

LTL2Action: Generalizing LTL Instructions for Multi-Task RL



Pashootan Vaezipoor*



Andrew Li*



Rodrigo Toro Icarte



Sheila McIlraith

University of Toronto and Vector Institute



Motivation

How can humans specify high-level objectives to an RL agent?



Motivation

How can humans specify high-level objectives to an RL agent?

via Reward Function?

Infeasible for humans to program for every possible task.

Motivation

How can humans specify high-level objectives to an RL agent?

via Reward Function?

Infeasible for humans to program for every possible task.

via Natural Language?

Unclear what reward to optimize.


This Work

- Formal language of Linear Temporal Logic (LTL) [*Pnueli, 1977*]

This Work

- Formal language of Linear Temporal Logic (LTL) [*Pnueli, 1977*]
 - Expressive temporally extended goals

eventually (pickup_coal **or** pickup_wood)
and (**always** (**not** on_lava))



This Work

- Formal language of Linear Temporal Logic (LTL) [*Pnueli, 1977*]
 - Expressive temporally extended goals
 - **Automatic mapping** *Instruction* → *Reward*

This Work

- Formal language of Linear Temporal Logic (LTL) [*Pnueli, 1977*]
 - Expressive temporally extended goals
 - **Automatic mapping** *Instruction* → *Reward*
 - Procedural sampling from > 10^{39} unique LTL instructions


This Work

- Formal language of Linear Temporal Logic (LTL) [*Pnueli, 1977*]
 - Expressive temporally extended goals
 - **Automatic mapping** *Instruction* \rightarrow *Reward*
 - Procedural sampling from $> 10^{39}$ unique LTL instructions
- **Theoretical** benefits
 - **non-myopic** composition
 - **non-Markovian** rewards
- **Empirical** benefits
 - **Discrete & Continuous** (MuJoCo) environments

This Work


- Formal language of Linear Temporal Logic (LTL) [*Pnueli, 1977*]
 - Expressive temporally extended goals
 - **Automatic mapping** *Instruction* → *Reward*
 - Procedural sampling from > 10^{39} unique LTL instructions
- **Theoretical** benefits
 - **non-myopic** composition
 - **non-Markovian** rewards
- **Empirical** benefits
 - **Discrete & Continuous** (MuJoCo) environments
 - **Zero-shot generalization** to unseen and larger instructions

What tasks can be expressed in LTL?



Assume agent can detect primitive high-level events:


What tasks can be expressed in LTL?



Assume agent can detect primitive high-level events:

- `pickup_coal`
- `pickup_wood`

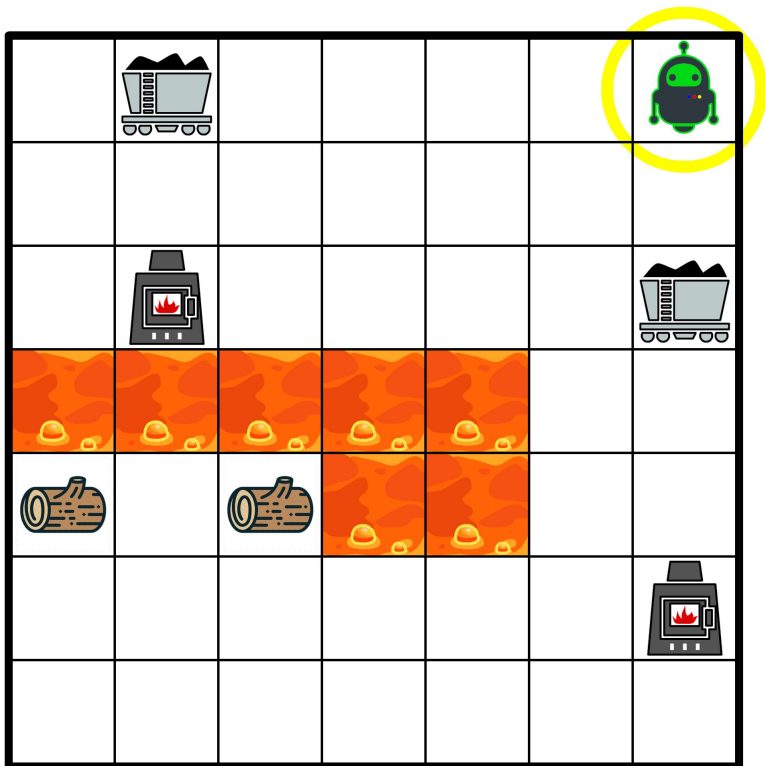
What tasks can be expressed in LTL?



Assume agent can detect primitive high-level events:

- pickup_coal
- pickup_wood
- use_furnace

What tasks can be expressed in LTL?



Assume agent can detect primitive high-level events:

- `pickup_coal`
- `pickup_wood`
- `use_furnace`
- `on_lava`

Tasks in LTL

Task Type	LTL Formula	English
Single Goal	eventually pickup_coal	“Get coal”
Ordered Goals	eventually (pickup_coal and (eventually use_furnace))	“Get coal, then use the furnace”
Unordered Goals	(eventually pickup_coal) and (eventually pickup_wood)	“Get coal and get wood, in any order”
Disjunctive Goals	(eventually pickup_coal) or (eventually pickup_wood)	“Get coal or get wood”
Safety	(eventually pickup_wood) and (always (not on_lava))	“Get wood while avoiding lava”

Tasks in LTL

Task Type	LTL Formula	English
Single Goal	eventually pickup_coal	“Get coal”
Ordered Goals	eventually (pickup_coal and (eventually use_furnace))	“Get coal, then use the furnace”
Unordered Goals	(eventually pickup_coal) and (eventually pickup_wood)	“Get coal and get wood, in any order”
Disjunctive Goals	(eventually pickup_coal) or (eventually pickup_wood)	“Get coal or get wood”
Safety	(eventually pickup_wood) and (always (not on_lava))	“Get wood while avoiding lava”

Tasks in LTL

Task Type	LTL Formula	English
Single Goal	eventually pickup_coal	“Get coal”
Ordered Goals	eventually (pickup_coal and (eventually use_furnace))	“Get coal, then use the furnace”
Unordered Goals	(eventually pickup_coal) and (eventually pickup_wood)	“Get coal and get wood, in any order”
Disjunctive Goals	(eventually pickup_coal) or (eventually pickup_wood)	“Get coal or get wood”
Safety	(eventually pickup_wood) and (always (not on_lava))	“Get wood while avoiding lava”

LTl Instruction \rightarrow Reward

$$R = \begin{cases} 1 & \text{if instruction is **satisfied**} \\ -1 & \text{if instruction is **falsified**} \\ 0 & \text{otherwise} \end{cases}$$


Task Decomposition

Task Decomposition

LTL Progression [*Bacchus & Kabanza, 2000*]

- Formally defined for all LTL formulas
- **Simplify instructions** once parts of the task are solved


LTL Progression — Example



“Get coal or wood, then use the furnace.”

```
eventually ((pickup_coal or pickup wood)
and (eventually use_furnace))
```

LTL Progression — Example



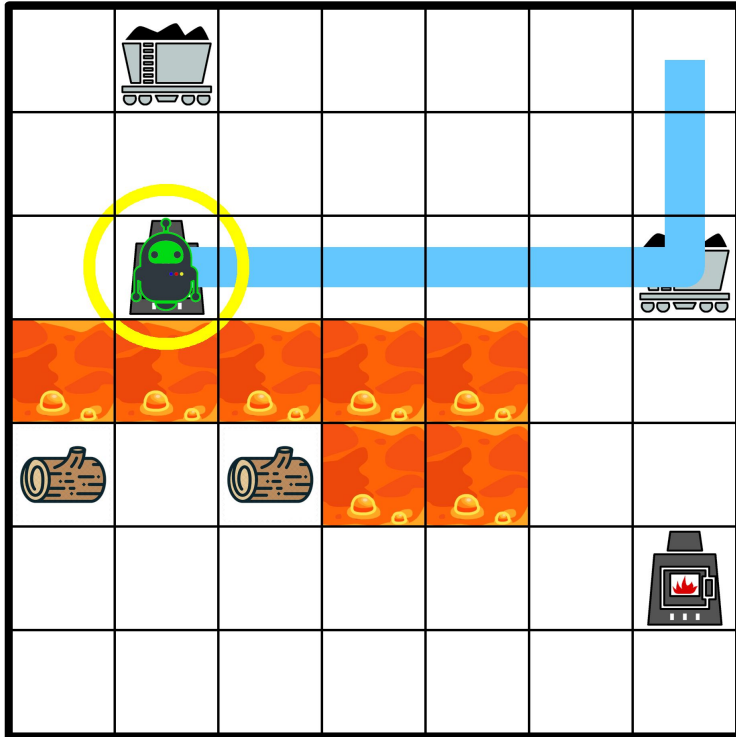
“Get coal or wood, then use the furnace.”

```
eventually ((pickup_coal or pickup wood)
and (eventually use_furnace))
```

“Use the furnace.”

```
eventually use_furnace
```

LTL Progression — Example



“Get coal or wood, then use the furnace.”

```
eventually ((pickup_coal or pickup wood)
and (eventually use_furnace))
```

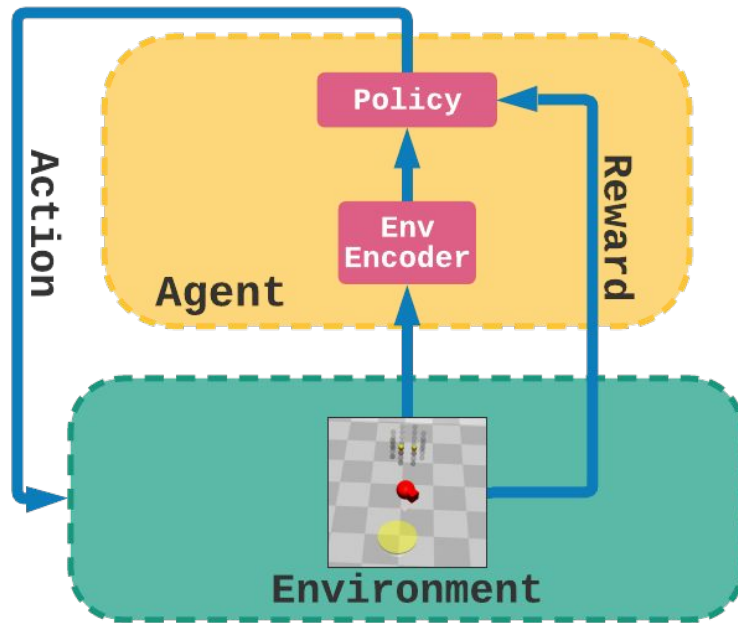
“Use the furnace.”

```
eventually use_furnace
```

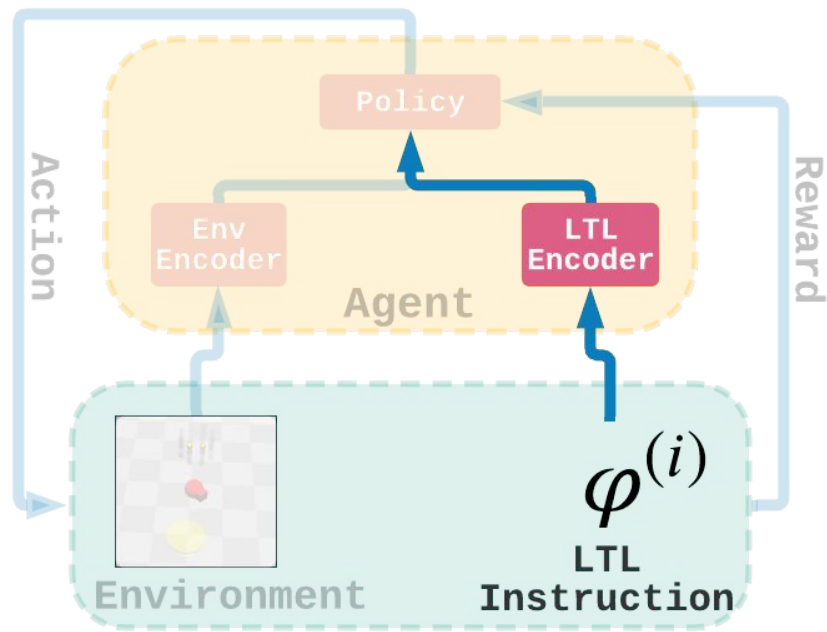
Complete

True

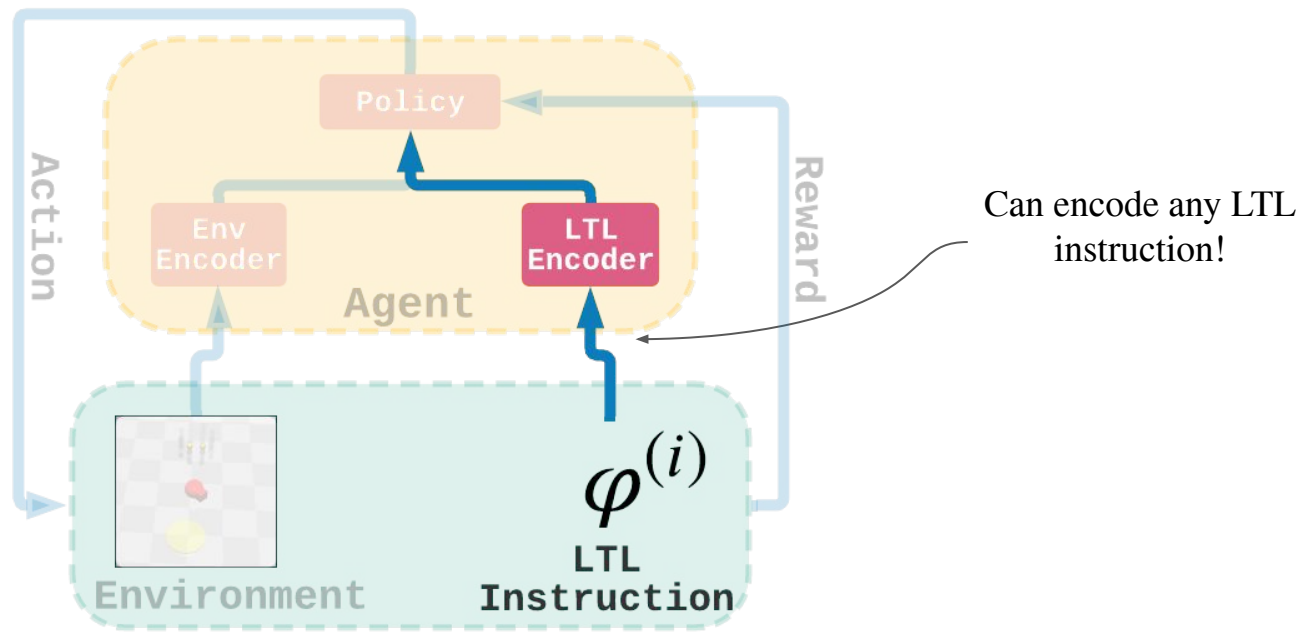
Our Approach



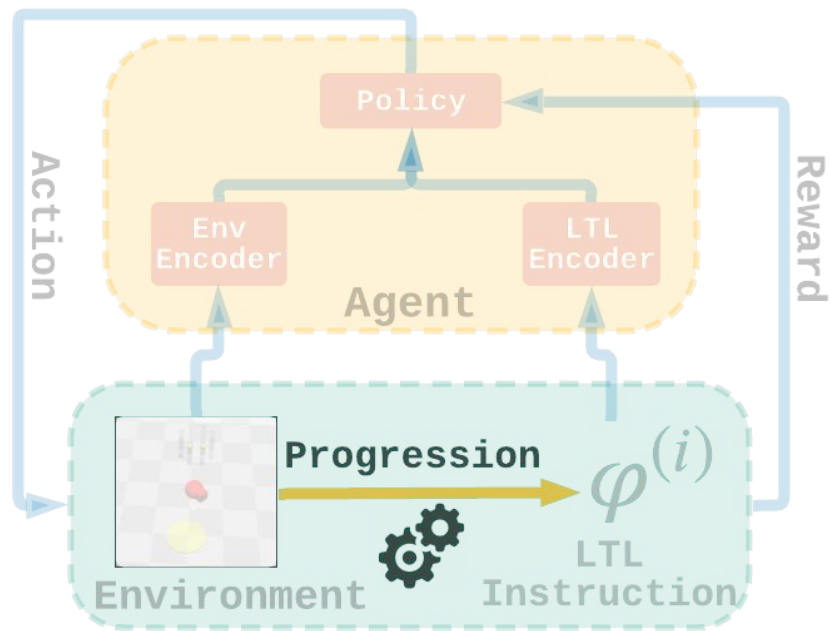
Our Approach



Our Approach



Our Approach



Theorems

Theorems

- ✓ Retains optimal convergence guarantees

Theorems

- ✓ Retains optimal convergence guarantees
- ✓ Markov assumptions hold

Results

Avoidance Tasks

(> 970 million possible tasks)

formula := **sequence** \wedge **formula** | **sequence**
sequence := \neg prop(prop \wedge **sequence**) | \neg propprop


Partially-Ordered Tasks

(> 10^{39} possible tasks)

formula := **sequence** \wedge **formula** | **sequence**
sequence := \diamond (**term** \wedge **sequence**) | \diamond **term**
term := prop | prop \vee prop


Results

Gridworld (Discrete)




Results

Gridworld (Discrete)

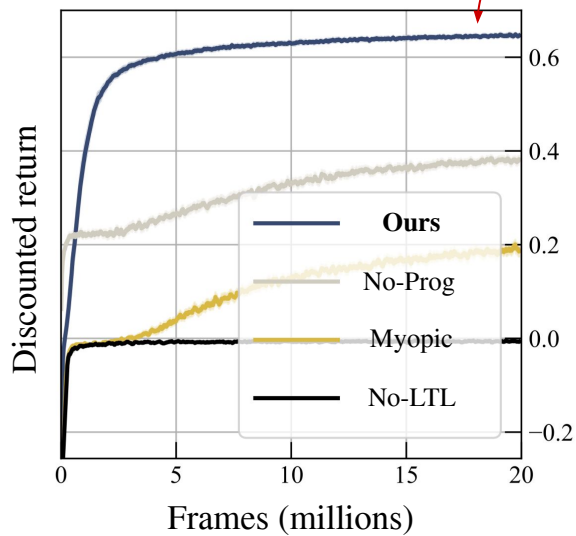



MuJoCo (Continuous)





Results

Gridworld (Discrete)

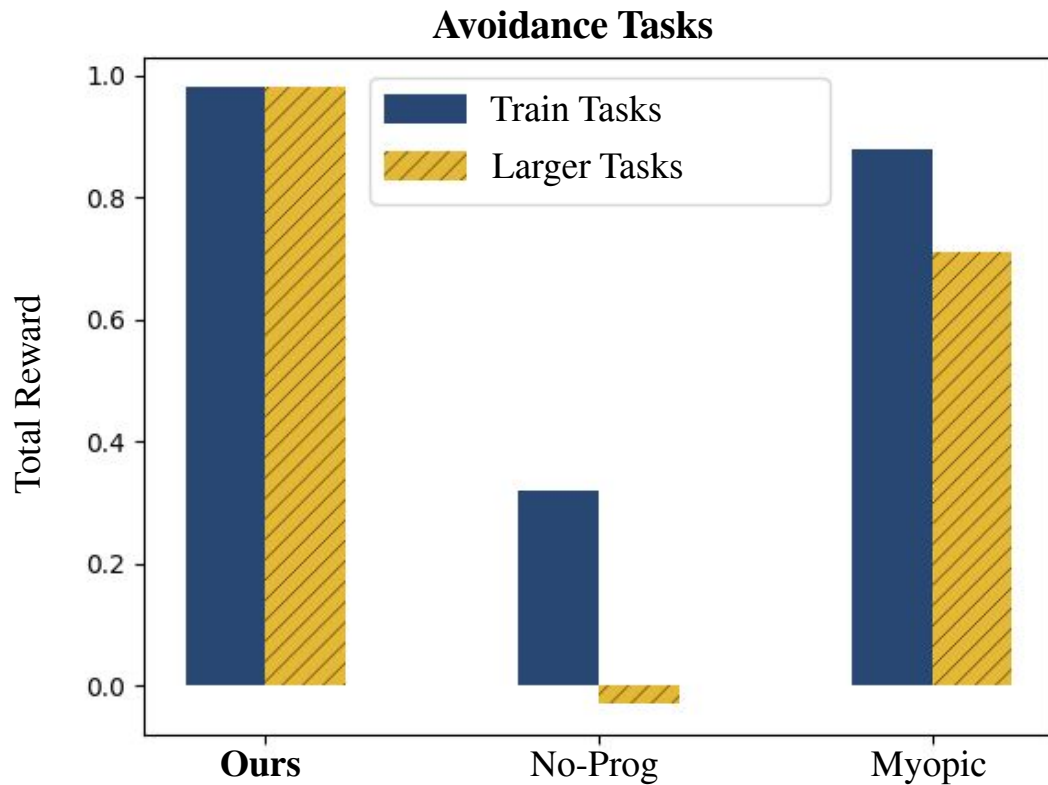


MuJoCo (Continuous)



Ours

Zero-Shot Generalization



Other topics ...

Other topics ...

- A environment-agnostic **pre-training** scheme to learn LTL semantics

Other topics ...

- A environment-agnostic **pre-training** scheme to learn LTL semantics
- What neural architecture best encodes LTL instructions?

Other topics ...

- A environment-agnostic **pre-training** scheme to learn LTL semantics
- What neural architecture best encodes LTL instructions?
- Procedural generation of meaningful LTL instructions

Other topics ...

- A environment-agnostic **pre-training** scheme to learn LTL semantics
- What neural architecture best encodes LTL instructions?
- Procedural generation of meaningful LTL instructions

Come to our poster!

Code is available at:

<https://github.com/LTL2Action/LTL2Action>