Applied Energy 427 (2027) 128946

Contents lists available at ScienceDirect

Applied Energy

journal homepage: www.elsevier.com/locate/apen

Surrogate modeling of interconnector flows: A machine learning alternative to full-scale power system simulations with application to cross-border electricity exchange

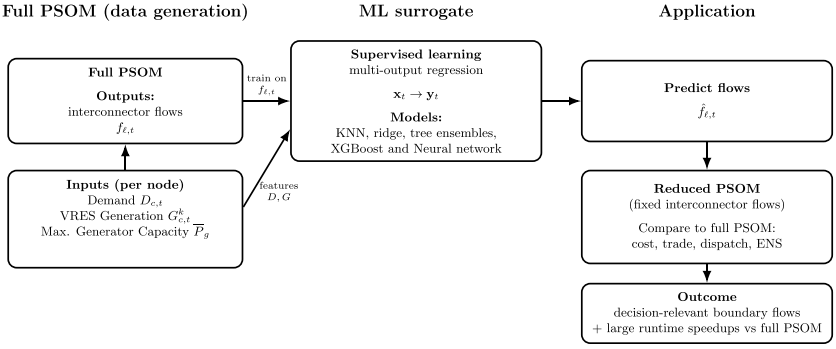
Robert Gaugl [a, b, *], Eloy Insunza [c], José Portela [c], Sonja Wogrin [a, b]

[a] Institute of Electricity Economics and Energy Innovation (IEE), Graz University of Technology, Inffeldgasse 18, Graz, 8010, Austria
[b] Research Center ENERGETIC, Graz University of Technology, Rechbauerstraße 12, Graz, 8010, Austria
[c] Institute for Research in Technology (IIT), ICAI School of Engineering, Universidad Pontificia Comillas, Alberto Aguilera 23, Madrid, 28015, Comunidad de Madrid, Spain

HIGHLIGHTS

- Core contribution: ML predicts PSOM-generated interconnector boundary flows.
- Neural networks generalize robustly relative to six full-feature ML baselines.
- A soft feasibility-aware loss discourages implausible net-export patterns.
- ML flows yield single-country PSOM results close to full European simulations.
- Large amortized runtime reductions enable repeated reduced-PSOM studies.

GRAPHICAL ABSTRACT



ARTICLE INFO

Keywords:
Surrogate modeling
Machine learning
Power system optimization
DC optimal power flow
Renewable integration
Energy system modeling

ABSTRACT

Cross-border electricity exchanges are crucial for operating and planning highly renewable power systems. Many studies reduce spatial granularity to keep models tractable and prescribe cross-border exchanges exogenously, often by reusing historical import/export time series. Such assumptions become inconsistent as renewable penetration changes the magnitude and timing of flows. The core contribution of this work is a machine-learning (ML) surrogate framework that maps nodal time-series data (e.g., hourly demand and renewable generation) to synthetic interconnector-level flow time series. The resulting flows serve as fixed boundary conditions in reduced power system optimization models (PSOMs). This addresses studies focused on a subset of an interconnected system, for which repeatedly solving the full model across multiple climate years or scenarios is computationally expensive. We demonstrate the framework on a pan-European single-node-per-country DC optimal power flow model. We benchmark k-nearest neighbors (KNN), ridge regression, random forests, Extra Trees, gradient boosting, XGBoost, and feedforward neural-network (SQU) surrogates. As an additional contribution, we introduce a novel custom neural-network loss that softly penalizes physically implausible flow patterns. This soft penalty is not a guarantee of full power-flow feasibility. Across benchmarks, the SQU variants provide the most consistent out-of-sample and downstream optimization performance while substantially outperforming scaled historical profiles. The ML-generated flows closely reproduce country-level results of the full European model while greatly

* Corresponding author at: Institute of Electricity Economics and Energy Innovation (IEE), Graz University of Technology, Inffeldgasse 18, Graz, 8010, Austria.
Email addresses: robert.gaugl@tugraz.at (R. Gaugl), einsunza@comillas.edu (E. Insunza), jportela@comillas.edu (J. Portela), wogrin@tugraz.at (S. Wogrin).

https://doi.org/10.1016/j.apenergy.2026.128946
Received 28 April 2026; Received in revised form 15 September 2026; Accepted 25 September 2026
Available online 29 September 2026
0306-2619/© 2026 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

reducing computation time (up to ~500×). The proposed framework enables scenario-consistent reduced PSOM studies across climate years after a single full-model run, eliminating repeated pan-European simulations and supporting more robust climate-year analyses under fixed system assumptions.

1. Introduction

1.1. Motivation

The European power system is characterized by a high degree of interconnectivity among national electricity grids, making cross-border power flows a critical element of both system operation and long-term planning. In future power systems with high shares of variable renewable energy sources (VRES), cross-border electricity trading will become even more essential for balancing fluctuations in VRES generation across regions. Recognizing this, the European Union (EU) has introduced regulatory measures mandating that at least 70 % of cross-border transmission capacity be made available for electricity trading, as outlined in the Clean Energy for All Europeans package [1].

Power system optimization models (PSOMs) are widely used to analyze such systems and guide infrastructure and policy decisions. However, PSOMs face inherent trade-offs due to computational limitations, particularly across three key dimensions: (1) temporal resolution, (2) technical detail, and (3) spatial scope. Even with access to modern high-performance computing, it remains challenging to fully capture all three dimensions simultaneously. Consequently, many studies restrict the spatial domain to a single country or region to preserve detail in other aspects. Yet, electricity imports and exports are crucial for capturing cross-border dynamics. Therefore, broader spatial modeling is often needed to ensure system-wide interactions are adequately represented, especially those involving international electricity trade.

To reduce the computational burden of large-scale spatial PSOMs, cross-border exchanges are frequently treated exogenously, for example by prescribing historical or otherwise fixed import and export time series. This type of simplification is common in geographically restricted modeling studies, where cross-border trade is represented in a stylized way to retain tractability, but it can materially affect results as shown by Mertens et al. [2]. However, as the European power system transitions toward a higher share of VRES, the patterns of cross-border electricity flows are already undergoing significant changes. Imports and exports increasingly reflect fluctuating renewable availability rather than predictable demand-driven patterns [3,4]. As a result, scaling historical time series—a common practical workaround when full-system simulations are computationally prohibitive—can lead to misleading assumptions and inaccurate model outcomes if these structural changes are ignored.

This modeling problem occurs frequently in applied energy-system research: the question of interest may concern only one country, region, or a detailed lower-level grid, while that subsystem remains influenced by the surrounding interconnected system. Explicitly co-optimizing a detailed regional model together with the full European system for every climate year or sensitivity can substantially increase computational requirements. The proposed framework therefore targets energy-system and scenario analysis rather than real-time system operation, security assessment, or detailed grid-feasibility studies. Applications that require guaranteed AC or DC power-flow feasibility, security constraints, or operational set points still require the corresponding physics-based network model.

1.2. Literature review

Over the past decade, machine learning (ML) techniques have become central to forecasting tasks in the energy sector. Among the most established applications are electricity demand and market price forecasting, where numerous studies in high-impact journals have demonstrated the reliability and accuracy of ML models. For instance, deep learning and ensemble methods have shown strong performance in predicting short-term electricity demand [5–9], while similar approaches have been widely applied for electricity price forecasting [10–14].

More recently, ML has gained traction as a surrogate for computationally intensive power system optimization tasks, including optimal power flow (OPF), unit commitment (UC), and economic dispatch (ED), where repeated solves for uncertainty analysis or scenario screening can be prohibitive. In the OPF context, Mohammadi et al. [15] review both analytical surrogates (e.g., linearized or reduced-order models) and learned surrogates trained on large sets of OPF solutions, highlighting that data-driven approaches can yield substantial runtime reductions while retaining high accuracy. Existing ML-based OPF surrogates commonly predict dispatch decisions, identify active constraints, approximate optimal solutions, or provide warm starts for the original optimization problem [16]. For UC and ED, ML models are similarly used to approximate commitment decisions or accelerate optimization, although their performance must be assessed against strong heuristic baselines [17]. Recent work nonetheless reports encouraging results, including support vector machine-based warm-start strategies [18] and reinforcement-learning-assisted methods embedded in decomposition frameworks [19,20].

The core contribution of the present work lies in a specific use of surrogate modeling: rather than predicting the complete solution of an OPF or accelerating the solution of the original optimization problem, we predict interconnector-level cross-border flows generated by a full PSOM and subsequently use these learned flows as fixed boundary conditions in geographically reduced PSOMs. This addresses a recurring problem in power-system studies in which the research question concerns only a country or region, while the surrounding interconnected system still materially influences the modeled region through imports and exports. Explicitly representing the complete surrounding system for every climate year or scenario can substantially increase computational requirements, whereas omitting it neglects relevant external-system interactions.

Methodologically, this core contribution is implemented within a predict-then-optimize structure [21]: interconnector flows are first predicted and subsequently imposed as fixed parameters in a reduced PSOM. Unlike decision-focused learning, where the predictive model is trained according to the quality of the resulting downstream decision [22], the surrogate developed here is trained to reproduce boundary flows generated by the full PSOM. Predictive accuracy and downstream optimization performance are therefore evaluated separately.

Beyond the application-oriented development of ML surrogates for OPF, another important methodological research direction is physics-informed machine learning. Rather than relying solely on data-driven training objectives, physics-informed approaches embed physical knowledge directly into the learning process through governing equations or physically motivated constraints [23,24]. The feasibility-aware custom loss proposed in this work is inspired of the same general philosophy by incorporating domain-specific operational knowledge directly into the learning objective. Rather than enforcing the full DC-OPF equations or power-flow physics, it introduces a soft penalty for net-export patterns that are inconsistent with an approximate domestic supply bound and may induce artificial ENS when imposed in a reduced PSOM. The penalty is not a projection onto the DC-OPF feasible set and does not provide a formal feasibility guarantee. Extending this idea toward richer physics-informed formulations that incorporate additional power-system constraints represents a promising direction for future research.

Table 1
Comparison of the proposed framework with representative literature most closely related to its core contribution. A checkmark indicates that the corresponding element is an explicit part of the cited approach.

Study	ML used	Cross-border flow prediction	Full-PSOM flow targets	External-system representation	Learned flows as reduced-PSOM boundaries	Downstream reduced-PSOM validation
Mertens et al. [2]	–	–	–	✓	–	✓
Frey et al. [25]	–	–	–	✓	–	✓
Aravena and Papavasiliou [26]	–	–	–	✓	–	–
Kledzik and Haker [30]	✓	✓	–	–	–	–
Fioretto et al. [27]	✓	–	–[a]	–	–	–
Zamzam and Baker [28]	✓	–	–[a]	–	–	–
Donti et al. [29]	✓	–	–[a]	–	–	–
This work	✓	✓	✓	✓	✓	✓

[a] These studies learn variables associated with the original constrained optimization problem rather than interconnector-flow targets that are transferred as boundary conditions to a geographically reduced PSOM.

In parallel, the power-system modeling literature documents that cross-border exchanges are often simplified exogenously to retain tractability. For example, Frey et al. [25] employ historical import/export profiles within a coupled modeling framework and explicitly note that such assumptions can become problematic as the energy transition alters both the magnitude and temporal structure of cross-border exchanges. Similarly, Aravena and Papavasiliou [26] model exchanges between the CWE region and non-CWE countries as fixed flows in a zonal market context. Mertens et al. [2] directly demonstrate the importance of this modeling problem for long-term optimization models with limited geographical scope: the treatment of cross-border electricity trade can affect technology choices and welfare estimates, motivating explicit representations of the external system. Their work establishes the modeling problem addressed here, but uses model-based representations of foreign electricity trade rather than learning interconnector-level flow time series from a full-system PSOM.

The comparison with adjacent ML-based optimization literature further sharpens the gap. Representative OPF-surrogate studies learn variables of the original optimization problem and seek to approximate or accelerate the original OPF, including feasibility-aware approaches that incorporate engineering constraints [27–29]. These methods address an important but different task than transferring the effect of an external interconnected system to a geographically reduced PSOM. Thus, the individual elements of the proposed framework have precedents in different research streams, but our literature review did not identify previous work that combines them in the same way: full-system PSOM results provide interconnector-flow training targets, the learned interconnector flows become fixed boundary conditions of a geographically reduced PSOM, and their quality is assessed through downstream optimization outcomes. This combination constitutes the core contribution of the proposed framework.

Table 1 makes the specific literature gap explicit. We do not claim that ML-based OPF surrogates, ML-based cross-border-flow prediction, or geographically reduced energy-system models are individually new. The core contribution is their connection in a reusable boundary-flow surrogate framework: interconnector-level flows generated by a full PSOM are learned for subsequent use as boundary conditions in a geographically reduced PSOM, and the resulting reduced-model costs, dispatch, storage operation, trade, and ENS are compared with the corresponding full-system outcomes.

1.3. Original contributions and proposed ML models

Given the background outlined above, the core contribution is a framework for learning cross-border boundary information from a full European PSOM and transferring this information to geographically reduced PSOM studies. The framework is intended for prospective energy-system planning and scenario analysis in which demand and renewable deployment are specified exogenously and weather variability is represented through different climate years. It is not proposed as a replacement for detailed power-flow, security, or real-time operational analysis.

Therefore, this paper makes the following research contributions:

- **Core contribution:** We propose an ML-based surrogate framework that predicts hourly interconnector-level cross-border flows generated by a full PSOM and uses them as fixed boundary conditions in geographically reduced PSOMs. This enables repeated reduced-model studies across different climate years without repeatedly solving the full European PSOM under unchanged structural assumptions.
- We demonstrate the decision relevance of the predicted boundary flows by evaluating not only statistical flow-prediction accuracy but also their effects on downstream reduced-PSOM results, including system costs, generation dispatch, storage operation, cross-border trade, and ENS, relative to the corresponding country-level results of the full European model.
- **Additional methodological contribution:** We introduce a feasibility-aware neural-network loss that applies a soft penalty to net-export patterns that are inconsistent with an approximate domestic supply bound when the predicted interconnector flows are intended for subsequent use in a reduced PSOM. The loss is explicitly treated as a soft plausibility mechanism, not as a guarantee of complete power-flow feasibility.
- We benchmark the proposed approach against KNN, ridge regression, random forests, Extra Trees, gradient boosting, XGBoost, multiple neural-network variants, and historical/scaled import–export profiles across Austria, Germany, and Spain and out-of-sample climate years.
- We generate training targets with the pan-European LEGO DC-OPF (single node per country) under TYNDP 2024 *National Trends* 2030 assumptions, train on climate year (CY) 2009, and evaluate both predictive accuracy and downstream single-country PSOM impacts on the out-of-sample CYs 1995 and 2008 to test robustness across weather realizations.

The remainder of the paper is structured as follows. Section 2 presents the generic surrogate workflow, including the PSOM formulation used to generate interconnector-flow targets, the ML problem definition, and the KNN, ridge regression, random forests, Extra Trees, gradient boosting, XGBoost and SQU surrogate model variants (including the feasibility-aware custom-loss formulation). Section 3 then instantiates this methodology for our empirical setting by describing the simplified European single-node-per-country DC-OPF setup, the scenario data and climate years used for training and out-of-sample evaluation, and the selected case-study countries. Section 4 reports predictive performance and quantifies how surrogate-generated boundary flows affect reduced single-country PSOM outcomes relative to the corresponding country components of the full European benchmark. Finally, Sections 5

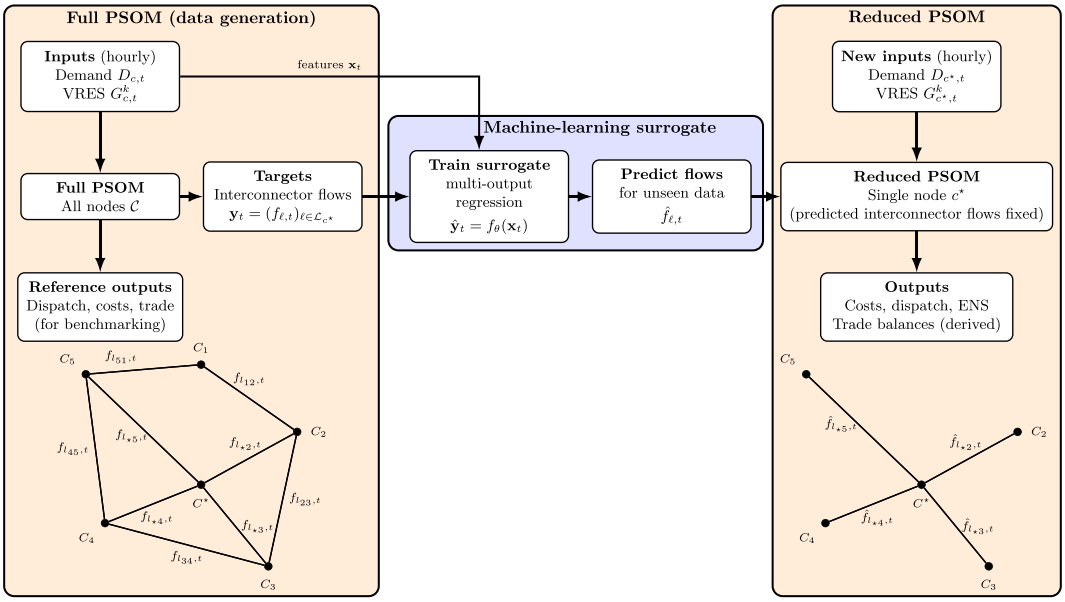


Fig. 1. Workflow of the proposed approach. A full PSOM is run to generate hourly interconnector-flow targets for the node of interest. Full-feature baselines (KNN, ridge regression, random forests, Extra Trees, gradient boosting, and XGBoost) and SQU neural-network variants are trained on demand and VRES inputs and predict flows for unseen data (e.g., a different climate year). The predicted flows are imposed as fixed boundary conditions in a reduced single-node PSOM, and the resulting costs, dispatch, ENS, and trade are compared against the corresponding outcomes from the full PSOM. This workflow enables repeated geographically reduced PSOM studies that retain a model-consistent representation of cross-border interaction while avoiding repeated pan-European optimization runs. The reduced PSOM provides a downstream consistency check; scenarios in which surrogate-generated boundary conditions induce unacceptable behavior can be rerun with the full European PSOM.

and 6 discuss the implications of the findings and outline directions for future work.

2. Methodology

We propose a surrogate framework (illustrated in Fig. 1) that maps node-level renewable generation and demand time series to hourly interconnector flows. We use the open-source LEGO PSOM [31] to generate training data for the ML models and to compute the reference results. The surrogates are trained on inputs comprising renewable generation profiles (including PV, wind, run-of-river hydro, and concentrated-solar-power (CSP)) and demand per node, with the target outputs being the interconnector flows produced by the PSOM. As an additional contribution, we develop a feasibility-aware custom loss for the neural-network surrogate. The loss softly penalizes predicted net exports that exceed an approximate domestic supply bound which may otherwise induce artificial energy-not-supplied (ENS) when the flows are imposed in a reduced PSOM. It does not enforce the full DC-OPF feasible region and is not presented as a guarantee of power-flow feasibility.

We evaluate predictive performance by comparing surrogate-generated interconnector flows to those obtained from the full PSOM. In a second step, the predicted interconnector flows are used as fixed inputs in a simplified PSOM with spatial resolution restricted to the node of interest. This allows for assessing decision relevance by comparing deviations in key outputs (objective value, trade balances, and dispatch indicators) between reduced- and full-scale PSOM runs.

The reduced PSOM also serves as a downstream consistency check for the predicted boundary conditions. If a predicted flow profile induces ENS or otherwise yields an unacceptable reduced-model solution, the corresponding climate year or scenario can be identified and evaluated with the full European PSOM instead. The surrogate is therefore intended to reduce the number of computationally expensive full-system runs, rather than to eliminate the full model or provide a formal feasibility guarantee for every predicted scenario.

To facilitate understanding of the workflow, the remainder of this section is organized as follows. Section 2.1 introduces the PSOM formulation used to generate interconnector-flow targets. Section 2.2 then defines the supervised learning problem, including target variables, feature construction, and the surrogate families, along with the feasibility-aware custom-loss variant. It also describes preprocessing, training, and the evaluation metrics. Finally, Section 2.3 explains how surrogate-predicted flows are imposed as fixed boundary conditions in a reduced PSOM and how decision relevance is assessed against the corresponding full-model benchmark outcomes.

2.1. Full PSOM

The open-source LEGO model is used in order to generate training data for the ML models and to compute the reference results. LEGO is formulated as a linear optimization model based on a DC-OPF approximation.

Let $\mathcal{G}$ denote the set of generators. We distinguish between dispatchable thermal generators $\mathcal{G}^{\text{th}}$, variable renewable generators $\mathcal{G}^{\text{vres}}$, and storage units $\mathcal{G}^{\text{st}}$, while $\mathcal{T}$ denotes the set of time steps:

$$\begin{aligned}\min\ f_{\text{obj}} = \sum_{t\in\mathcal{T}} \Bigg\{ & \sum_{g\in\mathcal{G}^{\text{th}}} \left(C_g^{\text{su}}\, su_{g,t} + C_g^{\text{com}}\, u_{g,t} + C_g^{\text{var}}\, p_{g,t} \right) \\ & + \sum_{r\in\mathcal{G}^{\text{vres}}} C_r^{\text{var}}\, p_{r,t} + \sum_{s\in\mathcal{G}^{\text{st}}} C_s^{\text{var}}\, p_{s,t}^{\text{dis}} \\ & + P^{\text{CO}_2} \sum_{g\in\mathcal{G}^{\text{th}}} \varepsilon_g \left(\alpha_g^{\text{su}}\, su_{g,t} + \alpha_g^{\text{int}}\, u_{g,t} + \alpha_g^{\text{var}}\, p_{g,t} \right) \Bigg\}. \end{aligned} \tag{1}$$

The objective function in (1) minimizes the total operational costs of all generation units across all nodes and time steps. It consists of two main components. The first component captures the conventional techno–economic operating costs. For dispatchable thermal units $g \in \mathcal{G}^{\text{th}}$, these costs include per time-step startup costs $C_g^{\text{su}}\, su_{g,t}$, commitment costs $C_g^{\text{com}}\, u_{g,t}$, and generation–dependent variable costs $C_g^{\text{var}}\, p_{g,t}$. The binary variables $su_{g,t} \in \{0,1\}$ and $u_{g,t} \in \{0,1\}$ indicate whether thermal unit g starts up or is committed at time t. For variable renewable generators $r \in \mathcal{G}^{\text{vres}}$ and storage units $s \in \mathcal{G}^{\text{st}}$, we account for variable operating

costs through $C_r^{\text{var}} p_{r,t}$ and $C_s^{\text{var}} p_{s,t}^{\text{dis}}$, respectively. The second component represents CO_2 emission costs associated with fossil fuel consumption and applies only to thermal units. These costs are computed by multiplying a uniform CO_2 price P^{CO_2} by the emission factor ε_g and the fuel consumption linked to each operating decision: startups $\alpha_g^{\text{su}} su_{g,t}$, commitment $\alpha_g^{\text{int}} u_{g,t}$, and energy production $\alpha_g^{\text{var}} p_{g,t}$.

The main constraints of the model represent standard power system relationships. For each node c and time-step t we impose a nodal power balance constraint:

$$0 = D_{c,t} - \sum_{g \in \mathcal{G}_c^{\text{th}}} p_{g,t} - \sum_{r \in \mathcal{G}_c^{\text{vres}}} p_{r,t} - \sum_{s \in \mathcal{G}_c^{\text{st}}} p_{s,t}^{\text{dis}} + \sum_{s \in \mathcal{G}_c^{\text{st}}} p_{s,t}^{\text{ch}} - \sum_{\ell \in \mathcal{L}_c} f_{\ell,t}, \quad \forall c, t. \tag{2}$$

where $D_{c,t}$ is the demand of node c at time t, $p_{g,t}$ denotes thermal generation, $p_{r,t}$ denotes variable renewable generation, $p_{s,t}^{\text{ch}}$ and $p_{s,t}^{\text{dis}}$ denote storage charging and discharging power, respectively, $f_{\ell,t}$ denotes the flow on interconnector ℓ at time t, and $\mathcal{L}_c$ denotes the set of lines from or to node c.[1] For each unit we impose generation and power limits:

$$0 \le p_{g,t} \le \overline{P}_g \qquad \forall g \in \mathcal{G}^{\text{th}},\ t, \tag{3}$$

$$0 \le p_{r,t} \le \overline{P}_{r,t}, \qquad \forall r \in \mathcal{G}^{\text{vres}},\ t, \tag{4}$$

$$0 \le p_{s,t}^{\text{ch}} \le \overline{P}_s^{\text{ch}}, \qquad \forall s \in \mathcal{G}^{\text{st}},\ t, \tag{5}$$

$$0 \le p_{s,t}^{\text{dis}} \le \overline{P}_s^{\text{dis}}, \qquad \forall s \in \mathcal{G}^{\text{st}},\ t. \tag{6}$$

where $\overline{P}_g$ is the installed capacity of thermal unit g, and $\overline{P}_s^{\text{ch}}$ and $\overline{P}_s^{\text{dis}}$ denote the maximum charging and discharging power of storage unit s, respectively. For VRES units, $\overline{P}_{r,t}$ is the time-dependent available power, given by installed capacity times the corresponding availability (capacity-factor) profile. For each interconnector ℓ and time-step t we impose line capacity constraints:

$$-\overline{F}_\ell \le f_{\ell,t} \le \overline{F}_\ell, \quad \forall \ell, t, \tag{7}$$

where $\overline{F}_\ell$ denotes the thermal limit of interconnector ℓ. The complete LEGO formulation, including the DC-OPF formulation, is given in Wogrin et al. [32]. For brevity, we restrict the presentation here to the components required for the surrogate modeling setup.

Running LEGO yields dispatch per time-step as well as the interconnector flows for all nodes.

2.2. Machine-learning surrogate

For the ML surrogate, we first define the input features and target variables (Section 2.2.1). We then compare six full-feature regression baselines (KNN, ridge regression, random forests, Extra Trees, gradient boosting, and XGBoost) (Section 2.2.2) with a feedforward sequential neural-network surrogate (SQU) (Section 2.2.3). The baselines are implemented using `scikit-learn` [33] and `XGBoost` [34], while the neural-network models are implemented in `TensorFlow/Keras` [35].

2.2.1. Target variables & feature construction

The full LEGO model computes for each interconnector ℓ and time-step t the power flow $f_{\ell,t}$. For a node of interest c^*, we define $\mathcal{L}_{c^*}$ as the set of all interconnectors that connect c^* to other nodes.

In this work, the machine-learning model is trained to reproduce the *complete vector of flows* on these interconnectors, which we obtain from the full LEGO run. Thus, for each time-step t, the target variable is the vector

$$\mathbf{y}_t = \left(f_{\ell,t}\right)_{\ell \in \mathcal{L}_{c^*}}, \qquad t \in \mathcal{T}^{\text{train}}, \tag{8}$$

which contains one entry for every interconnector linking the node of interest to other nodes.

[1] Positive $f_{\ell,t}$ denotes flow into node c.

The supervised target should not be interpreted as a unique physical mapping from demand and renewable-generation profiles to interconnector flows. In a linear OPF with aggregated technologies, identical exogenous inputs may admit multiple cost-equivalent or near-equivalent dispatch solutions, which can imply different net injections and interconnector flows. The target used here is therefore the deterministic solution produced by the specific full PSOM formulation, input data, solver configuration, numerical tolerances, and solver tie-breaking behavior. Consequently, the surrogate learns the mapping induced by this complete modeling and solution pipeline rather than a unique physical law.

The feature vector $\mathbf{x}_t$ is constructed from the time-step profiles per node of electricity demand together with the renewable generation time series, namely CSP, PV, onshore wind, offshore wind, and run-of-river hydropower. In the most general case, the feature vector for time-step t is given by

$$\mathbf{x}_t = \left[D_{c,t}, \left\{ G_{c,t}^k \right\}_{k \in \mathcal{K}},\ c \in C \right], \tag{9}$$

where C denotes the set of nodes and $\mathcal{K}$ denotes the set of renewable technologies (PV, onshore wind, offshore wind, run-of-river, CSP), and $G_{c,t}^k$ is the generation at time-step t of technology k at node c. This feature set represents the complete information available to the PSOM. We refer to this comprehensive feature representation as *Full*.

All baseline benchmark models use the complete *Full* feature set so that differences among them reflect model-class performance rather than different feature-selection pipelines. To investigate the trade-off between neural-network complexity and predictive performance, we additionally consider a reduced SQU feature set. This set, denoted as *Selected*, contains the most influential demand and renewable-generation variables identified through the neural-network sensitivity analysis described in Section 2.2.3.

We aim to learn a mapping

$$\hat{\mathbf{y}}_t = f_\theta(\mathbf{x}_t), \tag{10}$$

where θ denotes the ML model parameters (e.g., the weights and biases of the neural network in the SQU models; for the KNN baseline, θ represents the chosen hyperparameters together with the stored training instances used for neighbor lookup). The learned surrogate $f_\theta(\mathbf{x}_t)$ approximates the interconnector flows generated by the full LEGO model:

$$f_\theta(\mathbf{x}_t) \approx \mathbf{y}_t, \quad \forall t \in \mathcal{T}^{\text{train}}. \tag{11}$$

2.2.2. Full-feature baselines

The set of full-feature baseline models comprises KNN, ridge regression, random forests, Extra Trees, gradient boosting, and XGBoost. All baseline models are trained using the identical full input feature set to isolate differences attributable to the regression model itself rather than the feature-selection procedure. Their hyperparameter settings are summarized in Table 2.

As a non-parametric baseline, we employ KNN regression. For each time-step t, the model predicts the interconnector flow value $\hat{\mathbf{y}}_t$ as the average of the observed interconnector flow values associated with the k most similar historical system states. Similarity is measured using the Euclidean distance in the standardized feature space defined by $\mathbf{x}_t$. Ridge regression provides a regularized linear baseline by estimating a linear mapping between the input features and interconnector flows while penalizing large regression coefficients. Random forests represent a bagged ensemble of decision trees that reduce prediction variance through bootstrap aggregation. Extra Trees extend this ensemble approach by introducing additional randomness in the feature and split selection during tree construction. Gradient boosting builds an additive ensemble of decision trees sequentially, with each tree correcting the residual errors of the previous ensemble. XGBoost follows the same boosting

Table 2
Summary of ML model configurations used in the case study. Rows with a white background denote the baseline models, whereas grey-shaded rows denote the SQU variants.

Ctry.	Variant	#Inp.	#Outp.	Key hyperparameters	Notes
AT	KNN_FULL	249	6	$k = 20$, weights = uniform, alg. = auto	distance-based baseline
AT	RIDGE_FULL	249	6	$\alpha = 1$	linear baseline
AT	RANDOM_FOREST_FULL	249	6	400 trees, unrestricted depth	bagged ensemble
AT	EXTRA_TREES_FULL	249	6	400 trees, unrestricted depth	randomized ensemble
AT	GRADIENT_BOOSTING_FULL	249	6	300 iter., lr = 0.05	histogram gradient boosting
AT	XGBOOST_FULL	249	6	500 trees, depth = 6, lr = 0.04	subsample/columns = 0.85
AT	SQU_FULL	249	6	epochs = 100, batch = 32	MSE loss
AT	SQU_SELECTED	35	6	epochs = 100, batch = 32	top 20/output
AT	SQU_OPTIMIZED	35	6	epochs = 142, batch = 8	tuned
AT	SQU_CUSTOM_LOSS	35	6	epochs = 142, batch = 8, + custom loss	feasibility-aware
DE	KNN_FULL	249	16	$k = 20$, weights = uniform, alg. = auto	distance-based baseline
DE	RIDGE_FULL	249	16	$\alpha = 1$	linear baseline
DE	RANDOM_FOREST_FULL	249	16	400 trees, unrestricted depth	bagged ensemble
DE	EXTRA_TREES_FULL	249	16	400 trees, unrestricted depth	randomized ensemble
DE	GRADIENT_BOOSTING_FULL	249	16	300 iter., lr = 0.05	histogram gradient boosting
DE	XGBOOST_FULL	249	16	500 trees, depth = 6, lr = 0.04	subsample/columns = 0.85
DE	SQU_FULL	249	16	epochs = 100, batch = 32	MSE loss
DE	SQU_SELECTED	43	16	epochs = 100, batch = 32	top 20/output
DE	SQU_OPTIMIZED	43	16	epochs = 142, batch = 8	tuned
DE	SQU_CUSTOM_LOSS	43	16	epochs = 142, batch = 8, + custom loss	feasibility-aware
ES	KNN_FULL	249	3	$k = 20$, weights = uniform, alg. = auto	distance-based baseline
ES	RIDGE_FULL	249	3	$\alpha = 1$	linear baseline
ES	RANDOM_FOREST_FULL	249	3	400 trees, unrestricted depth	bagged ensemble
ES	EXTRA_TREES_FULL	249	3	400 trees, unrestricted depth	randomized ensemble
ES	GRADIENT_BOOSTING_FULL	249	3	300 iter., lr = 0.05	histogram gradient boosting
ES	XGBOOST_FULL	249	3	500 trees, depth = 6, lr = 0.04	subsample/columns = 0.85
ES	SQU_FULL	249	3	epochs = 100, batch = 32	MSE loss
ES	SQU_SELECTED	25	3	epochs = 100, batch = 32	top 20/output
ES	SQU_OPTIMIZED	25	3	epochs = 142, batch = 8	tuned
ES	SQU_CUSTOM_LOSS	25	3	epochs = 142, batch = 8, + custom loss	feasibility-aware

principle while incorporating second-order optimization, regularization, and efficient tree construction to improve predictive performance and computational efficiency.

All models predict the complete vector of interconnector flows in a multi-output setting. Models that do not natively support multi-output regression are fitted using one regressor per interconnector output. No feature selection or hyperparameter optimization is applied to these baselines, providing a transparent comparison of different regression model classes under the same information set.

2.2.3. Neural network model (SQU)

The SQU model is a fully connected feedforward sequential neural network with multiple dense hidden layers and rectified linear unit (ReLU) activation functions.

The baseline configuration, SQU_FULL, is trained using the standard mean squared error (MSE) loss:

$$\mathcal{L}_{\mathrm{MSE}}(\theta) = \frac{1}{|\mathcal{T}^{\mathrm{train}}|} \sum_{t \in \mathcal{T}^{\mathrm{train}}} \left\| \mathbf{y}_t - f_\theta(\mathbf{x}_t) \right\|_2^2 . \tag{12}$$

In addition to SQU_FULL, we evaluate a reduced-feature variant (SQU_SELECTED) and a hyperparameter-tuned variant (SQU_OPTIMIZED).

For the neural networks, feature importance is derived from a sensitivity analysis based on the Jacobian of the trained model using the neuralsens Python package [36]. Specifically, we compute the partial derivatives of each output with respect to each standardized input feature across the training set. The standard deviation of these sensitivities serves as an importance measure, reflecting how strongly variations in a given feature influence predicted interconnector flows. For each output, the most influential features are selected, and their union defines the SQU_SELECTED feature set. Specifically, feature importance is evaluated separately for each interconnector output. The twenty highest-ranked features for each output are combined, and duplicate entries are removed to form a single reduced feature set used for training all outputs simultaneously. The threshold of twenty features per output is applied consistently across all case-study countries as a pragmatic compromise between retaining output-specific information and substantially reducing the dimensionality of the full 249-feature input space.

In addition, we formulate a novel feasibility-aware loss function (SQU_CUSTOM_LOSS) designed to discourage interconnector-flow predictions that are inconsistent with an approximate domestic supply bound. Unlike generic regularization terms, the proposed loss is specifically tailored to the surrogate prediction of interconnector boundary conditions for geographically reduced PSOMs. Its purpose is to improve the operational consistency of the predicted exchanges when they are imposed as fixed boundary conditions in reduced optimization models. In addition to minimizing the squared error of the interconnector-flow vector, the loss includes a penalty for excessive predicted net exports. In particular, it discourages situations in which the implied exports exceed the node's maximum available generation capacity plus inflows to the node minus demand, which would otherwise lead to artificial ENS events in the reduced PSOM. Rather than enforcing the complete DC-OPF equations, the proposed loss incorporates a physically motivated operational constraint directly into the learning objective. This follows the general philosophy of physics-informed machine learning, where prior physical knowledge is embedded into the optimization process through additional loss terms.

To operationalize this idea, we precompute a time-step-specific residual bound $R_t^{\max}$ (in MW) from exogenous demand and capacity information for c^*. Physically, $R_t^{\max}$ represents an upper bound on the net export that node c^* can sustain in hour t given demand and maximum available nodal supply. If the predicted net export exceeds this bound, the implied exchange profile cannot be balanced by the reduced single-node PSOM without creating artificial ENS:

$$R_t^{\max} = -D_{c^*,t} + \sum_{k \in \mathcal{K}} G^k_{c^*,t} + \overline{P}^{\mathrm{thermal}}_{c^*} + \overline{P}^{\mathrm{storage}}_{c^*} + \overline{P}^{\mathrm{OtherRES}}_{c^*,t}, \tag{13}$$

where $D_{c^*,t}$ is demand, $G^k_{c^*,t}$ denotes non-dispatchable renewable generation time series included in the input dataset, $\overline{P}^{\text{thermal}}_{c^*}$ is the aggregate maximum thermal production capacity, $\overline{P}^{\text{storage}}_{c^*}$ is the aggregate maximum storage discharge capacity, and $\overline{P}^{\text{OtherRES}}_{c^*,t}$ is the production of the *OtherRES* category.

We orient interconnector flows such that $\tilde{\hat{y}}_{\ell,t}$ is negative for inflow to c^* and positive for outflow.[2] Accordingly, the net export is given by $\sum_{\ell\in\mathcal{L}_{c^*}} \tilde{\hat{y}}_{\ell,t}$ (positive for net exports, negative for net imports). The feasibility penalty activates when this predicted net export exceeds the residual bound $R^{\max}_t$.

$$\mathcal{L}_{\text{custom}}(\theta) = \frac{1}{|\mathcal{T}^{\text{train}}|} \sum_{t\in\mathcal{T}^{\text{train}}} \left[\frac{1}{|\mathcal{L}_{c^*}|} \left\| \tilde{\mathbf{y}}_t - \tilde{\hat{\mathbf{y}}}_t \right\|_2^2 + \lambda \max\left\{ 0, \sum_{\ell\in\mathcal{L}_{c^*}} \tilde{\hat{y}}_{\ell,t} - R^{\max}_t \right\} \right]. \tag{14}$$

where $\lambda \geq 0$ controls the penalty strength.

2.2.4. Training procedure and evaluation metrics

All input features are standardized using the `StandardScaler` implementation from the `scikit-learn` library. The scaling parameters (mean and standard deviation) are estimated exclusively from the training set and subsequently applied unchanged to all validation and out-of-sample datasets. For the neural networks, early stopping based on validation loss is employed to prevent overfitting.

Predictive performance is evaluated using the normalized mean absolute error (NMAE) and the coefficient of determination R^2:

$$\text{NMAE} = \frac{\sum_{t=1}^{N} |y_t - \hat{y}_t|}{\sum_{t=1}^{N} |y_t|}, \tag{15}$$

$$R^2 = 1 - \frac{\sum_{t=1}^{N} (y_t - \hat{y}_t)^2}{\sum_{t=1}^{N} (y_t - \bar{y})^2}, \tag{16}$$

where y_t denotes the real interconnector flow obtained from the full PSOM, $\hat{y}_t$ the corresponding ML prediction, N the number of evaluated time-steps, and $\bar{y}$ the mean of the reference interconnector flow time series. Lower NMAE values indicate better predictive accuracy, whereas higher R^2 values (closer to 1) indicate that a larger share of the variance in y_t is explained by the model. The NMAE is used to make the error comparable across interconnectors with different flow magnitudes. The metrics are computed jointly across all interconnectors and evaluated time-steps rather than separately per line.

2.3. Integration into a reduced PSOM

Lastly, we apply the trained surrogate to *unseen* inputs (demand and VRES time series) to obtain hourly interconnector-flow predictions for the node of interest. We then construct a reduced, single-node LEGO model for the same node. In this reduced PSOM, cross-border exchanges are not optimized endogenously; instead, each interconnector flow is prescribed exogenously and fixed to the surrogate prediction $\hat{f}_{\ell,t}$ for all $\ell \in \mathcal{L}_{c^*}$ and $t \in \mathcal{T}$. This setup isolates the effect of the imposed flow time series on costs, dispatch, and feasibility (e.g., ENS), because the reduced model uses the same node-specific demand, renewable availability, technology portfolio, and cost parameters as in the full PSOM, while representing the rest of the system only through the fixed interconnector flows. The resulting reduced-model solution is conditionally optimal for the imposed boundary flows, but it is not guaranteed to reproduce the country-level component of the full-system optimum if the predicted flows deviate from the full-model reference.

The quantities evaluated here reflect the intended application of the framework: energy-system planning and scenario analysis, including generation by technology, storage operation, system cost, annual trade, and ENS. The framework is not intended to determine real-time operating points, grid-security margins, or detailed transmission- or distribution-network feasibility, for which the corresponding physics-based network formulation remains mandatory.

We solve the single-node LEGO model and we record the objective function value $f^{c^*}_{\text{obj}}$, the annual electricity imports and exports, the generation mix by technology with a particular focus on the dispatchable technologies available at the node of interest, and the ENS. We define the relative error in system cost with respect to the reference as:

$$\Delta f_{\text{obj}}[\%] = 100 \cdot \frac{f^{c^*}_{\text{obj}} - f^{c^*,\text{TRUE}}_{\text{obj}}}{f^{c^*,\text{TRUE}}_{\text{obj}}}. \tag{17}$$

And we define analogous relative deviations in total annual imports and exports.

3. Case study description

For the case study, we use a simplified European network using a DC-OPF approach, representing each bidding zone as a single node with hourly resolution in the PSOM (see Fig. 2, top left). All techno-economic input data are taken from the ENTSO-E TYNDP 2024 *National Trends* scenario for the year 2030 [37]. In most cases, a bidding zone coincides with a single country. However, some countries are subdivided into multiple bidding zones (e.g., Italy, Norway, and Sweden). For simplicity, we refer to bidding zones as *countries* in the remainder of the paper, and thus use c interchangeably as the index for the country/node. For every country c and hour t, the TYNDP dataset includes hourly electricity demand $D_{c,t}$ as well as generation time series $G^k_{c,t}$ for each non-dispatchable technology k, such as onshore wind, offshore wind, PV, CSP, run-of-river and other RES. In addition, we incorporate dispatchable generation technologies, including hydro (pumped–storage) units (with inflow profiles), nuclear power plants, gas–fired units, coal–and oil–based generation, and other non–renewable technologies. Furthermore, we include technical and economic parameters of all technologies such as installed capacity, variable and fixed costs and efficiencies. We also include interconnector capacities between all modeled countries. Since the TYNDP dataset does not provide transmission line reactances required for DC-OPF modeling, we construct a single-node-per-country network representation using the PyPSA data and the PyPSA spatial-resolution clustering method and match the resulting lines with the TYNDP interconnectors [38].

We select three countries for the case studies: Austria (AT), Germany (DE), and Spain (ES). These countries are chosen to represent distinct structural positions within the European power system. Austria is located at the geographical center of continental Europe and exhibits high levels of cross-border electricity exchange, making it strongly influenced by trans-European power flows. Germany represents one of the largest electricity systems in Europe, both in terms of installed capacity and total generation, and features a highly diversified technology mix with significant shares of VRES. Spain, in contrast, is comparatively weakly interconnected with the rest of the continental European grid, which limits its participation in cross-border electricity trading and results in more localized system dynamics. Together, these three case studies allow us to evaluate the proposed surrogate model under substantially different generation structures and degrees of cross-border integration.

Austria also illustrates the practical motivation for the proposed framework. In applied studies, a detailed regional or distribution-level model may be the actual object of interest, while its transmission-level boundary conditions remain influenced by Austria's strong coupling

[2] This orientation is applied to ensure a consistent sign convention across interconnectors, since interconnectors may be defined in either direction (from c^* to a neighbor or vice versa). By multiplying each interconnector flow by an interface-specific sign factor, we enforce that inflows to c^* are always negative and outflows are always positive, independent of the original interconnector definition.

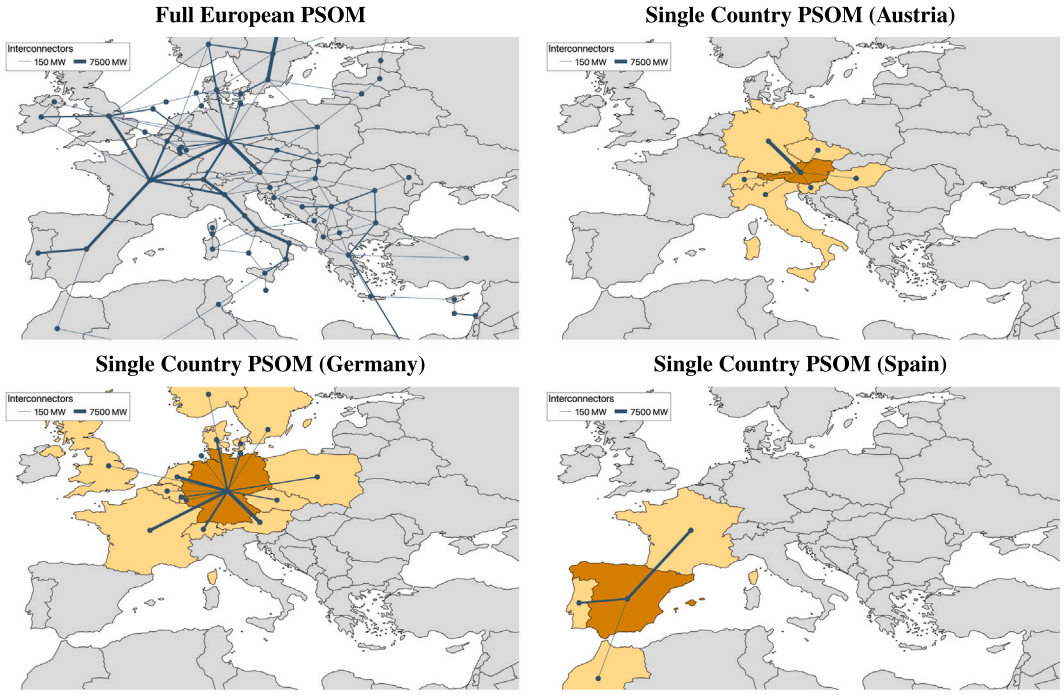


Fig. 2. Spatial scope of the optimization models. Top left: Full European PSOM, representing all modeled countries and cross-border interconnectors. Top right and bottom row: Reduced single-country PSOMs for Austria (top right), Germany (bottom left), and Spain (bottom right). In each reduced PSOM, only the country of interest is optimized explicitly; directly interconnected neighboring countries are shown for geographical context and to indicate the interfaces on which interconnector flows are prescribed exogenously as fixed boundary conditions. Line thickness indicates interconnector capacity (MW). The dark yellow shaded area marks the country of interest, while the light yellow shaded areas indicate its directly interconnected neighboring countries.

to the surrounding European system. Combining an already detailed regional or distribution-grid representation with repeated full pan-European optimizations for many climate years or sensitivities can substantially increase computation time. In such applications, learned cross-border flows can provide scenario-consistent transmission-level boundary information for the geographically reduced analysis. This use case concerns scenario assessment; it does not replace the detailed network calculations required for grid operation or security assessment.

Using the specified DC-OPF European PSOM, interconnector flows are generated based on the renewable generation profiles (including PV, wind, run-of-river hydro, and CSP) and demand per node for climate year (CY) 2009. This is selected for this case study because the ENTSO-E TYNDP 2024 Implementation Guidelines identify it as the most representative climate year [39]. Therefore, an 8 760 sample dataset is generated that includes 249 input variables (generation and demand variables). The output dimension depends on the country of interest and equals the number of its interconnectors: 6 targets for Austria, 16 for Germany, and 3 for Spain (each target representing the signed flow on one interconnector; see Table 2). We split these 8 760 samples into an 80 %/20 % training/test partition to evaluate predictive performance within CY 2009. In addition, we evaluate predictive accuracy and downstream PSOM impacts on the two out-of-sample CYs used in TYNDP 2024 (1995 and 2008), which allows us to assess whether the ranking of flow specifications is robust across different weather realizations.

These out-of-sample tests vary the weather-dependent demand and renewable-generation profiles while keeping the structural 2030 scenario assumptions fixed. The surrogate therefore learns the interconnector flows of the full PSOM under the assumed fuel and carbon prices, installed generation and storage capacities, market formulation, transmission capacities, and network topology.

As shown in Fig. 2, we construct a single-node LEGO model for each case study country that uses the same technology portfolio, the same demand time series and the same cost parameters as in the full European LEGO model, but does not explicitly model neighboring nodes. Instead for all cross-border connections $f_{\ell,t}$, the interconnection flows are fixed exogenously at the values predicted by the respective machine learning model.

To verify the proposed surrogate modeling approach, we evaluate the single-country PSOM using interconnector flows derived from different sources. As a reference, we define the *True* case, which employs the interconnector flows obtained from the full European LEGO simulation for the CY 1995 & 2008. This reference represents the internally consistent outcome of the pan-European model and serves as the benchmark against which all alternatives are assessed.

For the Austrian case study, we further include historical interconnector flows from 2024 as published by the Austrian transmission system operator (`APG_2024`). This allows us to assess whether historical cross-border flow time series are suitable boundary conditions for reduced single-country PSOMs. To reflect common practices in single-country modeling studies, we also consider two scaled variants of the APG data: one scaled to match the total annual electricity demand in the 2030 TYNDP scenario (`APG_Scaled_Demand`), and one scaled such that the annual imports/exports equal those of the *True* reference case (`APG_Scaled_ImpExp`).

4. Results

This section presents the results in four parts. First, Section 4.1 summarizes the evaluated surrogate variants, their effective input dimensionality, and the tuned hyperparameters used in the case study. Second, Section 4.2 evaluates predictive performance on the climate year 2009 training/test split. Third, Section 4.3 assesses out-of-sample predictive accuracy on the unseen climate years 1995 and 2008. Finally, Section 4.4 investigates *decision relevance* by embedding the predicted flow time series as fixed interconnector flows in reduced single-country PSOMs for Austria, Germany, and Spain and comparing objective values (total system cost), annual trade balances, dispatch patterns (including storage operation), feasibility indicators such as ENS, and computational

Table 3
Overall NMAE and R^2 on the 2009 training and test sets for the three case-study countries. Lower NMAE and higher R^2 values (closer to 1) indicate better predictive accuracy. Cell background indicates performance (green: good, yellow: intermediate, red: weak).

Model		AT		DE		ES	
		Train	Test	Train	Test	Train	Test
KNN_FULL	NMAE	0.37	0.40	0.38	0.41	0.53	0.55
	Overall R^2	0.75	0.71	0.73	0.70	0.58	0.54
RIDGE_FULL	NMAE	0.38	0.40	0.37	0.39	0.53	0.54
	Overall R^2	0.72	0.70	0.77	0.75	0.56	0.54
RANDOM_FOREST_FULL	NMAE	0.09	0.25	0.09	0.26	0.13	0.35
	Overall R^2	0.98	0.85	0.98	0.85	0.97	0.77
EXTRA_TREES_FULL	NMAE	0.00	0.22	0.00	0.21	0.00	0.32
	Overall R^2	1.00	0.88	1.00	0.89	1.00	0.79
GRADIENT_BOOSTING_FULL	NMAE	0.12	0.23	0.11	0.22	0.17	0.32
	Overall R^2	0.96	0.87	0.98	0.89	0.95	0.79
XGBOOST_FULL	NMAE	0.08	0.22	0.07	0.21	0.12	0.32
	Overall R^2	0.98	0.87	0.99	0.90	0.97	0.79
SQU_FULL	NMAE	0.16	0.24	0.18	0.24	0.19	0.33
	Overall R^2	0.93	0.85	0.92	0.85	0.90	0.70
SQU_SELECTED	NMAE	0.19	0.26	0.20	0.24	0.21	0.32
	Overall R^2	0.91	0.82	0.90	0.85	0.88	0.75
SQU_OPTIMIZED	NMAE	0.19	0.26	0.20	0.24	0.19	0.31
	Overall R^2	0.91	0.82	0.91	0.86	0.90	0.74
SQU_CUSTOM_LOSS	NMAE	0.21	0.29	0.30	0.31	0.21	0.32
	Overall R^2	0.91	0.82	0.64	0.63	0.88	0.71

runtimes against the corresponding country components of the full European benchmark.

4.1. Model configurations and hyperparameters

Before discussing predictive performance, Table 2 summarizes the evaluated model configurations. KNN_FULL, RIDGE_FULL, RANDOM_FORESTS_FULL, EXTRA_TREES_FULL, GRADIENT_BOOSTING_FULL, XGBOOST_FULL, and SQU_FULL all use the full input matrix of 249 variables. For SQU_SELECTED, the top 20 features per interconnector are retained and their union forms the country-specific reduced feature set. For SQU_OPTIMIZED, the table reports the best-performing neural-network hyperparameters found by the tuning procedure.

4.2. Training and testing results of interconnector flows for CY 2009

Table 3 reports the overall NMAE and R^2 values for the 2009 training and test sets.

The baselines reveal that nonlinear tree ensembles are highly competitive on the within-year test split. For Austria (AT) and Germany (DE), EXTRA_TREES_FULL and XGBOOST_FULL achieve the lowest test NMAE values (0.21–0.22) and test R^2 values between 0.87 and 0.90. GRADIENT_BOOSTING_FULL is similarly strong, with test NMAE of 0.23 for Austria and 0.22 for Germany. Random forests reach test NMAE of 0.25–0.26, whereas ridge regression remains close to KNN_FULL, indicating that a purely linear mapping is insufficient.

For Spain, the strongest tree-based baselines reach test NMAE values of approximately 0.32–0.35 and test R^2 values of 0.77–0.79. These results are comparable to the SQU variants, whose test NMAE ranges from 0.31 to 0.33. The perfect training fit of EXTRA_TREES_FULL ($R^2 = 1.00$) and the pronounced train–test gap in several tree ensembles nevertheless indicate substantial in-sample memorization, motivating the out-of-sample climate-year evaluation in Section 4.3.

Overall, the SQU variants achieve predictive performance comparable to the strongest tree-based ensembles on the 2009 test split. While EXTRA_TREES_FULL, XGBOOST_FULL, and GRADIENT_BOOSTING_FULL obtain the lowest statistical errors for Austria and Germany, the SQU models consistently rank among the best-performing approaches across all three case studies. SQU_CUSTOM_LOSS illustrates the trade-off between feasibility-aware training and statistical accuracy. For AT and ES its predictive performance remains close to the other SQU variants, but for DE it is notably worse (test $R^2 = 0.63$, test NMAE 0.31), indicating that the additional penalty can distort the learned mapping in systems where the reference data from the full European PSOM itself contains small ENS for Germany that the custom loss tends to suppress.

Consequently, the within-year results demonstrate that SQU provides a robust approximation of interconnector flows that is competitive with state-of-the-art tree ensembles, while the relative merits of the different approaches must ultimately be assessed using unseen climate years and their downstream impact on reduced PSOM results.

4.3. Prediction of cross-border interconnector flows for the CY 1995 & 2008

Fig. 3 shows hourly cross-border flows between Austria and Germany for one representative week. The **black** curve is the `Full European Model` (full LEGO PSOM) reference for the out-of-sample climate year 1995, i.e., the benchmark flow trajectory produced endogenously by the pan-European DC-OPF.

To illustrate the limitations of common “fixed-flow” practice, the **red** curves use historical exchanges published by the Austrian TSO APG for the calendar year 2024. These APG-based series are *not* climate-year-consistent with the 1995 weather realization; they are included purely as benchmark boundary conditions that practitioners often reuse when running reduced (single-country) models. Besides the original 2024 series, we also show two scaled variants: one scaled to match the 2030 scenario demand level (“Scaled (Demand)”) and one scaled to match annual import/export totals (“Scaled (Imp/Exp)”).

The remaining panels show the ML surrogates applied to the same climate year 1995 inputs. In the **orange** row (Baselines) and the **green** row (SQU), the dashed curves are the predicted flows for the different model variants, each compared against the same **black** full-model reference. This visualization highlights that matching annual totals via scaling does not recover the correct hourly structure, whereas the ML surrogates better reproduce the timing and magnitude of the 1995 flow pattern. The tree ensembles reproduce the principal flow peaks well, although noticeable deviations remain during low-flow periods and around sign changes, which are particularly relevant to the downstream reduced PSOM.

Across the two out-of-sample CYs (1995 and 2008), Table 4 confirms these qualitative observations. For Austria, the APG baselines perform very poorly in both years, with NMAE around 1.09–1.32 and strongly

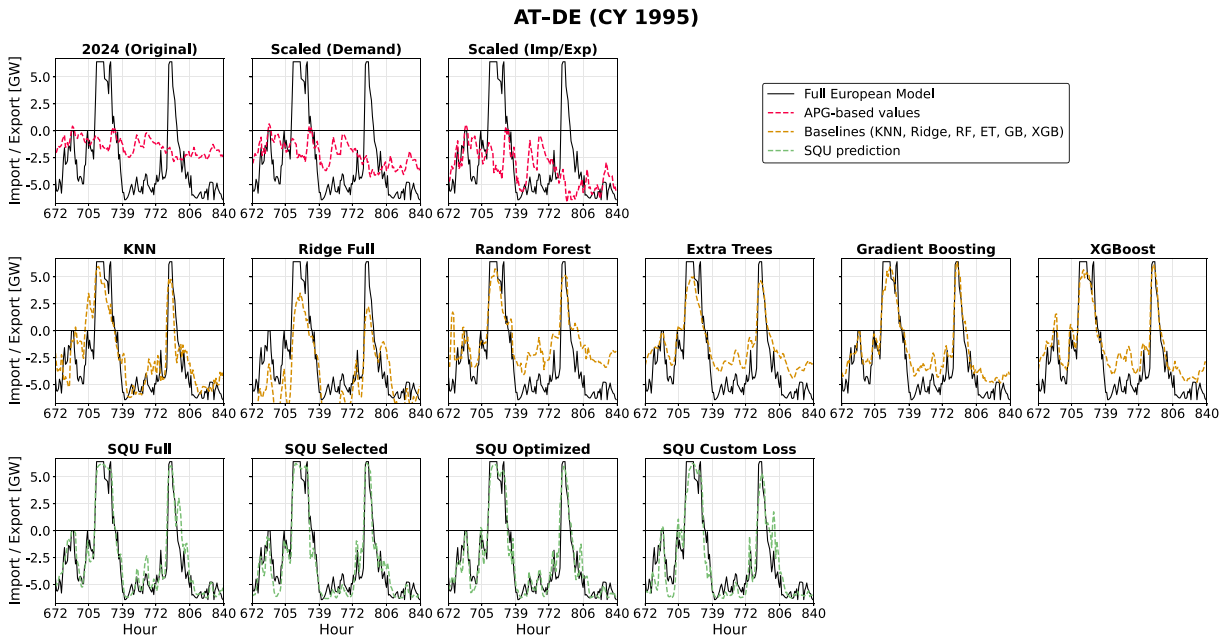


Fig. 3. Comparison of hourly cross-border interconnector flows between Austria and Germany for week 5 of climate year 1995. Positive values indicate imports into Austria and negative values indicate exports. The black line is the full European LEGO PSOM reference. Red lines show the APG 2024 historical profile and two scaled variants. Orange lines show the full-feature tabular baselines: KNN, ridge regression, random forest, Extra Trees, gradient boosting, and XGBoost. Green lines show the SQU variants using full, selected, optimized, and custom-loss configurations.

Table 4
Overall NMAE and R^2 on the out-of-sample climate years 1995 and 2008 for the three case-study countries, including the five additional full-feature baselines. Lower NMAE and higher R^2 (closer to 1) indicate better performance. Cell background indicates performance (green: good, yellow: intermediate, red: weak).

Model	Score	1995			2008		
		AT	DE	ES	AT	DE	ES
APG_2024	NMAE	1.09	n.a.	n.a.	1.10	n.a.	n.a.
	Overall R^2	-1.21	n.a.	n.a.	-1.44	n.a.	n.a.
APG_Scaled_Demand	NMAE	1.24	n.a.	n.a.	1.26	n.a.	n.a.
	Overall R^2	-2.57	n.a.	n.a.	-2.96	n.a.	n.a.
APG_Scaled_ImpExp	NMAE	1.30	n.a.	n.a.	1.32	n.a.	n.a.
	Overall R^2	-3.79	n.a.	n.a.	-4.51	n.a.	n.a.
KNN_FULL	NMAE	0.54	0.58	0.71	0.51	0.56	0.72
	Overall R^2	0.47	0.45	0.29	0.47	0.47	0.25
RIDGE_FULL	NMAE	0.54	0.52	0.67	0.52	0.51	0.67
	Overall R^2	0.39	0.50	0.19	0.42	0.53	0.28
RANDOM_FOREST_FULL	NMAE	0.47	0.42	0.46	0.44	0.41	0.48
	Overall R^2	0.57	0.62	0.68	0.57	0.64	0.64
EXTRA_TREES_FULL	NMAE	0.45	0.42	0.47	0.43	0.41	0.49
	Overall R^2	0.60	0.69	0.68	0.61	0.65	0.63
GRADIENT_BOOSTING_FULL	NMAE	0.38	0.34	0.41	0.36	0.33	0.44
	Overall R^2	0.66	0.73	0.72	0.67	0.75	0.67
XGBOOST_FULL	NMAE	0.39	0.35	0.43	0.37	0.34	0.46
	Overall R^2	0.65	0.73	0.70	0.66	0.74	0.66
SQU_FULL	NMAE	0.39	0.43	0.46	0.38	0.42	0.49
	Overall R^2	0.59	0.61	0.54	0.59	0.63	0.48
SQU_SELECTED	NMAE	0.38	0.41	0.37	0.37	0.38	0.39
	Overall R^2	0.57	0.65	0.68	0.57	0.67	0.64
SQU_OPTIMIZED	NMAE	0.38	0.41	0.36	0.37	0.37	0.41
	Overall R^2	0.59	0.65	0.69	0.57	0.70	0.62
SQU_CUSTOM_LOSS	NMAE	0.42	0.44	0.36	0.40	0.41	0.40
	Overall R^2	0.55	0.45	0.69	0.57	0.54	0.62

negative R^2 values (e.g., $R^2 = -1.21$ in 1995 and $R^2 = -1.44$ in 2008 for APG_2024), indicating performance worse than that of a trivial mean predictor. The benchmark shows a clear hierarchy among the baseline models. RIDGE_FULL and KNN_FULL generalize least robustly, whereas the tree ensembles substantially reduce the error. GRADIENT_BOOSTING_FULL is the strongest conventional baseline overall: its NMAE ranges from 0.33 to 0.44 and its R^2 from 0.66 to 0.75 across the six country–year combinations. XGBOOST_FULL follows closely, with NMAE of 0.34–0.46 and R^2 of 0.65–0.74.

While no single model class consistently achieves the highest predictive accuracy across all countries, the SQU variants deliver the most consistently strong performance across all three test regions. They

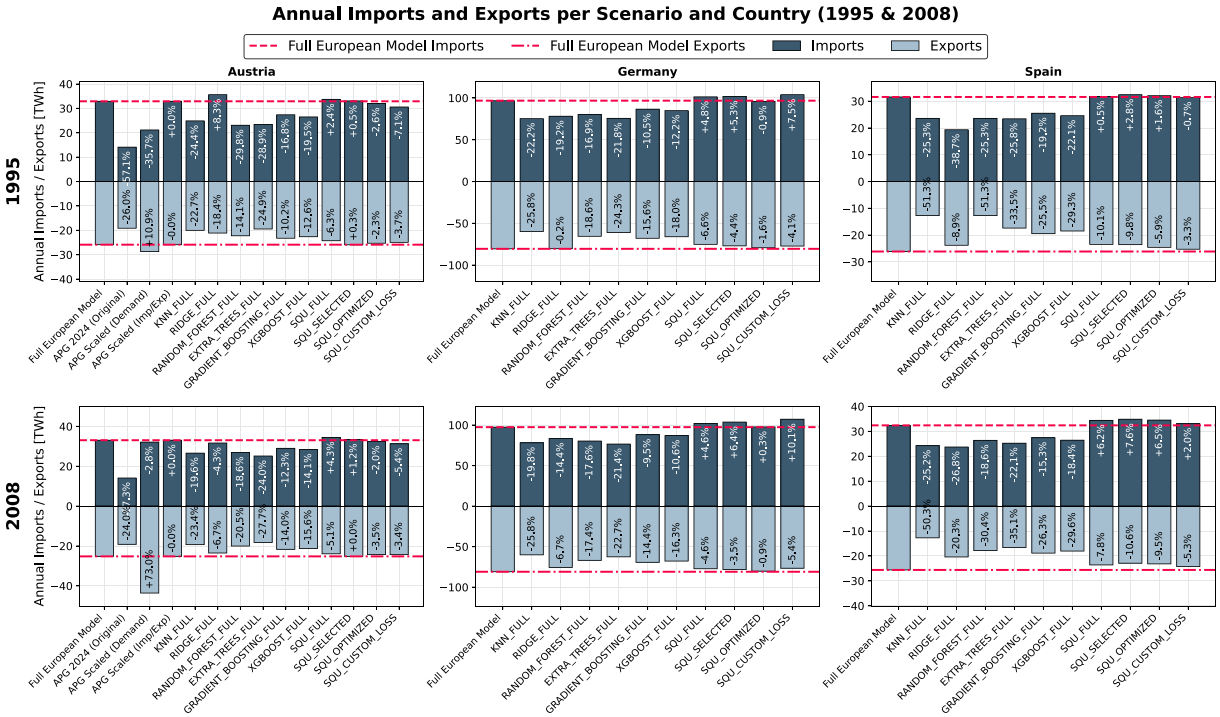


Fig. 4. Annual imports and exports in the single-country PSOM for Austria (AT), Germany (DE), and Spain (ES) under alternative cross-border flow specifications for the two out-of-sample climate years 1995 and 2008. Imports are shown as positive values and exports as negative values. Dashed lines indicate the `Full European PSOM model` reference.

remain competitive with the best-performing tree ensembles in Austria and Germany while achieving the highest accuracy in Spain.

4.3.1. Annual import and export balances

Annual import and export volumes provide a complementary consistency check on the implied cross-border interactions when cross-border exchanges are prescribed exogenously. Fig. 4 compares annual imports (positive bars) and exports (negative bars) of the country of interest in the reduced single-node PSOM against the corresponding country-level imports and exports from the `Full European PSOM Model` benchmark (dashed lines) for the two out-of-sample climate years 1995 and 2008 and the three case-study countries.

The APG-based benchmark series for Austria again highlights the limitations of reusing historical exchange patterns under structurally different system conditions. For both out-of-sample years, the unscaled APG 2024 profile produces very large deviations from the benchmark (e.g., imports are underestimated by roughly −57 % in both 1995 and 2008). The two scaled variants illustrate different limitations: `APG_Scaled_Demand` adjusts the historical profile to the 2030 demand level but does not necessarily reproduce annual trade volumes, while `APG_Scaled_ImpExp` matches annual import/export totals by construction. As shown earlier in Fig. 3, even when annual trade totals are forced to coincide, the hourly trajectory can remain inconsistent with the benchmark, and it is precisely this temporal distribution that drives dispatch, storage cycling, and cost formation.

Across both out-of-sample years, the ML-based surrogates preserve annual trade volumes more consistently than the historical APG-based profiles. `RIDGE_FULL` and the tree ensembles sometimes reproduce annual totals closely, but no single conventional baseline dominates across all countries, years, imports, and exports. The SQU variants remain comparatively close to the full-model reference in most panels. For Austria, the SQU variants keep import and export deviations within single-digit percentages. In 1995, import deviations range from approximately +0.5 % to −7.1 %, while export deviations range from approximately +0.3 % to −3.7 %. A similar pattern holds in 2008, where the SQU variants remain close to the benchmark, whereas the full-feature baseline models generally exhibit larger negative import deviations.

For Germany and Spain, the same two-year comparison shows that SQU variants typically reproduce annual imports and exports more closely than the conventional full-feature baselines. In Germany, the baseline models imply sizeable import and export shortfalls in both years, whereas SQU variants reduce these deviations to small single-digit values in most cases (with `SQU_CUSTOM_LOSS` standing out as an exception due to its systematically higher implied imports). In Spain, the baseline models again exhibit larger import underestimation, while SQU variants cluster nearer to the benchmark for both imports and exports across 1995 and 2008.

Overall, the two out-of-sample years reinforce the central point: matching annual totals ex post (as in scaling approaches) is not sufficient for economic consistency, whereas the ML-based surrogates, particularly the SQU family, better preserve the cross-border interaction implied by the full European benchmark in both magnitude and (as shown by the weekly comparison) temporal structure.

4.4. Impact on PSOM results

The statistical evaluation in the previous sections shows that ML-based surrogates can reproduce interconnector-flow patterns substantially better than by reusing historical (and scaled) profiles, which is valuable in its own right for generating scenario-consistent flow time series. Beyond prediction accuracy, however, the same surrogates can also serve as a practical modeling device in reduced spatial PSOMs: by prescribing ML-generated flows as boundary conditions, we aim to retain the key influence of cross-border exchanges while avoiding repeated full European solves.

We therefore assess the surrogate models based on their impact on key optimization outputs. To this end, we run single-country LEGO models for Austria, Germany, and Spain under alternative cross-border flow specifications. The comparison includes three classes of flow representations: (i) the `True` reference derived from the full European LEGO model

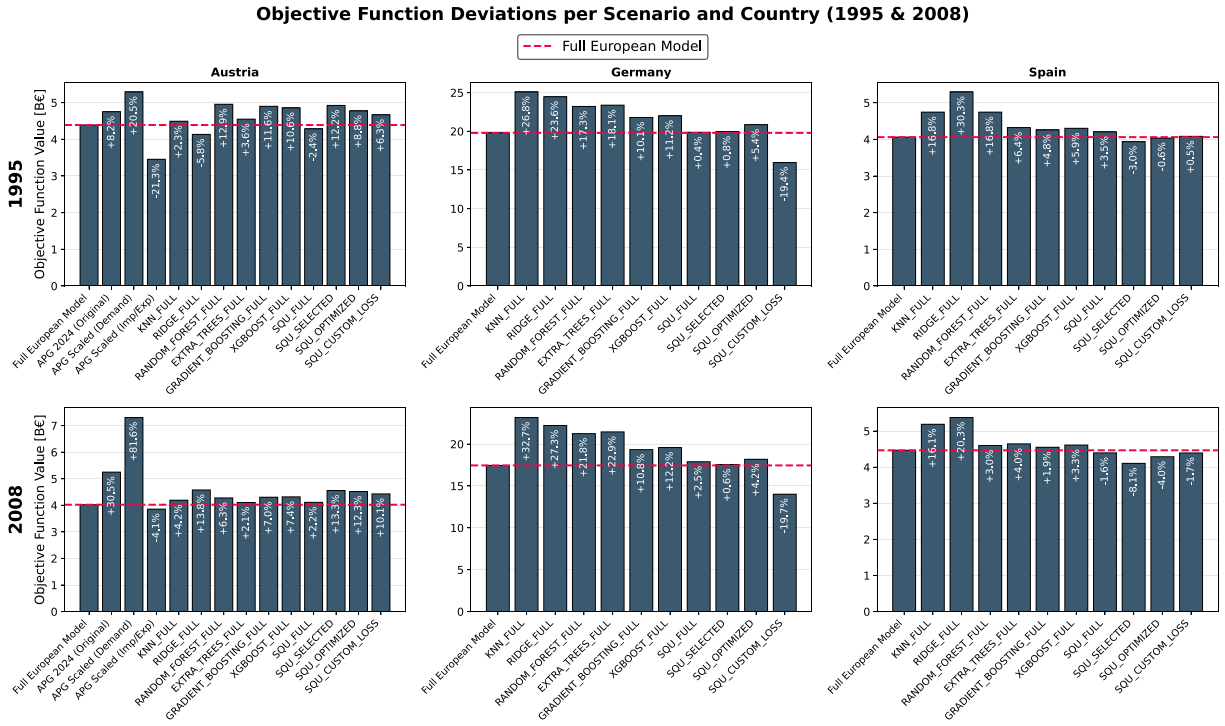


Fig. 5. Total system cost for Austria (AT), Germany (DE), and Spain (ES) under alternative cross-border flow specifications for the two out-of-sample weather years 1995 and 2008. Dashed lines indicate the `Full European PSOM model` reference.

for climate years 1995 and 2008, (ii) ML-generated interconnector flows from KNN, ridge regression, random forests, Extra Trees, gradient boosting, XGBoost, and the SQU variants for climate years 1995 and 2008, and (iii) for Austria only, APG-based benchmark series constructed from observed 2024 flows (original and scaled variants). For each country, we assess deviations in total system cost (objective function) and selected indicators of system operation and dispatch.

4.4.1. Objective function

Fig. 5 compares the objective function values (i.e., total annual system cost) of the single-country PSOMs against the corresponding `Full European Model` benchmark. To ensure a consistent comparison, the benchmark objective is evaluated for the same country only (i.e., using the country-specific component of the full European solution) and is shown as a dashed line. The deviations therefore quantify the economic distortion introduced when cross-border exchanges are prescribed exogenously rather than endogenously determined in the full European optimization. Results are reported for the two out-of-sample climate years (1995 and 2008).

In 1995, the APG-based benchmarks generate large cost distortions: the unscaled 2024 profile increases the objective by +8.2 %, scaling with demand increases it to +20.5 %, while scaling to match annual import/export totals produces a pronounced decrease of −21.3 %. In 2008, these distortions become even more pronounced (+30.5 % for the unscaled APG series and +81.6 % when demand-scaled), whereas scaling-to-Imp/Exp reduces the deviation to −4.1 %. This widening spread across climate years reinforces that matching annual aggregates can be misleading: the objective is driven by the temporal alignment of imports and exports with domestic residual demand and renewable availability, which directly affects hydro flexibility usage and the need for thermal dispatch.

The full-feature baseline models for Austria substantially reduce the objective distortion compared with the APG-based approaches, although their performance varies considerably. In 1995, `EXTRA_TREES_FULL` provides the closest conventional approximation (+3.6 %), followed by `KNN_FULL` (+2.3 %), whereas `RIDGE_FULL`, `RANDOM_FOREST_FULL`, `GRADIENT_BOOSTING_FULL`, and `XGBOOST_FULL` exhibit larger deviations ranging from −5.8 % to +12.9 %. The SQU variants remain competitive, with `SQU_FULL` slightly underestimating the objective (−2.4 %) and the remaining variants deviating between +6.3 % and +12.2 %. A similar picture emerges for 2008. Among the conventional baselines, `EXTRA_TREES_FULL` again performs best (+2.1 %), followed by `KNN_FULL` (+4.2 %) and `RANDOM_FOREST_FULL` (+6.3 %), while `RIDGE_FULL`, `GRADIENT_BOOSTING_FULL`, and `XGBOOST_FULL` show larger deviations. The SQU variants span from +2.2 % to +13.3 %. Overall, Austria demonstrates that several ML approaches substantially improve the economic consistency of the reduced PSOM compared with historical exchange profiles, while the apparent ranking of individual models remains climate-year dependent.

Germany exhibits a clearer separation between the conventional baselines and the SQU variants. In 1995, all full-feature baseline models overestimate system cost, with deviations ranging from +10.1 % for `GRADIENT_BOOSTING_FULL` to +26.8 % for `KNN_FULL`. `GRADIENT_BOOSTING_FULL` and `XGBOOST_FULL` provide the strongest conventional performance, followed by `RANDOM_FOREST_FULL` and `EXTRA_TREES_FULL`, whereas `RIDGE_FULL` and `KNN_FULL` perform worst. In contrast, the standard SQU variants remain much closer to the benchmark, with deviations between +0.4 % and +5.4 %. The same ranking is observed in 2008. The best conventional baselines, `GRADIENT_BOOSTING_FULL` and `XGBOOST_FULL`, reduce the deviation to approximately +11 %, but remain substantially less accurate than `SQU_FULL`, `SQU_SELECTED`, and `SQU_OPTIMIZED`, whose deviations range from only +0.6 % to +4.2 %. In contrast, `SQU_CUSTOM_LOSS` consistently underestimates the objective by approximately −20 % in both climate years. This systematic deviation indicates that the feasibility-oriented penalty alters the implied import–dispatch balance, favoring higher imports and therefore lower domestic generation costs than observed in the full European benchmark.

For Spain `GRADIENT_BOOSTING_FULL` and `XGBOOST_FULL` provide the closest conventional approximations for 1995, with deviations of

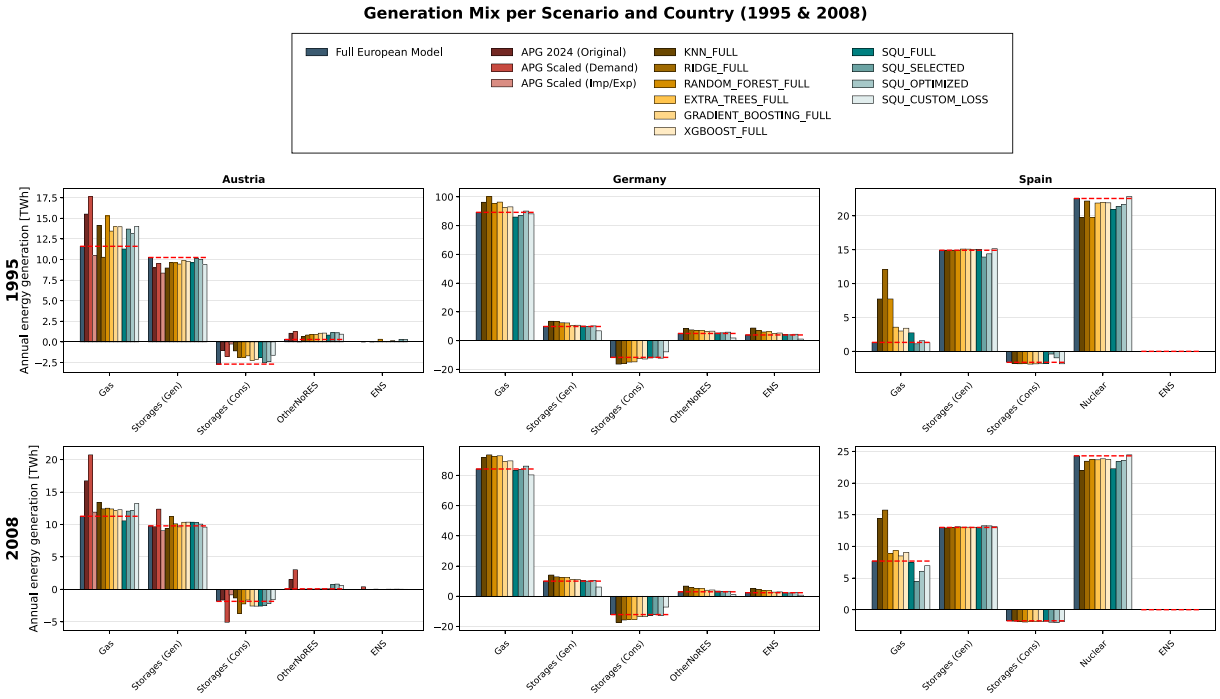


Fig. 6. Annual generation of dispatchable technologies, storage generation and charging, and ENS for Austria, Germany, and Spain under alternative cross-border flow specifications for climate years 1995 and 2008. The comparison includes the APG profiles for Austria, six full-feature tabular baselines, and the four SQU variants. Storage charging is shown as negative energy; dashed lines indicate the corresponding full European PSOM reference.

+4.8 % and +5.9 %, respectively, whereas `KNN_FULL`, `RIDGE_FULL`, and `RANDOM_FOREST_FULL` overestimate the objective by approximately 17 % to 30 %. The SQU variants remain closest to the benchmark, ranging from −3.0 % to +3.5 %. In 2008, the conventional tree ensembles again outperform `KNN_FULL` and `RIDGE_FULL`, with `EXTRA_TREES_FULL`, `GRADIENT_BOOSTING_FULL`, and `XGBOOST_FULL` limiting deviations to between +1.9 % and +4.0 %. Nevertheless, the SQU variants remain closest overall, with deviations between −8.1 % and −1.6 %. These results indicate that even for a less interconnected system, accurately reproducing the timing and direction of cross-border exchanges has a measurable impact on cost-optimal system operation.

Across all three countries and both out-of-sample climate years, the SQU-based surrogates generally provide the most consistent approximation of the benchmark objective function, outperforming the conventional full-feature baselines in terms of downstream PSOM accuracy. Although Gradient Boosting and XGBoost achieve competitive prediction accuracy, particularly for Germany, the SQU variants translate these predictions more reliably into economically consistent reduced-model solutions.

4.4.2. Generation mix and storage operation

While the objective function condenses the system response into a single cost figure, the dispatch decomposition reveals *where* and *why* alternative flow specifications affect the reduced model. Fig. 6 reports annual energy volumes by dispatchable technology group (positive values denote generation, negative values denote storage charging), with dashed red segments indicating the `Full European Model` reference for each group. Non-dispatchable renewable generation (e.g., wind, PV, and run-of-river) is omitted because it is prescribed via capacity factors and thus identical across scenarios.

A key operational outcome in this comparison is ENS, because ENS directly signals infeasibilities (or scarcity) induced by the imposed interconnector-flow time series. The `SQU_CUSTOM_LOSS` formulation is designed to reduce such physically implausible flow patterns.

Austria's operational response is dominated by gas-fired generation and pumped-hydro/storage operation, reflecting the country's role as a highly interconnected system with substantial flexible capacity. In both out-of-sample years, the APG-based benchmarks produce the largest distortions in dispatch. In particular, the APG (original) and demand-scaled variants imply a substantially different balancing requirement than the benchmark, which translates into higher gas generation and altered storage cycling relative to the dashed reference markers. These effects persist across 1995 and 2008 and reinforce that transferring a historical exchange pattern from a structurally different system year can misrepresent the flexibility needed by dispatchable generators. The ML-based surrogates cluster more closely around the benchmark for gas and storage charging/discharging, indicating that they reproduce the exchange dynamics sufficiently well to preserve the intended role of Austrian hydro and storage flexibility in system balancing. Importantly, Austria also illustrates the ENS mechanism most clearly: when the single-country PSOM is driven by ML-generated flows, a non-zero ENS appears for essentially all ML variants, indicating that small but system-relevant inconsistencies in the imposed exchange profile can translate into scarcity hours in the reduced model. By contrast, `SQU_CUSTOM_LOSS` eliminates ENS in Austria, consistent with its design goal of discouraging flow profiles that violate basic feasibility limits. This shows that, for Austria, feasibility-aware training improves *decision relevance*: the surrogate not only matches flows statistically, but also produces boundary trajectories that avoid inducing artificial shortage in the downstream optimization.

Quantitatively, `SQU_CUSTOM_LOSS` reduces the artificial Austrian ENS induced by `SQU_FULL` to zero in both CY 1995 and CY 2008, corresponding to a 100 % reduction in each year. This feasibility improvement is accompanied by a modest loss in flow-prediction accuracy: the Austrian NMAE increases from 0.39 to 0.42 in CY 1995 (a relative increase of approximately 7.7 %) and from 0.38 to 0.40 in CY 2008 (approximately 5.3 %). The result illustrates the intended trade-off of the custom loss: a small deterioration in statistical fit can materially improve the operational feasibility of the downstream reduced model.

In an additional feasibility stress test, we applied the Austrian `SQU_CUSTOM_LOSS` flow predictions to the reduced Austrian LEGO model for every climate year from 1982 through 2016. Across all 35 climate years, annual ENS was zero. This extends the two detailed case-study

years and provides evidence that the custom-loss formulation avoids surrogate-induced shortage in Austria across a wide range of historical weather realizations. It remains an empirical robustness result for the fixed 2030 model configuration, not a formal feasibility guarantee under changed capacities, prices, topology, or market rules.

Germany exhibits much larger absolute dispatch volumes (note the different axis scale), but the qualitative pattern is similar across both climate years: deviations concentrate in gas generation and storage cycling, while *OtherNoRES* remain comparatively stable. The full-feature baseline models tend to induce systematically higher gas output than the benchmark, consistent with their less accurate representation of imports/exports, whereas the standard SQU variants remain closer to the reference markers in both 1995 and 2008. A notable difference from Austria is that the `Full European Model` benchmark already contains a small amount of ENS for Germany. In this setting, `SQU_CUSTOM_LOSS` reduces ENS even further in the single-country PSOM compared to the benchmark, which is consistent with the feasibility penalty pushing the flow profile toward states that avoid shortage.

Spain, as a more weakly interconnected system, shows smaller relative sensitivity in storage operation and nuclear output across most specifications and in both climate years. The main visible differences are concentrated in gas generation, where the full-feature baseline models deviate more strongly from the benchmark than the SQU variants. This indicates that even for a less coupled system, correctly capturing exchange timing still matters for the thermal balancing requirement, especially in periods where imports would otherwise displace domestic gas generation. Spain exhibits no ENS in these experiments, and the ENS bars remain at zero across scenarios and both out-of-sample years. Consequently, the custom-loss formulation has limited room to improve feasibility and therefore changes little in the operational outcomes relative to the other SQU variants.

Overall, the dispatch comparison reinforces two points. First, reproducing annual trade totals is not sufficient: small timing inconsistencies in the imposed interconnector flows can manifest as operational distortions and even ENS in reduced PSOMs. Second, feasibility-aware training can be valuable, but its downstream impact is system-dependent: it is clearly beneficial in Austria, where it removes surrogate-induced ENS. However, in countries where the full European PSOM already contains a small amount of ENS, the same feasibility penalty can alter the trade-off between statistical fidelity and operational feasibility and may therefore not improve overall consistency.

4.4.3. Model generation and solution times

Using ML surrogates changes the PSOM runtime profile in a beneficial way. First, PSOM model generation becomes much faster because the single-country formulation contains far fewer components than the full European model. Second, PSOM solution times decrease sharply since the solver operates on a smaller and less tightly coupled optimization problem. The remaining overhead stems from the ML pipeline (data preparation, training, and prediction). For well-behaved surrogates, this overhead is small compared to the time saved in optimization. The requirement to solve the Full European Model for generating training targets does not imply that a new full-system run is required for every additional climate year. Under the same structural scenario assumptions, including generation capacities, transmission capacities, fuel and carbon prices, and market formulation, the surrogate is trained once and can then be applied to the demand and renewable-generation profiles of additional climate years without retraining. Consequently, for a fixed structural scenario, only a single Full European Model simulation is required. Additional climate years can be evaluated using the trained surrogate without rerunning the pan-European optimization. Thus, for a robustness analysis covering, for example, 35 climate years, the Full European Model has to be solved only once for the training climate year, whereas the remaining 34 climate years require only ML prediction and reduced single-country PSOM runs. The one-time full-model and training costs are therefore amortized over the multi-year analysis. Retraining would be required only when the underlying structural scenario assumptions change materially. The speed-ups reported below compare the surrogate preparation, prediction, and reduced-PSOM run with one additional Full European Model simulation; they do not repeatedly assign the one-time full-model training-data generation cost to every surrogate application.

All surrogate-based approaches substantially reduce computational effort compared with the Full European Model, which requires approximately 5 906 s per annual simulation. Depending on the country and surrogate model, total runtimes range from approximately 10 s to 867 s, corresponding to speed-ups between 6.8× and 582.4×.

Among the conventional baselines, ridge regression is consistently the fastest, achieving speed-ups of 255.2×, 254.7×, and 582.4× for Austria, Germany, and Spain, respectively. KNN exhibits similarly low runtimes, while the tree-based ensemble methods incur higher one-time training costs. Random forests require approximately 20 s– 22 s for model training, Extra Trees reduce this to about 8 s, and XGBoost and gradient boosting exhibit the highest training costs, particularly for Germany where training requires up to 89.9 s. Nevertheless, all tree ensembles still reduce total computation time by more than one order of magnitude relative to the Full European Model, with overall speed-ups ranging from approximately 52× to 333×.

Within the SQU family, `SQU_FULL` provides the most attractive runtime characteristics, achieving speed-ups between 161× and 242×. The additional feature-selection and hyperparameter-optimization stages increase the one-time offline effort for `SQU_SELECTED`, while `SQU_OPTIMIZED` and `SQU_CUSTOM_LOSS` require approximately 730 s– 840 s for automated hyperparameter optimization and network training, reducing the total speed-up to around 7×–8×. However, these additional costs are incurred only once during model development.

Consequently, once a surrogate has been trained, repeated analyses for additional climate years require only the inexpensive prediction step together with the reduced PSOM solution. No retraining and no additional Full European Model runs are necessary (Table 5).

5. Discussion

The core contribution of this work should not be interpreted as a new generic regression algorithm. Rather, it is the use of learned interconnector-level flows as transferable boundary information between a computationally expensive full-system PSOM and repeated geographically reduced PSOM studies. This distinction is practically relevant whenever the research question is regional but the region remains strongly coupled to the surrounding power system. The literature comparison in Table 1 shows that previous studies address individual elements of this problem, such as cross-border-flow prediction, external-system representation, or ML approximation of the original OPF, whereas the present framework combines full-PSOM-generated flow targets, learned boundary conditions, and downstream reduced-PSOM validation.

The intended scope is also important when interpreting feasibility. The surrogate is not proposed for real-time power-system operation, security-constrained dispatch, or detailed transmission- or distribution-grid planning, for which a feasible power-flow solution and the relevant network constraints are mandatory. Instead, the framework targets energy-system planning and scenario studies in which the principal quantities of interest include generation mixes, storage operation, system costs, curtailment, ENS, and energy exchange with neighboring systems. In such applications, geographically reduced models can provide substantial computational advantages, provided that the influence of the external system is represented sufficiently well for the research question under investigation.

Accordingly, the proposed framework should be interpreted as a computational model-reduction and screening approach rather than a feasibility-guaranteed replacement for the full PSOM. Predicted interconnector flows are evaluated by solving the downstream reduced

Table 5
Speed comparison of the full European model and single-country PSOM runs using ML flow surrogates. The ML column includes one-time surrogate preparation and training together with mean annual prediction time, but excludes the separate full European model run required to generate the training targets. The reported speed-ups therefore represent marginal or amortized speed-ups for repeated applications under unchanged structural scenario assumptions. For additional climate years, only prediction and reduced-PSOM solution times recur.

Ctry.	Case	ML	PSOM generation	PSOM solution	Total time	Speed- up
		[s]	[s]	[s]	[s]	×
–	Full European Model	0	445	5 462	5 906	1.0
AT	KNN_FULL	2	7	16	24	244.1
AT	RIDGE_FULL	0.14	7	16	23.14	255.2
AT	RANDOM_FOREST_FULL	20.25	7	16	43.25	136.6
AT	EXTRA_TREES_FULL	7.76	7	16	30.76	192.0
AT	GRADIENT_BOOSTING_FULL	25.71	7	16	48.71	121.3
AT	XGBOOST_FULL	19.42	7	16	42.42	139.2
AT	SQU_FULL	13	7	11	32	186.0
AT	SQU_SELECTED	38	7	16	61	96.7
AT	SQU_OPTIMIZED	775	7	11	792	7.5
AT	SQU_CUSTOM_LOSS	786	7	16	810	7.3
DE	KNN_FULL	4	9	14	27	215.1
DE	RIDGE_FULL	0.18	9	14	23.18	254.7
DE	RANDOM_FOREST_FULL	21.48	9	14	44.48	132.8
DE	EXTRA_TREES_FULL	8.11	9	14	31.11	189.8
DE	GRADIENT_BOOSTING_FULL	89.92	9	14	112.92	52.3
DE	XGBOOST_FULL	54.84	9	14	77.84	75.9
DE	SQU_FULL	14	8	15	37	160.9
DE	SQU_SELECTED	94	8	15	117	50.4
DE	SQU_OPTIMIZED	832	8	14	855	6.9
DE	SQU_CUSTOM_LOSS	844	9	15	867	6.8
ES	KNN_FULL	1	3	7	11	525.5
ES	RIDGE_FULL	0.14	3	7	10.14	582.4
ES	RANDOM_FOREST_FULL	21.57	3	7	31.57	187.1
ES	EXTRA_TREES_FULL	7.72	3	7	17.72	333.3
ES	GRADIENT_BOOSTING_FULL	9.37	3	7	19.37	304.8
ES	XGBOOST_FULL	9.08	3	7	19.08	309.5
ES	SQU_FULL	14	3	7	24	242.2
ES	SQU_SELECTED	71	3	7	82	72.2
ES	SQU_OPTIMIZED	728	3	7	739	8.0
ES	SQU_CUSTOM_LOSS	737	3	7	747	7.9

PSOM. Cases exhibiting surrogate-induced ENS or otherwise unacceptable behavior can subsequently be evaluated with the full European model. This workflow is particularly attractive for multi-climate-year or large-sensitivity studies, because the computationally expensive full-system model needs to be revisited only for cases for which the surrogate representation proves inadequate.

The results show that the choice of cross-border flow representation is not a minor modeling detail, but a structural assumption that can materially affect both statistical fit and optimization outcomes. In particular, the comparison against APG-based profiles highlights that "fixed-flow" practices can fail even when they appear reasonable at an aggregate level. For Austria, the 2024 APG series and its scaled variants deliver strongly negative R^2 values in both out-of-sample climate years, indicating that they are not only inaccurate but systematically inconsistent with the flow patterns implied by the 2030 system assumptions. The key issue is not the annual net balance per se, but the intra-annual timing of imports and exports relative to residual demand and renewable availability. Scaling can force annual totals to align by construction, yet still produces flow shapes that shift scarcity hours, storage cycling, and thermal dispatch in ways that distort the objective value and operational indicators. This reinforces a central point: for reduced PSOMs, the temporal structure of cross-border exchanges is decision-relevant.

However, the surrogate does not learn the real European power system. It learns the flow behavior of a specific simplified pan-European PSOM DC-OPF formulation under a fixed scenario, climate-year set, and solver configuration. The resulting flow profiles are therefore model-consistent synthetic boundary conditions, not forecasts of actual future market outcomes.

A key implication is that statistical improvements translate into more consistent optimization outcomes when the flows are used as fixed boundary conditions. In the single-country PSOMs, SQU-based flow specifications cluster closer to the corresponding country components of the full European benchmark in terms of total system cost, annual trade balances, and dispatch patterns. This indicates that the surrogate is not merely fitting point-wise errors, but preserving enough of the timing and sign structure of exchanges to maintain the role of domestic flexibility resources. Austria illustrates this mechanism clearly: deviations in annual imports are typically compensated by gas generation and storage operation, so mis-specified flow timing quickly propagates into distorted dispatch. Germany shows smaller relative sensitivity in many cases due to system scale and portfolio diversity, yet even modest percentage shifts remain economically meaningful. Spain is less driven by cross-border exchanges in magnitude, but improved flow timing still affects the residual demand met by dispatchable units and hence the cost outcome.

Feasibility-aware training provides an additional, operationally meaningful lever. The custom-loss formulation augments the prediction error with a penalty for physically implausible net-export profiles, thereby targeting surrogate-induced infeasibilities that manifest as ENS once the flows are imposed in the reduced PSOM. The results indicate that this mechanism can work as intended when the benchmark itself is feasible: in Austria, where the full European benchmark exhibits no ENS, several ML-based flow specifications induce small surrogate-driven ENS, whereas SQU_CUSTOM_LOSS removes these artifacts and improves operational consistency. The additional robustness check, where the Austrian SQU_CUSTOM_LOSS surrogate was evaluated for every available climate year from 1982 to 2016 shows, that across all 35 climate years,

the reduced PSOM remained free of artificial ENS, indicating that the feasibility-aware penalty generalizes well across weather variability under fixed structural assumptions. For Germany, where the full European benchmark contains a small amount of ENS, the custom loss can reduce ENS further in the reduced model, but this comes with the risk of shifting imports and dispatch away from the benchmark trajectory and thereby affecting the objective. In Spain, ENS does not occur, and the custom loss has limited impact. Overall, feasibility-aware training is most effective for eliminating surrogate-induced infeasibilities; its benefits are system-dependent and may require adaptive designs to avoid unintended economic biases when the benchmark itself contains scarcity.

A first methodological limitation concerns uniqueness. In linear OPF and energy-system models with aggregated technologies, identical exogenous demand and renewable-generation profiles may admit several cost-equivalent or near-equivalent dispatch solutions, which can imply different interconnector flows. The supervised target is therefore the solution produced by a specific full PSOM instance and solver configuration. The surrogate learns the deterministic mapping induced by the formulation, data, solver, numerical tolerances, and tie-breaking behavior rather than a unique physical law. This is appropriate when the objective is to reproduce a particular modeling workflow.

A second limitation concerns generalization. The surrogate implicitly learns the dispatch and net-injection response of the full PSOM under fixed structural assumptions. The out-of-sample tests on CY 1995 and CY 2008 therefore evaluate robustness to unseen weather realizations, but not to changes in fuel prices, carbon prices, installed generation or storage capacities, market rules, transmission capacities, network topology, or bidding-zone configurations. Because such changes alter the hidden mapping from exogenous inputs to dispatch, net injections, and flows, transfer to substantially different system conditions would require new full-model simulations followed by retraining the ML models. Future work could therefore investigate scenario-conditioned or transfer-learning surrogates that remain applicable across multiple structural scenarios.

A third limitation concerns feasibility and optimality. Feasibility is indispensable when ML outputs are intended to replace physics-based operating or security calculations. That is not the role claimed for the present framework. The custom loss can reduce net-export predictions that exceed an approximate domestic supply bound and thereby reduce surrogate-induced ENS, but it remains a soft penalty: it is neither a DC-OPF feasibility projection nor a guarantee of reduced-model or full-system feasibility. This limitation is why downstream PSOM feasibility and ENS are explicitly evaluated and why problematic scenarios can be returned to the full European model.

Relative to decomposition-based and learning-assisted optimization methods, the proposed approach intentionally makes a different trade-off. Decomposition approaches preserve an explicit link to the original optimization problem and can provide convergence or optimality guarantees, but they generally require iterative coordination or repeated subproblem solves. Our surrogate workflow is non-iterative once trained and is designed for repeated reduced-model studies where a fully coupled European solve is impractical. It therefore improves practical tractability and boundary-condition consistency, but does not solve the original full OPF more efficiently or provide formal guarantees.

The runtime results should therefore be interpreted in an amortized sense. The full-model solve needed to generate training data, feature selection, hyperparameter tuning, and model training must be counted as preparation costs. The approach becomes attractive when these costs are reused over many reduced-model runs, for example across climate years.

Overall, the framework is most appropriate when (i) a full-system PSOM can be solved for a representative training scenario, (ii) many reduced-model studies need consistent cross-border boundary conditions, and (iii) the target scenarios remain within or near the distribution represented by the training data. It is less appropriate when formal feasibility guarantees, market-clearing guarantees, or extrapolation to structurally different systems are required.

6. Conclusions and outlook

The core contribution of this work is an ML surrogate framework for transferring cross-border interactions from a full-system PSOM to geographically reduced PSOM studies through predicted interconnector-level boundary flows. The surrogate learns cross-border flows generated by the full European PSOM and provides them as fixed boundary conditions to reduced models, thereby avoiding repeated full-European optimization runs when multiple climate years are analyzed under unchanged structural assumptions. The framework was implemented with the open-source LEGO DC-OPF model under ENTSO-E TYNDP 2024 *National Trends* assumptions for 2030. Surrogates were trained on climate year 2009 and evaluated on two fully out-of-sample climate years (1995 and 2008) for Austria, Germany, and Spain.

The results show, within the considered LEGO DC-OPF setup, that commonly used fixed-flow approximations based on historical import/export profiles can become highly inconsistent under changing weather realizations and future system configurations. In the Austrian case, APG 2024 baselines (including scaled variants) perform poorly in both out-of-sample years and can induce substantial distortions in reduced-model outcomes. By contrast, ML-based surrogates provide robust improvements across countries and in both out-of-sample years. In particular, the neural-network family (SQU) generalizes more robustly than the conventional full-feature baselines to the out-of-sample climate years and yields single-country PSOM outcomes that are markedly closer to the corresponding country components of the full European benchmark in total system cost, annual trade balances, and dispatch patterns. This indicates that, under fixed scenario assumptions, the surrogates capture part of the nonlinear dependence of exchanges on renewable availability and demand sufficiently well to preserve the role of domestic flexibility resources in reduced optimization models.

As an additional contribution, the feasibility-aware custom loss softly penalizes net-export profiles that exceed an approximate domestic supply bound. It does not guarantee complete power-flow feasibility, but it is effective at suppressing *surrogate-induced* ENS in the examined application. In Austria, where the full European benchmark exhibits no ENS, the custom loss eliminates ENS in the reduced single-country runs and thereby improves operational consistency. This result is further supported by the Austrian climate-year stress test for 1982–2016, in which the reduced LEGO model produced zero ENS for all 35 years when driven by `SQU_CUSTOM_LOSS` flows.

From a computational perspective, the workflow offers large speedups for repeated reduced-model solves because reduced single-country PSOMs are substantially smaller than the full European model and surrogate prediction is inexpensive once trained. In addition, the approach is naturally parallelizable: countries and climate years can be evaluated independently, enabling efficient multi-country and multi-year analyses.

Building on these results, a key next step is to increase spatial granularity by predicting flows per interconnector or transmission line in a fully European setting with an explicit high-voltage grid representation, allowing reduced models to reflect congestion and regional trade patterns more realistically. Further work should assess the approach in generation expansion planning models, broaden the surrogate output space (e.g., zonal prices, congestion indicators, reserve requirements), and develop probabilistic surrogates to quantify uncertainty and propagate it into planning decisions. Another promising direction is decision-focused learning, in which the surrogate is trained directly with respect to downstream optimization quality rather than prediction accuracy alone. Finally, feasibility-aware training warrants deeper investigation, including hybrid or country-specific constraints that reduce ENS without introducing systematic biases in implied trade and dispatch patterns. Future work includes extending the feasibility-aware loss

toward more comprehensive physics-informed surrogate models that embed additional network and operational constraints directly into the learning process.

The framework does not guarantee power-flow feasibility and is not intended to replace physics-based operational or security analyses. Instead, the reduced PSOM serves as a downstream consistency check, with the full European model available when the surrogate representation is insufficient. Future work could build on advances in physics-informed learning and feasibility-restoring methods to extend the current soft constraint toward stricter or hard feasibility enforcement.

Overall, the proposed framework demonstrates that learning boundary conditions, rather than complete OPF solutions, provides a practical and computationally efficient pathway for geographically reduced power system optimization under fixed structural scenarios.

CRediT authorship contribution statement

Robert Gaugl: Writing – original draft, Visualization, Software, Methodology, Formal analysis, Conceptualization. **Eloy Insunza:** Writing – review & editing, Validation, Methodology. **José Portela:** Writing – review & editing, Validation, Supervision, Methodology. **Sonja Wogrin:** Writing – review & editing, Validation, Supervision.

Disclaimer

The authors acknowledge the use of OpenAI's ChatGPT (version GPT-5.2, February 2026) for assistance in improving the grammar, spelling, and clarity of the manuscript. All content and conclusions are the sole responsibility of the authors.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Robert Gaugl reports that financial support was provided by Rudolf Chaudoire Foundation. Robert Gaugl reports that financial support was provided by Erasmus+. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

R. Gaugl gratefully acknowledges the Institute for Research in Technology (IIT) – School of Engineering (ICAI) at the Universidad Pontificia Comillas for hosting his research visit and for the funding granted by the Rudolf Chaudoire Foundation and the Erasmus+ program.

Data availability

Data will be made available on request.

References

[1] European Union. Regulation (EU) 2019/943 of the European parliament and of the council of 5 June 2019 on the internal market for electricity; 2019. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32019R0943.

[2] Mertens T, Poncelet K, Duerinck J, Delarue E. Representing cross-border trade of electricity in long-term energy-system optimization models with a limited geographical scope. Appl Energy 2020;261:114376. https://doi.org/10.1016/j.apenergy.2019.114376. https://www.sciencedirect.com/science/article/pii/S030626191932063X.

[3] Crozier C, Baker K. The effect of renewable electricity generation on the value of cross-border interconnection. Appl Energy 2022;324:119717. https://doi.org/10.1016/j.apenergy.2022.119717. https://www.sciencedirect.com/science/article/pii/S030626192201008X.

[4] Unger EA, Ulfarsson GF, Gardarsson SM, Matthiasson T. The effect of wind energy production on cross-border electricity pricing: the case of Western Denmark in the Nord pool market. Econ Anal Policy 2018;58:121–30. https://doi.org/10.1016/j.eap.2018.01.006. https://www.sciencedirect.com/science/article/pii/S0313592617302011.

[5] Bedi J, Toshniwal D. Deep learning framework to forecast electricity demand. Appl Energy 2019;238:1312–26. https://doi.org/10.1016/j.apenergy.2019.01.113. https://www.sciencedirect.com/science/article/pii/S0306261919301217.

[6] Gajowniczek K, Zabkowski T. Two-stage electricity demand modeling using machine learning algorithms. Energies 2017;10:1547. https://doi.org/10.3390/en10101547. https://www.mdpi.com/1996-1073/10/10/1547, Number: 10.

[7] Nowotarski J, Weron R. Recent advances in electricity price forecasting: a review of probabilistic forecasting. Renew Sustain Energy Rev 2018;81:1548–68. https://doi.org/10.1016/j.rser.2017.05.234. https://www.sciencedirect.com/science/article/pii/S1364032117308808.

[8] Lago J, Marcjasz G, De Schutter B, Weron R. Forecasting day-ahead electricity prices: a review of state-of-the-art algorithms, best practices and an open-access benchmark. Appl Energy 2021;293:116983. https://doi.org/10.1016/j.apenergy.2021.116983. https://www.sciencedirect.com/science/article/pii/S0306261921004529.

[9] Billé AG, Gianfreda A, Del Grosso F, Ravazzolo F. Forecasting electricity prices with expert, linear, and nonlinear models. Int J Forecast 2023;39:570–86. https://doi.org/10.1016/j.ijforecast.2022.01.003. https://www.sciencedirect.com/science/article/pii/S0169207022000036.

[10] Tschora L, Pierre E, Plantevit M, Robardet C. Electricity price forecasting on the day-ahead market using machine learning. Appl Energy 2022;313:118752. https://doi.org/10.1016/j.apenergy.2022.118752. https://www.sciencedirect.com/science/article/pii/S0306261922002057.

[11] Nitsch F, Schimeczek C, Bertsch V. Applying machine learning to electricity price forecasting in simulated energy market scenarios. Energy Rep 2024;12:5268–79. https://doi.org/10.1016/j.egyr.2024.11.013. https://www.sciencedirect.com/science/article/pii/S2352484724007327.

[12] Hong T, Fan S. Probabilistic electric load forecasting: a tutorial review. Int J Forecast 2016;32:914–38. https://doi.org/10.1016/j.ijforecast.2015.11.011. https://www.sciencedirect.com/science/article/pii/S0169207015001508.

[13] Eren Y, Küçükdemiral I. A comprehensive review on deep learning approaches for short-term load forecasting. Renew Sustain Energy Rev 2024;189:114031. https://doi.org/10.1016/j.rser.2023.114031. https://www.sciencedirect.com/science/article/pii/S1364032123008894.

[14] Biswal B, Deb S, Datta S, Ustun TS, Cali U. Review on smart grid load forecasting for smart energy management using machine learning and deep learning techniques. Energy Rep 2024;12:3654–70. https://doi.org/10.1016/j.egyr.2024.09.056. https://www.sciencedirect.com/science/article/pii/S2352484724006346.

[15] Mohammadi S, Bui V-H, Su W, Wang B. Surrogate modeling for solving OPF: a review. Sustainability 2024;16:9851. https://doi.org/10.3390/su16229851. https://www.mdpi.com/2071-1050/16/22/9851, Number: 22.

[16] Khaloie H, Dolányi M, Toubeau J-F, Vallée F. Review of machine learning techniques for optimal power flow. Appl Energy 2025;388:125637. https://doi.org/10.1016/j.apenergy.2025.125637. https://www.sciencedirect.com/science/article/pii/S0306261925003678.

[17] Pineda S, Morales JM. Is learning for the unit commitment problem a low-hanging fruit? Electr Power Syst Res 2022;207:107851. https://doi.org/10.1016/j.epsr.2022.107851. https://www.sciencedirect.com/science/article/pii/S0378779622000815.

[18] Pourahmadi F, Kazempour J. Unit commitment predictor with a performance guarantee: a support vector machine classifier. IEEE Trans Power Syst 2025;40:715–27. https://doi.org/10.1109/TPWRS.2024.3400405. https://ieeexplore.ieee.org/document/10530167.

[19] Zhu Y, Cui G, Liu A, Jia Q-S, Guan X, Zhai Q, Guo Q, Guo X. A reinforcement learning embedded surrogate Lagrangian relaxation method for fast solving unit commitment problems. IEEE Trans Power Syst 2025;1–12. https://doi.org/10.1109/TPWRS.2025.3529700. https://ieeexplore.ieee.org/abstract/document/10847790.

[20] Huang G, Mao T, Zhang B, Cheng R, Ou M. An intelligent algorithm for solving unit commitments based on deep reinforcement learning. Sustainability 2023;15:11084. https://doi.org/10.3390/su151411084. https://www.mdpi.com/2071-1050/15/14/11084, Number: 14.

[21] Elmachtoub AN, Grigas P. Smart "Predict, then optimize". Manag Sci 2022;68:9–26. https://doi.org/10.1287/mnsc.2020.3922

[22] Wilder B, Dilkina B, Tambe M. Melding the data-decisions pipeline: decision-focused learning for combinatorial optimization. In: Proceedings of the AAAI conference on artificial intelligence, vol. 33; 2019. p. 1658–65. https://doi.org/10.1609/aaai.v33i01.33011658

[23] Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. J Comput Phys 2019;378:686–707. https://doi.org/10.1016/j.jcp.2018.10.045. https://www.sciencedirect.com/science/article/pii/S0021999118307125.

[24] Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. Physics-informed machine learning. Nat Rev Phys 2021;3:422–40. https://doi.org/10.1038/s42254-021-00314-5. https://www.nature.com/articles/s42254-021-00314-5.

[25] Frey UJ, Sasanpour S, Breuer T, Buschmann J, Cao K-K. Tackling the multitude of uncertainties in energy systems analysis by model coupling and high-performance computing. Front Environ Econ 2024;3. https://doi.org/10.3389/frevc.2024.1398358. https://www.frontiersin.org/journals/environmental-economics/articles/10.3389/frevc.2024.1398358/full.

[26] Aravena I, Papavasiliou A. Renewable energy integration in zonal markets. IEEE Trans Power Syst 2017;32:1334–49. https://doi.org/10.1109/TPWRS.2016.2585222. https://ieeexplore.ieee.org/document/7501844.

[27] Fioretto F, Mak TWK, Hentenryck PV. Predicting AC optimal power flows: combining deep learning and Lagrangian dual methods. Proc AAAI Conf Artif Intell 2020;34:630–7. https://doi.org/10.1609/aaai.v34i01.5403. https://ojs.aaai.org/index.php/AAAI/article/view/5403.

[28] Zamzam AS, Baker K. Learning optimal solutions for extremely fast AC optimal power flow. In: 2020 IEEE international conference on communications, control, and computing technologies for smart grids (SmartGridComm); 2020. p. 1–6. https://doi.org/10.1109/SmartGridComm47815.2020.9303008. https://ieeexplore.ieee.org/document/9303008.
[29] Donti PL, Rolnick D, Kolter JZ. DC3: a learning method for optimization with hard constraints. In: International conference on learning representations; 2021.
[30] Kledzik V, Haker J. Predicting cross-border power flow using weather data [Ph.D. thesis]. KTH Royal Institute of Technology; 2024. https://www.diva-portal.org/smash/record.jsf?pid=diva2%3A1946084&dswid=4245.
[31] Wogrin S, Tejada-Arango D, Gaugl R, Klatzer T, Bachhiesl U. LEGO: the open-source low-carbon expansion generation optimization model. Softwarex 2022;19. https://doi.org/10.1016/j.softx.2022.101141
[32] Wogrin S, Tejada-Arango D, Delikaraoglou S, Botterud A. Assessing the impact of inertia and reactive power constraints in generation expansion planning. Appl Energy 2020;280:115925. https://doi.org/10.1016/J.APENERGY.2020.115925
[33] Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay E. Scikit-Learn: machine learning in Python. J Mach Learn Res 2011;12:2825–30.
[34] Contributors X. XGBoost Python package; 2025. https://xgboost.readthedocs.io/.
[35] Chollet F, et al. Keras; 2015. https://keras.io.
[36] Pizarroso J, Portela J, Muñoz A. NeuralSens: sensitivity analysis of neural networks. J Stat Softw 2022;102:1–36. https://doi.org/10.18637/jss.v102.i07
[37] ENTSO-E. TYNDP 2024 - scenario report. Technical Report ENTSO-E; 2025. https://2024.entsos-tyndp-scenarios.eu/.
[38] Brown T, Victoria M, Zeyen E, Hofmann F, Neumann F, Frysztacki M, Hampp J, Schlachtberger D, Hörsch J, Schledorn A, Schauß C, van Greevenbroek K, Millinger M, Glaum P, Xiong B, Seibold T. PyPSA-Eur: an open sector-coupled optimisation model of the European energy system; 2026. https://github.com/pypsa/pypsa-eur.
[39] ENTSO-E. TYNDP 2024 -implementation guidelines; 2024. https://tyndp.entsoe.eu/resources/tyndp-2024-implementation-guidelines-march.