APPROVED FOR PUBLIC RELEASE; DISTRIBUTION IS UNLIMITED

Deep Learning

S&T Campaign: Computational Sciences
Data Intensive Sciences

David Doria, (410) 278-2310
david.l.doria.civ@mail.mil

Research Objective

- Develop adaptive algorithms to enable autonomous systems to handle the dynamic and unpredictable battlefield environment
- Improve state-of-the-art machine learning techniques to automatically learn action policies in video games

Autonomous agents on the battlefield of the future will require real-time analysis of distributed arrays of heterogeneous sensors to inform strategy recommendation. Video games are an excellent proxy to develop these techniques in a controlled environment.

By taking actions and observing their consequences and the response of the system, the algorithm can learn optimal policies.

Army Battlefield Strategy	Video Game
1) Collect data from distributed sensors	1) Collect pixels from game output
2) Analyze data	2) Analyze data
3) Develop/decide on strategy	3) Develop/decide on strategy
4) Take action (troop movement, etc.)	4) Take action (move Mario, etc.)

At a high level, video games are an excellent proxy for complex learning and decision making problems.

Challenges

- Significantly delayed reward signals are difficult for existing reinforcement learning techniques
- Model-free learning systems cannot distinguish between the agents contribution to state transitions and system dynamics
- "Black box" nature of neural network makes it hard to understand representations that are learned

ARL Facilities and Capabilities Available to Support Collaborative Research

- Training deep networks is very computationally intensive - 10 days of training on 2496 GPU cores for simple games
- Separated learning representations from learning action policies for significant decreases in retraining time
- Reward-biased sampling for experience replay methods
- Repetition inhibition during training to avoid excessively repeated states in replay memory

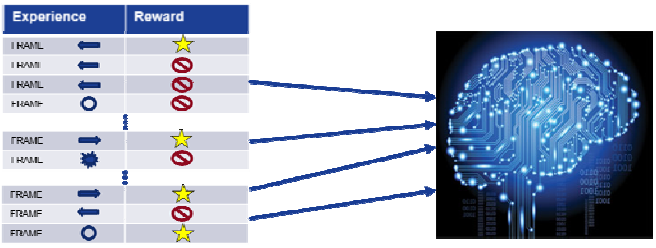

Experience reply takes advantage of previous state transitions multiple times, reducing the amount of observations necessary.

Complementary Expertise/ Facilities/ Capabilities Sought in Collaboration

- Reinforcement learning expertise
- Neural network expertise

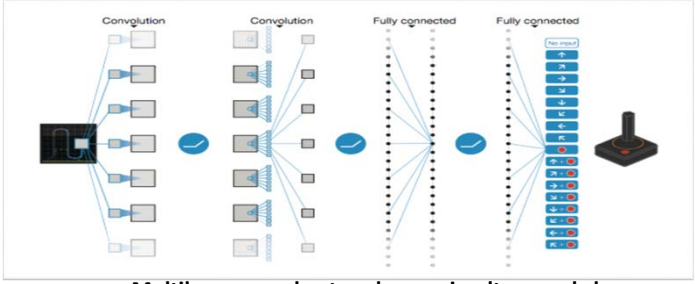

Multilayer neural networks can simultaneously learn representations and action policies

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION IS UNLIMITED