

D2.1. Identify the most relevant supervised learning models able to handle the complex (non-linear, high-dimensional with discrete and continuous variables) characteristics of the problem

Authors: Jean-François Toubeau, Hooman Khaloie, François Vallée

This deliverable is based on the following publication:

Khaloie, H., Dolanyi, M., Toubeau, J. F., & Vallée, F. (2025). Review of machine learning techniques for optimal power flow. Applied Energy, 388, 125637.



Table of Contents

| | |
|--------------------------------------|---|
| Table of Contents | 2 |
| 1. Purpose within the project | 3 |
| 2. Methodological contribution | 3 |
| 2.1. End-to-end learning | 4 |
| 2.2. Learning-to-Optimize | 4 |
| 3. Main results | 5 |

1. Purpose within the project

This paper was carried out to identify the most relevant supervised learning models capable of handling the complex characteristics of power-system optimisation problems. Future transmission-system operation increasingly requires solving Optimal Power Flow (OPF) and related optimisation problems under stringent computational constraints, while accounting for nonlinear power-flow relationships, large numbers of decision variables, network constraints, renewable uncertainty, and discrete operational decisions.

Conventional optimisation methods provide high-quality solutions but often become computationally demanding when applied repeatedly in real-time operation. The objective of this task was therefore to assess the suitability of modern supervised learning approaches for approximating optimisation outcomes and accelerating operational decision-making while maintaining acceptable levels of optimality and feasibility.

The work was motivated by the observation that power-system digitalisation is generating large quantities of operational data from SCADA systems, PMUs, smart meters, and optimisation tools. These datasets create opportunities to transfer part of the computational effort from online optimisation to offline training. Instead of repeatedly solving computationally intensive optimisation problems, supervised learning models can be trained to learn the relationship between system conditions and optimal operational decisions. Once trained, these models can provide solutions within milliseconds, making them attractive candidates for near-real-time applications.

| Aspect | Summary |
|----------------------|--|
| Main objective | Identify supervised learning models suitable for OPF and related optimisation problems |
| Methodological focus | End-to-End supervised learning, constrained learning, graph-based learning |
| Operational value | Faster decision-making while maintaining acceptable optimality and feasibility |

2. Methodological contribution

The task first established a classification framework for supervised learning approaches applied to optimisation problems. The review distinguishes between End-to-End learning schemes, where a machine-learning model directly learns the mapping between system inputs and optimisation outputs, and Learning-to-Optimize, where machine-learning algorithms assist conventional optimisation solvers by accelerating specific stages of the optimisation process, such as identifying active constraints, generating high-quality warm-start solutions, or supporting distributed optimisation procedures.

While End-to-End approaches seek to replace the optimisation process with a learned surrogate model, Learning-to-Optimize methods preserve the underlying optimisation framework and use machine learning to improve computational efficiency and scalability. This distinction provides a useful basis for assessing the suitability of different supervised learning models for complex power-system optimisation problems.

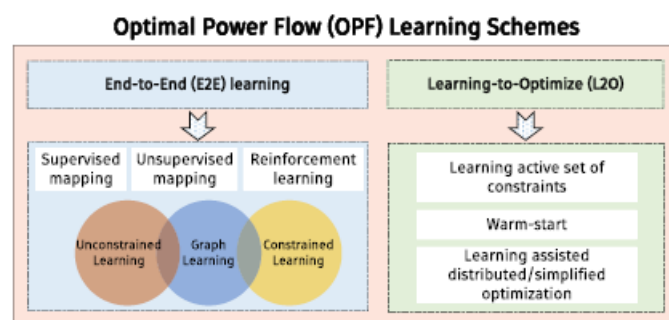


Figure. Overview of OPF learning schemes.

2.1. End-to-end learning

In the end-to-end framework, historical optimisation results are used as labelled training data. Input variables typically include load levels, renewable generation forecasts, network topology information, and operational conditions, while outputs correspond to optimisation decisions such as generation dispatch, bus voltages, power flows, or operating costs. The resulting supervised learning model effectively acts as a surrogate optimiser capable of producing decisions without explicitly solving the underlying optimisation problem.




Figure – End-to-End supervised learning framework for optimisation problems

A first category identified in the review consists of unconstrained supervised learning models. These methods focus primarily on prediction accuracy and learn direct mappings between system states and optimisation results. Neural networks represent the dominant approach within this category due to their ability to approximate highly nonlinear relationships and scale to large datasets. Several studies demonstrate that deep neural networks can accurately predict generator dispatch, voltage profiles, and other optimisation variables across a wide range of operating conditions. Random forests, decision trees, and other regression models have also been investigated, particularly for medium-sized systems where interpretability remains important. The results indicate that neural-network-based approaches generally provide the best balance between scalability and predictive accuracy for large optimisation problems.

A second category consists of constrained supervised learning models. Unlike unconstrained approaches, these methods incorporate physical and operational constraints directly into the learning process. The review identifies several mechanisms for achieving this objective, including penalty-based formulations, Lagrangian duality approaches, and differentiable optimisation layers. By embedding feasibility considerations during training, constrained learning models significantly reduce the risk of producing operationally infeasible solutions. This represents an important advancement because prediction accuracy alone is insufficient for operational deployment; solutions must also satisfy network security limits, generator capabilities, voltage constraints, and power-balance equations. The review shows that constrained learning frameworks consistently outperform unconstrained approaches in terms of feasibility while maintaining similar levels of prediction accuracy.

The third major methodological contribution concerns graph-based supervised learning. Power systems are naturally represented as graphs consisting of buses and transmission lines. Traditional neural networks do not explicitly exploit this topological structure, which can limit their ability to generalise across different network configurations. Graph Neural Networks (GNNs) address this limitation by embedding network connectivity directly into the learning architecture. The review demonstrates that graph-based models are particularly attractive for large-scale power systems and for applications involving changing network topologies, contingencies, or reconfiguration actions. By preserving topological information throughout the learning process, GNNs improve scalability, transferability, and robustness compared with conventional neural-network architectures.

2.2. Learning-to-Optimize

In parallel to End-to-End learning, the review identifies **Learning-to-Optimize (L2O)** as a second major family of supervised learning approaches for optimisation problems. Rather than replacing the optimisation process entirely, L2O methods aim to accelerate conventional optimisation algorithms by learning specific elements of the solution procedure. This hybrid paradigm seeks to combine the computational efficiency of machine learning with the reliability and physical consistency of optimisation-based methods.

The review classifies L2O approaches into three main categories.

The first consists of **active-set prediction methods**, where supervised learning models are trained to identify which operational constraints are likely to be binding at the optimum. Since only a limited subset of all network constraints is

typically active in a given operating condition, predicting these constraints enables the optimisation problem to be significantly reduced before being solved. The resulting reduction in problem size can substantially decrease computational effort while preserving solution quality.

A second category concerns **warm-start approaches**, where machine-learning models are trained to predict an initial solution that is close to the optimal operating point. These predicted operating points are subsequently provided as starting values to conventional optimisation solvers. Because iterative optimisation algorithms often spend a significant amount of time converging toward the optimum, high-quality warm-start solutions can reduce the number of iterations required and thereby accelerate the overall solution process. The review highlights that neural-network-based warm-start methods have demonstrated significant reductions in computational time while maintaining near-optimal solutions.

The third category comprises **learning-assisted distributed and simplified optimisation methods**. In these approaches, machine-learning models are integrated with existing optimisation frameworks to improve convergence behaviour, parameter selection, or coordination between decomposed subproblems. Such methods are particularly relevant for large-scale transmission systems where decentralised or distributed optimisation techniques are required to preserve scalability and data privacy. By supporting the optimisation process rather than replacing it, these approaches retain strong feasibility guarantees while benefiting from the adaptability of machine learning.

A key conclusion of the review is that Learning-to-Optimize methods often provide a more conservative but operationally robust pathway toward the adoption of machine learning in power-system optimisation. Whereas End-to-End approaches offer the highest potential computational gains, they may encounter challenges related to feasibility and generalisation. In contrast, L2O approaches preserve the underlying optimisation structure and therefore inherit many of the reliability properties of conventional optimisation algorithms. The review consequently identifies Learning-to-Optimize as a particularly promising framework for real-time operational applications where computational speed, solution quality, and security requirements must be simultaneously satisfied.

3. Main results

The comparative analysis performed in the task highlights several important findings.

First, the review confirms that **supervised learning provides a credible pathway for accelerating the solution of complex optimisation problems**, either by directly approximating optimisation outcomes through End-to-End learning or by enhancing conventional optimisation algorithms through Learning-to-Optimize frameworks. While both approaches offer substantial computational benefits, they differ in the trade-off they establish between speed, optimality, feasibility, and interpretability.

For End-to-End approaches, neural-network-based models consistently emerge as the most promising family of supervised learning techniques for high-dimensional optimisation problems because of their ability to capture highly nonlinear relationships between system conditions and optimal decisions. The review further shows that incorporating physical constraints directly into the learning process significantly improves solution feasibility and operational reliability, while graph-based architectures provide additional benefits for large-scale transmission systems by explicitly exploiting network topology and accommodating topology changes.

For Learning-to-Optimize approaches, the results demonstrate that machine learning can effectively support optimisation models without replacing them. Active-set prediction, warm-start generation, and learning-assisted optimisation all provide meaningful reductions in computational effort while preserving the underlying model structure. As a result, these methods often offer stronger guarantees regarding feasibility and operational security than purely End-to-End approaches, making them particularly attractive for real-time operational applications where reliability remains a primary concern.

Overall, the review indicates that End-to-End and Learning-to-Optimize approaches should be viewed as complementary rather than competing paradigms. End-to-End methods offer the largest potential computational gains by directly approximating optimisation decisions, whereas Learning-to-Optimize methods provide a more conservative pathway that retains the robustness and physical consistency of traditional optimisation frameworks. The most suitable approach therefore depends on the operational context, required computation speed, and acceptable level of approximation.

The task also identifies several challenges that remain relevant for future research. These include guaranteeing feasibility under all operating conditions, improving generalisation to rare or previously unseen system states, handling mixed



continuous-discrete decision spaces, and maintaining performance under changing network topologies. Additional challenges concern the integration of uncertainty information into learning-based optimisation frameworks and the development of models that remain reliable under increasingly variable renewable-generation conditions. These research directions directly motivate the subsequent work packages on reinforcement learning, uncertainty-aware optimisation, and risk-based operational decision support.

Overall, this paper contributes to the project by establishing the methodological foundations for supervised-learning-based optimisation. The review provides a structured assessment of both End-to-End and Learning-to-Optimize approaches and identifies the machine-learning techniques that are most suitable for addressing the nonlinear, high-dimensional, and constrained nature of power-system optimisation problems. By clarifying the strengths and limitations of the different supervised-learning paradigms, the task creates the foundation for the development of scalable, computationally efficient, and operationally reliable decision-support tools in the subsequent stages of the DISCRETE project.