Abstract

The growing demand for electricity, coupled with the increasing integration of renewable energy sources (RES), has created a critical need for more efficient energy management solutions. Demand Response (DR) plays a vital role in balancing supply and demand, particularly in the residential sector, where flexible energy consumption can help mitigate grid stress and optimize resource use. Home Energy Management Systems (HEMSs) have become central to facilitating residential DR by automating and optimizing household energy usage. Various methods have been employed to optimize HEMSs, including traditional optimization techniques, metaheuristic methods, Model Predictive Control, and robust and stochastic optimization approaches. While these methods offer valuable advantages, they also face significant limitations, such as handling uncertainty, high time complexity, modeling challenges, and limited adaptivity to real-time changes in energy conditions. In contrast, Deep Reinforcement Learning (DRL) has emerged as a promising solution that addresses many of these limitations by offering adaptive, scalable, and real-time decision-making capabilities in dynamic environments. DRL's ability to learn from interactions with the energy system allows it to handle uncertainties and changing conditions more effectively than traditional approaches. This dissertation further discusses different coordination paradigms for multi-HEMS operation using DRL, the integration of adaptive comfort preferences, and the application of federated learning (FL) as a scalable and privacy-preserving framework for coordinated peak-demand mitigation. Building on this federated formulation, the dissertation also addresses the problem of agent dropout, a practical challenge that arises when households become temporarily unavailable due to communication, computation, or scheduling constraints, and which biases the federated aggregation toward the surviving households. To mitigate this effect, a Confidence-Aware Gradient Substitution (CAGS) algorithm is proposed, combining an exponentially weighted moving average similarity tracker with an adaptive confidence gate that substitutes a dropped agent's update with that of a trusted neighbour only when the estimated similarity is sufficiently reliable. CAGS is validated

on standard federated supervised benchmarks (MNIST, CIFAR-10) and on a heterogeneous federated HEMS deployment, where it consistently outperforms existing dropout-handling baselines, particularly under aggressive dropout ratios.