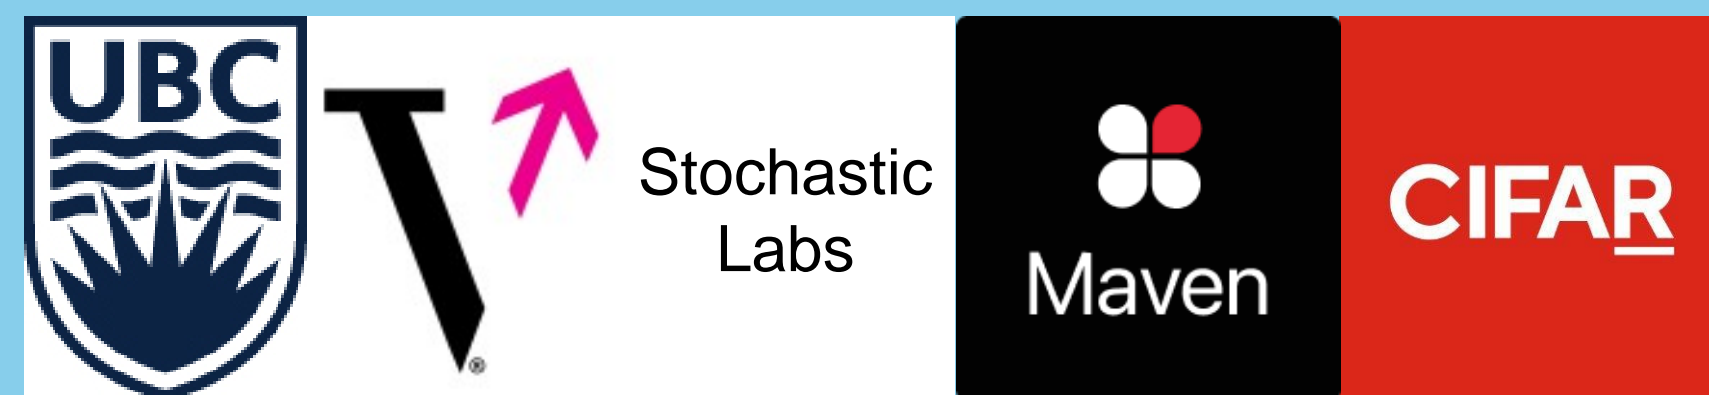


# OMNI

## Open-endedness via Models of human Notions of Interestingness

Jenny Zhang<sup>1,2</sup>, Joel Lehman<sup>3</sup>, Kenneth Stanley<sup>4</sup>, Jeff Clune<sup>1,2,5</sup>

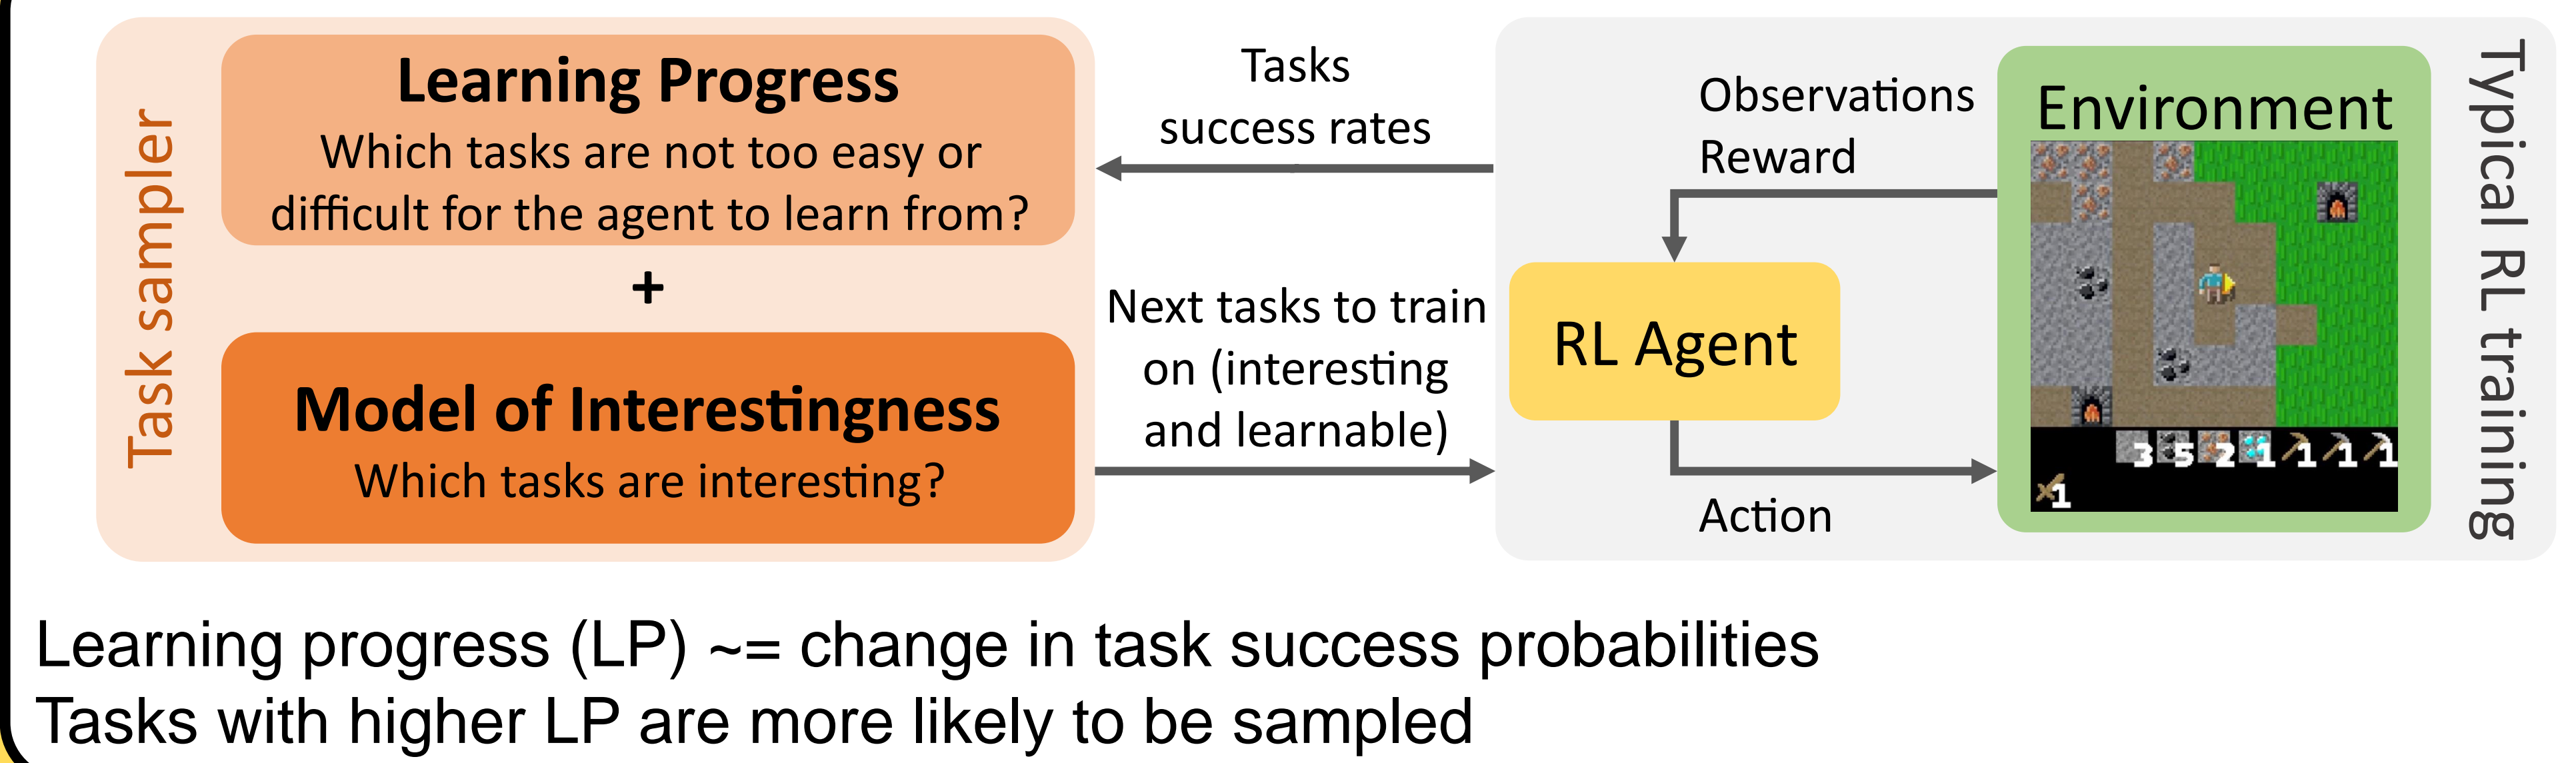
<sup>1</sup>University of British Columbia <sup>2</sup>Vector Institute <sup>3</sup>Stochastic Labs <sup>4</sup>Maven <sup>5</sup>Canada CIFAR AI Chair



### Motivation

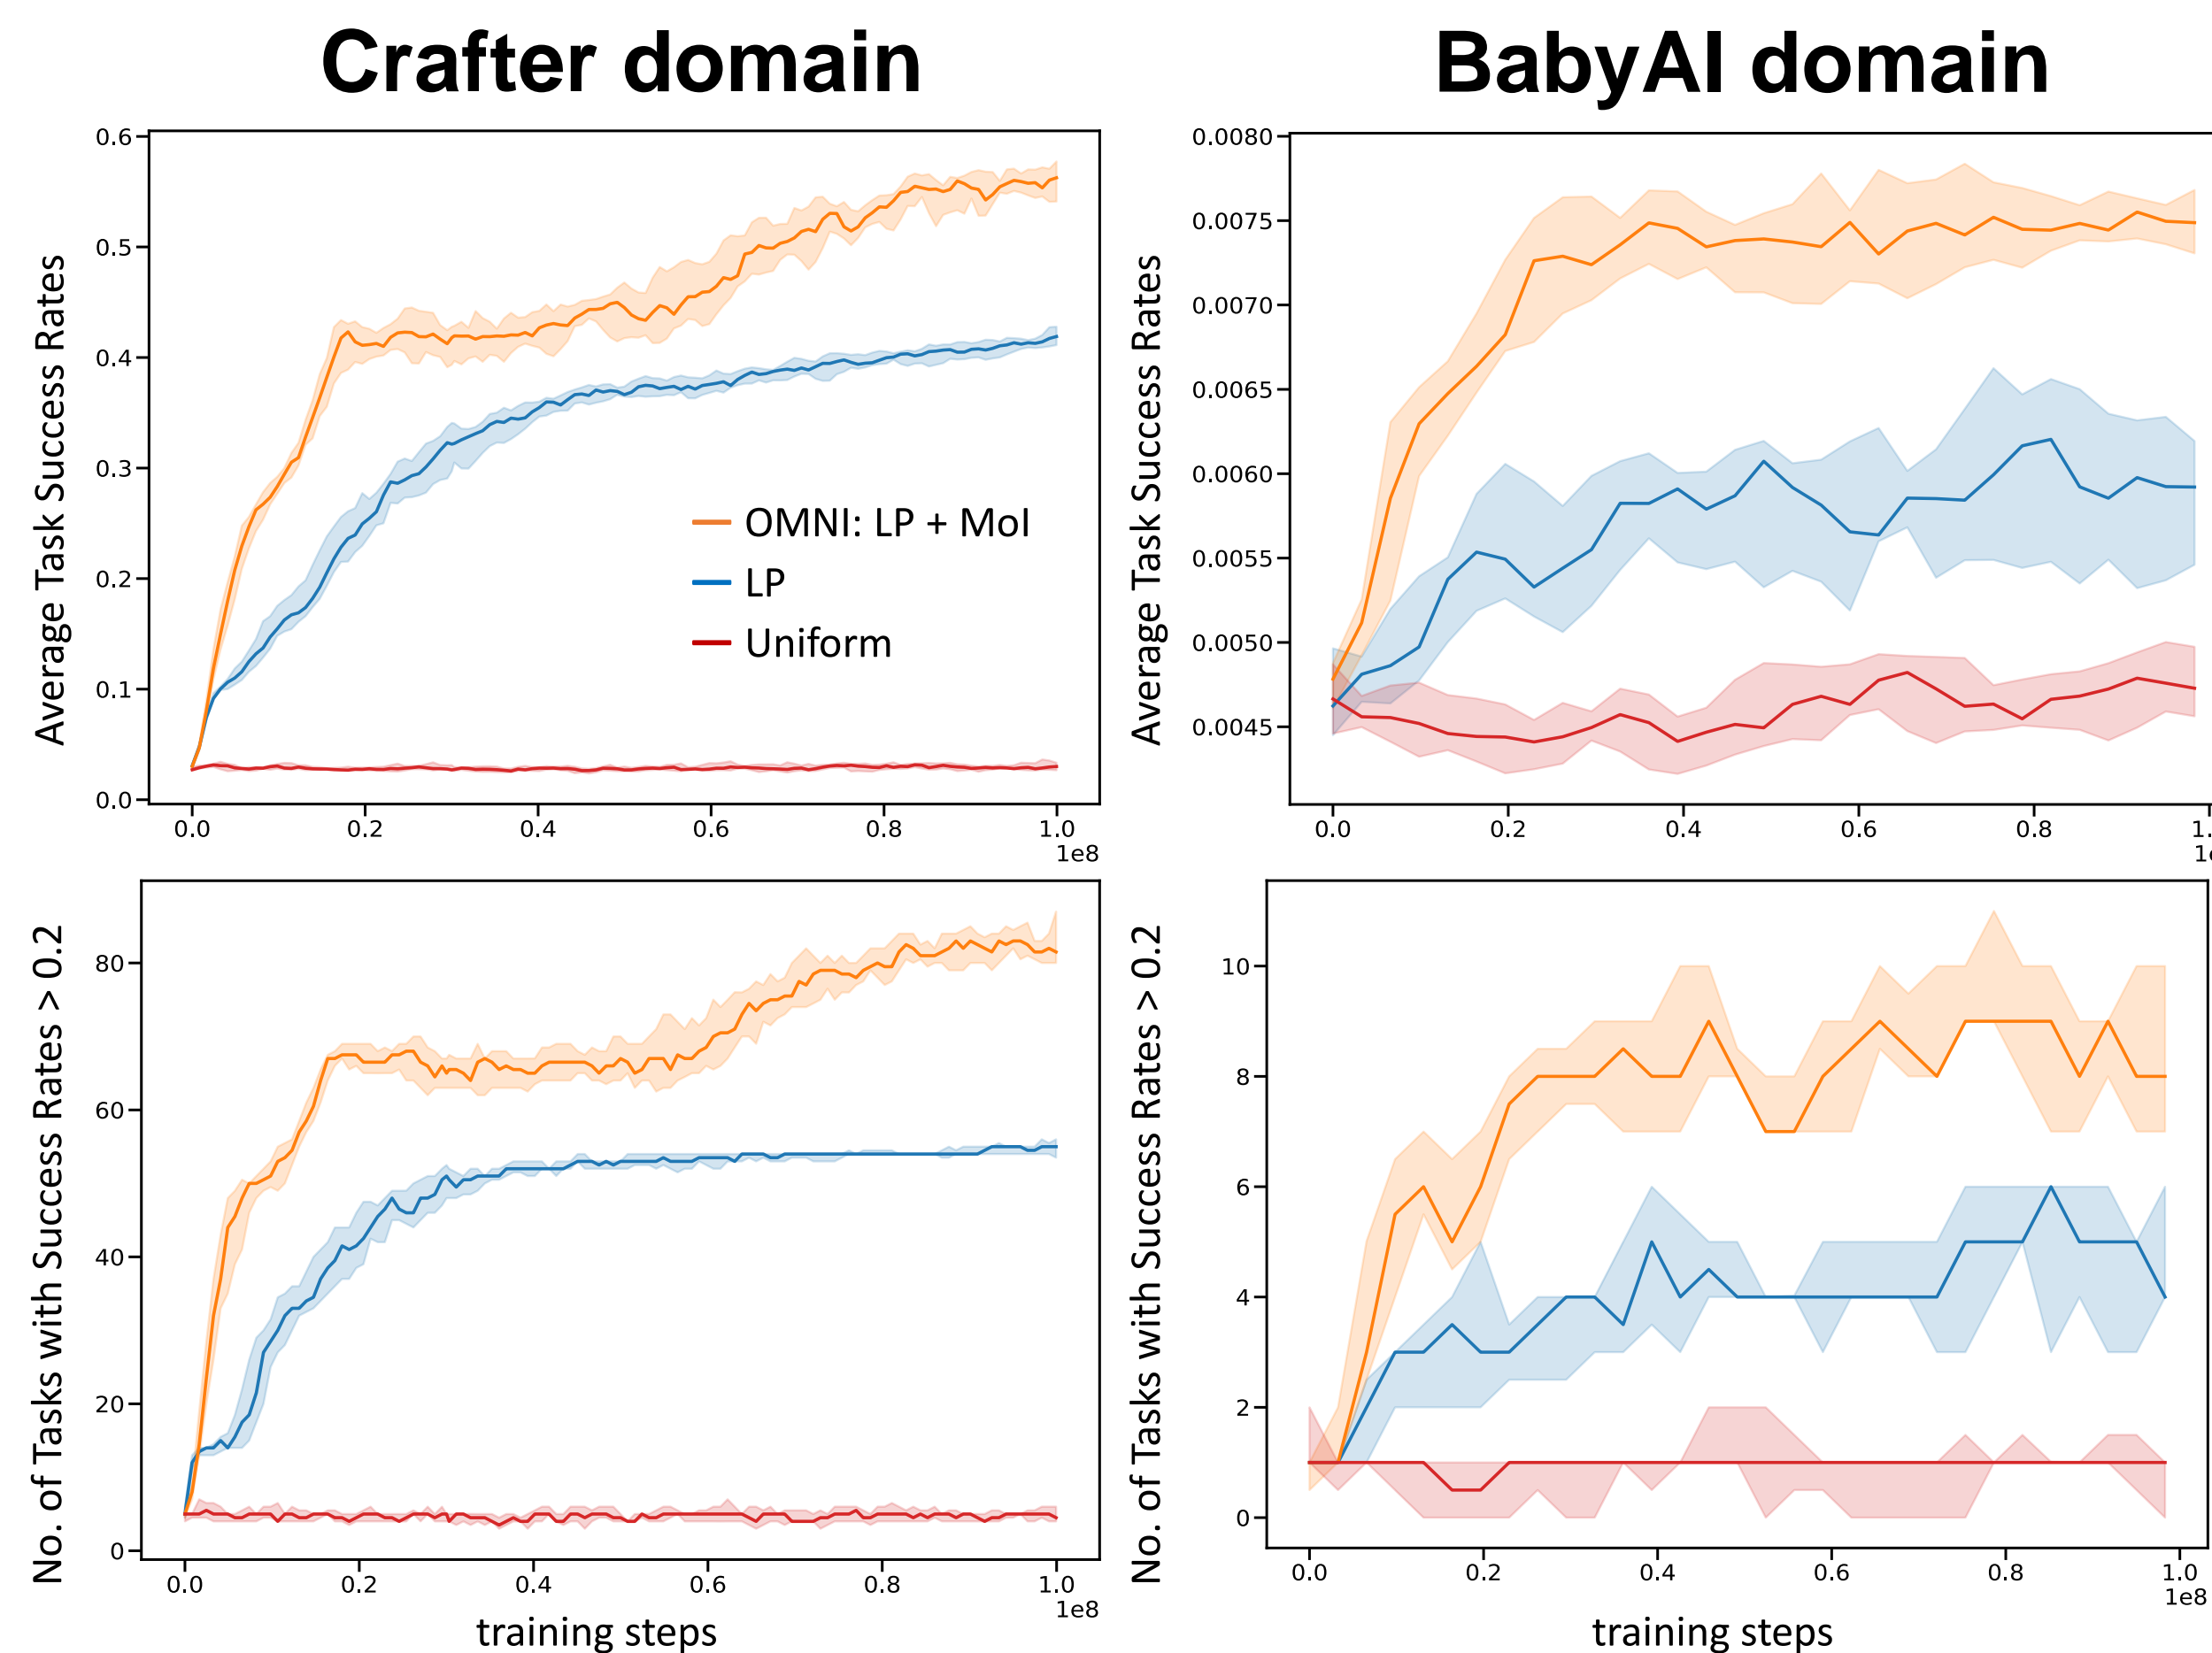
- **Open-ended** algorithms seek to learn new, interesting, diverse tasks **forever**
- Requires a truly **vast** space of possible tasks
- But most tasks are not **interesting** or **learnable**
- Past attempts to quantify new/ interesting/ novel produce pathologies
- **OMNI solves these problems with Foundation Models, which intriguingly understand what is interesting and learnable**

### Methods

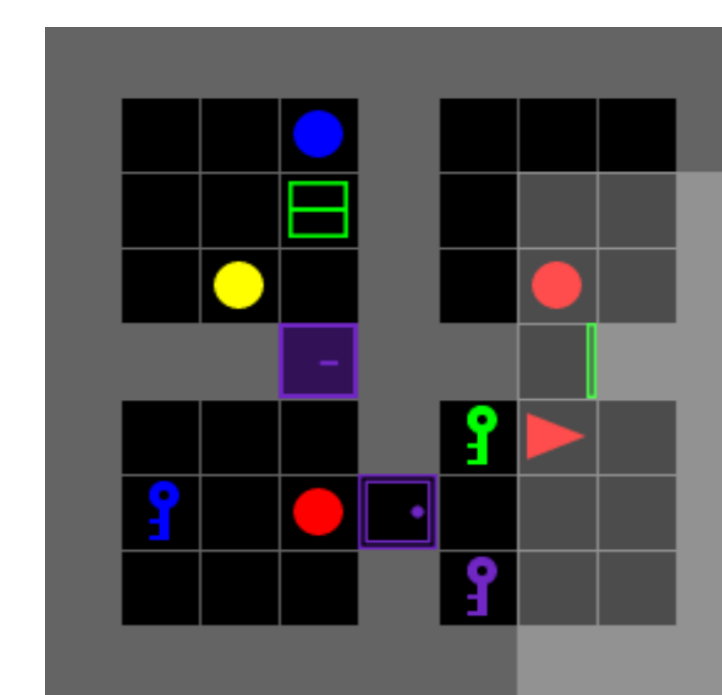


### Experiments (Predefined Task Space)

GPT-3 / GPT-4 predicts whether the tasks are interesting, conditioned on the agent's current skills



- 2D Minecraft
- 15 interesting tasks
- 90 boring tasks – repetitive tasks
- 1023 impossible tasks

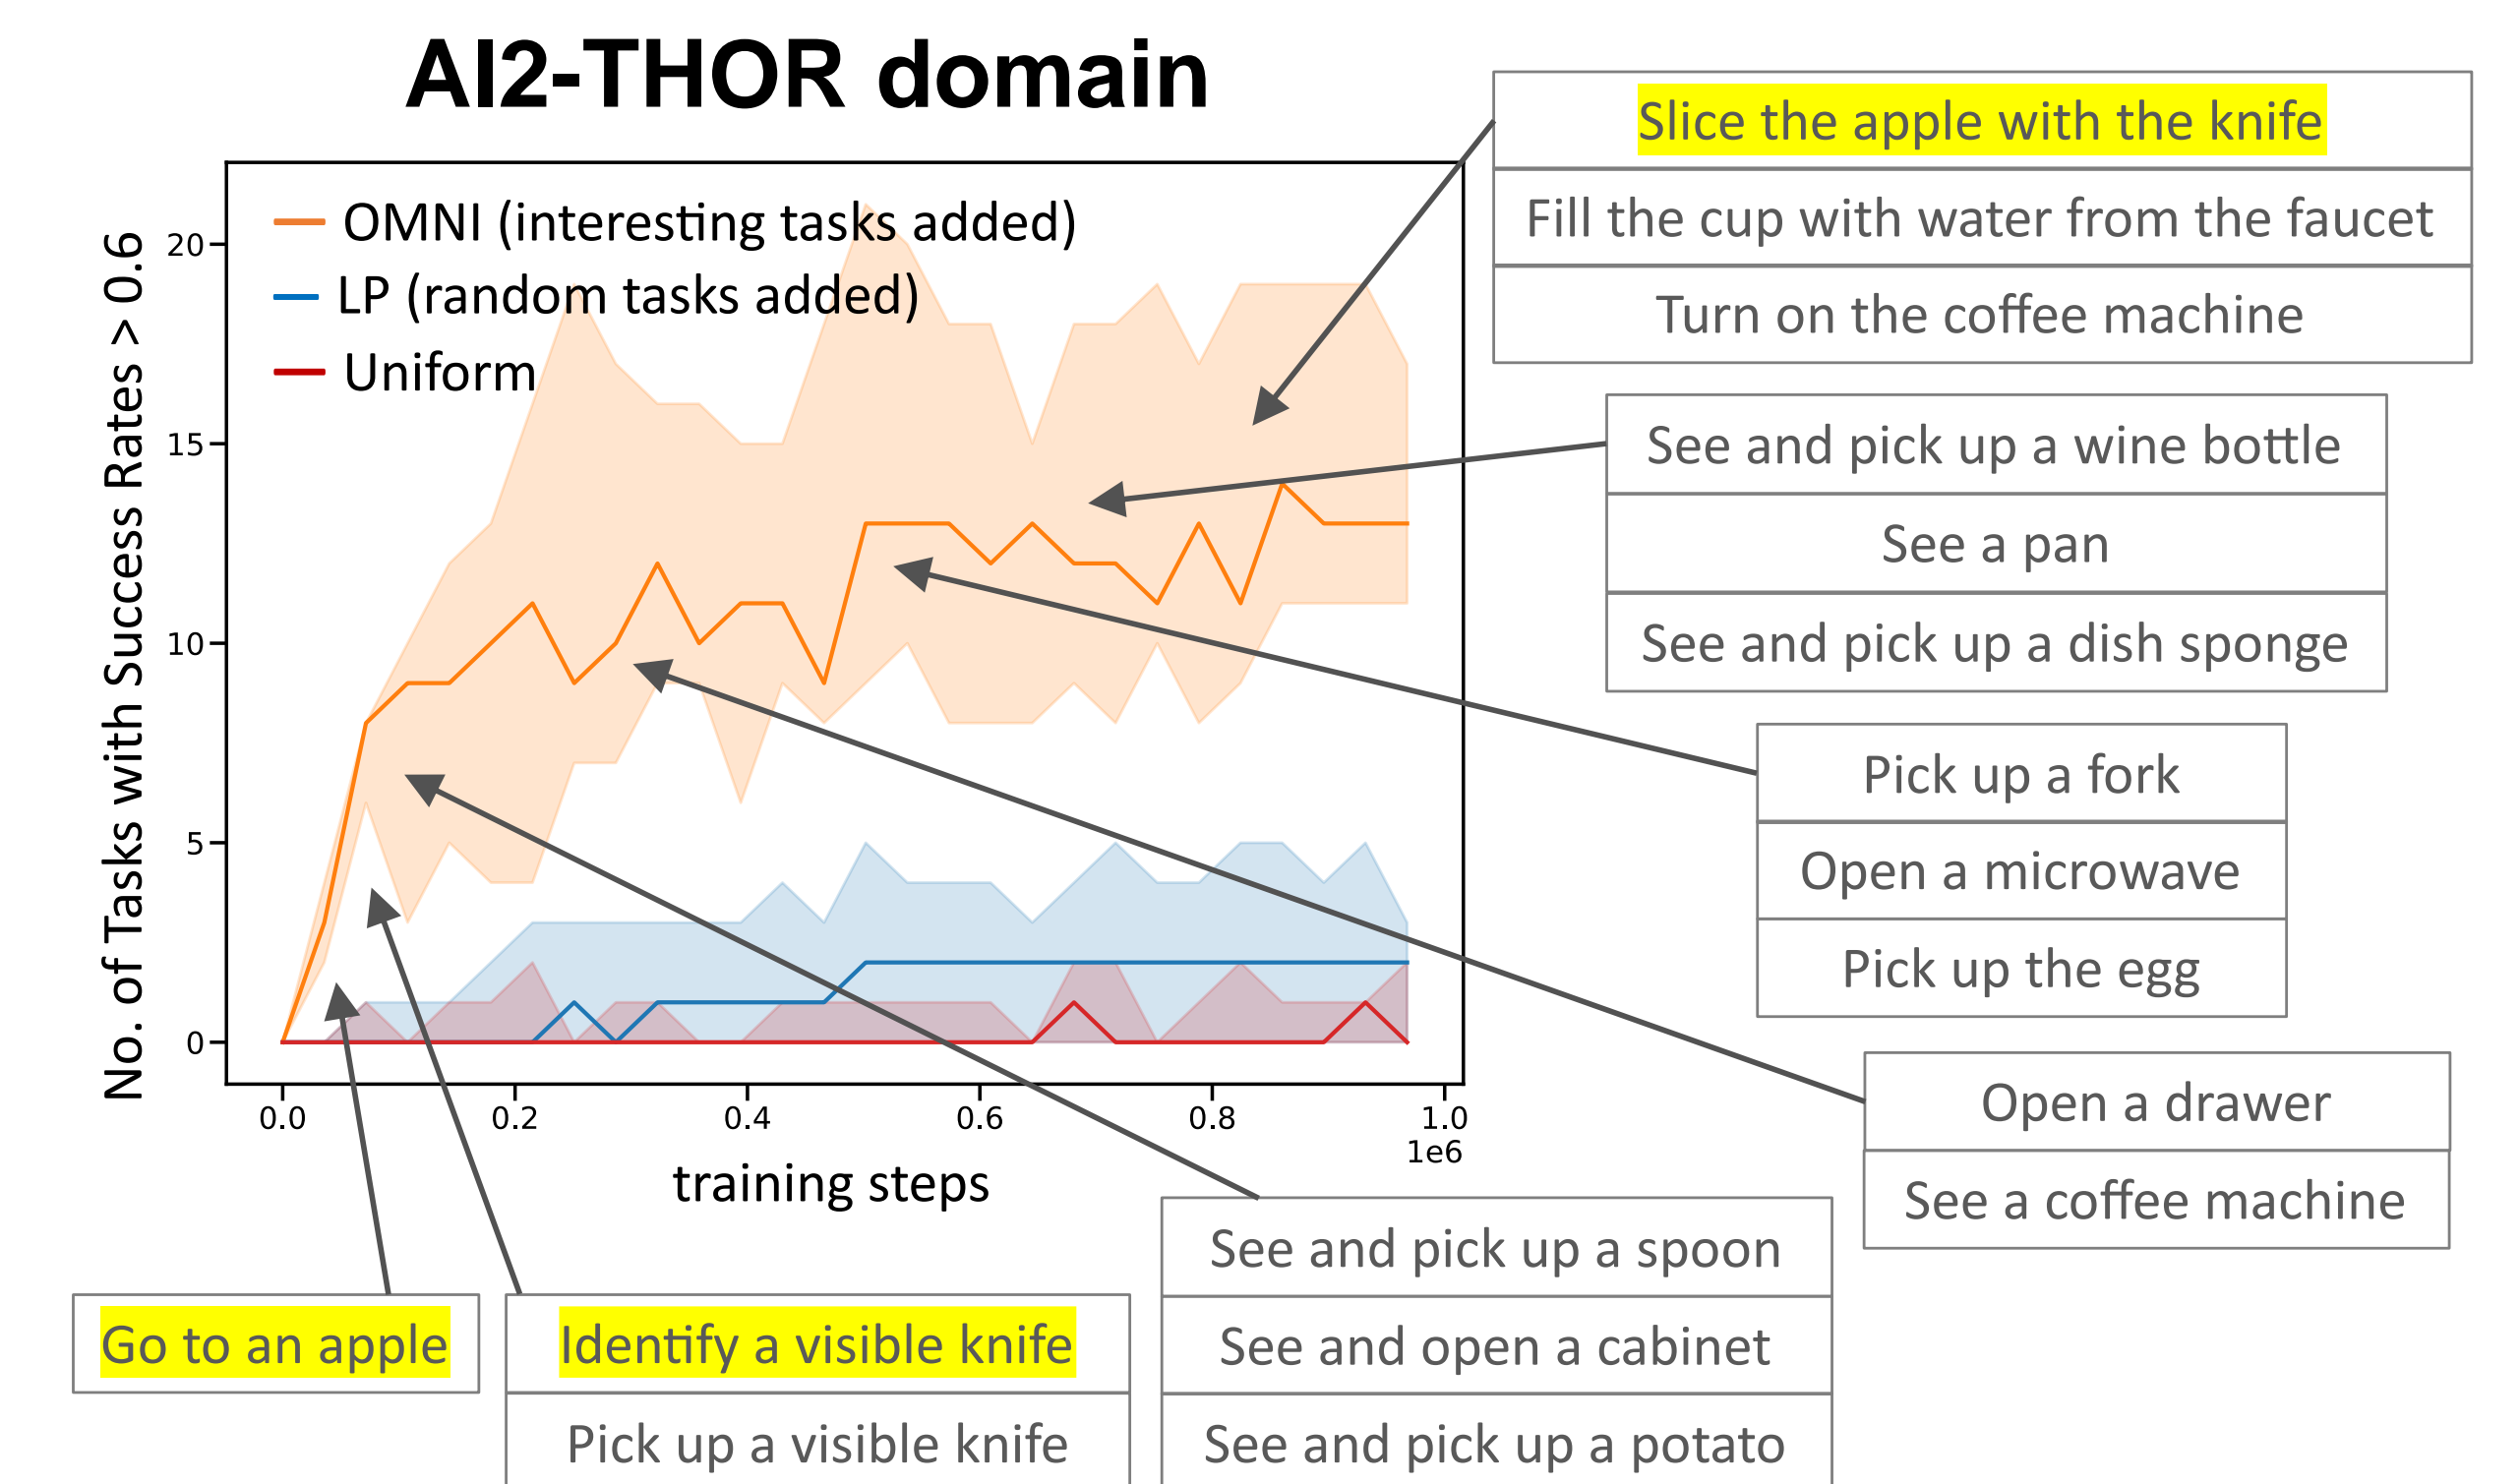


- 1364 tasks
- Sequence of up to 5 instructions
- Each instruction – GoTo, PickUp, OpenDoor, PutNextTo

### Experiments (Infinite Task Space)

GPT-4 suggests the next learnable and interesting tasks for the agent to learn on, conditioned on the agent's current skills

- Use code as a *Universal Reward Function* to identify if any task proposed has been successfully completed



- Embodied kitchen environment
- All tasks specifiable in natural language
- No restrictions on task space

#### Input Prompt Example

**Encourage suggestions of interestingly different tasks**  
The ultimate goal is to learn as many interestingly different tasks as possible ...

**Environment description**  
The agent is in a kitchen ...

**Tasks the agent currently does well:**

- Go to an apple
- Go to the knife
- See a potato

**Tasks the agent cannot do yet:**

- Go to the bread
- Pick up the apple
- Open the fridge
- See a plate

#### Output Example

**Reasoning:**  
The next tasks should involve simple interactions with objects that agent can already recognize ...

**Next tasks in natural language:**

1. Pick up the knife
2. Slice the potato with the knife
3. Recognize and locate the coffee machine

**Next tasks as sequence of environment states:**

1. [[obj\_attributes("Knife", "visible": True)], [obj\_attributes("Knife", "isPickedUp": True)]]
2. [[obj\_attributes("Potato", "visible": True)], [obj\_attributes("Knife", "visible": True)], [obj\_attributes("Potato", "isSliced": True)]]
3. [[obj\_attributes("CoffeeMachine", "visible": True)]]

Part of the agent's observations

Code as reward functions

### Discussion & Future Work

- OMNI improves as FMs improve
- Multi-modal Model of Interestingness – Visual inputs?
- Scale – what happens with more compute?
- Can we train across domains?