

Neuro-symbolic Action Masking for Deep Reinforcement Learning

Shuai Han
Utrecht University
Utrecht, the Netherland
s.han@uu.nl

Mehdi Dastani
Utrecht University
Utrecht, the Netherland
m.m.dastani@uu.nl

Shihan Wang
Utrecht University
Utrecht, the Netherland
s.wang2@uu.nl

ABSTRACT

Deep reinforcement learning (DRL) may explore infeasible actions during training and execution. Existing approaches assume a symbol grounding function that maps high-dimensional states to consistent symbolic representations and a manually specified action masking techniques to constrain actions. In this paper, we propose Neuro-symbolic Action Masking (NSAM), a novel framework that automatically learn symbolic models, which are consistent with given domain constraints of high-dimensional states, in a minimally supervised manner during the DRL process. Based on the learned symbolic model of states, NSAM learns action masks that rules out infeasible actions. NSAM enables end-to-end integration of symbolic reasoning and deep policy optimization, where improvements in symbolic grounding and policy learning mutually reinforce each other. We evaluate NSAM on multiple domains with constraints, and experimental results demonstrate that NSAM significantly improves sample efficiency of DRL agent while substantially reducing constraint violations.

KEYWORDS

Deep reinforcement learning, neuro-symbolic learning, action masking

ACM Reference Format:

Shuai Han, Mehdi Dastani, and Shihan Wang. 2026. Neuro-symbolic Action Masking for Deep Reinforcement Learning. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 10 pages. <https://doi.org/10.65109/JWPH6906>

1 INTRODUCTION

With the powerful representation capability of neural networks, deep reinforcement learning (DRL) has achieved remarkable success in a variety of complex domains that require autonomous agents, such as autonomous driving [23, 24, 78], resource management [25, 75], algorithmic trading [39, 70] and robotics [6, 41, 60]. However, in real-world scenarios, agents face the challenges of learning policies from few interactions [6] and keeping violations of domain constraints to a minimum during training and execution [79]. To address these challenges, an increasing number of neuro-symbolic reinforcement learning (NSRL) approaches have been proposed,

aiming to exploit the structural knowledge of the problem to improve sample efficiency [33, 55, 77] or to constrain agents to select actions [15, 66, 74].

Among these NSRL approaches, a promising practice is to exclude infeasible actions for the agents.¹ This is typically achieved by assuming a predefined symbolic grounding [57] or label function [33] that maps high-dimensional states into symbolic representations and manually specify action masking techniques [8, 45, 67]. However, predefining the symbolic grounding function is often expensive [63], as it requires complete knowledge of the environmental states, and could be practically impossible when the states are high-dimensional or infinite. Learning symbolic grounding from environmental state is therefore crucial for NSRL approaches and remains a highly challenging problem [63].

In particular, there are three main challenges. First, real-world environments should often satisfy complex constraints expressed in a domain specific language, which makes learning the symbolic grounding function difficult [1]. Second, obtaining full supervision for learning symbolic representations in DRL environments is unrealistic, as those environments rarely provide the ground-truth symbolic description of every state. Finally, even if symbolic grounding can be learned, integrating it into reinforcement learning to achieve end-to-end learning remains a challenge.

To address these challenges, we propose Neuro-symbolic Action Masking (NSAM), a framework that integrates symbolic reasoning into deep reinforcement learning. The basic idea is to use probabilistic sentential decision diagrams (PSDDs) to learn symbolic grounding. PSDDs serve two purposes: they guarantee that any learned symbolic model satisfies domain constraints expressed in a domain specific language [37], and they allow the agent to represent probability distributions over symbolic models conditioned on high-dimensional states. In this way, PSDDs bridge the gap between numerical states and symbolic reasoning without requiring manually defined mappings. Based on the learned PSDDs, NSAM combines action preconditions with the inferred symbolic model of numeric states to construct action masks, thereby filtering out infeasible actions. Crucially, this process only relies on minimal supervision in the form of action explorability feedback, rather than full symbolic description at every state. Finally, NSAM is trained end-to-end, where the improvement of symbolic grounding and policy optimization mutually reinforce each other.

We evaluate NSAM on four DRL decision-making domains with domain constraints, and compare it against a series of state-of-the-art baselines. Experimental results demonstrate that NSAM not only learns more efficiently, consistently surpassing all baselines, but



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/JWPH6906>

¹We use the term infeasible actions throughout the paper, which can also be considered as unsafe, unethical or in general undesirable actions.




Figure 1: Example states in the Visual Sudoku environment

also substantially reduces constraint violations during training. The results further show that the symbolic grounding plays a crucial role in exploiting underlying knowledge structures for DRL.

2 PROBLEM SETTING

We study reinforcement learning (RL) on a Markov Decision Process (MDP) [59] $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{T}, R, \gamma)$ where \mathcal{S} is a set of states, \mathcal{A} is a finite set of actions, $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is a transition function, $\gamma \in [0, 1]$ is a discount factor and $R : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is a reward function. An agent employs a policy π to interact with the environment. At a time step t , the agent takes action a_t according to the current state s_t . The environment state will transfer to next state s_{t+1} based on the transition probability \mathcal{T} . The agent will receive the reward r_t . Then, the next round of interaction begins. The goal of this agent is to find the optimal policy π^* that maximizes the expected return: $\mathbb{E}[\sum_{t=0}^T \gamma^t r_t | \pi]$, where T is the terminal time step.

To augment RL with symbolic domain knowledge, we extend the normal MDP with the following modules $(\mathcal{P}, \mathcal{AP}, \phi)$ where $\mathcal{P} = \{p_1, \dots, p_K\}$ is a finite set of atomic propositions (each $p \in \mathcal{P}$ represents a Boolean property of a state $s \in \mathcal{S}$), $\mathcal{AP} = \{(a, \varphi) | a \in \mathcal{A}, \varphi \in L(\mathcal{P})\}$ is the set of actions with their preconditions, and $L(\mathcal{P})$ denotes the propositional language over \mathcal{P} . We use (a, φ) to state that action a is *explorable*² in a state if and only if its precondition φ holds in that state, $\phi \in L(\mathcal{P})$ is a domain constraint. We use $[[\phi]] = \{\mathbf{m} | \mathbf{m} \models \phi\}$ to denote the set of all possible symbolic models of ϕ ³.

To illustrate how symbolic domain knowledge $(\mathcal{P}, \mathcal{AP}, \phi)$ is reflected in our formulation, we consider the Visual Sudoku task as a concrete example. In this environment, each state is represented as a non-symbolic image input. The properties of a state can be described using propositions in \mathcal{P} . For example, the properties of the state in Figure 1(a) include ‘position (1,1) is number 1’, ‘position (1,2) is empty’, etc. Each action a of filling a number in a certain position corresponds to a symbolic precondition φ , represented by $(a, \varphi) \in \mathcal{AP}$. For example, the action ‘filling number 1 at position (1,1)’ requires that both propositions ‘position (1,2) is number 1’ and ‘position (2,1) is number 1’ are false. Finally, ϕ is used to constrain the set of possible states, e.g., ‘position (1,1) is number 1’ and ‘position (1,1) is number 2’ cannot both be simultaneously true for a given state. To leverage this knowledge, challenges arise due to the following problems.

²All actions $a \in \mathcal{A}$ can in principle be chosen by the agent. However, we use the term *explorable* to distinguish actions whose preconditions are satisfied (safe, ethical, desirable actions) from those whose preconditions are not satisfied (unsafe, unethical, undesirable actions).

³a model is a truth assignment to all propositions in \mathcal{P}

(P1) Numerical–symbolic gap. Knowledge is based on symbolic property of states, but only raw numerical states are available.

(P2) Constraint satisfaction. The truth values of propositions in \mathcal{P} mapped from a DRL state s must satisfy domain constraints ϕ .

(P3) Minimal supervision. The RL environment cannot provide full ground truth of propositions at each state.

(P4) Differentiability. The symbolic reasoning with φ introduces non-differentiable process, which could be conflicting with gradient-based DRL algorithms that require differentiable policies.

(P5) End-to-end learning. Achieving end-to-end training on prediction of propositions, symbolic reasoning over preconditions and optimization of policy is challenging.

In summary, the above challenges fall into three categories. (P1–P3) concern learning symbolic models from high-dimensional states in DRL, which we address in Section 3. (P4) relates to the differentiability barrier when combining symbolic reasoning with gradient-based DRL, which we tackle in Section 4. (P5) raises the need for an end-to-end training, which we present in Section 5.

3 LEARNING SYMBOLIC GROUNDING

This section introduces how NSAM learns symbolic grounding. At a high level, the goal is to learn whether an action is explorable in a state. Specifically, the agent receives minimal supervision from state transitions after executing an action a . Using this supervision, NSAM learns to estimate the symbolic model of the high-dimensional input state that is in turn used to check the satisfiability of action preconditions. To achieve this, Section 3.1 presents a knowledge compilation step to encode domain constraints into a symbolic structure, while Section 3.2 explains how this symbolic structure is parameterized and learned from minimal supervision.

3.1 Compiling the Knowledge

To address P2 (Constraint satisfaction), we introduce the Probabilistic Sentential Decision Diagram (PSDD) [37]. PSDDs are designed to represent probability distributions $Pr(\mathbf{m})$ over possible models, where any model \mathbf{m} that violates domain constraints is assigned zero probability [53]. For example, consider the distribution in Figure 2(a). The first step in constructing a PSDD is to build a Boolean circuit that captures the entries whose probability values are always zero, as shown in Figure 2(b). Specifically, the circuit evaluates to 0 for model \mathbf{m} if and only if $\mathbf{m} \not\models \phi$. The second step is to parameterize this Boolean circuit to represent the (non-zero) probability of valid entries, yielding the PSDD in Figure 2(c).

To obtain the Boolean circuit in Figure 2(b), we represent the domain constraint ϕ using a general data structure called a Sentential Decision Diagram (SDD) [20]. An SDD is a normal form of a Boolean formula that generalizes the well-known Ordered Binary Decision Diagram (OBDD) [12, 13]. SDD circuits satisfy specific syntactic and semantic properties defined with respect to a binary tree, called a vtrees, whose leaves correspond to propositions (see Figure 2(d)). Following Darwiche’s definition [20, 52], an SDD normalized for a vtrees v is a Boolean circuit defined as follows: If v is a leaf node labeled with variable p , the SDD is either p , $\neg p$, \top , \perp , or an OR gate with inputs p and $\neg p$. If v is an internal node, the SDD has the structure shown in Figure 2(e), where $prime_1, \dots, prime_n$ are SDDs normalized for the left child v^l , and sub_1, \dots, sub_n are SDDs

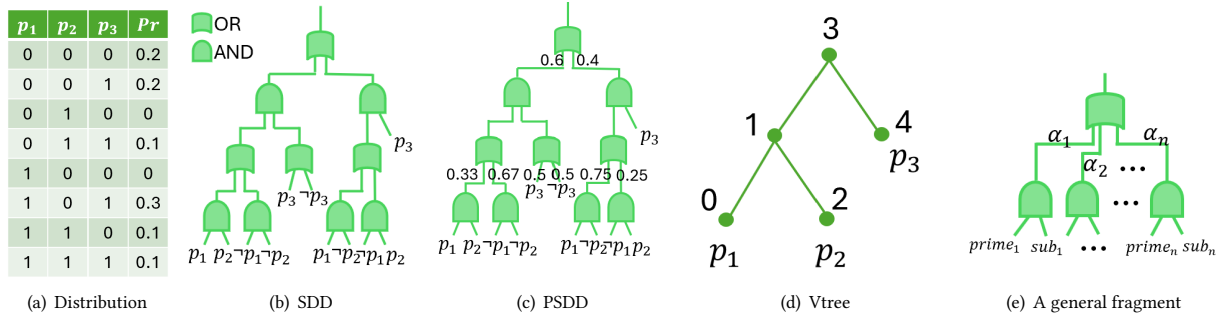


Figure 2: (a) An example of joint distribution for three propositions p_1, p_2 and p_3 with the constraint $(p_1 \leftrightarrow p_2) \vee p_3$. (b) A SDD circuit with ‘OR’ and ‘AND’ logic gate to represent the constrain $(p_1 \leftrightarrow p_2) \vee p_3$. (c) The PSDD circuit to represent the distribution in Fig. 2(a). (d) The vtree used to group variables. (e) A general fragment to show the structure of SDD and PSDD.

normalized for the right child v^r . SDD circuits alternate between OR gates and AND gates, with each AND gate having exactly two inputs. The OR gates are mutually exclusive in that at most one of their inputs evaluates to true under any circuit input [20, 52].

A PSDD is obtained by annotating each OR gate in an SDD with parameters $(\alpha_1, \dots, \alpha_n)$ over its inputs [37, 52], where $\sum_i \alpha_i = 1$ (see Figure 2(e)). The probability distribution defined by a PSDD is as follows. Let m be a model that assigns truth values to the PSDD variables, and suppose the underlying SDD evaluates to 0 under m ; then $Pr(m) = 0$. Otherwise, $Pr(m)$ is obtained by multiplying the parameters along the path from the output gate.

The key advantage of using PSDDs in our setting is twofold. First, PSDDs strictly enforce domain constraints by assigning zero probability to any model m that violates ϕ [53], thereby ensuring logical consistency (P2). Second, by ruling out impossible truth assignment through domain knowledge, PSDDs effectively reduce the scale of the probability distribution to be learned [1].

Besides, PSDDs also support tractable probabilistic queries [16, 52]. While PSDD compilation can be computationally expensive as its size grows exponentially in the number of propositions and constraints, it is a one-time offline cost. Once compilation is completed, PSDD inference is linear-time, making symbolic reasoning efficient during both training and execution [52].

3.2 Learning the parameters of PSDD in DRL

To address P1 (Numerical–symbolic gap), we need to learn distributions of models that satisfy the domain constraints. Inspired by recent deep supervised learning work on PSDDs [1], we parameterize the PSDD using the output of gating function g . This gating function is a neural network that maps high-dimensional RL states to PSDD parameters $\Theta = g(s)$. This design allows the PSDD to represent state-conditioned distributions over propositions through its learned parameters, while strictly adhering to domain constraints (via its structure defined by symbolic knowledge ϕ). The overall process is shown in Figure 3. We use $Pr(m | \Theta = g(s), m \models \phi)$ to denote the probability of model m that satisfy the domain constraints ϕ given the state s (this is calculated by PSDD in Figure 3).

After initializing g and the PSDD according to the structure in Figure 3, we obtain a distribution over m such that for all $m \not\models \phi$,




Figure 3: The architecture design to calculate the probability of symbolic model m given DRL state s .

$Pr(m | \Theta = g(s), m \not\models \phi) = 0$. However, for the probability distribution over m that does satisfy ϕ , we still need to learn from data to capture the probability of different m by adjusting parameters of gating function g . To train the PSDD from minimal supervision signals (for problem (P3) in Section 2), we construct the supervision data from Γ_ϕ , which consists of tuples (s, a, s', y) where transitions (s, a, s') are explored from the environment and y is calculated by:

$$y = \begin{cases} 1, & \text{if } s \text{ and } s' \text{ do not violate } \phi, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

That is, the action a is labeled as explorable (i.e., $y = 1$) in state s if it does not lead to a violation of the domain constraint ϕ ; otherwise the action a is not explorable (i.e., $y = 0$).

Unlike a fully supervised setting that expensively requires labeling every propositional variable in P , Eq. (1) only requires labeling whether a given state violates the domain constraint ϕ , which is a minimal supervision signal. In practice, the annotation of y can be obtained either (i) by providing labeled data on whether the resulting state s' violates the constraint ϕ [10], or (ii) via an automatic constraint-violation detection mechanism [26, 47].

We emphasize that action preconditions φ and the domain constraints ϕ are two separate elements and treated differently. We first automatically generate training data to learn PSDD parameters by constructing tuples (s, a, s', y) as defined in Equation (1). The argument y in tuples (s, a, s', y) is then used as an indicator for action preconditions. Specifically, we use y to label whether action a is executable in state s , i.e., if transition (s, a, s') is explored by DRL policy in a non-violating states s and s' , then $y = 1$, meaning that action a is explorable in s ; otherwise $y = 0$. We thus use y in

(s, a, s', y) as a minimal supervision signal to estimate the probability of the precondition of action a being satisfied in non-violating s during PSDD training.

By continuously rolling out the DRL agent in the environment, we store (s, a, s', y) into a buffer \mathcal{D} . After collecting sufficient data, we sample batches from \mathcal{D} and update g via stochastic gradient descent [11, 36]. Concretely, the update proceeds as follows. Given samples (s, a, s', y) , we first use the current PSDD to estimate the probability that action a is explorable in state s , i.e., the probability that s satisfies the precondition φ associated with a in \mathcal{AP} :

$$\hat{P}(a|s) = \sum_{\mathbf{m} \models \varphi} Pr(\mathbf{m} | \Theta = g(s), \mathbf{m} \models \varphi) \quad (2)$$

Note that $\hat{P}(a|s)$ here does not represent a policy as in standard DRL; rather, it denotes the estimated probability that action a is explorable in state s . As shown in Equation (2), this probability is calculated by aggregating the probabilities of all models \mathbf{m} that satisfy the precondition φ . In addition, to evaluate if $\mathbf{m} \models \varphi$, we assign truth values to the leaf variables of φ 's SDD circuit based on \mathbf{m} and propagate them bottom-up through the 'OR' and 'AND' gates, where the Boolean value at the root indicates the satisfiability.

Given the probability estimated from Equation (2), we compute the cross-entropy loss [21] by comparing it with the explorability label y . Specifically, for a single data (s, a, s', y) , the loss is:

$$L_g = -[y \cdot \log(\hat{P}(a|s)) + (1 - y) \cdot \log(1 - \hat{P}(a|s))] \quad (3)$$

The intuition of this loss is straightforward: at each s it encourages the PSDD to generate higher probability to actions that are explorable (when $y = 1$), and generate lower probability to those that are not explorable (when $y = 0$).

4 COMBINING SYMBOLIC REASONING WITH GRADIENT-BASED DRL

Through the training of the gating function defined in Equation (3), the PSDD in Figure 3 can predict, for a given DRL state, a distribution over the symbolic model \mathbf{m} for atomic propositions in \mathcal{P} . This distribution then can be used to evaluate the truth values of the preconditions in \mathcal{AP} and to reason about the explorability of actions. However, directly applying symbolic logical formula of preconditions to take actions results in non-differentiable decision-making [5], which prevents gradient flow during policy optimization. This raises a key challenge on integrating symbolic reasoning with gradient-based DRL training in a way that preserves differentiability, i.e., problem (P4) in Section 2.

To address this issue, we employ the PSDD to perform maximum a posteriori (MAP) query [16], obtaining the most likely model $\hat{\mathbf{m}}$ for the current state. Based on $\hat{\mathbf{m}}$ and the precondition φ of each action a , we re-normalize the action probabilities from a policy network. In this way, the learned symbolic representation from the PSDD can be used to constrain action selection, while the underlying policy network still provides a probability distribution that can be updated through gradient-based optimization.

Concretely, before the DRL agent makes a decision, we first use the PSDD to obtain the most likely model describing the state:

$$\hat{\mathbf{m}} = \operatorname{argmax}_{\mathbf{m}} Pr(\mathbf{m} | \Theta = g(s), \mathbf{m} \models \varphi) \quad (4)$$

Importantly, the argmax operation on the PSDD does not require enumerating all possible \mathbf{m} . Instead, it can be computed in linear time with respect to the PSDD size by exploiting its structural properties on decomposability and Determinism (see [52]). This linear-time inference makes PSDDs particularly attractive for DRL, where efficient evaluation of candidate actions are essential [4].

After obtaining the symbolic model of the state, we renormalize the probability of each action a according to its precondition φ :

$$\pi^+(s, a, \phi) = \frac{\pi(s, a) \cdot C_\varphi(\hat{\mathbf{m}})}{\sum_{a' \in \mathcal{A}} \pi(s, a') \cdot C_{\varphi'}(\hat{\mathbf{m}})} \quad (5)$$

where $\pi(s, a)$ denotes the probability of action a at state s predicted by the policy network, $C_\varphi(\hat{\mathbf{m}})$ is the evaluation of the SDD encoding from φ under the model $\hat{\mathbf{m}}$, and φ' is the precondition of action a' . The input of Equation (5) explicitly includes ϕ , as ϕ is required for evaluating the model $\hat{\mathbf{m}}$ in Equation (4). Intuitively, $C_\varphi(\hat{\mathbf{m}})$ acts as a symbolic mask. It equals to 1 if $\hat{\mathbf{m}} \models \varphi$ (i.e., the precondition is satisfied) and 0 otherwise. As a result, actions whose preconditions are violated are excluded from selection, while the probabilities of the remaining actions are renormalized as a new distribution. It is important to note that during the execution, we use the PSDD (trained by y in Equation (2) and (3)) to infer the most probable symbolic model of the current state (in Equation (4)), and therefore can formally verify whether each action's precondition is satisfied with this symbolic model (happened in C_φ in Equation (5)).

According to prior work, such 0-1 masking and renormalization still yield a valid policy gradient, thereby preserving the theoretical guarantees of policy optimization [30]. In practice, we optimize the masked policy π^+ using the Proximal Policy Optimization (PPO) objective [51]. Concretely, the loss is:

$$\mathcal{L}_{\text{PPO}}(\pi^+) = \mathbb{E}_t \left[\min \left(\mathbf{r}_t(\pi^+) \hat{A}_t, \operatorname{clip}(\mathbf{r}_t(\pi^+), 1 - \epsilon, 1 + \epsilon) \hat{A}_t \right) \right] \quad (6)$$

where $\mathbf{r}_t(\pi^+)$ denotes the probability ratio between the new and old masked policies, 'clip' is the clip function and \hat{A}_t is the advantage estimate [51]. In this way, the masked policy can be trained with PPO to effectively exploit symbolic action preconditions, leading to safer and more sample-efficient learning.

5 END-TO-END TRAINING FRAMEWORK

After deriving the gating function loss of PSDD in Equation (3) and the DRL policy loss in Equation (6), we now introduce an end-to-end training framework that combines the two components.

Before presenting the training procedure, we first summarize how the agent makes decisions, as illustrated in Figure 4. At each time step, the state s is first input into the symbolic grounding module, whose internal structure is shown in Figure 3. Within this module, the PSDD produces the most probable symbolic description of the state, i.e., a model $\hat{\mathbf{m}}$, according to Equation (4). The agent then leverages the preconditions in \mathcal{AP} (following Equation (5)) to mask the action distribution from policy network, and samples an action from the renormalized distribution to interact with the environment.

To achieve end-to-end training, we propose Algorithm 1. The key idea for this training framework is to periodically update the gating function of the PSDD during the agent's interaction with the environment, while simultaneously training the policy network




Figure 4: An illustration of the decision process of our agent, where the symbolic grounding module is as in Figure 3 and \hat{m} is calculated via the PSDD by Equation (4).

Algorithm 1 Training framework.

```

1: Compile  $\phi$  as SDD to obtain structure of PSDD
2: Initialize gating network  $g$  according to the structure of PSDD
3: Initialize policy network  $\pi$ , total step  $T \leftarrow 0$ 
4: Initialize a data buffer  $\mathcal{D}$  for learning PSDD
5: for  $Episode = 1 \rightarrow M$  do
6:   Reset  $Env$  and get  $s$ 
7:   while not terminal do
8:     Calculate action distribution before masking  $\pi(s, a)$ 
9:     Calculate  $\Theta = g(s)$  and assign parameter  $\Theta$  to PSDD
10:    Calculate  $\hat{m}$  in Equation (4)
11:    Calculate action distribution after masking  $\pi^+(s, a, \phi)$ 
12:    Sample an action  $a$  from  $\pi^+(s, a, \phi)$ 
13:    Execute  $a$  and get  $r, s'$  from  $Env$ 
14:    Obtain the truth-value  $y$  according to Equ. (1)
15:    Store  $(s, a, \phi)$  into  $\mathcal{D}$ 
16:    if terminal then
17:      Update policy  $\pi^+(s, a, \phi)$  using the trajectory of this episode
with Equation (6)
18:    end if
19:    if  $(T + 1) \% freq_g == 0$  then
20:      Sample batches from  $\mathcal{D}$ 
21:      Update gating function  $g$  with Equation (3)
22:    end if
23:     $s \leftarrow s', T \leftarrow T + 1$ 
24:  end while
25: end for

```

under the guidance of action masks. As the RL agent explores, it continuously generates minimally supervised feedback for the PSDD via Γ_ϕ , thereby improving the quality of the learned action masks. In turn, the improved action masking reduces infeasible actions and guides the agent toward higher rewards and more informative trajectories, which accelerates policy learning.

Concretely, before the start of training, the domain constrain ϕ is compiled into an SDD in Line 1, which determines both the structure of the PSDD and the output dimensionality of the gating function. Lines 2~4 initialize the gating function, the RL policy network, and a replay buffer \mathcal{D} that stores minimally supervised feedback for PSDD training. In Lines 5~25, the agent interacts with the environment and jointly learns the gating function and policy network. At each step (lines 8~11), the agent computes the masked action distribution based on the current gating function and policy network, which is crucial to minimizing the selection of infeasible actions during training. At the end of each episode (lines 16~18), the policy network is updated using the trajectory of

this episode. In this process, the gating function is kept frozen. In addition, the gating function is periodically updated (lines 19~22) with frequency $freq_g$. This periodically update enables the PSDD to provide increasingly accurate action masks in subsequent interactions, which simultaneously improves policy optimization and reduces constraint violations.

6 RELATED WORK

Symbolic grounding in neuro-symbolic learning. In the literature on neuro-symbolic systems [9, 65], symbol grounding refers to learning a mapping from raw data to symbolic representations, which is considered as a key challenge for integrating neural networks with symbolic knowledge [63]. Various approaches have been proposed to address this challenge in deep supervised learning. [19, 32, 62] leverage logic abduction or consistency checking to periodically correct the output symbolic representations. To achieve end-to-end differentiable training, the most common methods are embedding symbolic knowledge into neural networks through differentiable relaxations of logic operators, such as fuzzy logic [43] or Logic Tensor Networks (LTNs) [5]. These methods approximate logical operators with smooth functions, allowing symbolic supervision to be incorporated into gradient-based optimization [43, 63, 64, 71]. More recently, advances in probabilistic circuits [16, 52] give rise to efficient methods that embed symbolic knowledge via differentiable probabilistic representations, such as PSDD [37]. In these methods, symbolic knowledge is first compiled into SDD [20] to initialize the structure, after which a neural network is used to learn the parameters for predicting symbolic outputs [1]. This class of approaches has been successfully applied to structured output prediction tasks, including multi-label classification [1] and routing [54].

Symbolic grounding is also crucial in DRL. NRM [63] learn to capture the symbolic structure of reward functions in non-Markovian reward settings. In contrast, our approach learns symbolic properties of states to constrain actions under Markovian reward settings. KCAC [40] has extended PSDDs to MDP with combinatorial action spaces, where symbolic knowledge is used to constrain action composition. Our work also uses PSDDs but differs from KCAC. We use preconditions of actions as symbolic knowledge to determine the explorability of each individual action in a standard DRL setting, whereas KCAC incorporates knowledge about valid combinations of actions in a DRL setting with combinatorial action spaces.

Action masking. In DRL, action masking refers to masking out invalid actions during training to sample actions from a valid set [30]. Empirical studies in early real-world applications show that masking invalid actions can significantly improve sample efficiency of DRL [8, 31, 45, 50, 67, 76]. Following the systematic discussion of action masking in DRL [35], [29] investigates the impact of action masking on both on-policy and off-policy algorithms. Works such as [46, 58] extend action masking to continuous action spaces. [30] proves that binary action masking have Valid policy gradients during learning. In contrast to these approaches, our method does not assume that the set of invalid actions is predefined by the environment. Instead, we learn the set of invalid actions in each state for DRL using action precondition knowledge.

Another line of work employs a logical system (e.g., linear temporal logic [14]) to restrict the agent’s actions [2, 7, 34, 66]. These




Figure 5: Four tasks with logical constraints

approaches require a predefined symbol grounding function to map states into its symbolic representations, whereas our method learn such function (via PSDD) from data. PLPG [74] learns the probability of applying shielding with action constraints formulated in probabilistic logics. By contrast, our preconditions are hard constraints expressed in propositional logic: if the precondition of an action is evaluated to be false, the action is strictly not explorative.

Cost-based safe reinforcement learning. In addition to action masking, a complementary approach is to jointly optimize rewards and a safety-wise cost function to improve RL safety. In these cost-based settings, a policy is considered safe if its expected cumulative cost remains below a pre-specified threshold [49, 72]. A representative foundation of such cost-based approach is the constrained Markov decision process (CMDP) framework [3], which aims to maximize expected reward while ensuring costs below a threshold. Subsequent works often adopt Lagrangian relaxation to incorporate constraints into the optimization objective [17, 42, 49, 61, 73]. However, these methods often suffer from unsafe behaviors in the early stages of training [27]. To address such issues, safe exploration approaches emphasize to control the cost during exploration in unknown environments [2, 34]. Recently, SaGui [72] employed imitation learning and policy distillation to enable agents to acquire safe behaviors from a teacher agent during early training. RC-PPO [56] augmented unsafe states to allow the agent to anticipate potential future losses. While constraints can in principle be reformulated as cost functions, our approach does not rely on cost-based optimization. Instead, we directly exploit them to learn masks to avoid the violation of actions constraints.

7 EXPERIMENT

The experimental design aims to answer the following questions:

Q1: Without predefined symbolic grounding, can NSAM leverage symbolic knowledge to improve the sample efficiency of DRL?

Q2: By jointly learning symbolic grounding and masking strategies, can NSAM significantly reduce constraint violations during exploration, thereby enhancing safety?

Q3: In NSAM, is symbolic grounding with PSDDs more effective than replacing it with a module based on standard neural network?

Q4: In what ways does symbolic knowledge contribute to the learning process of NSAM?

7.1 Environments

We evaluate ASG on four highly challenging reinforcement learning domains with logical constraints as shown in Figure 5. Across all these environments, agents receive inputs in the form of unknown representation such as vectors or images.

Sudoku. Sudoku is a decision-making domain with logical constraints [44]. In this domain, the agent fills one cell with a number at each step until the board is complete. Action preconditions naturally arise: filling a number to a cell requires that the same number does not already exist in this row and column. Prior work [22, 44] shows that existing DRL algorithms struggle to solve Sudoku without predefined symbolic grounding functions.

N-Queens. The N-Queens problem requires placing N queens on a chessboard so that no two attack each other, making it a classic domain with logical constraints [48]. The agent places one queen per step on the chessboard until all N queens are placed safely. This task fits naturally within our extended MDP framework, where action preconditions arise: a queen can be placed at a position if and only if no already-placed queen can attack it.

Graph Coloring. Graph coloring is a NP-hard problem and serves as a reinforcement learning domain with logical constraints [18]. The agent sequentially colors the nodes of an undirected graph with a limited set of colors. Action preconditions arise naturally: a node may be colored with a given color if and only if none of its neighbors have already been assigned that color.

Visual Sudoku. Visual Sudoku follows the same rules as standard Sudoku but uses image-based digit representations instead of vector inputs. Following prior visual Sudoku benchmark [69], we generate boards by randomly sampling digits from the MNIST dataset [38]. Visual Sudoku 5×5 is with a 140×140 -dimensional state space and a 125-dimensional action space. In addition, the corresponding PSDD contains 125 atomic propositions and 782 clauses as constrains. This environment poses an additional challenge, as the digits are high-dimensional and uncertain representations, which increases the difficulty of symbolic grounding.

7.2 Hyperparameters

We run NSAM on a 2.60 GHz AMD Rome 7H12 CPU and an NVIDIA GeForce RTX 3070 GPU. For the policy and value functions, NSAM uses three-layer fully connected networks with 64 neurons per layer, while the gating function (Figure 3) uses a three-layer fully connected network with 128 neurons per layer. In the Visual Sudoku environment with image-based inputs, the policy and value functions are equipped with a convolutional encoder consisting of two convolutional layers (kernel sizes of 5 and 3, stride 2), followed by ReLU activations. The encoder output is then connected to a two-layer fully connected network with 256 neurons per layer. Similarly, the gating function in Visual Sudoku incorporates the same convolutional encoder, whose output is connected to a three-layer fully connected network with 128 neurons per layer. The learning rates for the actor, critic, and gating networks are set to 3×10^{-4} , 3×10^{-4} , and 2×10^{-4} , respectively. The gating function is trained using the Adam optimizer with a batch size of 128, updated every 1,000 time steps with 1,000 gradient updates per interval. Other hyperparameters follow the standard PPO settings⁴.

7.3 Baselines

We compare ASG with representative baselines from four categories. (1) Classical RL: Rainbow [28] and PPO [51], two standard DRL methods, where Rainbow integrates several DQN enhancements

⁴The code of NSAM is open-sourced at: <https://github.com/shan0126/NSRL>

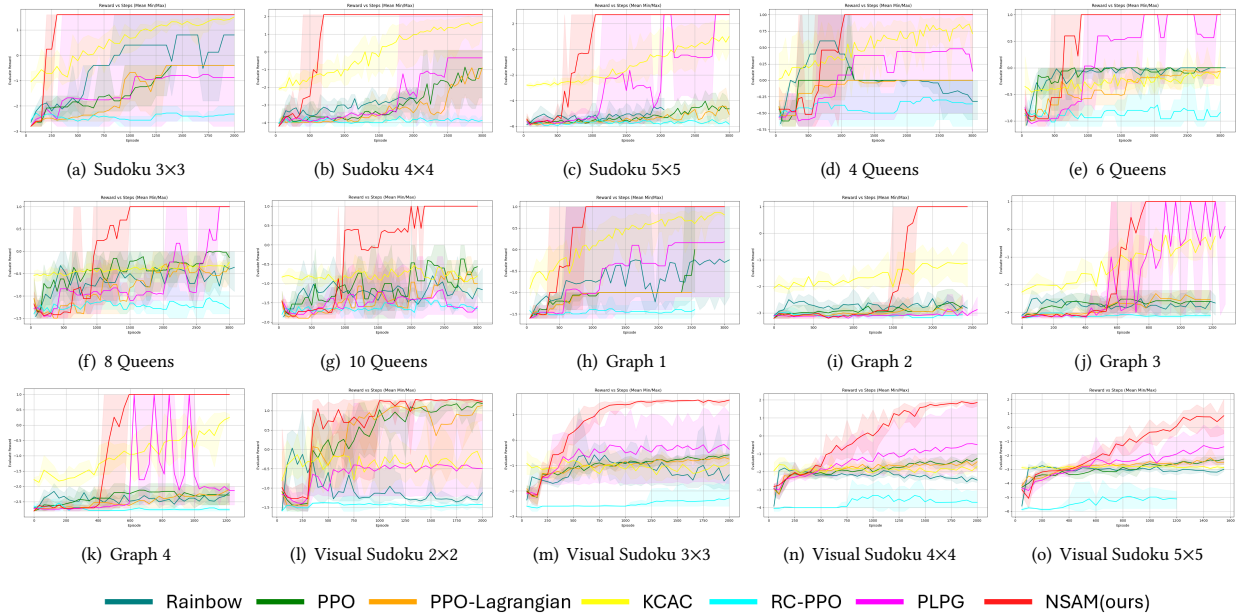


Figure 6: Comparison of learning curves on 4 domains. As the size and queen number increase in Sudoku and N-Queens, the learning task becomes more challenging. Graphs 1~4 denote four graph coloring tasks with different topologies.

		Rainbow		PPO		PPO-lagrangian.		KCAC		RC-PPO		PLPG		NSAM(ours)	
		Rew.	Viol.	Rew.	Viol.	Rew.	Viol.	Rew.	Viol.	Rew.	Viol.	Rew.	Viol.	Rew.	Viol.
Sudoku	2x2	1.3	36.2%	1.3	14.9%	1.3	0.7%	1.3	9.2%	0.3	11.3%	-0.8	5.9%	1.3	0.1%
	3x3	0.8	88.2%	-0.4	99.6%	-0.4	99.9%	1.5	57.1%	-2.3	99.0%	-0.9	8.9%	1.6	0.3%
	4x4	-2.6	100%	-2.6	100%	-3.4	100%	1.0	91.2%	-3.8	99.9%	-2.2	15.7%	2.1	0.6%
	5x5	-4.5	100%	-5.2	100%	-5.3	100%	-0.5	94.9%	-5.8	100%	-3.3	18.3%	2.7	4.3%
N-Queens	N=4	-0.3	97.8%	0.0	99.8%	0.0	100%	0.7	78.7%	-0.3	93.6%	0.1	10.4%	1.0	2.3%
	N=6	0.0	100%	0.0	100%	-0.1	100%	-0.1	100%	-0.8	100%	1.0	12.1%	1.0	1.3%
	N=8	-0.4	100%	-0.1	100%	-0.3	100%	-0.2	100%	-1.3	100%	1.0	41.6%	1.0	1.1%
	N=10	-1.2	100%	-0.6	100%	-1.0	100%	-0.8	100%	-1.7	100%	-1.6	98.2%	1.0	1.5%
Graph Coloring	G1	-0.2	88.7%	0.0	98.6%	-1.0	100%	0.8	43.9%	-1.4	99.1%	0.2	5.9%	1.0	0.7%
	G2	-2.8	100%	-2.8	100%	-2.8	100%	-1.1	98.7%	-3.1	98.7%	-2.8	18.7%	1.0	0.7%
	G3	-2.7	100%	-2.6	100%	-2.5	100%	-0.3	90.0%	-3.1	98.9%	0.1	7.5%	1.0	0.4%
	G4	-2.2	100%	-2.3	100%	-2.1	100%	0.3	85.4%	-2.8	75.8%	-2.1	11.2%	1.0	0.2%
Visual Sudoku	2x2	-1.1	61.8%	1.2	60.0%	1.1	25.0%	-0.4	32.2%	-1.4	68.3%	-0.5	16.1%	1.2	6.1%
	3x3	-1.4	96.0%	-0.6	99.6%	-0.8	100%	-0.9	40.3%	-2.3	95.4%	-0.4	34.6%	1.5	1.0%
	4x4	-2.5	100%	-1.3	100%	-1.4	100%	-1.7	39.4%	-3.7	96.3%	-0.5	38.2%	1.9	0.6%
	5x5	-3.0	100%	-2.2	100%	-2.5	100%	-2.9	88.5%	-5.1	99.7%	-1.4	53.7%	0.8	2.5%

Table 1: Comparison of final reward (Rew.) and violation (Viol.) rate during training.

and PPO is known for robustness and sample efficiency. (2) Neuro-symbolic RL: KCAC [40], which incorporates domain constraint knowledge to reduce invalid actions and represents the state-of-the-art in action-constrained DRL. (3) Action masking: PLPG [74], a state-of-the-art method that learns soft action masks for automatic action filtering. (4) Cost-based safe DRL: PPO-Lagrangian [49] and RC-PPO [56], which jointly optimize rewards and costs under the CMDP framework. For these methods, we construct the cost function using $(s, a, s', y) \in \Gamma_\phi$, where $y = 0$ increases the cost by 1 [68].

7.4 Learning efficiency and final performance

To answer Q1, we compare our method with all baselines on 16 tasks across four domains. The learning curves are shown in Figure 6. The error bounds (i.e., shadow shapes) indicate the upper and lower bounds of the performance with 5 runs using different random seeds. During the learning process of NSAM, we apply the normalization process defined in Eq. (5), which plays a critical role in excluding unsafe or unexplorable actions from the RL policy. On the Sudoku tasks, as the size increases and the action-state space grows, NSAM exhibits a slight decrease in sample efficiency but




Figure 7: Violation rate during training




Figure 8: Ablation study on PSDD in NSAM

consistently outperforms all baselines. A similar trend is observed in the N-Queens domain as the number of queens increases. In the Graph Coloring tasks, NSAM is able to converge to the optimal policy regardless of the graph topology among Graph 1~4. For the Visual Sudoku tasks, NSAM shows small fluctuations in performance after convergence due to the occurrence of previously unseen images. These fluctuations remain minor, indicating that NSAM’s symbolic grounding and learned policy generalize well to unseen digit images. The final converged reward values are reported in Table 1. Overall, NSAM achieves stable and competitive performance, consistently matching or outperforming the baselines.

7.5 Less violation

To answer Q2, we record constraint violations during training for each task. An episode is considered to violate constraints if any transition (s, a, s', y) is labeled with $y = 0$, i.e., the action a is not explorable in that state. The violation rate is computed as the ratio of the episodes violating constraints to the total number of episodes. The final violation rates are summarized in Table 1. MSAM achieves significantly lower violation rates than all other methods.

We further compare the change of violation rates during training for NSAM, PPO, and PPO-Lagrangian, as shown in Figure 7. In the early stages of training, NSAM exhibits a slightly higher violation rate because the PSDD parameters is not trained, which may cause inaccurate evaluation of action preconditions ϕ . However, as training progresses, the violation rate rapidly decreases to near zero. During the whole training process, the violation rate of NSAM is consistently lower than that of PPO and PPO-Lagrangian.

7.6 Ablation study

To answer Q3, we conduct an ablation study on the symbolic grounding module by replacing the PSDD with a standard three-layer fully connected neural network (128 neurons per layer). The experiment results are shown in Figure 8. Unlike PSDDs, neural networks




Figure 9: Policy training result from a single transition.

struggle to efficiently exploit symbolic knowledge and cannot guarantee logical consistency in their predictions. As a result, the policy trained with this ablated grounding module exhibits highly unstable performance, confirming the importance of PSDD for reliable symbolic grounding in NSAM.

7.7 Exploiting knowledge structure

To answer Q4, we design a special experiment where NSAM and PPO-Lagrangian are trained using only a single transition, as illustrated on the left side of Fig. 9. The right side of the figure shows the heatmaps of action probabilities after training. With the cost function defined via Γ_ϕ , PPO-Lagrangian can only leverage this single negative transition to reduce the probability of the specific action in this transition. As a result, in the heatmap of action probabilities after training, only the probability of this specific action decreases. In contrast, our method can exploit the structural knowledge of action preconditions to infer from the explorability of one action that a set of related actions are also infeasible. Consequently, in the heatmap, our method simultaneously reduces the probabilities of four actions. This demonstrates that our approach can utilize knowledge to generalize policy from very limited experience, which can significantly improve sample efficiency of DRL.

8 CONCLUSIONS AND FUTURE WORK

In this paper, we proposed NSAM, a novel framework that integrates symbolic reasoning with deep reinforcement learning through PSDDs. NSAM addresses the key challenges of learning symbolic grounding from high-dimensional states with minimal supervision, ensuring logical consistency, and enabling end-to-end differentiable training. By leveraging action precondition knowledge, NSAM learns effective action masks that substantially reduce constraint violations while improving the sample efficiency of policy optimization. Our empirical evaluation across four domains demonstrates that NSAM consistently outperforms baselines in terms of sample efficiency and violation rate. In this work, the symbolic knowledge in propositional form. A promising research direction for future work is to investigate richer forms of symbolic knowledge as action preconditions, such as temporal logics or to design learning framework when constraints are unknown or incorrect. Extending NSAM to broader real-world domains is also an important future direction.

ACKNOWLEDGMENTS

We sincerely thank the anonymous reviewers. This work is partly funded by the China Scholarship Council (CSC).

REFERENCES

- [1] Kareem Ahmed, Stefano Teso, Kai-Wei Chang, Guy Van den Broeck, and Antonio Vergari. 2022. Semantic probabilistic layers for neuro-symbolic learning. *Advances in Neural Information Processing Systems* 35 (2022), 29944–29959.
- [2] Mohammed Alshiekh, Roderick Bloem, Rüdiger Ehlers, Bettina Könighofer, Scott Niekum, and Ufuk Topcu. 2018. Safe reinforcement learning via shielding. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.
- [3] Eitan Altman. 1994. Denumerable constrained Markov decision processes and finite approximations. *Mathematics of operations research* 19, 1 (1994), 169–191.
- [4] Ivan Anokhin, Rishav Rishav, Matthew Riemer, Stephen Chung, Irina Rish, and Samira Ebrahimi Kahou. [n.d.]. Handling Delay in Real-Time Reinforcement Learning. In *The Thirteenth International Conference on Learning Representations*.
- [5] Samy Badreddine, Artur d’Avila Garcez, Luciano Serafini, and Michael Spranger. 2022. Logic tensor networks. *Artificial Intelligence* 303 (2022), 103649.
- [6] Bahador Beigomi and Zheng H Zhu. 2024. Towards Real-World Efficiency: Domain Randomization in Reinforcement Learning for Pre-Capture of Free-Floating Moving Targets by Autonomous Robots. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 11753–11759.
- [7] Francesco Belardinelli, Alexander W Goodall, et al. 2025. Probabilistic Shielding for Safe Reinforcement Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 16091–16099.
- [8] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębniak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. 2019. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680* (2019).
- [9] Bikram Pratim Bhuyan, Amar Ramdane-Cherif, Ravi Tomar, and TP Singh. 2024. Neuro-symbolic artificial intelligence: a survey. *Neural Computing and Applications* 36, 21 (2024), 12809–12844.
- [10] Christopher M Bishop and Nasser M Nasrabadi. 2006. *Pattern recognition and machine learning*. Vol. 4. Springer.
- [11] Léon Bottou. 2010. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT’2010: 19th International Conference on Computational Statistics/Paris France, August 22-27, 2010 Keynote, Invited and Contributed Papers*. Springer, 177–186.
- [12] Simone Bova. 2016. SDDs are exponentially more succinct than OBDDs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 30.
- [13] Randal E Bryant. 1986. Graph-based algorithms for boolean function manipulation. *Computers, IEEE Transactions on* 100, 8 (1986), 677–691.
- [14] John P Burgess and Yuri Gurevich. 1985. The decision problem for linear temporal logic. *Notre Dame Journal of Formal Logic* 26, 2 (1985), 115–128.
- [15] Satchit Chatterji and Erman Acar. 2024. Think Smart, Act SMAR! Analyzing Probabilistic Logic Shields for Multi-Agent Reinforcement Learning. *arXiv preprint arXiv:2411.04867* (2024).
- [16] Yoojung Choi, Antonio Vergari, and Guy Van den Broeck. [n.d.]. Probabilistic Circuits: A Unifying Framework for Tractable Probabilistic Models. ([n. d.]).
- [17] Yinlam Chow, Mohammad Ghavamzadeh, Lucas Janson, and Marco Pavone. 2018. Risk-constrained reinforcement learning with percentile risk criteria. *Journal of Machine Learning Research* 18, 167 (2018), 1–51.
- [18] Chase Cummins and Richard Veras. 2024. Reinforcement learning for graph coloring: Understanding the power and limits of non-label invariant representations. *arXiv preprint arXiv:2401.12470* (2024).
- [19] Wang-Zhou Dai, Qiuling Xu, Yang Yu, and Zhi-Hua Zhou. 2019. Bridging machine learning and logical reasoning by abductive learning. *Advances in Neural Information Processing Systems* 32 (2019).
- [20] Adnan Darwiche. 2011. SDD: A new canonical representation of propositional knowledge bases. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, Vol. 22. 819.
- [21] Pieter-Tjerk De Boer, Dirk P Kroese, Shie Mannor, and Reuven Y Rubinfeld. 2005. A tutorial on the cross-entropy method. *Annals of operations research* 134, 1 (2005), 19–67.
- [22] Haoyu Du, Lanhe Gao, and Xiaoyue Hu. 2021. A 4 × 4 Sudoku Solving Model Based on Multi-layer Perceptron. In *2021 International Conference on Electronic Information Engineering and Computer Science (EIECS)*. IEEE, 306–310.
- [23] Qiming Gao, Fangle Chang, Jiahong Yang, Yu Tao, Longhua Ma, and Hongye Su. 2024. Deep reinforcement learning for autonomous driving with an auxiliary actor discriminator. *Sensors* 24, 2 (2024), 700.
- [24] Xinwei Gao, Arambam James Singh, Gangadhar Royyuru, Michael Yuhas, and Arvind Easwaran. 2025. CRLK: Constrained Reinforcement Learning for Lane Keeping in Autonomous Driving. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems*. 3026–3028.
- [25] Reyhane Ghafari and Najme Mansouri. 2025. Reinforcement learning-based solution for resource management in fog computing: A comprehensive survey. *Expert Systems with Applications* (2025), 127214.
- [26] Giorgio Gnecco, Marco Gori, Stefano Melacci, and Marcello Sanguineti. 2014. Learning with mixed hard/soft pointwise constraints. *IEEE transactions on neural networks and learning systems* 26, 9 (2014), 2019–2032.
- [27] Tairan He, Weiye Zhao, and Changliu Liu. 2023. Autocost: Evolving intrinsic cost for zero-violation reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 14847–14855.
- [28] Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. 2018. Rainbow: Combining improvements in deep reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.
- [29] Yueqi Hou, Xiaolong Liang, Jiaqiang Zhang, Qisong Yang, Aiwu Yang, and Ning Wang. 2023. Exploring the use of invalid action masking in reinforcement learning: A comparative study of on-policy and off-policy algorithms in real-time strategy games. *Applied Sciences* 13, 14 (2023), 8283.
- [30] Shengyi Huang and Santiago Ontañón. 2020. A closer look at invalid action masking in policy gradient algorithms. *arXiv preprint arXiv:2006.14171* (2020).
- [31] Shengyi Huang, Santiago Ontañón, Chris Bamford, and Lukasz Grela. 2021. Gymµrts: Toward affordable full game real-time strategy games research with deep reinforcement learning. In *2021 IEEE Conference on Games (CoG)*. IEEE, 1–8.
- [32] Yu-Xuan Huang, Wang-Zhou Dai, Le-Wen Cai, Stephen H Muggleton, and Yuan Jiang. 2021. Fast abductive learning by similarity-based consistency optimization. *Advances in Neural Information Processing Systems* 34 (2021), 26574–26584.
- [33] Rodrigo Toro Icarte, Torny Q Klassen, Richard Valenzano, and Sheila A McIlraith. 2022. Reward machines: Exploiting reward function structure in reinforcement learning. *Journal of Artificial Intelligence Research* 73 (2022), 173–208.
- [34] Nils Jansen, Bettina Könighofer, Sebastian Junges, Alex Serban, and Roderick Bloem. 2020. Safe reinforcement learning using probabilistic shields. In *31st International Conference on Concurrency Theory (CONCUR 2020)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 3–1.
- [35] Anssi Kanervisto, Christian Scheller, and Ville Hautamäki. 2020. Action space shaping in deep reinforcement learning. In *2020 IEEE conference on games (CoG)*. IEEE, 479–486.
- [36] Diederik P Kingma. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [37] Doga Kisa, Guy Van den Broeck, Arthur Choi, and Adnan Darwiche. 2014. Probabilistic sentential decision diagrams. In *Fourteenth International Conference on the Principles of Knowledge Representation and Reasoning*.
- [38] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 2002. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (2002), 2278–2324.
- [39] Namyeong Lee and Jun Moon. 2023. Transformer actor-critic with regularization: automated stock trading using reinforcement learning. In *Proceedings of the 2023 International conference on autonomous agents and multiagent systems*. 2815–2817.
- [40] Jiajing LING, Moritz Lukas SCHULER, Akshat KUMAR, and Pradeep VARAKANTHAM. 2023. Knowledge compilation for constrained combinatorial action spaces in reinforcement learning. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems, London, Great Britain*.
- [41] Jianlan Luo, Zheyuan Hu, Charles Xu, You Liang Tan, Jacob Berg, Archit Sharma, Stefan Schaal, Chelsea Finn, Abhishek Gupta, and Sergey Levine. 2024. Serl: A software suite for sample-efficient robotic reinforcement learning. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 16961–16969.
- [42] Haitong Ma, Yang Guan, Shegno Eben Li, Xiangteng Zhang, Sifa Zheng, and Jianyu Chen. 2021. Feasible actor-critic: Constrained reinforcement learning for ensuring statewise safety. *arXiv preprint arXiv:2105.10682* (2021).
- [43] Robin Manhaeve, Sebastijan Dumancic, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. 2018. DeepProlog: Neural probabilistic logic programming. *Advances in neural information processing systems* 31 (2018).
- [44] Anav Mehta. 2021. Reinforcement learning for constraint satisfaction game agents (15-puzzle, minesweeper, 2048, and sudoku). *arXiv preprint arXiv:2102.06019* (2021).
- [45] Branka Mirchevska, Christian Pék, Moritz Werling, Matthias Althoff, and Joschka Boedecker. 2018. High-level decision making for safe and reasonable autonomous lane changing using reinforcement learning. In *2018 21st international conference on intelligent transportation systems (ITSC)*. IEEE, 2156–2162.
- [46] Michael Neuner, Abbas Abdolmaleki, Markus Wulfmeier, Thomas Lampe, Tobias Springenberg, Roland Hafner, Francesco Romano, Jonas Buchli, Nicolas Heess, and Martin Riedmiller. 2020. Continuous-discrete reinforcement learning for hybrid control in robotics. In *Conference on Robot Learning*. PMLR, 735–751.
- [47] Eduardo HM Pena, Eduardo C de Almeida, and Felix Naumann. 2021. Fast detection of denial constraint violations. *Proceedings of the VLDB Endowment* 15, 4 (2021), 859–871.
- [48] Patnala Prudhvi Raj, Preet Shah, and Pragnya Suresh. 2019. Faster convergence to N-queens problem using reinforcement learning. In *International Conference on Modeling, Machine Learning and Astronomy*. Springer, 66–77.
- [49] Alex Ray, Joshua Achiam, and Dario Amodei. 2019. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708* 7, 1 (2019), 2.
- [50] Thomas Rudolf, Mingxiao Gao, Tobias Schürmann, Stefan Schwab, and Sören Hohmann. 2022. Fuzzy Action-Masked Reinforcement Learning Behavior Planning for Highly Automated Driving. In *2022 8th International Conference on Control, Automation and Robotics (ICCAR)*. IEEE, 264–270.
- [51] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*

- (2017).
- [52] Yujia Shen, Arthur Choi, and Adnan Darwiche. 2016. Tractable operations for arithmetic circuits of probabilistic models. *Advances in Neural Information Processing Systems* 29 (2016).
- [53] Yujia Shen, Arthur Choi, and Adnan Darwiche. 2018. Conditional PSDDs: Modeling and learning with modular knowledge. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [54] Yujia Shen, Anchal Goyanka, Adnan Darwiche, and Arthur Choi. 2019. Structured bayesian networks: From inference to learning with routes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 7957–7965.
- [55] Hikaru Shindo, Quentin Delfosse, Devendra Singh Dhami, and Kristian Kersting. 2024. Blendrl: A framework for merging symbolic and neural policy learning. *arXiv preprint arXiv:2410.11689* (2024).
- [56] Oswin So, Cheng Ge, and Chuchu Fan. 2024. Solving minimum-cost reach avoid using reinforcement learning. *Advances in Neural Information Processing Systems* 37 (2024), 30951–30984.
- [57] Sarath Sreedharan and Michael Katz. 2023. Optimistic exploration in reinforcement learning using symbolic model estimates. *Advances in Neural Information Processing Systems* 36 (2023), 34519–34535.
- [58] Roland Stolz, Hanna Krasowski, Jakob Thumm, Michael Eichelbeck, Philipp Gassert, and Matthias Althoff. 2024. Excluding the irrelevant: Focusing reinforcement learning through continuous action masking. *Advances in Neural Information Processing Systems* 37 (2024), 95067–95094.
- [59] Richard S Sutton, Andrew G Barto, et al. 1998. *Reinforcement learning: An introduction*. Vol. 1. MIT press Cambridge.
- [60] Sho Takeda, Satoshi Yamamori, Satoshi Yagi, and Jun Morimoto. 2025. An empirical evaluation of a hierarchical reinforcement learning method towards modular robot control. *Artificial Life and Robotics* 30, 2 (2025), 245–251.
- [61] Chen Tessler, Daniel J Mankowitz, and Shie Mannor. 2018. Reward constrained policy optimization. *arXiv preprint arXiv:1805.11074* (2018).
- [62] Efthymia Tsamoura, Timothy Hospedales, and Loizos Michael. 2021. Neural-symbolic integration: A compositional perspective. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 5051–5060.
- [63] Elena Umili, Francesco Argenziano, and Roberto Capobianco. 2024. Neural Reward Machines. In *ECAI 2024 - 27th European Conference on Artificial Intelligence, 19-24 October 2024, Santiago de Compostela, Spain - Including 13th Conference on Prestigious Applications of Intelligent Systems (PAIS 2024)*, Vol. 392. 3055–3062.
- [64] Elena Umili, Roberto Capobianco, and Giuseppe De Giacomo. 2023. Grounding LTLf specifications in image sequences. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, Vol. 19. 668–678.
- [65] Michael van Bekkum, Maaik de Boer, Frank van Harmelen, André Meyer-Vitali, and Annette ten Teije. 2021. Modular design patterns for hybrid learning and reasoning systems: a taxonomy, patterns and use cases. *Applied Intelligence* 51, 9 (2021), 6528–6546.
- [66] Giovanni Varricchione, Natasha Alechina, Mehdi Dastani, Giuseppe De Giacomo, Brian Logan, and Giuseppe Perelli. 2024. Pure-past action masking. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 21646–21655.
- [67] Oriol Vinyals, Timo Ewalds, Sergey Bartunov, Petko Georgiev, Alexander Sasha Vezhnevets, Michelle Yeo, Alireza Makhzani, Heinrich Küttler, John Agapiou, Julian Schrittwieser, et al. 2017. Starcraft ii: A new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782* (2017).
- [68] Akifumi Wachi, Wataru Hashimoto, Xun Shen, and Kazumune Hashimoto. 2023. Safe exploration in reinforcement learning: A generalized formulation and algorithms. *Advances in Neural Information Processing Systems* 36 (2023), 29252–29272.
- [69] Po-Wei Wang, Priya Donti, Bryan Wilder, and Zico Kolter. 2019. Satnet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver. In *International Conference on Machine Learning*. PMLR, 6545–6554.
- [70] Qinghao Wang and Yaodong Yang. 2024. Carbon trading supply chain management based on constrained deep reinforcement learning. *Autonomous Agents and Multi-Agent Systems* 38, 2 (2024), 38.
- [71] Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, and Guy Broeck. 2018. A semantic loss function for deep learning with symbolic knowledge. In *International conference on machine learning*. PMLR, 5502–5511.
- [72] Qisong Yang, Thiago D Simão, Nils Jansen, Simon H Tindemans, and Matthijs TJ Spaan. 2023. Reinforcement learning by guided safe exploration. *arXiv preprint arXiv:2307.14316* (2023).
- [73] Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J Ramadge. 2020. Projection-based constrained policy optimization. *arXiv preprint arXiv:2010.03152* (2020).
- [74] Wen-Chi Yang, Giuseppe Marra, Gavin Rens, and Luc De Raedt. 2023. Safe reinforcement learning via probabilistic logic shields. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*. 5739–5749.
- [75] Xi Yao, Yingdong Hu, Yicheng Xu, and Ruifeng Gao. 2024. Deep reinforcement learning-based resource management in maritime communication systems. *Sensors* 24, 7 (2024), 2247.
- [76] Deheng Ye, Zhao Liu, Mingfei Sun, Bei Shi, Peilin Zhao, Hao Wu, Hongsheng Yu, Shaojie Yang, Xipeng Wu, Qingwei Guo, et al. 2020. Mastering complex control in moba games with deep reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 34. 6672–6679.
- [77] Zihan Ye, Oleg Arenz, and Kristian Kersting. 2025. Learning from Less: Guiding Deep Reinforcement Learning with Differentiable Symbolic Planning. *arXiv preprint arXiv:2505.11661* (2025).
- [78] Dongkun Zhang, Jiaming Liang, Ke Guo, Sha Lu, Qi Wang, Rong Xiong, Zhenwei Miao, and Yue Wang. 2025. Carplanner: Consistent auto-regressive trajectory planning for large-scale reinforcement learning in autonomous driving. In *Proceedings of the Computer Vision and Pattern Recognition Conference*. 17239–17248.
- [79] Rongliang Zhou, Jiakun Huang, Mingjun Li, Hepeng Li, Haotian Cao, and Xiaolin Song. 2025. Knowledge transfer from simple to complex: A safe and efficient reinforcement learning framework for autonomous driving decision-making. *Advanced Engineering Informatics* 65 (2025), 103188.