Enhanced imitation learning algorithms using human gaze data
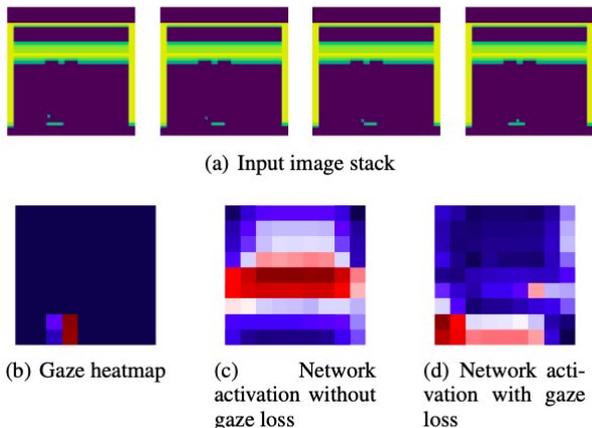
17 March 2020, by Ingrid Fadelli

(a) Input image stack

(b) Gaze heatmap (c) Network activation without gaze loss (d) Network activation with gaze loss

Input image stack fed to the algorithms. Credit: Saran et al.

Past psychology studies suggest that the human gaze can encode the intentions of humans as they perform everyday tasks, such as making a sandwich or a hot beverage. Similarly, human gaze has been found to enhance the performance of imitation learning methods, which allow robots to learn how to complete tasks by imitating human demonstrators.

Inspired by these previous findings, researchers at the University of Texas at Austin and Tufts University have recently devised a novel strategy to enhance imitation learning algorithms using human gaze-related data. The method they developed, outlined in a paper pre-published on arXiv, uses the gaze of a human demonstrator to direct the attention of imitation learning algorithms toward areas that they believe are important, based on the fact that human users attended to them.

"Deep-learning algorithms have to learn to identify

important features in visual scenes, for instance, a video game character or an enemy, while also learning how to use those features for decision making," Prof. Scott Niekum from the University of Texas at Austin told TechXplore. "Our approach makes this easier, using the human's gaze as a cue that indicates which visual elements of the scene are most important for decision-making."

The approach devised by the researchers entails the use of human gaze-related information as guidance, directing the attention of a deep learning model to particularly important features in the data it is analyzing. This gaze-related guidance is encoded in the loss function applied to deep learning models during training.

"Prior research exploring the use of gaze data to enhance imitation learning approaches typically integrated gaze data by training algorithms with more learnable parameters, making the learning computationally expensive and requiring gaze information at both train and test time," Akanksha Saran, a Ph.D. student at the University of Texas at Austin who was involved in the study, told TechXplore. "We wanted to explore alternative avenues to easily augment existing imitation learning approaches with human gaze data, without increasing learnable parameters."

The strategy developed by Niekum, Saran and their colleagues can be applied to most existing convolutional neural network (CNN)-based architectures. Using an auxiliary gaze loss component that guides the architectures toward more effective policies, their approach can ultimately enhance the performance of a variety of deep-learning algorithms.

The new approach has several advantages over other strategies that use gaze-related data to guide deep learning models. The two most notable ones are that it does not require access to gaze data at test time and the addition of supplementary

learnable parameters.

The researchers evaluated their approach in a series of experiments, using it to enhance different deep learning architectures and then testing their performance on the Atari games. They found that it significantly improved the performance of three different imitation learning algorithms, outperforming a baseline method that uses human gaze data. Moreover, the researchers' approach matched the performance of another strategy that uses gaze-related data both during training and at test time, but that entails increasing the number of learnable parameters.

"Our findings suggest that the benefits of some previously proposed approaches come from an increase in the number of learnable parameters themselves, not from the use of gaze data alone," Saran said. "Our method shows comparable improvements without adding parameters to existing imitation learning techniques."

While conducting their experiments, the researchers also observed that the motion of objects in a given scene alone does not fully explain the information encoded by gaze. In the future, the strategy they developed could be used to enhance the performance of imitation learning algorithms on a variety of different tasks. The researchers hope that their work will also inform further studies aimed at using human gaze-related data to advance computational techniques.

"While our method reduces computational needs during test time, it requires tuning of hyperparameters during training to get good performance," Saran said. "Alleviating this burden during training by encoding other intuitions of human gaze behavior will be an aspect of future work."

The approach developed by Saran and her colleagues has so far proved to be highly promising, yet there are several ways in which it could be improved further. For instance, it does not currently model all aspects of human gaze-related data that could be beneficial for imitation learning applications. The researchers hope to focus on some of these other aspects in their future studies.

"Finally, temporal connections of gaze and action have not yet been explored and might be critical to attain more benefits in performance," Saran said. "We are also working on using other cues from human teachers to enhance imitation learning, such as human audio accompanying demonstrations."

More information: Efficiently guiding imitation learning algorithms with human gaze. arXiv: 2002.12500 [cs.LG]. arxiv.org/abs/2002.12500

APA citation: Enhanced imitation learning algorithms using human gaze data (2020, March 17) retrieved 6 May 2021 from https://techxplore.com/news/2020-03-imitation-algorithms-human.html