CoRe Essay 8
Possible Ways to Learn and Represent Human Utility

Haochen Zhao*
Yuanpei College
Peking University
2100017417@stu.pku.edu.cn

Abstract

Understanding and representing human utility is a multifaceted challenge crucial for the advancement of artificial intelligence (AI). Human decision-making is inherently subjective, complex, and varies across individuals, making the task of estimating utility functions a daunting endeavor. This essay explores various methodologies employed by AI to learn and represent human utility, analyzing their advantages and disadvantages. From computer vision approaches to reinforcement learning or robotics techniques, we delve into the diverse landscape of methodologies aiming to capture the nuanced nature of human preferences. By critically examining these methods, we aim to contribute to the ongoing discourse on the intersection of AI and human decision-making, paving the way for more informed and ethically grounded applications of AI technologies.

1 Introduction

The ability of artificial intelligence (AI) systems to comprehend and adapt to human preferences is fundamental for their successful integration into diverse domains, ranging from personalized recommendations to human-robot interaction. At the heart of this integration lies the challenge of understanding and representing human utility—a complex and subjective concept that encapsulates the satisfaction or desirability associated with various choices and outcomes. The diverse and individualistic nature of human preferences poses a formidable hurdle for AI systems seeking to model utility functions accurately.

In fact, utility is a concept from economics at the beginning. By definition, utility is a measure of the satisfaction that a certain person has from a certain state of the world.[2] The term was introduced initially as a measure of pleasure or happiness as part of the theory of utilitarianism. In this context, the utilities of different agents in the same state are comparable, which makes utility a more computationally feasible concept when it comes to the field of machine social intelligence. However, utility is not an explicit representation to learn. As a result, the modeling of utility and utility function becomes a key challenge. In this essay, we are trying to learn and represent human utility from explicit human sensory and actions via different ways, including computer vision, robotics and psychology.

2 Computer Vision Method

Vision makes a large part in human sensory. We often perceive others' utility in daily live via visual observation. Thus, learning utility from video seems to be a reasonable solution. As early as 2016, a work has been done following this idea.[3]

In this work, they propose an innovative computer vision method aimed at capturing the intricate nuances of human interaction with real-world objects. The approach centers around a novel notion of

*Thanks to course instructor Yixin Zhu, TA Guangyuan Jiang and Yuyang Li for their helpful suggestions.

affordance, where physical quantities generated during the human body's interaction with objects are considered. Unlike traditional methods that rely on geometric compatibility between body poses and 3D objects, this framework introduces a learning paradigm that incorporates the concept of human utilities. This not only provides a deeper understanding of object affordance but also offers a finer-grained account of how individuals interact with these objects.

Crucially, this method goes beyond merely capturing physical comfort. They extend the concept of human utilities to account for a spectrum of preferences, transcending comfort intervals. This includes considerations for meaningful tasks within scenes and spatiotemporal constraints inherent in motion planning. By training the system on observed choices and incorporating human utilities, it is enabled to discern not only physical comfort but also preferences aligned with diverse scenarios, providing a foundation for applications in robot task planning and beyond. This holistic approach marks a significant step toward a more nuanced understanding of human-object interactions in real-world contexts.

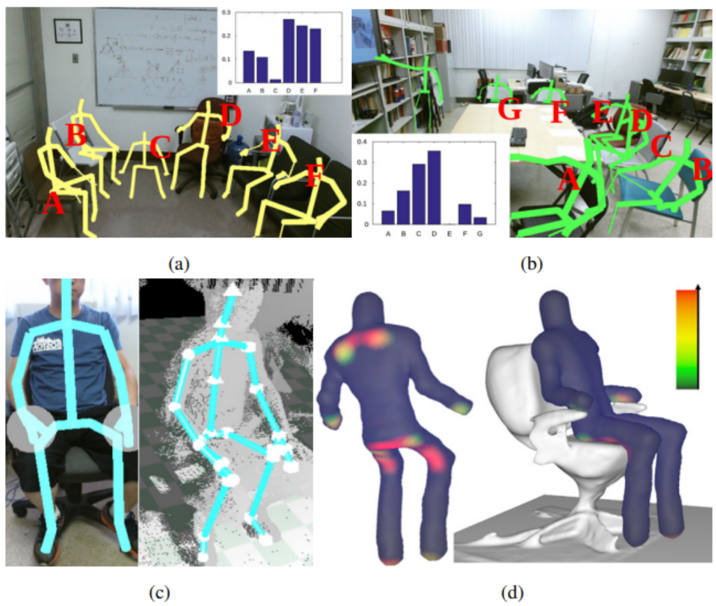

Figure 1: Examples of learning utility from video[3]

Computer vision is a good starting point in general. On one hand, there are many computer vision models and video datasets available as of 2023. On the other hand, current models have achieved a remarkable performance in efficiency. However, speaking to generalization, this kind of method may be poor. Because there are many ultilities that can not be perceived from merely vision.

3 Robotics Method

Using computer vision as the only method may lack generalization, but adding robotics as supplement seems to perform well in generalization. A work in 2017 shows that robot learns how to fold t-shirts from visual demonstrations, and proposes a plan when folding never-before-seen articles of clothing.[1]

This work presents a pioneering proof-of-concept that integrates computer vision with robotics to develop a comprehensive framework for representing and utilizing human utility. The proposed pipeline addresses the complex task of planning a cloth-folding action by enabling the agent to learn the external utility associated with human actions. The key innovation lies in the incorporation of a learned utility function that guides the agent's decision-making process.

This kind of combined methods outperform pure computer vision methods in terms of generalization. It is able to handle unseen missions with same setting, but it may perform poorly in a completely new setting. Also, its robotics learning can be further improved if this part is taken out and trained separately. Considering robotics training datasets are coming one after another, collecting training data will soon become easier.

4 Method Based on Large Language Model

Besides computer vision and robotics, natural language as another modal can play an important role in utility representation. As large language models(LLMs) like ChatGPT and Llama2 show outstanding performance on general understanding task, it is natural to apply LLMs in utility understanding.

Using NLP method to tackle utility understanding task may not suffer from data collection, due to many available text materials such as scripts can be used directly in training. In scripts, utilities and actions are explicitly noted to help actors to better act as the character, which makes it very high-quality data. As a result, constructing a dataset for this task may not be of great difficulty.

To be optimistic, powerful LLMs are able to do some few-shot or even zero-shot learning on various tasks, so it is possible that LLMs can perform well on utility tasks with high efficiency. But this idea need more evidence to validate.

But the method based on LLMs will adopt drawbacks from LLMs as well, such as interpretality problem.

5 Embodied and Multimodal Method

In fact, the methods mentioned above are all trying to model human utility without understanding why. This will be a theoretical bottleneck and ceiling of the overall performance.

Utility is a measurement of satisfaction. If an embodied and multimodal agent can perceive the enviornment around it via many kind of sensors, and the agent has its own reward system, it will have chance to understand what the concept of comfortable and utility means to itself, instead of just being told.

Based on our current understanding of biological evolution, many neurobiological developments and consciousness-related changes in ancient organisms were shaped by lifestyle and natural selection. Even in organisms with simple structures and lifestyles, their behavioral patterns are to some extent governed by utility. Therefore, simplifying the inquiry to comparatively basic issues related to embodied intelligence and reinforcement learning may offer new insights and understanding.

However, this novel method suffers from the lack of data. At the same time, reinforcement learning is complicated and does not guarantee to output a stable and good enough policy.

6 Conclusion

As an invisible feature and dark force[4] behind actions, utility is hard to represent and learn in general. In this essay, we have discussed several possible methods to learn and represent human utility. From computer vision, robotics to large language models and embodied agents, these methods have different levels of performance and potentiality. We should step further in model and data together in order to work out a more powerful solution of utility.

References

[1] Nishant Shukla, Yunzhong He, Frank Chen, and Song-Chun Zhu. Learning human utility from video demonstrations for deductive planning in robotics. In Sergey Levine, Vincent Vanhoucke, and Ken Goldberg, editors, *Proceedings of the 1st Annual Conference on Robot Learning*, volume 78 of *Proceedings of Machine Learning Research*, pages 448–457. PMLR, 13–15 Nov 2017. URL `https://proceedings.mlr.press/v78/shukla17a.html`. 2

[2] Wikipedia contributors. Utility — Wikipedia, the free encyclopedia, 2023. URL `https://en.wikipedia.org/w/index.php?title=Utility&oldid=1185181934`. [Online; accessed 3-December-2023]. 1

[3] Yixin Zhu, Chenfanfu Jiang, Yibiao Zhao, Demetri Terzopoulos, and Song-Chun Zhu. Inferring forces and learning human utilities from videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 1, 2

[4] 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