

Do It for HER: First-Order Temporal Logic Reward Specification in Reinforcement Learning

Pierriccardo Olivieri¹, Fausto Lasca¹, Alessandro Gianola², Matteo Papini³

¹DEIB, Politecnico di Milano, Milan, Italy

²INESC-ID/Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal

³Dipartimento di Informatica, Università degli Studi di Milano, Milan, Italy

pierriccardo.olivieri@polimi.it, fausto.lasca@mail.polimi.it,
alessandro.gianola@tecnico.ulisboa.pt, matteo.papini@unimi.it

Abstract

In this work, we propose a novel framework for the logical specification of non-Markovian rewards in Markov Decision Processes (MDPs) with large state spaces. Our approach leverages Linear Temporal Logic Modulo Theories over finite traces (LTLfMT), a more expressive extension of classical temporal logic in which predicates are first-order formulas of arbitrary first-order theories rather than simple Boolean variables. This enhanced expressiveness enables the specification of complex tasks over unstructured and heterogeneous data domains, promoting a unified and reusable framework that eliminates the need for manual predicate encoding. However, the increased expressive power of LTLfMT introduces additional theoretical and computational challenges compared to standard LTLf specifications. We address these challenges from a theoretical standpoint, identifying a fragment of LTLfMT that is tractable but sufficiently expressive for reward specification in an infinite-state-space context. From a practical perspective, we introduce a method based on reward machines and Hindsight Experience Replay (HER) to translate first-order logic specifications and address reward sparsity. We evaluate this approach to a continuous-control setting using Non-Linear Arithmetic Theory, showing that it enables natural specification of complex tasks. Experimental results show how a tailored implementation of HER is fundamental in solving tasks with complex goals.

1 Introduction

In Reinforcement Learning (RL), an agent interacts with an environment and receives a scalar reward signal as feedback, which it uses to improve its decision-making over time. In principle, under the *reward hypothesis* (Sutton and Barto 2018), any task can be modeled by specifying an appropriate reward function. However, reward engineering is a difficult task, often heuristic in nature and reliant on domain knowledge. Furthermore, many tasks of interest are inherently non-Markovian, making commonly used Markovian reward functions less effective for learning. Introducing structure in the reward via logic specification helps address these challenges by bringing reward engineering closer to human language and improving interpretability. In *logic reward specification*, the reward the agent receives at each time step depends on whether a formula is satisfied based on

the set of true propositional symbols at that time step. This approach allows expressing complex tasks using *temporal logic*. To better picture this, consider a robot operating in a warehouse, having to deliver a box to a specific location. This task has two sequential sub-goals: (A) picking up the box in a given position (x_o, y_o) , then (B) delivering it to the designated location (x_g, y_g) . This behavior can be enforced via a logic specification that encodes the temporal relationship between events A and B, such as "Do A, then B".

Currently, most works in this line employ logic specification based on Linear Temporal Logic with finite traces (LTLf) (De Giacomo and Vardi 2013; Brafman, De Giacomo, and Patrizi 2018). Stemming from its counterpart with infinite traces (LTL (Pnueli 1977a)), which is the most common language for specifying system properties in the field of formal verification. In the last decade, LTLf has also gained popularity in AI, where reasoning over finite executions, such as in planning tasks, is more intuitive (Giacomo and Rubín 2018). LTLf offers sufficient expressiveness for many *discrete* domains, and benefits from a solid toolchain for translating such specifications into automata or "reward machines" (Icarte et al. 2018). However, the expressiveness of LTLf is limited to Boolean predicates, which constrain the complexity of the reward specification. Evaluating the truth of these predicates typically requires defining *encoding functions* manually. This limitation is often not an issue in discrete domains or for simple reward functions. However, in many real-world scenarios, such as business process management (Gianola, Montali, and Winkler 2024) or robotics (Li, Vasile, and Belta 2017), a richer reward specification *language* allows providing equally or more informative feedback to the learner *with less human intervention* and can make the specification of new tasks for the same environment significantly faster. The importance of an expressive specification language is even more evident in the presence of *heterogeneous* data types, as shown in the following extended example.

Warehouse robot example. The robot must reach the object's position, identify a specific object by its ID, and satisfy a weight constraint. The resulting predicate might be " $A = (x - x_o)^2 + (y - y_o)^2 < r^2 \wedge i = \text{"H123"} \wedge \text{weight} < 10$ ". The robot has a *heterogeneous* state described by variables: $x, y, w \in \mathbb{R}^+$, representing position and object weight ($w = 0$ if no object is being carried) and $i \in \{\text{H123}, \text{S456}, \dots\}$

is a given set of IDs.

Typically, evaluating such a predicate requires some hand-crafted encoders, one to determine when the Euclidean distance between the agent’s position (x, y) and the objects location (x_o, y_o) falls below a threshold r , and another to check the ID and find the object’s weight in an inventory database. Even in this relatively simple example, implementing the specification involves (i) manually encoding each condition, and likely (ii) producing a specification that is not reusable across different tasks. Without such user-defined encodings, handling continuous or heterogeneous states is impossible in LTLf, which operates at the propositional level.

To overcome these limitations, we propose using Linear Temporal Logic Modulo Theories over finite traces (LTLfMT) (Geatti, Gianola, and Gigante 2022; Geatti et al. 2023; Faella and Parlato 2024; Rodríguez and Sánchez 2024, 2023; Rodríguez et al. 2025), a extension of LTLf in which the atomic propositions are replaced by first-order formulas interpreted over (combinations of) arbitrary first-order theories. LTLfMT enables reasoning about reward specifications using diverse first-order theories such as arithmetic over the reals and/or integers, uninterpreted functions and relations (to represent, e.g., relational databases (Ghilardi et al. 2023),(Gianola 2023)), complex datatypes (e.g., lists, arrays) and even custom domain-specific predicates, providing the expressiveness needed to naturally encode more complex specifications. From a practical perspective, instead of manually defining separate encoders for each condition, we propose a unified framework that requires only a formula for the task and an assignment of values to constants. The formula can be expressed using different theories and evaluated via a universal Satisfiability Modulo Theories (SMT) solver (Barrett et al. 2021), without the need to explicitly encode it as a propositional predicate. Continuing with the delivery robot example, the predicate A discussed earlier can be expressed and evaluated in LTLfMT using the (non-linear) arithmetic theory over the reals, and evaluated via a standard SMT solver without the need for hand-crafted, domain-dependent encodings. By shifting the complexity to the logical specification layer, this framework allows for the definition of more expressive behaviors than LTLf, while maintaining generality and reusability across different tasks and domains though the use of off-the-shelf SMT solvers (Barbosa et al. 2022; de Moura and Bjørner 2008). Theoretically, the LTLfMT framework can be viewed as a composable method for reward specification where each theory acts as a building block, and adding more theories (or combinations thereof) increases the expressiveness of the framework. If no additional theories are included beyond Boolean predicates, we recover the standard LTLf semantics. To achieve greater expressivity, we can incorporate and combine arbitrary theories such as arithmetic and database theories. However, the choice of theories added to LTLfMT (some theories may lack available solvers) and some advanced features of the full logic (such as quantifiers and variable-lookahead operators) affect both decidability and solver support, impacting the framework’s efficiency and practical usability. It is thus important to identify sound

but sufficiently expressive fragments of the logic and well-supported theories that can handle the data types characteristic of the domain of interest.

Contributions and paper structure. After discussing related works in Section 2 and introducing the necessary technical background in Section 3, we present *two main contributions*— In **Section 4**, we define a theoretical framework for reward specifications by extending the classical LTLf logic to the more expressive LTLfMT, in which atomic propositions are replaced with first-order formulas over arbitrary theories. Then, we identify a fragment of LTLfMT that (i) is free from all decidability issues other than those introduced by the theory itself (ii) allows for easy translation of formulas to automata, and yet (iii) is expressive enough to capture most goals of practical interest— In **Section 5** we apply the proposed framework when the integrated theory is Non-Linear Real Arithmetic Theory (NRA) to handle reward specifications in a continuous-control scenario. In this domain, we address the problem of reward sparsity that is typical of logic specifications, exploiting the intrinsically “structured” nature of goals in our first-order framework. We propose a solution that combines Counterfactual Experiences for Reward Machines (CRM) and Hindsight Experience Replay (HER) (Andrychowicz et al. 2017) in a nontrivial way. Empirical results on continuous control tasks with complex goals suggest an improvement w.r.t. using only CRM (Section 6). The extended version of this paper can be found at: <https://bit.ly/4i20C4Z>.

2 Related Works

We discuss related works in terms of the two key challenges addressed in this paper: dealing with non-boolean data and dealing with reward sparsity.

Propositional languages equipped with labeling functions. Given the intrinsic “discrete” nature of LTL-based rewards, most works focus on MDPs with finite states (Bacchus, Boutilier, and Grove 1996; Thiébaux, Kabanza, and Slaney 2002; Gretton, Price, and Thiébaux 2003; Gretton et al. 2011; Gretton 2014; Brafman, De Giacomo, and Patrizi 2018; Bozkurt et al. 2020). There are notable exceptions: Cai et al. (2021) consider continuous MDP states, but logic formulas are defined using propositional letters corresponding to higher level concepts such as “regions” of the state space. Their definition is extra-logical and left to the user’s implementation. This approach is more explicit in the framework of “restraining bolts” (De Giacomo et al. 2019, 2020a,b) in terms of “features” mapping properties of the “world” (not necessarily observed by the agent) to propositional atoms. When this mapping is applied directly to the states of the MDP, it is usually called *labeling function* (Lacerta, Parker, and Hawes 2015; Hasanbeig, Kroening, and Abate 2020, 2023; Jiang et al. 2021).

Reward Machines (RM) (Icarte et al. 2018, 2022) are a programmatic alternative to logic specification where the user directly builds a finite-state automaton. In fact, we can identify the automaton generated by an LTLf formula as a type of RM called *simple reward machine* with a binary reward function. As in many works on logic discussed above,

the mapping between states of the MDP and propositional letters must be implemented separately and is also called a labeling function in this context.

The use of labeling functions has allowed the application of both logically-specified rewards and reward machines to MDPs with continuous states (and actions), opening the door to *deep RL* experiments (e.g., Hasanbeig, Kroening, and Abate 2020; Li et al. 2024). However, in all the aforementioned examples, the labeling function (or equivalent) is treated as a black box. This approach complicates generalization between tasks, is prone to errors, and undermines the interpretability of the logical formulas. A notable exception is the temporal logic proposed by Li, Vasile, and Belta (2017) for robotic control, which allows first-order terms of the kind $f(s) < c$, where s is a continuous state of the MDP, c a constant and f a function. We will similarly argue that inequality terms of this kind are enough to encode most goals of interest for the setting of continuous control. However, the first-order terms of Li, Vasile, and Belta (2017) are mere propositional atoms for which the form of the labeling function is prescribed in advance. We will show in which sense LTLfMT is more versatile.

Existing remedies to sparse rewards. By definition, goals specified by means of logical formulas tend to generate *sparse* (i.e., rare) rewards. These are notoriously hard to optimize in RL and led to the early development of general-purpose *reward shaping* strategies (Dorigo and Colombetti 1994; Ng, Harada, and Russell 1999)—the addition of extra feedback to guide the agent towards the goal. Within the recent *deep RL* literature, in particular regarding robotics applications, Hindsight Experience Replay (HER) (Andrychowicz et al. 2017) has established itself as a powerful alternative, if not the standard approach to goal-based tasks. However, HER works by exploiting the state-generalization capabilities of neural networks applied to goals defined as states or functions of state variables. This shared structure of goals and states is typically absent from specifications based on propositional logics or finite-state automata, or hidden inside black-box components such as labeling functions. The problem of sparsity was explicitly addressed in the RM literature, where several ad-hoc approaches were already proposed by (Icarte et al. 2022), namely Counterfactual experiences for Reward Machines (CRM), Hierarchical approaches (HRM), and Automated Reward Shaping (ARS). They also suggested to explore synergies between RM and HER, but to the best of our knowledge the idea was not explored further. Hierarchical approaches, in the form of temporal abstractions, can be found also in the LTL literature (Araki et al. 2021). Notably, Jothimurugan, Alur, and Bastani (2019) jointly propose a (propositional) reward-specification language and a reward-shaping method.

3 Preliminaries

In this section, we review the main concepts of RL and reward specification based on temporal logic.

Reinforcement learning. The reinforcement learning problem (Sutton and Barto 2018) is usually defined as the

solution to a Markov Decision Process (MDP) (Puterman 1990). An MDP consists of a tuple $\langle S, A, R, P, \gamma \rangle$ where S is the set of states, A is the set of actions and $R : S \times A \rightarrow \mathbb{R}$ is the reward function. The environment’s dynamics are described by the transition function $P : S \times A \times S \rightarrow [0, 1]$, which represents the probability of transitioning to state s' given we are in state $s \in S$ and take action $a \in A$. The discount factor $\gamma \in [0, 1)$ weighs the importance of future rewards versus immediate ones. The agent acts according to a policy $\pi : S \rightarrow \Delta(A)$, which assigns a probability distribution over actions depending on the observed state $s \in S$. The optimal policy is the one yielding the largest *return*, that is the expected discounted sum of rewards collected from the initial state following the policy. Tasks specifications based on temporal logic or reward machines naturally induce non-Markovian reward functions. This more expressive framework is commonly referred to as a Non-Markovian-Reward Decision Process (NMRDP). Formally an NMRDP is a tuple $\langle S, A, \bar{R}, P, \gamma \rangle$ equivalent to an MDP except for the non-Markovian reward function $\bar{R} : (S \times A)^* \rightarrow \mathbb{R}$, which depends on the entire history of state-action pairs (Bacchus, Boutilier, and Grove 1996). Classical reinforcement learning methods cannot be applied directly to NMRDPs.

Linear temporal logic over finite traces. The LTLf logic (De Giacomo and Vardi 2013) is a variant of LTLf (Pnueli 1977b) defined over finite traces. Logic formulas specified in LTLf are evaluated from propositional Boolean symbols defined in the alphabet \mathcal{P} . A finite trace $\sigma = \langle \sigma_0, \sigma_1, \dots, \sigma_n \rangle$ is a sequence of *states* (not to be confused with the states of the MDP). Each state $\sigma_i \subseteq \mathcal{P}$ is a set of propositional symbols true at the i -th timestep, where the finite integer $n \in \mathbb{N}$ specifies the last time step. In LTLf, a formula φ is defined by the grammar $\varphi := p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid X\varphi \mid \varphi_1 \mathcal{U} \varphi_2$, where $p \in \mathcal{P}$, X is the *next* operator and \mathcal{U} is the *until* operator. Other operators can be defined via standard abbreviations, for instance $\varphi_1 \vee \varphi_2 := \neg(\neg\varphi_1 \wedge \neg\varphi_2)$, the *weak next* operator $\tilde{X}\varphi := \neg X\neg\varphi$ and the *eventually* operator $F\varphi := \top \mathcal{U} \varphi$. A formula φ is evaluated over trace σ at a time step t , which is formally denoted as $(\sigma, t) \models \varphi$. We adopt the standard semantic of LTLf, defined by induction on the formula (De Giacomo and Vardi 2013). The key aspects on top of the semantic of propositional logic are: $(\sigma, t) \models X\varphi$ if and only if $(\sigma, t + 1) \models \varphi$, and $(\sigma, t) \models \varphi_1 \mathcal{U} \varphi_2$ if and only if $(\sigma, j) \models \varphi_2$ for some $j \geq t$ and $(\sigma, k) \models \varphi_1$ for all $t \leq k < j$. In reward specification based on temporal logic, we interpret propositional symbols in \mathcal{P} as world features, which the agent receive at each time step t . At each time step t , the (binary) reward obtained by the agent is decided by evaluating φ over the current logical state σ_t . Typically, the propositional symbols or fluents in \mathcal{P} represent truth values about the MDP state. For example, (De Giacomo et al. 2019) refer to this encoding as a *feature* function $f_i : W \rightarrow D_i$ mapping a world state w from the set of world states W to a finite domain D such as $\{\text{True}, \text{False}\}$. A state of the trace $\sigma_t = \langle f_1(w_t), f_2(w_t) \dots \rangle$ is obtained by extracting the features f_i from the current world state w_t at time t . The clear non-Markovian nature of this reward signal makes the decision problem a NMRDP. A widely adopted

solution (e.g., De Giacomo et al. 2020b) is to convert φ into a finite automaton $\mathcal{A}_\varphi = \langle Q, 2^{\mathcal{P}}, \delta, q_0, F \rangle$, where Q is a set of automaton states, $2^{\mathcal{P}}$ is the input alphabet, δ is the transition function, $q_0 \in Q$ is the initial state, and $F \subseteq Q$ the set of accepting (terminal) states. The automaton accepts a trace $\sigma = \langle \sigma_0, \sigma_1, \dots, \sigma_n \rangle$ if, starting from q_0 and sequentially consuming inputs from σ , it reaches an accepting state $q \in F$. By combining the NMRDP \mathcal{M} with the automaton \mathcal{A}_φ , we construct a product MDP $\mathcal{M}' := \langle S', A', R', P', \gamma \rangle$ where the state space $S' = S \times Q$ tracks both the environment state and the automaton state. Transition and rewards in the product MDP are Markovian, as the automaton state capture the necessary information about the history. Hence, \mathcal{M}' is equivalent to the original NMRDP, but can be solved using standard RL techniques.

4 Reward Specification via LTLfMT

LTLfMT (Geatti, Gianola, and Gigante 2022) extends the expressiveness of temporal logic by allowing predicates in \mathcal{P} to be first-order formulas over arbitrary theories. A predicate $p \in \mathcal{P}$ is no longer a Boolean symbol but can instead represent a first-order formula, interpreted according to a chosen background theory. In this section, we formally define the syntax of LTLfMT showing how different theories are integrated into this framework. Then we propose a propositionalizable *fragment* of LTLfMT retaining sufficient expressiveness for reward specification in RL. Finally, we show how to translate a formula of this fragment into a Deterministic Finite Automaton (DFA) that can be used to define a product MDP for reinforcement learning.

4.1 Linear Temporal Logic Modulo Theories with finite traces

To represent predicates as first order formulas, in LTLfMT, the alphabet \mathcal{P} is replaced by a multi-sorted first-order signature $\Sigma = \mathcal{S} \cup \mathcal{P} \cup \mathcal{C} \cup \mathcal{F} \cup \mathcal{V} \cup \mathcal{W}$ where, \mathcal{S} is a set of sort symbols (i.e., data types), \mathcal{C} is a set of constants, \mathcal{P} is the set of predicates,¹ \mathcal{F} the set of functions, \mathcal{V} and \mathcal{W} are the sets of variables and quantified variables respectively. A Σ -term t serves as a building block for constructing atomic formulas in LTLfMT, and is generated by the following grammar:

$$t := v \mid w \mid c \mid f(t_1, \dots, t_k) \mid \bigcirc v \mid \odot v \quad (1)$$

where $c \in \mathcal{C}$, $v \in \mathcal{V}$, $w \in \mathcal{W}$ and $f \in \mathcal{F}$ is a function symbol of arity k . The operators \bigcirc and \odot denote the lookahead and weak lookahead operators, respectively, similar to the next X and weak next \tilde{X} in LTLf, but acting on Σ -variables; they are used to relate the value of variables at a given time to their value at the following instant. The complete grammar of LTLfMT is defined as:

$$\begin{aligned} \alpha &:= p(t_1, \dots, t_k), \\ \lambda &:= \alpha \mid \neg \alpha \mid \lambda_1 \vee \lambda_2 \mid \lambda_1 \wedge \lambda_2 \mid \exists v \lambda \mid \forall v \lambda, \\ \varphi &:= \top \mid \lambda \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \wedge \varphi_2 \mid X\varphi \mid \tilde{X}\varphi \mid \varphi_1 \mathcal{U} \varphi_2. \end{aligned} \quad (2)$$

¹Note that in LTLfMT there is no alphabet. Hence, from now on \mathcal{P} denotes the predicate symbols of the first-order signature Σ .

Compared to LTLf, this grammar introduces two additional layers α and λ . Intuitively, α defines atomic predicates over terms, and λ allows first-order combinations over these atoms. The top-layer formulas φ are standard temporal formulas over these first-order expressions, similar to LTLf. Theories are integrated into this framework via the interpretation of Σ -terms. A theory \mathcal{T} can be defined as a set of Σ -structures, i.e., sets that interpret all the symbols in Σ (so, also Σ -terms) which are called models of the theory. Intuitively, a Σ -signature defines the available symbols (e.g., 1, 2, +, >, ...), a Σ -structure provides a concrete interpretation (e.g., $+(1, 2)$ is interpreted as integer addition). A theory \mathcal{T} provides interpretations for each symbol in Σ . For example, the Linear Integer Arithmetic (LIA) theory introduces a *sort* for Integer data types into \mathcal{S} , arithmetic operations such as + and its inverse - in the set \mathcal{F} , constants and variables valued in \mathbb{Z} via \mathcal{C} , \mathcal{V} and \mathcal{W} , and comparison predicates (e.g., =, <) into \mathcal{P} . Multiple theories can be combined, including their Σ -signatures, and their interpretations are given through combined Σ -structures. Formally a \mathcal{T} -state $s = (M, \mu)$, contains two elements, $M \in \mathcal{T}$ (i.e., a \mathcal{T} -model) is a first-order Σ -structure that interprets all the symbols in Σ and $\mu : \mathcal{V} \rightarrow \text{dom}(M)$ a function mapping non-quantified variables to their domain of reference. Quantified variables are interpreted by a separate environment function $\xi : \mathcal{W} \rightarrow \text{dom}(M)$. A trace (or word), in LTLfMT is a finite sequence of \mathcal{T} -state over time-steps, $\sigma = \langle (M, \mu_0), \dots, (M, \mu_n) \rangle$ where $M \in \mathcal{T}$ is fixed across the trace, and μ_i may differ, reflecting the evolution of variables over time. The evaluation of a term t at time i in a trace σ under variable environment ξ is defined using the functions μ_i and ξ in the base case of variables in \mathcal{V} and \mathcal{W} , resp., and defined inductively in the usual way for constants and functional terms. The satisfaction relation for an LTLfMT formula φ over a trace σ , at time-step i with environment function ξ is denoted as $\sigma, \xi, i \models \varphi$ and follows the usual definition of first-order semantics. This satisfaction relation is used to define the *satisfaction modulo \mathcal{T}* of an LTLfMT formula φ over the trace σ at time point time-step i , denoted as $\sigma, i \models \varphi$ and inductively defined in the standard way to deal with the temporal operators (see (Geatti, Gianola, and Gigante 2022) for the formal definition of the semantics).

4.2 Fragment for Reward Specification in RL

LTLfMT framework allows for very general expressions, and it is in general undecidable (Geatti, Gianola, and Gigante 2022). It is possible to define a fragment of LTLfMT that is semi-decidable. Formally, for decidable underlying first-order theories, LTLfMT has been proven to be semi-decidable (i.e., a positive answer can always be obtained for satisfiable instances, while reasoning might not terminate over unsatisfiable instances), thanks to the semi-decision procedure shown by Geatti, Gianola, and Gigante (2022), and also implemented in the BLACK solver.² Intuitively, for identifying a decidable fragment of LTLfMT, the underlying theory \mathcal{T} must necessarily be decidable, but this is not

²<https://www.black-sat.org/en/stable/>

sufficient in general: decidability in first-order extensions of LTLf may still be broken by the interaction of temporal operators and quantifiers, i.e., in the case of LTLfMT, by the presence of the lookaheads operators. We propose a tractable fragment of LTLfMT for reward specification in reinforcement learning by removing the (weak) lookahead operator \bigcirc (\odot) from the Σ -layer. While the (weak) next operator X (\tilde{X}) in the temporal logic layer is essential for specifying temporal constraints between “subtasks” that the agent may have to complete, the addition of (weak) lookahead operator on the Σ -term, \bigcirc (\odot), which applies to variables, is neither natural nor beneficial for reward specification. The Σ -layer can be regarded as a data collection interface for capturing unstructured and heterogenous environmental information. We argue that defining temporal constraints on data collection is not useful for reward specification, and that the Σ -layer should just be used to capture immediate information available to the agent. Note that our fragment of LTLfMT still supports arbitrary first-order theories.

$$\begin{aligned}
t &:= v \mid w \mid c \mid f(t_1, \dots, t_k) \\
\alpha &:= p(t_1, \dots, t_k) \\
\lambda &:= \alpha \mid \neg\alpha \mid \lambda_1 \vee \lambda_2 \mid \lambda_1 \wedge \lambda_2 \mid \exists v\lambda \mid \forall v\lambda \\
\varphi &:= \top \mid \lambda \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \wedge \varphi_2 \mid X\varphi \mid \tilde{X}\varphi \mid \varphi_1 \mathcal{U} \varphi_2
\end{aligned} \tag{3}$$

In this fragment, the source of undecidability becomes the theory itself (and the use of quantifiers). Assuming the theory itself is decidable, we can leverage any tool implementing a decision procedure for the theory itself to deal with the first-order dimension: this is the case if the quantifier-free fragment of the theory is decidable and the quantifiers do not affect decidability because, for instance, a quantifier elimination procedure is available (as in the case of real arithmetic (Tarski 1951)). In such scenarios, we can employ a state-of-the-art QE module for eliminating the quantifiers (Collins 1974), and then universal tools such as SMT solvers (Barrett et al. 2021) to evaluate the obtained quantifier-free formula. Depending on the theory, in cases where the use of quantifiers leads to undecidability, we can further restrict our fragment by removing them. Doing so, we obtain a LTLfMT fragment similar to “data-LTLf” by Gianola, Montali, and Winkler (2024).

4.3 Translation to DFA and Product MDP

Geatti, Gianola, and Gigante (2022) give a non-constructive proof that a Boolean abstraction of any language definable in lookahead-free LTLfMT can also be defined in LTLf, and vice versa. We take a more practical viewpoint and provide a propositionalization of our fragment in the following sense: (i) a syntactic transformation of lookahead-free LTLfMT formulas into LTLf formulas, and (ii) a transformation of LTLfMT traces into LTLf traces. This is enough to define an automaton and a product MDP that can be simulated efficiently using only a theory solver for the elected theory (such as an SMT solver if the theory is quantifier-free). Refer to the grammar of our fragment in Equation (3). Given a formula φ from our language, we first substitute each α containing only free variables from \mathcal{V} with a distinct propositional letter. Then, we substitute each remaining λ

with a distinct propositional letter. The result is a LTLf formula φ' . Let \mathcal{P} now denote the set of all the employed propositional letters—our new alphabet. A LTLfMT trace $\sigma = (s_0, \dots, s_t)$ is mapped to an LTLf trace $\sigma' = (\sigma'_1, \dots, \sigma'_n)$ by assigning a truth value to each propositional letter in \mathcal{P} at each step. This is done, e.g., with a theorem prover (Kovács and Voronkov 2013) or a decision procedure for the theory of choice (or an SMT solver if the theory is quantifier-free, as detailed in Section 4.2). Crucially, the problem given to the solver at a given timestep is purely first-order, without any temporal operator. For example, assuming the theory is real arithmetic, if “ $\forall y \exists z (y * z > 0 \implies z < y)$ ” from ϕ was replaced with letter P in ϕ' , the solver decides $P \in \sigma'_t$ if and only if $s_t \models \forall y \exists z (y * z > 0 \implies z < y)$, applying efficient quantifier elimination first. At runtime, the propositional abstraction φ' is evaluated on the LTLf trace σ' . The SMT solver plays the role of an extremely general *labeling function* (cf. Section 2). For an RL-friendly implementation, we can translate φ' into a Deterministic Finite Automaton (DFA) $\mathcal{A}_{\varphi'} = \langle Q, 2^{\mathcal{P}}, \delta, q_0, F \rangle$ and use it to define a product MDP $M' := \langle S', A', R', P', \gamma \rangle$ with extended state space $S' = S \times Q$, exactly like in LTLf (cf. Section 3). We can do this automatically using the translation technique defined in (Zhu, Pu, and Vardi 2019)³ At training time, state transitions in the DFA are resolved by the SMT solver. This is a special kind of reward machine (specifically, a *simple* reward machine with a Boolean-valued reward function, cf. (Icarte et al. 2018)), again with the SMT solver playing the role of the labeling function.

5 Reinforcement Learning with LTLfMT Rewards

In this section, we provide an RL solution to the tasks defined using our proposed specification language, with a focus on the continuous control-domain.

Let us first recap our reward specification pipeline, stressing that it reduces significantly manual interventions compared to other methods. The end goal of the specification is to construct a product MDP to which standard RL algorithms (and more sophisticated techniques inspired by reward machines) can be applied. First, depending on the domain of interest (e.g., continuous control), we pick a theory (e.g., NRA) with an associated SMT solver (with all the precautions of Section 4.2). Several different tasks can be specified within the domain of interest. From the user defining the task, or *reward engineer*, the only input we require is a formula φ written in the lookahead-free LTLfMT fragment defined in Section 4.2. Compared to existing reward specification frameworks (cf. Section 2), the user need not manually implement a labeling function or anything similar. At most, to enable *multi-goal RL* (for example, to use HER as we will detail later) the user needs to provide a dictionary assignment of first-order constants appearing in the formula to numerical values. In this way, the task is “parametrized” by (a subset of) the constants. In the following, we will focus on continuous control, both to provide an intuitive, realistic

³In our implementation, we use the LTLf2DFA library (<https://github.com/whitemech/LTLf2DFA>)

example of our reward specification pipeline and study the effectiveness of HER in dealing with sparse rewards generated by logic specifications.

Continuous Control Case Study Common tasks in continuous control consist in controlling a variable to a target value or range, or a vector of variables to a target point. For example, a target position in space for a robot, a desired temperature for a HVAC system, a speed for cruise control. These "subtasks" can be described in terms of L^p distances between state vectors (e.g., Euclidean distance, box constraints) This can be done using real arithmetic, and subtasks can be combined into more complex tasks using temporal operators. Hence, for continuous control, we can select Non-Linear Real Arithmetic (NRA) as our theory for lookahead-free LTLfMT, with all the computational advantages discussed in Section 4.2. For example, consider an autonomous car moving in 2-dimensional space within a parking area. A state $s \in \mathbb{R}^2$ represents spatial coordinates. We specify a (parametric) task where the agent must reach a point A in space first, then a point B , without ever going through a certain "unsafe" area. The task can be specified in our lookahead-free LTLfMT fragment as $\varphi = \neg F \neg (x \geq x_{min} \wedge x \leq x_{max}) \wedge F((x - x_a)^2 + (y - y_a)^2 < r_a^2 \wedge F((x - x_b)^2 + (y - y_b)^2 < r_b^2))$. In this case, we use an SMT solver suited for Linear Real Arithmetic (LRA) and Non-Linear Real Arithmetic (NRA) to propositionalize φ and obtain $\varphi' := \neg F \neg \alpha_1 \wedge F(\alpha_2 \wedge F(\alpha_3))$.

We use this example to stress the advantages of our reward specification framework. First, the input we require to the user to specify the task is very simple. It consists of a simple string representing the specification formula φ , along with an optional assignment of constants to numerical values to parametrize the task.⁴ In contrast, existing approaches often require complex inputs. For example, reward machines as proposed by (Icarte et al. 2018) require explicit specification of the automaton. They also require an implementation of the labeling function, just as several logic-based approaches discussed in 2. Moreover, defining new tasks for the same domain can be very quick in our framework. In the continuous control domain we have outlined, NRA allows defining many new tasks by *just* writing new formulas. Even more substantial modifications to the domain just require changing the theory (for example, from NRA to linear arithmetic), which could be almost effortless depending on the availability of efficient solvers. Finally, in the case of continuous control, our framework makes it easy to handle sparse rewards with HER, as shown in the next section.

5.1 Addressing Sparsity in the Product MDP

A well known issue of logic-based reward specification methods is the sparsity of rewards. For the continuous-control domain, we propose to combine the idea of counterfactual reasoning (CRM) from reward machines (Icarte et al. 2022) with Hindsight Experience Replay (HER). Both techniques work at the data collection level, generating artificial experiences to improve learning. Both can be applied

⁴This can be as simple as a json configuration file specifying the names of the constants and the values.

to trajectories generated from the product MDP $\tau = \{ \langle s_0, q_0 \rangle, \langle a_0, r_1 \rangle, \langle s_1, q_1 \rangle, \dots, \langle s_T, q_T \rangle \}$. Counterfactual reasoning exploits the underlying reward automaton \mathcal{A}_φ , producing additional "fake" experience by replacing the DFA state with other possible ones. This is always possible since the DFA, differently from the MDP, is simulated by the agent itself. Namely, given a single experience $\{ \langle s_t, q_t \rangle, \langle a_t, r_{t+1} \rangle, \langle s_{t+1}, q_{t+1} \rangle \}$, we can create $|Q|$ artificial transitions replacing q_t with any other $q \in Q$ and recomputing the reward. HER is designed for goal-based scenarios with sparse binary reward feedback. It improves sample efficiency by learning from unsuccessful trajectories, exploiting the state-generalization capabilities of neural networks. HER requires the definition of a "goal-based problem" consisting of a goal space \mathcal{G} , a predicate reward function $f_g : S \rightarrow \{0, 1\}$ that checks whether goal $g \in \mathcal{G}$ is achieved in state $s \in S$, and a mapping $m : S \rightarrow \mathcal{G}$ s.t. $\forall s \in S, f_{m(s)}(s) = 1$ used to generate artificial goals $g' = m(s)$. A goal-based transition is a tuple $((s_t, g), \langle a_t, r_{t+1} \rangle, (s_{t+1}, g))$. First, we show how to adapt CRM and HER to our framework, then we combine them to exploit their synergy to mitigate reward sparsity.

Adapting HER to the product MDP. HER requires a goal specification, which is not intuitive to define over the states of the product MDP and is usually provided by the user. To automate HER, we define as "true" goal of the product MDP the state $\langle s, q \rangle \in S'$, where q is the predecessor of an accepting state $q_f \in F$ on the automaton $\mathcal{A}_{\varphi'}$ and s is such that $q_f = \delta(q, L(s, \cdot, \cdot))$, that is a state triggering the transition to q_f . Then, the goal-based specification can be defined as follows: $\mathcal{G} = S'$, $f_g(s) = \mathbb{I}[m(s) = g]$, and $m(s) = s$, where $s \in S'$ is a state of the product MDP.

Combining CRM and HER. Given a trajectory τ , we add to the replay buffer $|Q|$ artificial experiences obtained by copying τ and applying counterfactual reasoning as described above. Then, we apply HER to each (real or counterfactual) trajectory generating a total of $2|Q|$ artificial experiences, all added to the replay buffer. The intermediate goals of HER are defined using the last state visited in each trajectory τ (Andrychowicz et al. 2017). In the next section we compare both the separate and combined approaches.

6 Experiments

In this section we use Deep Deterministic Policy Gradients (DDPG) (Lillicrap et al. 2016) to learn a policy from the product MDP constructed for the *parking environment*. We use as baselines different approaches, operating at the *replay buffer* level, presented in Section 5.1. **Baseline:** the task is learned directly from the product MDP using plain DDPG. **CRM:** CRM is applied to the DFA that specifies the task generating virtual experiences for DDPG. **HER:** HER is applied directly to the augmented state of the product MDP, without using CRM. **CRM-HER:** our main proposed approach, where HER is applied to real experiences and counterfactual experiences generated with CRM. Our goal is to assess how HER and CRM-HER compare to prior methods across tasks of varying complexity, and whether the sup-




Figure 1: Success rate of each baselines for parking environment in 4 different tasks, with 95% bootstrap confidence interval over 20 independent experiments. The x-axis reports the training steps of DDPG. Task names are specified above each plot and are ordered left-to-right by task complexity.

posed synergy of CRM and HER translates into practical advantages.

Environments and tasks. We conducted experiments on a range of tasks (i.e., different φ formulae) defined within two distinct continuous control environments: *Parking* and *Reacher*. The *Parking* environment, from the *HighwayEnv* collection (Leurent 2018), involves an agent controlling a car in an empty parking lot. We modified this environment to accommodate task specification defined by arbitrary LTLfMT-NRA formulas. In each of the two underlying environments, we defined multiple tasks of increasing complexity using our specification pipeline detailed in Section 5. The complexity of these tasks stems from two main factors: the size and structure of the formula used to define the task (reflected in the complexity of the resulting DFA/product MDP) and the difficulty of discovering the goal, which becomes particularly challenging due to the presence of sparse rewards. Complete details about task specification (including LTLfMT formulas) and results for *Reacher* are in appendix.

Analysis of results. Figure 1 presents the results obtained by averaging performance over multiple training runs for each method. Performance is measured as the *success rate* during evaluation episodes conducted periodically throughout training. An episode is considered successful if the agent reaches the goal within the episode’s time limit. Our results show that, in most tasks, the baseline method struggles to consistently learn a policy capable of reaching the objective, lagging behind the other approaches. This trend holds across all difficulty levels but becomes especially pronounced in the more complex tasks, where it completely fails to discover a successful policy (e.g., see results for tasks *parking_1*, *parking_2*). In most cases, CRM is more effective at handling the increase in complexity of the DFA. This is particularly evident in *tasks_1* and *task_2* of the Parking environment, where its performance is on par with the best-performing methods. However, when the goal is difficult to discover, CRM alone offers little benefit, and its performance falls behind that of HER-based approaches (see especially *parking_1*). RM-HER exhibits mixed performance, with results varying depending on the nature of the task. In tasks where the primary challenge lies in goal discovery, it performs well, clearly demonstrating the benefits of HER

(e.g., see *task_2*, *parking_2*). In other tasks, however, the improvement over the baseline is less pronounced or, in some cases, barely noticeable. CRM-HER consistently performs well across all tasks, either achieving the best results or coming close to the top-performing method. The results indicate that CRM-HER effectively combines the strengths of both CRM and HER, enabling it to handle increasing task complexity. The performance advantage becomes especially clear in the more challenging tasks—most notably in *parking_2* task, where it is the only method able to learn a successful policy. The code and the experimental results can be found at: <https://bit.ly/4i20C4Z>.

7 Conclusion

We proposed a new framework for reward specification in RL based on a fragment of the LTLfMT logic. Compared to LTLf, our language is more expressive as it allows predicates to be first-order logic formulas over arbitrary theories. Although the definition of the logic-augmented MDP follows the standard procedure of LTLf, the increased expressiveness of the language has two important consequences for the user responsible of specifying the reward: (i) the user is relieved of the burden of *coding* labeling functions, which are replaced by off-the shelf SMT solvers, only having to *write a logical formula* and (ii) in continuous-control domains, with minimal effort, they can provide a goal specification for HER using the same variables used in the formula. We exploit this second advantage to address the inevitable sparsity of the rewards resulting from such a logic-based specification, combining the CRM technique of reward machines with HER. Experiments in benchmark continuous-control domains on tasks specified with lookahead-free LTLf-modulo-NRA show promising results. Future work should explore integration of these techniques with hierarchical RL methods (Icarte et al. 2022). Finally, our reincorporation of first-order terms *inside* the specification language paves the way to future further combinations of symbolic aspects with *continuous-space* RL, such as the specification of safety constraints, of continuous rewards, of multi-task problems, ultimately to the formal verification of learning processes in continuous domains.

Acknowledgments

This work was partially supported by the ‘OptiGov’ project, with ref. n. 2024.07385.IACDC (DOI: 10.54499/2024.07385.IACDC), fully funded by the ‘Plano de Recuperação e Resiliência’ (PRR) under the investment ‘RE-C05-i08 - Ciência Mais Digital’ (measure ‘RE-C05-i08.m04’), framed within the financing agreement signed between the ‘Estrutura de Missão Recuperar Portugal’ (EMRP) and Fundação para a Ciência e a Tecnologia, I.P. (FCT) as an intermediary beneficiary. This work was also partly supported by Portuguese national funds through Fundação para a Ciência e a Tecnologia, I.P. (FCT), under projects UID/50021/2025 and UID/PRR/50021/2025.

References

- Andrychowicz, M.; Crow, D.; Ray, A.; Schneider, J.; Fong, R.; Welinder, P.; McGrew, B.; Tobin, J.; Abbeel, P.; and Zaremba, W. 2017. Hindsight Experience Replay. In *NIPS*, 5048–5058.
- Araki, B.; Li, X.; Vodrahalli, K.; DeCastro, J. A.; Fry, M. J.; and Rus, D. 2021. The Logical Options Framework. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, 307–317. PMLR.
- Bacchus, F.; Boutilier, C.; and Grove, A. 1996. Rewarding behaviors. In *Proceedings of the National Conference on Artificial Intelligence*, 1160–1167.
- Barbosa, H.; Barrett, C. W.; Brain, M.; Kremer, G.; Lachnitt, H.; Mann, M.; Mohamed, A.; Mohamed, M.; Niemetz, A.; Nötzli, A.; Ozdemir, A.; Preiner, M.; Reynolds, A.; Sheng, Y.; Tinelli, C.; and Zohar, Y. 2022. cvc5: A Versatile and Industrial-Strength SMT Solver. In Fisman, D.; and Rosu, G., eds., *Tools and Algorithms for the Construction and Analysis of Systems - 28th International Conference, TACAS 2022, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Munich, Germany, April 2-7, 2022, Proceedings, Part I*, volume 13243 of *Lecture Notes in Computer Science*, 415–442. Springer.
- Barrett, C. W.; Sebastiani, R.; Seshia, S. A.; and Tinelli, C. 2021. Satisfiability Modulo Theories. In Biere, A.; Heule, M.; van Maaren, H.; and Walsh, T., eds., *Handbook of Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, 1267–1329. IOS Press.
- Bozkurt, A. K.; Wang, Y.; Zavlanos, M. M.; and Pajic, M. 2020. Control Synthesis from Linear Temporal Logic Specifications using Model-Free Reinforcement Learning. In *ICRA*, 10349–10355. IEEE.
- Brafman, R. I.; De Giacomo, G.; and Patrizi, F. 2018. LTLf/LDLf Non-Markovian Rewards. In *AAAI*, 1771–1778. AAAI Press.
- Cai, M.; Hasanbeig, M.; Xiao, S.; Abate, A.; and Kan, Z. 2021. Modular Deep Reinforcement Learning for Continuous Motion Planning With Temporal Logic. *IEEE Robotics Autom. Lett.*, 6(4): 7973–7980.
- Collins, G. E. 1974. Quantifier elimination for real closed fields by cylindrical algebraic decomposition-preliminary report. *SIGSAM Bull.*, 8(3): 80–90.
- De Giacomo, G.; Favorito, M.; Iocchi, L.; Patrizi, F.; and Ronca, A. 2020a. Temporal Logic Monitoring Rewards via Transducers. In *KR*, 860–870.
- De Giacomo, G.; Iocchi, L.; Favorito, M.; and Patrizi, F. 2019. Foundations for restraining bolts: Reinforcement learning with LTLf/LDLf restraining specifications. In *Proceedings of the international conference on automated planning and scheduling*, volume 29, 128–136.
- De Giacomo, G.; Iocchi, L.; Favorito, M.; and Patrizi, F. 2020b. Restraining Bolts for Reinforcement Learning Agents. In *AAAI*, 13659–13662. AAAI Press.
- De Giacomo, G.; and Vardi, M. Y. 2013. Linear Temporal Logic and Linear Dynamic Logic on Finite Traces. In *Ijcai*, volume 13, 854–860.
- de Moura, L. M.; and Bjørner, N. S. 2008. Z3: An Efficient SMT Solver. In Ramakrishnan, C. R.; and Rehof, J., eds., *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings*, volume 4963 of *Lecture Notes in Computer Science*, 337–340. Springer.
- Dorigo, M.; and Colombetti, M. 1994. Robot Shaping: Developing Autonomous Agents Through Learning. *Artif. Intell.*, 71(2): 321–370.
- Faella, M.; and Parlato, G. 2024. A Unified Automata-Theoretic Approach to LTL_f Modulo Theories. In Endriss, U.; Melo, F. S.; Bach, K.; Diz, A. J. B.; Alonso-Moral, J. M.; Barro, S.; and Heintz, F., eds., *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)*, volume 392 of *Frontiers in Artificial Intelligence and Applications*, 1254–1261. IOS Press.
- Geatti, L.; Gianola, A.; and Gigante, N. 2022. Linear Temporal Logic Modulo Theories over Finite Traces. In *IJCAI*, volume 22, 2641–2647.
- Geatti, L.; Gianola, A.; Gigante, N.; and Winkler, S. 2023. Decidable Fragments of LTL_f Modulo Theories. In Gal, K.; Nowé, A.; Nalepa, G. J.; Fairstein, R.; and Radulescu, R., eds., *ECAI 2023 - 26th European Conference on Artificial Intelligence, September 30 - October 4, 2023, Kraków, Poland - Including 12th Conference on Prestigious Applications of Intelligent Systems (PAIS 2023)*, volume 372 of *Frontiers in Artificial Intelligence and Applications*, 811–818. IOS Press.
- Ghilardi, S.; Gianola, A.; Montali, M.; and Rivkin, A. 2023. Safety Verification and Universal Invariants for Relational Action Bases. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China*, 3248–3257. ijcai.org.
- Giacomo, G. D.; and Rubin, S. 2018. Automata-Theoretic Foundations of FOND Planning for LTLf and LDLf Goals. In Lang, J., ed., *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI*

- 2018, July 13-19, 2018, Stockholm, Sweden, 4729–4735. ijcai.org.
- Gianola, A. 2023. *Verification of Data-Aware Processes via Satisfiability Modulo Theories*, volume 470 of *Lecture Notes in Business Information Processing*. Springer.
- Gianola, A.; Montali, M.; and Winkler, S. 2024. Linear-Time Verification of Data-Aware Processes Modulo Theories via Covers and Automata. In *AAAI*, 10525–10534. AAAI Press.
- Gretton, C. 2014. A More Expressive Behavioral Logic for Decision-Theoretic Planning. In *PRICAI*, volume 8862 of *Lecture Notes in Computer Science*, 13–25. Springer.
- Gretton, C.; Kabanza, F.; Price, D.; Slaney, J. K.; and Thiébaux, S. 2011. Decision-Theoretic Planning with non-Markovian Rewards. *CoRR*, abs/1109.2355.
- Gretton, C.; Price, D.; and Thiébaux, S. 2003. Implementation and Comparison of Solution Methods for Decision Processes with Non-Markovian Rewards. In *UAI*, 289–296. Morgan Kaufmann.
- Hasanbeig, H.; Kroening, D.; and Abate, A. 2023. Certified reinforcement learning with logic guidance. *Artif. Intell.*, 322: 103949.
- Hasanbeig, M.; Kroening, D.; and Abate, A. 2020. Deep Reinforcement Learning with Temporal Logics. In *FORMATS*, volume 12288 of *Lecture Notes in Computer Science*, 1–22. Springer.
- Icarte, R. T.; Klassen, T.; Valenzano, R.; and McIlraith, S. 2018. Using reward machines for high-level task specification and decomposition in reinforcement learning. In *International Conference on Machine Learning*, 2107–2116. PMLR.
- Icarte, R. T.; Klassen, T. Q.; Valenzano, R.; and McIlraith, S. A. 2022. Reward machines: Exploiting reward function structure in reinforcement learning. *Journal of Artificial Intelligence Research*, 73: 173–208.
- Jiang, Y.; Bharadwaj, S.; Wu, B.; Shah, R.; Topcu, U.; and Stone, P. 2021. Temporal-Logic-Based Reward Shaping for Continuing Reinforcement Learning Tasks. In *AAAI*, 7995–8003. AAAI Press.
- Jothimurugan, K.; Alur, R.; and Bastani, O. 2019. A Composable Specification Language for Reinforcement Learning Tasks. In *NeurIPS*, 13021–13030.
- Kovács, L.; and Voronkov, A. 2013. First-Order Theorem Proving and Vampire. In *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, volume 8044 of *Lecture Notes in Computer Science*, 1–35. Springer.
- Lacerda, B.; Parker, D.; and Hawes, N. 2015. Optimal Policy Generation for Partially Satisfiable Co-Safe LTL Specifications. In *IJCAI*, 1587–1593. AAAI Press.
- Leurent, E. 2018. An Environment for Autonomous Driving Decision-Making. <https://github.com/eleurent/highway-env>.
- Li, A. C.; Chen, Z.; Klassen, T. Q.; Vaezipoor, P.; Icarte, R. T.; and McIlraith, S. A. 2024. Reward Machines for Deep RL in Noisy and Uncertain Environments. In *NeurIPS*.
- Li, X.; Vasile, C. I.; and Belta, C. 2017. Reinforcement learning with temporal logic rewards. In *IROS*, 3834–3839. IEEE.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2016. Continuous control with deep reinforcement learning. In *ICLR (Poster)*.
- Ng, A. Y.; Harada, D.; and Russell, S. 1999. Policy Invariance Under Reward Transformations: Theory and Application to Reward Shaping. In *ICML*, 278–287. Morgan Kaufmann.
- Pnueli, A. 1977a. The Temporal Logic of Programs. In *18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October - 1 November 1977*, 46–57. IEEE Computer Society.
- Pnueli, A. 1977b. The temporal logic of programs. In *18th annual symposium on foundations of computer science (sfcs 1977)*, 46–57. ieeec.
- Puterman, M. L. 1990. Markov decision processes. *Handbooks in operations research and management science*, 2: 331–434.
- Rodríguez, A.; Amir, G.; Corsi, D.; Sánchez, C.; and Katz, G. 2025. Shield Synthesis for LTL Modulo Theories. In Walsh, T.; Shah, J.; and Kolter, Z., eds., *AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA*, 15134–15142. AAAI Press.
- Rodríguez, A.; and Sánchez, C. 2023. Boolean Abstractions for Realizability Modulo Theories. In Enea, C.; and Lal, A., eds., *Computer Aided Verification - 35th International Conference, CAV 2023, Paris, France, July 17-22, 2023, Proceedings, Part III*, volume 13966 of *Lecture Notes in Computer Science*, 305–328. Springer.
- Rodríguez, A.; and Sánchez, C. 2024. Adaptive Reactive Synthesis for LTL and LTLf Modulo Theories. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*, 10679–10686. AAAI Press.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning - an introduction, 2nd Edition*. MIT Press.
- Tarski, A. 1951. A decision method for elementary algebra and geometry. *University of California Press*.
- Thiébaux, S.; Kabanza, F.; and Slaney, J. K. 2002. Anytime State-Based Solution Methods for Decision Processes with non-Markovian Rewards. In *UAI*, 501–510. Morgan Kaufmann.
- Zhu, S.; Pu, G.; and Vardi, M. Y. 2019. First-Order vs. Second-Order Encodings for LTLf-to-Automata Translation. arXiv:1901.06108.