

Distributional Probabilistic Model Checking

Ingy Elsayed-Aly¹ , David Parker² , and Lu Feng¹ 

¹ University of Virginia, Charlottesville VA 22904, USA
`{ie3ne,lu.feng}@virginia.edu`

² University of Oxford, Oxford, UK
`david.parker@cs.ox.ac.uk`

Abstract. Probabilistic model checking provides formal guarantees for stochastic models relating to a wide range of quantitative properties, such as runtime, energy consumption or cost. But this is typically with respect to the *expected* value of these quantities, which can mask important aspects of the full probability distribution, such as the possibility of high-risk, low-probability events or multimodalities. We propose a *distributional* extension of probabilistic model checking, for discrete-time Markov chains (DTMCs) and Markov decision processes (MDPs). We formulate distributional queries, which can reason about a variety of distributional measures, such as variance, value-at-risk or conditional value-at-risk, for the accumulation of reward or cost until a co-safe linear temporal logic formula is satisfied. For DTMCs, we propose a method to compute the full distribution to an arbitrary level of precision, based on a graph analysis and forward analysis of the model. For MDPs, we approximate the optimal policy using distributional value iteration. We implement our techniques and investigate their performance and scalability across a range of large benchmark models.

1 Introduction

Computer systems are increasingly being integrated seamlessly with sensing, control and actuation of the physical world. Many of these systems (e.g., robotics) exhibit probabilistic and non-deterministic behavior due to inherent uncertainty (e.g., sensor noise, human interactions), which pose significant challenges for ensuring their safe, reliable, timely and resource-efficient execution.

Probabilistic model checking offers a collection of techniques for modelling systems that exhibit probabilistic and non-deterministic behavior. It supports not only their verification against specifications in temporal logic, but also synthesis of optimal controllers (policies). Commonly used models include discrete-time Markov chains (DTMCs) and Markov decision processes (MDPs). A range of verification techniques for these, and other models, are supported by widely used probabilistic model checkers such as PRISM [22] and Storm [11].

To capture the range of quantitative correctness specifications needed in practice, it is common to reason about *rewards* (or, conversely, *costs*). Examples include checking the worst-case execution time of a distributed coordination

algorithm, or synthesizing a controller that guarantees the minimal energy consumption for a robot to complete a sequence of tasks. Typically the *expected* value of these quantities is computed, but in some situations it is necessary to consider the full probability distribution. Notably, in safety-critical applications, it can be important to synthesize *risk-sensitive* policies, that avoid high-cost, low-probability events, which can still arise when minimizing expected cost. Risk-aware distributional measures such as *conditional value-at-risk* (CVaR) [25] address this by minimizing the costs that occur above a specified point in the tail of the distribution. Within probabilistic model checking, the use of *quantiles* has been proposed [32,28,18,16] to reason about cost or reward distributions.

In this paper, we develop a *distributional probabilistic model checking* approach, which computes and reasons about the full distribution over the reward associated with a DTMC or MDP. More precisely, we consider the reward accumulated until a specification in co-safe LTL is satisfied, the latter providing an expressive means to specify, for example, a multi-step task to be executed by a robot [19], or a sequence of events leading to a system failure. We propose a temporal logic based specification for such distributional queries.

For a DTMC, we perform model checking of these queries by generating a precise representation of the distribution, up to an arbitrary, pre-specified level of accuracy (the distribution is discrete, but often has countably infinite support, so at least some level of truncation is typically required). This is based on a graph analysis followed by a forward numerical computation. From this, we can precisely compute a wide range of useful properties, such as the mean, variance, mode or various risk-based measures.

For an MDP, we instead aim to optimize such properties over all policies. In this paper, we focus on optimizing the expected value or CVaR, whilst generating the full reward distribution for each state of the MDP. This is done using *distributional value iteration* (DVI) [3], which can be seen as a generalization of classical value iteration. Rather than computing a single scalar value (e.g., representing the optimal expected reward) for each MDP state, DVI associates a full distribution with each state, replacing the standard Bellman equation with a distributional Bellman equation.

We consider two types of DVI algorithms, namely *risk-neutral* DVI for optimizing the expected value and *risk-sensitive* DVI for optimizing CVaR. Risk-neutral DVI can be shown to converge to a deterministic, memoryless optimal policy, if a unique one exists [3]. For CVaR, memoryless policies do not suffice for optimality, but risk-sensitive DVI does converge for a product MDP that incorporates a (continuous) slack variable representing a cost/reward budget [2]. For computational tractability, we present a risk-sensitive DVI algorithm based on a discretization of the slack variable, and show that the algorithm converges to a CVaR optimal policy for increasingly precise discretizations.

For both DVI algorithms, in practice it is necessary to use approximate distributional representations. We consider the use of categorical and quantile representations. This can impact both optimality and the precision of computed distributions but, for the latter, we can construct the DTMC induced by generated

MDP policies and use our precise approach to generate the correct distribution. Finally, we implement our distributional probabilistic model checking framework as an extension of the PRISM model checker [22] and explore the feasibility and performance of the techniques on a range of benchmarks.

An extended version of this paper, with proofs, is available as [12].

1.1 Related Work

Distributional properties. Some existing probabilistic model checking methods consider distributional properties beyond expected values, notably *quantiles* [32,28,18,16], i.e., optimal reward thresholds which guarantee that the maximal or minimal probability of a reward-bounded reachability formula meets a certain threshold. While [32] and [28] focus on complexity results, [18] and [16] consider practical implementations to compute quantiles, for single- and multi-objective variants, respectively, using model unfoldings over “cost epochs”; [16] also proposes the use of interval iteration to provide error bounds. By contrast, our methods derive the full distribution, rather than targeting quantiles specifically, and our DTMC approach derives error bounds from a forward computation. We also mention [7], which computes probability distributions in a forwards manner, but for infinite-state probabilistic programs and using generating functions, and [6], which proposes an algorithm (but not implementation) to compute policies that trade off expected mean payoff and variance.

Risk-aware objectives. For MDPs, we focus in particular on *conditional value-at-risk* (CVaR). There are alternatives, such as mean-variance [31] and value-at-risk [14] but, as discussed in [25], these are not *coherent risk metrics*, which may make them unsuitable for rational decision-making. Other work on the CVaR objective includes: [20], which studies decision problem complexity, but for mean-payoff rewards and without implementations; [9], which repeatedly solves piecewise-linear maximization problems, but has limited scalability, taking over 2 hours to solve an MDP with about 3,000 states; and [26], which proposes both linear programming and value iteration methods to solve CVaR for MDPs and DTMCs. Other, not directly applicable, approaches tackle constrained problems that incorporate the CVaR objective [29,5,8]. Again, our approach differs from all these in that it computes the full distribution, allowing multiple distributional properties to be considered. We also work with temporal logic specifications. Alternative temporal logic based approaches to risk-aware control include [10], which proposes risk-aware verification of MDPs using cumulative prospect theory, and [17] which proposes chance constrained temporal logic for control of deterministic dynamical systems.

Distributional reinforcement learning. Our work is based on probabilistic model checking, which fully explores known models, but our use of DVI is inspired by *distributional reinforcement learning* [3], which can be used to learn risk-sensitive policies and improve sample efficiency (see [24] for a comparison of expected and distributional methods). We take a formal verification approach and use numerical solution, not learning, but adopt existing categorical and

quantile distributional approximations and our risk-neutral DVI algorithm is a minimization variant adapted from [3]. Risk-sensitive DVI is also sketched in [3], based on [2], but only a theoretical analysis of the method is given, without considering practical implementation aspects, such as how to discretize slack variables for computational efficiency, and how such approximations would affect the correctness of model checking. We extend risk-sensitive DVI with a discretized slack variable and show its effects theoretically in Section 4.3 and empirically via computational experiments in Section 5.

2 Background

We begin with some background on random variables, probability distributions, and the probabilistic models used in this paper. We let \mathbb{N} , \mathbb{R} , and \mathbb{Q} denote the sets of naturals, reals and rationals, respectively, and write $\mathbb{N}_\infty = \mathbb{N} \cup \{\infty\}$.

2.1 Random Variables and Probability Distributions

Let $X : \Omega \rightarrow \mathbb{R}$ be a random variable over a probability space $(\Omega, \mathcal{F}, \Pr)$. The *cumulative distribution function* (CDF) of X is denoted by $\mathcal{F}_X(x) := \Pr(X \leq x)$, and the inverse CDF is $\mathcal{F}_X^{-1}(\tau) := \inf\{x \in \mathbb{R} : \mathcal{F}_X(x) \geq \tau\}$. Common properties of interest for X include, e.g., the *expected value* $\mathbb{E}(X)$, the *variance* $\text{Var}(X)$ which is the square of the *standard deviation* (s.d.), or the *mode*.

In this paper, we also consider several *risk*-related measures. The *value-at-risk* of X at level $\alpha \in (0, 1)$ is defined by $\text{VaR}_\alpha(X) := \mathcal{F}_X^{-1}(\alpha)$, which measures risk as the minimum value encountered in the tail of the distribution with respect to a risk level α . The *conditional value-at-risk* of X at level $\alpha \in (0, 1)$ is given by $\text{CVaR}_\alpha(X) := \frac{1}{1-\alpha} \int_\alpha^1 \text{VaR}_\nu(X) d\nu$, representing the expected loss given that the loss is greater or equal to VaR_α . Figure 1a illustrates an example probability distribution of a random variable X , annotated with its expected value $\mathbb{E}(X)$, value-at-risk $\text{VaR}_{0.9}(X)$ and conditional value-at-risk $\text{CVaR}_{0.9}(X)$.

When working with the probability distributions for random variables, we write distributional equations as $X_1 \stackrel{D}{=} X_2$, denoting equality of probability laws (i.e., the random variable X_1 is distributed according to the same law as X_2). We use δ_θ to denote the Dirac delta distribution that assigns probability 1 to outcome $\theta \in \mathbb{R}$. In practice, even when distributions are discrete, we require approximate, finite representations for them. In this paper, we consider *categorical* and *quantile* distributional representations, both of which provide desirable characteristics such as tractability and expressiveness [3].

Definition 1 (Categorical representation). A categorical representation *parameterizes the probability of m atoms as a collection of evenly-spaced locations* $\theta_1 < \dots < \theta_m \in \mathbb{R}$. *Its distributions are of the form* $\sum_{i=1}^m p_i \delta_{\theta_i}$ *where* $p_i \geq 0$ *and* $\sum_{i=1}^m p_i = 1$. *We define the stride between successive atoms as* $\varsigma_m = \frac{\theta_m - \theta_1}{m-1}$.

Definition 2 (Quantile representation). A quantile representation *parameterizes the location of m equally-weighted atoms. Its distributions are of the form* $\frac{1}{m} \sum_{i=1}^m \delta_{\theta_i}$ *for* $\theta_i \in \mathbb{R}$. *Multiple atoms may share the same value.*

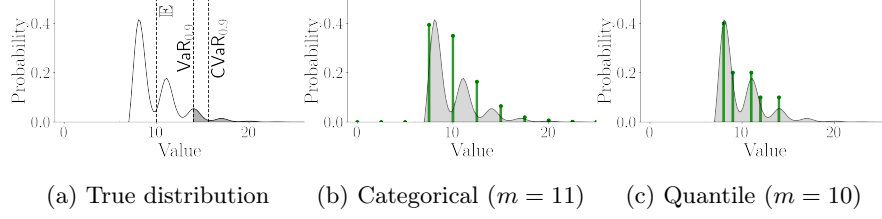


Fig. 1: An example distribution with its categorical and quantile representations.

Figures 1b and 1c show categorical and quantile representations, respectively, approximating the distribution shown in Figure 1a. When performing operations (e.g., during DVI), the intermediate result might not match the chosen representation parameters. In that case, the result is *projected* back onto the chosen representation as described in [3].

2.2 Markov Chains and Markov Decision Processes

In this paper, we work with both discrete-time Markov chains (DTMCs) and Markov decision processes (MDPs).

Definition 3 (DTMC). A discrete-time Markov chain (DTMC) is a tuple $\mathcal{D} = (S, s_0, P, AP, L)$, where S is a set of states, $s_0 \in S$ is an initial state, $P : S \times S \rightarrow [0, 1]$ is a probabilistic transition matrix satisfying $\forall s \in S : \sum_{s' \in S} P(s, s') = 1$, AP is a set of atomic propositions and $L : S \rightarrow 2^{AP}$ is a labelling function.

A DTMC \mathcal{D} evolves between states, starting in s_0 , and the probability of taking a transition from s to s' is $P(s, s')$. An (infinite) *path* through \mathcal{D} is a sequence of states $s_0 s_1 s_2 \dots$ such that $s_i \in S$ and $P(s_i, s_{i+1}) > 0$ for all $i \geq 0$, and a finite path is a prefix of an infinite path. The sets of all infinite and finite paths in \mathcal{D} are denoted $IPaths_{\mathcal{D}}$ and $FPaths_{\mathcal{D}}$, respectively. We define a probability measure $\Pr_{\mathcal{D}}$ over the set of paths $IPaths_{\mathcal{D}}$.

Definition 4 (MDP). A Markov decision process (MDP) is a tuple $\mathcal{M} = (S, s_0, A, P, AP, L)$, where states S , initial state s_0 , atomic propositions AP and labelling L are as for a DTMC, A is a finite set of actions, and $P : S \times A \times S \rightarrow [0, 1]$ is a probabilistic transition function satisfying $\forall s \in S, \forall a \in A : \sum_{s' \in S} P(s, a, s') \in \{0, 1\}$.

In each state s of an MDP \mathcal{M} , there are one or more *available* actions which can be taken, denoted $A(s) = \{a \in A \mid P(s, a, s') > 0 \text{ for some } s'\}$. If action a is taken in s , the probability of taking a transition from s to s' is $P(s, a, s')$, also denoted $P(s'|s, a)$. Paths are defined in similar fashion to DTMCs but are now alternating sequences of states and actions $s_0 a_0 s_1 a_1 s_2 \dots$ where $a_i \in A(s_i)$ and $P(s_i, a_i, s_{i+1}) > 0$ for all $i \geq 0$, and the sets of all infinite and finite paths are $IPaths_{\mathcal{M}}$ and $FPaths_{\mathcal{M}}$, respectively.

The choice of actions in each state is resolved by a *policy* (or *strategy*), based on the execution of the MDP so far. Formally, a policy takes the form $\pi : FPaths \rightarrow A$. We say that π is *memoryless* if the mapping $\pi(\omega)$ depends only on $last(\omega)$, the final state of ω , and *finite-memory* if it depends only on $last(\omega)$ and the current memory value, selected from a finite set and updated at each step of execution. The set of all policies for MDP \mathcal{M} is denoted $\Sigma_{\mathcal{M}}$.

Under a given policy π , the resulting set of (infinite) paths has, as for DTMCs, an associated probability measure, which we denote $\Pr_{\mathcal{M}}^{\pi}$. Furthermore, for both memoryless and finite-memory policies, we can build a (finite) *induced DTMC* which is equivalent to \mathcal{M} acting under π .

Definition 5 (Reward structure). A reward structure is, for a DTMC \mathcal{D} , a function $r : S \rightarrow \mathbb{N}$ and, for an MDP \mathcal{M} , a function $r : S \times A \rightarrow \mathbb{N}$.

For consistency with the literature on probabilistic model checking and temporal logics, we use the terminology *rewards* although in practice these can (and often do) represent *costs*, such as time elapsed or energy consumed. For the purposes of our algorithms, we assume that rewards are integer-valued, but we note that these could be defined as rationals, using appropriate scaling. For an infinite path ω , we also write $r(\omega, k)$ for the sum of the reward values over the first k steps of the path, i.e., $r(s_0s_1s_2\dots, k) = \sum_{i=0}^{k-1} r(s_i)$ for a DTMC and $r(s_0a_0s_1a_1s_2\dots, k) = \sum_{i=0}^{k-1} r(s_i, a_i)$ for an MDP.

To reason about rewards, we define random variables over the executions (infinite paths) of a model, typically defined as the total reward accumulated along a path, up until some event occurs. Formally, for a DTMC \mathcal{D} , such a random variable is defined as a function of the form $X : IPaths_{\mathcal{D}} \rightarrow \mathbb{R}$, with respect to the probability measure $\Pr_{\mathcal{D}}$ over $IPaths_{\mathcal{D}}$. For an MDP \mathcal{M} and policy $\pi \in \Sigma_{\mathcal{M}}$, a random variable is defined as a function $X : IPaths_{\mathcal{M}} \rightarrow \mathbb{R}$, with respect to the probability measure $\Pr_{\mathcal{M}}^{\pi}$.

3 Distributional Probabilistic Model Checking

We formulate our approach as a *distributional* extension of probabilistic model checking, which is a widely used framework for formally specifying and verifying quantitative properties of probabilistic models. In particular, we build upon existing temporal logics in common use. The core property we consider is the probability distribution over the amount of reward (or cost) that has been accumulated until some specified sequence of events occurs (which could constitute, for example, the successful completion of a task by a robot).

To represent events, we use the co-safe fragment [21] of linear temporal logic (LTL) [27]. LTL formulae are evaluated over infinite paths of a model labelled with atomic propositions from the set AP but, for use with cumulative reward, we restrict our attention to the *co-safe* fragment, containing formulae which are satisfied in finite time. Formally, this means any satisfying path ($\omega \models \psi$) has a *good prefix*, i.e., a finite path prefix ω' such that $\omega'\omega'' \models \psi$ for any suffix ω'' .

The key ingredient of our temporal logic specifications is a *distributional query*, which gives a property (such as the expected value, or variance) of the distribution over the accumulated reward until an event’s occurrence.

Definition 6 (Distributional query). *For a DTMC, a distributional query takes the form $\mathbf{R}_{=?}^{f(r)}[\psi]$, where r is a reward structure, f is a random variable property (e.g., \mathbb{E} , Var , s.d., mode, VaR , CVaR), and ψ is a formula in co-safe LTL.*

Examples of distributional queries for a DTMC are:

- $\mathbf{R}_{=?}^{\text{Var}(r_{\text{energy}})}[\mathbf{F}(goal_1 \wedge \mathbf{F} goal_2)]$ – the variance in energy consumption until a robot visits location $goal_1$ followed by location $goal_2$;
- $\mathbf{R}_{=?}^{\text{mode}(r_{\text{coll}})}[\mathbf{F} sent_1 \vee \mathbf{F} sent_2]$ – the most likely number of packet collisions before a communication protocol successfully sends one of two messages.

For an MDP, the goal is to optimize a random variable property f over its policies, which we call *distributional optimization queries*. In this paper, we focus on two particular cases, expected value (\mathbb{E}) and conditional value-at-risk (CVaR),

Definition 7 (Distributional optimization query). *For an MDP, a distributional optimization query takes the form $\mathbf{R}_{\text{opt}=?}^{f(r)}[\psi]$, where r is a reward structure, $f \in \{\mathbb{E}, \text{CVaR}\}$, $\text{opt} \in \{\min, \max\}$ and ψ is a formula in co-safe LTL. For the resulting policy, we can perform policy evaluation on the induced DTMC using one or more other distributional queries $\mathbf{R}_{=?}^{f'(r')}[\psi']$.*

An example optimization query is $\mathbf{R}_{\min=?}^{\text{CVaR}_{0.9}(r_{\text{time}})}[\mathbf{F} goal]$, which minimizes the conditional value-at-risk with respect to the time for a robot to reach its goal.

Semantics. A distributional query $\mathbf{R}_{=?}^{f(r)}[\psi]$ is evaluated on a DTMC \mathcal{D} , and a distributional optimization query $\mathbf{R}_{\text{opt}=?}^{f(r)}[\psi]$ on an MDP \mathcal{M} , in each case via a random variable for the reward accumulated from its initial state:

$$\begin{aligned} \mathbf{R}_{=?}^{f(r)}[\psi] &= f(X_{\mathcal{D}}^{r,\psi}) \\ \mathbf{R}_{\min=?}^{f(r)}[\psi] &= \inf_{\pi \in \Sigma_{\mathcal{M}}} f(X_{\mathcal{M},\pi}^{r,\psi}) \quad \text{or} \quad \mathbf{R}_{\max=?}^{f(r)}[\psi] = \sup_{\pi \in \Sigma_{\mathcal{M}}} f(X_{\mathcal{M},\pi}^{r,\psi}) \end{aligned}$$

where the random variables $X_{\mathcal{D}}^{r,\psi} : \text{IPaths}_{\mathcal{D}} \rightarrow \mathbb{R}$, $X_{\mathcal{M},\pi}^{r,\psi} : \text{IPaths}_{\mathcal{M}} \rightarrow \mathbb{R}$ are:

$$X_{\mathcal{D}}^{r,\psi}(\omega) = X_{\mathcal{M},\pi}^{r,\psi}(\omega) = \{r(\omega, k_{\psi}) - 1 \text{ if } \omega \models \psi; \infty \text{ otherwise}\}$$

and $k_{\psi} = \min\{k \mid (\omega, k) \models \psi\}$ is the length of the shortest good prefix for ψ .

Example 1. We illustrate our framework with an example of an autonomous robot navigating within a risky environment modelled as an MDP (Figure 2). The robot starts in the leftmost location (blue circle), and may pass through two types of terrain, mud (orange zigzag) and ground littered with nails (purple hatching). The default cost of navigation is 1 per step, obstacles (gray) incurring a cost of 35. In the “nails” terrain, there is a probability of 0.2 incurring a cost of

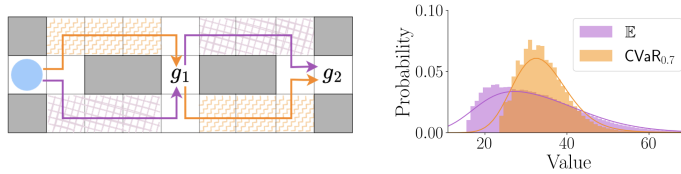


Fig. 2: The “mud & nails” example. Left: Map of the terrain to navigate, with two policies that minimize expected cost and conditional value-at-risk to visit g_1 and then g_2 . Right: The corresponding distributions over cost.

5; the “mud” terrain is safer but slower: a fixed cost of 3 per step. Consider the total cost to visit g_1 and then g_2 . Given a reward structure $cost$ encoding the costs as above, we can aim to minimize either the expected cost or the conditional value-at-risk, using queries $\mathbf{R}_{\min=?}^{\mathbb{E}(cost)}[\mathbf{F}(g_1 \wedge \mathbf{F} g_2)]$ or $\mathbf{R}_{\min=?}^{\text{CVaR}_{0.7}(cost)}[\mathbf{F}(g_1 \wedge \mathbf{F} g_2)]$. Figure 2 also shows the resulting policies, plotted on the map in purple and orange, respectively, and the corresponding probability distributions over cost. We can analyze each policy with further distributional queries, e.g., $\mathbf{R}_{=?}^{\mathbf{f}(cost)}[\mathbf{F} g_1]$ for $\mathbf{f} = \{\mathbb{E}, \text{Var}\}$ to evaluate the mean and variance of the cost to reach g_1 . ■

4 Distributional Model Checking Algorithms

We now describe algorithms for distributional probabilistic model checking, i.e., to evaluate distributional queries of the form $\mathbf{R}_{=?}^{\mathbf{f}(r)}[\psi]$ for a DTMC or $\mathbf{R}_{\text{opt}=?}^{\mathbf{f}(r)}[\psi]$ for an MDP. Following the semantics given in Section 3, for a DTMC \mathcal{D} , this necessitates generating the probability distribution of the random variable $X_{\mathcal{D}}^{r,\psi}$, corresponding to reward structure r and LTL formula ψ , on \mathcal{D} . The value $\mathbf{f}(X_{\mathcal{D}}^{r,\psi})$ can then be evaluated on the distribution for any \mathbf{f} . For an MDP \mathcal{M} , we aim to find a policy π^* which optimizes the value $\mathbf{f}(X_{\mathcal{M},\pi}^{r,\psi})$ over policies π .

For both classes of model, in standard fashion, we reduce the problem to the simpler case where ψ is a *reachability* formula by constructing an automaton product. More precisely, we build a deterministic finite automaton (DFA) \mathcal{A}_ψ representing the “good” prefixes of co-safe LTL formula ψ , and then construct a DTMC-DFA product $\mathcal{D} \otimes \mathcal{A}_\psi$ or MDP-DFA product $\mathcal{M} \otimes \mathcal{A}_\psi$ with state space $S \times Q$, where S is the state space of the original model and Q the states of the DFA. There is a one-to-one correspondence between paths (and, for MDPs, policies) in the original model and the product model [1].

Hence, in what follows, we restrict our attention to computing the probability distributions for random variables defined as the reward to reach a target set of states $T \subseteq S$, describing first the case for a DTMC and then the cases for risk-neutral ($\mathbf{f} = \mathbb{E}$) and risk-sensitive ($\mathbf{f} = \text{CVaR}$) optimization for an MDP. For the latter two, for presentational simplicity, we focus on the case of minimization, but it is straightforward to adapt the algorithms to the maximizing case.

Algorithm 1: Forward distribution generation for DTMCs

Input : DTMC $\mathcal{D} = (S, s_0, P, AP, L)$, rewards r , target $T \subseteq S$, accuracy $\varepsilon \in \mathbb{R}_{>0}$
Output: The discrete probability distribution μ for $X_{\mathcal{D}}^{r,F T}$.

```

1  $S_{\infty} \leftarrow \{s \in S \mid s \in \text{BSCC } C \subseteq S \text{ with } C \cap T = \emptyset\}; \mu_{\times} \leftarrow \delta_{(s_0,0)}; p_{\overline{T}} = 1; p_{\infty} = 0$ 
2 while  $p_{\overline{T}} - p_{\infty} > \varepsilon$  do
3    $\mu'_{\times} \leftarrow \{\}; p_{\overline{T}} \leftarrow 0$ 
4   for  $((s, i) \mapsto p_{s,i}) \in \mu_{\times}$  do
5     if  $s \in T$  then
6        $\mu'_{\times}(s, i) \leftarrow \mu'_{\times}(s, i) + p_{s,i}$ 
7     else
8       for  $(s' \mapsto p_{s'}) \in P(s, \cdot)$  do
9         if  $s' \notin T$  then
10           $p_{\overline{T}} \leftarrow p_{\overline{T}} + p_{s,i} \cdot p_{s'}$ 
11         if  $s' \notin S_{\infty}$  then
12           $\mu'_{\times}(s', i + r(s)) \leftarrow \mu'_{\times}(s', i + r(s)) + p_{s,i} \cdot p_{s'}$ 
13         else
14           $p_{\infty} \leftarrow p_{\infty} + \mu_{\times}(s, i) \cdot p_{s'}$ 
15    $\mu_{\times} \leftarrow \mu'_{\times}$ 
16 return  $\{i \mapsto p_i \mid p_i = \sum_s \mu_{\times}(s, i)\} \cup \{\infty \mapsto p_{\infty}\}$ 
    
```

4.1 Forward Distribution Generation for DTMCs

We fix a DTMC \mathcal{D} , reward structure r and set of target states T . In this section, we describe how to compute the probability distribution for the reward r accumulated in \mathcal{D} until T is reached, i.e., for the random variable $X_{\mathcal{D}}^{r,F T}$. We denote this distribution by μ . Note that, since individual rewards are integer-valued, and are summed along paths, μ is a discrete distribution.

We compute the distribution in a forward manner, up to a pre-specified accuracy ε , using Algorithm 1. First, note that the reward accumulated along a path that never reaches the target T is defined to be ∞ (see Section 3). Probabilistic model checking algorithms typically compute the *expected* reward to reach a target T from a state s , which is therefore infinite if s has a non-zero probability of not reaching T . Here, we have to take slightly more care since there may be states from which there is a non-zero probability of both accumulating finite and infinite reward. This means that μ is a distribution over \mathbb{N}_{∞} .

Algorithm 1 first identifies the states S_{∞} of \mathcal{D} from which the probability of accumulating infinite reward is 1, which are those in bottom strongly connected components (BSCCs) of \mathcal{D} that do not intersect with T . It then computes a discrete distribution μ_{\times} over $S \times \mathbb{N}_{\infty}$ where, at the k th iteration, $\mu_{\times}(s, i)$ is the probability of being in state s and having accumulated reward i after k steps. A new version μ'_{\times} is computed at each step. Abusing notation, we write distributions as lists $\{x_1 \mapsto p_1, \dots\}$ of the elements x_j of their support and their probabilities p_j . We also keep track of the probabilities $p_{\overline{T}}$ and p_{∞} of, by the k th iteration, *not* having reached the target set T and being in S_{∞} , respectively.

Algorithm 2: Risk-Neutral Distributional Value Iteration

Input : MDP $\mathcal{M} = (S, s_0, A, P, AP, L)$, rewards r , target $T \subseteq S$, and $\epsilon \in \mathbb{R}_{>0}$
Output: optimal policy π^* for query $\mathbb{R}_{\min=?}^{\mathbb{E}(r)}[\mathbf{F}T]$, distribution μ_{s_0} under π^*

- 1 $e = \infty; \mu_s \leftarrow \delta_0, \forall s \in S$
- 2 **while** $e > \epsilon$ **do**
- 3 **foreach** $s \in S \setminus T$ **do**
- 4 **foreach** $a \in A(s)$ **do**
- 5 $\eta(s, a) := \mathbb{D} \text{proj}(r(s, a) + \sum_{s' \in S} P(s, a, s') \cdot \mu_{s'})$
- 6 $\pi^*(s) \leftarrow \arg \min_{a \in A(s)} \mathbb{E}(X|X \sim \eta(s, a)) ; \mu'_s \leftarrow \eta(s, \pi^*(s))$
- 7 $e \leftarrow \sup_{s \in S \setminus T} d(\mu_s, \mu'_s)$
- 8 $\mu_s \leftarrow \mu'_s, \forall s \in S$
- 9 **return** π^* and μ_{s_0}

The distribution μ is finally computed by summing $\mu_{\times}(s, i)$ values over all states and can be analyzed with additional distributional properties.

Correctness and convergence. Let μ be the exact distribution for $X_{\mathcal{D}}^{r, \mathbf{F}T}$ and $\hat{\mu}$ be the one returned by Algorithm 1, using accuracy $\epsilon > 0$. We have:

$$\mu(i) \leq \hat{\mu}(i) \leq \mu(i) + \epsilon \quad \text{for all } i \in \mathbb{N}_{\infty} \quad (1)$$

Note that the support of μ may be (countably) infinite, but $\hat{\mu}$ is finite by construction. In this case, the total truncation error is also bounded by ϵ : if $\hat{k} \in \mathbb{N}$ is the maximum finite value in the support of $\hat{\mu}$, then $\sum_{\hat{k} < i < \infty} \mu(i) \leq \epsilon$.

To see the correctness of Equation (1), observe that $\hat{\mu}(i)$ is ultimately computed from the sum of the values $\sum_s \mu_{\times}(s, i)$ in Algorithm 1, the total value of which is non-decreasing since rewards are non-negative. In any iteration, at most $p_{\overline{T}} - p_{\infty}$ will be added to any value $\mu_{\times}(s, i)$ and, on termination, $p_{\overline{T}} - p_{\infty} \leq \epsilon$. Convergence is guaranteed for any $\epsilon > 0$: since we separate the states S_{∞} in non-target BSCCs, within k iterations, the combined probability of having reached T (i.e., $1 - p_{\overline{T}}$) or reaching S_{∞} (i.e., p_{∞}) tends to 1 as $k \rightarrow \infty$.

4.2 Risk-Neutral Distributional Value Iteration for MDPs

In this section, we present a *risk-neutral* DVI method, for computing value distributions of states of an MDP \mathcal{M} under an optimal policy that minimizes the *expected* cumulative reward to reach a target set $T \subseteq S$, i.e., minimizes $\mathbb{E}(X_{\mathcal{M}, \pi}^{r, \mathbf{F}T})$ for random variables $X_{\mathcal{M}, \pi}^{r, \mathbf{F}T}$ of MDP policies π . In contrast to the case for DTMCs, we now assume that there exists an optimal policy with finite expected reward, i.e., which reaches the target set T with probability 1. This can be checked efficiently with an analysis of the underlying graph of the MDP [4].

The risk-neutral DVI method is shown in Algorithm 2. For each MDP state $s \in S$, it initializes its value distribution μ_s to Dirac distribution δ_0 . The algorithm loops through lines 2-8 to update value distributions of any non-target state $s \in S \setminus T$ as follows. For each available action $a \in A(s)$ in state s , a value distribution is obtained via the distributional Bellman equation shown in line

5 then projected to $\eta(s, a)$ to match the chosen representation (see Sec. 2.1). The optimal action $\pi^*(s)$ in state s is the one that achieves the minimal expected value of $\eta(s, a)$. The updated value distribution μ'_s of state s is given by $\eta(s, \pi^*(s))$. The algorithm terminates when the supremum of distributional distance $d(\mu_s, \mu'_s)$ across all states (the choice of metrics is discussed below) is less than the convergence threshold ϵ . Unlike the accuracy ε for Algorithm 1, this threshold ϵ does *not* provide a guarantee on the precision of the result after convergence (similar issues occur in classical value iteration for MDPs [15]).

Distributional approximation. To enable a practical implementation of the algorithm, we need a probability distribution representation with finitely many parameters to store value distributions in memory. Here, we can adopt the categorical (see Definition 1) or quantile (see Definition 2) representations. Specifically, we need to apply the categorical or quantile projection (see [3]) after each update of the distributional Bellman equation (line 5). We use the supremum Cramér distance $\bar{\ell}_2$ for categorical representations and the supremum Wasserstein distance \bar{w}_1 for quantile representations as the distance metric in line 7 (see [3] for distributional distance definitions).

Policy convergence. When there exists a unique risk-neutral optimal policy, Algorithm 2 is guaranteed to converge to it (following [3, Theorem 7.9]). However, when there are multiple optimal policies, risk-neutral DVI may fail to converge (see [3, Section 7.5]). Furthermore, inaccuracies due the use of distributional approximations could potentially lead to a sub-optimal policy being chosen. To mitigate this, for either categorical or quantile representations, increasing the number m of atoms used yields tighter approximation error bounds [3].

4.3 Risk-Sensitive Distributional Value Iteration for MDPs

By contrast to risk-neutral policies that seek to minimize the expected reward, *risk-sensitive* policies make decisions accounting for risk properties. In this section, we present a risk-sensitive DVI method for minimizing the *conditional value-at-risk* of reaching a target set in an MDP \mathcal{M} , i.e., minimizing $\text{CVaR}_\alpha(X_{\mathcal{M}, \pi}^{r, F, T})$ for random variables $X_{\mathcal{M}, \pi}^{r, F, T}$ of MDP policies π . Our method follows a key insight from [2,30] that conditional value-at-risk can be represented as the solution of a convex optimization problem.

Lemma 1 (Dual Representation of CVaR [2,30]). *Let $[x]^+$ denote the function that is 0 if $x < 0$, and x otherwise. Given a random variable X over the probability space $(\Omega, \mathcal{F}, \text{Pr})$, it holds that:*

$$\text{CVaR}_\alpha(X) = \min_{b \in \mathbb{R}} \left\{ b + \frac{1}{1 - \alpha} \mathbb{E}([X - b]^+) \right\}, \quad (2)$$

and the minimum-point is given by $b^* = \text{VaR}_\alpha(X)$. □

Intuitively, the *slack variable* $b \in [V_{\min}, V_{\max}]$ encodes the risk budget and possible $\text{VaR}_\alpha(X)$ values. Since $\text{VaR}_\alpha(X) \in [V_{\min}, V_{\max}]$, the slack variable is similarly bounded by the minimum and maximum possible accumulated reward

Algorithm 3: Risk-Sensitive Distributional Value Iteration

Input : MDP $\mathcal{M} = (S, s_0, A, P, AP, L)$, reward structure r , target set $T \subseteq S$, risk level α , slack variable set B , convergence threshold $\epsilon \in \mathbb{R}_{>0}$

Output: optimal policy π^* for query $\mathbb{R}_{\min=?}^{\text{CVaR}_\alpha(r)}[FT]$, distribution μ_{s_0} under π^*

- 1 Construct product MDP $\mathcal{M}^b = (S \times B, \{s_0\} \times B, A, P^b, AP, L^b)$
- 2 $\mu_{\langle s, b \rangle} \leftarrow \delta_0, \forall \langle s, b \rangle \in S \times B$
- 3 **while** $e > \epsilon$ **do**
- 4 **foreach** $\langle s, b \rangle \in (S \setminus T) \times B$ **do**
- 5 **foreach** $a \in A(s)$ **do**
- 6 $\eta(\langle s, b \rangle, a) \stackrel{D}{=} \text{proj}(r(s, a) + \sum_{\langle s', b' \rangle \in S \times B} P^b(\langle s, b \rangle, a, \langle s', b' \rangle) \cdot \mu_{\langle s', b' \rangle})$
- 7 $\pi^b(\langle s, b \rangle) \leftarrow \arg \min_{a \in A(s)} \mathbb{E}([X - b]^+ | X \sim \eta(\langle s, b \rangle, a))$
- 8 $\mu'_{\langle s, b \rangle} \leftarrow \eta(\langle s, b \rangle, \pi^b(\langle s, b \rangle))$
- 9 $e \leftarrow \sup_{\langle s, b \rangle \in (S \setminus T) \times B} d(\mu_{\langle s, b \rangle}, \mu'_{\langle s, b \rangle})$
- 10 $\mu_{\langle s, b \rangle} \leftarrow \mu'_{\langle s, b \rangle}, \forall \langle s, b \rangle \in (S \setminus T) \times B$
- 11 $\bar{b}^* \leftarrow \arg \min_{\bar{b} \in B} \text{CVaR}_\alpha(X | X \sim \mu_{\langle s_0, \bar{b} \rangle}), \forall \bar{b} \in B$
- 12 $\pi^* \leftarrow$ policy π^b of the product MDP \mathcal{M}^b with initial state fixed to $\langle s_0, \bar{b}^* \rangle$
- 13 **return** π^* and $\mu_{\langle s_0, \bar{b}^* \rangle}$

within the MDP, respectively. We assume that the reward values are bounded and the probability of reaching the target states is 1, therefore V_{\min} and V_{\max} are also bounded. To enable efficient computation, we consider a discrete number of values for b . More precisely, we define a set B with n evenly-spaced atoms $b_1 < \dots < b_n$ such that $b_1 = V_{\min}$, $b_n = V_{\max}$, and the stride between two successive atoms is $\zeta_n = \frac{V_{\max} - V_{\min}}{n-1}$. Based on Lemma 1, determining the optimal slack variable value b^* requires computation of VaR_α for the distribution, which cannot be obtained *a priori*. Thus, we consider all possible risk budgets.

Algorithm 3 illustrates the proposed method. We construct a product MDP model $\mathcal{M}^b = (S \times B, \{s_0\} \times B, A, P^b, AP, L^b)$. Unlike the product MDP defined in Section 3, this MDP has multiple initial states, one state $\langle s_0, \bar{b} \rangle$ for each risk budget $\bar{b} \in B$, where s_0 is the initial state of the MDP \mathcal{M} . For each transition $s \xrightarrow{a} s'$ in \mathcal{M} with $P(s, a, s') > 0$, there is a corresponding transition $\langle s, b \rangle \xrightarrow{a} \langle s', b' \rangle$ in \mathcal{M}^b , where b' is obtained by rounding down the value of $b - r(s, a)$ to the nearest smaller atom in B and $P^b(\langle s, b \rangle, a, \langle s', b' \rangle) = P(s, a, s')$. The labelling function is given by $L^b(\langle s, b \rangle) = L(s)$. Next, in lines 2-12, Algorithm 3 initializes and updates the value distribution of each augmented state $\langle s, b \rangle \in S \times B$ in the product MDP \mathcal{M}^b in a similar fashion to the risk-neutral DVI described in Section 4.2. However, when choosing the optimal action (line 8), Algorithm 3 adopts a different criterion that minimizes $\mathbb{E}([X - b]^+)$ based on the dual representation of CVaR (see Equation 2).

Different choices of the initial risk budget \bar{b} lead to various value distributions. Once DVI on the product MDP \mathcal{M}^b converges, the algorithm selects the optimal risk budget, denoted by \bar{b}^* , that yields the minimum CVaR of all possible initial value distributions $\mu_{\langle s_0, \bar{b} \rangle}$. Finally, the algorithm returns the optimal policy π^*

resulting from the risk-sensitive DVI on the product MDP \mathcal{M}^b with initial state $\langle s_0, \bar{b}^* \rangle$, and returns the distribution $\mu_{\langle s_0, \bar{b}^* \rangle}$.

Correctness and convergence. Following [2, Theorem 3.6], when the slack variable b is continuous (i.e., $B = \mathbb{R}$), there exists a solution b^* of Equation 2 and the optimal policy π^b of product MDP \mathcal{M}^b with initial state fixed to $\langle s_0, b^* \rangle$ is the CVaR optimal policy of MDP \mathcal{M} . Algorithm 3, which uses a discretized slack variable (i.e., the set of atoms B is finite), converges to the same optimal policy π^b as $|B|$ increases, which is formalised below (and a proof can be found in the extended version of this paper [12]).

Lemma 2. *Let π_1 denote the optimal policy for minimizing $\text{CVaR}_\alpha(X_{\mathcal{M}, \pi}^{r, F^T})$, which is obtained with a continuous slack variable. Let π_2 denote the optimal policy returned by Algorithm 3 where B is a finite set of n evenly-spaced atoms with stride ς_n . It holds that $\text{CVaR}_\alpha(X_{\mathcal{M}, \pi_2}^{r, F^T}) - \text{CVaR}_\alpha(X_{\mathcal{M}, \pi_1}^{r, F^T}) = \mathcal{O}(\varsigma_n)$. As ς_n tends to 0 (i.e., $|B|$ increases), π_2 converges to the CVaR optimal policy. \square*

5 Experiments

We built and evaluated a prototype implementation³ of our distributional probabilistic model checking approach based on PRISM [22], extending its Java explicit-state engine. Our evaluation focuses initially on solving MDPs using the DVI methods (of Sections 4.2 and 4.3), then on solving the resulting policies using the DTMC method (of Section 4.1). All experiments were run on a machine with an AMD Ryzen 7 CPU and 14 GB of RAM allocated to the JVM. We set $V_{\min} = 0$ for all case studies; V_{\max} varies, as detailed below.

5.1 Case Studies

Betting Game. This case study is taken from [29]. The MDP models an agent with an amount of money, initially set to 5, which can repeatedly place a bet of amount $0 \leq \lambda \leq 5$. The probability of winning is 0.7, the probability of losing is 0.25, and the probability of hitting a jackpot (winning 10λ) is 0.05. The game ends after 10 stages. The reward function is given by the maximal allowance (e.g., 100) minus the final amount of money that the agent owns. We use $V_{\max} = 100$.

Deep Sea Treasure. This case study is also taken from [29]. The model represents a submarine exploring an area to collect one of several treasures. At each time step, the agent chooses to move to a neighbouring location; it succeeds with probability 0.6, otherwise moves to another adjacent location with probability 0.2. The agent stops when it finds a treasure or has explored for 15 steps. The reward function is defined based on the travel cost (5 per step) and opportunity cost (i.e., maximal treasure minus collected treasure value). We set $V_{\max} = 800$.

Obstacle. This case study is inspired by the navigation example in [9]. We consider an MDP model of an $N \times N$ gridworld with a set of scattered obstacles.

³ Code and models are at <https://www.prismmodelchecker.org/files/nfm24dpmc>.

Table 1: Experimental results: Timing and accuracy of each method.

Model	Method	MDP	Time (s)	E	CVaR $_{\alpha}$	Time _{dtmc} (s)	$\Delta_E^{\%}$	$\Delta_{CVaR}^{\%}$
Betting Game	risk-neut. VI	$8.9 \cdot 10^2$	< 1	61.9	-	< 1	-	-
	risk-neut. DVI	$8.9 \cdot 10^2$	< 1	61.9	98.0	< 1	0.0	0.0
	risk-sens. DVI	$9.0 \cdot 10^4$	36	85.3	92.2	< 1	0.0	0.0
DS Treasure	risk-neut. VI	$1.2 \cdot 10^3$	< 1	359.3	-	< 1	-	-
	risk-neut. DVI	$1.2 \cdot 10^3$	< 1	359.3	474.6	< 1	0.0	0.33
	risk-sens. DVI	$1.2 \cdot 10^5$	72	370.1	458.6	< 1	0.0	0.32
Obstacle ($N = 150$)	risk-neut. VI	$2.3 \cdot 10^4$	< 1	402.8	-	1,838	-	-
	risk-neut. DVI	$2.3 \cdot 10^4$	97	402.7	479.2	1,838	0.01	1.95
	risk-sens. DVI	$2.3 \cdot 10^6$	15,051	402.9	478.4	1,673	0.01	2.00
UAV	risk-neut. VI	$1.7 \cdot 10^4$	< 1	124.1	-	< 1	-	-
	risk-neut. DVI	$1.7 \cdot 10^4$	4	123.8	168.8	< 1	0.2	0.47
	risk-sens. DVI	$1.7 \cdot 10^6$	2,366	134.9	169.1	< 1	0.0	0.01
Energy ($N = 15$)	risk-neut. VI	$2.6 \cdot 10^4$	10	184.3	-	251	-	-
	risk-neut. DVI	$2.6 \cdot 10^4$	108	184.0	382.0	234	0.17	0.47
	risk-sens. DVI	$1.3 \cdot 10^6$	9,384	184.6	380.9	122	0.16	0.33

The agent’s goal is to navigate to a destination, while avoiding obstacles which cause a delay. At each time step, the agent moves in a selected direction with probability 0.9 and an unintended direction with probability 0.1. The reward function is given by the time spent to reach the destination. We use $V_{\max} = 600$.

UAV. This case study is adapted from the MDP model of the interaction between a human and an unmanned aerial vehicle (UAV) from [13]. A UAV performs road network surveillance missions with the assistance of a human operator, and is given a mission specified with LTL formula $\psi = (\mathbf{F} w_2) \wedge (\mathbf{F} w_5) \wedge (\mathbf{F} w_6)$, which translates into covering waypoints w_2 , w_5 and w_6 in any order. The reward function is given by the mission completion time. We pick $V_{\max} = 500$.

Energy. This case study considers a robot navigating an $N \times N$ gridworld with energy constraints. At each time step, the robot moves to an adjacent grid location with probability 0.7, or to an unintended adjacent location otherwise. It starts with a fixed amount of energy and consumes 1 unit per step. The robot can only recharge its battery in the charging station. When the energy is depleted, the robot is transported with a delay to the charging station. The robot is asked to complete a mission specified with LTL formula $\psi = (\mathbf{F} w_1) \wedge (\mathbf{F} w_2) \wedge (\mathbf{F} w_3)$. The reward function represents the mission completion time. We use $V_{\max} = 500$.

5.2 Results Analysis

Method comparison. Table 1 summarizes our experimental results across the benchmarks described above. For each MDP, we run both the risk-neutral and risk-sensitive variants of distributional value iteration (DVI), optimizing expected value and CVaR, as described in Section 4.2 and Section 4.3, respectively. For the risk neutral case we also run standard value iteration (VI), as implemented in PRISM. For all three methods, we then evaluate the resulting policy,

computing the full reward distribution using the forward distribution generation method described in Section 4.1, allowing us to compute more precise results for the expected value and CVaR on those policies.

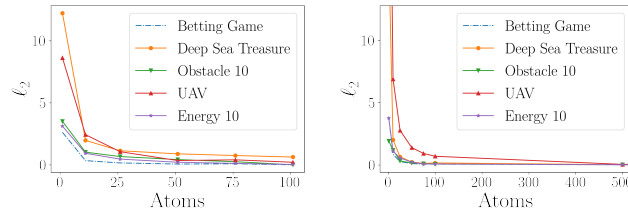
The table shows the time to run each algorithm and the values computed during optimization (the value for the objective being optimized is shown in bold). Additionally, the table shows the time to run the forward distribution method on the induced DTMC, and the (percentage) relative error when comparing the VI/DVI results with the forward distribution outcomes.

For each case study, we also report the number of states in the (product) MDP that is solved. The UAV and Energy benchmarks use non-trivial co-safe LTL formulae for the mission specification (the others are reachability specifications) and so the MDP is a MDP-DFA product. For risk-sensitive DVI, the state space is also augmented with a slack variable resulting in larger product MDPs. We set the slack variable size to $|B| = 51$ for the Energy model, and $|B| = 101$ for the rest. We use the categorical representation with $m = 201$ for DVI, with $\epsilon = 0.01$ for the convergence metric. For policy evaluation, we use precision $\epsilon = 10^{-3}$ for the Obstacle and Energy case studies and $\epsilon = 10^{-5}$ for the others.

Our DVI methods successfully optimize their respective objectives on a range of large MDPs. Generally, the policy resulting from the risk-neutral method has a lower expected value, while the policy from the risk-sensitive method has a lower CVaR_α , and the risk-neutral method yields the same optimal policy as baseline VI. As expected, DVI methods are more expensive than VI, since they work with distributions, but the DVI methods are successfully applied to MDPs with several million states. Additionally, the baseline VI method can only provide expected reward values, while the distribution returned by our methods can be used to compute additional distributional properties (variance, VaR, etc.). Comparing the two variants of DVI, the risk-sensitive version takes considerably longer to run. This is primarily due to the use of a larger product model, incorporating a slack variable, rather than the computation required for DVI itself. For the same reason, risk-neutral DVI scales to larger models, but for clarity Table 1 only includes models that all methods can solve.

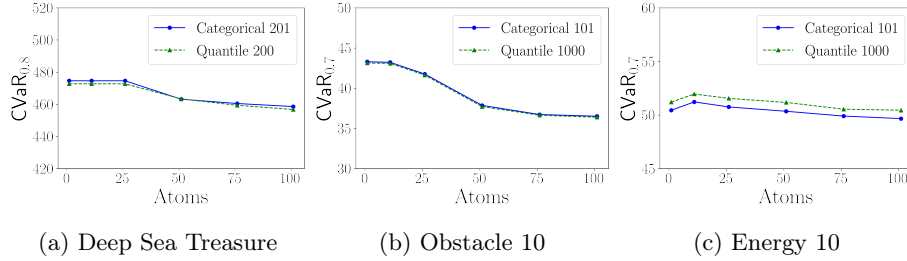
The DTMC forward computation also works on all models. It is often very fast (under a second in 3 cases), but grows expensive on models where the support of the distribution is large. From its results, we see that both DVI methods produce approximate distributions that are close to the true distribution.

Note that in the last three case studies, the V_{\max} value is higher, resulting in a larger stride and thus more coarse representations for both the value distributions and the slack variable (for risk-sensitive DVI). This results in more approximation errors when computing metrics from the value distributions generated using DVI. This can be seen in the case of the UAV model where the risk-neutral method underestimates CVaR_α (168.8 compared to 169.6 from the true distribution generated by the DTMC method for the same policy). The following experiments aim to evaluate how the parameters of the distributional representation affect the resulting approximate distributions generated by DVI.



(a) Categorical representation (b) Quantile representation

Fig. 3: Varying the numbers of atoms for distributional representations.



(a) Deep Sea Treasure

(b) Obstacle 10

(c) Energy 10

Fig. 4: Results for varying the numbers of atoms in risk-sensitive DVI.

Effects on distributional approximation. Figure 3 plots the effects of varying the number of atoms used for categorical and quantile representations for distributions, in terms of the ℓ_2 distance between the approximate distribution resulting from risk-neutral DVI and the ground truth (obtained via applying the DTMC forward distribution generation method with $\varepsilon = 10^{-5}$ on the resulting optimal policy). *For both representations, the ℓ_2 distance approaches 0 as the number of atoms increases, indicating that the approximate distributions become very close to the ground truth.* We observe similar effects with the risk-sensitive method and thus omit the resulting plot. Note that the Deep Sea Treasure model has a larger V_{\max} and thus the resulting ℓ_2 is higher than other models when using a maximum of 101 atoms in the categorical representation.

A larger number of atoms (m value) leads to a higher computational cost, thus we consider smaller models for the Obstacle and Energy case studies with $N = 10$ for plotting. As an illustration of accuracy/cost trade-off, for Energy 10, the runtime using categorical representations with 11 atoms (resp. 101 atoms) is 0.3s (resp. 0.63s), while the runtime when using quantile representations with 10 atoms (resp. 100 atoms) is 0.9s (resp. 5s). The quantile projection is more expensive than the categorical projection, resulting in higher runtimes.

Effects of slack variable atoms. Figure 4 illustrates the effects of varying the number of atoms used for the slack variables ($|B|$) in risk-sensitive DVI. *The results show that increasing $|B|$ generally leads to better policies with smaller CVaR values.* This is in part because the algorithm would check a larger set of initial risk budgets $\bar{b} \in B$. But there is a trade-off since the computational cost

Table 2: Performance comparison for DTMC forward computation

Model	Param.s	States	Transitions	V_{\max}	DVI(s)	DTMC(s)	$\Delta_{\mathbb{E}}^{\%}$	$\Delta_{\text{CVaR}}^{\%}$
EGL	N=8,L=3	$5.4 \cdot 10^6$	$5.5 \cdot 10^6$	40	439	1	0.4	0.5
	N=8,L=4	$7.5 \cdot 10^6$	$7.6 \cdot 10^6$	40	897	1	0.4	0.4
	N=8,L=5	$9.6 \cdot 10^6$	$9.7 \cdot 10^6$	50	4,345	1	0.3	0.4
Leader	N=8,K=5	$2.7 \cdot 10^6$	$3.1 \cdot 10^6$	20	41	2	0.0	0.0
	N=10,K=4	$9.4 \cdot 10^6$	$1.0 \cdot 10^7$	30	577	15	0.3	0.6
	N=8,K=6	$1.2 \cdot 10^7$	$1.3 \cdot 10^7$	20	163	9	0.0	0.1
Herman	N=13	$8.2 \cdot 10^3$	$1.6 \cdot 10^6$	100	4	14	0.6	0.9
	N=15	$3.3 \cdot 10^4$	$1.4 \cdot 10^7$	120	57	190	0.5	0.9
	N=17	$1.3 \cdot 10^5$	$1.3 \cdot 10^8$	140	1,234	2,369	0.8	1.2

grows with an increasing $|B|$. For example, in the Energy 10 model, the runtime using the categorical representation with 101 atoms for $|B| = 11$ (resp. $|B| = 101$) is 7.8s (resp. 78.6s), whereas the runtime of using the quantile representation with 1,000 atoms for $|B| = 11$ (resp. $|B| = 101$) is 477s (resp. 5,163s).

DTMC forward computation. Finally, we further evaluate the forward computation method for DTMCs from Section 4.1 on a range of common DTMC benchmarks from the PRISM benchmark suite [23]. In particular, we compare to an alternative computation using the risk-neutral DVI method method of Section 4.2, treating DTMCs as a special case of MDPs. Table 2 shows the performance of the two methods. For each model, we indicate the parameters used for the benchmark and the DTMC size (states and transitions). For the DVI method, we use the categorical representation with a stride of 1 and a value of V_{\max} large enough to represent the distribution (also shown in the table).

In two of the three models, the DTMC computation is much faster. This is because the DVI method calculates a reward distribution for every state. For the third model, where V_{\max} is significantly higher, DVI is actually faster (the same can be seen for the Obstacle and Energy models in Table 1). The DTMC method computes distribution to a pre-specified accuracy, but DVI may incur approximation errors, primarily due to convergence. Table 2 also shows (relative) errors for the expected value and CVaR metrics for each benchmark.

6 Conclusion

This paper presents a distributional approach to probabilistic model checking, which supports a rich set of distributional queries for DTMCs and MDPs. Experiments on a range of benchmark case studies demonstrate that our approach can be successfully applied to check various distributional properties (e.g., CVaR, VaR, variances) of large MDP and DTMC models. We believe that this work paves the way for applying distributional probabilistic model checking in many safety-critical and risk-averse domains. For future work, we will explore distributional queries with multiple objectives and under multi-agent environments.

Acknowledgements. This work was supported in part by NSF grant CCF-1942836 and the ERC under the European Union’s Horizon 2020 research and innovation programme (FUN2MODEL, grant agreement No. 834115).

References

1. Baier, C., Katoen, J.P.: Principles of Model Checking. MIT Press (2008)
2. Bäuerle, N., Ott, J.: Markov decision processes with average-value-at-risk criteria. *Mathematical Methods of Operations Research* **74**(3), 361–379 (2011)
3. Bellemare, M.G., Dabney, W., Rowland, M.: *Distributional Reinforcement Learning*. MIT Press (2023), <http://www.distributional-rl.org>
4. Bianco, A., de Alfaro, L.: Model checking of probabilistic and nondeterministic systems. In: Proc. 15th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS’95). LNCS, vol. 1026, pp. 499–513. Springer (1995)
5. Borkar, V., Jain, R.: Risk-constrained Markov decision processes. *IEEE Transactions on Automatic Control* **59**(9), 2574–2579 (2014)
6. Brazdil, T., Chatterjee, K., Forejt, V., Kucera, A.: Trading performance for stability in Markov decision processes. *Journal of Computer and System Sciences* **84**, 144–170 (2017)
7. Chen, M., Katoen, J., Klinkenberg, L., Winkler, T.: Does a program yield the right distribution? - verifying probabilistic programs via generating functions. In: Proc. 34th International Conference on Computer Aided Verification (CAV’22). LNCS, vol. 13371, pp. 79–101. Springer (2022)
8. Chow, Y., Ghavamzadeh, M.: Algorithms for CVaR optimization in MDPs. *Advances in neural information processing systems* **27** (2014)
9. Chow, Y., Tamar, A., Mannor, S., Pavone, M.: Risk-sensitive and robust decision-making: a cvar optimization approach. *Advances in neural information processing systems* **28** (2015)
10. Cubuktepe, M., Topcu, U.: Verification of Markov decision processes with risk-sensitive measures. In: Proc. Annual American Control Conference (ACC’18). pp. 2371–2377. IEEE (2018)
11. Dehnert, C., Junges, S., Katoen, J.P., Volk, M.: A storm is coming: A modern probabilistic model checker. In: Proc. 29th International Conference on Computer Aided Verification (CAV’17) (2017)
12. Elsayed-Aly, I., Parker, D., Feng, L.: Distributional probabilistic model checking. arXiv preprint arXiv:2309.05584 (2023)
13. Feng, L., Wiltsche, C., Humphrey, L., Topcu, U.: Synthesis of human-in-the-loop control protocols for autonomous systems. *IEEE Transactions on Automation Science and Engineering* **13**(2), 450–462 (2016)
14. Filar, J.A., Krass, D., Ross, K.W.: Percentile performance criteria for limiting average Markov decision processes. *IEEE Transactions on Automatic Control* **40**(1), 2–10 (1995)
15. Haddad, S., Monmege, B.: Reachability in MDPs: Refining convergence of value iteration. In: Proc. 8th International Workshop on Reachability Problems (RP’14). LNCS, vol. 8762, pp. 125–137. Springer (2014)
16. Hartmanns, A., Junges, S., Katoen, J.P., Quatmann, T.: Multi-cost bounded trade-off analysis in MDP. *Journal of Automated Reasoning* **64**(7), 1483–1522 (2020)

17. Jha, S., Raman, V., Sadigh, D., Seshia, S.A.: Safe autonomy under perception uncertainty using chance-constrained temporal logic. *Journal of Automated Reasoning* **60**(1), 43–62 (2018)
18. Klein, J., Baier, C., Chrszon, P., Daum, M., Dubslaff, C., Klüppelholz, S., Märcker, S., Müller, D.: Advances in probabilistic model checking with PRISM: variable re-ordering, quantiles and weak deterministic Büchi automata. *International Journal on Software Tools for Technology Transfer* **20**(2), 179–194 (2018)
19. Kress-Gazit, H., Fainekos, G.E., Pappas, G.J.: Temporal logic-based reactive mission and motion planning. *IEEE Transactions on Robotics* **25**(6), 1370–1381 (2009)
20. Křetínský, J., Meggendorfer, T.: Conditional value-at-risk for reachability and mean payoff in Markov decision processes. In: *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*. pp. 609–618 (2018)
21. Kupferman, O., Vardi, M.Y.: Model checking of safety properties. *Formal Methods in System Design* **19**(3), 291–314 (2001)
22. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) *Proc. 23rd International Conference on Computer Aided Verification (CAV'11)*. LNCS, vol. 6806, pp. 585–591. Springer (2011)
23. Kwiatkowska, M., Norman, G., Parker, D.: The PRISM benchmark suite. In: *Proc. 9th International Conference on Quantitative Evaluation of SysTems (QEST'12)*. pp. 203–204. IEEE CS Press (2012)
24. Lyle, C., Bellemare, M.G., Castro, P.S.: A comparative analysis of expected and distributional reinforcement learning. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 33, pp. 4504–4511 (2019)
25. Majumdar, A., Pavone, M.: How should a robot assess risk? Towards an axiomatic theory of risk in robotics. In: *Robotics Research: The 18th International Symposium ISRR*. pp. 75–84. Springer (2020)
26. Meggendorfer, T.: Risk-aware stochastic shortest path. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 36, pp. 9858–9867 (2022)
27. Pnueli, A.: The temporal semantics of concurrent programs. *Theoretical Computer Science* **13**, 45–60 (1981)
28. Randour, M., Raskin, J.F., Sankur, O.: Percentile queries in multi-dimensional Markov decision processes. *Formal methods in system design* **50**, 207–248 (2017)
29. Rigter, M., Duckworth, P., Lacerda, B., Hawes, N.: Planning for risk-aversion and expected value in MDPs. In: *Proceedings of the International Conference on Automated Planning and Scheduling*. vol. 32, pp. 307–315 (2022)
30. Rockafellar, R.T., Uryasev, S.: Conditional value-at-risk for general loss distributions. *Journal of banking & finance* **26**(7), 1443–1471 (2002)
31. Sobel, M.J.: The variance of discounted Markov decision processes. *Journal of Applied Probability* **19**(4), 794–802 (1982)
32. Ummels, M., Baier, C.: Computing quantiles in Markov reward models. In: *Proc. 16th International Conference on Foundations of Software Science and Computation Structures (FOSSACS'13)*. LNCS, vol. 7794, pp. 353–368. Springer (2013)