Docket #: S21-254

IQ-Learn: State-of-the-Art Imitation Learning for AI

Researchers at Stanford have developed an imitation learning method, IQ-Learn, shown to surpass existing methods in some applications. Imitation learning is an AI process of learning by observing an expert, and has been recognized as a powerful approach for sequential decision-making, with diverse applications like healthcare, autonomous driving and complex game playing. However, conventional imitation learning methodologies often utilize behavioral cloning, which has advantages of simplicity and stability, but fails to recognize any information involving an environment's dynamics. Conventional methods that do exploit dynamics information tend to be difficult to train in practice due to an adversarial optimization process over reward and policy approximators. To address these deficiencies, the researchers have introduced a method for dynamics-aware learning which avoids adversarial training by learning a single Q-function, implicitly representing both reward and policy. Inverse soft-Q learning (IQ-Learn) obtains state-of-the-art results in offline and online imitation learning settings, surpassing existing methods both in the number of required environment interactions and scalability in high-dimensional spaces.

Stage of Development
Proof of concept

Applications

- AI and robotics
- Autonomous driving

Advantages

- Unlike previous methods, the approach converge in a small number of steps recovering the optimal reward and agent policy
- Uses simple optimization and is easy to train
- Scales to high-dimensional inputs like images, enabling human-like gameplay on video games using video demonstrations of humans/experts
- State-of-the-art in imitating experts without requiring interactions with a simulator or the real world, enabling learning just from passive observations of experts
- Works with visual expert demonstrations of car driving or robotic simulation environments, successfully imitating the experts and reaching their level of performance
- Recovers learned rewards that show a high positive correlation with the ground-truth environment rewards, leading to the interpretability of learned behavior

Publications

- Garg, Divyansh, et al. "IQ-Learn: Inverse soft-Q Learning for Imitation." *arXiv preprint arXiv:2106.12142* (2021).
- Nikki Goth Itoi. Training Smarter Bots for the Real World *Stanford HAI News* (2022).

Patents

- Published Application: 20230045360
- Issued: 12,737,675 (USA)

Innovators

- Divyansh Garg

Licensing Contact

Imelda Oropeza

Senior Licensing Manager, Physical Sciences

Email