Abstract

Climate change is a legitimate concern in the world today and, through the increased occurrence and intensity of extreme weather and a string of record high temperatures, is a concern that we are already feeling the effects of. The biggest culprit of this phenomenon is the vast amounts of CO2 that are being pumped into the air primarily due to the carbon-intensive energy generation ubiquitous in the world today. Renewable energy supply (RES) is a key com- ponent to achieve sustainable clean energy production by shifting away from energy generation with harmful by-products and towards solar, wind, or hydro energy. However, currently, the high implementation and maintenance cost, unpredictability in energy generation, and the lack of an established system to manage the generation effectively when it occurs are proving to be the barriers to fully adopting RES. To overcome these obstacles it is important to have ef- ficient and effective energy management systems in place. Solutions are being developed to optimise devices such as PV panels and wind turbines to curtail their awkward generation schedule and efficiently store any generated energy for later use. The solutions are also being adopted on a large scale, across a group of RES generation devices and energy stores and loads called microgrids. Currently a machine learning model known as Reinforcement Learning (RL) is proving successful in optimising the management system for a microgrid. In RL, an agent (the microgrid) learns how to perform a task (energy management) by interacting with its environment (consumption/production profiles and devices connected to the grid). The issue, however, is that these RL models need a lot of time for training to find the optimal policy as each agent is being trained from scratch, (i.e., with no prior knowledge). Transfer learning (TL) is an approach towards reducing this training delay.

TL is a relatively new phenomenon where a pre-trained model, or a partially pre-trained model, is used to speed up the training of another model that is learning a separate task. The pre-trained model is referred to as the source model and the model to be trained is referred to as the target model. In this paper, I investigate the application of a TL approach on a Deep RL model [?] that looks to optimise the management of energy in a microgrid. The model contains production and consumption profiles from a residential customer in Belgium which are used as the training samples. The goal of the energy system is to find the best policy to handle the demand and supply of energy in the microgrid by managing a short-term and long-term storage connected to the grid. The TL approach used in this paper is known as Weight Initialisation and is implemented by taking the weights produced by the source model, after it has been trained, and using those values to initialise the weights for the target model. The scenarios tested were source models generated from different configurations of the microgrid (Small, Medium, Large) and used to initialise the weights of a different configuration. This took the approach a step further to analyse the value of the knowledge obtained by one configuration of microgrid compared to another. This paper looks to determine the effect of TL on a Deep RL model for energy management and to also investigate whether the knowledge learned from a smaller or larger microgrid is more useful to a different microgrid.