Machine Learning-Informed Optimization for Modeling Flexible Energy Resources

Pietro Favaro

Decarbonization of power systems through large-scale renewable integration creates unprecedented grid management challenges. Variable wind and solar generation introduce significant uncertainty and volatility, requiring flexible resources to balance supply and demand in real time. Demand-side flexibility from building HVAC systems represent critical flexibility providers, but their optimal operation remains challenging. These assets exhibit complex nonlinear dynamics, such as heat pump efficiency curves and thermal dynamics, that conventional linear approximations inadequately represent. These misrepresentations result in suboptimal scheduling decisions. A fundamental tension exists in energy system decision-making. Constrained optimization frameworks, particularly convex formulations, provide formal guarantees on solution quality and global optimality through well-established theory and efficient solvers. However, this mathematical rigor requires linearization of system dynamics, sacrificing modeling fidelity. Conversely, machine learning excels at capturing nonlinear system behavior from data but lacks the decision-making structure and solution quality guarantees that optimization provides. This thesis bridges this critical gap by developing a unified framework that harnesses the nonlinear expressiveness of machine learning while preserving the solution quality guarantees and tractability of constrained optimization. We introduce neural network-constrained optimization, which embeds trained neural networks directly into mixed-integer optimization frameworks as exact piecewise linear reformulations. This approach enables precise modeling of unknown nonlinear thermal dynamics of buildings for day-ahead scheduling. Later, a multi-fidelity analysis demonstrates that incorporating nonlinearities through high-fidelity simulation significantly improves solution quality. However, fully nonlinear formulations render the optimization problem computationally intractable for practical implementation. To bridge this gap, we propose decision-focused learning: a framework that learns asset representations optimized for decision quality rather than statistical accuracy (e.g., mean squared error minimization) on historical data. We first present a methodology for learning linear constraint parameters within convex optimization programs, directly aligning model learning with the downstream optimization objective. We subsequently extend this framework to non-differentiable programs (e.g., mixed-integer linear programs) using stochastic smoothing and score function gradient estimation (i.e., REINFORCE algorithm), enabling learning in inherently non-differentiable settings. This is particularly valuable for physical systems (or high-fidelity simulators of such systems) that are not differentiable. Crucially, this extended framework allows neural network parameters in neural network-constrained optimization to be learned end-to-end, accounting for true system nonlinearities while maintaining computational tractability.

Case studies on fictitious multi-zone office buildings achieve sixfold reductions in cost prediction errors. These results validate that machine learning informed optimization and decision-focused learning enable computationally tractable, and physically consistent decision-making strategies essential for grid flexibility in the renewable energy transition.