Proceedings of the Canadian Society for Mechanical Engineering International Congress 2022
CSME Congress 2022
June 5-8, 2022, Edmonton, AB, Canada

Deep Reinforcement Learning for Emission Control in Diesel Engines

David Gordon[1], Armin Norouzi[1*], Saeid Shahpouri[1], Mahdi Shahbakhti[1], Charles Robert Koch[1]
[1]Mechanical Engineering Department, University of Alberta, Edmonton, Canada
*norouziy@ualberta.ca

Abstract

Heavy-duty and medium-duty diesel engines power millions of vehicles that move goods around the globe. These engines have long provided a reliable method to convert chemical energy into useful rotational motion. The high efficiency, long life and durability of the diesel engine has resulted in its wide use in many of today's transportation systems. To balance engine performance, power output, fuel economy, and longevity, traditional engines are extensively tested where two-dimensional look-up tables, also known as calibration maps, are created and used for feed-forward engine control. These calibration maps are also assessed for engine emissions, making the process complex and time intensive. Using a model-based optimum controller is one possible solution to reduce the ever-increasing calibration effort. However, the requirement of ever increasing model accuracy leads to complex nonlinear models that require a complex control law for implementation. Machine Learning (ML) has been used to solve a variety of engineering challenges and has been shown to be especially useful in control engineering, where identifying the correct model and corresponding controller of a system is challenging.

The objective of this study is to develop a model-free deep reinforcement learning (RL) application to lower the diesel engine's NOx emissions while decreasing the fuel consumption while tracking the desired engine-out load. Using a physically-based Engine Simulation Model (ESM), in GT-power, a deep deterministic policy gradient (DDPG) is developed. The ESM is used to train and evaluate the RL agent. To reduce the possibility of an unsafe control output, the deep RL is enhanced using a safety filter. This filter is shown to have no influence on the output of the fully trained agent since deep RL is capable of learning output constraints; nonetheless, it is critical to enforce these limits during training where major constraint violations may occur. This safe RL is then compared to another model-free controller Iterative Learning Controller (ILC) and a model-based Nonlinear model predictive controller (NMPC). This comparison demonstrates that both safe RL and ILC perform similarly to the NMPC after sufficient engine cycles where they are able to learn the combustion process. The advantage RL has is any arbitrary reference can be tracked while ILC requires a repetitive reference. The comparison of RL and NMPC has demonstrated that for this application RL is capable of learning the optimum control action directly from the experiment without the requirement of a model.

Word count: 392