Abstract

Road transport is a major source of greenhouse-gas emissions, and the European Union legislation framework raises fleet CO_2 targets on the way to a 2035 zero-emission sales target, pushing both for vehicle efficiency and to find fleet-level solutions. This dissertation is focused on developing efficient controls and models for the automotive and mobility sectors, employing machine learning at three linked scales: control of urban fleets and hybrid vehicle energy management, and modeling the combustion process of internal combustion engines. At the fleet scale, Autonomous Mobility on Demand is posed as a network-flow control problem and addressed with a three-step framework, optimal dispatch, minimum-cost rebalancing, and a graph-aware Soft Actor–Critic policy that predicts region-level vehicle allocations. In an in-the-loop mesoscopic SUMO model of Luxembourg, this maximizes the operational profit for the operator, while reducing waiting time for passengers and fleet-level emissions, and remains robust across time of day, spatial aggregation, and simulator fidelity. At the vehicle scale, learning-based energy management systems are developed for a state-of-the-art plug-in hybrid electric vehicle and benchmarked against optimization strategies, such as dynamic programming and equivalent consumption minimization strategy. A dynamic programming-trained LSTM and an off-policy Soft Actor-Critic are trained on a backward model and validated on a detailed vehicle virtual test rig in both charge sustaining and charge depleting operation; they deliver competitive fuel/CO_2 results, respect SoC constraints, and generalize across cycles and initial conditions. At the combustion scale, two real-time data-driven models are introduced: a neural single-Wiebe parameterization for burn rate prediction, and a hybrid recurrent model that reconstructs the full burn rate profile and stays stable across speed, load, and dilution sweeps. Together, these results show that embedding domain knowledge in learning methods enables scalable, real-time decisions, from fleet rebalancing to power split and combustion prediction, supporting near-term emission cuts while meeting tightening policy targets.