Evento	Salão UFRGS 2018: SIC - XXX SALÃO DE INICIAÇÃO CIENTÍFICA DA UFRGS
Ano	2018
Local	Campus do Vale - UFRGS
Título	Blind Agent Playing Games through Deep Reinforcement Learning
Autor	LEONARDO BELTRÃO DUARTE
Orientador	LUIS DA CUNHA LAMB

Blind Agent Playing Games through Deep Reinforcement Learning

Aluno: Leonardo B. Duarte Orientador: Luís C. Lamb

Instituto de Informática, Universidade Federal do Rio Grande do Sul

The growth of computer processing power allied with a huge amount of data generated in the recent years allowed algorithms based on machine learning to succeed in tasks of great computational complexity such as language translation, text generation and image classification. The main technique used to achieve these results is called deep learning, which consists in creating high level abstractions from data, inspired by information processing and communication patterns in biological nervous systems. In order to build a general artificial intelligence algorithm, that can succeed in multiple scenarios (like the human brain), a slightly different idea is used. This modern approach involves using deep learning in a reinforcement learning context, which is now called deep reinforcement learning. Reinforcement learning is a machine learning paradigm in which the agent (algorithm) interacts with the environment trying to maximize a numerical signal.

This work attempts to implement existing deep reinforcement learning algorithms capable of playing classic games at a human-level solely with video input and modifying it so it can play using only the audio from the games as input. Our motivation comes from the fact that human players have shown to be able to obtain good scores in certain games (i.e. Street Fighter) while blindfolded and receiving only audio input. The idea is to create a high-level representation of the game screen, without the need of a high-dimensional input such as the screen pixels. To achieve this, we use a deep reinforcement learning idea: applying artificial neural networks as a function approximator in the context of reinforcement learning. Molding this work in the reinforcement learning framework we have the player as the agent, the game as the environment and the numerical signal as the score.

The process is the following: first we feed the neural network with the very first frame, creating a basis for what the network is supposed to see. Next, we feed the audio and hope that the network will learn to adapt its current state as the audio signal changes. We then use this approximate representation of the screen as the input to a network that will actively try to master the game.

In order to create an intelligent agent capable of thriving in different kinds of games with the same parameters, a reliable form of evaluation was needed. The Arcade Learning Environment, a platform widely used in the field, offers different classic games that we can use to easily benchmark our agents. Using this platform, we can graph the scores and see how the same agent is performing in different games. One drawback is that this framework is limited to games for the console Atari 2600, released in 1977. Part of this work involves extending these smooth transitions between games into consoles with more capabilities, so that we can extract more information from the game audio, such as the Game Boy Advance, released in 2001.

The current results suggest that the network may not converge into a good approximation. Even with the screen as input, the network suffers with a lot of instability due to the computing power limitations when applying state-of-the-art techniques. These techniques are used in order to eliminate the high correlation between subsequent game states but rely on a big amount of computer memory. The network may need simpler environments to learn a useful state representation.

Agradecemos ao doutorando Marcelo Prates pela colaboração nesta pesquisa.