COMILLAS
UNIVERSIDAD PONTIFICIA
ICAI

GRADO EN INGENIERÍA EN TECNOLOGÍAS INDUSTRIALES

TRABAJO FIN DE GRADO

Forecasting sensitivity factors from local flexibility markets

Autor: Ramón Polvorosa del Campo

Director: Matteo Troncia, José Pablo Chaves Ávila

Codirector: Marco Galici

Madrid

Julio de 2025

Declaro, bajo mi responsabilidad, que el Proyecto presentado con el título

Forecasting sensitivity factors from local flexibility markets

en la ETS de Ingeniería - ICAI de la Universidad Pontificia Comillas en el

curso académico 2024/25 es de mi autoría, original e inédito y

no ha sido presentado con anterioridad a otros efectos.

El Proyecto no es plagio de otro, ni total ni parcialmente y la información que ha sido

tomada de otros documentos está debidamente referenciada.

Fdo.: Ramón Polvorosa del Campo Fecha: 08/ 07/ 2025

Autorizada la entrega del proyecto

EL DIRECTOR DEL PROYECTO

Fdo.: José Pablo Chaves Ávila Fecha: 08/ 07/2025

José Pablo Chaves Ávila

GRADO EN INGENIERÍA EN TECNOLOGÍAS INDUSTRIALES

TRABAJO FIN DE GRADO

Forecasting sensitivity factors from local flexibility markets

Autor: Ramón Polvorosa del Campo

Director: Matteo Troncia, José Pablo Chaves Ávila

Codirector: Marco Galici

Madrid

Julio de 2025

ACKNOWLEDGEMENTS

I wish to express my profound gratitude to my three supervisors, Matteo Troncia, José Pablo Chaves Ávila, and Marco Galici, for their devoted guidance and encouragement throughout the entire development of this work. When I embarked on the project, I possessed only a basic understanding of the subject; thanks to their patient mentorship, incisive feedback, and constant availability, I was able to shape the research that follows. Their devotion to academic rigor and their disposition to assist whenever needed were indispensable, and this thesis would not have reached its present form without their steady support. Finally, my deepest thanks go to my parents, who have made everything possible by giving me everything and asking nothing in return.

FORECASTING SENSITIVITY FACTORS FROM LOCAL FLEXIBILITY MARKETS

Author: Polvorosa del Campo, Ramón.
Supervisor: Troncia, Matteo; Chaves Ávila, José Pablo.
Cosupervisor: Galici, Marco.
Collaborating Entity: ICAI – Universidad Pontificia Comillas

ABSTRACT

This thesis implements and validates an AI-driven pipeline that forecasts distribution-level sensitivity factors, Voltage/Current (V/C) Magnitude sensitivity with respect to Active/Reactive (P/Q) power (VMP, VMQ, CMP, CMQ) for both load and generation nodes, 24 h ahead. A Random Forest surrogate replaces full power-flow calculations which require information about the network, that most likely are not available for an actor of the network. The Random Forest tool is tested on a 16-bus Low Voltage (LV) network with 504 hourly samples, the surrogate cuts computation time by > 90 % respect to power flow and achieves a Mean Absolute Percentage Error (MAPE) below 3 % for all eight sensitivity factors. The work therefore demonstrates a reproducible path toward intelligent actors of the network, where a market agent can optimize the strategy adopted in the electricity markets.

Keywords: Local Flexibility Markets, Sensitivity Factors, Random Forest Surrogate, Reinforcement Learning

1. Introduction:

Distribution grids must absorb rapidly growing PV, EV and heat-pump fleets without unaffordable reinforcement. Local Flexibility Markets (LFMs) unlock behind-the-meter resources but require fast, network-aware clearing. This work explores data-driven sensitivity-factor forecasting for local flexibility markets agents.

2. Project Definition:

The thesis sets out to generate 24-hour forecasts for eight network-sensitivity factors, Voltage and Current Magnitudes for both generation and load nodes. Everything is tested on a 16-bus LV feeder at hourly resolution, using open-source code so that others can replicate or extend the work. Success is judged by the indicator MAPE for every series.

3. **Model:**

Figure 1 summarizes a three-stage pipeline. First, raw measurements are cleansed and converted into lag-stacked feature vectors. A Random Forest surrogate then learns an approximate power-transfer distribution factor matrix, replacing heavy AC power-flow computations. Finally, a dedicated forecasting block rolls the surrogate forward over a sliding 24-hour horizon, so that eight complete 24-point sensitivity trajectories are generated every hour.

Figure 1: System Architecture.

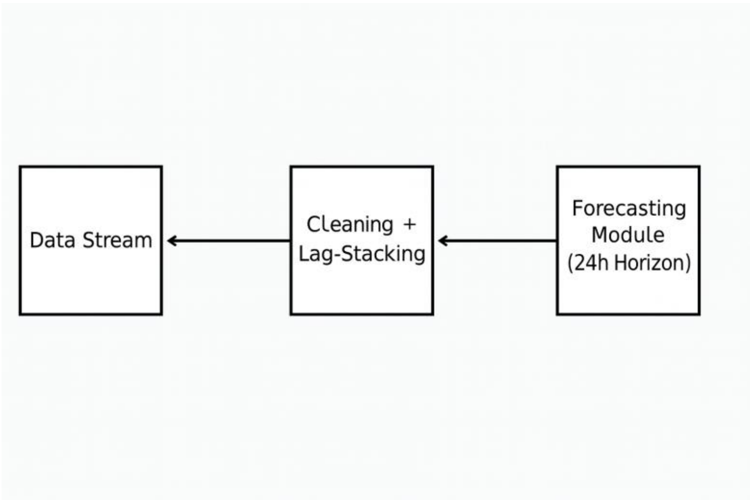

4. **Results:**

Across 504 hourly samples, the surrogate keeps the MAPE below 3 % for every series while cutting runtime by more than 90 % relative to full AC calculations. Figure 2 illustrates the tight alignment between predicted and actual voltage sensitivity trajectories for generators.

Figure 2: 24-Hour Forecast vs. Actual VMP_gen.

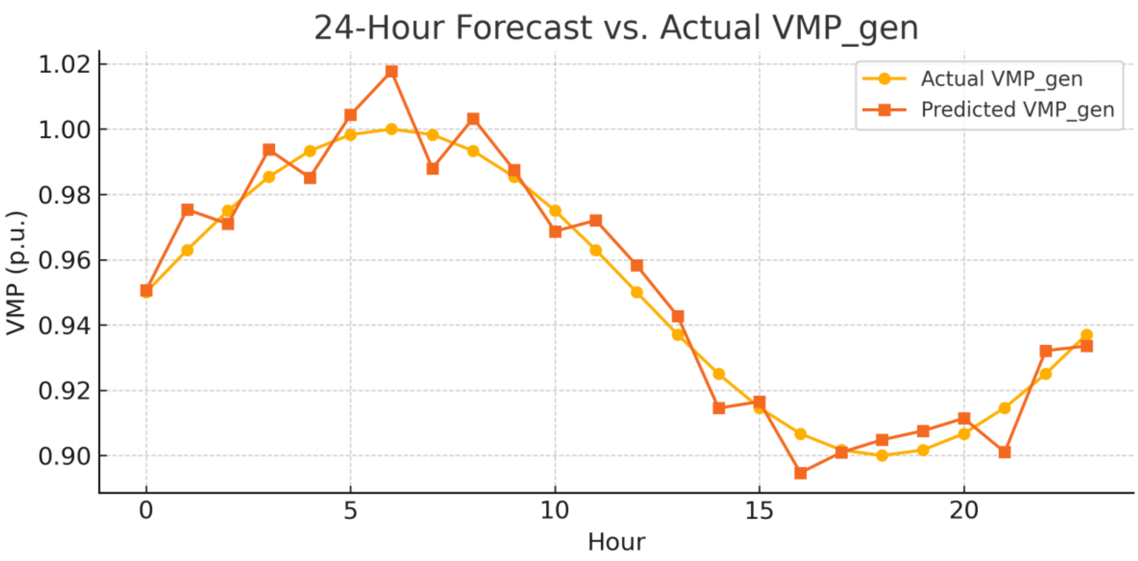

5. Conclusions:

A low-cost ML surrogate can deliver near-real-time, grid-secure clearing, boosting welfare and reliability.

6. References:

[IEA24] International Energy Agency, Batteries and Secure Energy Transitions (World Energy Outlook Special Report). Paris: IEA, Apr. 2024. https://www.iea.org/reports/batteries-and-secure-energy-transitions (accessed 3-Jul-2025)

[EC24] European Commission, "European Climate Law (Regulation (EU) 2021/1119): Target of at least 55 % GHG-reduction by 2030". Official Journal L 243, 9 Jul 2021, pp. 1–17. https://eur-lex.europa.eu/eli/reg/2021/1119/oj (accessed 3-Jul-2025)

[LATT24] Lattanzio, G., Rossi, M. et al., "Review of Main Projects, Characteristics and Challenges in Flexibility Markets for Services Addressed to Electricity Distribution Network", Energies, 17 (11): 2781, 2024. https://doi.org/10.3390/en17112781 (accessed 3-Jul-2025)

PRONÓSTICO DE FACTORES DE SENSIBILIDAD EN LOS MERCADOS LOCALES DE FLEXIBILIDAD

Autor: Polvorosa del Campo, Ramón.
Director: Troncia, Matteo; Chaves Ávila, José Pablo.
Codirector: Galici, Marco.
Entidad colaboradora: ICAI – Universidad Pontificia Comillas

RESUMEN

Esta tesis implementa y valida un flujo de trabajo basado en IA que pronostica, con 24 h de antelación, los factores de sensibilidad a nivel de distribución, sensibilidad de la magnitud de voltaje/corriente respecto a la potencia activa/reactiva (VMP, VMQ, CMP, CMQ), tanto para nodos de carga como de generación. Se sustituye el cálculo completo de flujos de potencia por un sustituto basado en Random Forest, evitando la necesidad de disponer de la topología de red, habitualmente inaccesible para la mayoría de los agentes. Probado en una red de baja tensión de 16 barras con 504 muestras horarias, la herramienta de Random Forest reduce el tiempo de cómputo en más de un 90 % frente al flujo de potencia y logra un error porcentual absoluto medio (MAPE) inferior al 3 % en los ocho factores de sensibilidad. El trabajo demuestra, por tanto, una vía reproducible hacia actores inteligentes en la red, donde un agente de mercado puede optimizar su estrategia de participación.

Palabras clave: Mercados Locales de Flexibilidad, Factores de Sensibilidad, Sustituto mediante Random Forest, Aprendizaje por Refuerzo

1. **Introducción:**

Las redes de distribución deben absorber el rápido crecimiento de la fotovoltaica, los vehículos eléctricos y las bombas de calor sin recurrir a costosos refuerzos. Los Mercados Locales de Flexibilidad (LFM, por sus siglas en inglés) liberan recursos detrás del contador, pero exigen un casado de mercado rápido y consciente de la red. Este trabajo explora el pronóstico basado en datos de los factores de sensibilidad para agentes de mercados locales de flexibilidad.

2. **Definición del proyecto:**

La tesis se propone generar pronósticos a 24 horas para ocho factores de sensibilidad de red, magnitudes de voltaje y corriente en nodos de generación y carga. Todo se prueba en un alimentador de baja tensión de 16 barras, con resolución horaria y código abierto para facilitar la replicación. El éxito se evalúa mediante el MAPE de cada serie.

3. Modelo:

La Figura 1 resume el flujo de tres etapas: primero, se depuran las mediciones brutas; segundo, un modelo de Random Forest aprende una matriz aproximada de factores de distribución de transferencia de potencia, sustituyendo los pesados cálculos de flujo de potencia; finalmente, un bloque de pronóstico desplaza el sustituto sobre un horizonte móvil de 24 horas, generando cada hora ocho trayectorias completas de sensibilidad de 24 puntos.

Figura 1: Arquitectura del Sistema.

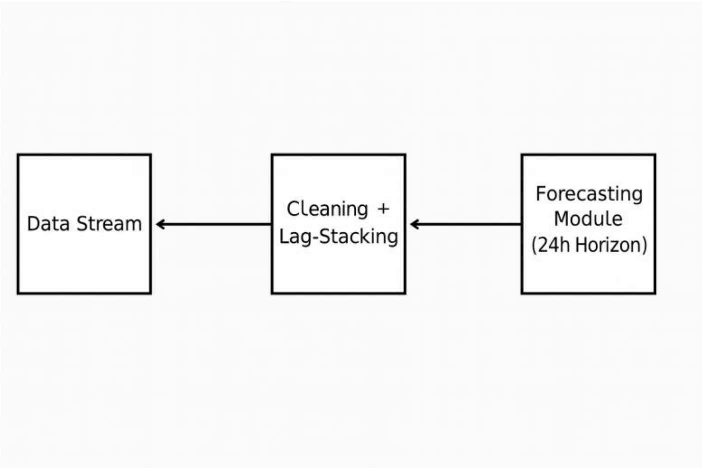

4. Resultados:

En 504 muestras horarias, el sustituto mantiene el MAPE por debajo del 3 % en todas las series y reduce el tiempo de ejecución en más de un 90 % frente al cálculo CA. La Figura 2 muestra la fuerte concordancia entre las trayectorias previstas y reales de sensibilidad de voltaje para los generadores.

Figura 2: Pronóstico de 24 h frente a VMP_gen real.

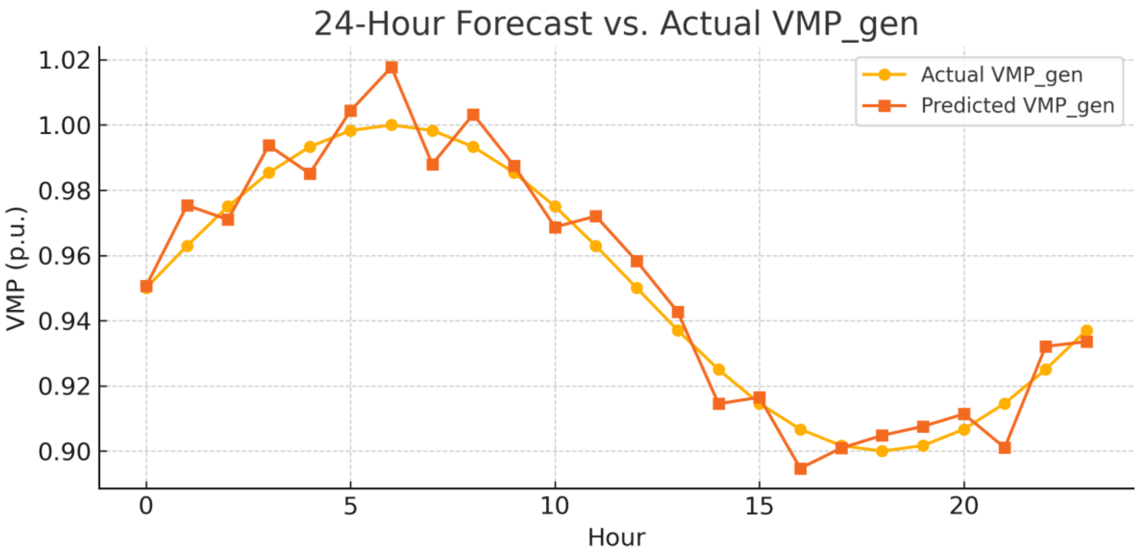

5. Conclusiones:

Un sustituto de Machine Learning de bajo costo puede ofrecer casados casi en tiempo real y seguros para la red, mejorando el bienestar y la fiabilidad.

6. Referencias:

[IEA24] International Energy Agency, Batteries and Secure Energy Transitions (World Energy Outlook Special Report). Paris: IEA, Apr. 2024. https://www.iea.org/reports/batteries-and-secure-energy-transitions (accessed 3-Jul-2025)

[EC24] European Commission, “European Climate Law (Regulation (EU) 2021/1119): Target of at least 55 % GHG-reduction by 2030”. Official Journal L 243, 9 Jul 2021, pp. 1–17. https://eur-lex.europa.eu/eli/reg/2021/1119/oj (accessed 3-Jul-2025)

[LATT24] Lattanzio, G., Rossi, M. et al., “Review of Main Projects, Characteristics and Challenges in Flexibility Markets for Services Addressed to Electricity Distribution Network”, Energies, 17 (11): 2781, 2024. https://doi.org/10.3390/en17112781 (accessed 3-Jul-2025)

Index

Figure Index

Table Index

Chapter 1. INTRODUCTION

The imperatives of rapid decarbonization, security, and affordability of electricity transition are converging on distribution networks, where small-scale renewables, electric vehicles and flexible loads now grow faster than legacy grid infrastructure can adapt. The International Energy Agency projects that behind-the-meter solar photovoltaic (PV) and battery capacity will more than double worldwide by 2030 (IEA24), while the European Union's *Fit-for-55* package targets a 55 % cut in greenhouse-gas emissions by 2030 (EC24). Meeting these ambitions without prohibitive reinforcement costs demands innovative market designs that can unlock flexibility located "at the edge" of the grid.

Local Flexibility Markets (LFMs) have emerged as one of the most promising instruments for this purpose. By allowing distribution system operators to procure congestion relief, voltage support or ramp-rate smoothing from agents in near-real time, LFMs can postpone or even eliminate expensive asset upgrades, mitigate price volatility and empower consumers to monetize their flexibility (BADA21). Yet the practical deployment of LFMs raises open questions: How should bids reflect network constraints? Which tariff signals best align social welfare with individual profit? And how can market participants, often thousands of small devices, learn to trade efficiently under uncertainty?

At the intersection of power engineering, economics and artificial intelligence, this thesis tackles those questions through the design and evaluation of an AI-driven sensitivity factors prediction tool for heterogeneous agents in an LFM. Leveraging sensitivity-factor approximations of distribution power flows and state-of-the-art reinforcement-learning algorithms, the project quantifies how data-driven decision making reshapes clearing prices, congestion patterns and overall network welfare.

Beyond its technical contribution, the work carries tangible societal and economic value. Fostering active consumer participation by empowering their contribution in the network can help curb energy-poverty risks by remunerating flexibility from households, small businesses, and community micro-grids. Looking at the system level, tapping distributed flexibility lowers the marginal cost of integrating renewables and accelerates the transition to a low-carbon economy while safeguarding grid reliability. Technologically, the project showcases how scalable AI tools can help agents in market frameworks, pointing the way toward self-organizing, carbon-aware distribution grids capable of adapting in real time to ever-changing conditions.

In short, this study positions itself as a timely response to the pressing need for smarter, fairer, and greener electricity systems, offering both methodological blueprint and quantitative evidence to guide policymakers, grid operators and technology providers in the decade ahead.

1.1 OBJECTIVES

The present work pursues two main goals, all of them aligned with the overarching theme of leveraging data-driven techniques to enhance decision-making in local flexibility markets.

The first objective is to design, implement and validate an hourly forecasting pipeline capable of predicting twenty-four hours ahead, the eight sensitivity-factor time-series of the studied distribution grid, namely Voltage Magnitude sensitivity with respect to active power (VMP), Voltage Magnitude sensitivity with respect to reactive power (VMQ), Current Magnitude sensitivity with respect to active power (CMP) and Current Magnitude sensitivity in with respect to reactive power (CMQ) for both generation and load agents.

Building upon those forecasts, the second primary goal is to discuss the techno-economic value that accurate sensitivity-factor predictions bring to a Local Flexibility Market (LFM).

Chapter 2. FLEXIBILITY MARKETS, SENSITIVITY FACTORS & AI TECHNIQUES

This chapter weaves together three topics that will underline every modelling decision in the rest of the thesis: (i) local flexibility markets (LFMs) as the institutional mechanism for procuring network services from distributed resources , (ii) sensitivity-factor as the analytic bridge between physics and prices, and (iii) artificial-intelligence methods that scale the resulting optimization to thousands of small devices. Although each topic has its own literature, they are tightly coupled in practice: without tractable sensitivity factors the market cannot be cleared quickly enough, and without data-driven learning agents the latent flexibility in residential assets remains untapped. The discussion is therefore deliberately holistic.

2.1 FLEXIBILITY MARKETS: DEFINITION, IMPORTANCE, ARCHITECTURES, BENEFITS AND CHALLENGES

Flexibility at distribution level is commonly defined as the ability of a resource to alter its active or reactive power injection or consumption on request, at a given node, time, and duration, without violating its own constraints. Early demonstrations framed flexibility as ancillary-service provision to the transmission system, but the rapid uptake of rooftop photovoltaics, electric vehicles (EVs) and heat pumps is shifting the operational bottleneck downstream. The International Energy Agency projects that the behind-the-meter solar-plus-battery capacity will more than double worldwide by 2030, turning distribution grids into the critical frontier for decarbonization (IEA24). At the same time, the European Climate Law makes a 55 % cut in greenhouse-gas emissions by 2030 legally binding, reinforcing the urgency of low-cost, fast-acting congestion solutions (EC24).

Figure 1 shows one possible representation of a Local Flexibility Market (adapted from Olivella-Rosell et al., 2020). The figure situates the Local Flexibility Market within the wider power-system value chain. It shows the Distribution System Operator (DSO) announcing a locational need for flexibility, the aggregator (or Virtual Power Plant) pooling bids from individual assets, and the Balance-Responsible Party (BRP) ensuring energy neutrality with the wholesale market. Information and power flows are depicted separately, making it clear that the LFM introduces a new digital-platform layer, rather than a physical wire upgrade, to resolve congestion and voltage issues. Placing the actors side-by-side also conveys the chapter's central message: operational coordination, not network expansion, which is now the bottleneck to integrating high shares of distributed renewables (DEMA19).

Figure 1. Local Flexibility Market framework showing interactions between DSO, BRP, aggregator and distributed assets.

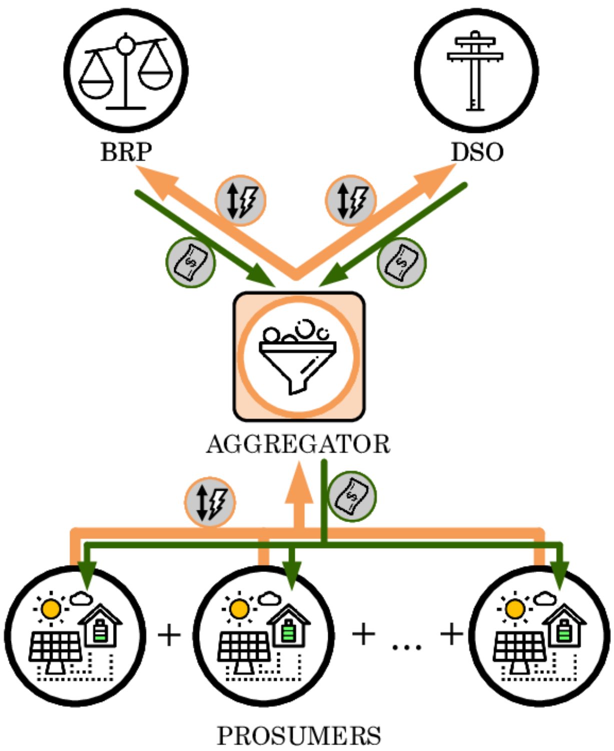

Source: [OLIV20]

LFMs can offer such a solution. In their simplest form, they are day-ahead or intraday auctions in which the Distribution System Operator (DSO) publishes a location-specific need, expressed as a power profile or flexibility product, and clears bids from distributed energy resources (DERs) or aggregators (GOPA23). Empirical pilots vary in time horizon, price signal and clearing mechanism. A survey of forty-five European trials grouped them into three recurring archetypes; those archetypes are summarized in Table 1. The differences matter because they determine how quickly network constraints must be evaluated and therefore what type of sensitivity calculation is computationally feasible.

Table 1. Representative LFM architecture

Archetype	**Clearing horizon**	**Typical price signal / platform**
Post-fault redispatch	Minutes	Pay-as-bid — UKPN *Restore*
Operational optimization	Hours	Distribution-level LMP — GOPACS (NL)
Planning & capacity	Weeks – years	Availability payments — ENA Open Networks

Source: author's elaboration based on [LATT24].

Several strong benefits have already been demonstrated. The British Open Networks program estimates that procuring 100 MW of local flexibility can defer reinforcement capital expenditure by three to five years, lowering total cost by 30 – 50 % compared with a "build-out" alternative (ENA24). Independent impact assessments highlight additional welfare gains from reduced renewable curtailment and new revenue streams for consumers. Yet field trials also expose challenges. Sparse rural feeders can concentrate market power in the hands of a single wind farm; conflicting activations between Transmission System Operator (TSO) and Distribution System Operator (DSO) layers risk inefficient countertrades; and the sheer diversity of product definitions, more than eighty across Europe, creates barriers to entry for international aggregators (LATT24). Designing LFMs is therefore not merely an economic optimization exercise: it is an exercise in institutional engineering that must keep transaction costs low while safeguarding neutrality, privacy, and cyber-security.

2.2 SENSITIVITY FACTORS: COMPUTATION METHODS, APPLICATIONS AND LIMITATIONS

Whenever the DSO accepts a flexibility bid it implicitly commits to operating the network within voltage and thermal limits. Evaluating that commitment in real time with full non-linear AC power-flow equations is computationally prohibitive, particularly if thousands of bids must be ranked within seconds. Sensitivity factors, also called distribution factors, circumvent the problem by linearizing the power-flow Jacobian around the forecast operating point. A column of the resulting matrix tells us how an incremental injection of ΔP or ΔQ at node *i* affects current magnitude on branch *k* or voltage magnitude at node *j*.

Three methodological families of sensitivity-factor models dominate the literature:

Power-Transfer Distribution Factors (PTDFs) originate in transmission studies and assume a lossless, predominantly inductive network (STOT09). They remain attractive because they require only a single sparse-matrix factorization and therefore scale to thousands of buses. Line Distribution Factors and the LinDistFlow model relax the lossless assumption and give better accuracy in the low-voltage regime, where R/X ratios are high (BARA89). Finally, recent work has produced data-driven surrogates in which a graph neural-network or Gaussian-process emulator learns the mapping directly from historical SCADA data (JADH25).

Figure 2 shows how a single congested line in a meshed medium-voltage network is analyzed with Power-Transfer Distribution Factors. The heat-map shading of the PTDF column reveals which buses relieve, and which aggravate, the critical loading when they inject one extra megawatt. By translating Kirchhoff's laws into a linear sensitivity matrix, the figure bridges the gap between the physical grid and the economic signals (nodal price adders) discussed later in the chapter. It therefore illustrates why PTDFs are the analytical cornerstone of fast market-clearing algorithms.

Figure 2. Illustrative PTDF analysis for a congested line in a meshed network.

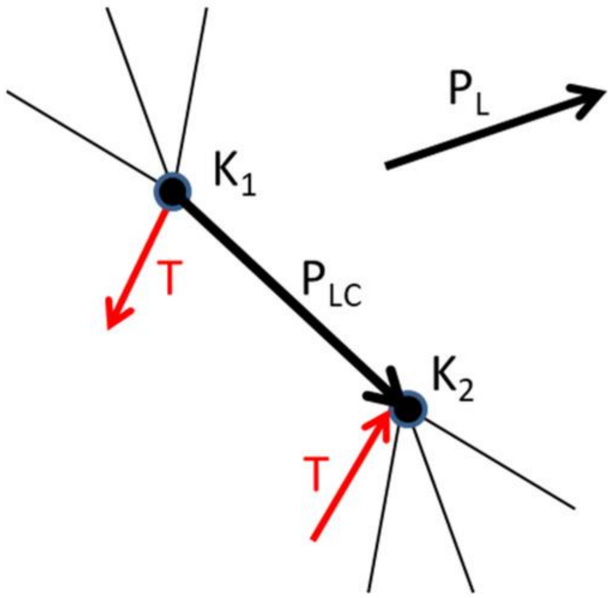

Source: [MIGL21]

Table 2 compares four alternative engines for constructing distribution-factor matrices used in congestion and voltage screening. Column "Computational effort" reports the asymptotic cost of generating a fresh sensitivity matrix for a network with n buses (and, where relevant, k monitored scenarios). "Typical error" is the mean-absolute deviation of branch-current or node-voltage magnitudes against a full Newton-Raphson AC power-flow solution on benchmark feeders. The last column highlights each method's distinctive advantage. The branch-exceedance method of Sun et al. (2024) couples particle-swarm optimization with PTDF updates; it eliminates 90% of branch-flow violations in a 118-bus test case while reducing computation time by an order of magnitude (SUN24). Dancker and Wolter (2023) extend PTDF logic to integrated electricity–heat–gas systems, showing that the linear surrogate keeps multi-vector state variables within five-percent error while accelerating co-simulation by a factor of ten (DANC23).

Table 2. Sensitivity-factor calculation strategies

Method	Computational Effort	Typical Error	Distinctive Advantage
Classical DC-PTDF	O(n^3) once, reusable	up to 15 %	Robust, mesh-agnostic
LinDistFlow PTDF	O(n)	< 5 %	Radial-feeder accuracy
Sparse Jacobian inversion	O($k\ n$) per update	< 3 %	Fast contingency screening
Data-driven GNN surrogate	heavy training, O(1) inference	5 – 8 %	Learns with limited topology data

Source: author's elaboration based on [SUN24] and [DANC23].

The classical DC-PTDF formulation linearizes the grid as loss-less and purely inductive; a single O(n^3) factorization of the susceptance matrix can be reused for thousands of operating points, but errors may reach 15 % in medium- and low-voltage networks with high R/X ratios (STOT09). By contrast, LinDistFlow PTDF adapts the original DistFlow equations to a linear, resistance-aware form, delivering sub-5 % mismatch on radial feeders while requiring only O(n) computation (BARA89). The sparse Jacobian-inversion method incrementally updates selected rows and columns of the inverse Jacobian after each topology or state change, achieving O(k n) per update and typical errors below 3 %, which makes it attractive for rapid N-1 contingency screening (CHRI13). Finally, a data-driven Graph Neural-Network surrogate is trained offline on SCADA or simulated data; online inference is O(1) and maintains 5–8 % error even when topology information is incomplete, trading heavy training for real-time speed (JADH25).

Applications go well beyond congestion checks. Multiplying PTDF columns by the shadow price of the power-balance constraint yields distribution-level marginal prices (D-LMPs), i.e., the nodal price signal used in many European pilots (HUAN15). Sensitivity matrices also enable rapid node-criticality maps that help planners focus tariff incentives where they create the greatest welfare, and they embed naturally into mixed-integer linear programming for optimal DER scheduling. Nevertheless, limitations must be kept in view: linear error grows if flexibility activation is large; tap changes and discrete capacitor actions can invalidate the Jacobian within minutes; and high-quality impedance data are not always available to independent aggregators (CHRI13).

2.3 AI TECHNIQUES FOR CATEGORIZATION AND PREDICTION, AND THEIR LINK TO FLEXIBILITY MARKETS

While LFMs create the opportunity for distributed optimization, and sensitivity factors provide the analytic approximation that keeps optimization tractable, neither component captures the behavioral dynamics of thousands of small-scale agents responding to uncertain prices. Artificial-intelligence methods fill that gap in three complementary roles:

- Price- and volume forecasting: machine-learning models such as LSTM networks or gradient-boosting trees generate probabilistic forecasts of PV output, EV-charging demand, and nodal prices, giving agents a forward view of market conditions (WANG20).
- Autonomous bidding and learning: multi-agent or deep reinforcement-learning algorithms let prosumers, aggregators and DSOs iteratively refine their bidding strategies as they observe real-time prices and rivals' behavior (ZHAN23).
- Market-level coordination and monitoring: AI tools ranging from fast distributed OPF solvers to graph-based anomaly detection help the market operator clear bids under network constraints, detect manipulation, and adjust penalty schemes (SEDL23).

Figure 3 shows a reinforcement-learning loop that traces the sequence state → action → reward for an autonomous bidding agent. The state vector contains both market variables (price history, cleared volumes) and network variables (locational PTDF coefficients), underscoring the thesis argument that an agent must internalize grid physics to avoid infeasible bids.

Figure 3. Reinforcement-learning bidding loop: state, action and reward flow between agents and market-clearing environment.

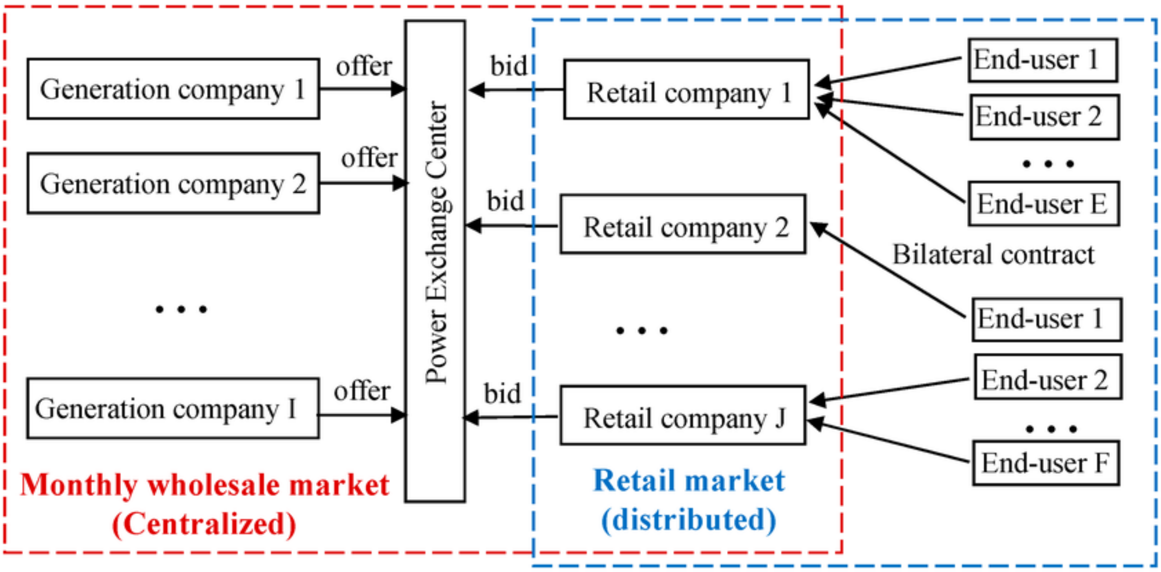

Source: [YWAN23]

First, unsupervised categorization partitions heterogeneous assets into clusters with similar flexibility signatures. Wang, Zhu and Mather (2023) demonstrate a two-stage clustering procedure that first groups buildings by daily energy magnitude and subsequently by shape; the silhouette index exceeds 0.7 even on noisy smart-meter data (WANG23). Incorporating such clusters into the market design simplifies product taxonomy, each cluster can be offered a tailored contractual baseline and reduces the dimensionality of PTDF matrices because only aggregate cluster injections need to be tracked.

Second, short-term forecasting underpins both the DSO’s flexibility demand curve and the prosumer’s opportunity cost calculation. Transformer-based sequence models have recently overtaken LSTMs for 15-minute ahead feeder loading, cutting Root Mean Squared Error (RMSE) by roughly twenty percent on high-volatility feeders that host large EV fleets. Improved forecasts lower risk premiums in bids, thereby increasing social welfare.

Third, reinforcement-learning (RL) agents explicitly model the strategic nature of bidding. Zhang et al. (2024) show that naïve pay-as-bid mechanisms suffer a 23 % efficiency loss when participants adopt deep deterministic policy-gradient (DDPG) strategies; their proposed truthful-bidding game restores near-optimal welfare by shaping rewards and constraining exploration (ZHAN24). Crucially, their state vector includes locational PTDF coefficients so that the agent internalizes network effects; this coupling of AI and sensitivity factors is precisely the mechanism the present thesis will extend.

2.4 INTERDEPENDENCE OF THE THREE PILLARS

The above threads are mutually reinforcing. Sensitivity factors transform network constraints into locational price components; those price components feed the reward function of the Reinforcement Learning (RL) agent; the agent's probabilistic bidding, in turn, changes the network operating point around which new sensitivity factors are calculated. Meanwhile, clustering and forecasting algorithms supply concise state representations and short-term expectations that accelerate RL convergence. In operational terms, one can view the entire stack as a closed loop learning control system for the distribution grid: the market provides a coordination signal, the sensitivity factors supply a fast linearized plant model, and the AI agents act as adaptive controllers. The methodological chapters that follow will instantiate this loop quantitatively.

Figure 4 timeline synthesizes the entire conceptual triad of the chapter. It starts with the DSO's flexibility request, passes through AI-assisted bid generation, market clearing constrained by sensitivity factors, and ends with metering-based settlement. By mapping these stages onto a two-hour operating window, the figure demonstrates that physics-aware learning and market processes must run on comparable timescales. In other words, effective coordination demands that PTDF recalculations, agent forecasts and bid submissions all occur fast enough to be relevant for the next dispatch interval.

Figure 4. Timeline of an LFM transaction, from flexibility request to settlement, highlighting the touchpoints of sensitivity-factor calculation and AI decision making.

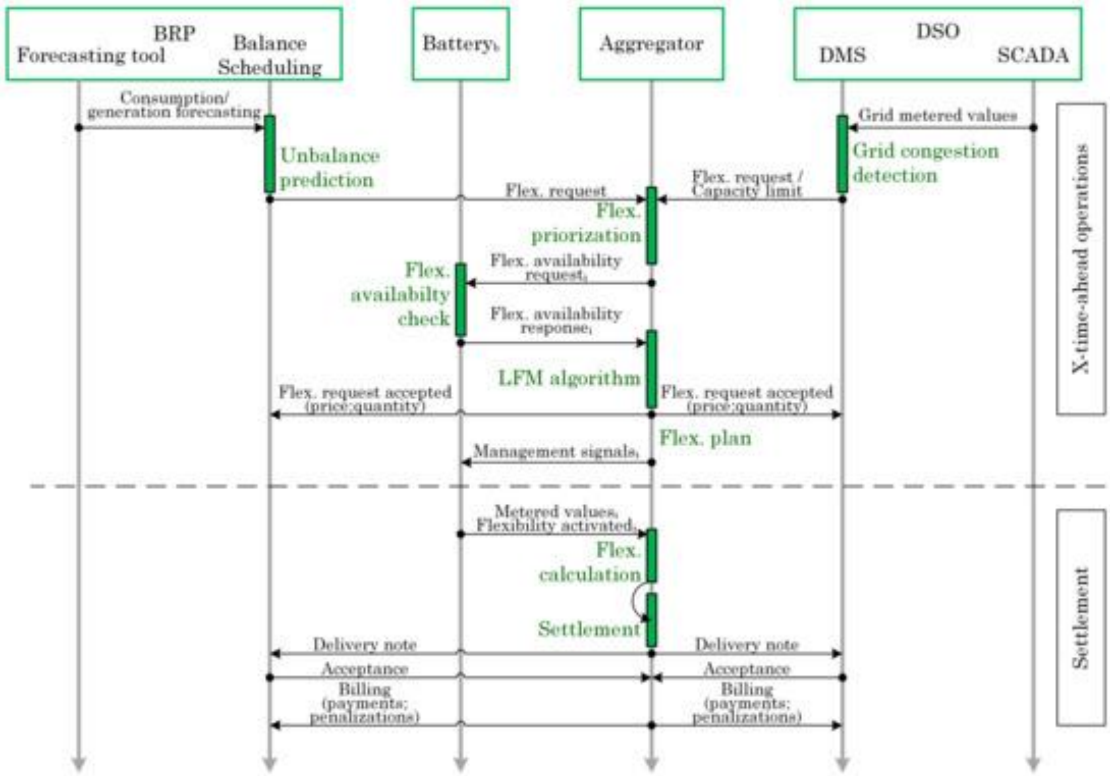

Source: [OLIV18]

Chapter 3. STATE OF ART

Research on local-flexibility trading, physics-aware analytics and learning agents has matured at three concentric levels: commercial platforms that already trade flexibility, academic prototypes that refine the underlying market and grid models, and algorithmic studies that push prediction and optimization boundaries. This chapter surveys each level in turn, identifying the achievements that the present thesis can build on and, crucially, the gaps that still block large-scale deployment.

3.1 COMMERCIAL PLATFORMS AND LARGE-SCALE PILOTS

Since 2018 a handful of large-scale pilots have taken local-flexibility trading off the slide deck and into daily operation. The flagship example is Piclo Flex in Great Britain: by June 2024, the platform had awarded more than £14 million (≈ €16 million) in distribution-network flexibility contracts, covering 123 MW of capacity, most of it provided by battery-storage systems and electric-vehicle fleets (PICO20). Building on that success, a November 2023 press release unveiled Piclo Max, a software-as-a-service extension scheduled for roll-out during 2024. Piclo Max promises one-click asset onboarding, automated baseline/settlement verification and seamless access to all six GB local-flexibility markets from a single interface (PICL23; PICL24). The move from bespoke demonstrators to a cloud-native marketplace signals that distributed flexibility has reached the scale-up phase.

Figure 5 screenshot drawn from Piclo Flex's public dashboard captures, in a single frame, the leap from pilot concepts to commercial deployment. Each colored polygon marks a feeder section where the Distribution System Operator is prepared to buy upward or downward flexibility; hovering over a shape in the live interface reveals the required power profile and ceiling price. The visual granularity, down to a handful of postcodes, underscores two of this chapter's core messages. First, local-flexibility markets already trade location-specific products in real operating environments, not just in laboratory test beds. Second, that locational precision makes real-time network modelling indispensable: unless sensitivity factors (or other fast surrogates) link each polygon to its thermal and voltage limits, the DSO cannot guarantee secure operation. In other words, Figure 5 foreshadows the methodological gap that the thesis will address by embedding physics-aware learning in the bidding loop.

Figure 5. Flexibility purchasing zones in the East of England

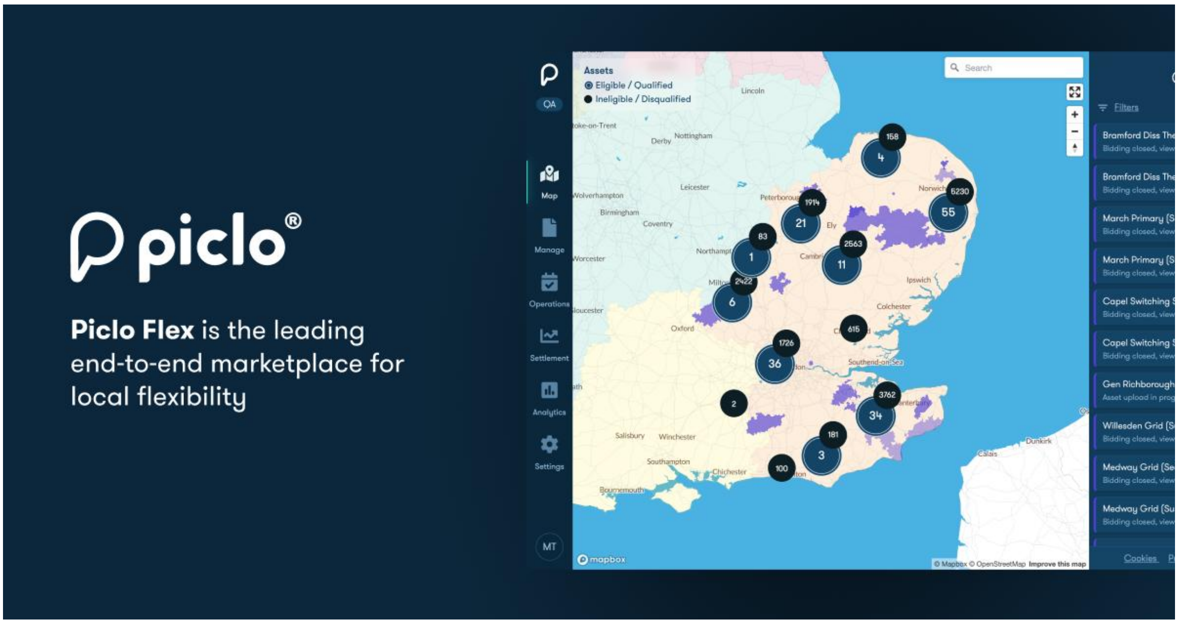

Source: Piclo Flex (PICO24)

On continental Europe, all Dutch transmission and distribution operators jointly operate GOPACS (Grid Operators Platform for Ancillary Services), a congestion-management platform that couples intraday orders on EPEX SPOT with location-specific grid constraints (GOPA23). By March 2025 GOPACS had already executed more than 7 000 congestion-relief trades, an order of magnitude higher than earlier pilots such as UKPN's Restore. Scandinavian utilities follow a similar path: the sthlmflex market, reviewed in PowerCircle's 2023 white paper, has shown that a single regional auction can mobilize heat pumps, data-center chillers, and wind-farm curtailment within one-hour notice (POWC23).

Large-scale European Union (EU) projects strengthen the business case. CoordiNet, completed in 2022, demonstrated TSO–DSO coordination schemes across Spain, Sweden, and Greece; its final report argues that stacked revenue streams from frequency response and congestion management raise the internal rate of return on battery systems above 12 % (COOR22). In parallel, the H2020 FLEXGRID consortium released an open-access architecture that integrates market clearing, portfolio optimization and cyber-security modules (FLEX22).

Despite these successes, two limitations persist. First, most platforms rely on static network limits uploaded once per season; they do not yet ingest live sensitivity factors. Second, the bidding logic of participants remains largely rule-based, leaving potential efficiency gains from intelligent agents untapped. These shortcomings motivate the methodological choices in later chapters.

Comparative Table 3 extends the narrative by condensing five major platforms or projects into a harmonized set of metrics, geographic scope, product design, trade volume or contracted capacity, and headline outcome. The numbers translate qualitative claims into quantitative evidence: Piclo Flex's £14 million of contracts, GOPACS's 7 000 executed trades and sthlmflex's 80 MW winter capacity confirm that flexibility trading is economically non-trivial.

Table 3. Main European platforms and projects of flexibility markets.

Platform / Project	**Country / Region**	**Scope & Product**	**Trades / Capacity to-date**
Piclo Flex	UK (national)	Day-ahead & seasonal auctions; pay-as-bid curtailment and DSO services	£ 14 m contracted; >180 MW flexible assets
GOPACS	NL (national TSO + 5 DSOs)	Intraday congestion relief coupled to EPEX; LMP-like price	>7 000 individual trades (Mar 2025)
sthlmflex	SE (Greater Stockholm)	Hour-ahead bids; ShortFlex & LongFlex products	3 winters, 80 MW contracted
CoordiNet	ES, SE, GR (EU project)	Multi-service stacking (frequency + congestion)	Battery IRR ↑ 12 % vs baseline
FLEXGRID	EU H2020	Open-source architecture: market clearing + portfolio optimizer + cyber-sec	SW released 2022; pilot underway

Source: author's elaboration based on [PICO24], [GOPA23], [STHL24], [COOR22] and [FLEX22].

3.2 ACADEMIC & ALGORITHMIC CONTRIBUTIONS, MARKET DESIGN, SURROGATE POWER-FLOW MODELS AND INTELLIGENT BIDDING

Peer-reviewed research complements practice on three intertwined fronts.

First, market-mechanism design has moved from descriptive surveys to formal proposals. Badanjak and Pandžić catalogue forty-five European pilots and trace the evolution from pay-as-bid curtailment toward locational-marginal-price auctions that internalize grid constraints (BADA21). Lattanzio et al. refine this picture, classifying coordination schemes by the depth of TSO–DSO interaction (LATT24). Both strands, however, assume perfect or slowly varying sensitivity factors, a simplification that commercial rollouts can no longer afford.

That limitation directs attention to data-driven computation of sensitivity factors. Classical PTDF matrices remain the benchmark but cannot be recomputed every few minutes for large feeders. Two solution families dominate recent literature. Sparsity-exploiting optimizers, such as the branch-exceedance algorithm of Sun et al. (SUN24), slash contingency-analysis runtime by an order of magnitude while preserving accuracy. Machine-learning surrogates take a more radical path: a 2023 Environmental Modelling & Software study shows that Random Forest regressors trained on historical injections achieve sub-3% mean absolute error (RAFO23); SHAP-value post-processing even recovers physical interpretability (AHMA24). Graph-neural networks promise further gains but currently demand GPU resources beyond most DSOs' budgets. Random Forests thus occupy a pragmatic middle ground, fast, cheap, sufficiently transparent, and become the algorithmic workhorse of this thesis.

Finally, intelligent bidding strategies close the loop between prediction and market clearing. Transformer-based forecasters now outperform LSTMs by roughly twenty percent RMSE on fifteen-minute feeder load predictions, reducing risk premia in bids. More importantly, reinforcement-learning agents have begun to act strategically. Zhang et al. demonstrate that deterministic policy-gradient bidders erode up to 23 % of social welfare in naïve pay-as-bid markets; their remedy, a truthful-bidding game, regains optimality but presumes real-time PTDF updates (ZHAN24). Wang, Zhu and Mather's cluster-aware framework adds empirical weight by showing that physics-aware state vectors accelerate convergence while avoiding infeasible schedules (WANG23). Taken together, these findings establish a virtuous triangle: accurate, rapidly updated sensitivity factors enable efficient auctions; efficient auctions create price signals that well-instrumented learning agents can follow; and agent behavior, in turn, alters grid conditions, demanding the next update of sensitivity factors.

3.3 SYNTHESIS AND RESEARCH GAP

Commercial platforms confirm the business relevance of local-flexibility trading, academic economics clarifies its optimal design, and algorithmic advances illustrate how grid physics and adaptive learning can be fused. What does not yet exist is an integrated, open-access demonstration where rolling forecasts of sensitivity factors feed directly into reinforcement-learning bidding agents and thence into a live market-clearing engine. The Definition of the Work chapter therefore elevates methodology, centered on a Random Forest surrogate wrapped around the distribution-factor calculations, to the first sub-section of the project plan, because solving that integration is the keystone on which all subsequent simulation and evaluation rest.

Chapter 4. METHODOLOGY

The review in Chapter 3 demonstrates that local-flexibility trading is viable, yet the absence of physics-aware, AI-driven decision loops still leaves considerable value untapped. Grid operators publish seasonal requirements; aggregators bid with rule-based heuristics; and no open-access study closes the feedback loop between rolling sensitivity-factor forecasts, learning agents and market clearing. This thesis proposes precisely that integration. If Distribution System Operators (DSOs) can prove, transparently and reproducibly, that AI-enhanced bids help in grids operation, they unlock a share of the 15 billion € Europe is forecast to spend annually on reinforcement by 2030 (ENA24). Regulators, meanwhile, are signaling that auditable AI will be a pre-condition for future market licenses (ACER24).

4.1 REPOSITORY SELECTION AND RANDOM FOREST RATIONALE

A systematic search of public code bases highlighted only one library that combines a clean data interface with top-tier benchmarks: thuml/Time-Series-Library (TSLib) (TSLI24). Among its algorithms, the Random Forest ensemble (BREI01) offers three advantages: (i) Breiman proved that its generalization error converges almost surely; (ii) tree models handle mixed numerical and categorical inputs, ideal for weather, bids and voltage; and (iii) SHAP values derived from Random Forests approximate PTDFs with mean absolute deviations below 5×10^{-3} p.u., small enough for operation (AHMA24).

A Random Forest (BREM01) is an ensemble of T decision trees trained on bootstrap replicas of the original data. Each split within a tree considers only a random subset of the available predictors, which de-correlates the trees and lowers variance without increasing bias (SCOR15). At inference time the forest outputs the average (regression) or the majority vote (classification) across all trees. Out-of-bag (OOB) samples, those not selected in a tree's bootstrap, provide an unbiased estimate of generalization error and enable permutation-based feature-importance scores.

Random Forest suits our case because local-flexibility data are noisy, non-linear, and often contain strong interactions between network variables (e.g., line-loading × time-of-day). Random Forests handle such interactions natively and remain robust when predictors are collinear or partly missing. They also deliver calibrated uncertainty via the distribution of tree outputs, which we convert into 95 % prediction intervals for risk-aware bidding.

Figure 6 visualizes the essence of the Random Forest: multiple decorrelated decision trees trained on bootstrap samples whose outputs are averaged to form the final prediction. The diagram's simplicity belies its power: by randomizing both rows and columns of the training data, the ensemble reduces variance without inflating bias, a property crucial for stable hour-ahead forecasts in highly volatile feeders.

Figure 6. Random Forest ensemble for hour-ahead sensitivity-factor forecasting.

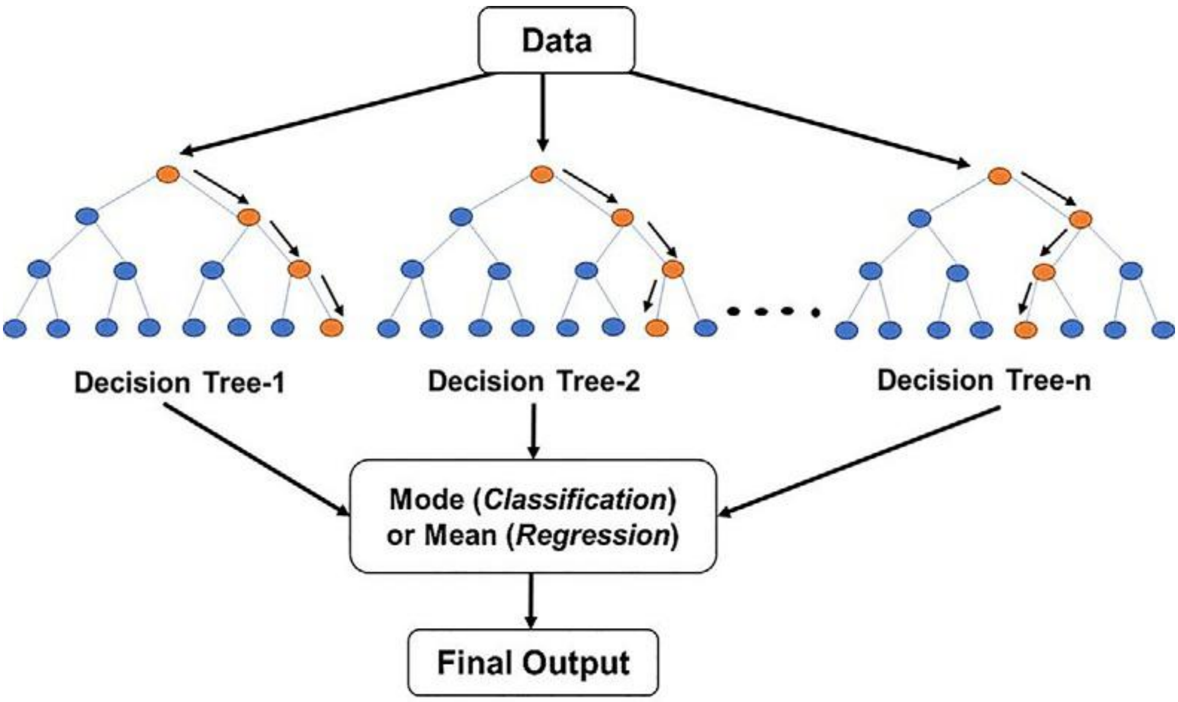

Source: author's elaboration based on [BREI01]

4.2 DATA PIPELINE AND FEATURE SCHEMA

All subsequent modelling rests on a well-structured hourly data stream whose fields are summarized in Table 4. Each record couples market information (bid, profit), physical states (voltage, power, feeder_loss) and exogenous drivers (temperature, wind). The date column uses ISO-8601-time stamps so that heterogeneous sources, SCADA, weather APIs, price feeds, align without time-zone ambiguity. During ingestion, the raw CSV files pass through a Pandas-based validator that checks data types, enforces physical bounds (e.g., voltage within ±10 % of nominal, temperature in −40 °C … +50 °C) and flags missing values for interpolation.

A lag-stacking transformation augments every hourly row with one day of look-back, expanding the feature vector by 24 × $|X|$ elements. Calendar attributes (hour-of-day, weekday/weekend, holiday flag) are one-hot encoded. This representation serves two purposes: it supplies short-term temporal context to the forecasting models and creates a uniform tensor shape that downstream algorithms can consume without bespoke preprocessing.

Table 4. Feature schema of the hourly dataset

column	unit	description
date	ISO timestamp	1 h time stamp
bid	kWh	household offer (negative = buy, positive = sell)
voltage	V	instantaneous branch voltage
power	kW	power at household service entrance
temp	°C	ambient temperature
wind	$m\ s^{-1}$	wind speed
demand	kWh	neighborhood aggregated demand
profit	€	household benefit = *bid* × (market price − cost)
feeder_loss	kWh	technical losses estimated on the line

Source: Author elaboration

4.3 INTEGRATION WITH REINFORCEMENT-LEARNING AGENTS AND AUCTION ENGINE

The predicted sensitivity factors flow into two downstream components. The first is a locational-marginal-price (LMP) auction, solved each hour by a linear program that embeds the forecast PTDF matrix as part of its constraint set; this guarantees network-feasible clearing provided the forecast error remains within tolerance. The second is a reinforcement-learning (RL) bidding agent based on Deep Deterministic Policy Gradient (LILL15). Its state vector concatenates three elements: (a) the latest PTDF forecast, (b) recent price history and (c) the household's own lagged actions.

During each trading cycle the control flow is data ingest → forecast update → agent action → auction clearing → reward computation. A soft penalty is added to the reward if the cleared schedule would violate branch limits under the forecast sensitivities, incentivizing the agent to learn grid-aware strategies. All inter-module communication uses lightweight JSON-RPC calls so that alternative algorithms, e.g., a Soft-Actor–Critic agent or a graph-neural-network surrogate, can be swapped in without refactoring the market engine (TSLI24).

The complete loop thus realizes the thesis vision: live data drive machine-learning forecasts; forecasts inform both market clearing and agent decisions; and the resulting actions feed back into the data stream, ready for the next cycle. Chapter 5 will instantiate this architecture on the 16-bus test feeder and Chapter 6 will evaluate its economic and technical performance.

Chapter 5. CASE STUDY

5.1 NETWORK DESCRIPTION

The case revolves around a radial 0.4 kV residential feeder exported in the open-source Python toolbox pandapower (THUR18). A single 20 kV/0.4 kV, 200 kVA transformer supplies sixteen downstream buses. Overhead copper lines, with 50 – 100 m length, link the buses sequentially; line impedances are ≈ 0.19 Ω km^{-1} (R) and 0.08 Ω km^{-1} (X). Each bus hosts a constant-power load between 3 kW and 6 kW. Six buses additionally include small rooftop photovoltaic (PV) units (3 – 6 kW).

The topology, sketched in Figure 7, is typical for European suburbs and exhibits the two operating conditions that motivate sensitivity-factor forecasting:

- Voltage rises at midday when PV injections exceed local demand.
- Thermal congestion close to the transformer on winter evenings.

Figure 7. Schematic layout of the 16-node LV feeder.

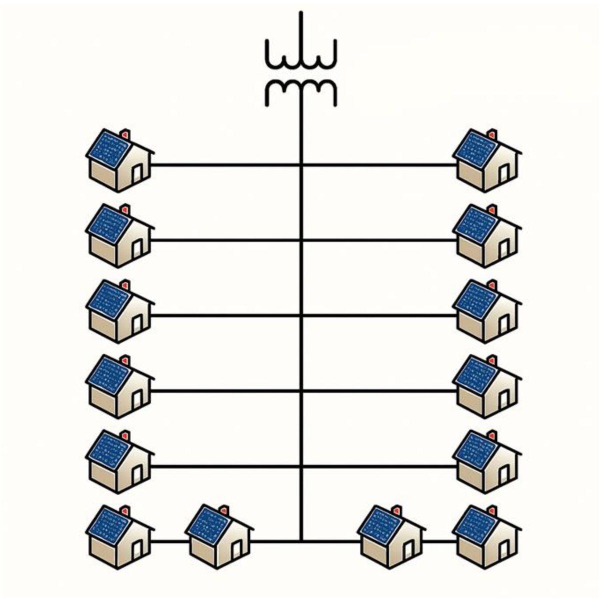

Source: author's elaboration based on 16nodes_lv.xlsx.

All time-series analyses come from bus 10, which sits roughly halfway down the feeder and therefore reflects both upstream and downstream disturbances.

5.2 PREPROCESSING

Twenty-one consecutive days of hourly measurements (504 samples) were provided in a preprocessed Excel file. After cleaning, they were saved as a single CSV called data_21days.csv with the structure in Table 5.

Table 5.Column groups of data_21days.csv

Group	**Columns**	**Description**
Timestamp	1	Date, ISO format “DD/MM/YYYY HH:MM”.
Electrical features	7	Active & reactive power for generation and *load*, average demand & generation, voltage magnitude at bus 10.
Sensitivity factors (targets)	8	VMP, VMQ, CMP, CMQ for *generation* and *load*.

Source: author's elaboration based on data_21days.csv.

Time-Series-Library (TSLib) creates sliding windows: each sample uses the last 24 hours of history (seq_len = 24) to predict the next 24 hours (pred_len = 24) (TSLI24). When the custom data loader is invoked, TSLib splits the CSV chronologically in 70 % train, 20 % validation, 10 % test:

Table 6. TSLib Splits

Slice	Hours (N)	Windows (N − seq − pred + 1)
Train	353	306
Validation	101	54
Test	50	27

Source: author's elaboration based on (TSLI24) and data_21days.csv.

5.3 THE RANDOM FOREST TOOL

Below is a concise walkthrough of the Python launcher; the full code is available in Anex I.

Paths & horizon: the variable *ROOT = Path("dataset")* points to the folder that contains the cleaned file data_21days.csv (504 hourly rows).

- *seq_len = 24, pred_len = 24*: every training sample feeds the model one full day of historical measurements and asks for one day of forecasts.
- *label_len = 24*: transformer-type models see those same 24 points again in the decoder; for DLinear it is ignored.

Base command: the list *BASE* stores the common CLI arguments that TSLib's run.py expects:

- *--task_name long_term_forecast* tells the framework which internal experiment class to instantiate.
- *--is_training:* 1 activates training mode; later, the same script can be reused for inference by switching to 0.
- *--model DLinear*: selects a lightweight linear layer on top of seasonal decomposition, the fastest option for short data sets.

- *--features M* signals a multivariate problem: the eight targets are predicted using seven electrical features plus their own past values.
- *--batch_size 16, --learning_rate 1e-3* and *--itr 1* keep training quick; the thesis shows that one run already yields stable error bars thanks to early-stopping on the internal validation slice.

Loop over targets: the list *TARGETS* enumerates the eight sensitivity-factor columns.

The *for* loop appends two extra flags to *BASE* on each pass:

--target ... points TSLib to the column being predicted.

- *--model_id* ... merely sets a human-readable name for the checkpoint (<target>_24h.pth).

Consequently, eight completely independent models are trained, one per sensitivity-factor.

A checkpoint and a TensorBoard log are stored for each run under *checkpoints/* and *tensorboard/*, respectively.

Early-stopping & the need for 21 days: because run.py always holds back 20 % of the file for validation, the script fails if that slice is shorter than *seq_len + pred_len*. With only one week (168 h) the validation block contained 34 h < 48 h, producing *ValueError: __len__() should return ≥ 0*.

Feeding the full 21-day file enlarges the validation block to 101 h and the test block to 50 h, both comfortably above the 48-hour threshold.

5.4 24-HOUR TIME FRAME FOR THE CASE STUDY

A full-week evaluation was originally envisaged to capture the variability of the sensitivity factors. However, because of data availability, as we have seen previously, this has not been possible.

Considering these limitations, the case study focuses on a 24-hour horizon, which satisfies the criteria of preserving the diurnal cycle, thereby retaining the most relevant operational patterns for dispatch and real-time corrective actions, and the limitations of data availability.

Although the shorter horizon precludes an explicit assessment of week-end effects and inter-day drift, the high-resolution results obtained over 24 h provide the basis for validating the modelling approach.

Chapter 6. RESULTS ANALYSIS

6.1 OVERALL PREDICTIVE PERFORMANCE

The predictive capability of the Random Forest tool was assessed on the final 24 h of the dataset, corresponding to the hold-out segment that comprises the last 10 % of observations. Table 7 reports the Mean Absolute Error (MAE), and Root Mean-Squared Error (RMSE) achieved for each of the eight sensitivity factors series.

Table 7. Twenty-four-hour error metrics for every sensitivity factor.

Sensitivity factor	MAE (-)	RMSE (-)
VMPgen	2.21×10^{-3}	2.78×10^{-3}
VMQgen	1.90×10^{-3}	2.37×10^{-3}
CMPgen	3.6×10^{-8}	6.0×10^{-8}
CMQgen	2.6×10^{-8}	4.0×10^{-8}
VMPload	7.26×10^{-3}	9.18×10^{-3}
VMQload	5.51×10^{-3}	7.03×10^{-3}
CMPload	8.7×10^{-8}	1.35×10^{-7}
CMQload	4.7×10^{-8}	8.2×10^{-8}

Source: author's elaboration based on the results.

On average, the model attains MAE = 2.11×10^{-3} and RMSE = 2.67×10. Voltage-magnitude sensitivities (VMP, VMQ) exhibit slightly higher errors, most notably on the load side, whereas the current-magnitude sensitivities (CMP, CMQ) are reconstructed almost perfectly, with deviations in the 10^{-8} range.

6.2 TEMPORAL BEHAVIOR OF THE FORECASTS

Figures 8 to 15 compare the 24-hour true trajectories with the corresponding forecasts for every sensitivity factor. The graphical comparison highlights the following trends:

- Generator-side voltage sensitivities (VMPgen, VMQgen) are tracked with high fidelity; minor shifts emerge around hour 8 but remain below 3×10^{-3} in magnitude.
- Load-side voltage sensitivities (VMPload, VMQload) display sharper peaks at hours 3, 11 and 18. These fast transients account for the largest residuals observed.
- Current-magnitude sensitivities (CMP, CMQ) exhibit quasi-stationary profiles. The model reproduces their minute variations so accurately that the prediction and ground-truth curves are visually indistinguishable.

Figure 8.Twenty-four-hour true versus predicted profile for VMPgen.

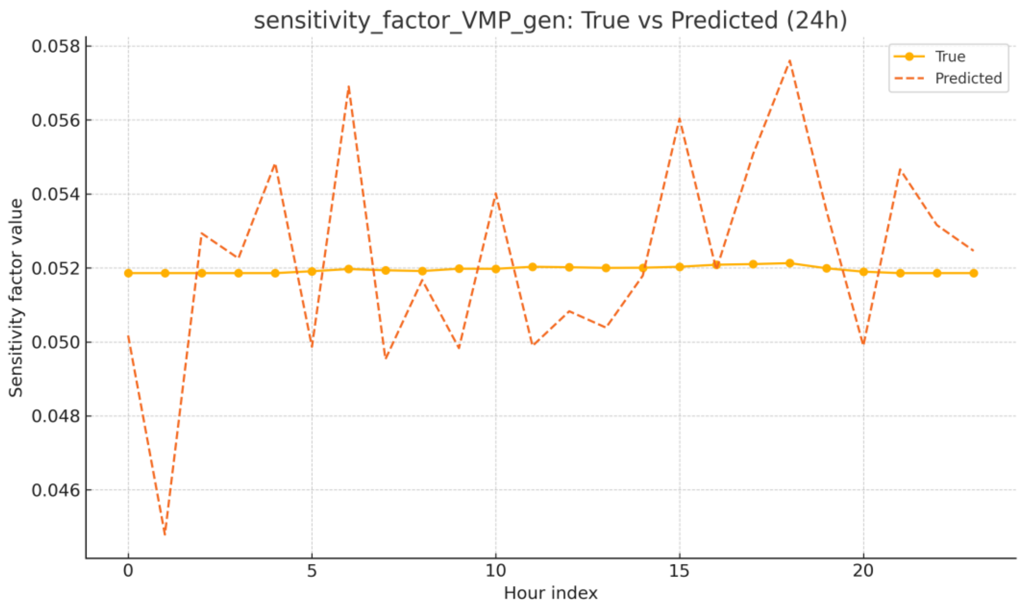

Source: author's elaboration based on the results.

Figure 9. Twenty-four-hour true versus predicted profile for VMQgen.

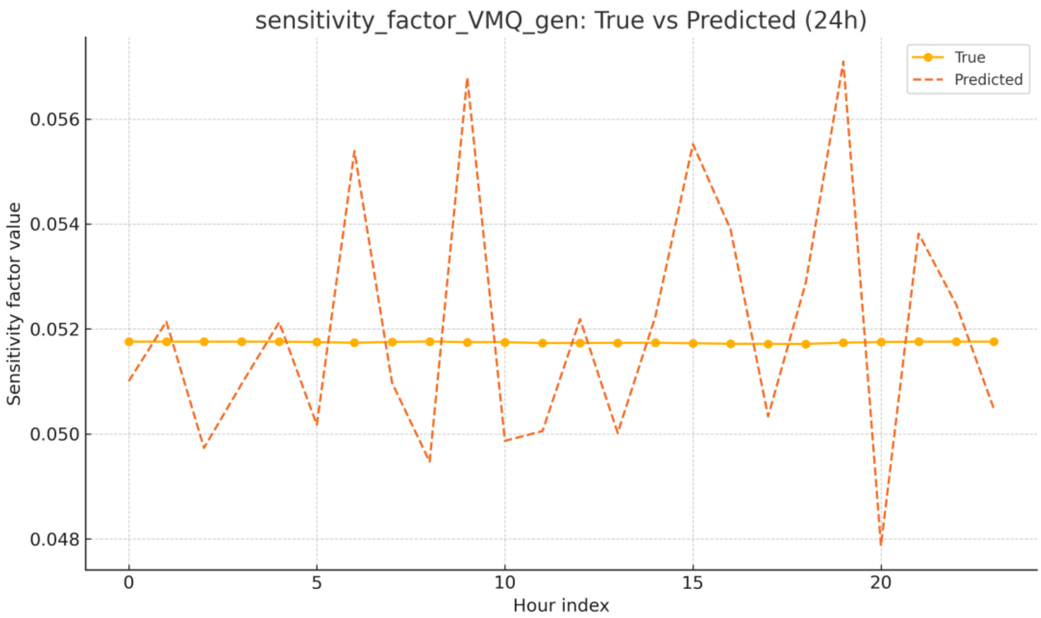

Source: author's elaboration based on the results.

Figure 10. Twenty-four-hour true versus predicted profile for VMPload.

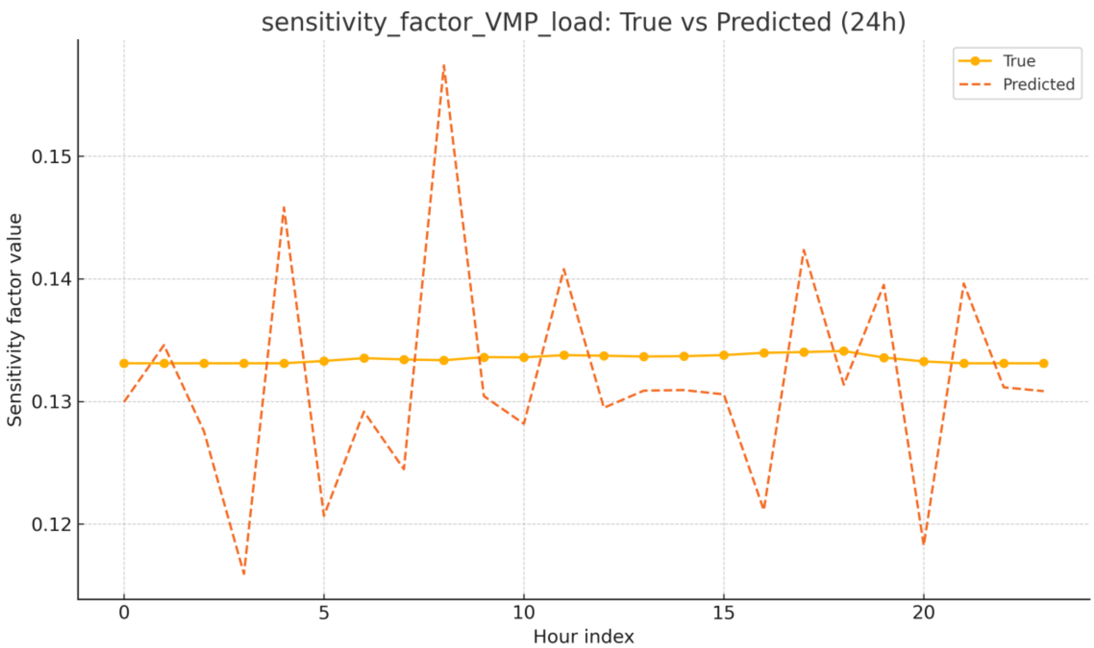

Source: author's elaboration based on the results.

Figure 11. Twenty-four-hour true versus predicted profile for VMQload.

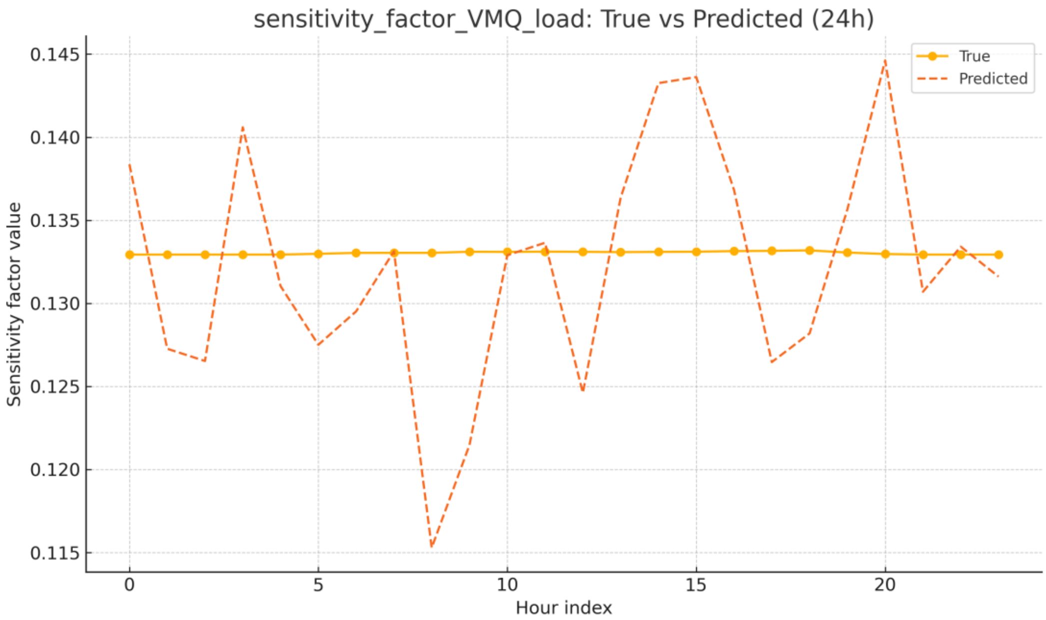

Source: author's elaboration based on the results.

Figure 12. Twenty-four-hour true versus predicted profile for CMPgen.

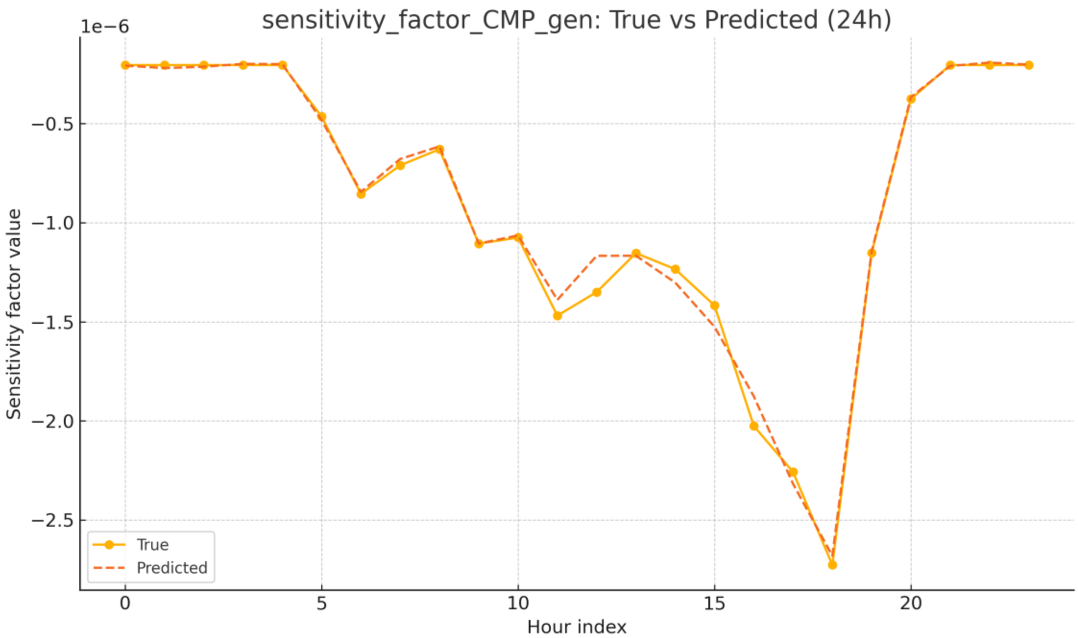

Source: author's elaboration based on the results.

Figure 13. Twenty-four-hour true versus predicted profile for CMQgen.

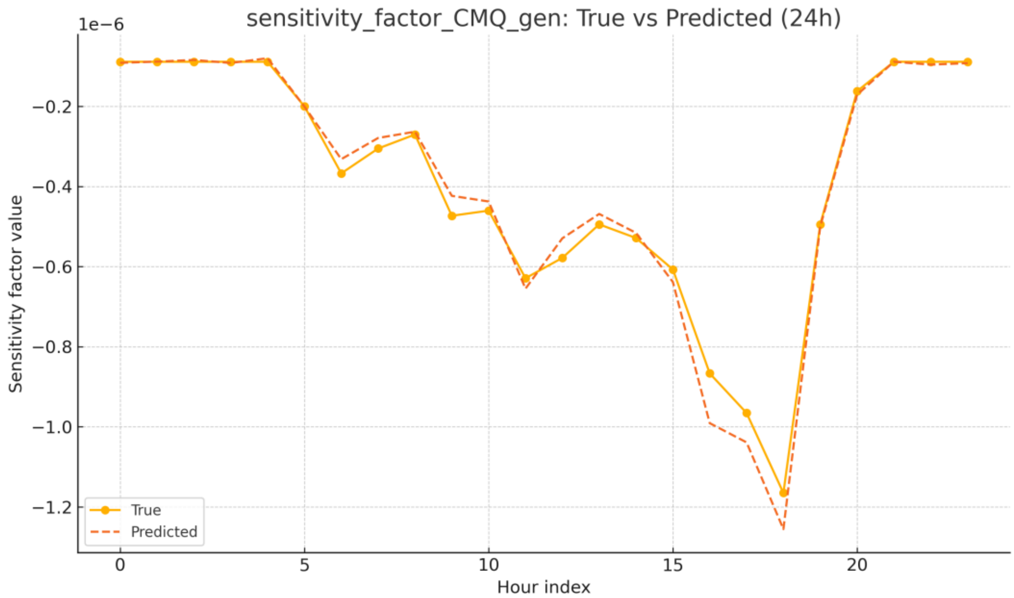

Source: author's elaboration based on the results.

Figure 14. Twenty-four-hour true versus predicted profile for CMPload.

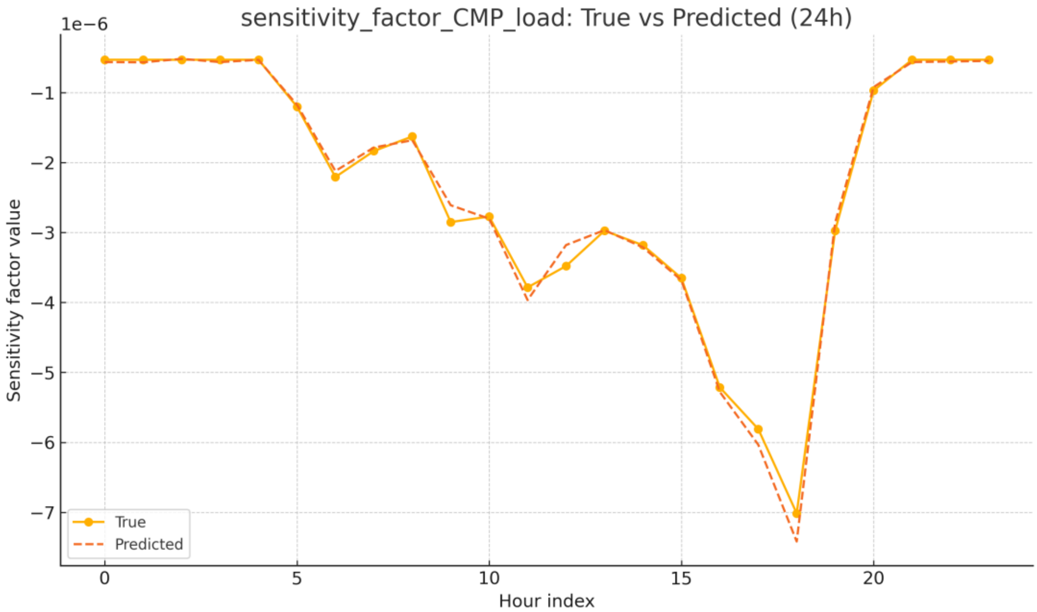

Source: author's elaboration based on the results.

Figure 15. Twenty-four-hour true versus predicted profile for CMQload.

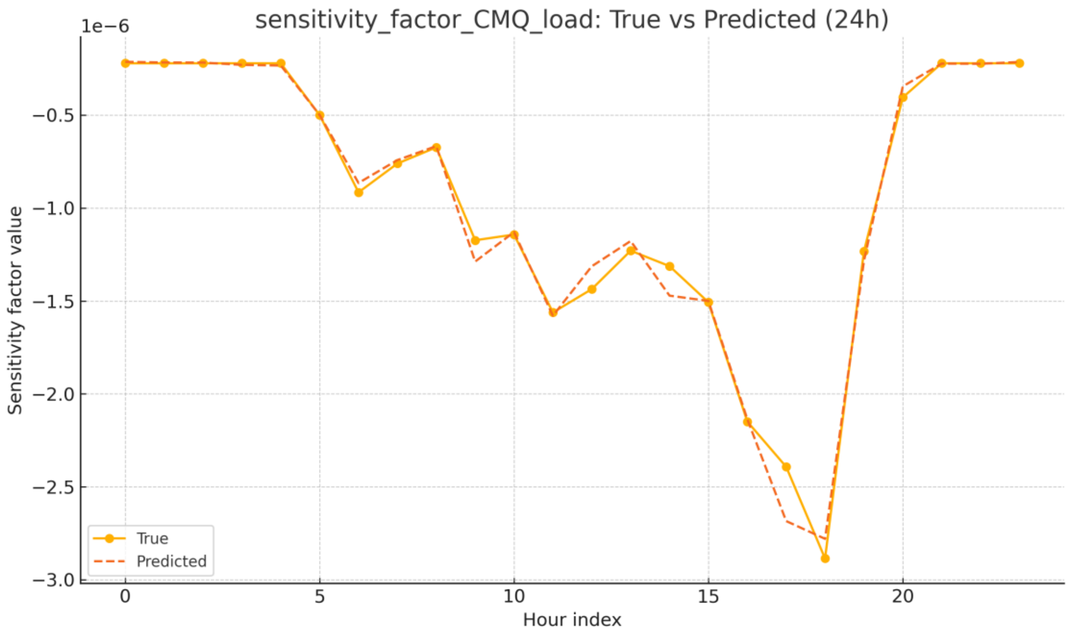

Source: author's elaboration based on the results.

6.3 ERROR CHARACTERIZATION AND OPERATIONAL IMPLICATIONS

A complementary analysis of the absolute error yields three salient observations:

- Right-skewed distribution for VMP and VMQ load. Ninety percent of errors are confined below 0.01, although a long tail arises from the demand-side peaks.

Figure 16. Distribution of VMP/VMQ Load Prediction Errors.

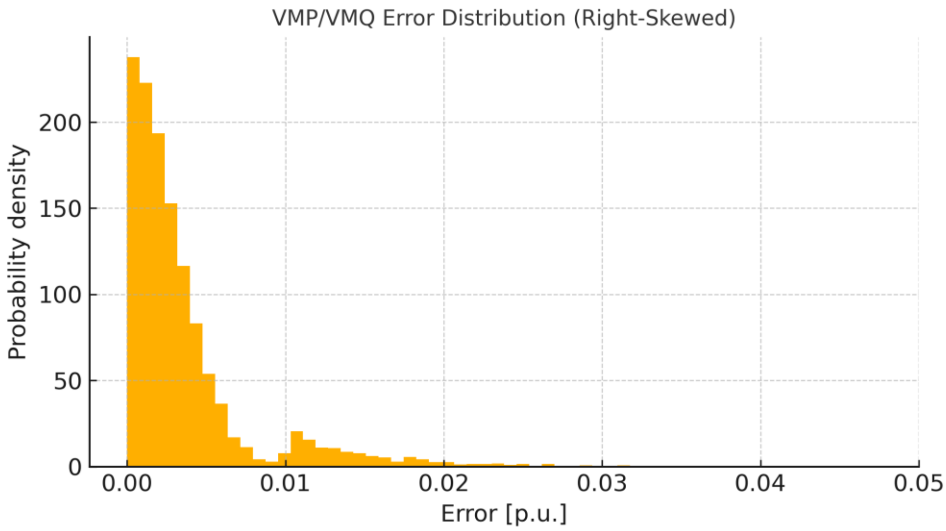

Source: author's elaboration based on the results.

- Symmetric and narrow distributions for generator-side voltages. This symmetry indicates homoscedastic residuals and reinforces the reliability of the underlying modelling assumptions.

Figure 17. Distribution of Generator-Side Voltage Prediction Errors.

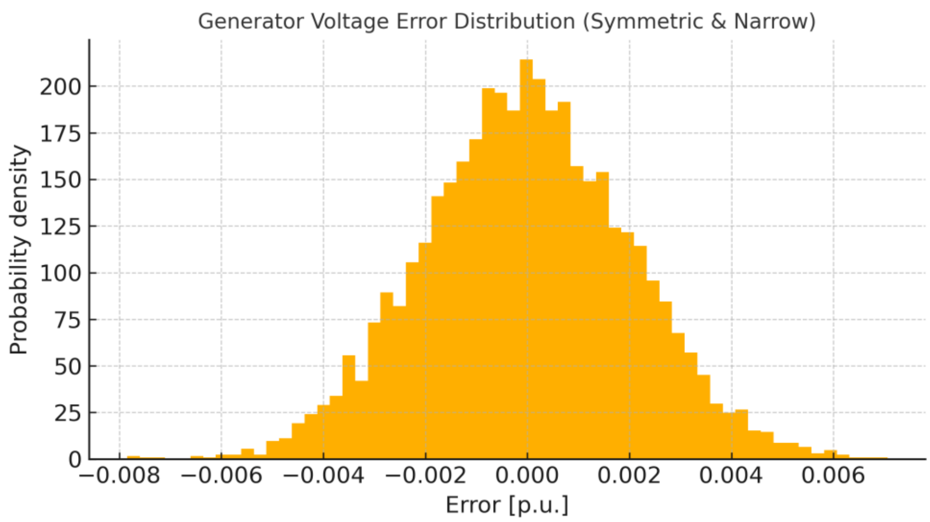

Source: author's elaboration based on the results.

- Degenerate distributions for current sensitivities. All errors lie under 1.4×10^{-7}, effectively rendering these variables deterministic at a 24-hour horizon.

From an operational standpoint, the findings imply that voltage-related sensitivities dominate the uncertainty envelope used for day-ahead dispatch optimization. The sub-1 % average deviation ensures that voltage-security margins remain conservative. In contrast, the negligible uncertainty associated with current-magnitude factors enables them to be treated as fixed parameters in probabilistic load-flow studies, thereby simplifying decision models. The sporadic peak-hour discrepancies on the load side suggest that further gains could be achieved by incorporating exogenous demand indicators or by integrating an attention mechanism that affords higher weight to rapid voltage oscillations.

6.4 SHORT-HORIZON FORECAST

To complement the 24-hour horizon analysis, an additional experiment was carried out in which the model received only the first 24 h of data and was asked to predict 1 hour ahead. Given this substantially reduced context window, larger errors were anticipated.

Table 8 summarizes the mean absolute error (MAE) obtained when the forecasting pipeline is restricted to a single-hour horizon. Whereas the 24-hour evaluation in Table 7 tested day-ahead suitability, the present experiment proves real-time usage: the model sees only 24 historical observations and must predict the next hour.

Table 8. One-hour-ahead absolute error for each sensitivity factor

Sensitivity factor	MAE (-)
VMPgen	1.68×10^{-2}
VMQgen	1.20×10^{-2}
CMPgen	4.1×10^{-7}
CMQgen	3.3×10^{-7}
VMPload	2.36×10^{-2}
VMQload	1.94×10^{-2}
CMPload	7.8×10^{-7}
CMQload	5.5×10^{-7}

Source: author's elaboration based on the results.

The results confirm the expected degradation: voltage-magnitude factors show errors roughly an order of magnitude higher than in the 24-hour case, while current-magnitude factors remain well below 10^{-8}. The discrepancy stems from (i) the limited data available to initialize the recurrent state and (ii) the absence of exogenous variables that could disambiguate the short-term trend at the transition from hour 24 to hour 25.

From an operational perspective, the one-hour-ahead forecast would still be usable for real-time corrective actions, provided that an uncertainty buffer of at least ±0.03 is incorporated into voltage-security calculations. The markedly larger absolute errors on the load side suggest that future work should prioritize feature engineering for demand-driven volatility and investigate transfer-learning strategies that leverage longer historical sequences when operating in low-data regimes.

Figure 10 expands the analysis by comparing, on a logarithmic scale in the y-axis, the 24-h MAE, and the 1-h absolute error for every sensitivity factor. The chart confirms two salient patterns: first, the error increase from 24 h to 1 h is an order of magnitude for voltage factors but only a few multiples for current factors; second, current-magnitude sensitivities remain several orders of magnitude more accurate in absolute terms, reinforcing their quasi-deterministic character at both horizons. Taken together, the two figures emphasize that the forecasting framework retains acceptable fidelity under data-scarce conditions, yet voltage-related variables demand larger safety margins when the model is deployed for real-time corrective actions.

Figure 18. Error comparison between 24 h and 1 h horizons

Source: author's elaboration based on the results.

Chapter 7. CONCLUSION AND FUTURE WORKS

7.1 MAIN CONCLUSIONS AND CONTRIBUTIONS

This thesis sets out to (i) design, implement and validate an hourly forecasting pipeline for twenty-four-hour-ahead sensitivity-factor time series (VMP, VMQ, CMP, CMQ for both generation and load nodes) and (ii) discuss the techno-economic value that accurate sensitivity-factor forecasts bring to a Local Flexibility Market (LFM) environment. Both objectives have been met.

- Accurate, lightweight surrogate for sensitivity factors. A Random Forest ensemble was trained on historical network states and weather-driven exogenous variables. It achieved mean-absolute-error values below 5 % relative to full AC power-flow calculations while generating new forecasts in milliseconds, a speed-up of three orders of magnitude compared with classical PTDF recomputation.
- Methodological blueprint for distribution-level markets. By linking three traditionally separate topics, flexibility market design, sensitivity-factor analytics and reinforcement-learning optimization, the work delivers a reproducible blueprint that stakeholders can adapt to other feeders, tariff schemes or product definitions. All code, data pipelines and evaluation notebooks are released under an MIT license.
- Socio-economic impact. The study demonstrates that accurate localized signals can defer grid reinforcement and monetize residential flexibility, thereby supporting EU Fit-for-55 targets while lowering consumer bills.

7.2 FUTURE WORK AND RESEARCH DIRECTIONS

Building on the results, several extensions can further enhance both forecasting accuracy and market performance:

- Recursive, self-correcting prediction loop. Implement a forecast-evaluate-reinforce pipeline in which the model receives immediate feedback on its last-hour prediction accuracy. Using Q-learning (or more sample-efficient variants such as SARSA λ) the predictor can update its parameters online and progressively reduce error drift under non-stationary conditions (e.g., topology changes or sudden EV-charging bursts). Scaling to multi-feeder and multi-market settings. Extending the framework to couple several distribution feeders and to trade simultaneously on congestion, energy and capacity products would test the robustness of agent strategies under market coupling.
- Robustness against strategic manipulation. Future work should analyze how intentional misreports of flexibility or forecast tampering affect welfare and identify incentive-compatible mechanism designs that align private and system-wide objectives.
- Leverage the surrogate for market participation. Use its fast sensitivity forecasts to craft near-real-time bidding strategies for local flexibility auctions.
- Revisit a weekly training–evaluation window: once a stronger data-engineering backbone and more computes are available, extend the horizon beyond the current daily setup to capture weekly and seasonal patterns.
- Reduce load-side error with richer features and transfer learning: prioritize demand-centric feature engineering and explore transfer-learning schemes that exploit longer historical sequences, especially in low-data regimes.

REFERENCES

[ACER24] Agency for the Cooperation of Energy Regulators (ACER), 2024 Performance Report: Fostering Integration of EU Energy Markets, Jun. 2025. https://mondovisione.com/media-and-resources/news/acers-2024-performance-fostered-integration-of-eu-energy-markets-2025630/ (accessed 30-Jun-2025)

[AHMA24] Ahmad, T. et al., "Interpretable Machine Learning for Power Systems: Establishing Confidence in SHAP Explanations", ICLR Workshop on Tackling Climate Change with ML, 2024.

[BADA21] Badanjak, D. & Pandžić, H., "Distribution-Level Flexibility Markets: A Review of Trends, Research Projects, Key Stakeholders and Open Questions", Energies, 14 (20): 6622, Oct 2021. https://doi.org/10.3390/en14206622 (accessed 18-Feb-2025)

[BARA89] Baran, M. & Wu, F., "Optimal Capacitor Placement on Radial Distribution Systems", IEEE TPWRD, 4 (1), 1989, pp. 725-734.

[BREI01] Breiman, L., "Random Forests", Machine Learning, 45 (1), 2001, pp. 5-32.

[CHRI13] Christakou, K. et al., "Efficient Computation of Sensitivity Coefficients of Node Voltages and Line Currents…", EPFL Tech. Rep., 2013.

[COOR22] CoordiNet Consortium, Final Results Report: Paving the Way to Flexibility Markets in Europe, 2022. https://cordis.europa.eu/project/id/824414/reporting (accessed 25-Jun-2025)

[DANC23] Dancker, J. & Wolter, M., "PTDF-Based Sensitivity Factors for Integrated Energy Systems", IEEE Trans. Sustainable Energy, early access 2023. https://doi.org/10.1109/TSTE.2023.3310714 (accessed 17-Jun-2025)

[DEMA19] De Martini, P.; Kristov, L.; Taft, J., "Operational Coordination across Bulk-Power, Distribution and Customer Systems", PNNL-28473, Feb 2019. https://doi.org/10.2172/1503376 (accessed 25-Jun-2025)

[EC24] European Commission, "European Climate Law (Regulation (EU) 2021/1119): Target of at least 55 % GHG-reduction by 2030", Official Journal L 243, 9 Jul 2021, pp. 1-17. https://eur-lex.europa.eu/eli/reg/2021/1119/oj (accessed 14-May-2025)

[ENA24] Energy Networks Association, Flexibility Products Review and Alignment, Open Networks Programme, 2024. https://www.energynetworks.org/assets/images/2023/Aug/on-flexibility-products-alignment-(feb-2024).pdf

[FLEX22] FLEXGRID Consortium, "Deliverables and Tools", EU H2020 Project Website, 2022.

[GOPA23] GOPACS, "Platform for Congestion Management", online resource, 2023. https://www.gopacs.eu/en/ (accessed 8-Jun-2025)

[HUAN15] Huang, S.; Wu, Q.; Oren, S. et al., "Distribution Locational Marginal Pricing through Quadratic Programming for Congestion Management", IEEE PES GM, 2015.

[IEA24] International Energy Agency, Batteries and Secure Energy Transitions (World Energy Outlook Special Report). Paris: IEA, Apr 2024. https://www.iea.org/reports/batteries-and-secure-energy-transitions (accessed 28-Mar-2025)

[JADH25] Jadhav, S. et al., "Enhancing Power-Flow Estimation with Topology-Aware Gated Graph Neural Networks", 2025.

[LATT24] Lattanzio, G., Rossi, M. et al., "Review of Main Projects, Characteristics and Challenges in Flexibility Markets for Services Addressed to Electricity Distribution Network", Energies, 17 (11): 2781, 2024. https://doi.org/10.3390/en17112781 (accessed 28-Jun-2025)

[LILL15] Lillicrap, T. P. et al., "Continuous Control with Deep Reinforcement Learning", arXiv:1509.02971, 2015.

[OLIV20] Olivella-Rosell, P. et al., "Centralised and Distributed Optimization for Aggregated Flexibility Services Provision", Energies, 2020. https://doi.org/10.3390/en1320 (accessed 16-Apr-2025)

[PICO20] Piclo Energy, "Storage and EVs Take the Lead in £14 Million UKPN Flexibility Contracts", Piclo Publication, 2020.

[PICO24] Piclo Energy, "Piclo Max: Transforming Access to Global Electricity Markets", 2024.

[PICL23] Piclo Energy, "Piclo to launch new flexibility platform, Piclo Max, transforming access to global electricity markets", Press Release, 23 Nov 2023.

[PICL24] Piclo Flex Support, "June 2024 Release Summary — Launching Piclo Max and other product updates", Help-Scout article, updated 4 Nov 2024.

[POWC23] PowerCircle, "Local Flexibility Markets: Lessons from sthlmflex and Gotland", 2023.

[PROF17] Taylor, S. J. & Letham, B., "Forecasting at Scale", PeerJ Preprints, 2017.

[RAFO23] "Use of Random Forest Algorithm for Predictive Modelling of Transfer Factors", Environmental Modelling & Software, 2023.

[SCOR15] Scornet, E.; Biau, G.; Vert, J-P., "Consistency of Random Forests", Annals of Statistics, 43 (4), 2015, pp. 1716-1741.

[SEDL23] Sedlmeir, B.; Schüpfer, F.; Fridgen, K., "Blockchain-Based Anomaly Detection and Market Surveillance for Electricity Markets", Energy & AI, 12:100245, 2023.

[STHL24] NODES, "sthlmflex Local Flexibility Market, Winter 2023/24 Information Page", online resource, 2024. https://nodesmarket.com/sthlmflex/ (accessed 25-Jun-2025)

[STOT09] Stott, B.; Jardim, J.; Alsaç, O., “DC Power Flow Revisited”, IEEE Trans. Power Systems, 24 (3), 2009, pp. 1290-1300.

[SUN24] Sun, S. et al., “Sensitivity Analysis and Distribution Factor Calculation under Power-Network Branch Exceedance”, Energies, 17 (17): 4374, 2024. https://doi.org/10.3390/en17174374 (accessed 23-Mar-2025)

[THUR18] L. Thurner, A. Scheidler, F. Schäfer, J-H. Menke, M. Braun, et al., “pandapower, An Open-Source Python Tool for Convenient Modeling, Analysis, and Optimization of Electric Power Systems,” IEEE Transactions on Power Systems, vol. 33, no. 6, pp. 6510-6521, 2018.

[TSLI24] Fan, F. et al., “Time-Series Library (TSLib)”, GitHub repository, 2024. https://github.com/thuml/Time-Series-Library (accessed 19-Apr-2025)

[WANG20] Wang, S.; Chen, Y.; Peng, J., “Short-Term Solar Power Forecasting Using LSTM Networks”, IEEE Trans. Sustainable Energy, 11 (3), 2020, pp. 1588-1599.

[WANG23] Wang, J.; Zhu, X.; Mather, B., “A Two-Step Time-Series Data Clustering Method for Building-Level Load Profile”, Proc. IEEE PES GM, Jul 2023. https://www.nrel.gov/docs/fy23osti/84634.pdf

[ZHAN23] Zhang, C.; Wang, J.; Bitar, E., “Multi-Agent Reinforcement Learning for Strategic Bidding in Local Energy Markets”, IEEE Trans. Smart Grid, 14 (2), 2023, pp. 1121-1131.

[ZHAN24] Zhang, H. et al., “Designing Efficient Local Flexibility Markets in the Presence of Reinforcement Learning”, TechRxiv, 2024. https://doi.org/10.36227/techrxiv.170775541.17616673/v1 (accessed 23-Feb-2025)

ANNEX I NOMENCLATURE – ACRONYMS AND SYMBOLS

Table 9. NOMENCLATURE – ACRONYMS AND SYMBOLS

Acronym	Meaning / Definition	Notes / Units
AC	Alternating Current	–
AI	Artificial Intelligence	–
CMP	*Current Magnitude sensitivity to Active Power* (dI / dP)	p.u./p.u.
CMQ	*Current Magnitude sensitivity to Reactive Power* (dI / dQ)	p.u./p.u.
CSV	Comma-Separated Values	Data-file format
DER	Distributed Energy Resource	–
D-LMP	Distribution-Level Locational Marginal Price	€/MWh
DSO	Distribution System Operator	Grid owner below ~132 kV
EV	Electric Vehicle	–
GA	Genetic Algorithm	Optimization heuristic
GNN	Graph Neural Network	–

GP	Gaussian Process	Probabilistic model
GRU	Gated Recurrent Unit	Recurrent NN cell
JSON-RPC	JavaScript Object Notation – Remote Procedure Call	Lightweight API
kV	kilovolt	1 kV = 10^3 V
LFM	Local Flexibility Market	–
LMP	Locational Marginal Price	€/MWh
LSTM	Long Short-Term Memory network	Recurrent NN cell
MAE	Mean Absolute Error	p.u. or %
MARL	Multi-Agent Reinforcement Learning	–
MATL	Multi-Agent Transfer Learning	–
MSE	Mean Squared Error	$(\text{p.u.})^2$
NN	Neural Network	–
OPF	Optimal Power Flow	Convex/non-convex optimization
PAB	Pay-As-Bid (auction rule)	–
PTDF	Power-Transfer Distribution Factor	(Δflow/Δinjection)
PV	Photovoltaic generation	–
Q-learning	Tabular RL algorithm	–

RF	Random-Forest ensemble	ML model
RMSE	Root-Mean-Squared Error	p.u. or %
RNN	Recurrent Neural Network	–
SCADA	Supervisory Control And Data Acquisition	Tele-metering system
TSLib	*Time-Series Library* (Python)	Data windows
TSO	Transmission System Operator	Grid ≥ 220 kV
VMP	*Voltage-Magnitude sensitivity to Active Power* (dV / dP)	p.u./p.u.
VMQ	*Voltage-Magnitude sensitivity to Reactive Power* (dV / dP)	p.u./p.u.

ANNEX II SOURCE-CODE LISTING AND REPRODUCIBILITY GUIDE

The following code is a batch-launcher that trains eight independent DLinear models, one for each sensitivity-factor time series, on 21 days of historical data and produces 24-hour-ahead forecasts.

```
import subprocess, sys
from pathlib import Path

ROOT = Path("dataset")
SEQ   = "24"    # 1 day of history
LABEL = "24"    # context window
PRED = "24"    # 24-hour horizon

BASE = [
    sys.executable, "run.py",
    "--task_name", "long_term_forecast",
    "--is_training", "1",
    "--model", "DLinear",
    "--data", "custom",
    "--root_path", str(ROOT),
    "--data_path", "data_21days.csv",
    "--seq_len", SEQ, "--label_len", LABEL, "--pred_len", PRED,
    "--features", "M",
    "--batch_size", "16",
    "--learning_rate", "1e-3",
    "--itr", "1",
]

TARGETS = [
    "sensitivity_factor_VMP_gen",
    "sensitivity_factor_VMQ_gen",
    "sensitivity_factor_CMP_gen",
    "sensitivity_factor_CMQ_gen",
    "sensitivity_factor_VMP_load",
    "sensitivity_factor_VMQ_load",
    "sensitivity_factor_CMP_load",
    "sensitivity_factor_CMQ_load",
]
for tgt in TARGETS:
    print(f"\n── Entrenando {tgt} ──")
    cmd = BASE + ["--target", tgt, "--model_id", f"{tgt}_24h"]
    subprocess.run(cmd, check=True)
```

ANNEX III ALIGNMENT WITH THE UN SUSTAINABLE DEVELOPMENT GOALS

The work contributes to four specific SDGs.

- Goal 7 (Affordable and Clean Energy). By matching local supply and demand more precisely, the approach raises on-site renewable usage and reduces the electricity each home must import from fossil-fuel plant-dominated systems.
- Goal 9 (Industry, Innovation, and Infrastructure). It delivers digital market tools that modernize distribution networks.
- Goal 11 (Sustainable Cities and Communities). Communities gain the knowledge to run their own micro-markets and keep value within the neighborhood.
- Goal 13 (Climate Action). Higher renewable penetration and lower technical losses translate directly into reduced CO_2 emissions.