



LCRL: Certified Policy Synthesis via Logically-Constrained Reinforcement Learning

Mohammadhosein Hasanbeig^{1(✉)}, Daniel Kroening², and Alessandro Abate¹

¹ Computer Science Department, University of Oxford,
Oxford, UK

{hosein.hasanbeig,alessandro.abate}@cs.ox.ac.uk

² Amazon Inc., London, UK

daniel.kroening@magd.ox.ac.uk



Abstract. LCRL is a software tool that implements model-free Reinforcement Learning (RL) algorithms over unknown Markov Decision Processes (MDPs), synthesising policies that satisfy a given linear temporal specification with maximal probability. LCRL leverages partially deterministic finite-state machines known as Limit Deterministic Büchi Automata (LDBA) to express a given linear temporal specification. A reward function for the RL algorithm is shaped on-the-fly, based on the structure of the LDBA. Theoretical guarantees under proper assumptions ensure the convergence of the RL algorithm to an optimal policy that maximises the satisfaction probability. We present case studies to demonstrate the applicability, ease of use, scalability and performance of LCRL. Owing to the LDBA-guided exploration and LCRL model-free architecture, we observe robust performance, which also scales well when compared to standard RL approaches (whenever applicable to LTL specifications). Full instructions on how to execute all the case studies in this paper are provided on a GitHub page that accompanies the LCRL distribution www.github.com/grockious/lcrl.

Keywords: Model-Free Reinforcement Learning · Policy Synthesis · Linear Temporal Logic · Limit Deterministic Büchi Automata

1 Introduction

Markov Decision Processes (MDPs) are extensively used for problems in which an agent needs to control a process by selecting actions that are allowed at the process' states and that affect state transitions. Decision making problems in MDPs are equivalent to resolving action non-determinism, and result in policy synthesis problems. Policies are synthesised to maximise expected long-term rewards

This work is in part supported by the HiClass project (113213), a partnership between the Aerospace Technology Institute (ATI), Department for Business, Energy and Industrial Strategy (BEIS) and Innovate UK.

© Springer Nature Switzerland AG 2022

E. Abraham and M. Paolieri (Eds.): QEST 2022, LNCS 13479, pp. 217–231, 2022.

https://doi.org/10.1007/978-3-031-16336-4_11

obtained from the process. This paper introduces a new software tool, LCRL, which performs policy synthesis for unknown MDPs when the goal is that of maximising the probability to abide by a task (or constraint) that is specified using Linear Temporal Logic (LTL). LTL is a formal, high-level, and intuitive language to describe complex tasks [9]. In particular, unlike static (space-dependent) rewards, LTL can describe time-dependent and complex non-Markovian tasks that can be derived from natural languages [16, 36, 46]. Any LTL specification can be translated efficiently into a Limit-Deterministic Büchi Automaton (LDBA), which allows LCRL to automatically shape a reward function for the task that is later employed by the RL learner for optimal policy synthesis. LCRL is implemented in Python, the *de facto* standard programming language for machine learning applications.

1.1 Related Work

There exists a few tools that solve control (policy) synthesis in a model-free fashion, but not under full LTL specifications. One exception is the work in [6] which proposes an interleaved reward and discounting mechanism. However, the reward shaping dependence on the discounting mechanism can make the reward sparse and small, which might negatively affect convergence. The work in [17] puts forward a tool for an average-reward scheme based on earlier theoretical work. Other model-free approaches with available code-bases are either (1) focused on fragments of LTL and classes of regular languages (namely finite-horizon specs) or (2) cannot deal with unknown black-box MDPs. The proposed approach in [29, 30] presents a model-free RL solution but for regular specifications that are expressed as deterministic finite-state machines. The work in [10, 11] takes a set of LTL_f/LDL_f formulae interpreted over finite traces as constraints, and then finds a policy that maximises an external reward function. The VSRL software tool [12–14, 28] solves a control synthesis problem whilst maintaining a set of safety constraints during learning.

1.2 Contributions

The LCRL software tool has the architecture presented in Fig. 1, and presents the following features:

- LCRL leverages **model-free RL algorithms**, employing only traces of the system (assumed to be an unknown MDP) to formally synthesise optimal policies that satisfy a given LTL specification with maximal probability. LCRL finds such policies by learning over a set of traces extracted from the MDP under LTL-guided exploration. This efficient, guided exploration is owed to reward shaping based on the automaton [18–21, 23, 26]. The guided exploration enables the algorithm to focus only on relevant parts of the state/action spaces, as opposed to traditional Dynamic Programming (DP) solutions, where the Bellman iteration is exhaustively applied over the whole state/action spaces [5]. Under standard RL convergence assumptions, the LCRL output is an optimal policy whose traces satisfy the given LTL specification with **maximal** probability.

- LCRL is **scalable** owing to LTL-guided exploration, which allows LCRL to cope and perform efficiently with MDPs whose state and action spaces are significantly large. There exist a few LDBA construction algorithms for LTL, but not all of resulting automata can be employed for quantitative model-checking and probabilistic synthesis [31]. The succinctness of the construction proposed in [39], which is used in LCRL, is another contributing factor to LCRL scalability. The scalability of LCRL is evaluated in an array of numerical examples and benchmarks including high-dimensional Atari 2600 games [3, 7].
- LCRL is the first RL synthesis method for LTL specifications in **continuous state/action** MDPs. So far no tool is available to enable RL, whether model-based or model-free, to synthesise policies for LTL on continuous-state/action MDPs. Alternative approaches for continuous-space MDPs [1, 34, 41, 44] discretise the model into a finite-state MDP, or alternatively propose a DP-based method with value function approximation [15].
- LCRL displays **robustness** features to hyper-parameter tuning. Specifically, we observed that LCRL results, although problem-specific, are not significantly affected when hyper-parameters are not tuned with care.




Fig. 1. The LCRL architecture: the inputs to the tool are the environment (MDP) - in particular its states s and corresponding labels $L(s)$ - as well as the LDBA generated from the user-defined LTL specification. The MDP (with state s) and the LDBA (with state q) are synchronised on-the-fly, thus generating the pair (s, q) . A reward r is then automatically generated by LCRL, based on the environment label $L(s)$ and the automaton state q ; actions a are selected accordingly and passed back to the environment (MDP). (Color figure online)

2 Logically-Constrained Reinforcement Learning (LCRL)

We assume the unknown environment is encompassed by an MDP, which in this work is a discrete-time stochastic control processes defined as the tuple $\mathfrak{M} = (\mathcal{S}, \mathcal{A}, s_0, P)$ over a set of continuous states $\mathcal{S} = \mathbb{R}^n$, and where $\mathcal{A} = \mathbb{R}^m$ is a set of continuous actions, and $s_0 \in \mathcal{S}$ is an initial state. $P : \mathcal{B}(\mathcal{S}) \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is a conditional transition kernel which assigns to any pair comprising a

state $s \in \mathcal{S}$ and an action $a \in \mathcal{A}$ a probability measure $P(\cdot|s, a)$ on properly measurable sets on \mathcal{S} [4]. A finite state/action MDP is a special case in which $|\mathcal{S}| < \infty$, $|\mathcal{A}| < \infty$, and $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is a transition probability matrix assigning a conditional probability to enter sets of states in \mathcal{S} . A variable $R(s, a) \sim \mathcal{Y}(\cdot|s, a) \in \mathcal{P}(\mathbb{R}^+)$ is defined over the MDP \mathfrak{M} , representing the reward obtained when action a is taken at a given state s , sampled from the reward distribution \mathcal{T} defined over the set of probability distributions $\mathcal{P}(\mathbb{R}^+)$ on subsets of \mathbb{R}^+ .

LCRL is a policy synthesis architecture for tasks that are expressed as specifications in LTL [18–26]. The LCRL Core in Fig. 1 is compatible with any general RL scheme that conforms with the environment state and action spaces. Inside the LCRL module the MDP and LDBA states are synchronised, resulting in an on-the-fly product MDP. Intuitively, the product MDP encompasses the extra dimension of the LDBA states, which is added to the state space of the original MDP. The role of the added dimension is to track the sequence of labels that have been read across episodes, and thus to act as a memory register for the given task. This allows to evaluate the (partial) satisfaction of the corresponding temporal property. More importantly, this synchronisation breaks down the non-Markovian LTL specification into a set of Markovian reach-avoid components, which facilitates the RL convergence to a policy whose traces satisfy the overall LTL specification. In practice, no product between the MDP and LDBA is computed: the LDBA simply *monitors* traces executed by the agents as episodes of the RL scheme.

Remark 1. The LDBA construction inherently introduces limited form of non-determinism, called ε -transitions, which is treated as an extra action over the original MDP action space [39, 40]. Namely, when there exists a non-deterministic transition in an LDBA state, the MDP action space is augmented with the non-deterministic transition predicate of the LDBA. These non-deterministic transitions are automatically handled by LCRL during the learning and appropriate on-the-fly modifications are carried out, so that the RL agent can learn to deal with those non-deterministic transitions in order to reach the accepting conditions of the LDBA. We emphasise that the underlying assumption in LCRL is that the MDP model is unknown (Fig. 1), and thus a single state is obtained as output when given a state and an action as input. \square

LCRL defines a reward function R for the RL Core, whose objective is to maximise the expected discounted return [42]:

$$\mathbb{E}^\pi \left[\sum_{n=0}^{\infty} \gamma^n R(s_n, a_n) | s_0 = s \right], \quad (1)$$

where $\mathbb{E}^\pi[\cdot]$ denotes the expected value given that the agent follows the policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ from state s ; parameter $\gamma \in [0, 1]$ is a discount factor; and

$s_0, a_0, s_1, a_1 \dots$ is the sequence of state/action pairs, initialised at $s_0 = s$. This reward is intrinsically defined over the product MDP, namely it is a function of the MDP state (describing where the agent is in the environment) and the state of the automaton (encompassing partial task satisfaction). For further details on the LCRL reward shaping, please refer to [18–21, 23, 26].

The discount factor γ affects the optimality of the synthesised policy and has to be tuned with care. There is standard work in RL on state-dependent discount factors [6, 35, 37, 45, 47], which is shown to preserve convergence and optimality guarantees. For LCRL the learner discounts the received reward whenever it’s positive, and leaves it un-discounted otherwise:

$$\gamma(s) = \begin{cases} \eta & \text{if } R(s, a) > 0, \\ 1 & \text{otherwise,} \end{cases} \quad (2)$$

where $0 < \eta < 1$ is a constant [20, 47]. Hence, (1) reduces to an expected return that is bounded, namely

$$\mathbb{E}^\pi \left[\sum_{n=0}^{\infty} \gamma(s_n)^{N(s_n)} R(s_n, \pi(s_n)) \mid s_0 = s \right], \quad 0 < \gamma(s) \leq 1, \quad (3)$$

where $N(s_n)$ is the number of times a positive reward has been observed at state s_n .

For any state $s \in \mathcal{S}$ and any action $a \in \mathcal{A}$, LCRL assigns a quantitative action-value $Q : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, which is initialised with an arbitrary and finite value over all state-action pairs. As the agent starts learning, the action-value $Q(s, a)$ is updated by a linear combination between the current $Q(s, a)$ and the target value:

$$R(s, a) + \gamma \max_{a' \in \mathcal{A}} Q(s', a'),$$

with the weight factors $1 - \mu$ and μ respectively, where μ is the learning rate.

An optimal stationary Markov policy synthesised by LCRL on the product MDP that maximises the expected return, is guaranteed to induce a finite-memory policy on the original MDP that maximises the probability of satisfying the given LTL specification [20]. Of course, in finite-state and -action MDPs, the set of stationary deterministic policies is finite and thus after a finite number of learning steps RL converges to an optimal policy. However, when function approximators are used in RL to tackle extensive or even infinite-state (or -action) MDPs, such theoretical guarantees are valid only asymptotically [21, 24].

2.1 Installation

LCRL can be set up by the `pip` package manager as easy as:

```
pip install lcrl
```

This allows to readily import LCRL as a package into any Python project

```
>>> import lcrl
```

Table 1. List of hyper-parameters and features that can be externally selected

Hyper-parameter	Default Value	Description
<code>algorithm</code>	'ql'	RL algorithm underlying LCRL Core, selected between (cf. Table 2): - 'ql': Q-learning, - 'nfq': Neural Fitted Q-iteration, - 'ddpg': Deep Deterministic Policy Gradient
<code>episode_num</code>	2500	number of learning episodes
<code>iteration_num_max</code>	4000	max number of iterations/steps within each episode
<code>discount_factor</code>	0.95	discounting coefficient η as in (2)
<code>learning_rate</code>	0.9	learning rate parameter μ
<code>epsilon</code>	0.1	value for epsilon-greedy exploration (= 0 for fully greedy)
<code>test</code>	true	run of closed-loop simulations to test the generated policy
<code>save_dir</code>	'./results'	directory address for saving the results
<code>average_window</code>	-1	number of episodes for moving-average window for plots (default value -1 for 30% of <code>episode_num</code>)

and employ its modules. Alternatively, the provided `setup` file found within the distribution package will automatically install all the required dependencies. The installation setup has been tested successfully on Ubuntu 18.04.1, macOS 11.6.5, and Windows 11.

2.2 Input Interface

LCRL training module `lcrl.src.train` inputs two main objects (cf. Fig. 1): an **MDP** black-box object that generates training episodes; and an **LDBA** object; as well as learning hyper-parameters¹ that are listed in Table 1.

MDP: An MDP is an object with internal attributes that are a priori unknown to the agent, namely the state space, the transition kernel, and the labelling function (respectively denoted by \mathcal{S} , P , and L). The states and their labels are observable upon reaching. To formalise the agent-MDP interface we adopt a scheme that is widely accepted in the RL literature [7]. In this scheme the learning agent can invoke the following methods from any state of the MDP:

- `reset()`: this resets the MDP to its initial state. This allows the agent to start a new learning episode whenever necessary.
- `step(action)`: the MDP `step` function takes an action (the yellow signal in Fig. 1) as input, and outputs a new state, i.e. the black signal in Fig. 1.

A number of well-known MDP environments (e.g., the stochastic grid-world) are embedded as classes within LCRL, and can be found within the module `lcrl.src.environments`. Most of these classes can easily set up an MDP object. However, note that the state signal output by the `step` function needs to be fed

¹ These parameters are called hyper-parameters since their values are used to control the learning process. This is unlike other parameters, such as weights and biases in neural networks, which are set and updated automatically during the learning phase.

to a labelling function `state_label(state)`, which outputs a list of labels (in string format) for its input `state` (in Fig. 1, the black output signal from the MDP is fed to the blue box, or labelling function, which outputs the set of label). For example, `state_label(state) = ['safe', 'room1']`. The labelling function `state_label(state)` can then be positioned outside of the MDP class, or it can be an internal method in the MDP class. The built-in MDP classes in `lctrl.src.environments` module have an empty `state_label(state)` method that are ready to be overridden at the instance level:

```

1 # create a SlipperyGrid object
2 gridworld_1 = SlipperyGrid()
3
4 # "state_label" function outputs the label of input state
5 # (input: state, output: string label)
6 def state_label(self, state):
7     # defines the labelling image
8     labels = np.empty([gridworld_1.shape[0], gridworld_1.shape[1]], dtype=object)
9     labels[0:40, 0:40] = 'safe'
10    labels[25:33, 7:15] = 'unsafe'
11    labels[7:15, 25:33] = 'unsafe'
12    labels[15:25, 15:25] = 'goal1'
13    labels[33:40, 0:7] = 'goal2'
14    # returns the label associated with input state
15    return labels[state[0], state[1]]
16
17 # now override the step function
18 SlipperyGrid.state_label = state_label.__get__(gridworld_1, SlipperyGrid)

```

Listing 1.1. Example of `state_label(state)` specification in the MDP object `lctrl.src.environments.gridworld_1`.

LDBA: An LDBA object is an instance of the `lctrl.src.automata.ldba` class. This class is structured according to the automaton construction in [39], and it encompasses modifications dealing with non-determinism, as per Remark 1. The LDBA initial state is numbered as 0, or can alternatively be specified using the class attribute `initial_automaton_state` once an LDBA object is created. The LDBA non-accepting sink state is numbered as -1 . Finally, the set of accepting sets, on which we elaborate further below, has to be specified at the instance level by configuring `accepting_sets` (Listing 1.2 line 1). The key interface methods for the LDBA object are:

- `accepting_frontier_function(state)`: this automatically updates an internal attribute of an LDBA class called `accepting_sets`. This is a list of accepting sets of the LDBA, e.g. $\mathcal{F} = \{F_1, \dots, F_f\}$. For instance, if the set of LDBA accepting sets is $\mathcal{F} = \{\{3, 4\}, \{5, 6\}\}$ then this attribute is a list of corresponding state numbers `accepting_sets = [[3, 4], [5, 6]]`. As discussed above, the `accepting_sets` has to be specified once the LDBA class is instanced (Listing 1.2 line 1). The main role of the accepting frontier function is to determine if an accepting set can be reached, so that a corresponding reward is given to the agent (cf. red signal in Fig. 1). Once an accepting set is visited it will be temporarily removed from the `accepting_sets` until the agent visits all the accepting sets within `accepting_sets`. After that,

`accepting_sets` is reset to the original list. To set up an LDBA class in LCRL the user needs to specify `accepting_sets` for the LDBA. LCRL then automatically shapes the reward function and calls the `accepting_frontier_function` whenever necessary. Further details on the `accepting_frontier_function` and the `accepting_sets` can be found in [18–21,23,26].

- `step(label)`: LDBA `step` function takes a label set, i.e. the blue signal in Fig. 1, as input and outputs a new LDBA state. The label set is delivered to the `step` function by LCRL. The `step` method is empty by default and has to be specified manually after the LDBA class is instantiated (Listing 1.2 line 5).
- `reset()`: this method resets the `state` and `accepting_sets` to their initial assignments. This corresponds to the agent starting a new learning episode.

```

1  goal1_or_goal2 = LDBA(accepting_sets=[[1, 2]])
2
3  # "step" function for the automaton transitions
4  # (input: label, output: automaton_state, non-accepting sink state is "-1")
5  def step(self, label):
6  # state 0
7  if self.automaton_state == 0:
8  if 'epsilon_1' in label:
9  self.automaton_state = 1
10 elif 'epsilon_2' in label:
11 self.automaton_state = 2
12 elif 'unsafe' in label:
13 self.automaton_state = -1 # non-accepting sink state
14 else:
15 self.automaton_state = 0
16 # state 1
17 elif self.automaton_state == 1:
18 if 'goal1' in label and 'unsafe' not in label:
19 self.automaton_state = 1
20 else:
21 self.automaton_state = -1 # non-accepting sink state
22 # state 2
23 elif self.automaton_state == 2:
24 if 'goal2' in label and 'unsafe' not in label:
25 self.automaton_state = 2
26 else:
27 self.automaton_state = -1 # non-accepting sink state
28 # step function returns the new automaton state
29 return self.automaton_state
30
31
32 # now override the step function
33 LDBA.step = step.__get__(goal1_or_goal2, LDBA)

```

Listing 1.2. Example of the specification of the `step(label)` method in the LDBA object `lcr1.automata.goal1_or_goal2` for the LTL specification $(\diamond\Box\text{goal1} \vee \diamond\Box\text{goal2}) \wedge \Box\neg\text{unsafe}$. The non-accepting sink state is numbered as -1 .

If the automaton happens to have ε -transitions, e.g. Fig. 2, they have to distinguishable, e.g. numbered. For instance, there exist two ε -transitions in the LDBA in Fig. 2 and each is marked by an integer. Furthermore, the LDBA class has an attribute called `epsilon_transitions`, which is a dictionary to specify which states in the automaton contain ε -transitions. In Fig. 2, only `state 0`

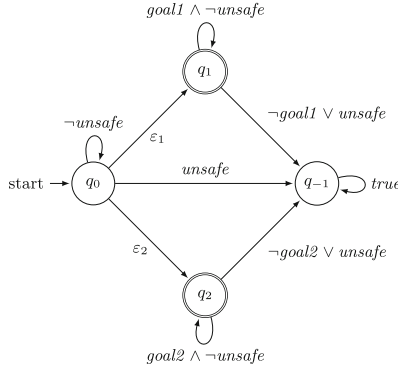


Fig. 2. LDBA for the LTL specification $(\diamond\Box\text{goal1} \vee \diamond\Box\text{goal2}) \wedge \Box\neg\text{unsafe}$.

has outgoing ε -transitions and thus, the attribute `epsilon_transitions` in the LDBA object `goal1_or_goal2` has to be set to

```
goal1_or_goal2.epsilon_transitions = {0:['epsilon.0', 'epsilon.1']}
```

2.3 Output Interface

LCRL provides the results of learning and testing as `.pkl` files. Tests are closed-loop simulations where we apply the learned policy over the MDP and observe the results. For any selected learning algorithm, the learned model is saved as `learned_model.pkl` and test results as `test_results.pkl`. The instruction on how to load these files is also displayed at the end of training for ease of re-loading data and for post-processing. Depending on the chosen learning algorithm, LCRL generates a number of plots to visualise the learning progress and the testing results. These plots are saved in the `save_dir` directory. The user has the additional option to export a generated animation of the testing progress, LCRL prompts this option to the user following the completion of the test. During the learning phase, LCRL displays the progress in real-time and allows the user to stop the learning task (in an any-time fashion) and save the generated outcomes.

3 Experimental Evaluation

We apply LCRL on a number of case studies highlighting its features, performance and robustness across various environment domains and tasks. All the experiments are run on a standard machine, with an Intel Core i5 CPU at 2.5 GHz and with 20 GB of RAM. The experiments are listed in Table 2 and discussed next.

The `minecraft` environment [2] requires solving challenging low-level control tasks (`minecraft-tX`), and features many sequential goals. For instance, in

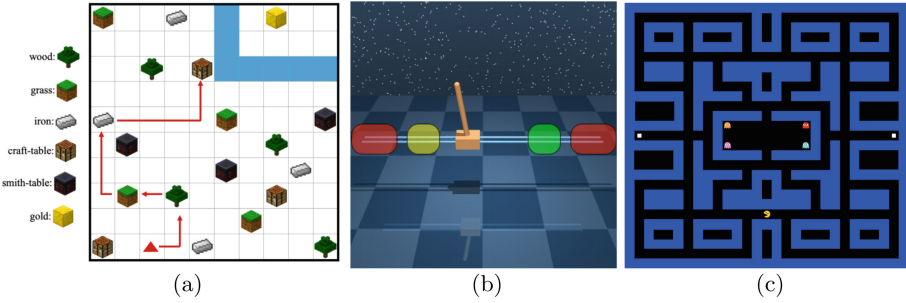


Fig. 3. (a) Synthesised policy by LCRL in `minecraft-t3`; (b) `cart-pole` experiment [43]; (c) `pacman-lrg` - the white square on the left is labelled as food 1 (f_1) and the one on the right as food 2 (f_2), the state of being caught by a ghost is labelled as (g) and the rest of the state space is labelled as neutral (n). (Color figure online)

`minecraft-t3` (Fig. 3a) the agent is tasked with collecting three items sequentially and to reach a final checkpoint, which is encoded as the following LTL specification: $\diamond(\text{wood} \wedge \diamond(\text{grass} \wedge \diamond(\text{iron} \wedge \diamond(\text{craft_table}))))$, where \diamond is the known *eventually* temporal operator.

The `mars-rover` problems are realistic robotic benchmarks taken from [21], where the environment features continuous state and action spaces.

The known `cart-pole` experiment (Fig. 3b) [8, 24, 43] has a task that is expressed by the LTL specification $\square \diamond y \wedge \square \diamond g \wedge \square \neg u$, namely, starting the pole in upright position, the goal is to prevent it from falling over ($\square \neg u$, namely *always not u*) by moving the cart, whilst in particular alternating between the yellow (y) and green (g) regions ($\square \diamond y \wedge \square \diamond g$), while avoiding the red (unsafe) parts of the track ($\square \neg u$).

The `robot-surve` example [38] has the task to repeatedly visit two regions (A and B) in sequence, while avoiding multiple obstacles (C) on the way: $\square \diamond A \wedge \square \diamond B \wedge \square \neg C$.

Environments `slp-easy` and `slp-hard` are inspired by the widely used stochastic MDPs in [42, Chapter 6]: the goal in `slp-easy` is to reach a particular region of the state space, whereas the goal in `slp-hard` is to visit four distinct regions sequentially in a given order.

The `frozen-lake` benchmarks are adopted from the OpenAI Gym [7]: the first three are reachability problems, whereas the last three require sequential visits of four regions, in the presence of unsafe regions to be always avoided.

Finally, `pacman-sml` and `pacman-lrg` are inspired by the well-known Atari game Pacman, and are initialised in a tricky configuration (`pacman-lrg` as in Fig. 3c), which is likely for the agent to be caught: in order to win the game, the agent has to collect the available tokens (food sources) without being caught by moving ghosts. Formally, the agent is required to choose between one of the two available foods and then find the other one ($\diamond[(f_1 \wedge \diamond f_2) \vee (f_2 \wedge \diamond f_1)]$), while

Table 2. Learning results with LCRL. MDP state and action space cardinalities are $|\mathcal{S}|$ and $|\mathcal{A}|$, the number of automaton states in LDBA is denoted by $|\mathcal{Q}|$, the optimal action value function in the initial state is denoted by “LCRL $\max_a Q(s_0, a)$ ”, which represents the LCRL estimation of the maximum satisfaction probability. For each experiment, the reported result includes the mean and the standard error of ten learning trials with LCRL. This probability is also calculated by the PRISM model checker [33] and, whenever the MDP model can be processed by PRISM, it is reported in column “max sat. prob. at s_0 ”. The closer “LCRL $\max_a Q(s_0, a)$ ” and “max sat. prob. at s_0 ” the better. Note that for continuous-state-action MDPs the maximum satisfaction probability cannot be precisely computed by model checking tools, unless abstraction approximation techniques are applied, hence “n/a”. Furthermore, if the MDP state (or action) space is large enough, e.g. `pacman`, the model checkers tools cannot parse the model and the model checking process times out, i.e. “t/o”. The column “LCRL conv. ep.” presents the episode number in which LCRL converged. Finally, “wall_clocktime” presents the average elapsed real time needed for LCRL to converge on a test machine. The rest of the columns provide the values of the hyper-parameters, as described in Table 1.

experiment	MDP	LDBA	LCRL \max_a	max sat.	alg.	episode.	iteration.	discount.	learning.	wall_clock
	$ \mathcal{S} , \mathcal{A} $	$ \mathcal{Q} $	$Q(s_0, a)$	prob. at s_0		num	num_max	factor*	rate [†]	time*(min)
minecraft-t1	100, 5	3	0.991 ± 0.009	1	‘ql’	500	4000	0.95	0.9	0.1
minecraft-t2	100, 5	3	0.991 ± 0.009	1	‘ql’	500	4000	0.95	0.9	0.1
minecraft-t3	100, 5	5	0.993 ± 0.007	1	‘ql’	1500	4000	0.95	0.9	0.25
minecraft-t4	100, 5	3	0.991 ± 0.009	1	‘ql’	500	4000	0.95	0.9	0.1
minecraft-t5	100, 5	3	0.995 ± 0.005	1	‘ql’	500	4000	0.95	0.9	0.1
minecraft-t6	100, 5	4	0.995 ± 0.005	1	‘ql’	1500	4000	0.95	0.9	0.25
minecraft-t7	100, 5	5	0.993 ± 0.007	1	‘ql’	1500	4000	0.95	0.9	0.5
mars-rover-1	∞, 5	3	0.991 ± 0.002	n/a	‘nfq’	50	3000	0.9	0.01	550
mars-rover-2	∞, 5	3	0.992 ± 0.006	n/a	‘nfq’	50	3000	0.9	0.01	540
mars-rover-3	∞, ∞	3	n/a	n/a	‘ddpg’	1000	18000	0.99	0.05	14
mars-rover-4	∞, ∞	3	n/a	n/a	‘ddpg’	1000	18000	0.99	0.05	12
cart-pole	∞, ∞	4	n/a	n/a	‘ddpg’	100	10000	0.99	0.02	1
robot-surve	25, 4	3	0.994 ± 0.006	1	‘ql’	500	1000	0.95	0.9	0.1
slp-easy-sml	120, 4	2	0.974 ± 0.026	1	‘ql’	300	1000	0.99	0.9	0.1
slp-easy-med	400, 4	2	0.990 ± 0.010	1	‘ql’	1500	1000	0.99	0.9	0.25
slp-easy-lrg	1600, 4	2	0.970 ± 0.030	1	‘ql’	2000	1000	0.99	0.9	2
slp-hard-sml	120, 4	5	0.947 ± 0.039	1	‘ql’	500	1000	0.99	0.9	1
slp-hard-med	400, 4	5	0.989 ± 0.010	1	‘ql’	4000	2100	0.99	0.9	5
slp-hard-lrg	1600, 4	5	0.980 ± 0.016	1	‘ql’	6000	3500	0.99	0.9	9
frozen-lake-1	120, 4	3	0.949 ± 0.050	0.9983	‘ql’	400	2000	0.99	0.9	0.1
frozen-lake-2	400, 4	3	0.971 ± 0.024	0.9982	‘ql’	2000	2000	0.99	0.9	0.5
frozen-lake-3	1600, 4	3	0.969 ± 0.019	0.9720	‘ql’	5000	4000	0.99	0.9	1
frozen-lake-4	120, 4	6	0.846 ± 0.135	0.9728	‘ql’	2000	2000	0.99	0.9	1
frozen-lake-5	400, 4	6	0.735 ± 0.235	0.9722	‘ql’	7000	4000	0.99	0.9	2.5
frozen-lake-6	1600, 4	6	0.947 ± 0.011	0.9467	‘ql’	5000	5000	0.99	0.9	9
pacman-sml	729,000, 5	6	0.290 ± 0.035	t/o [‡]	‘ql’	80e3	4000	0.95	0.9	1600
pacman-lrg	4,251,000, 5	6	0.282 ± 0.049	t/o [‡]	‘ql’	180e3	4000	0.95	0.9	3700

* coefficient η in (2) † learning rate μ ‡ timed out: too large for model-checking tools ★ on a machine running macOS 11.6.5 with Intel Core i5 CPU at 2.5 GHz and with 20 GB of RAM

avoiding the ghosts ($\Box\neg g$). We thus feed to the agent a conjunction of these associations, as the following LTL specification: $\Diamond[(f_1 \wedge \Diamond f_2) \vee (f_2 \wedge \Diamond f_1)] \wedge \Box\neg g$. Standard QL fails to find a policy generating satisfying traces for this experiment.

Table 3. Robustness of LCRL performance against hyper-parameter tuning, for the `frozen-lake-1` experiment. Maximum probability of satisfaction is 99.83% as calculated by PRISM (cf. Table 2). The reported values are the percentages of times that execution of LCRL final policy produced traces that satisfied the LTL property. Statistics are taken over 10 trainings and 100 testing for each training, namely 1000 trials for each hyper-parameter configuration.

$\eta \backslash \mu$	0.2	0.4	0.6	0.8	0.99
0.2	92.5 \pm 7.5%	96.7 \pm 3.2%	91.3 \pm 8.7%	98.8 \pm 1.1%	94.7 \pm 5.29%
0.4	98.6 \pm 1.4%	89.5 \pm 10.5%	94.5 \pm 5.5%	94.5 \pm 5.5%	99.2 \pm 0.74%
0.6	99.0 \pm 0.83%	94.5 \pm 5.5%	93.3 \pm 6.7%	96.4 \pm 3.59%	93.3 \pm 6.7%
0.8	95.8 \pm 4.2%	99.5 \pm 0.49%	99.5 \pm 0.49%	96.9 \pm 3.09%	97.7 \pm 2.2%
0.99	88.9 \pm 11.09%	98.4 \pm 1.55%	97.1 \pm 2.31%	96.1 \pm 3.73%	95.2 \pm 4.79%
overall avg.	95.676 \pm 4.268%				

We emphasise that the two tasks in `cart-pole` and `robot-surve` are not co-safe, namely require possibly infinite traces as witnesses.

Additionally, we have evaluated the LCRL robustness to RL key hyper-parameter tuning, i.e. discount factor η and learning rate μ , by training the LCRL agent for 10 times and testing its final policy for 100 times. The evaluation results and an overall rate of satisfying the given LTL specifications are reported for the `frozen-lake-1` experiments in Table 3. The statistics are taken across 10×100 tests, which results in 1000 trials for each hyper-parameter configuration.

4 Conclusions and Extensions

This paper presented LCRL, a new software tool for policy synthesis with RL under LTL and omega-regular specifications. There is a plethora of extensions that we are planning to explore. In the short term, we intend to: (1) directly interface LCRL with automata synthesis tools such as OWL [32]; (2) link LCRL with other model checking tools such as PRISM [33] and Storm [27]; and (3) embed more RL algorithms for policy synthesis, so that we can tackle policy synthesis problems for more challenging environments. In the longer term, we plan to extend LCRL such that (1) it will be able to handle other forms of temporal logic, e.g., signal temporal logic; and (2) it will have a graphical user-interface for the ease of interaction.

References

1. Abate, A., Prandini, M., Lygeros, J., Sastry, S.: Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica* **44**(11), 2724–2734 (2008)

2. Andreas, J., Klein, D., Levine, S.: Modular multitask reinforcement learning with policy sketches. In: ICML, vol. 70, pp. 166–175 (2017)
3. Bellemare, M.G., Naddaf, Y., Veness, J., Bowling, M.: The Arcade learning environment: an evaluation platform for general agents. *JAIR* **47**, 253–279 (2013)
4. Bertsekas, D.P., Shreve, S.: *Stochastic Optimal Control: The Discrete-Time Case*. Athena Scientific (2004)
5. Bertsekas, D.P., Tsitsiklis, J.N.: *Neuro-dynamic Programming*, vol. 1. Athena Scientific (1996)
6. Bozkurt, A.K., Wang, Y., Zavlanos, M.M., Pajic, M.: Control synthesis from linear temporal logic specifications using model-free reinforcement learning. arXiv preprint:1909.07299 (2019)
7. Brockman, G., et al.: OpenAI gym. arXiv preprint:1606.01540 (2016)
8. Cai, M., Hasanbeig, M., Xiao, S., Abate, A., Kan, Z.: Modular deep reinforcement learning for continuous motion planning with temporal logic. *IEEE Robot. Aut. Lett.* **6**(4), 7973–7980 (2021). <https://doi.org/10.1109/LRA.2021.3101544>
9. Clarke Jr, E.M., Grumberg, O., Kroening, D., Peled, D., Veith, H.: *Model Checking*. MIT Press, London (2018)
10. De Giacomo, G., Iocchi, L., Favorito, M., Patrizi, F.: Foundations for restraining bolts: Reinforcement learning with LTLf/LDLf restraining specifications. In: ICAPS, vol. 29, pp. 128–136 (2019)
11. Favorito, M.: Reinforcement learning framework for temporal goals. <https://github.com/whitemech/tempri> (2020)
12. Fulton, N.: Verifiably safe autonomy for cyber-physical systems. Ph.D. thesis, Carnegie Mellon University Pittsburgh (2018)
13. Fulton, N., Platzer, A.: Safe reinforcement learning via formal methods: Toward safe control through proof and learning. In: Proceedings of the AAAI Conference on Artificial Intelligence (2018)
14. Fulton, N., Platzer, A.: Verifiably safe off-model reinforcement learning. In: TACAS, pp. 413–430 (2019)
15. Gordon, G.J.: Stable function approximation in dynamic programming. In: Proceedings of the Twelfth International Conference on Machine Learning, pp. 261–268. Elsevier (1995)
16. Gunter, E.: From natural language to linear temporal logic: Aspects of specifying embedded systems in LTL. In: Workshop on Software Engineering for Embedded Systems: From Requirements to Implementation (2003)
17. Hahn, E.M., Perez, M., Schewe, S., Somenzi, F., Trivedi, A., Wojtczak, D.: Mungojerrie: reinforcement learning of linear-time objectives. arXiv preprint [arXiv:2106.09161](https://arxiv.org/abs/2106.09161) (2021)
18. Hasanbeig, M.: Safe and certified reinforcement learning with logical constraints. Ph.D. thesis, University of Oxford (2020)
19. Hasanbeig, M., Abate, A., Kroening, D.: Logically-constrained reinforcement learning. arXiv preprint:1801.08099 (2018)
20. Hasanbeig, M., Abate, A., Kroening, D.: Certified reinforcement learning with logic guidance. arXiv preprint:1902.00778 (2019)
21. Hasanbeig, M., Abate, A., Kroening, D.: Logically-constrained neural fitted Q-iteration. In: AAMAS. pp. 2012–2014. International Foundation for Autonomous Agents and Multiagent Systems (2019)
22. Hasanbeig, M., Abate, A., Kroening, D.: Cautious reinforcement learning with logical constraints. In: AAMAS. International Foundation for Autonomous Agents and Multiagent Systems (2020)

23. Hasanbeig, M., Kantaros, Y., Abate, A., Kroening, D., Pappas, G.J., Lee, I.: Reinforcement learning for temporal logic control synthesis with probabilistic satisfaction guarantees. In: Proceedings of the 58th Conference on Decision and Control, pp. 5338–5343. IEEE (2019)
24. Hasanbeig, M., Kroening, D., Abate, A.: Deep reinforcement learning with temporal logics. In: Bertrand, N., Jansen, N. (eds.) FORMATS 2020. LNCS, vol. 12288, pp. 1–22. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-57628-8_1
25. Hasanbeig, M., Kroening, D., Abate, A.: Towards verifiable and safe model-free reinforcement learning. In: Proceedings of Workshop on Artificial Intelligence and Formal Verification, Logics, Automata and Synthesis (OVERLAY), pp. 1–10. Italian Association for Artificial Intelligence (2020)
26. Hasanbeig, M., Yogananda Jeppu, N., Abate, A., Melham, T., Kroening, D.: DeepSynth: Program synthesis for automatic task segmentation in deep reinforcement learning. In: AAAI Conference on Artificial Intelligence. Association for the Advancement of Artificial Intelligence (2021)
27. Hensel, C., Junges, S., Katoen, J.-P., Quatmann, T., Volk, M.: The probabilistic model checker STORM. *Int. J. Softw. Tools Technol. Transfer* **22**, 1–22 (2021). <https://doi.org/10.1007/s10009-021-00633-z>
28. Hunt, N., Fulton, N., Magliacane, S., Hoang, N., Das, S., Solar-Lezama, A.: Verifiably safe exploration for end-to-end reinforcement learning. *arXiv preprint arXiv:2007.01223* (2020)
29. Icarte, R.T., Klassen, T., Valenzano, R., McIlraith, S.: Using reward machines for high-level task specification and decomposition in reinforcement learning. In: ICML, pp. 2107–2116 (2018)
30. Jothimurugan, K., Alur, R., Bastani, O.: A composable specification language for reinforcement learning tasks. In: NeurIPS, pp. 13041–13051 (2019)
31. Kini, D., Viswanathan, M.: Optimal translation of LTL to limit deterministic automata. In: Legay, A., Margaria, T. (eds.) TACAS 2017. LNCS, vol. 10206, pp. 113–129. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54580-5_7
32. Křetínský, J., Meggendorfer, T., Sickert, S.: Owl: a library for ω -words, automata, and LTL. In: Lahiri, S.K., Wang, C. (eds.) ATVA 2018. LNCS, vol. 11138, pp. 543–550. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01090-4_34
33. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: a. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 585–591. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22110-1_47
34. Lee, I.S., Lau, H.Y.: Adaptive state space partitioning for reinforcement learning. *Eng. Appl. Artif. Intell.* **17**(6), 577–588 (2004)
35. Newell, R.G., Pizer, W.A.: Discounting the distant future: how much do uncertain rates increase valuations? *J. Environ. Econ. Manag* **46**(1), 52–71 (2003)
36. Nikora, A.P., Balcom, G.: Automated identification of LTL patterns in natural language requirements. In: ISSRE, pp. 185–194. IEEE (2009)
37. Pitis, S.: Rethinking the discount factor in reinforcement learning: a decision theoretic approach. *arXiv preprint:1902.02893* (2019)
38. Sadigh, D., Kim, E.S., Coogan, S., Sastry, S.S., Seshia, S.A.: A learning based approach to control synthesis of Markov decision processes for linear temporal logic specifications. In: CDC, pp. 1091–1096. IEEE (2014)
39. Sickert, S., Esparza, J., Jaax, S., Křetínský, J.: Limit-deterministic Büchi automata for linear temporal logic. In: Chaudhuri, S., Farzan, A. (eds.) CAV 2016. LNCS, vol. 9780, pp. 312–332. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-41540-6_17

40. Sickert, S., Křetínský, J.: MoChiBA: probabilistic LTL model checking using limit-deterministic Büchi automata. In: Artho, C., Legay, A., Peled, D. (eds.) ATVA 2016. LNCS, vol. 9938, pp. 130–137. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46520-3_9
41. Soudjani, S.E.Z., Gevaerts, C., Abate, A.: FAUST²: Formal Abstractions of Uncountable-State STochastic processes. In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 272–286. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46681-0_23
42. Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction, vol. 1. MIT Press, Cambridge (1998)
43. Tassa, Y., et al.: Deepmind control suite. arXiv preprint:1801.00690 (2018)
44. Voronoi, G.: Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Deuxième mémoire. Recherches sur les paralléloèdres primitifs. *Journal für die reine und angewandte Mathematik* **134**, 198–287 (1908)
45. Wei, Q., Guo, X.: Markov decision processes with state-dependent discount factors and unbounded rewards/costs. *Oper. Res. Lett.* **39**(5), 369–374 (2011)
46. Yan, R., Cheng, C.H., Chai, Y.: Formal consistency checking over specifications in natural languages. In: Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition, pp. 1677–1682. EDA Consortium (2015)
47. Yoshida, N., Uchibe, E., Doya, K.: Reinforcement learning with state-dependent discount factor. In: ICDL, pp. 1–6. IEEE (2013)