

Originally presented at NeurIPS 2022

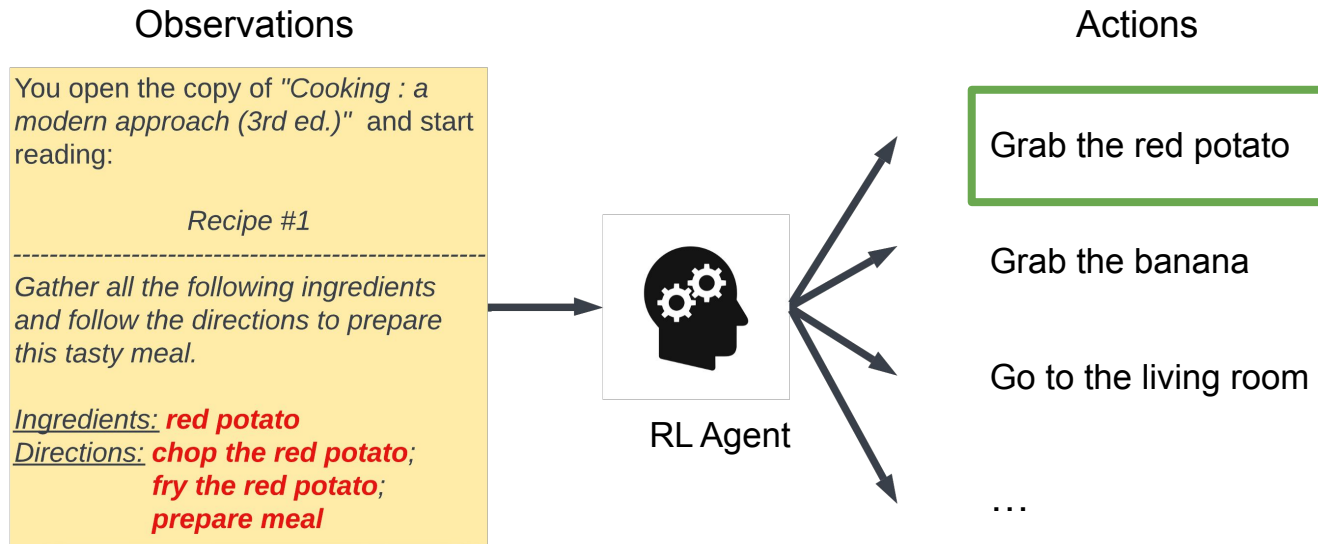
Learning to Follow Instructions in Text-Based Games

Mathieu Tuli - **Andrew C. Li** - Pashootan Vaezipoor - Toryn Q. Klassen
Scott Sanner - Sheila McIlraith

University of Toronto & Vector Institute

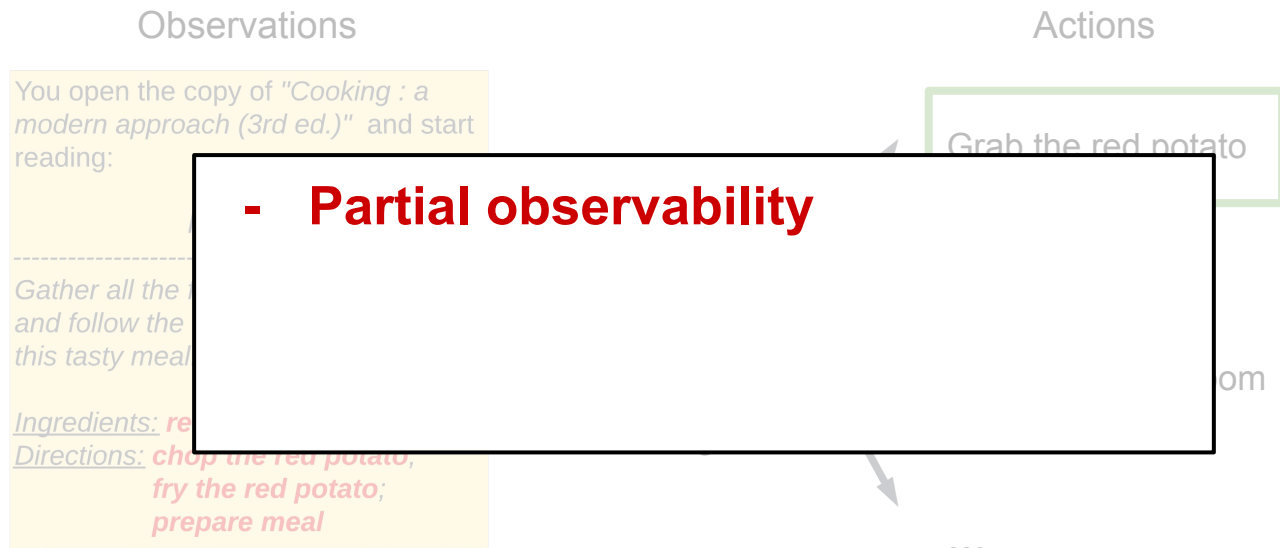


Reinforcement Learning for Text-Based Games



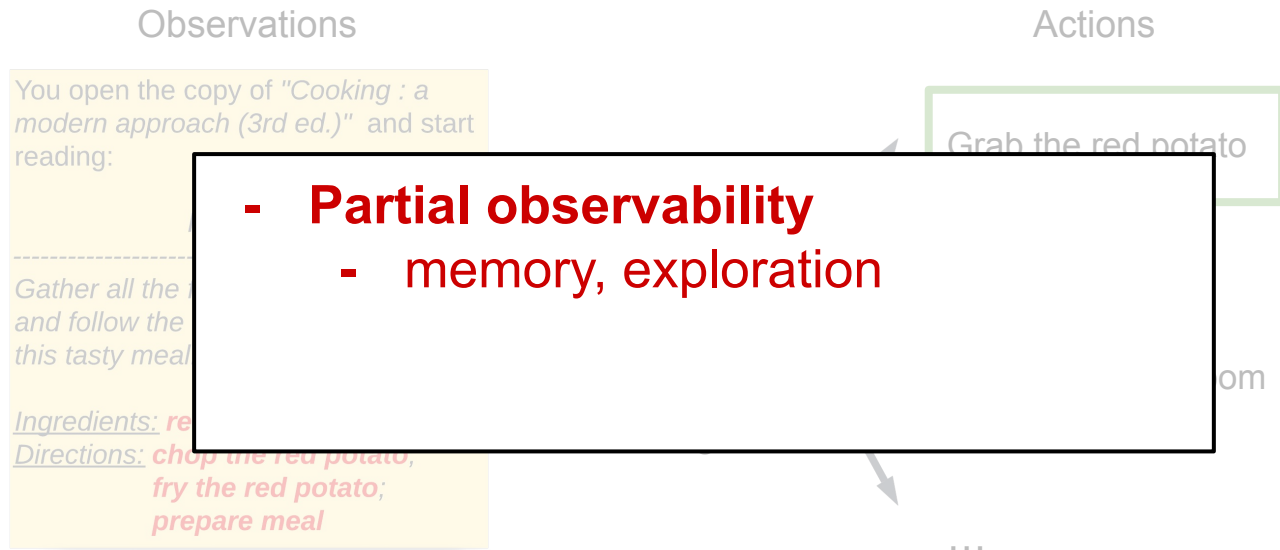
Player observes and interacts with the world through natural language.

Reinforcement Learning for Text-Based Games



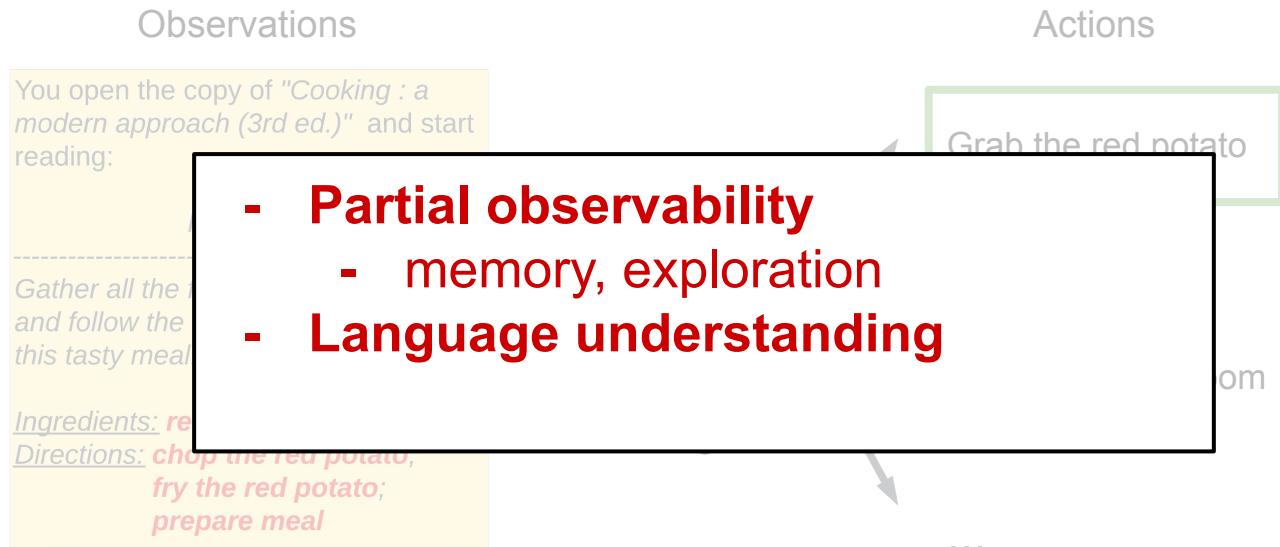
Player observes and interacts with the world through natural language.

Reinforcement Learning for Text-Based Games



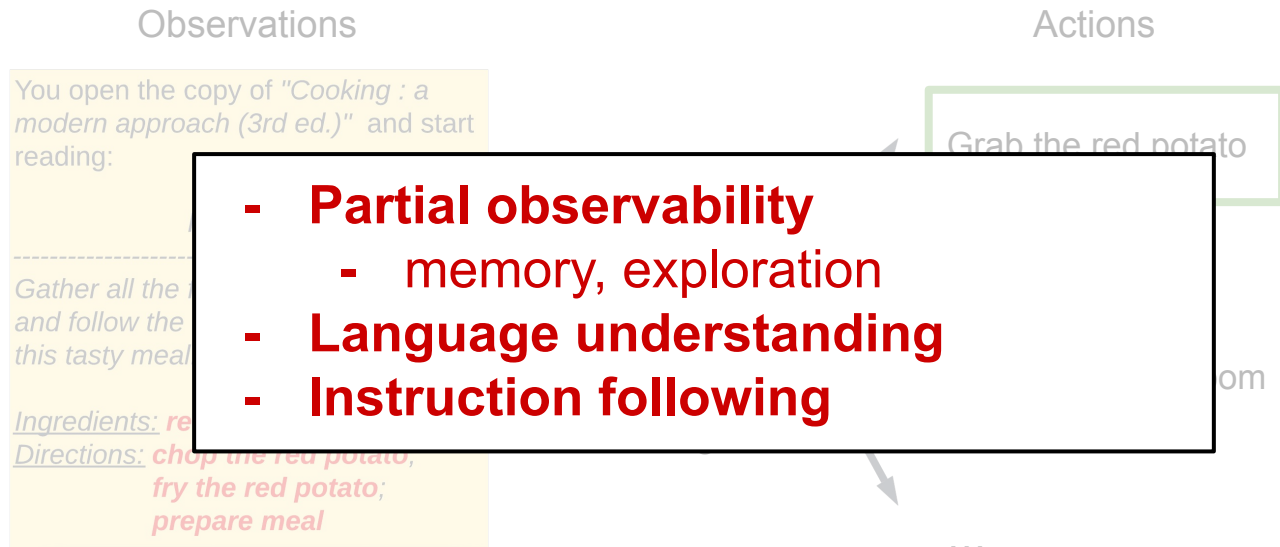
Player observes and interacts with the world through natural language.

Reinforcement Learning for Text-Based Games



Player observes and interacts with the world through natural language.

Reinforcement Learning for Text-Based Games



Player observes and interacts with the world through natural language.

Takeaways

- SoTA TextWorld RL agents **are impervious to instructions**

Takeaways

- SoTA TextWorld RL agents **are impervious to instructions**
- Translate natural language instructions into LTL
 - compositional syntax and semantics
 - monitor progress towards instruction completion

Takeaways

- SoTA TextWorld RL agents **are impervious to instructions**
- Translate natural language instructions into LTL
 - compositional syntax and semantics
 - monitor progress towards instruction completion
- Experiments with 500+ games show superior performance using LTL

The Cooking Domain of TextWorld^[1]

1. Find Cookbook

You open the copy of "Cooking : a modern approach (3rd ed.)" and start reading:

Recipe #1

Gather all the following ingredients and follow the directions to prepare this tasty meal.

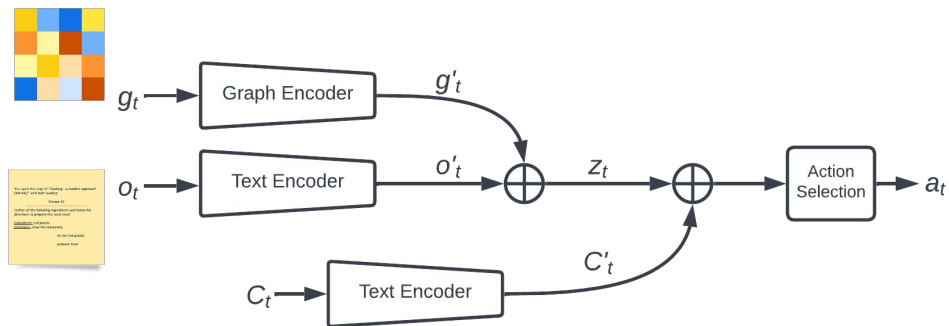
Ingredients: **red potato**

Directions: **chop the red potato;**
fry the red potato;
prepare meal


2. Prepare Recipe

- Grab the red potato
- Chop the red potato
- Fry the red potato
- Prepare meal

GATA

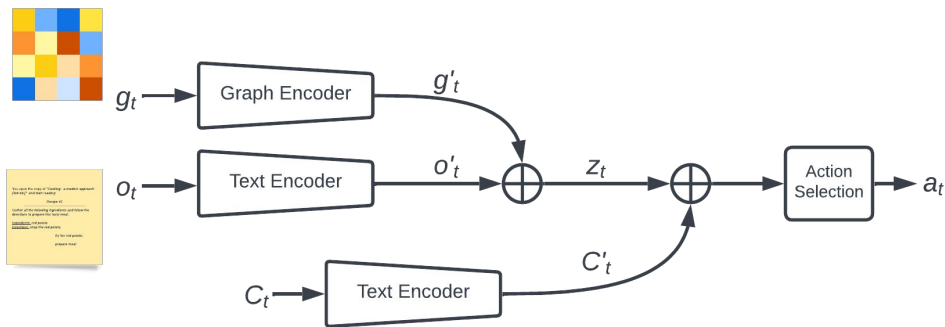


GATA



RL agent augmented with dynamic long-term memory.

GATA



Captures relations about objects:

- **player-has-apple**
- **player-at-kitchen**
- **apple-needs-chop**

RL agent augmented with dynamic long-term memory.

GATA Does Not Follow Instructions


- GATA often doesn't read the cookbook

GATA Does Not Follow Instructions

- GATA often doesn't read the cookbook
- GATA blindly grabs and prepares ingredients, without completing the task

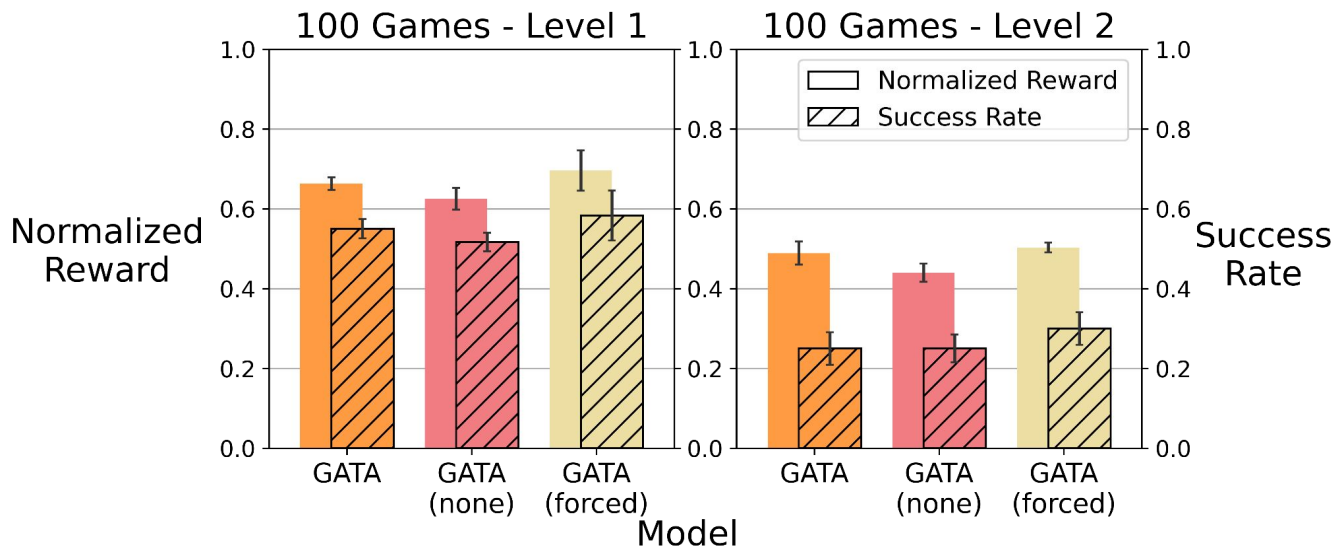
GATA Does Not Follow Instructions

- GATA often doesn't read the cookbook
- GATA blindly grabs and prepares ingredients, without completing the task
- The presence or absence of instructions does not change performance



GATA Does Not Follow Instructions

GATA is unable to understand and follow complex instructions in natural language.



1. Natural Language to LTL

You open the copy of "Cooking : a modern approach (3rd ed.)" and start reading:

Recipe #1

Gather all the following ingredients and follow the directions to prepare this tasty meal.

Ingredients: red potato

Directions: chop the red potato;
fry the red potato;
prepare meal

Natural Language Observation

φ_t

(EVENTUALLY red-potato-in-player)
 \wedge (EVENTUALLY red-potato-is-chopped)
 \wedge (EVENTUALLY red-potato-is-fried)
 \wedge (EVENTUALLY meal-in-player)
 \wedge (EVENTUALLY meal-is-consumed)

LTL Instructions

1. Natural Language to LTL

Automate with GPT-3!

You open a copy of "Cooking: a modern approach (3rd ed.)" and start reading:

Recipe #1

Gather all the following ingredients and follow the directions to prepare this tasty meal.

Ingredients: red potato

Directions: chop the red potato;
fry the red potato;
prepare meal

Natural Language Observation



φ_t

(EVENTUALLY red-potato-in-player)
 \wedge (EVENTUALLY red-potato-is-chopped)
 \wedge (EVENTUALLY red-potato-is-fried)
 \wedge (EVENTUALLY meal-in-player)
 \wedge (EVENTUALLY meal-is-consumed)

LTL Instructions

1. Natural Language to LTL

Automate with GPT-3!

- **Strong few-shot translation**
 - 93.2% correct after seeing six examples
 - 5.6% correct except for parentheses

You open the copy of "Cooking: a modern approach (3rd ed.)" and start reading:

Garnier's recipe #1

and follow the directions to prepare this tasty meal.

Ingredients: red potato

Directions: chop the red potato;
fry the red potato;
prepare meal

Natural Language Observation

φ_t

(EVENTUALLY red-potato-in-player)
 \wedge (EVENTUALLY red-potato-is-chopped)
 \wedge (EVENTUALLY red-potato-is-fried)
 \wedge (EVENTUALLY meal-in-player)
 \wedge (EVENTUALLY meal-is-consumed)

LTL Instructions

1. Natural Language to LTL

Automate with GPT-3!


- **Strong few-shot translation**
 - 93.2% correct after seeing six examples
 - 5.6% correct except for parentheses
- **OOD generalization**
 - unseen adjectives (e.g. `is_grilled`)
 - unseen nouns (e.g. `carrot`)
 - unseen compositions of LTL



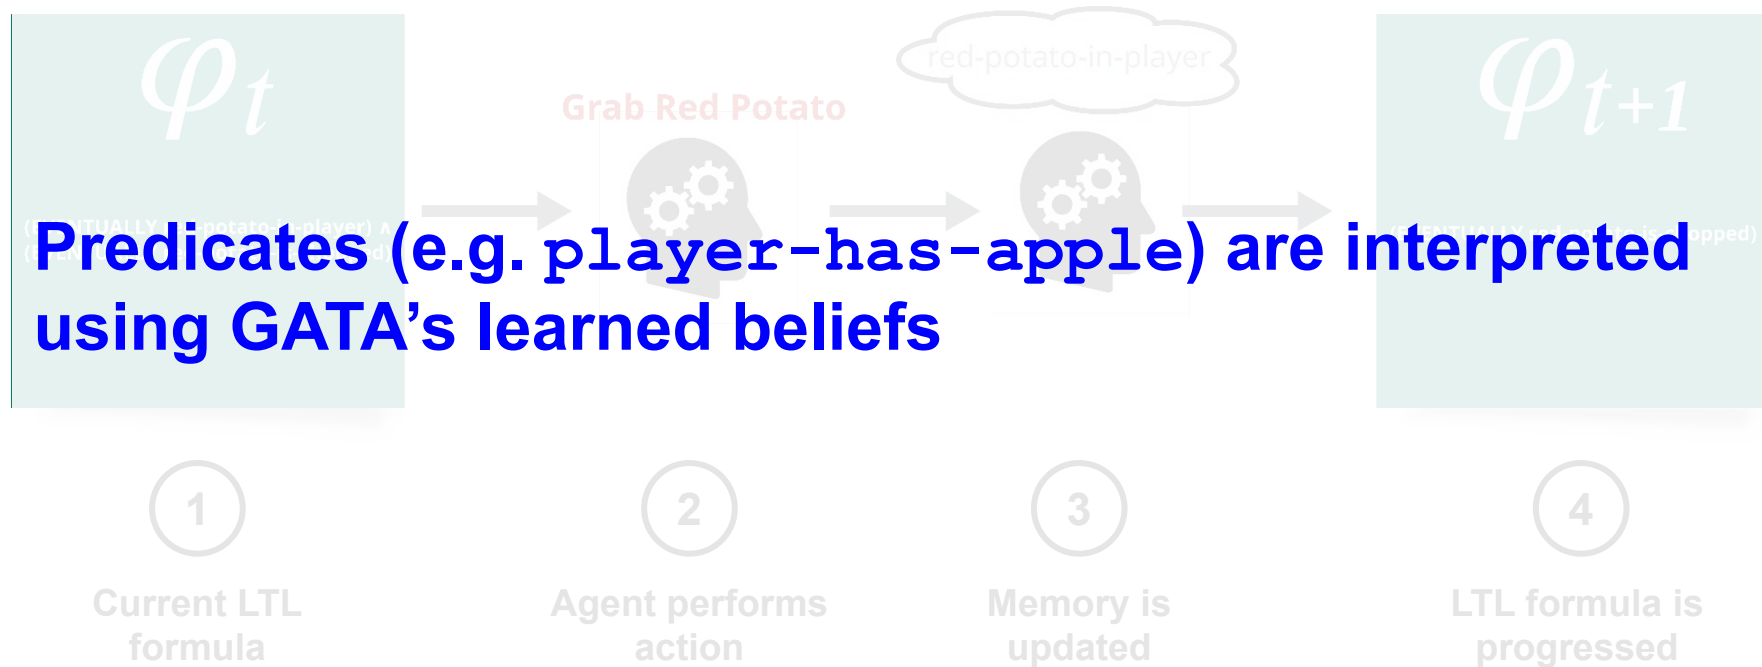
Natural Language Observation

LTL Instructions

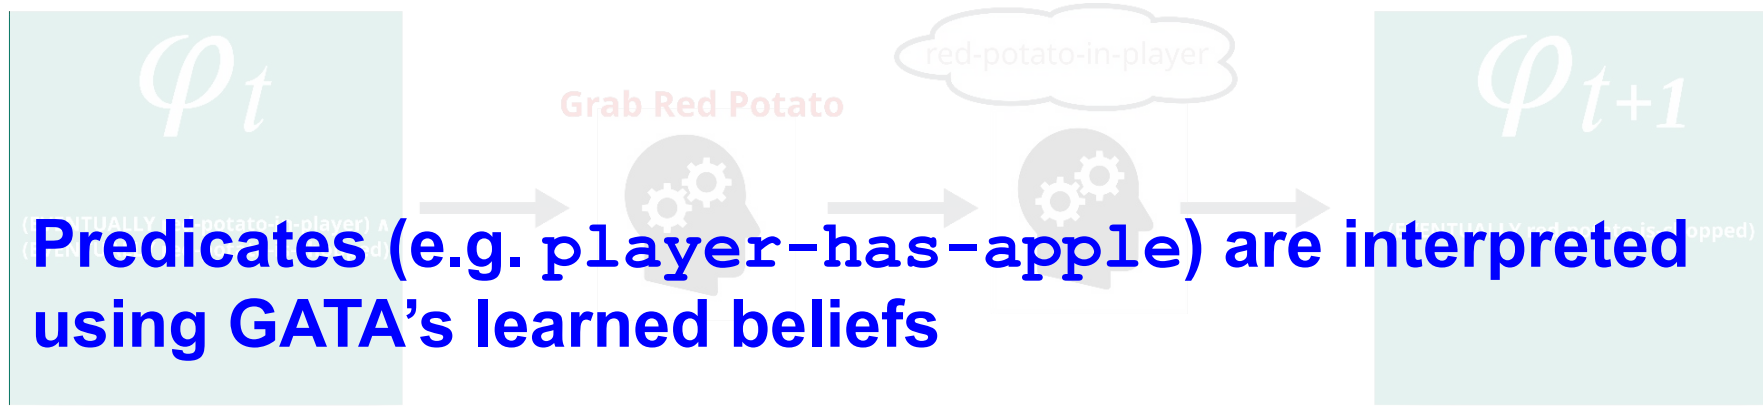
2. LTL Progression



2. LTL Progression

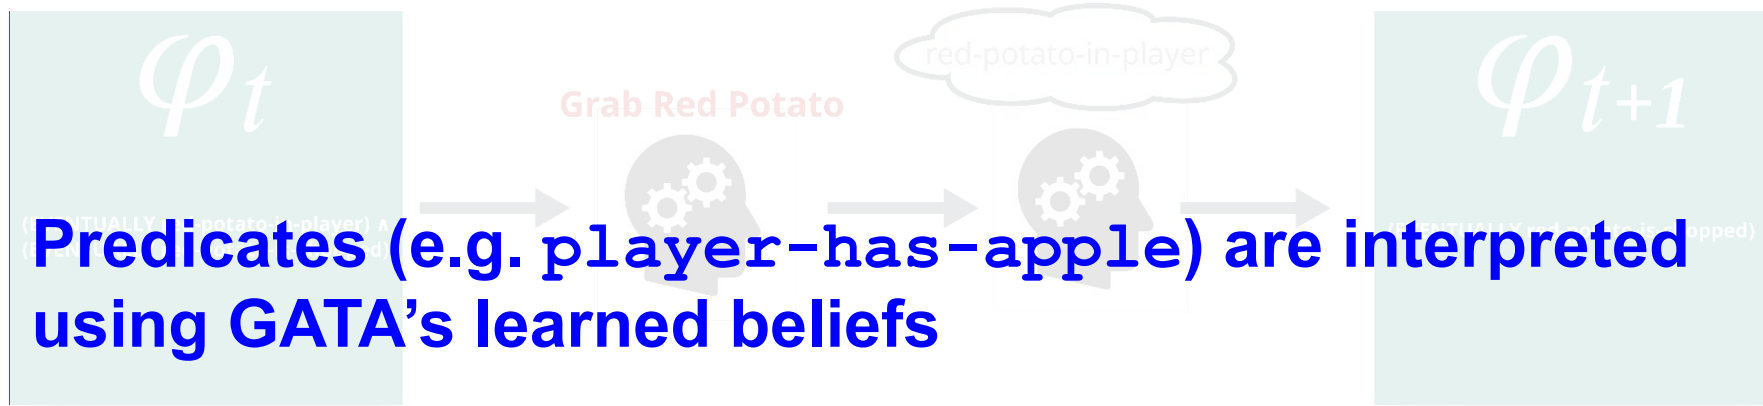


2. LTL Progression



- Does not require a ground-truth oracle/labelling function

2. LTL Progression



- Does not require a ground-truth oracle/labelling function
- LTL progression can be effective with noisy, learned models of predicates

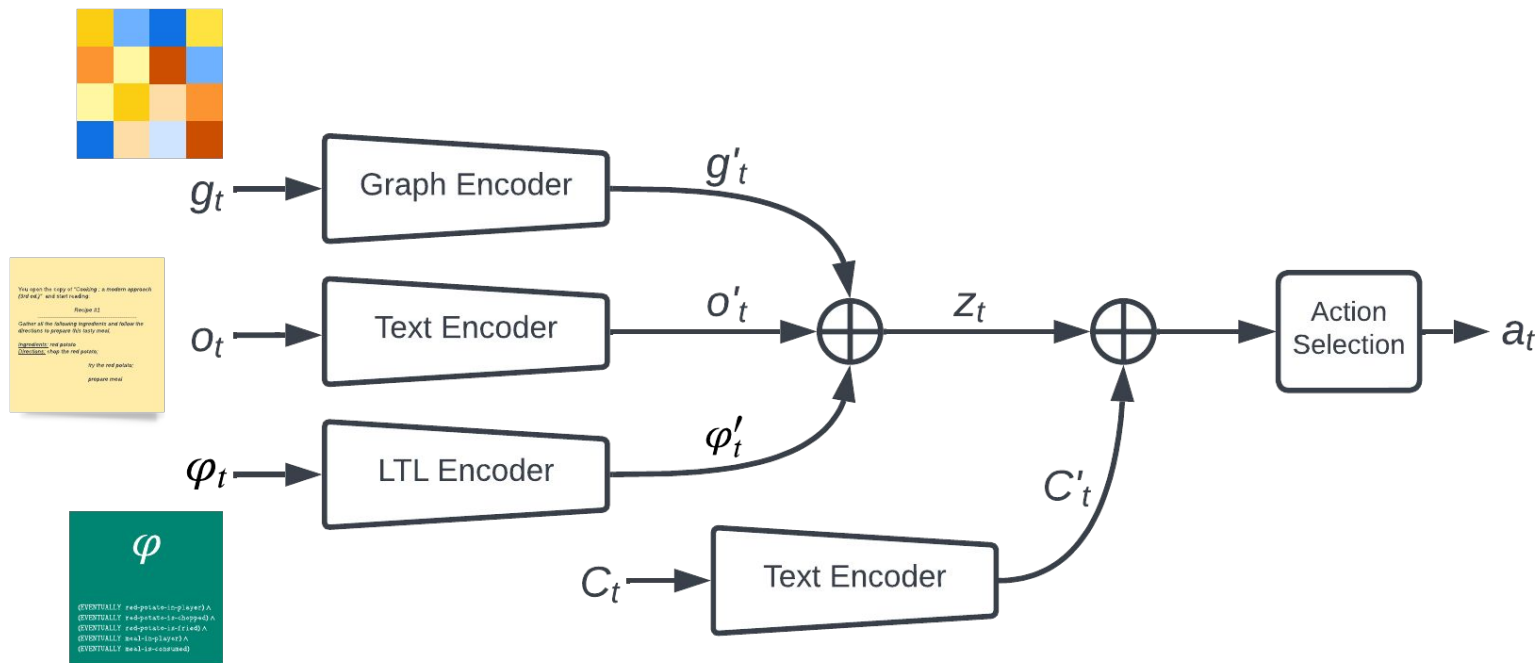
1
Initial LTL formula

2
Agent performs action


3
Memory is updated

4
LTL formula is progressed

3. LTL-conditioned Policy




Results



- Superior performance over baselines and SoTA
- Ablations highlight the importance of progression

Results



Takeaway

- SoTA TextWorld RL agents **are impervious to instructions**
- Translate natural language instructions into LTL
 - compositional syntax and semantics
 - monitor progress towards instruction completion
- Experiments with 500+ games show superior performance using LTL
- Applicable to a diversity of sequential decision-making tasks

Learning to Follow Instructions in Text-Based Games

Read the full paper!

Email us:

{mathieutuli, andrewli, pashootan, toryn, sheila}@cs.toronto.edu
ssanner@mie.utoronto.ca

Our code can be found at:

<https://github.com/MathieuTuli/LTL-GATA>