


CLIN: Continually Learning From INteractions



Bodhisattwa Prasad Majumder, Bhavana Dalvi Mishra,
Peter Jansen, Oyvind Tafjord, Niket Tandon, Li Zhang,
Chris Callison-Burch, Peter Clark



Sequential Decision Making (SDM)

- Real world decision-making tasks are **sequential** in nature.
E.g. navigation, shopping, communication
- **Text-worlds** are great simulations of real-world phenomena for testing systems ability to do **long-term planning, sequential decision making**



ScienceWorld challenges:

- Partially observable world
- Stochasticity in the environment
E.g. stove is broken
- Long/complex tasks
E.g. grow an orange

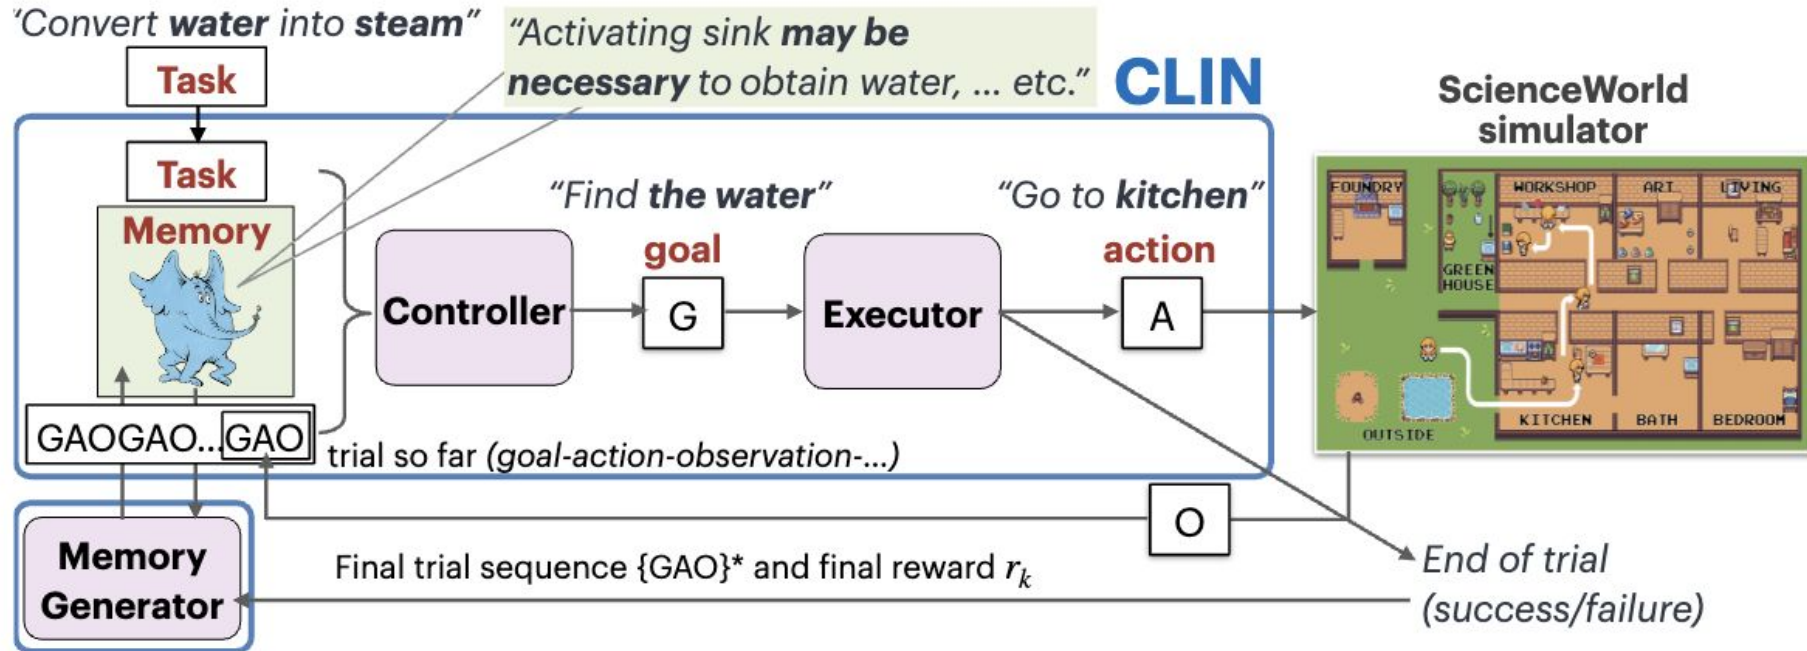
- **Existing methods are limited**
 - Require labeled data (Imitation learning),
trained for large number of epochs (Reinforcement learning)
 - Lack of interpretability, generalization

Research Questions

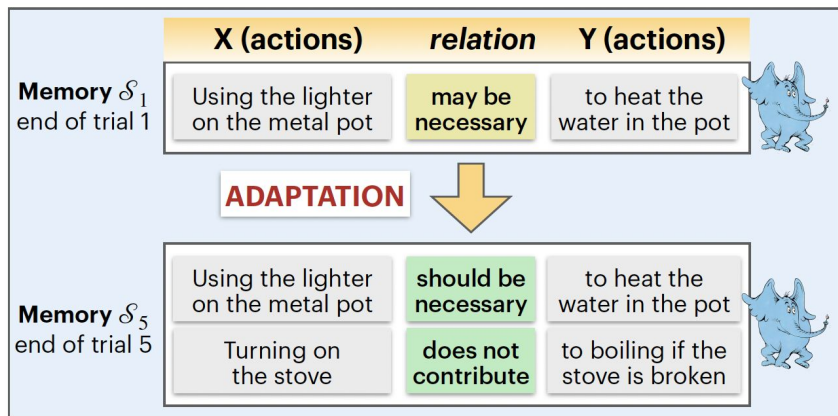
Can we build language agents that can
continually learn
from interacting and observing world changes?

How to **quickly adapt and generalize**
to a new task or environment at the test time?


CLIN: Continually Learning from INteractions



CLIN does Rapid Task Adaptation

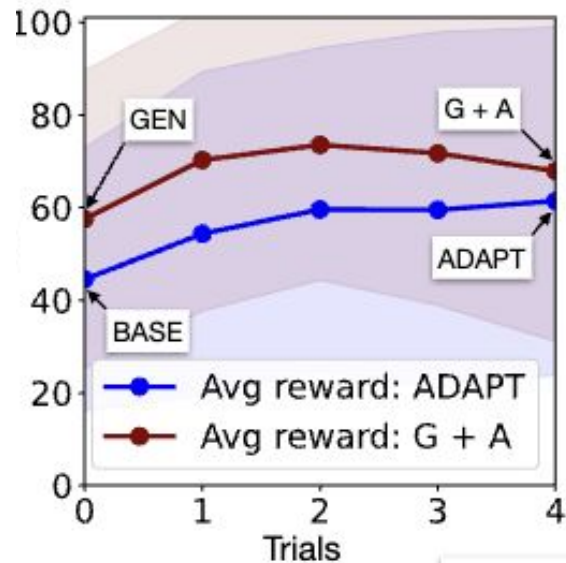
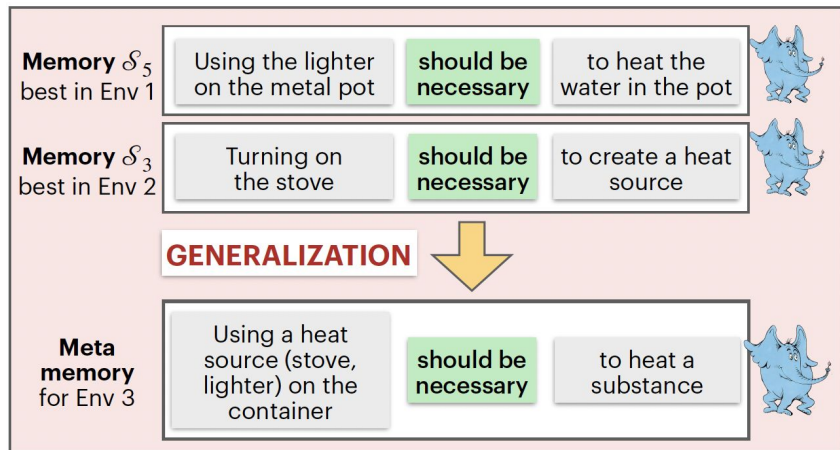


Quick adaptation, improved efficiency



- On avg. 17 points improv in reward
- Beats earlier SOTA: ReACT, Reflexion

CLIN generalizes across tasks



- Performance gain in the first trial of 38% episodes
- Agent achieves higher scores in lesser #steps


Are learnings always transferable to new tasks?

Roast a marshmallow
Melt wax

→

Melt cadmium

Using a lighter on a substance
should be necessary
to heat the substance



CLIN corrects its own memory over time

CLIN updates its memory

- **Delete:** 13% items
- **Edit:** 6% items

new environments

	GEN-Env (Trial 0)	G+A (best trial)
No. of insights	100	105
Correct insights	72.0%	91.4% ↑
Final score (on sampled tasks)	39.1	55.9 ↑

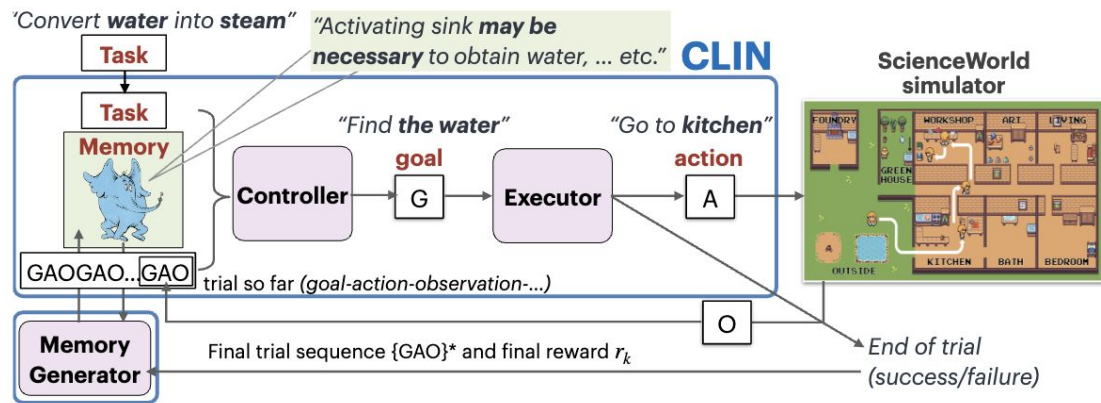
new tasks

	GEN-Task (Trial 0)	G+A (best trial)
No. of insights	98	107
Correct insights	73.9%	91.1% ↑
Final score (on sampled tasks)	43.7	58.1 ↑

Please visit our poster!



CLIN: Continually Learning From INteractions



Website: <https://allenai.github.io/clin/>