

Autonomous Task Completion Based on Goal-directed Answer Set Programming

Alexis R. Tudor¹, Gopal Gupta¹

¹University of Texas at Dallas, Richardson, USA

Abstract

Task planning for autonomous agents has typically been done using deep learning models and simulation-based reinforcement learning. This research proposes combining inductive learning techniques with goal-directed answer set programming to increase the explainability and reliability of systems for task breakdown and completion. Preliminary research has led to the creation of a Python harness that utilizes s(CASP) to solve task problems in a computationally efficient way. Although this research is in the early stages, we are exploring solutions to complex problems in simulated task completion.

Keywords

Logic in Computer Science, Software Engineering

1. Introduction

Task planning for autonomous agents has been an area of interest in recent years as robotics and deep learning have made major advances. Most approaches to task planning involve the use of deep learning models. The most popular approach is deep reinforcement learning, though recent work has used large-language models (LLMs) as well. Deep learning models generally achieve good results, however, they are uninterpretable and often produce flawed answers with no explanation. Much work has been done to improve the explainability of deep learning models, however, they remain untrustworthy.

A better solution is to use logic programming. Logic programming is a programming paradigm based primarily on the calculation of Horn clauses through the process of entailment. The most common logic programming language is Prolog, though most logic programming languages consist of a Prolog-like collection of facts and rules. One advantage of logic programs is that they are inherently interpretable and their errors can be logically understood and solved. The research proposed in this paper involves using answer set programming to complete tasks in a simulated environment. This will hopefully result in autonomous task planning that is both robust and trustworthy.

2. Background and Relevant Literature

As autonomous agents become more ubiquitous, the focus has turned to their ability to complete complex tasks in the real world, converting high-level instructions (like "fold laundry") to executable plans ("walk to clothes", "grab clothes", etc.). Autonomous task completion can mean anything from unmanned vehicles navigating from one point to another to robotic kitchen assistants designed to make certain foods. For the most part, modern autonomous systems use deep learning models to accomplish this [1]. This commonly takes the form of deep reinforcement learning and more recently LLMs. Deep learning has achieved excellent results on complex problems. However, most deep learning systems are black boxes that lack explainability and interpretability. This is especially dangerous given how dependent deep learning algorithms are on the (often flawed) data they are trained on. This makes it difficult to trust that their answers are correct and unbiased, as explored in DARPA's explainable AI retrospective [2]. This is important in critical systems, such as hospital diagnoses or

3rd Workshop on Goal-directed Execution of Answer Set Programs (GDE'24), October 12, 2024

✉ alexisrenee1@gmail.com (A. R. Tudor)

🆔 0000-0003-3992-6903 (A. R. Tudor); 0000-0001-9727-0362 (G. Gupta)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

military applications, where it has to be quickly apparent whether a model is correct or not. Deep learning as a whole is well explained in other high-quality survey papers [3]. Logic programming-based solutions focus more on the importance of explainability, which is often neglected in deep learning models.

Inductive Logic Programming (ILP) is a form of machine learning that codifies its learning in the form of first-order logic. Ever since the term was defined in 1991 [4] as the "intersection of Logic Programming and Machine Learning", ILP has served to solve machine learning problems. ILP can get results rivaling deep learning models while being inherently interpretable and explainable [5]. Recent advances in ILP, such as the FOLD family of algorithms, demonstrate that complex data can be represented in small logic programs using default rules. A more detailed description of default rules and the FOLD family of incremental learning algorithms can be found in the paper by Gupta et al. [6]. The research mentioned above uses a type of logic programming called Answer Set Programming (ASP). Unlike Prolog-based logic programming, which generates a true or false answer for a queried predicate, ASP is used to generate all entailable rules from a knowledge base. This collection is called an answer set. This can be used to generate "multiple worlds" where different answer sets are true.

Traditional ASP, like in Clingo [7], executes an answer set program through the use of a SAT solver and grounding. Grounding involves the generation of the program with all variables substituted with constants in the program. A disadvantage of this approach is that grounding is not always guaranteed to be feasible, which can leave some programs with no ASP solution. The s(CASP) system [8] solves this problem by performing a top-down goal-oriented search which eliminates the need for grounding. This advantage makes s(CASP) well-suited to the representation of complex world states and provides an advantage over other ASP systems [9].

One of the biggest weaknesses of the ILP approach to solving problems is the need for background information and 'program templates'. Program templates are a layout of how the generated information should look in the context of the logic program. A domain expert must provide this program template and explicitly logic program-based background knowledge for most ILP. Thus, for trivial examples, it would be just as easy to include the final found rules in the knowledge base at the start. Additionally, while ILP programs perform very well on data that can be represented in a logic program, logic programs have a difficult time representing complex data. These weaknesses can be overcome through the use of traditional machine learning algorithms to supplement a logic program. This approach increases explainability while utilizing the benefits of deep learning and other machine learning models, such as in the paper by Rajasekharan et al. that uses an s(CASP) knowledge base to constrain an LLM into providing more reliable results [10]. Other examples exist of using some form of knowledge base to improve deep learning algorithms [11] [12], but the use of logic programming to augment other algorithms merits further exploration.

3. Methodology

The research outlined in this paper seeks to explore the use of an s(CASP) knowledge base for autonomous task completion in a simulated virtual environment. To test our system, we use the VirtualHome simulator [13] as a playground for our s(CASP) agent to perform tasks in. The primary goal of this research is to achieve reasonably accurate task completion with goal-directed answer set programming. The end system would have a high level of explainability for decision-making, where the results are trustworthy and could be diagnosed if in error. We wish to further prove that even the very high-quality deep learning systems in use today could be augmented through the use of logic programming. Using logic in this way moves toward general artificial intelligence. Using s(CASP) to simulate how humans can perform common-sense logical interactions with the world brings us closer to reasoning AI.

An additional goal of this research is to make s(CASP) easier to use with simulators. A notable weakness of s(CASP) is that it does not have a Python API, which makes it difficult to run in line with other forms of machine learning. The software engineering goal of this research is to create a "harness"

for using s(CASP) in Python for interactions with simulators, as shown in Figure 1.

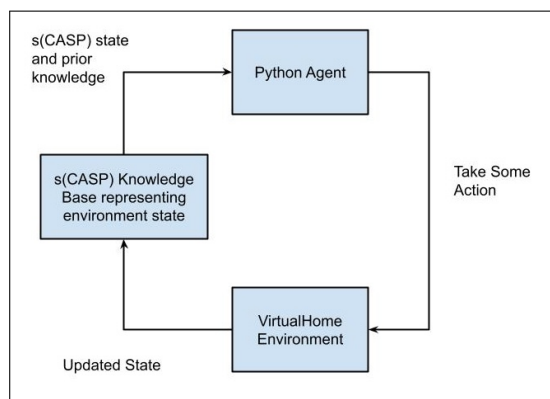


Figure 1: A diagram showing the high-level functionality of the Python harness for s(CASP). The Python harness can perform actions in the VirtualHome environment, and then convert the state of the environment to s(CASP) facts. These facts can then be used to inform the next action of the agent.

3.1. Preliminary Results

There have been promising results in producing executable actions for small-scale real-world tasks. Using the Python harness mentioned above, the simulated VirtualHome environment can be instantiated and transformed into an s(CASP) representation of the world state:

```

1  % With Time
2  current_time(1).
3  off(remotecontrol, 1).
4  inside([inside(remotecontrol, livingroom),
5          inside(character, livingroom)], 1).
6  % Without time
7  off(remotecontrol).
8  inside([inside(remotecontrol, livingroom),
9          inside(character, livingroom)]).
  
```

The above example represents a world state containing a single turned-off remote control sitting in a living room at time 1. The Python harness keeps track of a discretized world time where each action taken by the agent represents a step forward in time, however the addition of time greatly increases the complexity of the world state s(CASP) program. Using time naively in this manner results in intractable programs which loop over infinite time, and so when representing the world state we use the latter example where timestamps are not provided in the state facts. Even without the use of time, this representation of the world state easily grows to encompass a large amount of facts. The complexity of generating an answer set that accounts for all of these facts and possible worlds quickly becomes a computational obstacle. For testing purposes, the Python harness has a small-scale simulation environment built in. Still, the goal remains to execute plans in realistic environments.

To represent and complete tasks we treat task completion as a planning problem. We represent each task as a final state (i.e. if the task was to grab a remote control, the final state would include `holds(remotecontrol)`) and then formulate actions to reach that final state. The added complexity to this comes from the incorporation of the simulated world state when starting from an initial state. We use the following s(CASP) rules for the task planning problem:

```

1  % Planning
2  % Get the initial state of items close to the character
3  initial_state(List) :- close_to_character(List).
  
```

```

4 % Find a set of actions to reach the final state
5 transform(FinalState, Plan) :-
6     initial_state(State1),
7     transform(State1, FinalState, [State1], Plan).
8 transform(State1, FinalState,_,[]) :- subset(FinalState, State1).
9 transform(State1, State2, Visited, [Action|Actions]) :-
10     choose_action(Action, State1, State2),
11     update(Action, State1, State),
12     not member(State, Visited),
13     transform(State, State2, [State|Visited], Actions).
14 % We choose an action to take
15 choose_action(Action, State1, State2) :-
16     suggest(Action, State2), legal_action(Action, State1).
17 choose_action(Action, State1, _) :-
18     legal_action(Action, State1).
19 suggest(walk(X), State) :- member(close(X), State).
20 % Check if an action is legal given the state
21 legal_action(walk(X), State) :-
22     type(X, Y), Y \= character, not member(close(X), State).
23 % Update state
24 update(walk(X), State, [close(X) | State1]) :-
25     update_walking(X, State, State, [], State1).
26 % Tasks
27 complete_task(walk_to_remote, P) :-
28     type(Remote, remotecontrol), transform([close(Remote)], P).

```

These rules are a small representative subset of the rules used to generate actions to complete a task. In this very simple example, the task is to walk towards a remote control, which can be easily accomplished by the program. Using this knowledge base we can also achieve some inference. Given a final state where the agent is holding something, using the s(CASP) knowledge base constraints the agent can intuit that it first needs to walk to the item before attempting to pick it up. The most serious problem with this inference, however, is the large computation time of this program in a sufficiently large simulation.

We have made significant strides in reducing the impact of computation time on the program. To reduce computation time, we implemented a dynamic dependency graph that is used to remove facts and rules that are not relevant to the query. Figure 2 represents the dependency graph for the query ?- niece(X,Y). on the below knowledge base:

```

1 parent(tony, abe).
2 parent(tony, jill).
3 parent(abe, sarah).
4 male(tony).
5 male(abe).
6 female(jill).
7 female(sarah).
8
9 parent(Parent, Child) :- sibling(X, Child), parent(Parent, X).
10 grandparent(Grandparent, Child) :-
11     parent(Grandparent, Parent), parent(Parent, Child).
12 sibling(X,Y) :- parent(Parent, X), parent(Parent, Y), X\=Y.
13 auntuncle(AU,N) :- sibling(AU, Parent), parent(Parent, N).
14 niece(Niece, AU) :- auntuncle(AU, Niece), female(Niece).

```




Figure 2: An example dependency graph for a family tree program where niece is the queried rule.

Using the dependency graph the Python harness can simplify the above knowledge base, removing the male/1 and grandparent/2 predicates entirely. In a program of this size, the computational savings of such optimization is negligible. However, preliminary research has shown a significant time saving in the real-world environment. Table 1 demonstrates the time savings of using the dependency graph to prune the knowledge base for the specific task being accomplished on three semi-simple tasks that take one to four actions to fulfill. The computational time can be reduced from nearly thirty minutes to a fraction of a second using this approach, allowing for continued research into more complex tasks.

4. Future Work

Right now, the biggest issues facing this research concern the representation of the s(CASP) knowledge base. There are several outstanding questions.

Representing a Complex Real-World State Representing a simulation of any reasonable size leads to an exponential increase in the number of facts available in the world state. In addition to these facts, there also needs to be a set of rules adequate to perform tasks in the environment. This produces answer sets that are intractable to generate. The use of a dependency graph to pare down the knowledge base allows us to perform more complicated tasks, however there can be more optimization.

Another solution that will be explored is to keep groups of state facts and rules in different programs. The creation of modules that correspond to various tasks or locations would allow for faster calculation of relevant queries. This follows the human logic that one likely does not need their cooking knowledge if, for example, they need to walk their dog.

Table 1

Table of computational time in seconds for three tasks completed by the s(CASP) agent both optimized and unoptimized.

Task	Unoptimized Time to Complete (s)	Dependency Graph Optimized Time to Complete (s)
Grab Remote Control	13925.14	0.55
Grab Remote Control and Shirt	608.28	0.71
Grab Cell Phone and Sit on Couch	1771.21	0.64

The Passage of Time As mentioned above, the use of time in the knowledge base provides complications related to the ostensibly infinitely divisible nature of time (as posited by the famous Greek philosopher Zeno). This is a known problem with representing continuous time in logic programming and would require the inclusion of event calculus [14].

Large-scale Learning As deep learning and its applications for real-world task completion are already well explored, the value of this research lies in seeing how complex problems that the s(CASP) task planner can solve can get. To that end, explanation-based learning is a promising paradigm that would allow for generalized knowledge from a small number of examples [15] and works well with answer set programming.

Likely, s(CASP) by itself cannot encode all of the complexities of a real environment and remain tractable. Once that point is reached, there would still be benefits in combining s(CASP) with more traditional machine learning (and newer deep learning, such as LLMs) to improve performance in the former and explainability of the latter. We hope to leverage databases of task instructions and breakdowns, such as those provided by VirtualHome or ALFRED [16], to improve the performance of the s(CASP) agent at scale.

We expect to be able to answer these questions in a unified way to facilitate task completion in complex environments using s(CASP). Although solutions to these problems may always become intractable at certain levels of fidelity, there is valuable knowledge to be gained along the way.

5. Conclusion

In conclusion, this line of research could open up a broad number of solutions for challenging ILP problems. Simply creating a Python framework for the use of s(CASP) with simulated environments is an advancement for s(CASP), as it is currently lacking a Python API. Using the intersection of ILP and traditional machine learning is promising for improving the explainability and reliability of task-completing autonomous agents.

References

- [1] E. F. Morales, R. Murrieta-Cid, I. Becerra, M. A. Esquivel-Basaldua, A survey on deep learning and deep reinforcement learning in robotics with a tutorial on deep reinforcement learning, *Intelligent Service Robotics* 14 (2021) 773–805. doi:10.1007/s11370-021-00398-z.
- [2] D. Gunning, E. Vorm, J. Y. Wang, M. Turek, Darpa’s explainable ai (xai) program: A retrospective, *Applied AI Letters* 2 (2021) e61. doi:10.1002/ai12.61.
- [3] S. Dong, P. Wang, K. Abbas, A survey on deep learning and its applications, *Computer Science Review* 40 (2021) 100379. doi:10.1016/j.cosrev.2021.100379.
- [4] S. Muggleton, Inductive logic programming, *New Generation Computing* 8 (1991) 295–318. doi:10.1007/bf03037089.
- [5] Z. Zhang, L. Xu, L. Yilmaz, B. Liu, A critical review of inductive logic programming techniques for explainable ai, 2023. arXiv:2112.15319.
- [6] G. Gupta, H. Wang, K. Basu, F. Shakerin, E. Salazar, S. C. Varanasi, P. Padalkar, S. Dasgupta, Logic-based explainable and incremental machine learning, *Prolog: The Next 50 Years* (2023) 346–358. doi:10.1007/978-3-031-35254-6_28.
- [7] M. Gebser, R. Kaminski, B. Kaufmann, T. Schaub, Clingo = asp + control: Preliminary report, *ArXiv abs/1405.3694* (2014). doi:10.48550/arXiv.1405.3694.
- [8] J. Arias, M. Carro, E. Salazar, K. Marple, G. Gupta, Constraint answer set programming without grounding, *Theory and Practice of Logic Programming* 18 (2018) 337–354. doi:10.1017/S1471068418000285.

- [9] M. Gelfond, Y. Kahl, Knowledge representation, reasoning, and the design of Intelligent Agents: The answer-set programming approach, Cambridge University Press, 2014. doi:10.1017/CBO9781139342124.
- [10] A. Rajasekharan, Y. Zeng, P. Padalkar, G. Gupta, Reliable natural language understanding with large language models and answer set programming, Electronic Proceedings in Theoretical Computer Science 385 (2023) 274–287. doi:10.4204/EPTCS.385.27.
- [11] I. Uchendu, T. Xiao, Y. Lu, B. Zhu, M. Yan, J. Simon, M. Bennice, C. Fu, C. Ma, J. Jiao, S. Levine, K. Hausman, Jump-start reinforcement learning, 2023. arXiv:2204.02372.
- [12] Z. Hao, S. Liu, Y. Zhang, C. Ying, Y. Feng, H. Su, J. Zhu, Physics-informed machine learning: A survey on problems, methods and applications, 2023. arXiv:2211.08064.
- [13] X. Puig, K. Ra, M. Boben, J. Li, T. Wang, S. Fidler, A. Torralba, Virtualhome: Simulating household activities via programs, 2018. arXiv:1806.07011.
- [14] S. C. Varanasi, J. Arias, E. Salazar, F. Li, K. Basu, G. Gupta, Modeling and verification of real-time systems with the event calculus and s(casp), in: Practical Aspects of Declarative Languages: 24th International Symposium, PADL 2022, Philadelphia, PA, USA, January 17–18, 2022, Proceedings, Springer-Verlag, Berlin, Heidelberg, 2022, p. 181–190. doi:10.1007/978-3-030-94479-7_12.
- [15] J. Strout, Y. Zhang, R. Mooney, Do human rationales improve machine explanations?, in: Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP, Association for Computational Linguistics, Florence, Italy, 2019, pp. 56–62. doi:10.18653/v1/W19-4807.
- [16] M. Shridhar, J. Thomason, D. Gordon, Y. Bisk, W. Han, R. Mottaghi, L. Zettlemoyer, D. Fox, Alfred: A benchmark for interpreting grounded instructions for everyday tasks, 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 10737–10746. doi:10.1109/cvpr42600.2020.01075.