phase – Takeoff, Stable Flight, Maneuvering, Descent, and Landing. State transitions are governed by probabilities, while rewards represent energy consumption during these transitions. For example, transitioning from stable flight to maneuvering consumes more energy, whereas landing after descent requires minimal energy. This framework allows for evaluating the UAV’s operational efficiency across various scenarios. Standard PMC only considers the expected value (the grey line in Figure 1(b)) of the cumulative energy consumption, which provides limited information on the energy consumption pattern. In particular, it falls on the right-hand side of the head modality and far from the tail modality of the distribution. Recent categorical histogram-based methods [5] estimate the distribution of cumulative rewards, but they struggle to accurately capture the true distribution, especially in cases involving closely packed peaks in dense regions or isolated peaks in sparse regions. In dense regions, histograms tend to smooth out critical individual peaks, overestimating probabilities, while in sparse regions, they underestimate the significance of isolated high-reward transitions. Although quantile-based histogram methods [11] allow for more flexible binning to address the dense-sparse issue, they fail to capture peak and tail behaviors effectively. These limitations highlight the need for more precise and flexible methods to approximate the true distribution of the measure of interest.

In this work, we propose a distribution approximation method using moment matching on Erlang mixtures to efficiently approximate the full distributional profile of probability and reward distribution over DTMCs, enabling verification of robust properties based on the complete distribution rather than expected values. The Erlang mixture is chosen for its strong capability to approximate all positive random variables with any arbitrary level of accuracy [42] and its natural connection to Markov models, facilitating a detailed analysis of Markov reward distributions to support the verification of quantitative requirements. The approximated distribution addresses the challenges of capturing multi-modality and precise behaviors faced by histogram approximations, without relying on iterative forward computations or simulations and supporting continuous reward spaces without requiring discretization, while preserving statistical properties, such as skewness and kurtosis, through precise moment matching. Our approach enhances the robustness of PMC by capturing a broader spectrum of possible outcomes, thus improving system behavior verification under uncertainty. Our key contributions are as follows:

- 1) **Moment Matching Distribution Approximation:** We introduce a novel approach to approximate the cumulative reward distribution of stochastic systems using mixtures of Erlang distributions matched to its moments, providing a precise characterization of the full reward distribution with bounded error.
- 2) **Robust Property Model Checking:** We introduce a

procedure to verify *chance requirements* of the form $Pr(R \leq r^*) \geq \alpha$ on a measure of interest, which predicate on the probability of the threshold r^* being violated by a run of the system, rather than only constraining the measure’s expected value.

II. BACKGROUND

A. Probabilistic Model Checking

Probabilistic model checking is a formal method to reason about quantitative properties of systems, whose behavior is abstracted via convenient stochastic models [29]. In this work, we focus on discrete Markov models and on the verification of regular properties [3]. Regular properties can be reduced to predicates on the probability of reaching a target state or on the cumulative reward until reaching a target state. For example, the co-safe fragment of linear time logic (LTL) is used to define sequential plans for a robot [30] or probabilistic computation tree logic (PCTL) properties without nesting of the modal operators [3]. Without loss of generality, we will focus on properties of the forms $P_{\bowtie p}[\diamond\phi]$ and $R_{\bowtie r}[\diamond\phi]$, where $\bowtie \in \{\leq, \geq\}$, and $p \in [0, 1]$ and reward $r \in [0, \infty)$ represent bounds on the probability of eventually reaching a state where the Boolean predicate ϕ holds.

In its standard semantics [4], a formula $P_{\bowtie p}[\diamond\phi]$ holds in a state s if the *expected probability* of eventually reaching a state satisfying ϕ is $\bowtie p$. Similarly, $R_{\bowtie r}[\diamond\phi]$ is evaluated by constraining the expected cumulative reward until a state satisfying ϕ is reached. This formulation takes advantage of the linearity of the expected value to reduce the probabilistic model checking problem to the solution of a linear system of equations for DTMCs, while for MDPs [39] to efficient dynamic programming using the standard Bellman operator [3].

However, the expected value may not tell the full story about the system being analyzed [12]. For example, the dispersion of the distribution, its skewness, or the presence of multiple modalities may render the expected value a poor representative of the system’s behavior.

Example. For the UAV system in Figure 1, the expected cumulative reward is 7.51, which is rarely achieved across runs and does not capture the bi-modal behavior of the system.

Chance constraints. In this work, we aim to reinterpreting $P_{\bowtie p}(\cdot)$ and $R_{\bowtie r}(\cdot)$ as evaluations of the cumulative distribution function for the reachability and reward distributions. This enables verifying requirements of the form “a reward larger than r is obtained with at least probability c ”. Notice that 1) reasoning about the probability of reaching a target state (i.e., a state satisfying a desired property ϕ) can be reduced to reasoning about the cumulative rewards to reaching a target state via systematic model transformations [16] and, similarly, 2) the paths reaching a target state can be reduced to paths to absorption into a designated absorbing state, again via the identification of connected components and systematic model transformations [3]. Therefore, to simplify the presentation and without loss of generality, the remainder of this paper will focus on approximating and reasoning

about the probability distribution of the cumulative reward to absorption. More precisely, we adapt the definition of chance-constrained requirements from the robust control literature [14] and propose an approximate decision procedure, with bounded error, for requirements of the form $Pr(X \leq r^*) \geq \alpha$, where X represents the distribution of cumulative rewards to absorption over the possible runs of the Markov process, r^* is a threshold value, and α is the required probability level. Notably, chance-constrained requirements are also more general than the commonly used (Conditional)Value at Risk requirements used in previous literature [12].

B. Markov Reward Process

Discrete-Time Markov Chain (DTMC). A Discrete-Time Markov Chain is defined as a tuple $D = (S, s_0, P, r, AP, L)$, where S is a finite state space, $s_0 \in S$ is the initial state, $P: S \times S \rightarrow [0, 1]$ is the stochastic matrix representing the transition probabilities (such that $p_{ij} \geq 0$ for all $p_{ij} \in P$ and $\sum_j p_{ij} = 1$ for all i), and $r: S \rightarrow \mathbb{R}_{\geq 0}$ is a reward function that assigns a reward to each state¹, AP is a set of atomic propositions and $L: S \rightarrow 2^{AP}$ is a labeling function mapping each state to a set of propositions that hold true in that state. This stochastic system evolves as a sequence of random variables $\{X_n\}_{n \geq 0}$, where each X_n represents the state of the system at the n -th time step.

Notably, a specific *induced DTMC* $M_D = (S_D, s_0, P_D, r, AP, L)$ can be constructed by a Markov Decision Process (MDP) M and a given policy π , where S_D is a sub-space of the finite state space of the original MDP, $s_0 \in S_D$ is the initial state, $P_D: S \times A \times S \rightarrow [0, 1]$ is the stochastic transition matrix and $r: (S, A) \rightarrow \mathbb{R}_{\geq 0}$ is a reward function. The policy $\pi: S \times A \rightarrow \{0, 1\}$ is a specific action a , $\pi(a'|s) = 1$ and $\pi(a'|s) = 0$ for all $a' \neq a$. M_D is fully probabilistic with no open action choices.

Reward Process. Consider a non-negative reward function $r: S \rightarrow \mathbb{R}_{\geq 0}$, representing the instantaneous reward received in each state, together with a sequence of random variable $\{X_n\}_{n \geq 0}$. The total cumulative reward up to the time N over the times $0, \dots, N$ is $R(N) = \sum_{n=0}^N r(X_n)$. As the expectation of probabilities and rewards in Markov Models can be computed as the solutions of linear equations [9], the expected total cumulative reward $\mu = \mathbb{E}[\sum_{n=0}^T r(X_n)]$ can be easily computed by solving the linear system $\mu = r + P\mu$ if the system is transient. Furthermore, by utilizing the Moment Generating Function (MGF) of total reward to absorption, we can compute the k -th moments of total reward for higher moment analysis.

Moment-Generating Function (MGF). Consider a set of transient states $C \subseteq S$ and absorbing states $C^c \subseteq S$ (where $C \cap C^c = \emptyset$, $C \cup C^c = S$), with $T = \inf\{n \geq 0: X_n \in C^c\}$ denoting the first hitting time of a state in C^c . The total cumulative reward up to absorption is represented by $R_T =$

$\sum_{n=0}^{T-1} r(X_n)$. The MGF of R_T starting from $X_0 = x$, denoted by $M_{R_T}(\theta, x)$, is given by:

$$M_{R_T}(\theta, x) = \mathbb{E}_x[e^{\theta R_T}] = \mathbb{E}_x \left[\exp \left(\theta \sum_{n=0}^{T-1} r(X_n) \right) \right] \quad (1)$$

Using first-step analysis, this MGF M_{R_T} satisfies the linear system:

$$M_{R_T}(\theta) = f(\theta) + G(\theta)M(\theta),$$

where $f(\theta, x) = \sum_{y \in C^c} e^{\theta r(x)} P(x, y)$ for $x \in C$, and $G(\theta, x, y) = e^{\theta r(x)} P(x, y)$ for $x, y \in C$. The k -th moment $u_k(x) = \mathbb{E}_x [R_T^k]$ starting at $X_0 = x$ can be obtained by differentiating the MGF $u_k = \left. \frac{d^k}{d\theta^k} M_{R_T}(\theta) \right|_{\theta=0}$ with the following recurrence:

$$\begin{aligned} u_k(x) &= \sum_{y \in S} P(x, y) \mathbb{E}_y \left[\sum_{n=0}^{T-1} r(X_n) \right]^k \\ &= r(x)^k + \sum_{i=0}^{k-1} \binom{k}{i} r^{k-i}(x) \sum_{y \in S} P(x, y) u_i(y) \end{aligned} \quad (2)$$

Example. The first moment for our UAV computed with Equation (2) corresponds to the expected reward computed, e.g., by the PRISM model checker [31]: $\mu = 7.51$. However, the standard deviation (square root of the second moment) is $\sigma = 4.91$. Using, for example, Cantelli's inequality [19] (as known as one side Chebysev's inequality [40]), this implies only 50% of the trajectories are guaranteed to fall within the interval $[0, \mu + \sigma]$, leaving a significant portion of distribution not represented.

This result indicates that the expected value alone provides limited information, highlighting further refinement for accurate distribution approximations and robust model checking, particularly in risk-sensitive systems, to ensure reliable outcomes. We introduce next the base distribution we will use for such approximations.

Erlang Distribution. The Erlang distribution is a special case of both phase-type and Gamma distributions with integer shape parameters [13]. An Erlang distribution with shape parameter $a \in \mathbb{Z}^+$ and rate parameter $\lambda \in \mathbb{R}^+$ is defined by the probability density function:

$$f_X(x; a, \lambda) = \frac{\lambda^a x^{a-1} e^{-\lambda x}}{(a-1)!} \quad (3)$$

The Erlang cumulative distribution function is given by:

$$F_X(x; a, \lambda) = 1 - \sum_{j=0}^{a-1} \frac{(\lambda x)^j e^{-\lambda x}}{j!} \quad (4)$$

To approximate the cumulative reward to absorption distribution, we will employ a mixture of Erlang distributions, following its inherent connection with Markov models [1] and its strong ability to closely approximate the distributions of any random variable in continuous settings and discrete settings

¹transition rewards are sometimes used to make model representations more compact (e.g., Fig. 1) but they can be reduced to state rewards w.l.o.g. by systematically replacing a transition $(s, p/r, s')$ with two transitions (s, p, q) and $(q, 1, s')$ where q is a fresh state with associated state reward r .

with very fine granularity [23], [41], [42]. In particular, we will construct a mixture of Erlang distributions with a common rate parameter, which reduces the cost of moment matching without loss of approximation accuracy [8]. Such functional form allows to recover the full distributional behavior of the reward process with any arbitrary accuracy level, matching the first K moments for any $K > 0$.

III. ROBUST PROBABILISTIC MODEL CHECKING

We propose a robust PMC method for DTMCs by approximating the cumulative reward distribution towards absorption through moment matching with mixtures of Erlang distributions. The approximation is formulated as a truncated Stieltjes moment problem [18], to reconstruct the density function from a finite number of its moments on the semi-infinite interval $[0, +\infty)$. This approach specifically focuses on preserving the statistical properties of the distribution while addressing the challenge of having only partial information in the first K moments, rather than the full infinite moment sequence. This approach avoids the complexities of discretization and iterative updates of histogram-based methods limited to discrete rewards [12], thus avoiding potential issues of binning resolution trade-off and state-space explosion. Instead, our method focuses on refining the approximation based on a set of invariant statistical characteristics to approximate the density function with essential statistical information directly.

A. Markov Reward Distribution Approximation

Problem Formulation. Our objective is to approximate the distribution of the cumulative reward towards absorption $R(N)$ in a DTMC D , with initial state s_0 and transition matrix P . We define the approximating density $f_{\text{approx}}(x)$ of the approximate reward distribution X as a mixture of n Erlang densities:

$$f_{\text{approx}}(x) = \sum_{i=1}^n \omega_i f_{\text{Erlang}}(x; a_i, \lambda), \quad (5)$$

where $\omega_i \geq 0$ is the non-negative mixture weight of the i -th Erlang density, $a_i \in \mathbb{Z}^+$ is the shape parameter of the i -th Erlang density, and $\lambda \in \mathbb{R}^+$ is the common rate parameter.

We formalize the approximation process as solving a truncated Stieltjes moment problem [18] on the semi-infinite interval $[0, +\infty)$ with the first K -th moments of the random variable X . The parameters ω_i, a_i, λ are determined by solving an optimization problem that minimizes the difference between the first K -th moments μ_1, \dots, μ_K of the true reward distribution (cf. Sec. II-B) and the moments $\hat{\mu}_1, \dots, \hat{\mu}_K$ of the approximate distribution. We further incorporate the maximum entropy principle within our moment matching formulation [17], [36], to ensure that among all possible approximate distributions matching the given moments, the one with the least bias (i.e., the one maximizing entropy) is selected.

The objective function of the optimization problem is defined as:

$$\min_{\omega_i, a_i, \lambda} \mathcal{L} = \sum_{k=1}^K (\mu_k - \hat{\mu}_k)^2 - \gamma H(f_{\text{approx}}), \quad (6)$$

where $\gamma \geq 0$ is a weighting parameter controlling the trade-off between moment matching and entropy maximization, $H(f_{\text{approx}}) = -\int_0^\infty f_{\text{approx}}(x) \log f_{\text{approx}}(x) dx$ is the differential entropy of the approximated distribution, the moments μ_k of the true reward distribution are computed with the MGF in Equation (2), and the moments of the mixture of Erlang distributions is $\hat{\mu}_k = \sum_{i=1}^n \omega_i \frac{(a_i+k-1)!}{(a_i-1)!} \frac{1}{\lambda^k}$.

Main Algorithm. The approximation algorithm, detailed in Algorithm 1, takes as input a DTMC D along with an arbitrarily small error bound $\epsilon > 0$, the number K of moments to be matched, and the size of the mixture n , and it returns the approximate distribution. For a given number K of moments, the mixture of Erlang distributions can match them using approximately $\lfloor K/2 \rfloor + 1$ components, reducing the mixture size while preserving accuracy [28]. We first compute the true target moments μ_k , providing the necessary statistical information using Equation (2) (lines 1-2). Before the optimization phase, we initialize the mixture weights ω_i , shapes a_i , and rate parameter λ (detailed discussion on initialization and hyperparameters in Section V).

At each optimisation step (in the loop at line 4), we compute the first K moments of f_{approx} and the differential entropy $H(f_{\text{approx}})$. The objective function L balances moment matching with entropy maximization, controlled by the weighting parameter γ . The parameters ω_i, a_i , and λ are updated to minimize L subject to the constraints, until the convergence criteria are satisfied. The algorithm then returns the approximated distribution with the optimized parameters for further analysis via robust model checking. To improve the numerical stability of moment matching, we further utilize standardized moments $\bar{\mu}_k$ in the optimization process $\bar{\mu}_k = \mu_k/c$, with $c = \sqrt{\mu_2}$, for $k = 3, 4, \dots, K$.

Algorithm 1 produces an arbitrarily close approximation f_{approx} that closely matches the cumulative reward distribution associated with discrete Markov models. Arbitrary accuracy can be obtained by increasing the number of moments to be matched and the size of the mixture [18], [42]. The inclusion of entropy maximization ensures that, among all distributions matching the given moments, the algorithm selects the most unbiased one. This feature is particularly advantageous for robust probabilistic model checking to avoid overfitting on a small number of moments [36].

Complexity Observations. The main algorithm starts with a preliminary phase which computes the first K moments of the reward distribution based on the reward scheme and transition matrix of the DTMC. Each moment μ_k is computed by solving a system of linear equations involving the lower-order moments up to μ_{k-1} . Therefore, to compute the first K moments, it is sufficient to solve K systems of linear equations. The size ξ of each system of equations corresponds

Algorithm 1 Markov Reward Distribution Approximation

Input: DTMC D (or induced DTMC M_D) $= (S, s_0, P, r, AP, L)$, error bound $\epsilon > 0$, the number K of moments to be matched, and number n of Erlang distributions to mix

Output: Approximated mixture of n Erlang distributions $f_{\text{approx}}(x)$

- 1: $f = \sum_{y \in T} e^{kr(x)} P(x, y)$ $x \in S$ and $G = e^{kr(x)} P(x, y)$ $x, y \in S$
 - 2: $\mu_k = r^k f + \sum_{i=0}^k \binom{k}{i} G^{(k-i)} \quad \forall k \in \{1, \dots, K\}$ ▷ Compute analytical moments with Eq. (2)
 - 3: Initialize $\mathcal{L} = \infty$ and $\omega_i, a_i, \lambda \quad \forall i \in \{1, \dots, n\}$
 - 4: **repeat**
 - 5: $\mathcal{L}_{\text{prev}} = \mathcal{L}$
 - 6: $\hat{\mu}_k = \sum_{i=1}^n \omega_i \frac{(a_i+k-1)!}{(a_i-1)!} \frac{1}{\lambda^k} \quad \forall k \in \{1, \dots, K\}$ ▷ Compute approximate moments
 - 7: $H(f_{\text{approx}}) = - \int_0^\infty f_{\text{approx}}(x) \log f_{\text{approx}}(x) dx$ ▷ Calculate the differential entropy
 - 8: $\min_{\omega_i, a_i, \lambda} \mathcal{L} = \sum_{k=1}^K (\mu_k - \hat{\mu}_k)^2 - \gamma H(f_{\text{approx}})$
 - 9: subject to $\sum_{i=1}^n \omega_i = 1, \quad \omega_i \geq 0, \quad a_i \in \mathbb{N}^+, \quad \lambda > 0$
 - 10: $\omega_i, a_i, \lambda \leftarrow \arg \min_{\omega_i, a_i, \lambda} \mathcal{L}$ ▷ Update parameters
 - 11: **until** $|\mathcal{L}_{\text{prev}} - \mathcal{L}| < \epsilon$
 - 12: **return** $f_{\text{approx}}(x) = \sum_{i=1}^n \omega_i f_{\text{Erlang}}(x; a_i, \lambda)$
-

to the size of the transition matrix P , i.e., the number of states in the model. A worst-case complexity for solving a system of linear equations is $\mathcal{O}(\mathcal{E}^3)$, although standard numerical solvers (e.g., Numpy [22]) usually improve over such a bound and can exploit the sparsity and possible symmetries in the transition matrix for more efficient computations.

In the optimization loop, we separate the computation of the differential entropy of f_{approx} from the main optimization problem. While this step may increase the number of iterations to convergence, it reduces the complexity of the inner optimization problem at Line 8. This optimization problem is a mixed-integer non-linear form (the parameters a_i being positive integers), which places it in the NP class, which may be computationally intractable.

A common heuristic adopted in the literature to reduce the complexity of optimization is to fix a set of possible values for the shape parameters a_i [33], [44] while the optimizer identifies the optimal rate parameter and the weights of the mixture. While setting each $a_i = i$, for $i = 1, 2, 3, \dots$ ensures an optimal approximation by relying on the density of Erlang mixtures [42, Theorem 2.9.1], it may also require a large number of components to approximate the underlying distribution, especially in presence of heavy tails, sharp drops, or multiple modalities. Furthermore, using a large mixture to fit a small number of moments may lead to overfitting. A common recommendation when fitting a finite number of moments K is instead to spread the shape parameters [43, Sec 5.4.3] fed to the optimizer. We found in our experiments that an exponential spreading (i.e., $a_i = c^i$) of the fixed shapes yielded the best results, compared to a dense allocation (i.e., $a_i = i$) and the linear spreading (i.e., $a_i = c \cdot i$ for some $c > 1$) used in [43]. Fixing the shape parameters reduces the optimization problem to a non-convex quadratic form [6] that state-of-the-art optimizers – such as Gurobi [21] or SciPy [45] – can solve quickly, as demonstrated experimentally in Section V where we also report on the accuracy and computation time required for different spreading strategies and number of components.

Finally, in the next section we will describe how probability inequalities can be used to draw sufficient decisions on chance-constrained requirements directly from the moments computed during the preliminary phase of Algorithm 1, thus avoiding the need to execute the optimization loop entirely.

B. Distributional Analysis for Robust Model Checking

Robust Property Evaluation. The evolution of the random variable X can be treated as the outcome cumulative reward of the underlying stochastic system, with each transition governed by probabilistic dynamics. By considering the density f_{approx} of X , we formalize a robust property evaluation for chance-constrained requirements [14] of the form:

$$Pr(X \leq r^*) \geq \alpha, \quad (7)$$

where r^* is a critical threshold representing a performance or safety metric, and α is the required probability level. For example, in a safety-critical system, we can verify that the probability of a performance violation remains below a predefined risk level α , rather than asking whether just the expected reward is below the threshold. Notice that constraints of the form $X \geq r^*$ or, more in general, $r_* \leq X \leq r^*$, can be evaluated with the same method using the standard properties of the cumulative distribution function.

Almost sure properties, where the probability of satisfying a temporal formula is required to be 1 (or 0), can also be straightforwardly phrased as chance-constrained requirements, whereas chance-constrained formulations allow for the controlled relaxation of the requirements. This makes our distributional method well suited for the verification of probabilistic safety requirements [29], as well as for the robustness analysis of the policies synthesized for MDPs via reinforcement learning [20], by reasoning on the induced DTMC model [4, Ch. 10].

From f_{approx} , the evaluation of the chance-constrained requirement in Equation (7) reduces to computing the cumulative distribution function of the finite Erlang mixture X :

$$Pr(X \leq r^*) = \sum_{\omega_i} \omega_i F_{X_i}(r^*; a_i, \lambda_i) =: F_{\text{approx}}(r^*) \quad (8)$$

where the cumulative distribution functions $F_{X_i}(x; a_i, \lambda_i)$ are defined for each Erlang component of the mixture as in Equation (4).

Efficient Sufficient Conditions. After the moments of the rewards distribution are computed in Algorithm 1 (lines 1 and 2), a decision on the validity of a chance-requirement may already be possible, without computing f_{approx} , by relying on probabilistic inequality derived from Markov’s theorem [19]. Such inequalities can bound the probability of a random variable exceeding a threshold by computing only with a finite number of moments of the distribution. Being agnostic to the actual shape of the distribution, such inequalities tend to be over-conservative, as we will show later.

Nevertheless, probability inequalities, when they confirm the satisfaction of the chance-constrained requirements, can be used as an early termination criterion for Algorithm 1, allowing to skip the optimization loop entirely. Under our assumption of non-negative rewards, the sharpest inequality derived from Markov’s theorem is Cantelli’s [19] (also referred to as one-sided Chebyshev’s inequality). Its formulation generalized to higher moments is:

$$Pr(X - E[X] \geq a) \leq E\left[\left(\frac{X + b}{a + b}\right)^n\right] = \frac{E[(X + b)^n]}{(a + b)^n} \quad (9)$$

where $E[(X + b)^n]$ is the n -th moment of X centered in $-b$. When $n = 2$, Equation (9) reduces to the original formulation of Cantelli’s inequality by choosing $b = \sigma^2/a$, σ being the standard deviation of X (the low-tail correspondent, can be obtained by applying the inequality to $-X$, thus flipping the constraint to a \leq form).

After computing the first K moments of the reward distribution in the initialization phase of Algorithm 1, Equation (9) can provide a fast and robust decision on the probabilistic model checking problem by comparing its right-hand side with $1 - \alpha$. If the chance constraint holds, the decision procedure terminates. Otherwise, it is still possible that Cantelli’s inequality is not satisfied due to its conservativeness – the inequality must hold for *any* distribution. In such a case, reconstructing f_{approx} is required to obtain a more accurate answer.

Example. Figure 2 shows the cumulative distribution for our UAV running example. The solid green line shows the CDF F_{approx} (integral of f_{approx}) estimated with our method by matching the first three moments to construct a mixture of three Erlang distributions. The shaded bars show the empirical distribution drawn from one million simulations of the system, which we use to approximate the ground truth. The dashed vertical lines show the probability bounds computed using Cantelli’s inequality for the chance requirement of Equation (7), with values of r^* being on the x-axis and each dashed

line showing a different probability threshold α . Consider the requirement $Pr(X \leq 15) \geq 0.65$. Cantelli’s inequality ($n = 3$) ensures that $Pr(X \leq 15) \geq 0.68$, therefore the requirement holds and there is no need to compute f_{approx} (optimization loop in Algorithm 1) to draw a conclusion. If we were instead to verify $Pr(X \leq 15) \geq 0.8$, due to its conservative nature, Cantelli’s bound would not be able to prove that the system indeed satisfies the requirement, as correctly concluded by evaluating $F_{\text{approx}}(15)$.




Fig. 2. Cantelli Bound’s Coarseness in Model Checking

The moment matching process ensures that the approximation closely represents the true stochastic behavior of the system using only the first K moments [18]. The use of multiple moments, captures both central tendencies and tail behaviors, which are essential for assessing system robustness. The next section includes the main theoretical foundation underpinning our choice of approximants and moment-matching, while in Section V, we assess experimentally the quality of the approximants with different numbers of moments and mixture size, as well as the computation time required to construct them.

IV. THEORETICAL ANALYSIS

In this section, we demonstrate that it is acceptable and effective to approximate the cumulative reward distribution of a DTMC using mixtures of Erlang distributions. We use Erlangization to obtain a continuous approximation which is later approximated by a mixture of Erlang distribution on a modified DTMC with discretized reward. By applying these theorems, we justify the use of mixtures of Erlang distributions as a valid method for approximating the cumulative reward distribution and its exponential tail behavior.

Theorem IV.1 (Approximation of Cumulative Reward by Discrete Phase-Type (DPH) Distribution). *Consider a DTMC $D = (S, s_0, P, \mathbf{r})$ with a finite state space and non-negative rewards $r_i \in \mathbb{R}_{\geq 0}$ and a cumulative reward towards absorption $F(r) = R_T(s_0)$. For any $\epsilon > 0$, there exists a precision parameter $\delta > 0$ and a DTMC $\tilde{D} = (\tilde{S}, s_0, \tilde{P}, \delta \mathbf{e})$, where \mathbf{e} is the vector of all ones, such that $F(r)$ can be approximated*

by the DPH distribution $F_{\tilde{D}}(r) = \text{DPH}(s_0, \tilde{P})(r/\delta)$, in such a way that

$$\sup_{r \in [0, +\infty)} |F(r) - F_{\tilde{D}}(r)| \leq \epsilon.$$

Proof. For any given precision $\delta > 0$, we construct an extended DTMC $\tilde{D} = (\tilde{S}, \tilde{P}, \tilde{\mathbf{r}})$ by discretizing the rewards in D . Specifically, for each reward $r_i \in \mathbf{r}$, define the approximate integer reward $k_i = \lceil r_i/\delta \rceil$, where $k_i\delta$ is the closest integer multiple of δ that approximates r_i . The extended reward vector $\tilde{\mathbf{r}} = \{\delta, \dots, \delta\} = \delta \mathbf{e}$ assigns a constant reward of δ to each state. For each state $s_i \in S$ with approximate reward $k_i\delta$, we introduce $n_i = k_i - 1$ additional intermediate states, denoted as $s_{i,1}, s_{i,2}, \dots, s_{i,n_i}$. The extended state space becomes:

$$\tilde{S} = \{s_i, s_{i,1}, \dots, s_{i,n_i} : s_i \in S\}.$$

The extended transition matrix $\tilde{P} \in \mathbb{R}^{|\tilde{S}| \times |\tilde{S}|}$ is defined as follows: for each transition $s_i \rightarrow s_j$ in the original matrix P , we introduce intermediate transitions $s_i \rightarrow s_{i,1} \rightarrow s_{i,2} \rightarrow \dots \rightarrow s_{i,n_i} \rightarrow s_j$. Formally, the extended transition matrix \tilde{P} replaces each p_{ij} by \tilde{P}_{ij} , defined as:

$$\tilde{P}_{ij} = \begin{pmatrix} 0 & 1 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ p_{ij} & 0 & \dots & 0 \end{pmatrix}_{k_i \times k_i} = \begin{pmatrix} 0 & \mathbb{I}_{n_i \times n_i} \\ p_{ij} & 0 \end{pmatrix}.$$

The maximum error per reward is bounded by $0 \leq k_i\delta - r_i < \delta$. Therefore, the cumulative reward error after the termination time T satisfies $0 \leq \tilde{R}_T(s_0) - R_T(s_0) \leq T\delta$. Applying Markov's inequality yields:

$$\Pr \left(|\tilde{R}_T(s_0) - R_T(s_0)| \geq \epsilon \right) \leq \Pr \left(T \geq \frac{\epsilon}{\delta} \right) \leq \frac{E[T]\delta}{\epsilon}.$$

As $\delta \rightarrow 0$, the right-hand side approaches zero, leading to the weak convergence of the cumulative distribution functions: $\lim_{\delta \rightarrow 0} F_{\tilde{D}}(r) = F(r)$ for any continuity point r of F . \square

Thus, the extended DTMC \tilde{D} , constructed by introducing intermediate states and modifying the transition matrix P to accumulate integer rewards in steps of ϵ , approximates the cumulative reward distribution of the original DTMC D . The resulting transition matrix \tilde{P} provides a DPH representation of the reward distribution, enabling further approximation by other tractable distributions, such as Erlang mixtures, while ensuring accurate control over the approximation error.

Theorem IV.2 (Continuous Approximation of DPH by Erlangization). *Consider X be a DPH random variable with representation $\text{DPH}(\boldsymbol{\alpha}, P)$. Define the continuous phase-type random variable X_m by $\text{PH}_c(\boldsymbol{\gamma}_m, T_m)$, where $\boldsymbol{\gamma}_m = \boldsymbol{\alpha} \otimes \mathbf{e}(1)$ and $T_m = \lambda (I \otimes J_0 + P \otimes J_1)$ where*

$$J_0 = \begin{pmatrix} -1 & 1 & & & \\ & \ddots & \ddots & & \\ & & -1 & 1 & \\ & & & & -1 \end{pmatrix}, J_1 = \begin{pmatrix} 0 & & & & \\ & 0 & & & \\ & & \ddots & & \\ & & & & \\ 1 & 0 & \dots & 0 & \end{pmatrix}$$

such that by replacing the constant sojourn time in Markov chain of X with an Erlang random variable $E_{\lambda, m}$, where $\lambda = m$. Then X_m is a continuous approximation of X and converges to X in distribution and moments, such that:

$$\mathbb{E}[X_m] = \mathbb{E}[X]$$

$$\mathbb{E}[X_m^2] = \mathbb{E}[X^2] + \frac{\mathbb{E}[X]}{m}$$

$$\mathbb{E}[X_m^n] - \mathbb{E}[X^n] = \frac{\mathbb{E}[X^{n-1}]d_{n,n-1}}{m} + o\left(\frac{1}{m}\right)$$

X_m converges to X in Laplace-Stieltjes transformation hence also convergence in distribution as $m \rightarrow \infty$.

Proof. See ref. [23], Theorem 1. \square

Theorem IV.2 provides a method to obtain the continuous extension of the cumulative reward together with the previous theorem. The next theorem shows this continuous extension can be approximated by an Erlang mixture.

Theorem IV.3 (Mixture of Erlang Distributions Approximation). *Let $F(x)$ be the cumulative distribution function (CDF) of a positive random variable. For a fixed $\beta > 0$, define an approximate CDF $F_\beta(x)$:*

$$F_\beta(x) = \sum_{j=1}^{\infty} p_j(\beta) \cdot F_{\text{Erlang}(j, \beta)}(x), \quad x \geq 0$$

where $p_j(\beta) = F(j\beta) - F((j-1)\beta)$, $j = 1, 2, \dots$ and $F_{\text{Erlang}(j, \beta)}(x) = 1 - \sum_{k=0}^{j-1} \frac{e^{-x/\beta} (x/\beta)^k}{k!}$ is the CDF of an Erlang distribution with shape parameter j and scale parameter $\beta = 1/\lambda$. Then, we have the convergence:

$$\lim_{\beta \rightarrow 0} F_\beta(x) = F(x)$$

for any continuity point x of F .

Proof. We provide a simplified proof and a detailed proof see. ref [42], Theorem 3.9.1.

For fixed β , let $U_{\beta, x}$ be a Poisson-distributed random variable with

$$\Pr \{U_{\beta, x} = k\beta\} = e^{-x/\beta} \frac{(x/\beta)^k}{k!}, \quad k = 0, 1, \dots$$

with expectation $\mathbb{E}[U_{\beta, x}] = x$ and variance $\mathbb{E}[U_{\beta, x}^2] = x\beta$. As $\beta \rightarrow 0$, $U_{\beta, x}$ concentrates around x due to vanishing variance. For a bounded function $g(t)$, we show that:

$$\lim_{\beta \rightarrow 0} \mathbb{E}[g(U_{\beta, x})] = g(x)$$

by bounding the error $|E[g(U_{\beta, x})] - g(x)|$. Using Chebyshev's inequality, we control the tail probabilities, showing that as $\beta \rightarrow 0$, the contributions from $|U_{\beta, x} - x|$ become arbitrarily small. Applying this to $g(t) = F(t)$, the CDF of the random variable, we get:

$$F(x) = \lim_{\beta \rightarrow 0} E[F(U_{\beta, x})]$$

which leads to:

$$F(x) = \lim_{\beta \rightarrow 0} \sum_{j=1}^{\infty} p_j(\beta) \sum_{k=j}^{\infty} e^{-x/\beta} \frac{(x/\beta)^k}{k!}$$

yielding the desired convergence:

$$\lim_{\beta \rightarrow 0} F_{\beta}(x) = F(x)$$

□

The combination of discretizing rewards, approximating DPH distributions with continuous phase-type distributions, and representing these as Erlang mixtures provides a solid foundation for approximating the cumulative reward distribution of a DTMC. In real applications, selecting the appropriate number of Erlang distributions in the mixture to get an acceptable approximation is an empirical task [44].

V. EVALUATION

We evaluate our moment-matching approximation method with error analysis using case studies adapted from PRISM [31] and used also in previous literature [12]. Our implementation employs an Erlang mixture model with moment matching and maximum entropy using SciPy SLSQP [45], and Gurobi [21] for optimization, while the linear systems of equations required to compute the moments of the reward process are solved with Numpy [22]. The primary goal is to empirically assess the approximation performance of our method against the empirical distribution and traditional histogram-based approximation from the perspective of accuracy and computational tractability, as well as discuss hyper-parameters and overheads.

Benchmarks. The benchmarks, adapted from PRISM case studies [31], include scenarios like message synchronization and Herman protocol [25], which are modeled as DTMCs; for the MDP benchmarks (such as betting games, and grid world navigation), we evaluate the DTMC induced by the optimal policy produced by PRISM (reward schemes defined for the original models). We evaluate our method also on two models with continuous rewards, which cannot be analyzed by previous work [12]: UAV (our example model) and an adaptation of the Future Investor model [35]. The benchmark subjects cover a diverse range of state space sizes and transition structures, as summarized in Table I.

Experimental Settings. We construct the approximate reward distribution from the initial state of each model using SLSQP [45]. To evaluate the accuracy of the approximation, we construct a baseline empirical reward distribution by running one million simulations of the models and collecting the reward accumulated for each run. For models using integer rewards, we also compare with the histogram distributions constructed by [12].

To align the approximation problem with the truncated Stieltjes moment problem on $(0, +\infty)$, we shift the location of the Erlang distributions to $loc = \mu - \sigma$ (σ being the standard deviation of the reward distribution), following the location

Subject	States	Transitions	d_{KS} ([12])	d_{KS} (our)
GridWorld Navigation	396	6,291	0.42	0.14
Betting Game	891	4,216	0.16	0.09
DeepSea Treasure	1,756	5, 672	0.28	0.13
Herman	2,048	177,148	0.27	0.09
LeaderSync	1,050	1,292	0.01	0.07
UAV Energy	13	29	-	0.05
Future Investor	2,899	7,967	-	0.17

TABLE I
BENCHMARK SUBJECTS WITH DISCRETE REWARDS (TOP) AND
CONTINUOUS REWARDS (BOTTOM).

setting in [18]. Based on [28], the minimum number of Erlang components in the mixture required to match K moments is $\lfloor K/2 \rfloor + 1$. To allow more flexibility in the optimization, unless otherwise specified, we use $k = 3$ moments and $n = 3$ components. The common rate parameter λ is limited for optimization within the range $(0.01, 50)$ (based on our experience, the range was large enough to include the optimization results). The weights are initialized uniformly: $\omega_i = 1/n$. The parameter γ balancing the trade-off between moment matching and maximum entropy is set to 1 for all the experiments. As discussed in Section III-A, we restrict the Erlang shapes into a power-based sequence of shape parameters $a_i = 3^i$, $i \leq n$, which ensures a wide spread of options in the weighted selection of shapes [43] to cover head, multi-modal, and tail behaviors— we will evaluate the robustness of this choice of heuristics and compare with alternative choices in Section V-B.

A. Approximation Accuracy

To evaluate the accuracy of approximate distributions against the baseline empirical distributions, we use the Kolmogorov-Smirnov (KS) statistic metric. KS measures the largest difference between the empirical CDF of the true values and the approximated CDFs with $D_{KS} = \sup_x |F_{\text{empirical}}(x) - F_{\text{approx}}(x)|$. The KS statistic is chosen for its focus on capturing the maximum discrepancy between distributions, making it well-suited and widely used for comparing distributions and evaluating approximation quality [26], [34].

Discrete Reward Space. Figures 3 and 4 present an evaluation of our method on the discrete reward space subjects. The blue histograms represent the discrete approximation constructed by [12], referred to as *histogram method* in the figures. The gray histogram represents the empirical baseline, constructed from one million sample runs of the model. The KS comparison measures are reported in Table I.

Overall, the two methods show comparable performance, with our method achieving lower KS divergence for all subjects except *LeadSync*. In the *LeadSync* case, support of the probability distribution is in $(1, 10] \cap \mathbb{Z}$, allowing the histogram to fully capture the true distribution without loss of information. This result also shows that the accuracy of histogram-based methods is highly dependent on bin width, with finer binning improving precision but increasing the

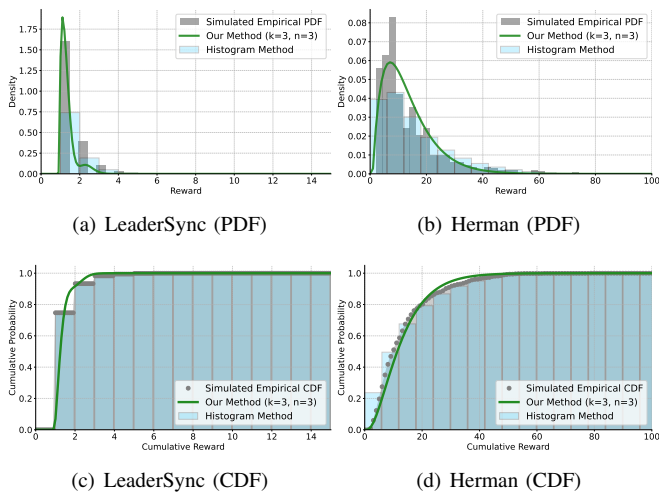


Fig. 3. PDF and CDF for discrete reward space subjects (figure 1 of 2)

computational cost. In the Herman scenario, our method shows a closer alignment with the empirical distribution, as also captured by the lower KS metric (Table I). Figure 4 includes the approximations for the remaining discrete reward space subjects, where the probability density function (PDF) includes multiple modalities, and the achievable reward values are generally more scattered. In these cases, the Erlang-based approximation offers a smoother fit compared to histograms and outperforms [12], especially in sparsely populated regions like GridWorld Navigation, as confirmed by the d_{KS} metric. Although the performance of any histogram-based method can be improved with finer bins [5], scalability and sparse-dense ranges issues remain problematic in continuous and larger finite state spaces.

Continuous Reward Space. Figure 5 shows the evaluation of our method on the two scenarios with continuous reward space: UAV Energy and Future Investor. The empirical distribution is represented by the gray histogram. For easier comparison across the subjects, we used $K = 3$ moments and $n = 3$ components of the Erlang mixture. In the case of Future Investor, as noticeable by visual inspection, and reflected in a KS metric of 0.17, this (K, n) settings tend to overapproximate the CDF around the head of the distribution, catching up after $x \geq 12$. In the next section, we will dive deeper into how such approximation can be improved by increasing K and n .

B. Hyper-parameters Effects on Accuracy and Runtime

Number of Moments and Erlang Components. Figure 6 shows the plots of the approximate distribution constructed for the Future Investor subject computed with $K \in \{3, 4, 5\}$ and $n \in [3..9]$. We selected Future Investor for this evaluation due to the presence of multiple, overlapping modalities in its PDF with a continuous reward space, which renders approximation particularly challenging. The KS metrics in Table II show that increasing the values of K and n

K	n	T_{opt} (s)	T_{total} (s)	Iterations	D_{KS}
3	3	0.17	1.42	36	0.17
	4	0.19	1.44	43	0.17
	5	0.31	1.56	57	0.26
	6	0.23	1.48	57	0.16
	7	0.44	1.69	63	0.15
	8	0.62	1.87	80	0.14
4	9	1.65	2.91	92	0.21
	3	0.59	2.20	69	0.12
	4	0.73	2.34	72	0.10
	5	0.87	2.48	74	0.10
	6	0.82	2.43	75	0.07
	7	0.39	2.00	71	0.09
5	8	0.92	2.53	79	0.09
	9	0.47	2.08	88	0.09
	3	2.14	4.24	92	0.17
	4	1.36	3.46	115	0.16
	5	0.80	2.90	108	0.13
	6	0.21	2.31	101	0.07
5	7	0.74	2.84	94	0.07
	8	0.83	2.93	84	0.07
	9	2.30	4.40	131	0.07

TABLE II
COMPUTATION TIMES AND D_{KS} FOR DIFFERENT COMBINATIONS OF MOMENTS K AND MIXTURE SIZE n ON THE FUTURE INVESTOR SUBJECT.

a_i	T_{opt} (s)	T_{total} (s)	D_{KS}
i	0.62	2.73	0.16
$2 \cdot i$	1.95	4.05	0.15
$3 \cdot i$	0.74	2.84	0.13
2^i	1.45	3.55	0.12
3^i	0.21	2.31	0.07
4^i	1.37	3.47	0.07

TABLE III
COMPARISON OF DIFFERENT SHAPE PARAMETER SETS a_i (FUTURE INVESTOR SUBJECT, $k = 5, n = 6$)

can achieve higher accuracy – the KS metric drops from 0.17 to 0.07. This is in line with the theoretical expectations from Section IV: larger mixtures theoretically enhance approximation accuracy by allowing Erlang distributions to approximate any positive random variable, while more moments capture additional more features of the reward process, including multiple modalities and tail behavior. For example, $K = 4$ and $n = 6$, already achieve the best KS metric (0.07), which is also reported for $K = 5$ and $n \in \{6, 7, 8, 9\}$.

However, numerical instability and limitations of the optimizer accuracy on this class of problems, means that, in practice, the quality of the (near-)optimal solutions found by the optimizer does not increase monotonically with K and n . Although larger mixtures subsume smaller ones by setting some weights to zero, and generally achieve better performance due to increased expressiveness, the non-convexity of the problem means that the optimization is not guaranteed to converge to a single global optimum. In particular, larger n in our experiments occasionally led to worse performance

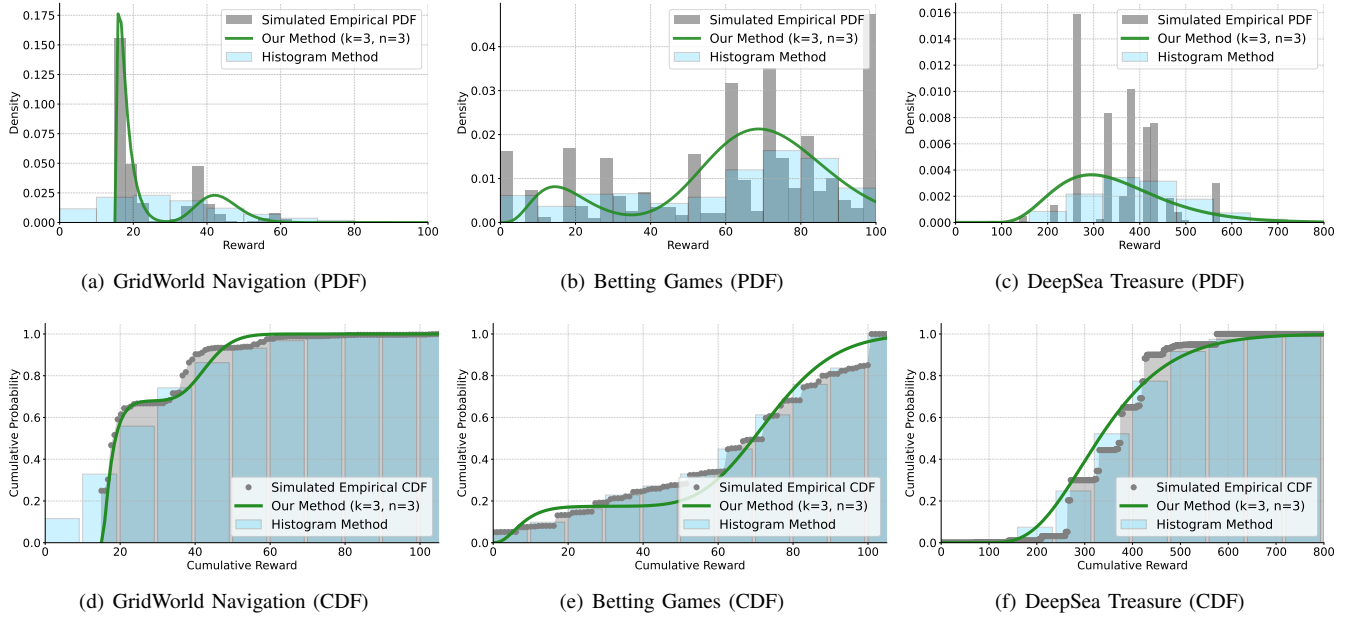


Fig. 4. PDF and CDF for discrete reward space subjects (figure 2 of 2)

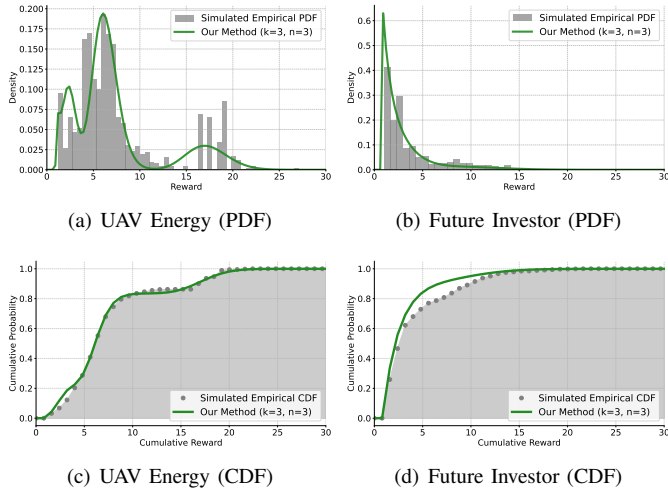


Fig. 5. PDF and CDF for continuous reward space subjects

due to the accumulation of numerical errors across the larger number of free variables (e.g., see $K = 4$ and $n \in \{7, 8, 9\}$, which performs marginally worse than with $n = 6$). We plan to experiment with different optimization strategies in the future, including the use of more robust methods such as Expectation Maximization as in [44] and studying different relaxation of the optimization problem to improve the stability of its numerical solution.

Runtime. Table II summarizes the runtime statistics for each of the (K, n) settings for `Future Investor` (2,899 states and 7,967 transitions). The total time (T_{total} , in seconds) shows the end-to-end time to construct the approximate distribution, which includes the computation of the moments of the reward distribution and the optimization process to

infer the Erlang mixture. The optimization time (T_{opt} , in seconds) indicates the total time spent on optimizer’s calls. The column `Iterations` shows the number of iterations of the optimization loop required to converge. While the number of iterations seems to weakly correlate with the value of K , the time for all the experiments remained consistently below 5 seconds on a Macbook Air M1 (8Gb of memory). Notably, both the evaluation of Cantelli’s inequality for early termination (Eq. (9)) and of the CDF of the Erlang mixture (Eq. (8)) only require negligible time to evaluate closed-form expressions.

Heuristic Restriction of Shape Values. In line with previous literature [33], [43], [44], we restricted the integer shape parameter values within a fixed set, in order to simplify the optimization problem (cf. Section III-A). Table III evaluates different heuristics on `Future Investor` with $K = 5$ and $n = 6$, which yielded an optimal KS metric for the subject: $a_i = i$ to the original assignment in [42], which converges to the optimal approximation only for large mixtures; $a_i = 2 \cdot i$ and $a_i = 3 \cdot i$ use a linear spreading of shape values [43]; finally, $a_i = 2^i$, $a_i = 3^i$ and $a_i = 4^i$ use the exponential spreading of our heuristics. The intuition behind spreading the shape parameter values is to allow capturing distribution features that may be different from the head to the tail, especially in the presence of multiple modalities. The exponential spreading heuristic consistently achieved the best fitting. The heuristic $a_i = 3^i$ consistently achieved the best performance across the board of our diverse subjects. Expectation Maximization has been used in literature to further refine the shape values around their initial heuristics points [43]; we plan to experiment with such finer tuning in future work.

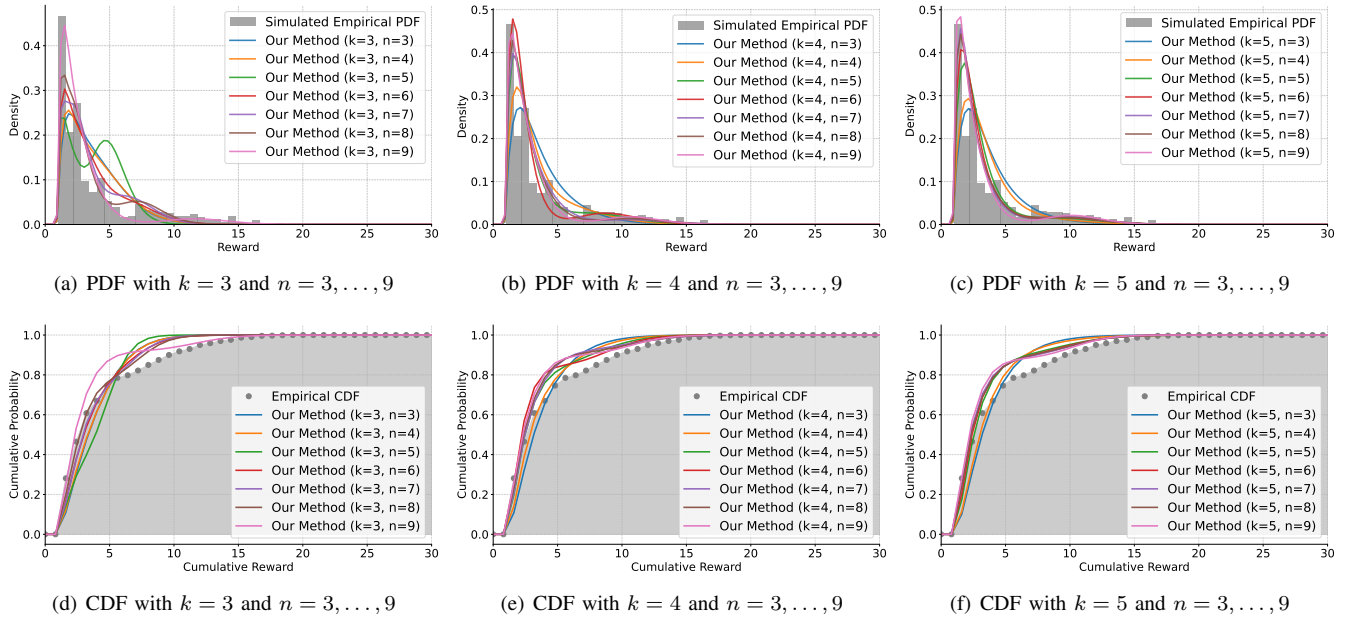


Fig. 6. Approximated Probabilities Distribution of Financial Market with different number of mixtures n

VI. RELATED WORK

Moment Matching Approximation. Moment matching is a classical method for approximating density functions [27], [37]. The main challenge of density approximation methods is to select an appropriate foundational distribution to represent the unknown distribution. Existing methods have focused on using powerful kernel density function [18] or Erlang mixtures [23], [28], given their capability in recovering a wide variety of distributions. However, the accuracy of these approximations depends heavily on how well the selected function or kernel aligns with the underlying distribution. We utilize the Erlang mixtures, not only due to their flexibility but also because of its inherent connection to Markov models, which allows us to approximate the reward distribution with bounded error.

Robust Model Checking. Robustness is one of the most important concerns in model checking methods. Some simulation-based approaches provide probabilistic bounds with confidence guarantees of model correctness [24], [38]. There are also several methods that provide evaluation beyond expected value with distributional information, including the quantile-based method [2] and histogram approximation method [12], but could lead to unbounded information loss due to the coarse approximation. In contrast, our approach provides smoother approximation while avoiding discretization errors and offering guaranteed accuracy in both discrete and continuous domains.

Distributional Reinforcement Learning. Recent work in distributional reinforcement learning has focused on estimating the full return distribution instead of just its expectation. A categorical histogram approximation method has been proposed in [5] to compute return distributions using discrete bins,

optimizing the Wasserstein distance between estimated and target distributions. Extensions of this work, such as quantile regression (QR-DQN) [11] and implicit quantile networks (IQN) [10], improve on this by learning the quantile-based approximation, offering greater flexibility and risk-sensitive policies. Our method can be further extended to the RL domain, providing precise estimation and flexibility in continuous reward environments than histogram methods, with a robust verification that can be used as guidance for distributional policy iteration.

VII. CONCLUSION

In this work, we addressed the limitations of current PMC methods that rely on expected values, by proposing a moment-matching approach using mixtures of Erlang distributions. This approach enables a more precise approximation of cumulative reward distributions in DTMCs by capturing multi-modality and higher-order statistical features. We enhance system verification by incorporating the full distribution of outcomes, enabling better decision support for adaptive systems. Next, we aim to extend our approach to robust model checking and learning in MDPs, where distributional robustness could guide policy search. We plan to leverage the prior knowledge of the accumulated reward distribution in MDPs, as demonstrated in this work, to provide a robust alternative to histogram-based approximation, which may struggle in distributional RL scenarios involving sparse or dense reward regions and continuous reward space. We also aim to explore the applicability of this approach in distributional learning settings beyond distributional model checking, to enhance policy iteration process and evaluate its potential for synthesizing more robust policies under uncertainty in practical applications.

REFERENCES

- [1] AYALA, D., JOFRÉ, L., GUTIÉRREZ, L., AND MENA, R. H. On a Dirichlet Process Mixture Representation of Phase-Type Distributions. *Bayesian Analysis* 17, 3 (2022), 765 – 790.
- [2] BAIER, C., DUBSLAFF, C., KLÜPPPELHOLZ, S., DAUM, M., KLEIN, J., MÄRCKER, S., AND WUNDERLICH, S. Probabilistic model checking and non-standard multi-objective reasoning. In *International Conference on Fundamental Approaches to Software Engineering* (2014), Springer, pp. 1–16.
- [3] BAIER, C., GRÖSSER, M., AND CIESINSKI, F. *Model Checking Linear-Time Properties of Probabilistic Systems*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 519–570.
- [4] BAIER, C., KATOEN, J., AND LARSEN, K. *Principles of Model Checking*. The MIT Press. MIT Press, 2008.
- [5] BELLEMARE, M. G., DABNEY, W., AND MUNOS, R. A distributional perspective on reinforcement learning. In *International conference on machine learning* (2017), PMLR, pp. 449–458.
- [6] BURER, S., AND LETCHFORD, A. N. Non-convex mixed-integer nonlinear programming: A survey. *Surveys in Operations Research and Management Science* 17, 2 (2012), 97–106.
- [7] CALINESCU, R., GHEZZI, C., KWIATKOWSKA, M., AND MIRANDOLA, R. Self-adaptive software needs quantitative verification at runtime. *Commun. ACM* 55, 9 (Sept. 2012), 69–77.
- [8] CHAKRAVARTHY, S., AND ALFA, A. S. *Matrix-analytic methods in stochastic models*. CRC Press, 1996.
- [9] CINLAR, E. *Introduction to stochastic processes*. Courier Corporation, 2013.
- [10] DABNEY, W., OSTROVSKI, G., SILVER, D., AND MUNOS, R. Implicit quantile networks for distributional reinforcement learning. In *International conference on machine learning* (2018), PMLR, pp. 1096–1105.
- [11] DABNEY, W., ROWLAND, M., BELLEMARE, M., AND MUNOS, R. Distributional reinforcement learning with quantile regression. In *Proceedings of the AAAI conference on artificial intelligence* (2018), vol. 32.
- [12] ELSAYED-ALY, I., PARKER, D., AND FENG, L. Distributional probabilistic model checking. In *NASA Formal Methods Symposium* (2024), Springer, pp. 57–75.
- [13] ERLANG, A. K. The theory of probabilities and telephone conversations. *Nyt. Tidsskr. Mat. Ser. B* 20 (1909), 33–39.
- [14] FARAHANI, S. S., MAJUMDAR, R., PRABHU, V. S., AND SOUDJANI, S. E. Z. Shrinking horizon model predictive control with chance-constrained signal temporal logic specifications. In *2017 American Control Conference (ACC)* (2017), pp. 1740–1746.
- [15] FILIERI, A., TAMBURRELLI, G., AND GHEZZI, C. Supporting self-adaptation via quantitative verification and sensitivity analysis at run time. *IEEE Transactions on Software Engineering* 42, 1 (2016), 75–99.
- [16] FOREJT, V., KWIATKOWSKA, M., NORMAN, G., PARKER, D., AND QU, H. Quantitative multi-objective verification for probabilistic systems. In *Tools and Algorithms for the Construction and Analysis of Systems* (Berlin, Heidelberg, 2011), P. A. Abdulla and K. R. M. Leino, Eds., Springer Berlin Heidelberg, pp. 112–127.
- [17] GANAPATHI, V., VICKREY, D., DUCHI, J., AND KOLLER, D. Constrained approximate maximum entropy learning of markov random fields. *arXiv preprint arXiv:1206.3257* (2012).
- [18] GAVRILIADIS, P., AND ATHANASSOULIS, G. The truncated stieltjes moment problem solved by using kernel density functions. *Journal of Computational and Applied Mathematics* 236, 17 (2012), 4193–4213.
- [19] GHOSH, B. K. Probability inequalities related to markov’s theorem. *The American Statistician* 56, 3 (2002), 186–190.
- [20] GROSS, D., JANSEN, N., JUNGES, S., AND PÉREZ, G. A. Cool-mc: a comprehensive tool for reinforcement learning and model checking. In *International Symposium on Dependable Software Engineering: Theories, Tools, and Applications* (2022), Springer, pp. 41–49.
- [21] GUROBI OPTIMIZATION, LLC. Gurobi Optimizer Reference Manual, 2024.
- [22] HARRIS, C. R., MILLMAN, K. J., VAN DER WALT, S. J., GOMMERS, R., VIRTANEN, P., COURNAPEAU, D., WIESER, E., TAYLOR, J., BERG, S., SMITH, N. J., ET AL. Array programming with numpy. *Nature* 585, 7825 (2020), 357–362.
- [23] HE, Q.-M., LIU, B., AND WU, H. Continuous approximations of discrete phase-type distributions and their applications to reliability models. *Performance Evaluation* 154 (2022), 102284.
- [24] HÉRAULT, T., LASSAIGNE, R., MAGNIETTE, F., AND PEYRONNET, S. Approximate probabilistic model checking. In *International Workshop on Verification, Model Checking, and Abstract Interpretation* (2004), Springer, pp. 73–84.
- [25] HERMAN, T. Probabilistic self-stabilization. *Information Processing Letters* 35, 2 (1990), 63–67.
- [26] HONG-ZHI, A., AND BING, C. A kolmogorov-smirnov type statistic with application to test for nonlinearity in time series. *International Statistical Review/Revue Internationale de Statistique* (1991), 287–307.
- [27] JOHN, V., ANGELOV, I., ÖNCÜL, A., AND THÉVENIN, D. Techniques for the reconstruction of a distribution from a finite number of its moments. *Chemical Engineering Science* 62, 11 (2007), 2890–2904.
- [28] JOHNSON, M. A., AND TAAFFE, M. R. Matching moments to phase distributions: Mixtures of erlang distributions of common order. *Stochastic Models* 5, 4 (1989), 711–743.
- [29] KATOEN, J.-P. The probabilistic model checking landscape. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science* (New York, NY, USA, 2016), LICS ’16, Association for Computing Machinery, pp. 31–45.
- [30] KRESS-GAZIT, H., FAINEKOS, G. E., AND PAPPAS, G. J. Temporal-logic-based reactive mission and motion planning. *IEEE Transactions on Robotics* 25, 6 (2009), 1370–1381.
- [31] KWIATKOWSKA, M., NORMAN, G., AND PARKER, D. Prism: Probabilistic symbolic model checker. In *International Conference on Modelling Techniques and Tools for Computer Performance Evaluation* (2002), Springer, pp. 200–204.
- [32] KWIATKOWSKA, M., NORMAN, G., AND PARKER, D. Probabilistic model checking and autonomy. *Annual Review of Control, Robotics, and Autonomous Systems* 5 (2022), 385–410.
- [33] LEE, S. C., AND LIN, X. S. Modeling dependent risks with multivariate erlang mixtures. *ASTIN Bulletin* 42, 1 (2012), 153–180.
- [34] MASSEY JR, F. J. The kolmogorov-smirnov test for goodness of fit. *Journal of the American statistical Association* 46, 253 (1951), 68–78.
- [35] MCIVER, A., AND MORGAN, C. Results on the quantitative μ -calculus $qm \mu$. *ACM Transactions on Computational Logic (TOCL)* 8, 1 (2007), 3–es.
- [36] MEAD, L. R., AND PAPANICOLAOU, N. Maximum entropy in the problem of moments. *Journal of Mathematical Physics* 25, 8 (1984), 2404–2417.
- [37] MNATSAKANOV, R. M., AND HAKOBYAN, A. S. Recovery of distributions via moments. *Lecture Notes-Monograph Series* (2009), 252–265.
- [38] PAPPAGALLO, A., MASSINI, A., AND TRONCI, E. Monte carlo based statistical model checking of cyber-physical systems: A review. *Information* 11, 12 (2020), 588.
- [39] PUTERMAN, M. L. Markov decision processes. *Handbooks in operations research and management science* 2 (1990), 331–434.
- [40] SAW, J. G., YANG, M. C., AND MO, T. C. Chebyshev inequality with estimated mean and variance. *The American Statistician* 38, 2 (1984), 130–132.
- [41] SCHMICKLER, L. Meda: mixed erlang distributions as phase-type representations of empirical distribution functions. *Communications in Statistics. Stochastic Models* 8, 1 (1992), 131–156.
- [42] TIJMS, H. *Stochastic Models: An Algorithmic Approach*. Wiley Series in Probability and Statistics - Applied Probability and Statistics Section. Wiley, 1994.
- [43] VERBELEN, R. Phase-type distributions & mixtures of erlangs. *A study of theoretical concepts, calibration techniques & actuarial applications university of leuven* (2013).
- [44] VERBELEN, R., GONG, L., ANTONIO, K., BADESCU, A., AND LIN, S. Fitting mixtures of erlangs to censored and truncated data using the em algorithm. *ASTIN Bulletin* 45, 3 (2015), 729–758.
- [45] VIRTANEN, P., GOMMERS, R., OLIPHANT, T. E., HABERLAND, M., REDDY, T., COURNAPEAU, D., BUROVSKI, E., PETERSON, P., WECKESSER, W., BRIGHT, J., ET AL. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods* 17, 3 (2020), 261–272.
- [46] WEYNS, D., IFTIKHAR, M. U., DE LA IGLESIA, D. G., AND AHMAD, T. A survey of formal methods in self-adaptive systems. In *Proceedings of the Fifth International C* Conference on Computer Science and Software Engineering* (New York, NY, USA, 2012), C3S2E ’12, Association for Computing Machinery, p. 67–79.