



Procedural Content Generation for Reinforcement Learning

Niels Justesen

AI researcher, Modl.ai Denmark

Abstract:

Reinforcement Learning (RL) has achieved marvellous results in various domains in the last

decade, e.g. in video games [1]. Most of the studied games are, however, structurally fixed and do not require agents to generalize to variations of the environment. We argue that testing RL models on the training environment is similar to testing supervised learning models on the training set. While such a methodology goes against the traditions in machine learning, it has become the standard practice in RL. To scale RL to dynamic real world problems, and even modern video games, agents must be able to generalize to environmental variations. At modl.ai, we incorporate Procedural Content Generation (PCG) into the RL training loop to learn a general policy across a variety of video game levels. Commercial video games typically have a small corpus of levels incapable of facilitating training of agents that can generalise. This makes PCG for RL an effective approach in this domain, while it increases the complexity of the task. Our research has shown that a progressive adaptation of the level difficulty allows RL to learn general policies in games that otherwise were infeasible to learn in fixed versions of the games [2]. PCG for RL can be applied by manually constructing a generator that produce representative levels. This is time consuming and is often quite difficult. Procedural Content Generation via Machine Learning (PCGML) [3] offers a new toolset to overcome such issues. Generative machine learning can be applied to model the distribution of existing video game levels to generate completely new levels [4]. While PCGML has several shortcomings, we use these techniques to generate novel levels of high quality for commercial video games. In addition, we employ such generators in RL and train

agents to do automatic testing of new game content in previously impossible ways.



Biography:

Niels is an AI researcher at modl.ai in Copenhagen, working on game-playing and content-generating algorithms. He did his Ph.D. on algorithms for adaptive game-playing agents at the IT University of Copenhagen where he was advised by Prof. Sebastian Risi and was part of the Center for Computer Games Research and the Robotics, Evolution, and Art Lab. Niels also has four years of experience as a software developer where he worked on the development of several intelligent systems for the public sector in Denmark. Niels is also running an AI competition called Bot Bowl, where bots compete in the board game Blood Bowl which is promoted to be the next board game challenge for AI.

Publication of speakers:

1. Deep learning for video game playing, Niels Justesen, 2019
2. Illuminating generalization in deep reinforcement learning through procedural level generation, Niels Justesen, 2018
3. Procedural content generation via machine learning (PCGML), Niels Justesen, 2018
4. Evolving Mario levels in the latent space of a deep convolutional generative adversarial network, Niels Justesen, 2018

8TH International conference on Robotics and Artificial Intelligence, August 10-11, 2020 London, UK

Citation: Niels Justesen, Procedural Content Generation for Reinforcement Learning, Euro Robotics 2020, August 10-11, 2020 London, UK