

# Neural Reward Machines

Elena Umili<sup>a,\*</sup>, Francesco Argenziano<sup>a</sup> and Roberto Capobianco<sup>a,b</sup>

<sup>a</sup>Sapienza University of Rome

<sup>b</sup>Sony AI

**Abstract.** Non-markovian Reinforcement Learning (RL) tasks are very hard to solve, because agents must consider the entire history of state-action pairs to act rationally in the environment. Most works use symbolic formalisms (as Linear Temporal Logic or automata) to specify the temporally-extended task. These approaches only work in finite and discrete state environments or continuous problems for which a mapping between the raw state and a symbolic interpretation is known as a symbol grounding (SG) function. Here, we define Neural Reward Machines (NRM), an automata-based neurosymbolic framework that can be used for both reasoning and learning in non-symbolic non-markovian RL domains, which is based on the probabilistic relaxation of Moore Machines. We combine RL with semisupervised symbol grounding (SSSG) and we show that NRMs can exploit high-level symbolic knowledge in non-symbolic environments without any knowledge of the SG function, outperforming Deep RL methods which cannot incorporate prior knowledge. Moreover, we advance the research in SSSG, proposing an algorithm for analysing the groundability of temporal specifications, which is more efficient than baseline techniques of a factor  $10^3$ .

## 1 Introduction

Reinforcement learning (RL) tasks are traditionally modeled as Markovian Decision Processes (MDPs) [29], where task feedback depends solely on the last state and action. However, this formulation is inadequate for many decision problems, which are inherently non-Markovian or temporally extended [1, 21]. Intelligent agents tackling such tasks must consider the entire history of state-action pairs to act rationally within the environment. Current research in this field involves addressing non-Markovianity by expanding the state space with features that encode the environment history, and then solving the augmented-state problem with established RL algorithms. The primary challenge lies in constructing these features. For non-symbolic-state problems, a popular approach combines RL algorithms with Recurrent Neural Networks (RNNs) [13, 17], which automatically extract features from data sequences. While this method does not guarantee extracting Markovian features, it does not require any prior knowledge on the task. In problems with a symbolic finite state space, most works utilize Linear Temporal Logic (LTL) [27] or LTLf [8] to specify the temporal task. This specification is then compiled into an automaton, and the automaton state is combined with the environment state to render the decision process Markovian. Reward Machines (RMs) [15] exemplify this approach. RMs can also be employed in non-symbolic-state environments if the symbol grounding (SG) function is known [4]. The latter maps the environment's raw state into

a boolean interpretation over the symbols used by the specifications in LTL, making the symbols observable. In summary, RM-like methods presuppose prior knowledge of: (i) the SG function, and (ii) the temporal task specification in a logical formalism. Despite previous work challenging the assumption of knowing the temporal specification [11, 38, 28, 10], no work thus far assumes a complete lack of knowledge regarding the SG function. We highlight that learning this function for realistic environments with raw and/or high-dimensional states can be challenging and necessitates a substantial amount of labeled data.

To achieve this, we propose a neurosymbolic (NeSy) approach to reward machines, which involves probabilistically relaxing logical constraints [31]. This approach integrates both logical reasoning and neural network-based perception into a unified framework, that we call "Neural Reward Machines" (NRM). We demonstrate the capability to relax the requirement of knowing the SG function while still leveraging available prior knowledge to enhance the performance of RL algorithms in non-symbolic non-Markovian environments. This achievement stems from combining RL on reward machines with semi-supervised symbol grounding (SSSG) [22, 2]. SSSG consists in connecting some prior logical knowledge with raw data by supervising the formula output on the data. SSSG is often challenging, particularly when logical knowledge is vague or when connected data doesn't cover a diverse range of scenarios [32, 23]. In such cases, multiple solutions may exist — each maintaining consistency between knowledge and data - where, only one is the intended solution, while the others are deemed "Reasoning Shortcuts" (RS) [23]. In this paper, we also advance SSSG research by presenting an algorithm for identifying all RSs of a given temporal specification, assuming a complete dataset of observations. We call this kind of RSs *unremovable*, since they solely depend on the structure of knowledge rather than the data used for SSSG or the specific symbol-perception task.

To the best of our knowledge, this is the first work exploiting temporal logical knowledge in single RL tasks under the assumption of *completely unknown* SG function, and proposing an algorithm for actually *discovering* RSs of temporal specifications.

## 2 Related works

**Non-Markovian Reinforcement Learning with Temporal Specifications** Temporal logic formalisms are widely used in Reinforcement Learning (RL) to specify non-Markovian tasks [21]. The large majority of works assumes the boolean symbols used in the formulation of the task are *perfectly* observable in the environment [5, 9, 11, 38, 28, 10]. For this reason they are applicable only in symbolic-state environments or when a perfect mapping between the environment state and a symbolic interpretation is known, also

\* Corresponding Author. Email: umili@diag.uniroma1.it

called labeled MDP [4]. Many works assume to know an *imperfect* SG function for the task [3, 35, 20]. Namely, a function that sometimes makes mistakes in predicting symbols from states or predicts a set of probabilistic *beliefs* over the symbol set instead of a boolean interpretation. These works represent a step towards integration with non-symbolic domains. However, they do not address the problem of learning the SG function, but only how to manage to use a pre-trained imperfect symbol grounder. Only one work in the literature assumes the same setting of ours [18], namely, that the agent observes sequences of non-symbolic states and rewards, and is aware of the LTL formula describing the task but without knowing the meaning of the symbols in the formula. This work employs a neural network composed of different modules, each representing a propositional symbol or a logical or temporal operator, which are interconnected in a tree-like structure, respecting the syntactic tree of the formula. This network learns a representation of states, that can be easily transferred to different LTL tasks in the same environment. However, the key distinctions between our and their work are: (i) Kuo et al. [18] learn a subsymbolic uninterpretable representation, while we learn *precisely* a mapping between states and symbols; (ii) their method provides benefits only in a multitask setting and is unable to expedite learning on a single task, while ours learns and leverages the symbol grounding in the individual task.

**Semisupervised Symbol Grounding** Much prior work approaches the SG problem by assuming as inputs: (i) some prior symbolic logical knowledge, and (ii) a set of raw data that are annotated with the output of the knowledge (for example data are divided in positive and negative sample accordingly to the formula)[37, 2, 22, 32, 7, 14, 30]. This practice is known as semi-supervised symbol-grounding (SSSG). SSSG is tackled mainly with two families of methods. The first approach [37, 2, 22, 32] consists in embedding a continuous relaxation of the available knowledge in a larger neural networks system, and training the system so to align the outputs of the relaxation with the high-level labels. The second approach [7, 14, 30] instead maintains a crisp boolean representation of the logical knowledge and uses a process of logic *abduction* to correct the current SG function, that is periodically retrained with the corrections. Much work in this area does not take into account temporal aspects, and only few works focus on learning symbol grounding using temporal logics. In particular, [32] employs an extension of Logic Tensor Networks (LTNs) to represent DFAs, which exploits for classifying sequences of images. Our approach can be categorized in the first family of mentioned methods. Differently from [32], which employs LTNs and apply to classification domains, we use probabilistic Moore Machines and our main application is to RL. Furthermore our framework is more versatile than [32], as the temporal component can be both imposed and learned through the network.

**Groundability and Reasoning Shortcuts** In [23], the authors define reasoning shortcuts and outline a method for counting deterministic optima, assuming knowledge of: the logical knowledge, the support dataset, and the SG function. This method is inapplicable in our scenario since we assume the SG unknown. Conversely, [32] introduces the concept of "ungroundability" as a property solely dependent on the logical knowledge and not on the data or the SG function. However, it gives only theoretical definitions and does not provide a method to actually discovering the RSs that cause the ungroundability. In this paper, we bridge these two concepts and we introduce the *Unremovable Reasoning Shortcuts* (URS), as the RSs identifiable when considering all possible observations. Beyond offering a theoretical definition, we develop an algorithm for actually calculating URS for a temporal

property, which is missing in the literature.

## 3 Background

### 3.1 Moore Machines

A Moore machine  $M$  is a tuple  $(P, Q, O, q_0, \delta_t, \delta_o)$ , where  $P$  is a finite alphabet of input propositional symbols,  $Q$  is a finite set of states,  $O$  is a finite set of output symbols,  $q_0 \in Q$  is the initial state,  $\delta_t : Q \times P \rightarrow Q$  is the transition function, and  $\delta_o : Q \rightarrow O$  is the output function. Let  $P^*$  be the set of all finite strings over  $P$ , and  $\epsilon$  the empty string. The transition function over strings  $\delta_t^* : Q \times P^* \rightarrow Q$  is defined recursively as

$$\begin{aligned} \delta_t^*(q, \epsilon) &= q \\ \delta_t^*(q, p+x) &= \delta_t^*(\delta_t(q, p), x) \end{aligned} \quad (1)$$

Where  $p \in P$  is a symbol,  $x \in P^*$  is a string, and  $p+x$  is the concatenation of  $p$  and  $x$ . Consequently, the output function over strings  $\delta_o^* : Q \times P^* \rightarrow O$  is defined as

$$\delta_o^*(q, x) = \delta_o(\delta_t^*(q, x)) \quad (2)$$

We also define the string of output  $\delta_o^{**} : Q \times P^* \rightarrow O^*$  as

$$\delta_o^{**}(q, x) = [\delta_o^*(q, x|_1), \delta_o^*(q, x|_2), \dots, \delta_o^*(q, x|_T)] \quad (3)$$

Where  $T$  is the length of string  $x$ , and  $x|_i$  is the string  $x$  up to index  $i$ . Let  $x_p = [p^{(1)}, p^{(2)}, \dots, p^{(T)}]$  be the input string, where  $p^{(t)}$  is the  $t$ -th character in the string, we denote as  $x_q = [q^{(0)}, q^{(1)}, \dots, q^{(T)}]$  and as  $x_o = [o^{(1)}, o^{(2)}, \dots, o^{(T)}]$  respectively the sequence of states and output symbols produced by the automaton while processing the string, namely  $q^{(0)} = q_0$  and  $q^{(t)} = \delta_t(q^{(t-1)}, x^{(t)})$  and  $o^{(t)} = \delta_o(q^{(t)})$  for all  $t > 0$ .

### 3.2 Non-Markovian Reward Decision Processes and Reward Machines

In Reinforcement Learning (RL) [29] the agent-environment interaction is generally modeled as a Markov Decision Process (MDP). An MDP is a tuple  $(S, A, t, r, \gamma)$ , where  $S$  is the set of environment states,  $A$  is the set of agent's actions,  $t : S \times A \times S \rightarrow [0, 1]$  is the *transition function*,  $r : S \times A \rightarrow \mathbb{R}$  is the *reward function*, and  $\gamma \in [0, 1]$  is the *discount factor* expressing the preference for immediate over future reward

In this classical setting, transitions and rewards are assumed to be Markovian – i.e., they are functions of the current state only. Although this formulation is general enough to model most decision problems, it has been observed that many natural tasks are non-Markovian [21]. A decision process can be non-markovian because markovianity does not hold on the reward function  $r : (S \times A)^* \rightarrow \mathbb{R}$ , or the transition function  $t : (S \times A)^* \times S \rightarrow [0, 1]$ , or both. In this work we focus on Non-Markovian Reward Decision Processes (NMRDP) [12]. Learning an optimal policy in such settings is hard, since the current environment outcome depends on the entire history of state-action pairs the agent has explored from the beginning of the episode; therefore, regular RL algorithms are not applicable. Rather than developing new RL algorithms to tackle NMRDP, the research has focused mainly on how to construct Markovian state representations of NMRDP. An approach of this kind are the so called Reward Machines (RMs).

RMs are an automata-based representation of non-Markovian reward functions [15]. Given a finite set of propositions  $P$  representing abstract properties or events observable in the environment, RMs

specify temporally extended rewards over these propositions while exposing the compositional reward structure to the learning agent. Formally, in this work we assume the reward can be represented as a Reward Machine  $RM = (P, Q, R, q_0, \delta_t, \delta_r, L)$ , where  $P$  is the automaton alphabet,  $Q$  is the set of automaton states,  $R$  is a finite set of continuous reward values,  $q_0$  is the initial state,  $\delta_t : Q \times P \rightarrow Q$  is the transition function,  $\delta_r : Q \rightarrow R$  is the reward function, and  $L : S \rightarrow P$  is the labeling (or symbol grounding) function, which recognizes symbols in the environment states.

Let  $x_s = [s^{(1)}, s^{(2)}, \dots, s^{(t)}]$  be a sequence of states the agent has observed in the environment up to the current time instant  $t$ , we define the labeling function over sequences  $L^* : S^* \rightarrow P^*$  as

$$L^*(x_s) = [L(s^{(1)}), L(s^{(2)}), \dots, L(s^{(t)})] \quad (4)$$

We denote with  $\delta_t^*$  and  $\delta_r^*$  the transition and reward function over strings, which are defined recursively, analogously to Equations 1 and 2. Given  $x_s$ , the RM produces an history-dependent reward value at time  $t$ ,  $r^{(t)} = \delta_r^*(q_0, L^*(x_s))$  and an automaton state at time  $t$ ,  $q^{(t)} = \delta_t^*(q_0, L^*(x_s))$ . The reward value can be used to guide the agent toward the satisfaction of the task expressed by the automaton, while the automaton state can be used to construct a Markovian state representation. In fact it was proven that the augmented state  $(s^{(t)}, q^{(t)})$  is a Markovian state for the task expressed by the RM [9].

**Example: Image-Based Minecraft-Like Environment** Consider the environment shown in Figure 1(a) consisting in a grid world containing different items: a pickaxe, a gem, a door and a lava-cell. The task to be accomplished in the environment is reaching at least once the pickaxe, the lava and the door cells, in any order. We can represent the task with the Moore Machine depicted in Figure 1(b), defined in terms of five symbols: one for each different item, plus a symbol indicating the absence of items. Therefore  $P = \{\text{'pickaxe'}, \text{'gem'}, \text{'door'}, \text{'lava'}, \text{'empty-cell'}\}$ . Each symbol is considered set to True when the agent is in the item-cell and false otherwise. The agent must learn how to navigate the grid in order to satisfy the task. At each step it receives as state  $s \in S$  an image showing the agent in the grid, similar to that depicted in the figure. This state representation is not Markovian because from the current image the agent can recover only if it is located on an item *now* and not which items it has visited in the past.

## 4 Neural Reward Machines

In this section, we formalize Neural Reward Machines, elaborate on their implementation using neural networks, and explore the reasoning and learning tasks achievable with NRMs. Subsequently, we will specifically delve into the integration of semi-supervised symbol grounding and RL through NRMs. Finally, we present our algorithm for identifying Unremovable Reasoning Shortcuts for a temporal property.

### 4.1 Definition and Notations

We define a Neural Reward Machine as a tuple  $NRM = (S, P, Q, R, q_0, \delta_{tp}, \delta_{rp}, sg)$ , where  $S$  is the set of environment states, possibly infinite and continuous, the machine can process;  $P$  is a finite set of symbols;  $Q$  is a finite set of states,  $R$  is a finite set of rewards;  $q_0$  is the initial state;  $\delta_{tp} : Q \times P \times Q \rightarrow [0, 1]$  is the machine transition probability function;  $\delta_{rp} : Q \times R \rightarrow [0, 1]$  is the reward probability function and  $sg : S \times P \rightarrow [0, 1]$  is the *symbol grounding* probability function.

Like an RM, an NRM produces a temporal state and a reward from a sequence of environment states. However, unlike an RM, the input sequences do not necessarily need to be symbolic; they can be of any type. The symbol grounding function maintains the link with the symbolic representation, assigning to a data instance  $s \in S$  a probability value for each symbol in the alphabet  $P$ . Given a time sequence of environment observations  $x_s = [s^{(1)}, s^{(2)}, \dots, s^{(T)}]$  the symbol grounding function grounds probabilistically each state in the set of symbols  $P$  producing a sequence of symbol probability vectors  $x_{p_p} = [p_p^{(1)}, p_p^{(2)}, \dots, p_p^{(T)}]$ , where the  $k$ -th component of  $p_p^{(t)}$  is equal to  $sg(s^{(t)}, p_k)$ . Note that probabilities are present both in the machine and the grounding: in our setting, the Reward Machine is a *probabilistic* Moore Machine taking as input a *probabilistic* symbols. We represent the stochastic Moore machine in matrix representation, as: (i) the initial state probability vector  $q_p^{(0)} \in [0, 1]^{|Q|}$ , containing at index  $i$  a 1 if  $q_0 = i$  and a 0 otherwise; (ii) a transition matrix  $M_t \in [0, 1]^{|P| \times |Q| \times |Q|}$  containing at index  $(p, q, q')$  the value of  $\delta_{tp}(q, p, q')$ ; (iii) a reward matrix  $M_r \in [0, 1]^{|Q| \times |R|}$  representing the reward function and containing at index  $(q, r)$  the value of  $\delta_{rp}(q, r)$ . The sequence of probabilistic symbols returned by the grounder is processed by the probabilistic Moore Machine, which produces a sequence of state probability vectors  $x_{q_p} = [q_p^{(1)}, q_p^{(2)}, \dots, q_p^{(T)}]$  and a sequence of reward probability vectors  $x_{r_p} = [r_p^{(1)}, r_p^{(2)}, \dots, r_p^{(T)}]$ .

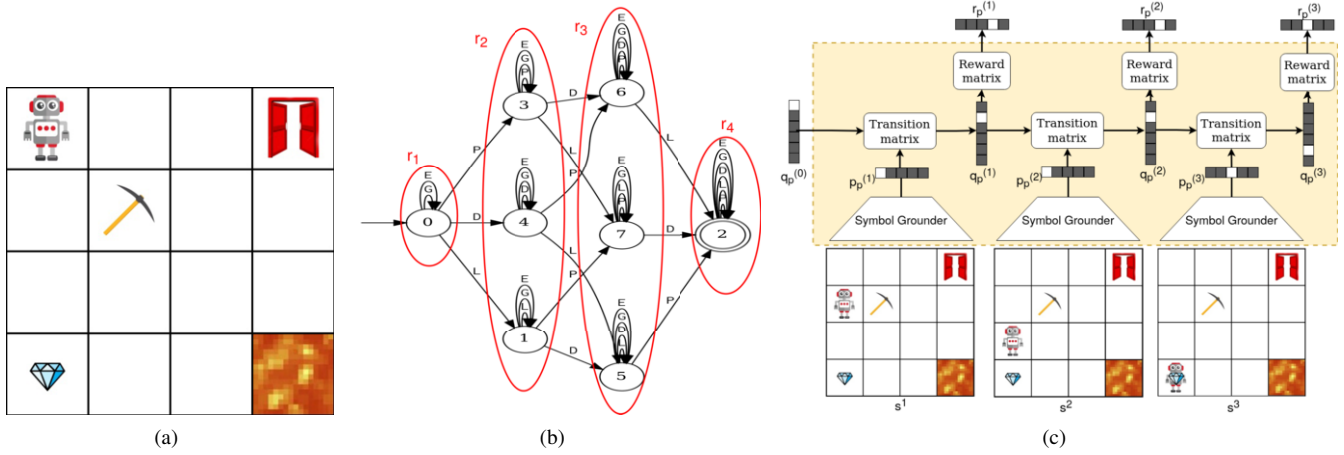
$$\begin{aligned} p_p^{(t)} &= sg(s^{(t)}) \\ q_p^{(t)} &= \sum_{i=1}^{|P|} p_p^{(t)}[i] (q_p^{(t-1)} M_t[i]) \\ r_p^{(t)} &= q_p^{(t)} M_r \end{aligned} \quad (5)$$

Where we denote with  $V[i]$  ( $v[i]$ ), the component  $i$  of matrix (vector)  $V$  ( $v$ ). We show the model in Figure 1(c). Summarizing, given a time sequence of environment states  $x_s$ , the model in Equation 5 can be applied recursively on  $x_s$  producing: a time sequence of input symbols  $x_{p_p}$ , one of machine states  $x_{q_p}$ , and a one of reward values  $x_{r_p}$ . Subscript  $p$  denotes that sequences are *probabilistically grounded*, namely, each point in the sequence is a probability vector. While we denote with  $x_p, x_q, x_r$ , without subscript  $p$ , the *symbolic* time sequences. Namely  $x_p = [p^{(1)}, p^{(2)}, \dots, p^{(T)}]$ ,  $x_q = [q^{(1)}, q^{(2)}, \dots, q^{(T)}]$ , and,  $x_r = [r^{(1)}, r^{(2)}, \dots, r^{(T)}]$ , with each point in the sequence being a symbol in the finite set  $P, Q$  and  $R$  respectively. Symbolic sequences can also be obtained from the probabilistic ones choosing the most probable symbol at each time step. Namely  $p^{(t)} = p_i \in P$ , with  $i = \arg \max_i (p_p^{(t)}[i])$ ;  $q^{(t)} = q_j \in Q$ , with  $j = \arg \max_j (q_p^{(t)}[j])$ ;  $r^{(t)} = r_k \in R$ , with  $k = \arg \max_k (r_p^{(t)}[k])$ . We next build our neural framework by using continuous parametric models based on this structure.

### 4.2 Implementation

In this section, we discuss the implementation of the formalism described above using Neural Networks (NN). Notably, we employ probabilistic relaxations and matrix representations of functions to facilitate this integration. Neural Networks are inherently continuous as they learn through a gradient-based optimization process, therefore they can face challenges in learning a ‘crispy’ boolean logical model like a Moore Machine. In our case, we aim for each function in the NRM definition—namely,  $sg$ ,  $\delta_{tp}$ , and  $\delta_{rp}$ —to be both *initialized* with prior knowledge and *learned* from data.

For the SG function, it can take any form, with the only constraint being that it must return a probability vector. Therefore, it can be



**Figure 1.** a) An example of non-Markovian navigation environment inspired by the Minecraft videogame. b) Moore Machine for the task: the agent has to visit the pickaxe (P), the lava (L) and the door (D) cells in any order. c) Implementation of NRM with neural networks.

implemented with any NN featuring softmax activation on the output layer. Conversely, learning the Moore Machine with backpropagation is more challenging. Our approach is based on the intuition that Probabilistic Moore Machines (PMM) are closely related to Recurrent Neural Networks (RNN), as they calculate the next state and output using multiplications between continuous vectors and matrices, similar to RNNs. Even though (deterministic) machines can be represented in matrix form with only one-hot row vectors. Following this idea, we define the recurrent module as a parametric PMM whose representation can be driven to be close to one-hot during training, enabling logical induction through backpropagation. We achieve this effect using an activation function that smoothly approximates discrete output, as seen in prior work [16, 36, 19]. In particular, we use a modified version of the classical softmax activation function:  $\text{softmax}_\tau(x, \tau) = \text{softmax}(x/\tau)$ , where  $0 < \tau \leq 1$  is a temperature value controlling the activation steepness. When  $\tau$  is close to 0, the  $\tau$ -softmax returns outputs close to one-hot vectors; when  $\tau$  is 1, the function is the standard softmax.

The NRM implementation with parametric models is represented as follows:

$$\begin{aligned} p_p^{(t)} &= \text{sg}(s^{(t)}; \theta_{sg}) \\ q_p^{(t)} &= \sum_{i=1}^{i=|P|} p_p^{(t)}[i] (q_p^{(t-1)}(\text{softmax}_\tau(M_t, \tau))[i]) \\ r_p^{(t)} &= q_p^{(t)} \text{softmax}_\tau(M_r, \tau) \end{aligned} \quad (6)$$

where the symbol grounding function,  $\text{sg}$  is now a parametric function with learnable parameters  $\theta_{sg}$ , and  $M_t$  and  $M_r$  are matrices of parameters. We process  $M_t$  and  $M_r$  with  $\text{softmax}_\tau$  before using them in order to ensure they tend to represent a (deterministic) Moore Machine as the temperature decreases. The discretization of parameters in  $M_t$  and  $M_r$  through temperature decreasing is unnecessary and can be omitted if the machine is initialized with a known Moore Machine. It is only required when *learning* the machine from data. A scheme of the proposed implementation is shown in Figure 1(c). The model in equation 6 can be seen as three parametric functions (NNs) producing the three sequences of probability vectors as output, when given a sequence of states  $x_s$  as inputs, namely:  $x_{p_p} = \mathcal{G}(x_s | \theta_{sg})$ ,  $x_{q_p} = \mathcal{Q}(x_s | \theta_{sg}, M_t)$ , and  $x_{r_p} = \mathcal{R}(x_s | \theta_{sg}, M_t, M_r)$ .

### 4.3 Reasoning and Learning with NRM

As anticipated the framework is designed to allow both reasoning and learning and both initialization with background knowledge and train-

ing with data of each module of the machine. Specifically, different configurations of prior knowledge on the NRM parameters ( $\theta_{sg}$ ,  $M_t$ , and  $M_r$ ) and of the possibility to observe rewards in the environment or not result in different possible uses of NRMs, that we describe here.

**Pure Reasoning** By feeding the machine with a time sequence of states  $x_s$  observed in the environment, we can calculate the sequence of probabilities on the machine states,  $x_{q_p} = \mathcal{Q}(x_s | \theta_{sg}, M_t)$ , and reward values,  $x_{r_p} = \mathcal{R}(x_s | \theta_{sg}, M_t, M_r)$  which can be used in the same manner as RMs do: for designing non-Markovian reward functions, for augmenting the environment state, for Counterfactual Experience or hierarchical RL [15]. This process requires prior knowledge of all the NRM parameters  $\theta_{sg}$ ,  $M_t$ , and  $M_r$  and does not require any reward feedback from the environment.

**Pure Learning** In case we do not know any of the NRM parameters but we can observe in the environment sequences of states  $x_s$  and ground truth sequences of rewards  $x_r$  we can exploit the latter to learn the NRM modules entirely from data. It is obtained by minimizing the cross-entropy loss (or other classification loss functions) between the rewards probabilities predicted by the NRM and the ground-truth rewards.

$$L(x_s, x_r) = \text{cross-entropy}(\mathcal{R}(x_s | \theta_{sg}, M_t, M_r), x_r) \quad (7)$$

**Learning and Reasoning Integration** If only *some* of the NRM modules are known, and sequences of states and rewards are observable in the environment, we can train the missing modules to align with both prior knowledge and data. This entails initializing the known modules with the available knowledge and training the remaining parameters using loss function in Equation 7. Various configurations are feasible in this scenario. Here, we specifically focus on **symbol grounding**: learning  $\theta_{sg}$  while assuming  $M_t$  and  $M_r$  are known. Another reasoning-learning configuration, which we are currently exploring, is that of **automata learning**, where  $\theta_{sg}$  is known while  $M_t$  and  $M_r$  are unknown, which we let for future research.

### 4.4 Exploiting NRMs for non-Markovian RL

In this paper, we assume that an agent has to learn to perform a temporally extended task by optimizing a non-Markovian reward signal. As is standard in RL, the agent interacts with the environment by taking an action and it receives a new state of the environment and a new reward in response. Additionally, we assume that the agent

has a *purely symbolic* knowledge of the task it is executing. For example, in a scenario like the Minecraft game depicted in Figure 1(a), the task for the agent could be "reach the door, the pickaxe, and the lava in any possible order." In this case, the agent would have access to the Moore Machine shown in Figure 1(b). However, it lacks knowledge on how to recognize a pickaxe, a gem, or a door within its environment. In this setting, traditional RMs are inapplicable because of the missing symbol grounding function, while purely deep-learning-based approaches have no way to exploit the symbolic knowledge and will rely only on rewards. With our framework we exploit both symbolic knowledge and data from the environment by interleaving RL and semi-supervised symbol grounding of the NRM. In particular, we initialize the parameters  $M_t$  and  $M_r$  with the task prior knowledge, and the parameters of the grounder  $\theta_{sg}$  with random weights. We record each episode of interaction, as a sequence of states  $x_s$  and a sequence of rewards  $x_r$ . We calculate the machine state sequence  $x_{qp}$  with the NRM as  $\mathcal{Q}(x_s | \theta_{sg}, M_t, M_r)$ , and we augment each environment state with the predicted machine state (**reasoning step**). The interaction is led by a RL algorithm running on the augmented state  $(s^{(t)}, q_p^{(t)})$ . Since the symbol grounder is initially randomized, the predicted machine state initially deviates from the ground truth, and the state representation is not perfect. However, at regular intervals, we update the symbol grounder by training on sequences  $x_t$  and  $x_r$  collected in the environment so to minimize the loss in Equation 7 (**learning step**). As the agent observes more scenarios in the environment, the symbol grounding function of the model gradually becomes more similar to the unknown ground truth, and so does the distribution of machine states. In case the perfect SG function is learned the NRM becomes equivalent to a RM.

**Reward Values as Supervision** Since the reward values serve as supervision for the NRM, we have observed that overly sparse rewards are not very effective for symbol grounding. For instance, consider the automaton shown in Figure 1(b), with 8 states, of which only one is final (state 2). If we employ a sparse signal, rewarding the agent solely upon task completion, this entails assigning a positive reward to state 2 and 0 reward to all the other states. Subsequently, all the unsuccessful episodes would provide no feedback to the symbol grounder. Moreover, in case the agent accomplishes the task, the grounder would glean feedback solely to discern that the *last* observation leading to winning reward ought to be associated with one of the three symbols bringing from a non rewarding state to state 2: lava (L), door (D), or pickaxe (P). However, it would lack the means to distinguish among these three symbols in the last observation, or discern which symbols the agent has encountered in the previous observations. We, therefore, assume a potential-based reward shaping [26], where the potential function varies based on the distance on the automaton from the nearest final state. This partitions the states of the machine in figure according to 4 reward values,  $r_1, r_2, r_3$  and  $r_4$ . Reward shaping is very commonly applied in RM applications to facilitate policy learning through RL [15]. Note that this type of reward generally does not provide distinct feedback for each single state of the machine. Consequently, while the resulting grounding process is feasible, it is by no means straightforward.

#### 4.5 Groundability Analysis of Temporal Specifications

In our approach RL and Symbol Grounding are closely intertwined and mutually influence each other. SG relies on RL as we collect data through RL exploration for SG training. Simultaneously, RL relies on SG, as it utilizes an estimated representation of the automaton's state, which becomes more accurate as SG accurately predicts symbols.

At the same time, regardless of the specific application, all types of semisupervised SG are affected by some reasoning shortcuts, especially when the knowledge is somewhat 'trivial' - including formulas that are either trivially false or trivially true across most instances - or contains specific symmetries [32].

Therefore, in this section, we delve into the issue of *ungroundability*. We first review the definitions of ungroundability and reasoning shortcuts and we then devise our own algorithm to find all the RS of a given temporal specification depending solely on the specifics, that we call for this reason 'unremovable'.

**Reasoning Shortcuts** Given the problem of grounding symbols of the alphabet  $P$  of a certain logical knowledge  $\phi$  exploiting the knowledge and a dataset  $D$  of couples (data, output of the formula), a RS is a renaming of symbols,  $\alpha : P \rightarrow P$ , different from the identity,  $\exists p \in P | \alpha(p) \neq p$ , which maintains perfect alignment of  $\phi \circ \alpha$  with the data in  $D$ . In [23] they estimate the number of reasoning shortcuts by assuming prior knowledge of the groundtruth SG function  $sg^*$ , the logical knowledge  $\phi$  and the support dataset  $D$ .

$$\#RS(\phi, D, sg^*) = \sum_{\alpha \in \mathcal{A}} \mathbf{1} \left\{ \bigwedge_{d \in D} (\phi \circ \alpha)(sg^*(d)) = \phi(sg^*(d)) \right\} \quad (8)$$

Where  $\mathcal{A}$  is the set of all possible mappings  $\alpha$ . Only the  $\alpha$  mapping each symbol to itself is indeed the correct solution, therefore the knowledge is said to admit RSs on  $D$  if  $\#RS(\phi, D, sg^*) > 1$

**Ungroundability** The following definition of ungroundability is given in [32]: a symbol  $p_i \in P$  is ungroundable through a formula  $\phi$  if there exists a mapping  $\alpha : P \rightarrow P$  that maps  $p_i$  to  $p_j$  with  $p_i \neq p_j$  such that  $\phi \equiv (\phi \circ \alpha)$ . Where  $\equiv$  denotes logical equivalence. Notice that ungroundability is a property of a *symbol* in the *formula*, and does not depend on the particular supporting dataset used in the application, nor it needs the knowledge of  $sg^*$  to be verified. However, [32] does not define a method to actually *find* the mappings  $\alpha$  which maintain logical equivalence between  $\phi$  and  $\phi \circ \alpha$  for a specific  $\phi$ . Nonetheless, these mappings could be identified using a naive brute-force method. This involves examining the semantic equivalence between  $\phi$  and  $\phi \circ \alpha$  for every possible mapping  $\alpha \in \mathcal{A}$  and returning the mapping that yields a positive check. Given that the number of possible mappings is  $|\mathcal{A}| = |P|^{|P|}$  and equivalence checking has an exponential cost, the brute-force approach is very time consuming and can become impracticable for long formulas defined over large alphabets. Here, we introduce a smarter algorithm for identifying knowledge-preserving mappings. This algorithm leverages certain stopping criteria, which we will delineate later, to enhance speed by a factor of  $10^3$ . We term the mappings discovered by this algorithm as *Unremovable Reasoning Shortcuts* (URS), referring to the RS that are *pathological* of the logical knowledge, which would be never eliminated, even considering an ideal support dataset including all the possible observations of the domain.

**Unremovable Reasoning Shortcuts (URS)** To calculate URSs we focus on the complete support  $D^*$ . Notice that, since the support is complete, we can directly create synthetically the complete *symbolic* dataset  $D_{sym}^*$ , and count the number of RS on that without knowing the ground truth symbol grounding function  $sg^*$ . We define the number of reasoning shortcuts under complete support assumption as

$$\#RS^*(\phi) = \sum_{\alpha \in \mathcal{A}} \mathbf{1} \left\{ \bigwedge_{d_{sym} \in D_{sym}^*} (\phi \circ \alpha)(d_{sym}) = \phi(d_{sym}) \right\} \quad (9)$$

**Finding URS of a Temporal Property** From now on, we assume that the logical knowledge  $\phi$  is a Moore Machine, or equivalently a Deterministic Finite Automaton (DFA) (DFAs indeed can be considered MMs with a binary output alphabet), an LTLf, or LDLf formula (which can be automatically translated into DFAs [8]). It is worth to notice that *temporal* formalisms (as LTLf, MMs, etc) do not admit a *finite* complete symbolic support  $D_{sym}^*$ , because strings in any support can always be extended at will in the time dimension. Therefore we cannot verify Equation 9 directly. We present in Algorithm 1 our algorithm for finding URS, which is basically an efficient way to verify Equation 9 on a *finite* symbolic support that can still be considered *complete* for  $\phi$ . We find this kind of support applying the following theorems.

**Theorem 1.** Let  $D_{sym}^*(L)$  denote the set of all strings over  $P$  with maximum length equal to  $L$ . If  $L_1 \leq L_2$ , then  $\#RS^*(\phi, D_{sym}^*(L_2)) \leq \#RS^*(\phi, D_{sym}^*(L_1))$ .

**Corollary 2.** The number of RS calculated on any complete support with finite horizon  $L$  is an upper bound for the number of URS.  $\#RS^*(\phi) \leq \#RS^*(\phi, D_{sym}^*(L)) \forall L < \infty$ .

**Theorem 3.** Let  $A \subseteq Q$  be the set of absorbing states of  $\phi$ .<sup>1</sup> If: (i) a string  $x$  has reached an absorbing state,  $\delta_t^*(q_0, x) \in A$ . (ii) for a certain mapping  $\alpha$ , also  $x \circ \alpha$  has reached an absorbing state,  $\delta_t^*(q_0, x \circ \alpha) \in A$ . (iii)  $\alpha$  is a working map for  $x$ , namely  $\delta_o^{**}(q_0, x) = \delta_o^{**}(q_0, x \circ \alpha)$ . Then  $\alpha$  is a working map also for all the strings having  $x$  as prefix, namely  $\delta_o^{**}(q_0, x + y) = \delta_o^{**}(q_0, (x + y) \circ \alpha) \forall y \in P^*$ .

**Theorem 4.** Given that: (i)  $\alpha$  is a working map, for the string  $x + z$ , with  $x, z \in P^*$ ,  $\delta_o^{**}(q_0, x + z) = \delta_o^{**}(q_0, (x + z) \circ \alpha)$ ; (ii)  $\exists p \in P$  such that both  $x$  and  $x + p$  reach the same state and  $x \circ \alpha$  and  $(x + p) \circ \alpha$  reach the same state,  $\delta_t^*(q_0, x) = \delta_t^*(q_0, x + p) = q \wedge \delta_t^*(q_0, x) = \delta_t^*(q_0, x + p)$ . Then  $\alpha$  works for the string  $x + y + z$ ,  $\delta_o^{**}(q_0, x + y + z) = \delta_o^{**}(q_0, (x + y + z) \circ \alpha)$ ,  $\forall y \in P^*$ .

We report the proofs of Theorems 1, 2, 3, and 4 in [33]. Algorithm 1 keeps a dataset  $D[\alpha]$  for each candidate URS  $\alpha$ . The set of candidates URS is initialized with  $\mathcal{A}$ .  $D[\alpha]$  is initialized with all the possible strings on  $P$  with length one for each  $\alpha$ . At each iteration we verify if  $\alpha$  is a working for dataset  $D[\alpha]$  (namely if Equation 9 with  $D[\alpha]$  in place of  $D_{sym}^*$  and  $\alpha$  in place of  $\mathcal{A}$  gives 1). If it is working it remains a candidate otherwise is discarded. Then we try to extend the datasets of surviving maps of one step. This follows Theorem 1, which says that if a map does not work for a shorter complete dataset has no chance to work in the longer one. Following theorems 3 and 4, we extend a string  $x$  in  $D[\alpha]$  with the symbol  $p \in P$  only if: (i) not both  $x$  and  $x \circ \alpha$  bring to an absorbing state (Theorem 3) and (ii) not both  $\text{state}(x) = \text{state}(x+p)$  and  $\text{state}(x \circ \alpha) = \text{state}((x + p) \circ \alpha)$  (theorem 4). We iterate until at least one dataset has grown with respect to the past iteration. Finally we return the  $\alpha$  remained in the set of candidates as URSs.

## 5 Experiments

In this section, we report the experiments validating our framework. Our code is available on github: <https://github.com/KRLGroup/NeuralRewardMachines>.

<sup>1</sup> A state  $q \in Q$  is called absorbing if  $\delta(q, p) = q \forall p \in P$

<sup>2</sup> with  $P^*$  being the set of strings obtained repeating the symbol  $p$  an arbitrary number of times.

---

### Algorithm 1 Find\_unremovable\_reasoning\_shortcuts( $\phi$ )

---

```

URScandidates  $\leftarrow \mathcal{A}$ 
for  $\alpha$  in URScandidates do
   $D[\alpha] \leftarrow \phi.P$ 
  while URScandidate not empty do
    for  $\alpha$  in URScandidate do
       $D_{next\_alpha} \leftarrow \{\}$ 
      if  $\alpha$  is working on  $D[\alpha]$  then
        for  $x$  in  $D[\alpha]$  do
          if not  $\phi(x).q.abs$  or not  $\phi(\alpha(x)).q.abs$  then
            for  $p$  in  $\phi.P$  do
               $x' \leftarrow x + p$ 
              if  $\phi(x).q \neq \phi(x').q$  or  $\phi(\alpha(x)).q \neq$ 
                 $\phi(\alpha(x')).q$  then
                   $D_{next\_alpha} \leftarrow D_{next\_alpha} + x'$ 
            if  $D_{next\_alpha}$  is empty then
              URScandidate  $\leftarrow$  URScandidate -  $\alpha$ 
            else
               $D[\alpha] \leftarrow D_{next\_alpha}$ 
    return URScandidates

```

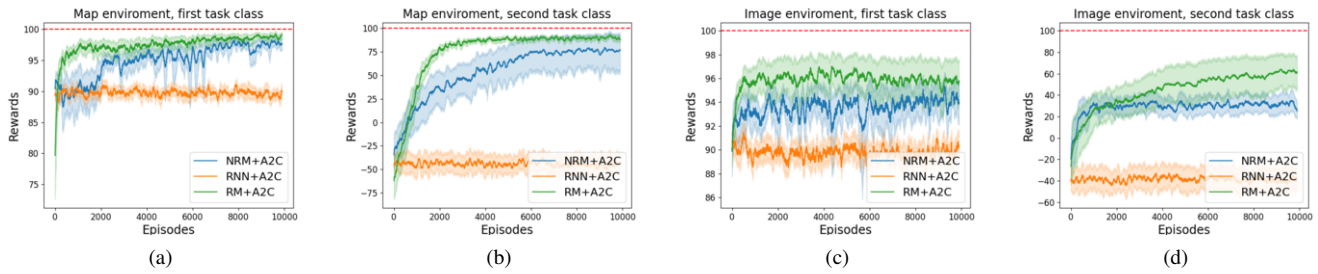
---

**Application** We test NRMs on two types of environment inspired by the Minecraft videogame, similar to those considered by [6] and [20] and described in the example of Section 3.2, in which we assume the SG function is unknown.

**States** Based on this application we construct two environments that showcase different levels of difficulty in terms of symbol grounding: (i) **The map environment**, in which the state is the 2D vector containing the  $x$  and  $y$  current agent location, (ii) **The image environment**, in which the state is an image of  $64 \times 64 \times 3$  pixels showing the agent in the grid, like the one shown in Figure 1(a).

**Rewards** We express the non-Markovian reward as an LTLf formula, that is then transformed into a DFA and, ultimately, into a Moore Machine. In the last process, each state  $q$  is assigned a reward value, which is maximum if  $q$  is a final states and gradually decreases as the distance from  $q$  to the final state increases (as detailed in Section 4.4 and shown in Figure 1(b)). Reward values are scaled so to give maximum cumulative reward always equal to 100. We focus on patterns of formulas that are popular in non-Markovian RL [15, 34], that we call here using the categorization given in [24]. **Visit** formulas: we denote as  $\text{Visit}(p_1, p_2, \dots, p_n)$  the LTLf formula  $F p_1 \wedge F p_2 \wedge \dots \wedge F p_n$ , expressing that the agent has to make true at least once each  $p_i$  in the formula in any possible order. **Sequenced Visit**: we denote with  $\text{Seq\_Visit}(p_1, p_2, \dots, p_n)$  the LTLf formula  $F(p_1 \wedge F(p_2 \wedge \dots \wedge F(p_n)))$ , expressing that we want the symbol  $p_{i+1}$  made true at least once *after* the symbol  $p_i$  has become True. **Global Avoidance**: we denote with  $\text{Glob\_Avoid}(p_1, p_2, \dots, p_n)$  the formula  $G(\neg p_1) \wedge G(\neg p_2) \wedge \dots \wedge G(\neg p_n)$ , expressing that the agent must always avoid to make True the symbols  $p_1, p_2, \dots, p_n$ .  $F$  and  $G$  are temporal operators called respectively ‘eventually’ and ‘globally’, we refer the reader to [27] for the formal semantics of these operators that we explain here only intuitively. Based on these patterns of formulas we construct tasks of increasing difficulty that we aggregate in two classes. The **first class** is a set of tasks obtained as a conjunction of visit and sequenced visit formulas. The **second class** contains conjunctions of visit, sequenced visit and global avoidance formulas. We report in [33] the formulas considered.

**Comparisons** Given the absence of alternative methods in the literature for leveraging ungrounded prior knowledge in non-Markovian



**Figure 2.** Results in the map environment when using a) tasks of the first class, and b) in the second class. Results in the image environment when training on c) the first task class d) the second task class.

Task	#RS*	Alg. 1 exec. time (s)	Brute force exec. time (s)
Class 1	$27.3 \pm 17.8$	$0.6 \pm 0.4$	$434.6 \pm 195.6$
Class 2	$6 \pm 2$	$0.2 \pm 0.0$	$399.7 \pm 24.5$

**Table 1.** Groundability analysis of formulas selected as RL tasks

RL, we compare NRM with the two most popular baselines in this domain: RNNs and RMs. However, it is important to note that the three methods assume different levels of knowledge. Specifically, RM’s performance serves as an upper bound for NRMs, as in the optimal scenario, where NRMs precisely learn the ground-truth SG function (which RMs possess from the outset), they become equivalent to RMs. We use Advantage Actor-Critic (A2C) [25] as RL algorithm for all the methods. Hence we denote the methods (and which information they can exploit for learning): **RNN+A2C** (rewards) the baseline based on RNNs, **NRM+A2C** (rewards+symbolic task) our method, and **RM+A2C** (rewards+symbolic task+grounding) the upper bound obtained with RMs.

**RL Results** In Figure 2, we show the training rewards obtained by the three methods on both the map environment for the first task class (a) and the second task class (b), and the image environment for the first (c) and second (d) task classes. For each task and method, we perform 5 runs with different random seeds. The figure clearly shows that the performance of NRMs falls between that of RNNs and RMs, closely resembling the latter across various experimental configurations, despite assuming less prior knowledge than RMs. This demonstrates that NRMs can effectively leverage the incomplete prior knowledge they possess. Specifically, NRMs converge almost to the same rewards as RMs, especially in the map environment. Generally, RNNs perform much worse than the other methods, often quickly converging to a local maximum of the reward, especially in challenging temporal tasks (class 2). All methods share the same hyperparameter settings for A2C and the neural networks used for the policy, the value, and the feature extraction (only used in the image environment), which we detail in [33] along with information on the network used as the grounder in the two environments.

**Groundability Analysis** Finally, we report the number of URS identified running the Algorithm 1 on the selected tasks, along with the execution time of our algorithm and that of a brute-force version that simply iterates over each possible map  $\alpha$  and checks whether  $\phi \equiv (\phi \circ \alpha)$ . Results are reported in Table 1. The number of URS found confirms our intuition that, as the tasks’ difficulty increases, the number of URS decreases, resulting in a reduction in groundability difficulty. This difficulty had never been investigated before and it is hard to estimate. We start exploring it using the URS tool, showing that the two task classes have significantly different average #RS\*. However, note that the number of URSs for the chosen tasks is always relatively low (27 out of a maximum of  $|P|^{|P|} = 3125$  for 5 symbols), which indicates the tasks chosen are meaningful for be used for SSSG even if they were not originally designed for that. Finally, we found

that our algorithm is very efficient, and it executes a thousand of times faster than the brute force implementation.

## 6 Conclusions and Future Work

In conclusion, we introduce Neural Reward Machines, a neurosymbolic framework providing non-Markovian rewards for non-symbolic-state environments. NRMs possess great versatility and applicability across various learning configurations, with a particular focus here on their capacity for grounding within environments. We address the challenge of semisupervised symbol grounding through temporal specifications and propose an algorithm for identifying Unremovable Reasoning Shortcuts in temporal tasks. We demonstrate the ability of NRMs to integrate the machine’s prior knowledge and leverage it to outperform Deep RL methods, achieving nearly equivalent rewards to RMs, even without possessing full prior knowledge. We leave exploration of other reasoning-learning scenarios, such as the integration of automata learning with RL, for future research.

## Acknowledgements

This work has been partially supported by PNRR MUR project PE0000013-FAIR.

The work of Francesco Argenziano was carried out when he was enrolled in the Italian National Doctorate on Artificial Intelligence run by Sapienza University of Rome.

## References

- [1] F. Bacchus, C. Boutilier, and A. Grove. Rewarding behaviors. pages 1160–1167, Portland, OR, 1996.
- [2] S. Badreddine, A. d’Avila Garcez, L. Serafini, and M. Spranger. Logic tensor networks. *Artificial Intelligence*, 303:103649, 2022. ISSN 0004-3702. doi: <https://doi.org/10.1016/j.artint.2021.103649>. URL <https://www.sciencedirect.com/science/article/pii/S0004370221002009>.
- [3] M. Cai, S. Xiao, B. Li, Z. Li, and Z. Kan. Reinforcement learning based reinforcement logic control with maximum probabilistic satisfaction. *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 806–812, 2020.
- [4] M. Cai, M. Hasanbeig, S. Xiao, A. Abate, and Z. Kan. Modular deep reinforcement learning for continuous motion planning with temporal logic. *IEEE Robotics and Automation Letters*, 6(4):7973–7980, 2021.
- [5] A. Camacho, R. Toro Icarte, T. Q. Klassen, R. Valenzano, and S. A. McIlraith. Ltl and beyond: Formal languages for reward function specification in reinforcement learning. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 6065–6073. International Joint Conferences on Artificial Intelligence Organization, 7 2019.
- [6] J. Corazza, I. Gavran, and D. Neider. Reinforcement learning with stochastic reward machines. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(6):6429–6436, Jun. 2022. doi: [10.1609/aaai.v36i6.20594](https://doi.org/10.1609/aaai.v36i6.20594). URL <https://ojs.aaai.org/index.php/AAAI/article/view/20594>.

- [7] W.-Z. Dai, Q. Xu, Y. Yu, and Z.-H. Zhou. Bridging machine learning and logical reasoning by abductive learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/9c19a2aa1d84e04b0bd4bc888792bd1e-Paper.pdf>.
- [8] G. De Giacomo and M. Y. Vardi. Linear temporal logic and linear dynamic logic on finite traces. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13*, page 854–860. AAAI Press, 2013. ISBN 9781577356332.
- [9] G. De Giacomo, L. Iocchi, M. Favorito, and F. Patrizi. Foundations for restraining bolts: Reinforcement learning with ltlf/ldlf restraining specifications. *Proceedings of the International Conference on Automated Planning and Scheduling*, 29(1):128–136, May 2021. URL <https://ojs.aaai.org/index.php/ICAPS/article/view/3549>.
- [10] D. Furelos-Blanco, M. Law, A. Jonsson, K. Broda, and A. Russo. Induction and exploitation of subgoal automata for reinforcement learning. *J. Artif. Int. Res.*, 70:1031–1116, may 2021. ISSN 1076-9757. doi: 10.1613/jair.1.12372. URL <https://doi.org/10.1613/jair.1.12372>.
- [11] M. Gaon and R. Brafman. Reinforcement learning with non-markovian rewards. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):3980–3987, Apr. 2020. doi: 10.1609/aaai.v34i04.5814. URL <https://ojs.aaai.org/index.php/AAAI/article/view/5814>.
- [12] G. D. Giacomo, L. Iocchi, M. Favorito, and F. Patrizi. Foundations for restraining bolts: Reinforcement learning with ltlf/ldlf restraining specifications, 2019.
- [13] D. Ha and J. Schmidhuber. Recurrent world models facilitate policy evolution. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [14] Y.-X. Huang, W.-Z. Dai, L.-W. Cai, S. H. Muggleton, and Y. Jiang. Fast abductive learning by similarity-based consistency optimization. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 26574–26584. Curran Associates, Inc., 2021. URL <https://proceedings.neurips.cc/paper/2021/file/d7e148cabfd9b608090fa5ee3348bfe-Paper.pdf>.
- [15] R. T. Icarte, T. Q. Klassen, R. A. Valenzano, and S. A. McIlraith. Reward machines: Exploiting reward function structure in reinforcement learning. *J. Artif. Intell. Res.*, 73:173–208, 2022. doi: 10.1613/JAIR.1.12440.
- [16] E. Jang, S. Gu, and B. Poole. Categorical reparameterization with gumbel-softmax. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=rkE3y85ee>.
- [17] S. Kapturowski, G. Ostrovski, W. Dabney, J. Quan, and R. Munos. Recurrent experience replay in distributed reinforcement learning. In *International Conference on Learning Representations*, 2019.
- [18] Y. Kuo, B. Katz, and A. Barbu. Encoding formulas as deep networks: Reinforcement learning for zero-shot execution of LTL formulas. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2020, Las Vegas, NV, USA, October 24, 2020 - January 24, 2021*, pages 5604–5610, 2020. doi: 10.1109/IROS45743.2020.9341325. URL <https://doi.org/10.1109/IROS45743.2020.9341325>.
- [19] R. Kusters, Y. Kim, M. Collery, C. de Sainte Marie, and S. Gupta. Differentiable rule induction with learned relational features. In A. d'Avila Garcez and E. Jiménez-Ruiz, editors, *Proceedings of the 16th International Workshop on Neural-Symbolic Learning and Reasoning as part of the 2nd International Joint Conference on Learning & Reasoning (IJCLR 2022), Cumberland Lodge, Windsor Great Park, UK, September 28-30, 2022*, volume 3212 of *CEUR Workshop Proceedings*, pages 30–44. CEUR-WS.org, 2022. URL <http://ceur-ws.org/Vol-3212/paper3.pdf>.
- [20] A. C. Li, Z. Chen, P. Vaezipoor, T. Q. Klassen, R. T. Icarte, and S. A. McIlraith. Noisy symbolic abstractions for deep RL: A case study with reward machines. *CoRR*, abs/2211.10902, 2022. doi: 10.48550/arXiv.2211.10902. URL <https://doi.org/10.48550/arXiv.2211.10902>.
- [21] M. L. Littman, U. Topcu, J. Fu, C. L. I. Jr., M. Wen, and J. MacGlashan. Environment-independent task specifications via GLTL. *CoRR*, abs/1704.04341, 2017. URL <http://arxiv.org/abs/1704.04341>.
- [22] R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, and L. De Raedt. Deepprolog: Neural probabilistic logic programming. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/file/dc5d637ed5e62c36ecb73b654b05ba2a-Paper.pdf>.
- [23] E. Marconato, S. Teso, A. Vergari, and A. Passerini. Not all neuro-symbolic concepts are created equal: Analysis and mitigation of reasoning shortcuts, 2023.
- [24] C. Menghi, C. Tsigkanos, P. Pelliccione, C. Ghezzi, and T. Berger. Specification patterns for robotic missions. *IEEE Trans. Software Eng.*, 47(10):2208–2224, 2021. doi: 10.1109/TSE.2019.2945329. URL <https://doi.org/10.1109/TSE.2019.2945329>.
- [25] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937. PMLR, 2016.
- [26] A. Y. Ng, D. Harada, and S. Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *Proceedings of the Sixteenth International Conference on Machine Learning (ICML 1999), Bled, Slovenia, June 27 - 30, 1999*, pages 278–287, 1999.
- [27] A. Pnueli. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October - 1 November 1977*, pages 46–57. IEEE Computer Society, 1977. doi: 10.1109/SFCS.1977.32. URL <https://doi.org/10.1109/SFCS.1977.32>.
- [28] A. Ronca, G. P. Licks, and G. D. Giacomo. Markov abstractions for PAC reinforcement learning in non-markov decision processes. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 3408–3415, 2022. doi: 10.24963/ijcai.2022/473. URL <https://doi.org/10.24963/ijcai.2022/473>.
- [29] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018. URL <http://incompleteideas.net/book/the-book-2nd.html>.
- [30] E. Tsamoura, T. Hospedales, and L. Michael. Neural-symbolic integration: A compositional perspective. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(6):5051–5060, May 2021. doi: 10.1609/aaai.v35i6.16639. URL <https://ojs.aaai.org/index.php/AAAI/article/view/16639>.
- [31] E. Umili, F. Argenziano, A. Barbin, and R. Capobianco. Visual reward machines. In *Proceedings of the 17th International Workshop on Neural-Symbolic Learning and Reasoning, La Certosa di Pontignano, Siena, Italy, July 3-5, 2023*, pages 255–267, 2023. URL <https://ceur-ws.org/Vol-3432/paper23.pdf>.
- [32] E. Umili, R. Capobianco, and G. De Giacomo. Grounding LTL Specifications in Image Sequences. In *Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning*, pages 668–678, 8 2023. doi: 10.24963/kr.2023/65. URL <https://doi.org/10.24963/kr.2023/65>.
- [33] E. Umili, F. Argenziano, and R. Capobianco. Neural reward machines, 2024. URL <https://arxiv.org/abs/2408.08677>.
- [34] P. Vaezipoor, A. C. Li, R. T. Icarte, and S. A. McIlraith. Lt2action: Generalizing LTL instructions for multi-task RL. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, pages 10497–10508, 2021.
- [35] C. K. Verginis, C. Köprülü, S. Chinchali, and U. Topcu. Joint learning of reward machines and policies in environments with partially known semantics. *CoRR*, abs/2204.11833, 2022. doi: 10.48550/arXiv.2204.11833. URL <https://doi.org/10.48550/arXiv.2204.11833>.
- [36] H. Walke, D. Ritter, C. Trimbach, and M. Littman. Learning finite linear temporal logic specifications with a specialized neural operator, 2021. URL <https://arxiv.org/abs/2111.04147>.
- [37] J. Xu, Z. Zhang, T. Friedman, Y. Liang, and G. Van den Broeck. A semantic loss function for deep learning with symbolic knowledge. In J. Dy and A. Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5502–5511. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/xu18h.html>.
- [38] Z. Xu, B. Wu, A. Ojha, D. Neider, and U. Topcu. Active finite reward automaton inference and reinforcement learning using queries and counterexamples. In *Machine Learning and Knowledge Extraction - 5th IFIP TC 5, TC 12, WG 8.4, WG 8.9, WG 12.9 International Cross-Domain Conference, CD-MAKE 2021, Virtual Event, August 17-20, 2021, Proceedings*, pages 115–135, 2021.