Expert Systems With Applications 302 (2026) 130554

Contents lists available at ScienceDirect

Expert Systems With Applications

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

Decision-Focused Learning Enhanced by Automated Feature Engineering for Energy Storage Optimisation

Nasser Alkhulaifi [a,*], Ismail Gokay Dogan [a], Timothy R. Cargan [a], Alexander L. Bowler [b], Direnc Pekaslan [c], Nicholas J. Watson [b], Isaac Triguero [d,e,a,*]

[a] Computational Optimisation and Learning (COL) Lab, School of Computer Science, University of Nottingham, NG8 1BB, UK
[b] School of Food Science and Nutrition, Faculty of Environment, University of Leeds, LS2 9JT, UK
[c] Lab for Uncertainty in Data and Decision Making (LUCID), School of Computer Science, University of Nottingham, NG8 1BB, UK
[d] Department of Computer Science and Artificial Intelligence, University of Granada, Spain
[e] DaSCI Andalusian Institute in Data Science and Computational Intelligence, Granada, Spain

ARTICLE INFO

Keywords:
Decision-focused learning
Predict-and-optimise
Predict-then-optimise
Automated feature engineering
Energy storage optimisation

ABSTRACT

Decision-making under uncertainty in energy management is complicated by unknown parameters hindering optimal strategies, particularly in Battery Energy Storage System (BESS) operations. Predict-Then-Optimise (PTO) approaches treat forecasting and optimisation as separate processes, allowing prediction errors to cascade into suboptimal decisions as models minimise forecasting errors rather than optimising downstream tasks. The emerging Decision-Focused Learning (DFL) methods overcome this limitation by integrating prediction and optimisation; however, they are relatively new and have been tested primarily on synthetic datasets with limited evidence of their practical viability. Real-world BESS applications present additional challenges, including greater variability and data scarcity due to collection constraints. Because of these challenges, this work leverages Automated Feature Engineering (AFE) to improve the nascent approach of DFL. This AFE–DFL integration automatically extracts decision-relevant features from limited energy data without requiring domain expertise, while ensuring features directly enhance BESS operational decisions rather than merely improving prediction accuracy metrics. We propose an AFE–DFL framework suitable for small datasets that forecasts electricity prices and demand while optimising BESS operations to minimise costs. We validate the framework's effectiveness on a novel real-world UK property dataset. The evaluation compares DFL methods against PTO, with and without AFE. Results show that DFL yields lower operating costs than PTO, and adding AFE further improves DFL performance by 22.9–56.5 % compared to models without AFE. These findings provide empirical evidence for DFL's practical viability, demonstrating that AFE-DFL integration reduces reliance on domain expertise while achieving superior economic outcomes for BESS optimisation.

1. Introduction and background

1.1. Introduction

Decision-making under uncertainty is common in real-world applications where unknown parameters significantly complicate the process (Ibrahim et al., 2020; Reza et al., 2023). For example, in residential energy systems with a Battery Energy Storage System (BESS), operators must make critical decisions about when to charge or discharge batteries to exploit time-varying tariffs (i.e., minimising electricity costs), and how much energy to store or release, while respecting physical and operational constraints of BESS (Yang et al., 2022; Yu et al., 2023). These decisions are made under uncertainty in both future electricity prices and household demand. In the literature, addressing these challenges typically involves a two-stage process: Machine Learning (ML) models forecast unknown variables, and then use these predictions as input parameters for Constrained Optimisation (CO) to determine optimal decisions within set boundaries (Bergmeir et al., 2025). This traditional sequential approach, also known as Predict-Then-Optimise (PTO), handles prediction and optimisation in isolation (Vanderschueren et al., 2022). The limitations of this approach manifest in two distinct ways: a) cascading errors arise from the sequential structure of PTO, where

* Corresponding authors.
E-mail addresses: nasser.alkhulaifi@nottingham.ac.uk (N. Alkhulaifi), ismail.dogan1@nottingham.ac.uk (I. Gokay Dogan), timothy.cargan@nottingham.ac.uk (T.R. Cargan), a.l.bowler@leeds.ac.uk (A.L. Bowler), direnc.pekaslan1@nottingham.ac.uk (D. Pekaslan), n.j.watson@leeds.ac.uk (N.J. Watson), triguero@decsai.ugr.es (I. Triguero).

https://doi.org/10.1016/j.eswa.2025.130554
Received 5 September 2025; Received in revised form 26 October 2025; Accepted 20 November 2025
Available online 26 November 2025

inaccuracies in the prediction stage may propagate and amplify through the optimisation stage, leading to suboptimal decisions (Mandi et al., 2022; Wilder et al., 2019); and b) PTO suffers from objective misalignment in the utility of information extraction, as it focuses on deriving features and patterns from data solely to minimise prediction errors, without prioritising the information most relevant or useful for the downstream decision task (Boettiger, 2022; Donti et al., 2017).

To overcome this challenge, an emerging approach known as Decision-Focused Learning (DFL) integrates predictive modelling and optimisation directly into the learning process (Wilder et al., 2019). In this approach, forecasts are selected or assessed based on their impact on the actual downstream cost of the optimisation problem, rather than on standard error metrics (i.e., error-based loss function) such as Mean Squared Error (MSE). In order to achieve this, a task-aware loss function, such as regret, can be used (Mandi et al., 2020). By incorporating regret-based loss functions, which measure the difference between realised outcomes under uncertainty and optimal outcomes under perfect foresight, DFL methods have the potential to enhance decision quality by aligning training with the end-use task and prioritising it over minimising forecasting errors (Anis Lahoud et al., 2025). Yet, DFL approaches remain largely untested in practical settings, with most evaluations conducted on artificial datasets and simplified scenarios (Geng et al., 2024; Kotary et al., 2021; Mandi et al., 2024). To establish their real-world effectiveness, DFL methods must be validated beyond basic synthetic benchmarks (Mandi et al., 2020; Zhou et al., 2024) using actual operational data and realistic constraints.

Recent DFL applications have used extensive datasets spanning multiple years (Bergmeir et al., 2025; Paredes et al., 2025; Sang et al., 2022; Wang et al., 2025), yet many practical implementations face data scarcity due to collection constraints, privacy limitations, or resource restrictions (Alkhulaifi et al., 2024a; Grinsztajn et al., 2022; Hollmann et al., 2022). This creates a critical research gap where DFL performance remains largely unvalidated under constrained, small-scale real-world conditions typical of operational energy systems. Feature engineering (FE) can maximise information extraction from limited datasets (Wang et al., 2022b), but FE remains a manual, expert-dependent task (Wang et al., 2022b; Wu et al., 2022). While Automated FE (AFE) shows promise for improving forecasting metrics (Alkhulaifi et al., 2025; Hollmann et al., 2024), its effectiveness for enhancing DFL in data-constrained environments remains unexplored.

These gaps in DFL validation under data-constrained real-world conditions and the unexplored potential of AFE for enhancing DFL while reducing reliance on domain expertise motivated this work. First, we propose a decision-aware, end-to-end ML framework that jointly forecasts electricity demand and prices while optimising BESS operations using regret-based objectives, specifically designed for small dataset sizes. Second, we enhance the nascent DFL approach by integrating domain-specific AFE to extract richer data representations without requiring extensive domain expertise, thereby streamlining DFL pipeline development for energy applications. Third, we validate this framework using novel real-world data from a UK-based property, providing a comprehensive comparative analysis between traditional PTO and DFL approaches to demonstrate the practical viability of DFL methods in operational BESS systems.

1.2. Related work

In gradient-based DFL literature, Smart “Predict, then Optimise” (SPO$^+$) (Elmachtoub & Grigas, 2022) and Differentiable Black-Box (DBB) (Pogančić et al., 2020) are two seminal methods. In brief, the SPO$^+$ method employs a convex surrogate loss function, derived via duality theory, that upper bounds the SPO loss, which measures the decision error (suboptimality gap) induced by predicted cost vectors in a linear, convex, or integer optimisation problem, enabling efficient gradient-based training tailored to optimise decision quality. DBB method implements an efficient backward pass for blackbox combinatorial solvers with linear objective functions by constructing a continuous interpolation function, whose gradient is computed using a single solver call on perturbed inputs, enabling the integration of combinatorial algorithms into neural network architectures. These methods have been applied to various classical DFL problems, such as the travelling salesman and shortest path problems (Pogančić et al., 2020). However, despite promising theoretical advances and growing interest in DFL research, these methods have predominantly been evaluated on synthetic benchmark problems, with a lack of real-world applications (Geng et al., 2024; Kotary et al., 2021; Mandi et al., 2024). Therefore, DFL methods need to be explored beyond small-scale synthetic problems (referred to as *toy-level* problems in Mandi et al. (2020), Zhou et al. (2024)) to demonstrate their practical viability using real-world data and constraints.

As the integration of renewables accelerates, prediction and optimisation approaches have attracted growing attention. For instance, as part of the IEEE-CIS Technical Challenge (Bergmeir et al., 2025), participants used different methods to forecast 15-minute power demand and solar production for six buildings and six solar arrays, while the optimisation task was to generate a schedule for a set of activities that minimised electricity costs across the buildings. Building on this, Abolghasemi and Bean (2022) reported strong positive Pearson correlations (0.81-0.9) between forecasting accuracy and optimisation cost across overforecast and underforecast scenarios (i.e., perturbed). Yet, they found that this correlation is asymmetric, meaning unequal effects between overforecasting and underforecasting, and that forecast-accuracy metrics may be sub-optimal for minimising complex optimisation costs.

A growing number of studies have started applying DFL frameworks to broader energy scheduling problems, including grid-level dispatch, renewable trading, and microgrid operations. These works report that aligning predictive models with operational objectives can yield measurable performance gains even in complex power systems (Zhang et al., 2025a). For example, Stratigakos et al. (2022) employed a prescriptive model tree to integrate forecasting and trading decisions for renewable energy, resulting in modest profit increases (approximately 3.8 % on average) in the French electricity market. Dias Garcia et al. (2025) demonstrated in large-scale grid scheduling cases (up to a 3000-bus system) that decision-focused training improved economic outcomes by approximately 11–13 % compared to traditional forecasts. In an IEEE 39-bus test system, Lu et al. (2022) reported that a task-tailored learning approach reduced the additional costs caused by forecast errors by about 5.5 %, while dramatically accelerating training convergence (by several orders of magnitude) relative to a standard indirect method. These studies underscore that even in multi-resource or network-level scheduling, optimising predictions for decision quality can translate to tangible cost savings or profit gains.

Similarly, the need for aligning optimisation with forecasting is particularly important in domains such as BESSs (Yang et al., 2022; Yu et al., 2023), where such problems serve as a testing ground for evaluating the interplay between forecasting accuracy and downstream decision quality. In such energy scheduling problems, where future energy demand and costs are unknown, prediction errors can significantly affect downstream decisions (Mandi & Guns, 2020), highlighting the limitations of traditional PTO methods and underscoring the potential of DFL to yield more practically valuable outcomes. In various contexts, literature on BESSs has studied PTO methods (Hannan et al., 2021; Song et al., 2024), while only recently have a few studies investigated DFL-based approaches for reserve-market participation (Paredes et al., 2025) and for developing bidding strategies for microgrids in day-ahead electricity markets (Alrasheedi et al., 2024). Sang et al. (2022) applied a decision-focused strategy to price forecasting for battery arbitrage in an energy storage system, and achieved nearly 47 % higher profits compared to a conventional error-minimising predictor, while the regret (suboptimality) of decisions dropped by over 90 %. These domain–specific applications confirm that DFL can materially enhance the economic performance of storage scheduling and bidding tasks by directly tying predictions to operational outcomes, rather than purely to forecast accuracy.

Integrated learning methods highlight a fundamental shift in how forecast performance is evaluated for decision-making. In some cases, models trained with decision-aware losses intentionally sacrifice accuracy on traditional metrics to improve the downstream objective. For instance,Zhang et al. (2024) found that DFL-trained policy (i.e., agent) yielded the lowest average system cost in a power system simulation, even though its forecast error was higher than that of standard predictors. This counterintuitive result underscores the idea that the “best” forecast in isolation may not necessarily yield the best overall decisions, particularly when downstream tasks are not accounted for in the learning process, a phenomenon known as the forecast trap (Boettiger, 2022).

Whilst there are recent attempts to use DFL in real-world settings with different dataset sizes (e.g., multi-year Bergmeir et al., 2025, one year (Paredes et al., 2025), two years (Wang et al., 2025), and six years (Sang et al., 2022)), these datasets are significantly larger than those found in many practical applications where data scarcity is common due to collection limitations, privacy concerns, or resource constraints (Alkhulaifi et al., 2024a; Grinsztajn et al., 2022; Hollmann et al., 2022). This highlights a critical gap in evaluating DFL’s performance under constrained, small-scale real-world conditions, such as the 55-day dataset used in this study. In such scenarios, FE can compensate for limited data by extracting informative features, thus maximising the utility of the available data as well as enhancing computational efficiency by eliminating noisy or irrelevant data (Wang et al., 2022b). However, manual FE for energy forecasting remains a time-consuming process that is prone to human error while relying heavily on domain expertise and iterative experimentation (Wang et al., 2022b; Wu et al., 2022). As a result, tasks such as FE and the integration of domain knowledge are largely left to human practitioners, which in turn has led to growing interest in AFE methods (Hollmann et al., 2024). Comparable evidence outside energy shows that automatically augmented representations and spatiotemporal feature integration can improve downstream tasks (Zhang et al., 2021, 2025b). While AFE traditionally focuses on improving energy forecasting error metrics (Alkhulaifi et al., 2025), it remains unclear whether such improvements yield better operational outcomes when integrated with DFL approaches.

1.3. Research gaps and contributions

This work, therefore, addresses the aforementioned gaps that challenge effective BESS optimisation, stemming from: A) DFL methods promise to align prediction with downstream objectives (Mandi et al., 2020; Wilder et al., 2019); however, they are relatively new and have been tested primarily on synthetic datasets or small-scale problems (i.e., simplified benchmarks) (Mandi, 2023; Zhou et al., 2024), highlighting the need to assess their practical viability in real-world applications such as BESS problems; and B) real-world datasets often exhibit greater variability and data scarcity due to practical constraints (Grinsztajn et al., 2022; Hollmann et al., 2022) which can compromise DFL performance and necessitate enhanced feature representations to extract richer information from limited data without the need for domain expertise. The novelty and contributions of this work are summarised as follows[1]:

- Proposing a decision-aware, end-to-end ML forecasting and optimisation framework for BESS problems and is suitable for small dataset sizes. Rather than treating forecasting and optimisation as separate tasks, the framework uses DFL to jointly forecast electricity demand and prices while optimising BESS operations using a regret-based objective.
- Improving the nascent approach of DFL by leveraging domain-specific AFE to extract richer representations without reliance on domain expertise, as presented in our previous work (Alkhulaifi et al., 2025), thereby streamlining the development of DFL pipelines for BESS problems.
- Using novel real-world data collected from a UK-based property to evaluate the proposed framework with a comprehensive comparative analysis of PTO versus DFL approaches, with and without AFE, thereby demonstrating the practical viability of DFL methods in real-world BESS applications.

The structure of the remaining sections of this paper is as follows: Section 2 outlines the methods used for the BESS problem in this work, Section 3 provides details of the experimental design, including the datasets used and evaluation criteria. Analysis and discussion of findings are presented in Section 4. To conclude, Section 5 summarises key insights.

2. Methods

This section outlines the problem under investigation in Subsection 2.1, presents the mathematical model for the BESS optimisation problem in Subsection 2.2, and describes the prediction methods in Subsection 2.3.

2.1. Problem definition

The BESS problem entails forecasting unknown parameters (electricity prices and household demand), which serve as inputs to an optimisation model that computes the optimal charge/discharge schedule for the battery over the planning horizon while minimising electricity costs and satisfying energy demand and operational constraints. As illustrated in Fig. 1, the proposed framework leverages AFE to enrich the dataset with domain-specific features (Alkhulaifi et al., 2025) while reducing domain knowledge requirements, thereby enhancing decision quality and improving the nascent approach of DFL for BESS scheduling problems.

2.1.1. Multiple unknowns, decision-making, and objectives

In this study, the BESS problem involves predicting two key unknowns (electricity demand and price) over a 24-hour planning horizon, mimicking real-world residential energy management where household BESS uses day-ahead predictions to optimise charging and discharging schedules. These uncertainties can be framed as: **(A)** What will the electricity prices be over the next 24 hours? and **(B)** What levels of electricity demand are expected? Given these forecasts, the core decision problem is to determine the optimal battery operation strategy across the 24-hour horizon, which can be formulated through the following operational questions: **(C)** At what times (i.e., hour of the day) should the battery be charged or discharged? and **(D)** By how much should it be charged or discharged? Successfully addressing these four prediction and operational questions enables the achievement of the following objectives:

- Exploit price differentials through optimal battery charging and discharging schedules to reduce overall electricity costs (i.e., cost minimisation).
- Ensure household energy requirements are consistently met through appropriate combinations of grid power and battery discharge (i.e., demand satisfaction).
- Generate a battery operation strategy that respects all battery physical constraints, including power limits, energy capacity, and state-of-charge boundaries (i.e., constraint compliance).

2.2. Mathematical model of the BESS optimisation problem

In this work, the optimisation problem of the BESS is formulated as a Mixed Integer Linear Programming (MILP) problem. MILP is widely adopted for BESS optimisation because it can efficiently handle both

[1] To support transparency and reproducibility, the historical electricity price and weather data, along with the code used for the experiments in this work, are publicly available on GitHub. See https://github.com/Nasser-Alkhulaifi/DFL

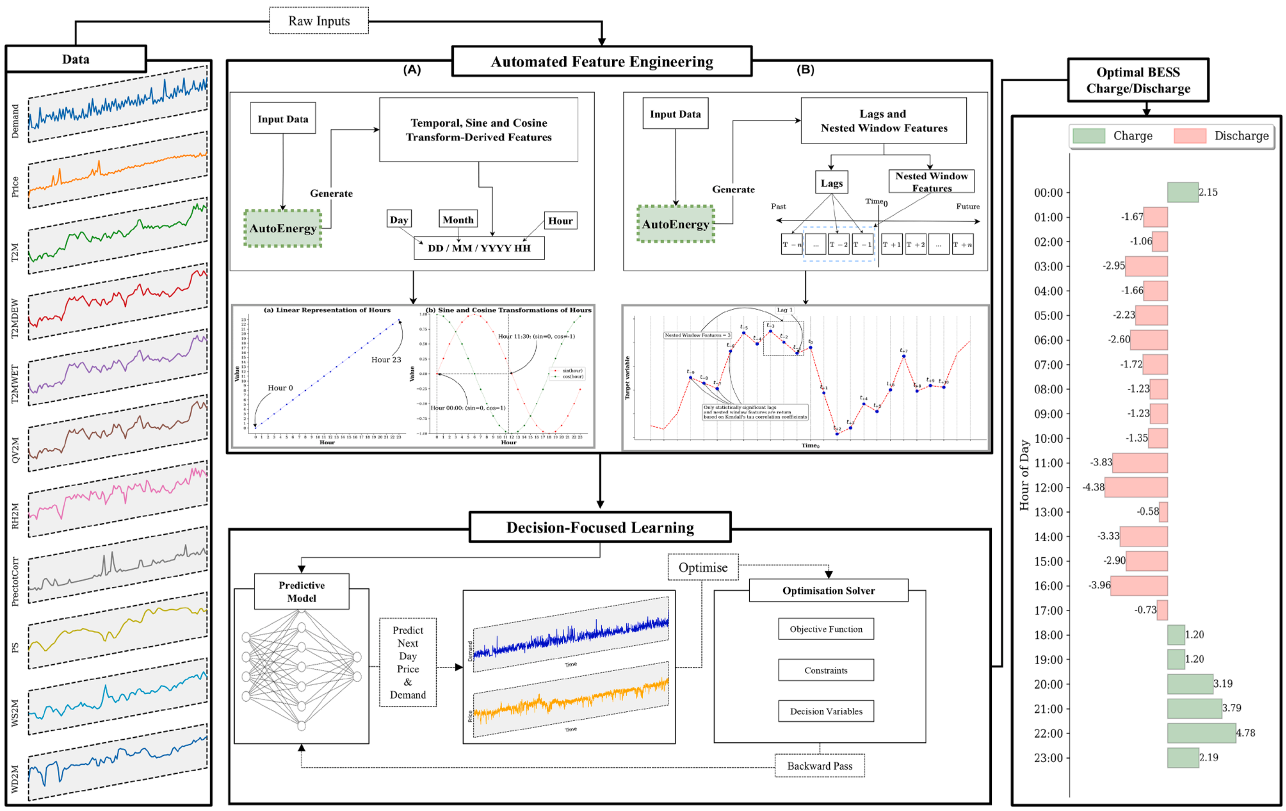

Fig. 1. The proposed framework jointly forecasts electricity prices and demand for the next day while optimising the downstream BESS task. It leverages AFE to enrich the dataset with domain-specific features using the *AutoEnergy* algorithm (Alkhulaifi et al., 2025). Further details on this algorithm, including how the features are extracted and selected, are provided in Subsection 2.4. This approach aims to minimise domain knowledge requirements and enhance decision quality (i.e., improve the nascent approach of DFL) for BESS scheduling problems under data scarcity. It is worth noting that while the problem formulation presented in Section 2, particularly the mathematical formulation in Subsection 2.2, is tailored specifically for BESS applications, the overarching framework methodology is adaptable to other energy management contexts. The core principles of jointly forecasting uncertain parameters while optimising operational decisions remain applicable, though the specific constraints, decision variables, and objective functions would require modification for different energy systems.

continuous variables (e.g., power flows and state of charge) and discrete decisions (e.g., on/off states or charge/discharge modes), enabling accurate modelling of operational constraints and system logic with computationally tractable and optimal solutions (Yang et al., 2022; Yu et al., 2023). The parameters are defined in Table 1, and the decision variables are presented in Table 2, followed by the objective function in Equation (1), subject to the listed constraints.

The BESS optimisation problem focuses on minimising the total electricity cost over the entire planning horizon. The objective is to determine, for each time interval, the amount of energy to be drawn from the battery or the grid to satisfy the predicted property demand, as well as the amount of energy to be charged into the battery. Due to physical constraints, the battery cannot be charged and discharged simultaneously. Moreover, to ensure reliable operation and maintain long-term battery health, a minimum state-of-charge level must be preserved at all times. The problem also incorporates limitations on the maximum charging and discharging rates of the battery. It is assumed that the initial state-of-charge, the unit price of electricity (per kWh), and the energy demand for each interval are known in advance. For more details on the studied system, see Subsection 3.2.

Let $T = \{1, \dots, T\}$ denote the set of all time intervals. For each interval $t \in T$, p_t and d_t represent the unit price of electricity (i.e., the cost of 1 kWh) and the energy demand, respectively. The parameters max_charge and $max\_discharge$ denote the maximum amount of energy that can be charged to or discharged from the battery in a given interval. The parameter e_0 indicates the initial battery level (in kWh). The decision variable e_t denotes the battery level at the end of interval t. The decision variables g_t, b_t, and c_t represent the amount of energy supplied from the grid to the property, from the battery to the property, and from the grid to the battery, respectively. The binary variable z_t takes the value 1 if energy is supplied from the battery to the property in interval t, and 0 otherwise.

$$\text{Min} \quad \sum_{t \in T} p_t \cdot (g_t + c_t) \tag{1}$$

Subject to

$$g_t + b_t = d_t \qquad \forall t \in T \tag{2}$$

$$e_t = e_{t-1} + c_t - b_t \qquad \forall t \in T \tag{3}$$

$$b_t \le e_{t-1} \qquad \forall t \in T \tag{4}$$

$$b_t \le M * z_t \qquad \forall t \in T \tag{5}$$

$$c_t \le M * (1 - z_t) \qquad \forall t \in T \tag{6}$$

$$c_t \le max\_charge \qquad \forall t \in T \tag{7}$$

$$b_t \le max\_discharge \qquad \forall t \in T \tag{8}$$

$$e_t \le E_{\max} \qquad \forall t \in T \tag{9}$$

$$e_t \ge SoC_{\min}\, E_{\max} \qquad \forall t \in T \tag{10}$$

$$g_t, b_t, e_t, c_t \in \mathbb{R}^+ \qquad \forall t \in T \tag{11}$$

$$z_t \in \{0, 1\} \qquad \forall t \in T \tag{12}$$

The objective function (1) minimises the total electricity cost. Constraint (2) ensures that the energy demand is satisfied for all intervals $t \in T$. Constraint (3) enforces energy flow conservation (i.e., updating the battery's energy level based on charging and discharging). Constraint (4) ensures that the energy drawn from the battery never exceeds

Table 1
Notations.

Set	Definition	
T	Set of (time) intervals	
max_charge	Maximum energy that can be added to the battery in an interval	
$max\_discharge$	Maximum energy that can be drained from the battery in an interval	
p_t	Price of 1 kWh energy at interval t	$\forall t \in T$
d_t	Energy demand at interval t	$\forall t \in T$
e_0	Initial battery level (kWh)	
$E_{\max}$	Usable battery capacity (kWh)	
$SoC_{\min}$	Minimum allowable state-of-charge (fraction of $E_{\max}$)	

Table 2
Decision Variables.

Dec. Var.	Definition
e_t	Amount of energy (kWh) in the battery at the end of interval t
g_t	Amount of energy (kWh) provided from grid to the property at interval t
b_t	Amount of energy (kWh) provided from battery to the property at interval t
c_t	Amount of energy (kWh) charged to battery at interval t
z_t	1, if energy is being sent from battery to the property at interval t 0, otherwise

the available energy at the beginning of the interval. Constraints (5) and (6) jointly ensure that the battery cannot be charged and discharged simultaneously (i.e., using a large positive constant M to enforce the binary logic of z_t). Constraints (7) and (8) impose upper bounds on charging and discharging rates. Constraint (9) limits the state of charge so that it never exceeds the battery's usable capacity $E_{\max}$. Constraint (10) enforces a floor of $SoC_{\min}\, E_{\max}$ that prevents deep cycling, thereby reducing degradation and safeguarding long-term operational longevity. Constraints (11) and (12) define the domains and the ranges of the decision variables.

Although the decision variables (see Table 2) are indexed by interval t, the MILP is solved jointly over the full horizon $T = \{1, \ldots, t., \ldots, T\}$. The solver find the optimum values for the entire sequence $\{g_t, b_t, c_t, e_t, z_t\}_{t \in T}$ while minimising the objective (1) subject to constraints (2) to (12). The intertemporal energy-balance (3), the capacity limit (9), the minimum state-of-charge bound (10), and the charge and discharge limits (7) and (8) couple decisions across time. In other words, the optimisation considers all decision variables for all intervals at once to minimise total electricity cost over the whole horizon while satisfying all constraints.

2.3. Prediction methods of the BESS forecasting problem

Three Artificial Neural Network (ANN)-based methods (PTO, SPO$^+$, DBB) are used for the BESS forecasting problem in this work, based on the following methodological rationale. First, the conventional PTO method serves as the established baseline, representing the predominant approach in BESSs where prediction and optimisation are handled separately (Vanderschueren et al., 2022). Second, two DFL methods, namely SPO$^+$ (Elmachtoub & Grigas, 2022) and DBB (Pogančić et al., 2020), as introduced in Section 1, are selected as the two seminal and most widely adopted DFL methods in gradient-based DFL literature. This selection enables a comprehensive comparison between the conventional PTO approach and the two main DFL methods, thereby addressing the research gap regarding DFL's practical viability in real-world BESS applications.

The DFL-based methods are centred on training predictive models in an end-to-end approach using decision loss derived from the associated optimisation task, and therefore, an optimisation problem needs to be solved in each training iteration. In short, the SPO$^+$ method uses a convex surrogate loss to upper-bound decision regret in downstream optimisation tasks, while the DBB method enables end-to-end training by approximating solver outputs for gradient computation. Further details on loss computation for the SPO$^+$ and DBB methods can be found in Elmachtoub and Grigas (2022) and Pogančić et al. (2020), respectively, and an explanation of the Python implementation is provided in Tang and Khalil (2024). In contrast, the PTO method trains the predictive model with MSE loss, which measures the average squared difference between the model's forecasts and the actual values, as shown in Equation (13). This approach, therefore, focuses solely on minimising forecasting errors during training, without considering the downstream optimisation task.

$$\mathrm{MSE} = \frac{1}{N} \sum_{i=1}^{N} \left(\hat{y}_i - y_i\right)^2 \tag{13}$$

where $\hat{y}_i$ is the model's prediction for the i^{th} observation, y_i denotes the true value for that observation, and N is the total number of observations.

2.4. Automated feature engineering

This study leverages an AFE method specifically designed for energy forecasting problems, which we introduced in our previous work (Alkhulaifi et al., 2025). This method, named *AutoEnergy*, aims to streamline ML pipeline development by automatically generating features from timestamps and historical energy consumption data, thereby minimising reliance on domain expertise for FE. It generates two primary categories of features: A) Temporal features: these features extract time-based patterns from timestamps such as hour of the day, day of the week, weekdays and weekends. To enhance the representation of cyclical patterns, these temporal features undergo sine and cosine transformations. Specifically, Fourier-based transformations are applied to capture periodic patterns that represent daily and weekly cycles inherent in energy data. B) Lag and nested window features: these features capture temporal dependencies and multi-scale statistical characteristics within the energy time series data. The method computes statistically significant lags and rolling statistics (e.g., mean and standard deviation) using Kendall's tau correlation coefficient. This approach ensures that only meaningful temporal relationships are incorporated into the feature set. Additional details of this method are explained in Alkhulaifi et al. (2025).

3. Experimental design

This section outlines the experimental design used in this work, including the datasets used in Subsection 3.1, BESS configuration in Subsection 3.2, experimental procedure in Subsection 3.3, and lastly the evaluation criteria in Subsection 3.4.

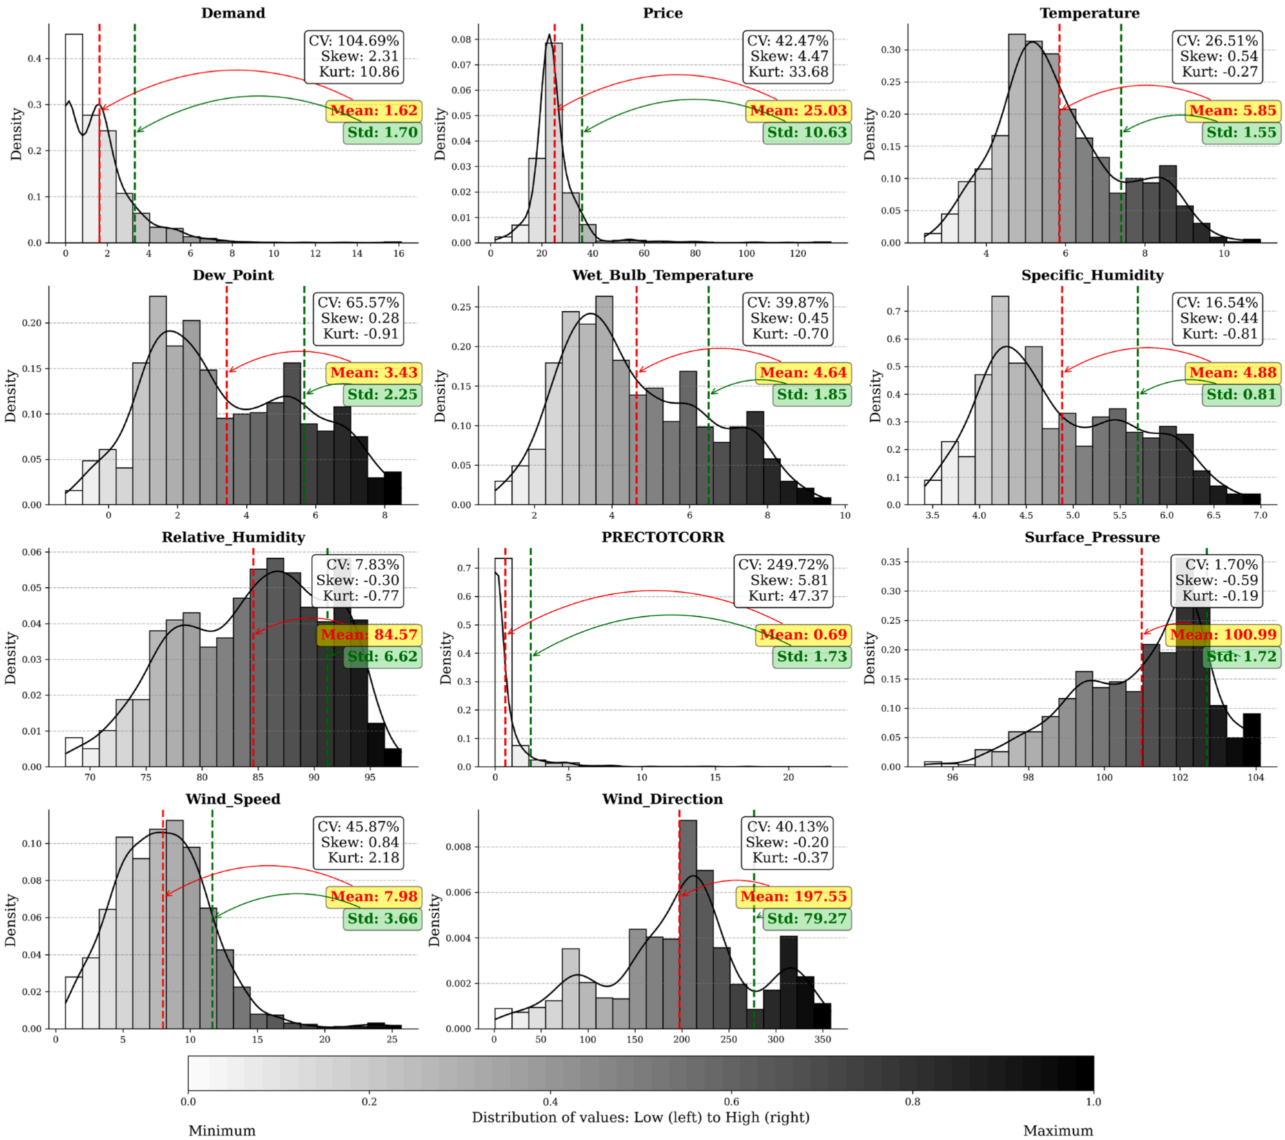

Fig. 2. The Dataset used in this study is depicted in histograms with colour-coded bars representing normalised bin positions. Red and green dashed lines indicate the mean and standard deviation, respectively. Annotated statistics include coefficient of variation (CV: relative variability), skewness (distribution asymmetry), and kurtosis (tailedness). The colour gradient in the histogram bars represents the distribution of values from low (left) to high (right).

3.1. Datasets

In this work, a dataset spanning 1 January 2025 to 24 February 2025 (55 days) is used. This scale is substantially smaller than in related works (e.g., multi-year datasets Bergmeir et al., 2025, one year (Paredes et al., 2025), two years (Wang et al., 2025), and six years (Sang et al., 2022)), thereby distinguishing the study's contribution further by rigorously evaluating DFL, enhanced by AFE, under limited real-world data. This hourly historical real-world electricity demand and BESS data were provided by the Intelligent Plant platform (Intelligent Plant Ltd, 2025). Historical electricity prices were sourced via the Octopus Energy API (Energy, 2025). Weather data was obtained from the NASA Langley Research Center's POWER Project (NASA, 2025), a repository of solar and meteorological datasets developed by NASA to support renewable energy and building energy efficiency research. Fig. 2 presents summary statistics of the collected data, while Fig. 3 displays patterns of average hourly electricity demand and price by day of the week.

3.2. BESS configuration

In this experiment, the maximum charging rate was set to 5 kWh per interval to comply with operational battery constraints, while the maximum discharging rate was limited to 4.5 kWh per interval to account for round-trip efficiency losses and system resistive losses inherent in BESSs. The minimum state of charge was maintained at 10 % to prevent battery degradation and ensure operational longevity. The capacity was set to 50 kWh, imposed by physical battery constraints. It is worth noting that the current BESS configuration operates without any solar panels. Table 3 provides a summary of the parameters and their values used in the experiment. In this system, as shown in Fig. 4, electrical power follows two primary pathways: first, power flows from the grid to charge the battery, which then supplies the property through inverters that perform DC-AC conversion to meet energy demand; second, power flows directly from the grid to the property when battery capacity is insufficient. This dual-path configuration ensures a continuous power supply through direct grid connection while enabling BESS optimisation for

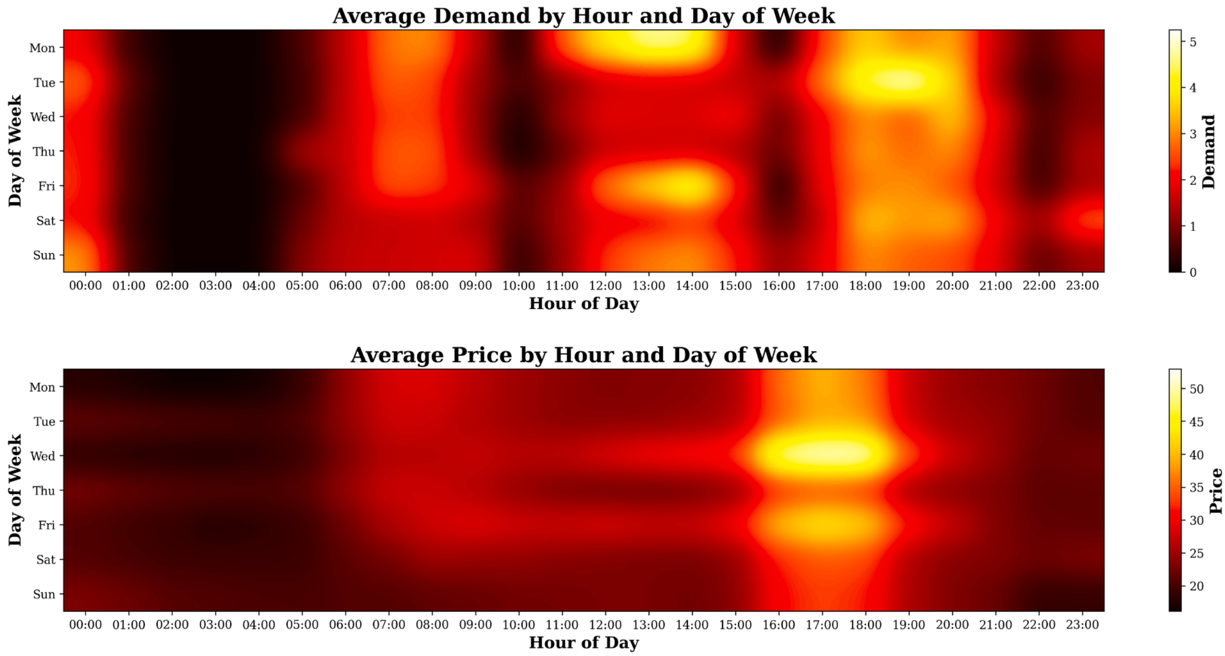

Fig. 3. Heatmaps of average hourly household electricity demand (top) and price (bottom) by day of week. Demand displays consistent peaks in the early morning (around breakfast), midday (lunch), and most prominently between 17:30 and 20:00 (dinner time), reflecting typical residential consumption patterns. Price peaks are concentrated between 16:00 and 19:30. Both demand and price are lowest during the early morning hours (01:00 to 05:00), reflecting reduced residential activity and system load during overnight periods.

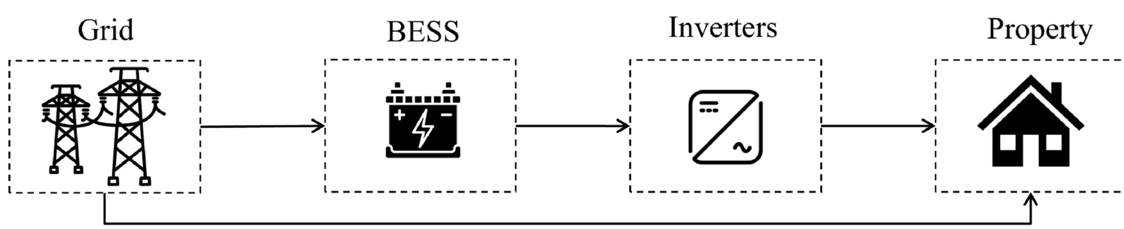

Fig. 4. Power flow diagram, where: A) power flows from the grid to charge the battery, which then supplies the property via inverters performing DC–AC conversion to meet energy demand; and B) power flows directly from the grid to the property when the battery capacity is insufficient.

Table 3
Summary of parameters and values used in the experiment. See Subsection 3.2 for BESS configuration details and Subsection 3.3 for the experimental procedure.

Parameter	Value
Planning horizon T [intervals]	24
Interval duration Δt [hours]	1
Usable capacity $E_{\max}$ [kWh]	50
Minimum SoC fraction $\mathrm{SoC}_{\min}$	0.10
Max charge per interval max_charge [kWh]	5.0
Max discharge per interval $max\_discharge$ [kWh]	4.5
Initial energy e_0 [kWh]	10
Electricity price p_t [£/kWh]	Variable by interval
Energy demand d_t [kWh]	Variable by interval

energy arbitrage (i.e., storing electricity during low-cost periods and supplying the property during high-cost periods), thereby reducing energy cost and contributing to grid stability.

3.3. Experimental procedure

In this work, the dataset is partitioned into three subsets: ∼50 % (27 days) allocated for initial training, ∼25 % (14 days) for validation to fine-tune hyperparameters, and the remaining ∼25 % (14 days) reserved for assessing model performance on previously unseen data (i.e., test dataset). ANN-based methods, as explained in Subsection 2.3, were used to predict the price and demand for the next 24 hours (i.e., the following day) based on input features such as weather conditions (e.g., outdoor temperature) and the engineered features described in Subsection 2.4, while accounting for the operational and physical constraints of the BESS explained in Subsection 2.2. This approach mimics real-world scenarios where BESS schedules charging and discharging operations daily to optimise cost savings by taking advantage of low electricity prices. It is important to note that, to better capture the interrelated dynamics between electricity price and demand, a multitask learning approach was used (Wang et al., 2022a), in which the same input features are processed through shared ANN layers to predict both outputs simultaneously, thereby leveraging shared information across tasks.

Five-fold cross-validation and grid search optimisation were employed to tune the ANN architecture and learning hyperparameters of all three methods (PTO, DBB, and SPO^{+}). During the hyperparameter tuning phase, many model configurations are trained exclusively on the training set and evaluated on the validation set until optimal hyperparameters are identified, after which the training and validation sets are combined (totalling 75 %, 41 days of the original dataset) to train the final model using the previously optimised hyperparameters, thereby maximising the utilisation of available data for final model training. Fig. 5 shows the experimental design and Table 4 shows the hyperparameter search space. Finally, the performance of the PTO, DBB, and SPO^{+} methods was evaluated on the test set using the evaluation metrics

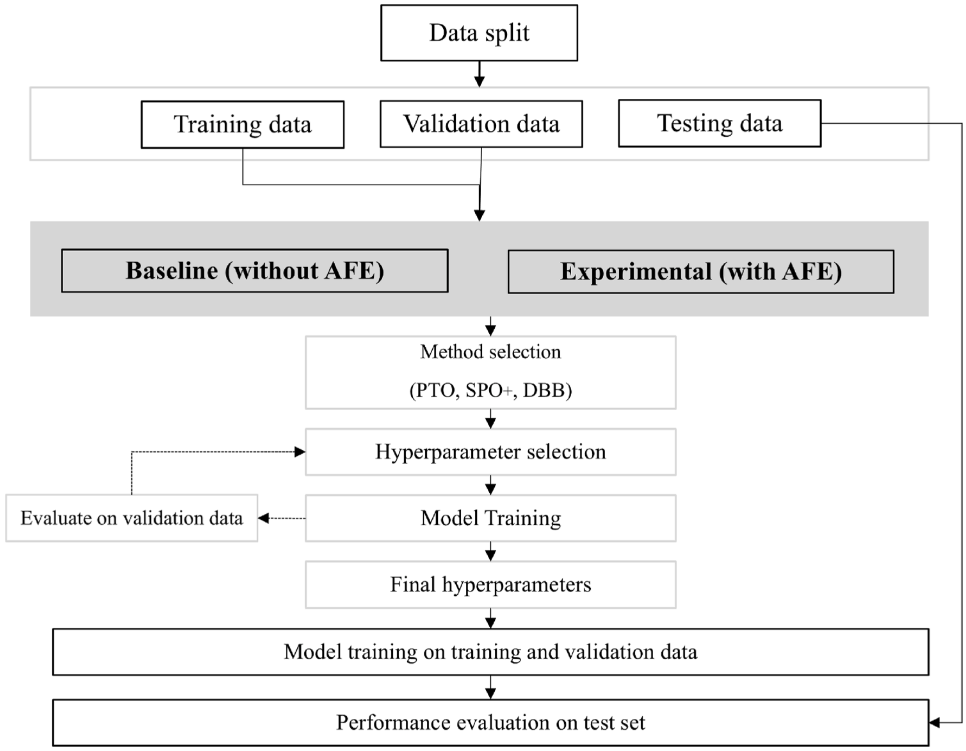

Fig. 5. Experimental design. See Subsection 3.3 for a detailed explanation of the experimental procedure.

Table 4
Hyperparameter search space and the corresponding best hyperparameters for the ANN models, determined through grid search on the validation dataset. To avoid over-parameterisation, the search space is constrained to the most critical parameters identified in related work (Paredes et al., 2025). Also, computational resource limitations further restrict the search space. However, it is worth mentioning that this may have restricted the discovery of optimal hyperparameter combinations that could further enhance model performance. The search includes architectures with one or two hidden layers; the best models use 256 and 128 neurons in the first and second hidden layers, respectively. The grid search explores 12 configurations with 5-fold cross-validation, yielding 60 candidate models per method and 180 ANNs in total (PTO, SPO$^+$, DBB); see Subsection 2.3 for details. For each training window, a 24-interval MILP (Subsection 2.2) is solved using Gurobi (Gurobi Optimization, 2024) to optimise the battery schedule. During training, the solver is called once per sample (i.e., one day) per epoch; therefore, the optimisation layer dominates the total run time. This procedure is repeated twice (see Fig. 5), once with AFE and once without AFE, as explained in Subsection 3.3. Unlike Paredes et al. (2025), which tunes core regressor hyperparameters on PTO and transfers them to DFL, we perform method-specific tuning: for PTO method, the models were trained with MSE loss on predictions; for DFL methods (SPO$^+$ and DBB), training integrates optimisation results via their respective decision-focused losses (see Mandi et al., 2020; Tang & Khalil, 2024 for more details). The total running time, including hyperparameter tuning, is approximately 39.97 hours.

Hyperparameter	Search Space	Best Hyperparameter					
		PTO		SPO$^+$		DBB	
		No AFE	AFE	No AFE	AFE	No AFE	AFE
Number of Layers	$\in \{1, 2\}$	2	2	2	2	2	1
Epochs	$\in \{10, 20, 30\}$	30	30	30	30	30	30
Learning Rate	$\in \{10^{-3}, 10^{-5}\}$	10^{-3}	10^{-3}	10^{-5}	10^{-5}	10^{-3}	10^{-3}

In this experiment, models were trained using Python 3.11, Gurobi API (Gurobi Optimization, 2024), and PyEPO API (Tang & Khalil, 2024). Computational experiments were performed on an Ubuntu system featuring an x86_64 architecture with 16 physical CPU cores (32 logical cores), 64GB of RAM.

described in Subsection 3.4. To ensure the robustness and reliability of the experimental results, models were trained and evaluated across ten independent runs with different random seeds. This practice mitigates variance introduced by stochastic training procedures (e.g., such as random weight initialisation) and enables statistically meaningful comparisons by reporting mean performance metrics together with their standard deviations. The results of all ten independent runs are reported in the supplementary material. It is also important to note that the data were not shuffled, as the energy data exhibit temporal patterns, making it essential to preserve the chronological order of the timestamps.

To evaluate the impact of AFE on decision quality, two conditions were established: A) a baseline scenario without AFE and B) an experimental scenario incorporating AFE. The baseline models serve to assess forecasting and optimisation capabilities with minimal inputs. This deliberately simplified configuration replicates model performance under conditions that simulate both aworst-case scenario (i.e., absence of

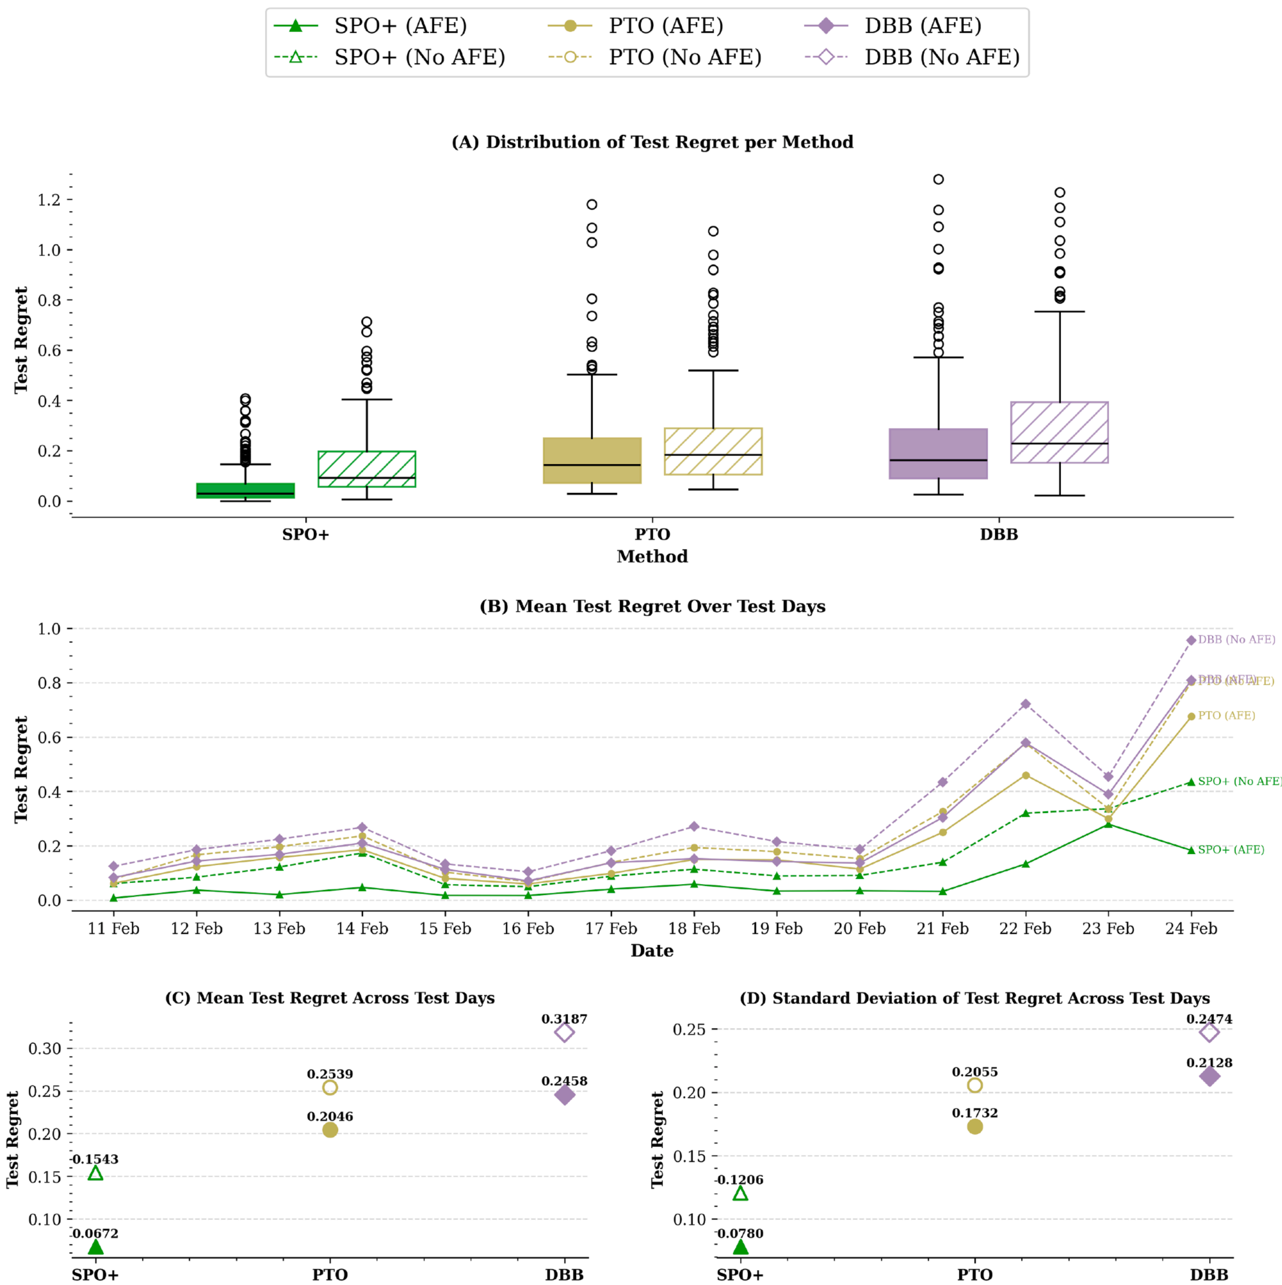

Fig. 6. Test regret across methods with and without AFE. (A) Box plots of test regret across all experimental runs and test days for PTO and DFL methods (SPO+ and DBB). The distributions reflect performance variability across multiple independent experimental runs and test days, providing a comprehensive view of method robustness and central tendency. (B) Daily mean test regret over the fourteen test days. (C) Mean test regret across test days for each method variant. (D) Standard deviation of test regret across test days..

domain-specific FE knowledge) and an initial testing phase where the model learns with limited data enhancement. Conversely, the experimental scenario applies the AFE method explained in Subsection 2.4, thereby enabling the assessment of the AFE impact on the downstream optimisation task (i.e., decision-making quality).

It is worth mentioning that the experimental setup with two well-established DFL methods and the traditional PTO approach (see Subsection 2.3) and systematic AFE evaluation (with vs. without; see Subsection 2.4), yielding a comprehensive experiment with six method conditions, is justified by: A) alignment with the stated research objective of demonstrating DFL's practical viability under data scarcity conditions rather than providing exhaustive benchmarking, where the 55-day dataset deliberately reflects real-world BESS deployment constraints faced by practitioners during initial operation (i.e., expanding to additional methods or datasets may obscure the core research questions); B) methodological clarity requirements, where incorporating more DFL methods may introduce confounding variables that complicate attribution of performance differences (i.e., results variations may reflect differences between DFL approaches rather than the contribution of AFE); and C) computational resource constraints associated with training additional DFL methods, where each training iteration necessitates solving an MILP optimisation problem for every sample in every epoch, such that additional methods would require multiplicative increases in computational demands, potentially exceeding resources available in practical deployment scenarios.

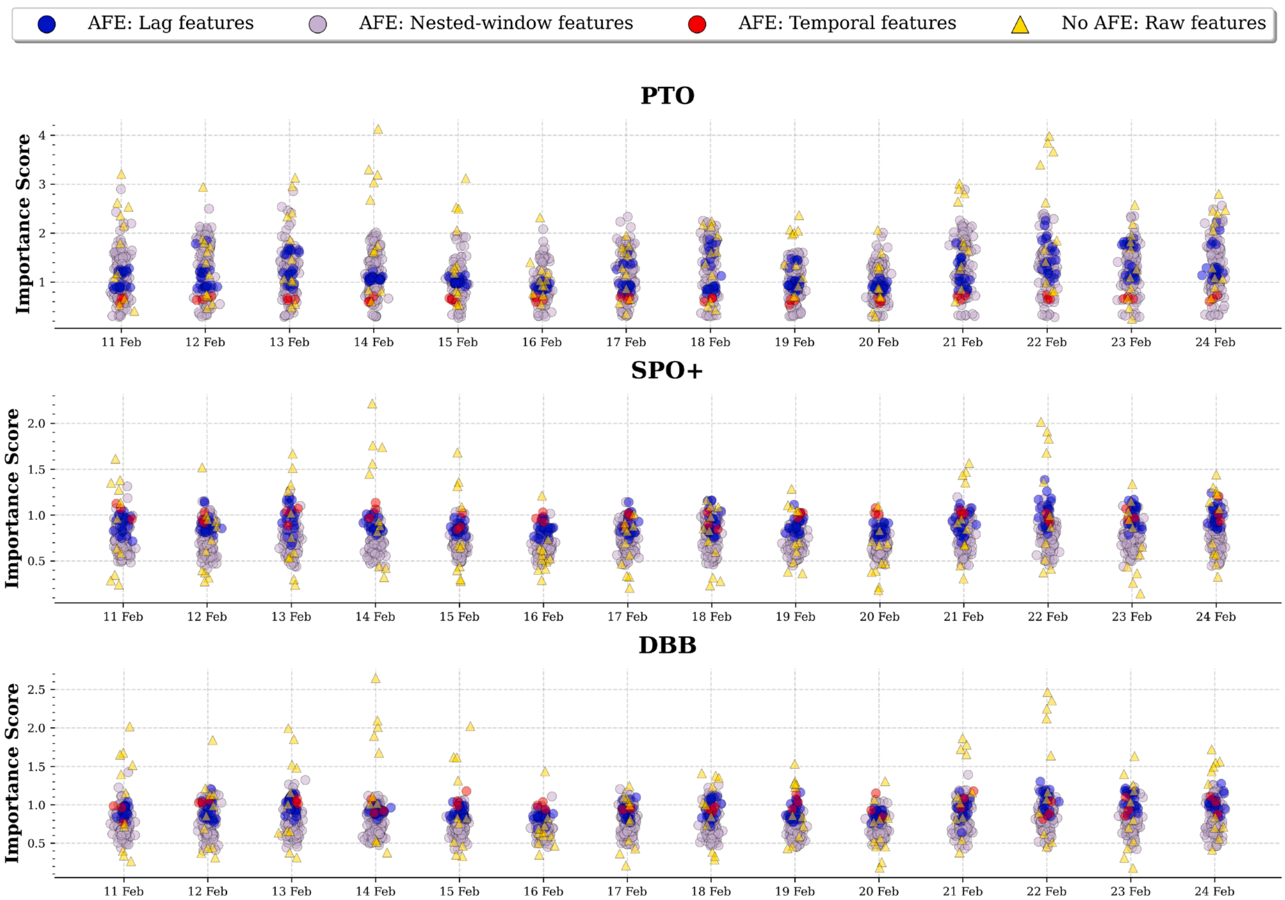

Fig. 7. Feature importance scores across all experimental runs for each method across test days. Each subplot corresponds to one method and shows the average SHAP-derived importance of individual features. Feature categories are distinguished by colour and marker shape. 'No AFE' indicates raw features, such as weather variables (e.g., outdoor temperature), while 'AFE' refers to automatically engineered features (e.g., lagged values, rolling-window statistics, temporal/calendar-based indicators). See Subsection 2.4 and our previous work (Alkhulaifi et al., 2024b, 2025) for further details on what these features represent and how they are generated and selected. Each marker represents a single feature from the corresponding category, allowing the distribution of importance values within each category to be visualised. Data points are jittered along the x-axis to reduce overlap. It is worth mentioning that the features shown are those utilised within a multitask predictive approach, in which a single input set is processed through shared layers of the ANN and then used to simultaneously generate forecasts for both electricity price and property demand.

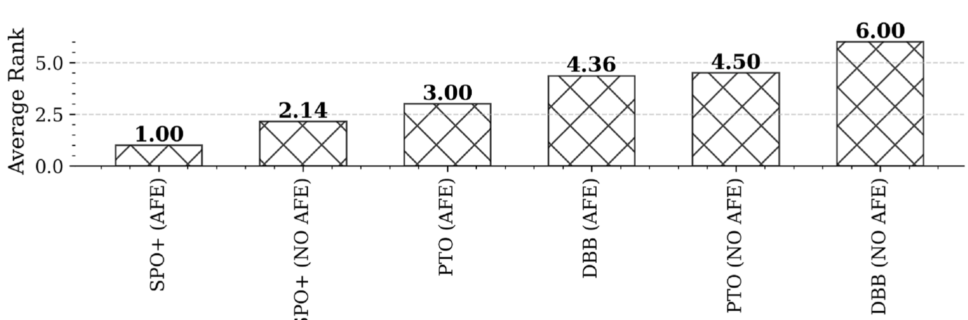

Friedman's Two-Way Analysis of Variance by Ranks Summary

Total N:	14
Test Statistic:	65.306
Degree Of Freedom:	5
Asymptotic Sig. (2-sided test):	< .001

Fig. 8. Result of comparisons on the test set between the PTO and DFL (SPO+ and DBB) methods, with and without AFE, using the Friedman average ranking (lower is better). The significance level is 0.05.

Table 5
Daily test regrets (lower is better) averaged across ten independent runs to improve the robustness and reliability of the results (see supplementary materials for detailed results of each of the ten runs across the test set days). The lowest mean and standard deviation values are bolded to indicate the best performance. The comparison includes PTO and DFL methods (SPO^+, DBB), each evaluated with and without AFE (see Subsection 3.3 for more details on the experimental procedure). Regret values are measured in monetary units (GBP) and represent the additional electricity cost incurred due to imperfect predictions when making battery scheduling decisions, where a regret of zero indicates optimal decision-making. Because regret equals additional electricity cost, the differences in mean regret quantify cost savings. Averaged over the 14-day test horizon, PTO (AFE) incurred a regret of 0.2046, whereas SPO^+ (AFE) achieved 0.0672, which is an absolute reduction of 0.1374 and a 67.16 % relative reduction. Within each method, AFE further reduced regret: SPO^+ from 0.1543 to 0.0672 (56.48 %), PTO from 0.2539 to 0.2046 (19.42 %), and DBB from 0.3187 to 0.2458 (22.87 %). These results indicate that training the predictor with a decision-focused loss yields materially cheaper day-ahead battery schedules and that AFE consistently improves decision quality. While absolute regret values may appear small in monetary terms because the BESS is sized for a single household, they correspond to sizeable percentage reductions, accumulate over time, and scale with asset size; using the worst and best mean regrets among the compared methods (DBB without AFE: £0.3187; SPO^+ with AFE: £0.0672) gives a daily difference of £0.2515, which is approximately £91.80 per year if sustained. Appendix A provides detailed daily BESS optimisation schedules for the best run of each day using the best-performing method.

Date	Test Set Regrets					
	PTO		SPO^+		DBB	
	AFE	No AFE	AFE	No AFE	AFE	No AFE
11 Feb	0.0620	0.0780	0.0079	0.0611	0.0832	0.1247
12 Feb	0.1234	0.1669	0.0369	0.0845	0.1440	0.1858
13 Feb	0.1573	0.1968	0.0205	0.1222	0.1685	0.2243
14 Feb	0.1856	0.2363	0.0469	0.1736	0.2101	0.2678
15 Feb	0.0800	0.1030	0.0173	0.0571	0.1133	0.1335
16 Feb	0.0599	0.0681	0.0169	0.0496	0.0708	0.1042
17 Feb	0.0985	0.1384	0.0406	0.0880	0.1377	0.1818
18 Feb	0.1495	0.1939	0.0586	0.1142	0.1525	0.2711
19 Feb	0.1483	0.1784	0.0334	0.0891	0.1420	0.2156
20 Feb	0.1144	0.1530	0.0344	0.0909	0.1366	0.1862
21 Feb	0.2495	0.3261	0.0322	0.1395	0.3041	0.4344
22 Feb	0.4604	0.5771	0.1334	0.3199	0.5789	0.7216
23 Feb	0.2994	0.3363	0.2791	0.3367	0.3896	0.4549
24 Feb	0.6756	0.8021	0.1832	0.4343	0.8098	0.9551
Mean	0.2046	0.2539	**0.0672**	0.1543	0.2458	0.3187
Std	0.1732	0.2055	**0.0780**	0.1206	0.2128	0.2474

3.4. Evaluation metric and statistical tests

In this work, model performance is evaluated using the normalised regret metric (Tang & Khalil, 2024). The notion of regret is used to measure the error in decision-making. It is characterised as the difference in the objective value between the true optimal solution and the optimal solution obtained by utilising the predicted coefficients. For minimisation problems, the normalised regret is defined as:

$$\frac{\sum_{i=1}^{n_{\text{test}}} \mathcal{L}_{\text{Regret}}(\hat{c}^i, c^i)}{\sum_{i=1}^{n_{\text{test}}} |z^*(c^i)|} \tag{14}$$

where $\mathcal{L}_{\text{Regret}}(\hat{c}^i, c^i) = c^{i\top} w^*(\hat{c}^i) - z^*(c^i)$ denotes the regret for instance i. Here, $w^*(\hat{c}^i)$ represents the optimal solution obtained using the predicted cost vector $\hat{c}^i$, $c^{i\top} w^*(\hat{c}^i)$ is the objective value achieved by this solution when evaluated under the true cost vector c^i, and $z^*(c^i)$ is the true optimal objective value under the actual cost vector c^i. The regret thus measures the difference between the objective value of the solution derived from predicted costs (evaluated under true costs) and the true optimal objective value. The denominator normalises the aggregated regret by the sum of absolute values of true optimal objectives across the test set, providing a scale-invariant measure of decision quality degradation (Tang & Khalil, 2024).

Non-parametric hypothesis tests were used to identify significant differences between the DFL and PTO methods, as well as between AFE versus without AFE and to support the experimental findings statistically (Sheskin, 2003). The Friedman Aligned-ranks test (García et al., 2010) first evaluated overall differences among the methods with the significance threshold set at $\alpha = 0.05$. Additionally, for pairwise comparisons, the Wilcoxon Signed-Rank test (Demšar, 2006; Garcia & Herrera, 2008) was applied at $\alpha = 0.05$ to explore potential differences between method pairs that the preceding test did not flag as statistically significant, thereby ensuring a comprehensive statistical analysis.

4. Results and discussion

This section presents an analysis and discussion of the results. Subsection 4.1 discusses the overall performance of the PTO and DFL methods for the investigated BESS problem. Subsequently, Subsection 4.2 provides a comparative analysis of the impact of AFE on the downstream optimisation task (i.e. decision quality) supported by feature importance analysis. Finally, Subsection 4.4 evaluates the statistical significance of the observed performance differences between the methods, to verify that the reported results are unlikely to be due to chance. Supplementary materials present the detailed results of each of the ten runs across the test days, while Appendix A illustrates the daily BESS optimisation schedules for the best run of each day using the best-performing method, showing how predicted prices and demand, subject to operational constraints, shaped the BESS scheduling decisions.

4.1. Performance overview of PTO compared to DFL

The performance comparison results between the PTO and DFL methods across the test set are presented in Fig. 6 and Table 5. Across the fourteen-day test horizon, the SPO^+ method outperformed both the traditional PTO and DBB methods. On average, the PTO method with

Table 6
Results of pairwise comparisons on the test set between the PTO and DFL (SPO^+ and DBB) methods, with and without AFE, using the Wilcoxon signed-rank test. Each row tests the null hypothesis that the distributions of Sample 1 (i.e., Method 1) and Sample 2 (i.e., Method 2) are the same. Asymptotic p-values (2-sided tests) are displayed. The significance level is 0.05. P-values below the significance level indicate rejection of the null hypothesis. Values reported as < 0.001 signify p-values less than 0.001, providing strong statistical evidence. [a] Significance values have been adjusted by the Bonferroni correction for multiple tests.

Method 1 vs Method 2	Test Statistic	Std. Test Statistic	Sig.	Adj. Sig.[a]
SPO^+ (AFE) vs SPO^+ (No AFE)	-1.143	-1.616	0.106	1.000
SPO^+ (AFE) vs PTO (AFE)	2.000	2.828	0.005	0.070
SPO^+ (AFE) vs DBB (AFE)	-3.357	-4.748	< 0.001	0.000
SPO^+ (AFE) vs PTO (No AFE)	3.500	4.950	< 0.001	0.000
SPO^+ (AFE) vs DBB (No AFE)	-5.000	-7.071	< 0.001	0.000
SPO^+ (No AFE) vs PTO (AFE)	0.857	1.212	0.225	1.000
SPO^+ (No AFE) vs DBB (AFE)	-2.214	-3.131	0.002	0.026
SPO^+ (No AFE) vs PTO (No AFE)	2.357	3.334	< 0.001	0.013
SPO^+ (No AFE) vs DBB (No AFE)	-3.857	-5.455	< 0.001	0.000
PTO (AFE) vs DBB (AFE)	-1.357	-1.919	0.055	0.824
PTO (AFE) vs PTO (No AFE)	-1.500	-2.121	0.034	0.508
PTO (AFE) vs DBB (No AFE)	-3.000	-4.243	< 0.001	0.000
DBB (AFE) vs PTO (No AFE)	0.143	0.202	0.840	1.000
DBB (AFE) vs DBB (No AFE)	-1.643	-2.323	0.020	0.302
PTO (No AFE) vs DBB (No AFE)	-1.500	-2.121	0.034	0.508

AFE incurred a regret of 0.2046, while the SPO^+ method achieved a remarkably low regret of 0.0672. This improvement corresponds to a substantial reduction of approximately 67.16 %, underscoring the significant benefits of integrating the optimisation layer during training with the SPO^+. In other words, when the model is trained to predict electricity price and property demand while minimising downstream cost using the SPO^+ loss, rather than a generic pointwise error metric (i.e., error-minimising forecasts), it yields cheaper battery schedules. Interestingly, the DBB method, which is considered a DFL approach, achieved the highest average regret (i.e., the highest cost in BESS scheduling). This observation is broadly in line with recent related work (Wang et al., 2025; Zharmagambetov et al., 2023), where DBB underperforms SPO^+, despite differences in datasets, experimental settings and problem formulations. Although DBB can differentiate through black-box combinatorial solvers, the findings of this work indicate that this method is less suitable for BESS optimisation tasks under data scarcity, as evidenced by the higher regrets compared with other methods, potentially due to its original design for problems with linear objectives (Mandi et al., 2024). Additionally, the poor performance of DBB compared to SPO^+ and PTO methods may indicate that DBB requires more training data to learn effective gradient approximations, making it particularly vulnerable in data-scarce environments such as the 55-day dataset examined in this study.

An examination of day-to-day performance reveals that SPO^+ with AFE delivers notably consistent superior performance, demonstrating reliability on a per-instance basis. This consistency is reflected in lower variance (0.0780 vs 0.1732 for PTO and 0.2128 for DBB), as alternative approaches exhibited greater sensitivity to volatile market conditions such as those observed on February 22nd and 24th. The stability demonstrated by SPO^+ suggests that training with regret-based loss may produce models that are more robust to diverse price-demand patterns compared to both traditional PTO and DBB methods. Such findings, while data-dependent, indicate that SPO^+ could offer practical advantages for BESS strategies by potentially reducing exposure to costly high-regret scheduling decisions, though broader validation across different operational contexts would strengthen these conclusions. The consistently higher variance observed in DBB's performance compared to SPO^+ further indicates potential instability in DBB's learning process under data constraints. This instability likely stems from the inherent challenges of approximating gradients through black-box solvers when training data is limited, resulting in less reliable gradient signals during optimisation.

4.2. Impact of AFE on PTO and DFL methods

The addition of AFE significantly enhanced the performance for all methods, evidenced by the results presented in Fig. 6 and Table 5, with particularly pronounced benefits for the DFL approaches. Without AFE, SPO^+ outperformed the other methods with a mean regret of 0.1543, compared to the PTO approach (0.2539) and DBB (0.3187). This initial superiority suggests that the SPO^+ formulation is inherently effective at learning the mapping between inputs and optimal decisions for the investigated problem, even with less feature representations. However, the introduction of AFE amplified these advantages, with SPO^+ achieving 56.48 % improvement in mean regret, while the PTO improved by 19.42 % and DBB by 22.87 %. The overall reductions in regrets indicate that the AFE, while minimising the need for expert-driven FE, successfully generates useful features not only for maximising forecasting accuracy (i.e., the PTO approach) but also for optimising downstream decision-making tasks within DFL frameworks (i.e., SPO^+ and DBB). Moreover, the notable responsiveness of SPO^+ to enhanced feature representations may suggest the synergistic effect between this method and AFE in BESS optimisation problems. It implies that AFE provides SPO^+ with the contextual information necessary to better approximate the discrete optimisation problem via its convex surrogate, leading to more cost-effective battery scheduling decisions. The enriched feature representations could enable the convex surrogate loss to better capture the relationship between prediction errors and their downstream decision costs. The differential response to AFE across methods provides additional insight into DBB's limitations. While SPO^+ achieved a higher improvement with AFE, DBB's improvement was noticeably lower, suggesting that DBB's gradient approximation mechanism may be less capable of effectively leveraging enhanced feature representations. This reduced feature sensitivity could stem from DBB's reliance on perturbation-based gradient estimation, which may not capture the complex relationships between engineered features and optimal decisions as effectively as the convex surrogate approach of SPO^+.

Nevertheless, the magnitude of AFE's influence varied across test set days, yet no instances of AFE-induced performance degradation were observed across any method. For example, adding AFE to the SPO^+ method improved its performance, leading to regret reductions of 83.22 % and 62.2 % on February 13, and 20, respectively. However, on some days, the improvements were less pronounced. For instance, on 23 February, SPO^+ improved by 17.1 % (from 0.3367 to 0.2791), indicating that while AFE consistently enhances performance (i.e., yielding lower regrets than without AFE), its effectiveness may vary depending on the relevance of the features generated to the underlying data patterns.

Beyond average improvements, AFE also enhanced the consistency and reliability of performance across all methods. The standard deviations of regrets decreased for all approaches with AFE: from 0.2055 to 0.1732 for PTO (15.72 % reduction), from 0.1206 to 0.0780 for SPO^{+} (35.32 % reduction), and from 0.2474 to 0.2128 for DBB (13.9 % reduction). Notably, SPO^{+} achieved the largest proportional variance reduction from AFE, indicating potential suitability for BESS optimisation problems where data scarcity heightens the importance of methods that respond strongly to AFE.

4.3. Feature importance analysis across PTO and DFL methods

SHAP values (SHapley Additive exPlanations) were used to interpret the relationships learned by the PTO and DFL methods for predicting electricity price and property demand simultaneously. Across all feature combinations, SHAP computes each input's average marginal effect on model outputs (i.e., predictions), thereby quantifying both the magnitude and the direction (positive or negative) of feature influence (Lundberg & Lee, 2017). As shown in Fig. 7, features generated through AFE demonstrate considerable temporal variation in their impact on model outputs across different test days. This day-to-day variation in feature importance likely reflects the dynamic nature of energy consumption patterns and market conditions during the test period. For instance, while engineered features dominated the importance rankings on certain days (e.g., 17 and 18 February), raw features such as weather data maintained higher importance scores on others (e.g., 14 and 22 February). This temporal variability suggests that the relative value of different information sources fluctuates based on underlying consumption patterns, weather conditions, and market dynamics, highlighting the importance of adaptive AFE in operational BESS systems.

Within the AFE categories generated by the AutoEnergy algorithm (see Subsection 2.4), statistical nested rolling-window and lag features consistently demonstrate higher impact scores compared to temporal features (e.g., hour of day and its sine/cosine transformations) across all methods. This dominance of historical pattern-based features may reflect strong autocorrelation in electricity demand time series, particularly due to building thermal inertia (Martínez Comesaña et al., 2020), where past usage influences near-future demand through gradual thermal changes, making historical patterns more predictive than calendar-based patterns for next-day forecasting. The consistent importance of lag and rolling-window features across different methods (PTO, SPO^{+}, and DBB) suggests that these feature types capture fundamental underlying relationships in energy data that remain valuable regardless of the learning approach employed.

While feature importance analysis is inherently dataset-dependent and these observations may vary across different datasets and environments, a key practical insight emerges in real-world BESS optimisation scenarios where additional exogenous features are unavailable (e.g., when weather data are not accessible), prediction performance can still be substantially improved by generating richer inputs from timestamp and historical data alone through AFE. This finding has potential implications for deployment in data-constrained environments where external data sources may be unreliable or unavailable.

4.4. Statistical significance analysis

As shown in Fig. 8, Friedman's test applied to the fourteen-day test horizon reveals a statistically significant global difference among the six methods ($p < 0.001$), providing compelling evidence that some of the observed performance variations reflect genuine methodological distinctions in handling BESS optimisation under data scarcity rather than random chance. The mean ranks demonstrate that methods with AFE occupy the top three positions out of the top four ranks, with SPO^{+} achieving the best rank, suggesting the superior decision-making efficacy of DFL approaches when combined with AFE for BESS optimisation problems.

Pairwise comparisons using the Wilcoxon signed-rank test with Bonferroni correction, as shown in Table 6, pinpoint where differences are statistically meaningful. SPO^{+} (AFE) demonstrates statistically significant superiority over multiple competing methods. Most notably, SPO^{+} (AFE) significantly outperforms both DBB variants and the PTO approach without AFE (adjusted $p < 0.001$), reflecting substantial and consistent reductions in regret across the test period. Intriguingly, differences between SPO^{+} (AFE) and PTO (AFE) or between SPO^{+} (AFE) and its own no-AFE baseline are not significant after correction (adjusted $p = .070$ and 1.000, respectively), suggesting that the superiority of SPO^{+} (AFE) over PTO (AFE) or improvements from AFE alone cannot be firmly established (i.e., are not statistically robust under correction). None of the within-method comparisons (i.e., PTO (AFE) vs PTO (No AFE), DBB (AFE) vs DBB (No AFE), or SPO^{+} (AFE) vs SPO^{+} (No AFE)) remain significant after Bonferroni adjustment, despite unadjusted p-values below 0.05. Nevertheless, AFE consistently provided substantial practical improvements across all methods, with performance gains of 56.48 %, 19.42 %, and 22.87 % for SPO^{+}, PTO, and DBB, respectively, without degrading performance on any test day. The conservative Bonferroni correction reduces the power to detect within-method AFE benefits, and the limited test horizon of fourteen days, while sufficient to detect strong global differences, constrains the statistical power to identify more subtle pairwise differences after stringent multiple comparison corrections. Despite these statistical limitations, the consistent directional improvements and substantial magnitude of performance gains across all methods provide compelling evidence for the practical value of AFE in BESS optimisation problems.

5. Conclusion, limitations and future work

5.1. Conclusion

This work addresses gaps that challenge effective BESS optimisation, stemming from: A) DFL methods promise to align prediction with downstream objectives (Mandi et al., 2020; Wilder et al., 2019); however, they are relatively new and have been tested primarily on synthetic datasets or small-scale problems (i.e., simplified benchmarks) (Mandi, 2023; Zhou et al., 2024), highlighting the need to assess their practical viability in real-world applications such as BESS problems; and B) real-world datasets often exhibit greater variability and data scarcity due to practical constraints (Grinsztajn et al., 2022; Hollmann et al., 2022) which can compromise DFL performance and necessitate enhanced feature representations to extract richer information from limited data without the need for domain expertise. This work proposes a decision-aware, end-to-end ML framework that directly addresses each gap: first, to overcome data scarcity limitations, the framework leverages domain-specific AFE (Alkhulaifi et al., 2025) to extract richer representations without heavy reliance on domain expertise; second, rather than treating forecasting and optimisation as separate tasks, the framework jointly learns to forecast electricity demand and prices while optimising battery operations using a regret-based objective, ensuring prediction errors directly inform decision quality; and third, the study evaluates this framework using a novel dataset collected from a UK household property, providing empirical assessment and evidence of the practical viability of DFL methods in real-world BESS applications.

In this framework, two DFL methods (SPO^{+} and DBB) and a PTO approach are compared under baseline and AFE-enhanced conditions. The results show that AFE improves the performance of all three methods, with particularly pronounced benefits for the DFL approaches (SPO^{+}: 56.48 %, DBB: 22.87 %, PTO: 19.42 %). Statistical analysis further confirms significant global differences among methods ($p < 0.001$), with SPO^{+} (AFE) ranking the best (i.e., lowest regret), demonstrating its superior decision-making efficacy for BESS optimisation under data scarcity.

The significance of this work extends beyond the proposed framework and methodological comparisons to practical implications for energy management systems, especially pertinent for real-world BESS deployments, where acquiring large, high-quality datasets is difficult and manual FE is time-consuming and error-prone. The findings indicate that in energy cost-aware scheduling, where future demand and prices are uncertain, aligning prediction and optimisation via DFL's regret-based training with AFE yields tangible economic benefits through more effective charge/discharge schedules. However, the study also shows that some DFL methods may slightly underperform conventional PTO (i.e., DBB has higher regrets, on average, than PTO), underscoring that sophisticated, emerging learning paradigms alone do not guarantee superior decision quality in data-scarce, real-world BESS applications. It is worth noting that while the problem formulation presented in Section 2, particularly the mathematical formulation in Subsection 2.2, is tailored specifically for BESS applications, the overarching framework methodology is adaptable to other energy management contexts. The core principles of jointly forecasting uncertain parameters while optimising operational decisions remain applicable, though the specific constraints, decision variables, and objective functions would require modification for different energy systems.

5.2. Limitations and future work

Despite the promising results, certain limitations warrant consideration. The study's relatively small dataset (55 days) collected from a single UK property during winter months (i.e., from 1 January to 24 February 2025) may present several generalisation limitations: A) findings may not extend to other regions or countries with different energy market structures, regulatory frameworks, or pricing mechanisms; and B) the winter-period data may not capture seasonal variations in consumption patterns, renewable generation potential, or temperature-dependent demand behaviours that occur across different seasons and climatic conditions. Furthermore, hyperparameter optimisation was constrained to a limited search space due to computational resource limitations, as detailed in Table 4, which may have prevented the discovery of optimal model architectures and learning parameters that could have further improved performance. More extensive hyperparameter optimisation (e.g., larger grid searches) might yield additional performance gains. Another limitation worth acknowledging is that the BESS optimisation model uses several assumptions (e.g., treating all operational parameters as deterministic and modelling the system as a single aggregated unit) that may not fully capture the complexities and uncertainties inherent in real-world BESS installations, which may comprise multiple battery units with varying characteristics and operational constraints. Finally, although the proposed framework and conducted experiments used a real-world dataset, the conclusions drawn, though insightful, remain data-dependent and would require validation across diverse geographical locations, temporal periods, and energy system configurations to establish broader generalisability and robustness of the DFL-AFE approach

Future research should explore several promising directions to build upon the current study's findings. First, extending the framework to incorporate renewable generation sources (e.g., solar panels) would provide a more comprehensive assessment of DFL's potential in integrated BESSs. Additionally, exploring other DFL approaches, such as noise-contrastive estimation (Mulamba et al., 2021), may yield further improvements. Moreover, testing the framework across longer forecasting horizons, beyond the current one-day-ahead predictions, would strengthen the evidence for DFL's applicability in real-world BESS operations. While the current regret metric focuses solely on cost minimisation, future work could incorporate multi-objective optimisation frameworks that balance economic objectives with environmental considerations, such as carbon emissions reduction, or operational constraints, such as user comfort requirements. Investigating the proposed AFE-DFL framework further with other learning paradigms, such as ensemble methods that combine multiple DFL approaches or transfer learning techniques (Gunduz et al., 2023) that leverage knowledge from larger energy datasets, could further enhance decision quality for BESS problems. Additionally, exploring hybrid approaches that combine AFE-DFL with reinforcement learning (Shen et al., 2024) or meta-learning frameworks may enable more adaptive and robust BESS optimisation strategies that can quickly adjust to changing energy market conditions.

Data availability

Electricity price and weather data will be available on GitHub. Demand data cannot be shared due to permission restrictions.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

The authors gratefully acknowledge Intelligent Plant Ltd (Steve Aitken, Josh Plumbly and Bruce Nicolson) for generously providing the battery and demand data used in this study. This work was supported by U.K. Engineering and Physical Sciences Research Council (EPSRC) for the University of Nottingham under Grant EP/S023305/1; in part by the EPSRC Centre for Doctoral Training in Horizon: Creating Our Lives in Data. I. Triguero's work is part of the TSI-100927-2023-1 Project, funded by the Recovery, Transformation and Resilience Plan from the European Union Next Generation through the Ministry for Digital Transformation and the Civil Service as well as supported by grant PID2023-149128NB-I00 funded by MICIU/AEI 10.13039/501100011033 and by the ERDF of the European Union.

Appendix A. Results of the Optimal BESS Scheduling Strategy Across the Test Set Days

The following figures present daily BESS optimisation results from multiple experimental runs over all test set days, with each subplot showing the best-performing method for that day based on the lowest regret values. As shown in Fig. A.9, and Fig. A.10, the visualisation uses a lollipop chart design where battery energy flows are represented by coloured circles connected to stems: green circles above the zero line indicate charging periods (energy stored), while red circles below represent discharging periods (energy released to meet property demand). Orange background bars show direct grid-to-property energy supply where needed. Predicted electricity prices are overlaid as dashed blue lines referenced to the secondary y-axis. Each subplot title identifies the test set date, the best-performing method, and the corresponding experimental run. The horizontal time axis spans 24 hours (00–23), enabling assessment of temporal scheduling patterns in response to price signals.

Fig. A.9. Daily BESS optimisation results showing best-performing methods per day. Green circles indicate battery charging, red circles show discharging, and orange background bars represent direct grid supply. The dashed blue line shows predicted electricity prices (right y-axis).

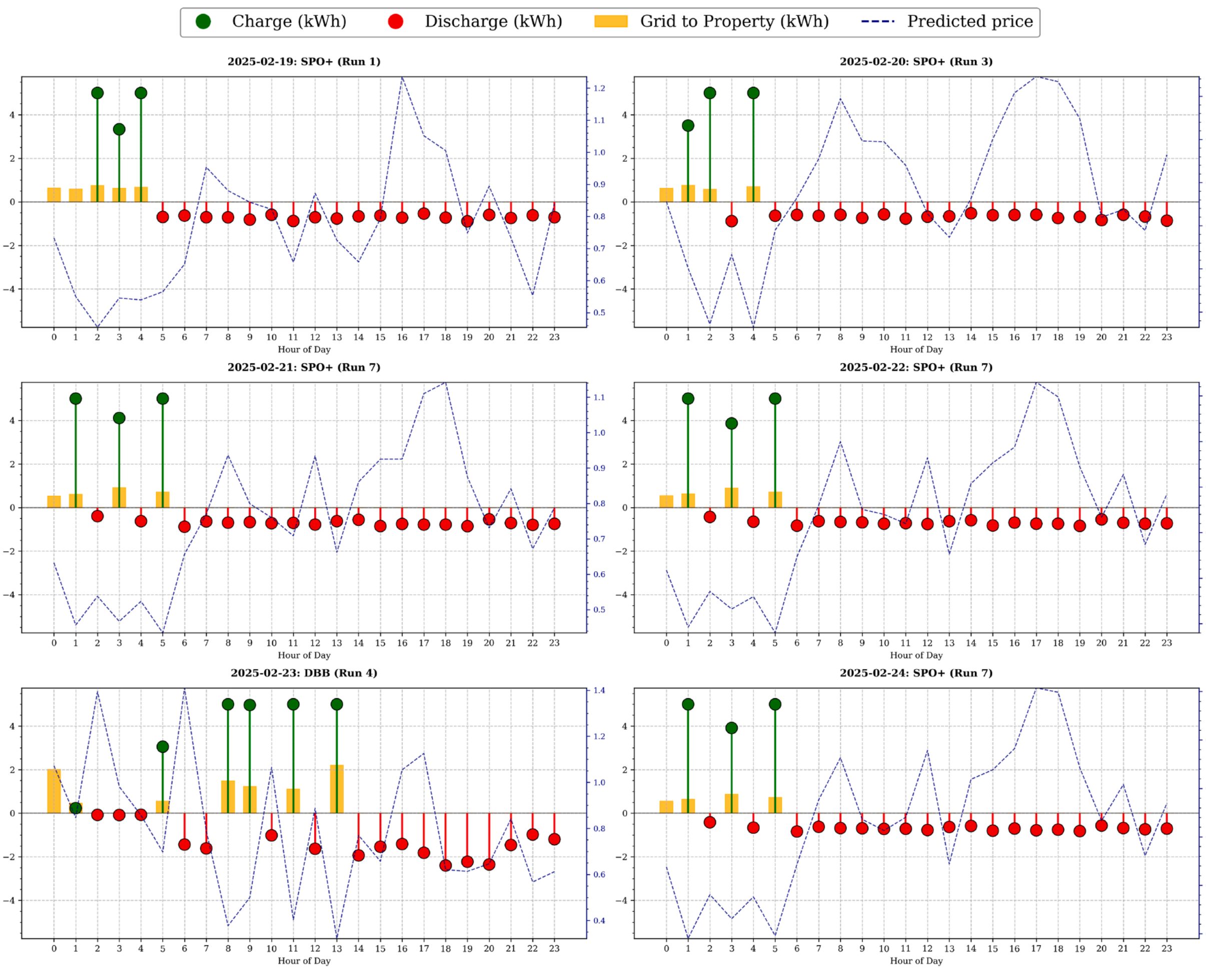

Fig. A.10. Daily BESS optimisation results showing best-performing methods per day. Green circles indicate battery charging, red circles show discharging, and orange background bars represent direct grid supply. The dashed blue line shows predicted electricity prices (right y-axis).

References

Abolghasemi, M., & Bean, R. (2022). How to predict and optimise with asymmetric error metrics. https://arxiv.org/abs/2211.13586.

Alkhulaifi, N., Bowler, A. L., Pekaslan, D., Serdaroglu, G., Closs, S., Watson, N. J., & Triguero, I. (2024a). Machine learning pipeline for energy and environmental prediction in cold storage facilities. *IEEE Access*, *12*, 153935–153951. https://doi.org/10.1109/ACCESS.2024.3482572

Alkhulaifi, N., Bowler, A. L., Pekaslan, D., Triguero, I., & Watson, N. J. (2024b). Exploring automated feature engineering for energy consumption forecasting with autoML. In *2024 IEEE International conference on systems, man, and cybernetics (SMC)* (p. 2993–2998). IEEE. https://doi.org/10.1109/SMC54092.2024.10831959

Alkhulaifi, N., Bowler, A. L., Pekaslan, D., Watson, N. J., & Triguero, I. (2025). Autoenergy: An automated feature engineering algorithm for energy consumption forecasting with autoML. *Knowledge-Based Systems*, (p. 114300). https://doi.org/10.1016/j.knosys.2025.114300

Alrasheedi, A. F., Alnowibet, K. A., & Alshamrani, A. M. (2024). A smart predict-and-optimize framework for microgrid's bidding strategy in a day-ahead electricity market. *Electric Power Systems Research*, *228*, 110016. https://doi.org/10.1016/j.epsr.2023.110016

Anis Lahoud, A., Khan, A. S., Schaffernicht, E., Trincavelli, M., & Stork, J. A. (2025). Predict-and-optimize techniques for data-driven optimization problems: A review. *Neural Processing Letters*, *57*(2). https://doi.org/10.1007/s11063-025-11746-w

Bergmeir, C., de Nijs, F., Genov, E., Sriramulu, A., Abolghasemi, M., Bean, R., Betts, J., Bui, Q., Dinh, N. T., Einecke, N., Esmaeilbeigi, R., Ferraro, S., Galketiya, P., Glasgow, R., Godahewa, R., Kang, Y., Limmer, S., Magdalena, L., Montero-Manso, P., Peralta, D., Kumar, Y. P. S., Rosales-Pérez, A., Ruddick, J., Stratigakos, A., Stuckey, P., Tack, G., Triguero, I., & Yuan, R. (2025). Predict+optimize problem in renewable energy scheduling. *IEEE Access*, *13*, 60064–60087. https://doi.org/10.1109/ACCESS.2025.3555393

Boettiger, C. (2022). The forecast trap. *Ecology Letters*, *25*(7), 1655–1664. https://doi.org/10.1111/ele.14024

Demšar, J. (2006). Statistical comparisons of classifiers over multiple data sets. *The Journal of Machine learning research*, *7*, 1–30. https://www.jmlr.org/papers/volume7/demsar06a/demsar06a.pdf.

Dias Garcia, J., Street, A., Homem-de Mello, T., & Muñoz, F. D. (2025). Application-driven learning: A closed-loop prediction and optimization approach applied to dynamic reserves and demand forecasting. *Operations Research*, *73*(1), 22–39. https://doi.org/10.1287/opre.2023.0565

Donti, P., Amos, B., & Kolter, J. Z. (2017). Task-based end-to-end model learning in stochastic optimization. *Advances in Neural Information Processing Systems*, *30*. https://proceedings.neurips.cc/paper/2017/hash/3fc2c60b5782f641f76bcefc39fb2392-Abstract.html.

Elmachtoub, A. N., & Grigas, P. (2022). Smart "predict, then optimize". *Management Science*, *68*(1), 9–26. https://doi.org/10.1287/mnsc.2020.3922

Energy, O. (2025). Octopus energy API endpoints. https://developer.octopus.energy/rest/guides/endpoints. [Accessed 19-may-2025].

Garcia, S., & Herrera, F. (2008). An extension on" statistical comparisons of classifiers over multiple data sets" for all pairwise comparisons. *Journal of Machine Learning Research*, *9*(12). https://www.jmlr.org/papers/volume9/garcia08a/garcia08a.pdf.

García, S., Fernández, A., Luengo, J., & Herrera, F. (2010). Advanced nonparametric tests for multiple comparisons in the design of experiments in computational intelligence and data mining: Experimental analysis of power. *Information Sciences*, *180*(10), 2044–2064. https://doi.org/10.1016/j.ins.2009.12.010

Geng, H., Ruan, H., Wang, R., Li, Y., Wang, Y., Chen, L., & Yan, J. (2024). Benchmarking pto and pno methods in the predictive combinatorial optimization regime. *Advances in Neural Information Processing Systems*, *37*, 65944–65971. https://doi.org/10.5555/3737916.3740024

Grinsztajn, L., Oyallon, E., & Varoquaux, G. (2022). Why do tree-based models still outperform deep learning on typical tabular data? *Advances in Neural Information Processing Systems*, *35*, 507–520. https://doi.org/10.5555/3600270.3600307

Gunduz, S., Ugurlu, U., & Oksuz, I. (2023). Transfer learning for electricity price forecasting. *Sustainable Energy, Grids and Networks*, *34*, 100996. https://doi.org/10.1016/j.segan.2023.100996

Gurobi Optimization, L. (2024). Python API Reference. https://docs.gurobi.com/projects/optimizer/en/current/reference/python.html. Accessed: 2025-04-01.

Hannan, M. A., Wali, S. B., Ker, P. J., Rahman, M. S. A., Mansor, M., Ramachandaramurthy, V. K., Muttaqi, K. M., Mahlia, T. M. I., & Dong, Z. Y. (2021). Battery energy-storage system: A review of technologies, optimization objectives, constraints, approaches, and outstanding issues. *Journal of Energy Storage*, *42*, 103023. https://doi.org/10.1016/j.est.2021.103023

Hollmann, N., Müller, S., & Hutter, F. (2024). Large language models for automated data science: introducing CAAFE for context-Aware automated feature engineering. *Advances in Neural Information Processing Systems*, *36*. https://doi.org/10.5555/3666122.3668060

Hollmann, N., Müller, S., Eggensperger, K., & Hutter, F. (2022). TabPFN: A Transformer That Solves Small Tabular Classification Problems in a Second. https://arxiv.org/abs/2207.01848. https://doi.org/10.48550/ARXIV.2207.01848

Ibrahim, M. S., Dong, W., & Yang, Q. (2020). Machine learning driven smart electric power systems: current trends and new perspectives. *Applied Energy*, *272*, 115237. https://doi.org/10.1016/j.apenergy.2020.115237

Intelligent Plant Ltd (2025). Intelligent plant – industrial software solutions. Accessed: 2025-05-19 https://www.intelligentplant.com/.

Kotary, J., Fioretto, F., Van Hentenryck, P., & Wilder, B. (2021). End-to-end constrained optimization learning: A survey. arXiv:2103.16378 https://arxiv.org/abs/2103.16378.

Lu, C., Jiang, W., & Wu, C. (2022). Effective end-to-end learning framework for economic dispatch. *IEEE Transactions on Network Science and Engineering*, *9*(4), 2673–2683. https://doi.org/10.1109/TNSE.2022.3168845

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in neural information processing systems*. Curran Associates, Inc. (*vol. 30*). https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html.

Mandi, J. (2023). Towards Scalable Decision-Focused Learning for Combinatorial Optimization Problems. Ph.D. thesis. Vrije Universiteit Brussel. https://cris.vub.be/ws/portalfiles/portal/105442104/Jayanta_Mandi_PhD_Dissertation.pdf.

Mandi, J., Bucarey, V., Tchomba, M. M. K., & Guns, T. (2022). Decision-focused learning: Through the lens of learning to rank. In *International conference on machine learning* (pp. 14935–14947). PMLR. https://proceedings.mlr.press/v162/mandi22a.html.

Mandi, J., Demirovi, E., Stuckey, P. J., & Guns, T. (2020). Smart predict-and-optimize for hard combinatorial optimization problems. *Proceedings of the AAAI Conference on Artificial Intelligence*, *34*(02), 1603–1610. https://doi.org/10.1609/aaai.v34i02.5521

Mandi, J., & Guns, T. (2020). Interior point solving for lp-based prediction + optimisation. *Advances in Neural Information Processing Systems*, *33*, 7272–7282. https://doi.org/10.48550/arXiv.2010.13943

Mandi, J., Kotary, J., Berden, S., Mulamba, M., Bucarey, V., Guns, T., & Fioretto, F. (2024). Decision-focused learning: Foundations, state of the art, benchmark and future opportunities. *Journal of Artificial Intelligence Research*, *80*, 1623–1701. https://doi.org/10.1613/jair.1.15320

Martínez Comesaña, M., Febrero-Garrido, L., Troncoso-Pastoriza, F., & Martínez-Torres, J. (2020). Prediction of building's thermal performance using LSTM and MLP neural networks. *Applied Sciences*, *10*(21), 7439. https://doi.org/10.3390/app10217439

Mulamba, M., Mandi, J., Diligenti, M., Lombardi, M., Bucarey, V., & Guns, T. (2021). Contrastive losses and solution caching for predict-and-optimize. In *Proceedings of the thirtieth international joint conference on artificial intelligence* IJCAI-2021 (pp. 2833–2840). International Joint Conferences on Artificial Intelligence Organization.https://doi.org/10.24963/ijcai.2021/390

NASA (2025). Prediction of worldwide energy resources (POWER). https://power.larc.nasa.gov/data-access-viewer/. [Accessed 19-may-2025].

Paredes, Á., Toubeau, J.-F., Aguado, J. A., & Vallée, F. (2025). On the participation of energy storage systems in reserve markets using decision focused learning. *Sustainable Energy, Grids and Networks*, *42*, 101677. https://doi.org/10.1016/j.segan.2025.101677

Pogančić, M. V., Paulus, A., Musil, V., Martius, G., & Rolinek, M. (2020). Differentiation of blackbox combinatorial solvers. In *International conference on learning representations*. https://openreview.net/forum?id=BkevoJSYPB.

Reza, M. S., Hannan, M. A., Ker, P. J., Mansor, M., Lipu, M. S. H., Hossain, M. J., & Mahlia, T. M. I. (2023). Uncertainty parameters of battery energy storage integrated grid and their modeling approaches: A review and future research directions. *Journal of Energy Storage*, *68*, 107698. https://doi.org/10.1016/j.est.2023.107698

Sang, L., Xu, Y., Long, H., Hu, Q., & Sun, H. (2022). Electricity price prediction for energy storage system arbitrage: A decision-focused approach. *IEEE Transactions on Smart Grid*, *13*(4), 2822–2832. https://doi.org/10.1109/TSG.2022.3166791

Shen, H., Shen, X., & Chen, Y. (2024). Real-time microgrid energy scheduling using meta-reinforcement learning. *Energies*, *17*(10), 2367. https://doi.org/10.3390/en17102367

Sheskin, D. J. (2003). Handbook of Parametric and Nonparametric Statistical Procedures: Third Edition. Chapman and Hall/CRC. https://doi.org/10.1201/9781420036268

Song, H., Liu, C., Amani, A. M., Gu, M., Jalili, M., Meegahapola, L., Yu, X., & Dickeson, G. (2024). Smart optimization in battery energy storage systems: an overview. *Energy and AI*, *17*, 100378. https://doi.org/10.1016/j.egyai.2024.100378

Stratigakos, A., Camal, S., Michiorri, A., & Kariniotakis, G. (2022). Prescriptive trees for integrated forecasting and optimization applied in trading of renewable energy. *IEEE Transactions on Power Systems*, *37*(6), 4696–4708. https://doi.org/10.1109/TPWRS.2022.3152667

Tang, B., & Khalil, E. B. (2024). PyEPO: A pytorch-based end-to-end predict-then-optimize library for linear and integer programming. *Mathematical Programming Computation*, *16*(3), 297–335. https://doi.org/10.1007/s12532-024-00255-x

Vanderschueren, T., Verdonck, T., Baesens, B., & Verbeke, W. (2022). Predict-then-optimize or predict-and-optimize? an empirical evaluation of cost-sensitive learning strategies. *Information Sciences*, *594*, 400–415. https://doi.org/10.1016/j.ins.2022.02.021

Wang, G., Lin, Z., Chen, Y., Ji, H., Chen, D., Zhang, H., Li, P., & Yan, J. (2025). Ai-optimized decision-making in energy systems: toward a decision-aware machine learning framework. *Engineering*, . https://doi.org/10.1016/j.eng.2025.06.037

Wang, L., Tan, M., Chen, J., & Liao, C. (2022a). Multi-task learning based multi-energy load prediction in integrated energy system. *Applied Intelligence*, *53*(9), 10273–10289. https://doi.org/10.1007/s10489-022-04054-6

Wang, Z., Xia, L., Yuan, H., Srinivasan, R. S., & Song, X. (2022b). Principles, research status, and prospects of feature engineering for data-driven building energy prediction: A comprehensive review. *Journal of Building Engineering*, *58*, 105028. https://doi.org/10.1016/j.jobe.2022.105028

Wilder, B., Dilkina, B., & Tambe, M. (2019). Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. *Proceedings of the AAAI Conference on Artificial Intelligence*, *33*(01), 1658–1665. https://doi.org/10.1609/aaai.v33i01.33011658

Wu, Y., Xi, X., & He, J. (2022). AFGSL: Automatic feature generation based on graph structure learning. *Knowledge-Based Systems*, *238*, 107835. https://doi.org/10.1016/j.knosys.2021.107835

Yang, Y., Bremner, S., Menictas, C., & Kay, M. (2022). Modelling and optimal energy management for battery energy storage systems in renewable energy systems: A review. *Renewable and Sustainable Energy Reviews*, *167*, 112671. https://doi.org/10.1016/j.rser.2022.112671

Yu, S., You, L., & Zhou, S. (2023). A review of optimization modeling and solution methods in renewable energy systems. *Frontiers of Engineering Management*, *10*(4), 640–671. https://doi.org/10.1007/s42524-023-0271-3

Zhang, H., Li, R., Du, Q., Tao, J., Pineda, S., Kariniotakis, G., Camal, S., Monroc, C. B., Sun, M., Wan, C., Xu, W., & Teng, F. (2025a). Decision-focused learning for power system decision-making under uncertainty. *IEEE Transactions on Power Systems*, (p. 1–18). https://doi.org/10.1109/tpwrs.2025.3597806

Zhang, J., Li, Y., & Xiao, W. (2021). Integrated multiple kernel learning for device-free localization in cluttered environments using spatiotemporal information. *IEEE Internet of Things Journal*, *8*(6), 4749–4761. https://doi.org/10.1109/JIOT.2020.3028574

Zhang, J., Xue, J., Li, Y., & Cotton, S. L. (2025b). Leveraging online learning for domain-adaptation in wi-fi-based device-free localization. *IEEE Transactions on Mobile Computing*, *24*(8), 7773–7787. https://doi.org/10.1109/TMC.2025.3552538

Zhang, Y., Wen, H., & Wu, Q. (2024). A contextual bandit approach for value-oriented prediction interval forecasting. *IEEE Transactions on Smart Grid*, *15*(2), 2271–2281. https://doi.org/10.1109/TSG.2023.3296577

Zharmagambetov, A., Amos, B., Ferber, A., Huang, T., Dilkina, B., & Tian, Y. (2023). Landscape surrogate: Learning decision losses for mathematical optimization under partial information. *Advances in Neural Information Processing Systems*, *36*, 27332–27350. https://proceedings.neurips.cc/paper_files/paper/2023/hash/574f145eac328cc4aaf9358e27120eb5-Abstract-Conference.html.

Zhou, H., Huang, R., Li, S., Jiang, G., Zheng, J., Cheng, B., & Lin, W. (2024). Decision focused causal learning for direct counterfactual marketing optimization. In *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining* (pp. 6368–6379). ACM. https://doi.org/10.1145/3637528.3672353