

LLM-ENHANCED DYNAMIC FLEET PLANNING WITH HIERARCHICAL MULTI-AGENT REINFORCEMENT LEARNING FRAMEWORK

Lijing LIU⁽¹⁾, James M Shihua⁽¹⁾, Qiyu Yan⁽²⁾, Rhea P. Liem⁽³⁾

⁽¹⁾The Hong Kong University of Science and Technology, Department of Mechanical and Aerospace Engineering

⁽²⁾China Unicom Global, Department of Cloud & Network Operation

⁽³⁾Imperial College London, Department of Aeronautics

Keywords: Fleet planning, Hierarchical multi-agent reinforcement learning, LLM-based hyperparameter tuning

Abstract

Civil aviation fleet planning is a typical long-horizon operations research (OR) problem, for which the problem formulation and algorithms could play critical roles in either the computational efficiency or the solution quality. However, conventional meta-heuristic approaches have some inherent drawbacks: during the optimization process, the model is unable to “learn” efficiently from infeasible solutions and gradually concentrate on high feasibility regions, leading to an exponential increase in complexity with increasing planning horizon, and their randomness prevents an optimality guarantee on individual instances. In this paper, we introduce a hierarchical multi-agent proximal policy optimization (PPO) framework to tackle these problems. To effectively handle ultra-long episodes, our approach builds up the optimization architecture by reward shaping and environment designing within a base PPO, which is then boosted through multi-agent coordination and hierarchical decomposition. Large language model (LLM)-based agent workflows are deployed to automatically tune hyperparameters toward the specific domains’ performances. Validated on realistic flight data from 2019 shared by our airline partner, Cathay Pacific Airways Limited (CX), our results show that the proposed method can reduce total airlines’ operational costs—including direct operating cost (DOC) and capital cost (CC) and achieves a computation speedup in comparison with a conventional optimization baseline.