"Crow-level" Tool Use

Hengli Li
Yuanpei College
Peking University
2000017754@stu.pku.edu.cn

Abstract

This paper delves into the fascinating challenge of enabling robots to acquire the capacity to employ novel tools solely through the observation of another agent's behavior. Unlike traditional imitation learning, this problem necessitates high-level cognitive abilities and the correct alignment of observed behaviors with the robot's physical embodiment. We discuss the complexities involved, emphasizing the importance of causal reasoning and correct grounding to the robot's body. To address this challenge, we propose key aspects for consideration, including intuitive physics and causal reasoning, reward shaping, and the development of an unlimited action space with robust grounding methods. The exploration of this challenge holds great promise for the field of artificial intelligence and robotics, offering exciting possibilities for intelligent and adaptive machines in the future.

Figure 1: From Nat Geo Wild Video, a crow making tool

1 Introduction

As delineated in Figure 1, the depiction captures a remarkable scene in which a crow ingeniously fashions a tool from a seemingly ordinary twig. Initially, the twig presents itself as a straight and unadorned specimen, replete with minor twigs and leaves. However, through astute observation and a keen appreciation for the superior functionality of hooked twigs over their straight counterparts in the hands of humans, the crow adapts its approach. It learns to skillfully carve a hook at the twig's extremity, thus enhancing its capacity to procure sustenance. This remarkable capacity to exploit available materials as implements for problem-solving is commonly referred to as "tool use." Traditionally, tool use has been regarded as a hallmark of human superiority in the animal kingdom [9]. Notably, Jane Goodall's observations of ant-dipping chimpanzees challenged the presumed uniqueness of tool use within the purview of human beings. An extensive body of literature underscores the pivotal role of tool use in distinguishing humans from their evolutionary kin [9]. This body of work introduces various crucial cognitive facets, including but not limited to causal reasoning, function representation, and social intelligence. This paper, however, is not intended to engage in a definitive judgment regarding the criterion for human exceptionalism. Instead, its primary focus is to explore specific facets of tool use within the context of the current discourse in artificial intelligence (AI). In the domain of AI, a concept analogous to tool use is identified as "robot tool use" [6]. In their comprehensive survey, Qin et al. (2022) define this concept across three core dimensions:

(I) Perception, (II) Manipulation, and (III) High-level Cognition. Notably, they enumerate seven key challenges associated with robot tool use, most of which predominantly pertain to the manipulation aspect, with relatively limited emphasis on learning. This paper endeavors to introduce an eighth challenge to the domain of robot tool use, focusing specifically on the perspective of learning: "How can a robot acquire the capacity to employ a novel tool through the sole observation of another agent's behavior?" This problem closely parallels the setting of imitation learning [2], wherein an agent is tasked with acquiring a behavioral policy solely from visual or recorded demonstrations of human actions. However, the problem at hand deviates in two key dimensions:

1. High-Level Cognitive Ability: The challenge of enabling a robot to learn tool use through observational learning necessitates a degree of high-level cognitive capability that surpasses mere mimicry.
2. Correct Grounding to Robot Body: Unlike standard imitation learning, the robot must correctly translate and adapt the observed behavior to its own embodied form.

The structure of this paper unfolds as follows: We shall commence by delineating the aforementioned challenge and expound upon the formidable difficulties it presents. Subsequently, we shall offer insightful perspectives on potential strategies for addressing this challenge within the purview of AI and robotics research.

2 The Challenge

As articulated earlier, we present the fundamental challenge that this work aims to address: "How can a robot acquire the capacity to employ a novel tool solely through the observation of another agent's behavior?" It is worth noting that this challenge diverges from traditional imitation learning in two essential dimensions: (I) the requisite integration of high-level cognitive capabilities and (II) the accurate alignment and adaptation of observed behaviors to the robot's physical embodiment. These distinctive aspects significantly underscore the inherent complexities of the challenge at hand.

2.1 High-Level Cognitive Ability

In their comprehensive work, Vaesen [9] aptly assert that the comprehension of causality transcends mere observation of the statistical association between a cause and its corresponding effect. They emphasize that a complete understanding of causality necessitates the ability to infer the mechanistic link that underlies this cause-effect relationship, thereby elucidating the factors responsible for the observed covariance.

Moreover, kyoung Ahn et al. [3] underscores a critical aspect of human cognition by positing that, rather than merely possessing knowledge of covariation between events, humans exhibit a distinct preference for grasping the underlying causal relationships that govern these events. This emphasis on causal relationships as opposed to mere correlation underscores the importance of developing robust causal reasoning abilities for the attainment of significant competence in tool use.

However, it is noteworthy that current research, such as the recent study by Gao et al. [1], reveals limitations in the realm of causal reasoning, even within state-of-the-art Language Model (LLM) systems, including but not limited to ChatGPT. These findings underscore the pressing need for advancements in the realm of causal reasoning, particularly concerning AI models, to bridge this cognitive gap and further their capacity to engage effectively in complex tasks such as tool use.

2.2 Correct Grounding to Robot Body

As elucidated by Rizzolatti and Craighero [7], mirror neurons in monkeys exhibit a remarkable propensity for generalization. Specifically, they manifest robustness in the face of diverse visual stimuli representing the same action. This robustness implies that these neurons can effectively imitate or learn behaviors by observing the grasping hand of a human agent. This observation leads to a significant inference: for a robot to attain a level of tool use proficiency comparable to that of a monkey, it must operate within an action space that is either extensive or even unbounded. Additionally, the robot must possess a grounding function that accurately maps the behavior of other agents onto its own action space.

While contemporary reinforcement learning algorithms have demonstrated considerable success in various domains, including game playing [5, 8], they are inherently constrained by finite action spaces. Similarly, imitation learning algorithms, as highlighted in Kai Arulkumaran [2], primarily excel at learning from well-defined and pre-mapped behaviors. In the context of encountering novel scenarios demanding tool use, the limitations posed by finite action spaces and reliance on restricted mapping functions become evident shortcomings. To surmount these challenges, there is a pressing need for innovative approaches that can empower robots to navigate uncharted tool-use scenarios with greater adaptability and flexibility, transcending the constraints of finite action spaces and predefined mappings.

3 A Possible Solution

As demonstrated above, the integration of existing modules falls short of yielding a proficient tool-using agent. In this section, we endeavor to delineate crucial aspects that warrant dedicated attention in the pursuit of addressing this formidable challenge. We posit that the following key facets merit focused consideration:

1. Intuitive Physics and Causal Reasoning [10]: The capacity to discern the causal relationship underpinning a given tool usage scenario is contingent upon the presence of robust intuitive physics and causal reasoning modules. These modules are instrumental in assessing the outcomes of observed behaviors and pinpointing the pivotal causal factors that lead to specific outcomes.
2. Reward Shaping [4]: Tool use scenarios often entail a transformation from the current situation to one that is considerably different, as exemplified in Figure 1, where the crow must transform a raw twig into a hooked one to achieve its goal. This transformation necessitates a guided path, which can be achieved through the application of reward shaping. The development of effective and general reward-shaping methodologies is crucial to guide the learning process.
3. Unlimited Action Space and Grounding Method: The attainment of a tool use capability akin to the monkey's ability to learn from a human's hand requires two interconnected attributes: an unbounded or expansive action space and a robust grounding method. Agents need the capacity to accommodate a diverse range of actions, extending beyond predefined limitations. Furthermore, these agents must possess the ability to effectively map other agents' actions, organs, or body parts to their own, facilitating comprehensive understanding and adaptation.

4 Conclusion

In conclusion, this paper has explored the intriguing challenge of robot tool-using. This problem goes beyond traditional imitation learning, requiring high-level cognitive abilities and correct grounding to the robot's physical embodiment. To address this formidable challenge, several key facets have been proposed for consideration, including intuitive physics and causal reasoning, reward shaping, and the development of an unlimited action space and robust grounding methods. These elements are pivotal in guiding the development of proficient tool-using agents in the context of artificial intelligence and robotics research. The exploration of this challenge opens up exciting possibilities in the field of AI, as it delves into the intersection of cognition, perception, and manipulation, paving the way for robots to harness observational learning and advance their capabilities in utilizing tools. As researchers continue to tackle this challenge, it promises to have profound implications for the future of AI and robotics, driving us closer to a new era of intelligent and adaptive machines.

References

[1] Jinglong Gao, Xiao Ding, Bing Qin, and Ting Liu. Is chatgpt a good causal reasoner? a comprehensive evaluation, 2023. 2

[2] Dan Ogawa Lillrank Kai Arulkumaran. A pragmatic look at deep imitation learning. 2023. 2, 3

[3] Woo kyoung Ahn, Charles W. Kalish, Douglas L. Medin, and Susan A. Gelman. The role of covariation versus mechanism information in causal attribution. *Cognition*, 54(3):299–352, 1995. ISSN 0010-0277. doi: https://doi.org/10.1016/0010-0277(94)00640-7. URL `https://www.sciencedirect.com/science/article/pii/0010027794006407`. 2

[4] Ofir Marom and Benjamin Rosman. Belief reward shaping in reinforcement learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), Apr. 2018. doi: 10.1609/aaai.v32i1.11741. URL `https://ojs.aaai.org/index.php/AAAI/article/view/11741`. 3

[5] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning, 2013. 3

[6] Meiying Qin, Jake Brawer, and Brian Scassellati. Robot tool use: A survey. *Frontiers in robotics and AI*, 9:1009488, 2022. ISSN 2296-9144. doi: 10.3389/frobt.2022.1009488. URL `https://europepmc.org/articles/PMC9885045`. 1

[7] Giacomo Rizzolatti and Laila Craighero. The mirror-neuron system. *Annual review of neuroscience*, 27:169–92, 02 2004. doi: 10.1146/annurev.neuro.27.070203.144230. 2

[8] David Silver, Aja Huang, Christopher Maddison, Arthur Guez, Laurent Sifre, George Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529:484–489, 01 2016. doi: 10.1038/nature16961. 3

[9] Krist Vaesen. The cognitive bases of human tool use. *Behavioral and Brain Sciences*, 35(4): 203–218, 2012. doi: 10.1017/S0140525X11001452. 1, 2

[10] Yixin Zhu, Tao Gao, Lifeng Fan, Siyuan Huang, Mark Edmonds, Hangxin Liu, Feng Gao, Chi Zhang, Siyuan Qi, Ying Nian Wu, et al. Dark, beyond deep: A paradigm shift to cognitive ai with humanlike common sense. *Engineering*, 6(3):310–345, 2020. 3