Online parameter control for carbon-aware scheduling using deep reinforcement learning

Andrea Mencaroni

The transition to net-zero greenhouse gas (GHG) emissions presents a major challenge for the manufacturing sector, which is heavily reliant on electricity. Throughout the day, grid electricity shows a variable carbon intensity due to shifts in the energy mix used for its production, while on-site renewables, despite being cleaner, have limited and variable availability. To achieve net-zero emissions, it is crucial to dynamically adjust electricity usage accordingly. Carbon-aware scheduling addresses this challenge by aligning electricity consumption with both grid carbon intensity and renewable availability to minimize emissions.

While metaheuristics have been used to solve this complex scheduling problem, their performance is highly sensitive to parameter settings, often requiring frequent re-tuning to adapt to daily variability in energy data. In this study, we use deep reinforcement learning (DRL) to dynamically control the parameters of a memetic algorithm for carbon-aware scheduling. Rather than relying on pre-tuning, a DRL agent is trained offline to learn an optimal policy, adjusting parameters in real-time during the algorithm's execution.

This approach not only accelerates convergence to superior solutions but also eliminates the need for manual parameter (re-)tuning, enabling rapid deployment on new instances and optimizing parameter values at each iteration.