Fido: A Universal Robot Control System Using Reinforcement Learning with Limited Feedback

Truell, Michael

Gruenstein, Joshua

A robot control system was developed that could be taught tasks through reinforcement learning. The system, nicknamed “Fido”, was designed to be universal regardless of inputs and outputs, robot kinematics, and processing capability. In addition, Fido was built to learn with limited feedback, allowing humans to train Fido in a minimal amount of time. This was achieved through the training of artificial neural networks with a wire-fitted interpolator following the Q-learning reinforcement learning algorithm and an intelligent action selection policy that utilizes a probabilistic approach to exploration. Functionality was first tested and evaluated in simulation. Next, hardware implementations of differing kinematics, sensors, and central processors were constructed. Fido successfully converged on all given tasks in simulation and in hardware within very few reward iterations while maintaining impressively low latency, demonstrating its potential as a comprehensive robot control system.

Awards Won:

Second Award of $2,000

European Organization for Nuclear Research-CERN: Second Award of $1,500