

Motivation

- Learning tabula rasa is hard and requires an unreasonable number of samples.
- When teaching humans a new task, it is natural to lend advice to speed-up learning. This advice typically contains partial information about goal states, rewards, best actions, relevant intermediate tasks, etc.
- We sought to formalize advice-giving so that humans can easily provide meaningful guidance to RL agents.

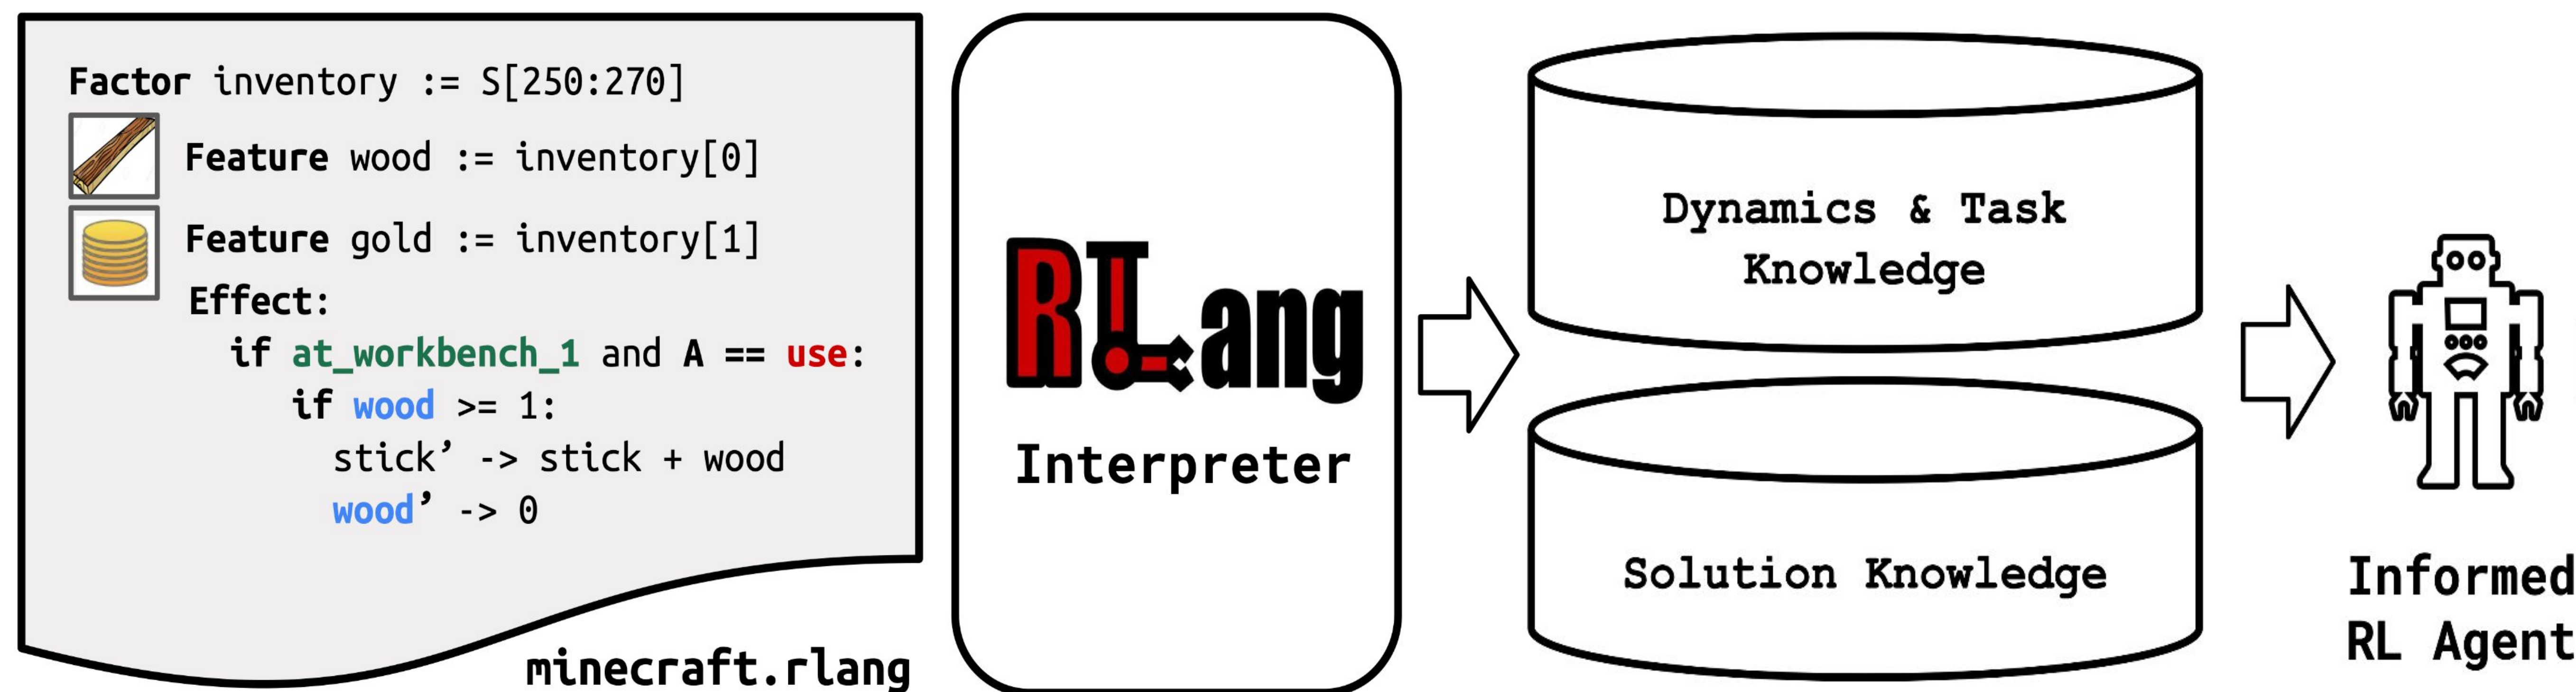
What can we express with RLang?

- RLang provides a formal, unambiguous, and unifying framework for expressing task-specific information.
- RLang provides syntax to specify information about an MDP's:
 - ◆ **Model:** Rewards and Dynamics,
 - ◆ **Solution:** Policy Hints and Policy Priors,
 - ◆ **Abstractions and Features:** Subpolicies (as Options) and State Features.

What's next?

- RLang unifies the varied features of other DSLs previously proposed; thus, it can more effectively be used as a universal store of symbol-oriented knowledge for RL agents in natural language grounding research;
- RLang enables research on **neuro-symbolic RL**; agents can reason and behave using both symbolic *and* latent knowledge/policies.
- RLang enables research in general **informed RL** methods.

RLang proposes a unified system for providing RL agents with task-specific, grounded advice that helps them learn faster than tabula rasa




Demonstrations

2D-Minecraft: Hierarchical Structure

```

1 Option go_to_workshop_0:
2   init(any):
3     Execute go_to_workshop_0_learnable_policy
4   until(at_workshop_0)
5 Option go_to_workshop_1:
6   init(any):
7     Execute go_to_workshop_1_learnable_policy
8   until(at_workshop_1)
9 Option get_wood:
10  init(there_is_wood):
11    Execute get_wood_learnable_policy
12  until(delta_wood >= 1)
13 Option build_plank:
14  init(wood >= 1 and at_workshop_1):
15    Execute use
16  until(delta_plank >= 1)
17 Option build_stick:
18  init(wood >= 1 and at_workshop_1):
19    Execute use
20  until(delta_stick >= 1)
21 Option build_ladder:
22  init(stick >= 1 and plank >= 1)
23    Execute use
24  until(delta_ladder >= 1)
    
```



Lunarlander: Policy Prior

```

1 Policy land:
2   if (left_leg_in_contact == 1.0) or (right_leg_in_contact == 1.0)
3     if (velocity_y/2 * -1.0) > 0.05:
4       Execute main_engine
5     else:
6       Execute do_nothing
7   elif remaining_hover > remaining_angle and remaining_hover > -1
8   * remaining_angle and remaining_hover > 0.05:
9     Execute main_engine
10  elif remaining_angle < -0.05:
11    Execute right_thruster
12  elif remaining_angle > 0.05:
13    Execute left_thruster
14  else:
15    Execute do_nothing
    
```

