Towards building Autonomous AI Agents and Robots for Open World Environments

Doctoral Consortium

Shivam Goel
Tufts University
Medford, USA
shivam.goel@tufts.edu

ABSTRACT

The shift of AI agents from controlled laboratory environments to real-world applications, such as autonomous vehicles and service robots, demands robust algorithms for navigating the intricacies of open-world scenarios. While traditional AI agents show proficiency in predictable, closed-world settings, their performance often diminishes in the dynamic and unforeseen conditions of real-world environments. My dissertation focuses on developing methods, frameworks, and domains that push the boundaries of open-world problem-solving in AI agents and robots. The central thesis question explores how AI agents can rapidly learn and adapt in open-world settings while tackling long-horizon, complex tasks. My work proposes integrative frameworks that combine reinforcement learning with symbolic planning, enabling on-the-fly adaptation of agents. Furthermore, we also propose environments designed for developing and assessing agent architectures adept at handling novelty. These advancements in open-world learning are pivotal in enhancing adaptability, speed, and robustness in AI agents and robots, laying a critical foundation for the development of resilient autonomous systems.

KEYWORDS

open-world AI, neurosymbolic learning, continual learning

ACM Reference Format:
Shivam Goel. 2024. Towards building Autonomous AI Agents and Robots for Open World Environments: Doctoral Consortium. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 3 pages.

1 INTRODUCTION

In the realm of artificial intelligence research, there is a growing interest in exploring environments that extend beyond the traditional "closed-world" assumption [10]. In conventional closed-world settings, agents are pre-equipped with a comprehensive understanding of all relevant concepts pertinent to their tasks. However, as we progress into more dynamic and unpredictable domains, the ability to not only recognize but also to swiftly learn and adeptly adapt to situations that are conceptually novel is becoming critical.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand.

This shift towards "open-world" paradigms has prompted a substantial investment of research efforts [2–5, 8, 9, 12]. These recent works underscore the necessity for AI systems that can operate in environments where new information can emerge and where the agents must evolve their abilities over time. This evolution in AI research demands a paradigmatic shift in how we conceive and develop agents. It compels us to rethink the very fundamentals of knowledge acquisition, concept formation, and adaptability in AI. The ability of an AI system to not only process but also to intuitively grasp and adapt to unforeseen circumstances (i.e., novelties) is paramount to developing AI agents for the real world.

In my dissertation thesis, I am exploring the realm of AI through the lens of open-world learning. The central thesis question of the dissertation explores how AI agents can rapidly learn and adapt in open-world settings while tackling long-horizon, complex tasks. As the research in open-world learning is relatively new, so there needs to be a theoretical understanding of what open-world means with respect to the agent, how we define and characterize novelty, how we evaluate such agents, and most importantly, how we can develop algorithms and sophisticated agent architectures to make robust and useful agents to operate in the real-world environments. In the next sections, I will first provide some background on the concept of open-world and novelty, then briefly describe my existing works. I will then describe the future directions I intend to develop as a part of my dissertation, followed by a conclusion.

2 BACKGROUND

Most AI research considers novelty and open-world scenarios as a form of concept drift, as outlined in various works [8, 13, 15]. We respect the views of the authors; however, in our formulation of open-world and novelty, we diverge from these views. We argue that novelty is neither a property of an object [1] nor a concept drift. Instead, we propose that novelty should be seen as an inherently agent-relative concept. This perspective views novelty as a relationship between the elements in the environment and an agent's cognitive system. The relativity of novelty is evident when considering that what may be a new object or concept for one agent, such as a Spitzer Space Telescope for a layman who has never encountered it or lacks knowledge about it, may not be novel for another agent, like an astronomer who uses it routinely.

Therefore, in the discussions of novelty within this dissertation, it is essential to understand that we are always considering it from the perspective of an epistemic agent. For this agent, something is considered novel if it is beyond their range of experience and if they cannot form representations of it based on their existing knowledge base [11].

3 EXISTING WORKS

3.1 Testbeds and Evaluation Systems: Novelgidworlds & NovelGym

To build agents that can function in open-world scenarios, it is important to have test beds, evaluation protocols, and evaluation methodologies for standardizing the research in this direction. To address this gap in the literature, we developed an openAI gym environment – Novelgridworlds [6] that was specifically designed to inject novelties in a grid world environment and evaluate agent architectures that can handle novelties. We further enhanced this system to introduce an ecosystem – NovelGym [7], that offered an easy injection of novelties and enhanced the capabilities of the system to implement hybrid training architectures that can develop novelty-aware agents. We also theorized the concepts of novelty and proposed novel training protocols and evaluation metrics for novelty-aware agents [7].

3.2 Algorithms and Architectures: RapidLearn & Cognitive architecture

In developing agent architectures that are novelty-aware and have faster adaptability to open-world scenarios, it is important to focus on information reuse. We developed a framework – RAPid-Learn [5] that integrates symbolic planning with reinforcement learning to adapt to novelty scenarios. The proposed approach utilized reinforcement learning to learn executors that repair the failed plan induced by novelties. These executors are abstracted in the form of options [14] and can be reused to plan again. While learning these executors, our framework also utilized knowledge-guided exploration techniques to perform efficient exploration. We achieved this due to the synergy between the symbolic planning and reinforcement learning systems. Our results, compared with the adapted existing techniques (transfer learning), improved multiple orders of magnitude in learning efficiency. We have also developed a cognitive architecture that merges symbolic planning and counterfactual reasoning with the principles of reinforcement learning and advanced computer vision methods. This architecture aims at efficient novelty detection and adaptation. The novel neurosymbolic cognitive architecture we proposed is complemented by general algorithms designed for handling novelty and performing explorations. These explorations utilize inference and reasoning methods, incorporating machine learning techniques like Behavior Cloning (BC), cyclic Generative Adversarial Networks (Cycle-GANs), and reinforcement learning.

4 FUTURE DIRECTIONS

The goal of the research is to move beyond closed-world assumptions in AI and develop novelty-aware agents that can act in the real world. In my existing works, we created testbeds and developed frameworks that are still limited to grid worlds and simulated environments such as games. I aim to delve deeper into developing frameworks that can be applicable to robotic systems. To that end, I have identified three important things: (1) Considering the robots' huge state and action spaces, developing methods to explore faster and learn better policies is important. (2) How can we learn perception models from raw data to aid neurosymbolic architectures for novelty handling? (3) How can we utilize human-guided hints in learning efficient policies for novelty handling?

4.1 Novelty adaptation in robotic systems

I aim to build algorithms that can perform open-world novelty adaptation in robotic systems. However, the problem becomes harder as the state space and the action space become continuous, and the domain becomes stochastic in nature. In my previous works, I have assumed discrete action spaces and deterministic domains. Moreover, in robotic systems, another aspect is the safety constraints. To enable research in this direction, we have developed robotic simulation environments that feature robots such as locobot, and UR5 arms and have built long-horizon tasks and test-beds in which novelties can induce action execution failures for the robot. My aim is to use symbolic inferences and Bayesian techniques to learn executors in constrained and efficient manner to adapt to novelties that occur in these environments.

4.2 Perception models for open-world learning

Modern AI systems' state representations are either handcrafted through humans or subject experts or due to advanced deep learning methods; the state is represented as raw sensor inputs (primarily using images). However, to design a system capable of novelty (sudden changes in agent's knowledge) handling, It is important to learn perception models automatically through learning. In this work, we propose a neurosymbolic framework that takes input as an image and performs symbolic inference and knowledge-guided reasoning to aid in learning a perception model.

4.3 Constrained based representations and human-guided learning

In this direction, I am focusing on a neurosymbolic approach that leverages provided hints—expressed as logical constraints to inform and adapt agents' decision-making processes at test times. By incorporating these hints into a RL framework, the agent can adapt to novelties, enhancing its performance efficiently. The approach aims to integrate symbolic reasoning with reinforcement learning, allowing for an efficient policy adaptation in response to new environmental situations. The logical constraints constrain the state space, directing the policy towards favorable actions while maintaining adaptability. The eventual goal of such frameworks is to incorporate human-guided learning into the robotic systems, as we believe that a human can enhance learning efficiency, especially in novelty-induced scenarios.

5 CONCLUSION

The primary objective of my research is to enhance the capabilities of AI agents and robots, enabling them to be effective in real-world environments. My approach centers on developing methods that are not only safe and reliable but also sample-efficient. I am particularly focused on hybrid methodologies, as I believe that integrating diverse AI techniques can enable efficient learning, improve knowledge reuse, and yield solutions that are directly applicable to real-world AI challenges.

REFERENCES

[1] Terrance E Boult, Steve Cruz, Akshay Raj Dhamija, Manuel Gunther, James Henrydoss, and Walter J Scheirer. 2019. Learning and the unknown: Surveying steps toward open world recognition. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 9801–9807.

[2] Annie Chen, Archit Sharma, Sergey Levine, and Chelsea Finn. 2022. You only live once: Single-life reinforcement learning. *Advances in Neural Information Processing Systems* 35 (2022), 14784–14797.

[3] Annie S Chen, Govind Chada, Laura Smith, Archit Sharma, Zipeng Fu, Sergey Levine, and Chelsea Finn. 2023. Adapt On-the-Go: Behavior Modulation for Single-Life Robot Deployment. *arXiv preprint arXiv:2311.01059* (2023).

[4] Evana Gizzi, Lakshmi Nair, Sonia Chernova, and Jivko Sinapov. 2022. Creative problem solving in artificially intelligent agents: A survey and framework. *Journal of Artificial Intelligence Research* 75 (2022), 857–911.

[5] Shivam Goel, Yash Shukla, Vasanth Sarathy, Matthias Scheutz, and Jivko Sinapov. 2022. RAPid-Learn: A Framework for Learning to Recover for Handling Novelties in Open-World Environments.. In *2022 IEEE International Conference on Development and Learning (ICDL)*. IEEE, 15–22.

[6] Shivam Goel, Gyan Tatiya, Matthias Scheutz, and Jivko Sinapov. 2021. Novelgridworlds: A benchmark environment for detecting and adapting to novelties in open worlds. In *AAMAS Adaptive Learning Agents (ALA) Workshop*.

[7] Shivam Goel, Yichen Wei, Panagiotis Lymperopoulos, Matthias Scheutz, and Jivko Sinapov. 2024. NovelGym: A Flexible Ecosystem for Hybrid Planning and Learning Agents Designed for Open Worlds. arXiv:2401.03546 [cs.AI]

[8] Khimya Khetarpal, Matthew Riemer, Irina Rish, and Doina Precup. 2022. Towards continual reinforcement learning: A review and perspectives. *Journal of Artificial Intelligence Research* 75 (2022), 1401–1476.

[9] Sahisnu Mazumder and Bing Liu. 2024. Open-World Continual Learning: A Framework. In *Lifelong and Continual Learning Dialogue Systems*. Springer, 21–47.

[10] Jack Minker. 1982. On indefinite databases and the closed world assumption. In *6th Conference on Automated Deduction: New York, USA, June 7–9, 1982 6*. Springer, 292–308.

[11] Faizan Muhammad, Vasanth Sarathy, Gyan Tatiya, Shivam Goel, Saurav Gyawali, Mateo Guaman, Jivko Sinapov, and Matthias Scheutz. 2021. A novelty-centric agent architecture for changing worlds. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*. 925–933.

[12] Vasanth Sarathy, Daniel Kasenberg, Shivam Goel, Jivko Sinapov, and Matthias Scheutz. 2020. Spotter: Extending symbolic planning operators through targeted reinforcement learning. *arXiv preprint arXiv:2012.13037* (2020).

[13] Andrés L Suárez-Cetrulo, David Quintana, and Alejandro Cervantes. 2023. A survey on machine learning for recurring concept drifting data streams. *Expert Systems with Applications* 213 (2023), 118934.

[14] Richard S Sutton, Doina Precup, and Satinder Singh. 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence* 112, 1-2 (1999), 181–211.

[15] Qiuyan Xiang, Lingling Zi, Xin Cong, and Yan Wang. 2023. Concept Drift Adaptation Methods under the Deep Learning Framework: A Literature Review. *Applied Sciences* 13, 11 (2023), 6515.