Bridging the Development–Production Gap in Industrial Energy Management Systems: An Experimental Study for Data Drifts

A. Zacharatos[1], C. Giannakis[1], and A. Bachoumis[1]

[1] *AI laboratory of enakronIC PC, Aspasias 72, 15561, Cholargos, Athens*

Keywords: DATA DRIFT, INDUSTRIAL EMS, CONTINUAL LEARNING, ROBUST OPTIMIZATION

Abstract

Industrial Energy Management Systems (EMSs) combine data-driven forecasting with optimisation to schedule Behind-the-Meter (BTM) batteries, Photovoltaic (PV) generation, grid exchange, and flexible loads. However, when industrial operating conditions change after deployment, offline-trained models degrade, affecting the quality of operational decisions. This paper studies the impact of data drift on an industrial warehouse EMS with a BTM battery, PV generation, and mandatory forklift charging. Drift regimes are detected using a multivariate Maximum Mean Discrepancy (MMD) permutation test and grouped into non-drift, forklift-drift, and demand-drift periods. Six prescriptive configurations are compared: static Predict-then-Optimize (PTO) and Predict-and-Optimize (PAO) baselines, their continual-learning (CL) variants with replay memory, and two robust optimisation (RO) formulations based on box and ellipsoidal uncertainty sets. Results show that forecasting accuracy is an insufficient proxy for decision quality, as lower MAE does not necessarily yield better schedules. CL improves performance under critical-load drift, reducing the optimality gap by up to 83%, while ellipsoidal RO achieves the best aggregate performance, reducing the optimality gap by 87% relative to PTO. These findings motivate drift-aware, decision-centric evaluation for deployed industrial EMSs.

1 Introduction

Industrial Energy Management Systems (EMSs) coordinate Distributed Energy Resources (DERs), principally Photovoltaic (PV) generation, Behind-the-Meter (BTM) Battery Energy Storage Systems (BESS), and heterogeneous loads, to minimise operating costs in industrial facilities. Operational decisions are taken over a finite scheduling horizon by combining forecasts of demand, DER generation, and flexible-load constraints with an optimisation layer that determines battery operation, grid exchange, and controllable load setpoints. Decision quality directly affects electricity costs, operational feasibility, BESS utilisation, and load satisfaction, particularly when critical shiftable loads, such as forklift charging, must be served within daily operational constraints.

A central challenge in deployment is that the data distribution observed during offline development may no longer hold over time [1]. Industrial facilities are subject to changing production schedules, PV degradation, equipment additions, and other factors. Such persistent temporal shifts, i.e., data drift, cause static forecasters to degrade under new operating regimes and render downstream schedules suboptimal, even when the EMS formulation itself remains unchanged. This degradation is not necessarily reflected in forecasting accuracy: a model may retain a low mean absolute error while producing schedules that are substantially more costly or infeasible. Model performance after deployment must, therefore, be evaluated through the realised operational outcomes of the resulting schedules, rather than through forecasting accuracy alone, a requirement that motivates this paper.

2 Research gap and contribution

A key limitation of the current literature is that offline-trained predictors for EMS cannot adapt to changes, such as modified production schedules, equipment additions and shifting operational patterns, without expensive full retraining [1]. Online learning offers one possible mechanism for adapting to incoming data streams, but it also introduces a competing risk, i.e., sequential updates without any replay mechanism may overwrite previously learned behavior, a well-established phenomenon known as *catastrophic forgetting* [1, 2]. Therefore, memory-augmented rehearsal strategies, such as Experience Replay (ER) and its extensions, have been proposed to mitigate forgetting by re-exposing the model to past samples stored in a fixed-size buffer, thereby stabilising training in non-stationary regimes [2, 3].

A further limitation concerns the coupling between prediction accuracy and operational performance. The standard Predict-then-Optimize (PTO) paradigm treats forecasting and optimization as separate design objectives [4]. However, as noted in [1], the separate design of the predictive and optimization layers can lead to sub-optimal control decisions when prediction errors are heterogeneous across time steps or load types. End-to-end (E2E), or *decision-aware*, training optimizes the forecaster directly with respect to the downstream cost objective and has been shown to improve EMS scheduling performance under certain settings [1]. However, whether this advantage persists, weakens, or becomes more pronounced under structured data drift remains an open empirical question. Drift-aware EMS evaluation must, therefore, jointly assess predictive accuracy and the downstream decision quality it induces.

Robust optimization (RO) offers a complementary, nonadaptive mechanism for handling forecasts uncertainty. By hedging against worst-case realizations within a structured uncertainty set [5], RO EMS frameworks can maintain operational feasibility without requiring model updates, although this is typically achieved at the cost of increased conservatism. Under non-stationary conditions, it remains unclear whether performance gains are best sought through adaptive updates, decision-focused training, or robust optimization, a systematic comparison across drift regimes using operational KPIs as the primary criterion is still absent from the literature.

This paper contributes a drift-aware experimental evaluation framework for industrial EMS operation under chronologically evolving data distributions, compares adaptive continual-learning, decision-focused PAO, and robust optimization strategies under the same EMS use case and downstream KPIs, and demonstrates that forecasting accuracy is not a reliable proxy for EMS decision quality, motivating the use of operational metrics such as optimality gap and unsatisfied critical load.

3 Problem formulations

We consider a grid-connected industrial facility equipped with a BTM BESS, local PV generation, and a mixed portfolio of electrical loads. The EMS solves a day-ahead scheduling problem over T intervals of length Δt, optimizing BESS charge/discharge profiles and controllable load setpoints to minimize total operational cost. Loads are classified by flexibility and criticality: **(i)** inflexible base demand ($\hat{D}_t^{\text{base}}$), which must be fully served at all times; **(ii)** critical shiftable loads (forklift chargers, $\hat{D}_t^{\text{fork}}$), whose total daily energy is operationally mandatory but can be time-shifted; therefore, unserved demand accumulates as a penalised backlog; and **(iii)** non-critical shiftable loads (L2/L3 EV chargers, $\hat{D}_t^{\text{l2}}$, $\hat{D}_t^{\text{l3}}$), served opportunistically based on grid prices and PV availability. The optimization layer relies on multi-series point forecasts $\hat{y}$ generated by an MLP with two 168-unit hidden layers and a non-negative output. The interplay between the predictive module and the optimization layer is the core subject of our evaluation under data drift. Table 1 summarizes the six evaluated configurations by optimization layer type and test-time adaptation strategy. The time horizon is indexed $t \in \mathcal{T} = \{1, \ldots, T\}$ with step Δt; all notation shared across formulations is defined in Table 2.

Table 1: Taxonomy of the evaluated modeling approaches.

Approach	Optimization Layer	Runtime Model Update
PTO Baseline	Deterministic MILP	None (Static parameters)
PAO Baseline	Deterministic Convex LP	None (Static parameters)
PTO CL	Deterministic MILP	Monthly retraining
PAO CL	Deterministic Convex LP	Monthly retraining
Box-RO	Rob. Convex LP – interval box	None (Static parameters & margins)
Ellipsoidal-RO	Rob. Convex LP – horizon covariance	None (Static parameters & margins)

Table 2: Notation shared across all formulations.

Symbol	Type	Unit	Description
Forecasted Inputs			
$\hat{D}_t^{\text{base}}$	Param.	kWh	Inflexible base energy demand
$\widehat{PV}_t$	Param.	kW	Photovoltaic power generation
$\hat{D}_t^{\text{fork}}$	Param.	kWh	Forklift charging energy demand
$\hat{D}_t^{\text{l2}}$	Param.	kWh	L2 EV charging energy demand
$\hat{D}_t^{\text{l3}}$	Param.	kWh	L3 EV charging energy demand
System Parameters			
$E_{\text{cap}}^{\text{BESS}}$	Param.	kWh	Nominal battery capacity
E_0, E_r	Param.	kWh	Initial SOC; terminal reserve target
$E_{\min}$, $E_{\max}$	Param.	kWh	$0.1\,E_{\text{cap}}^{\text{BESS}}$, $0.9\,E_{\text{cap}}^{\text{BESS}}$
$P_s^{\max}$	Param.	kW	Max battery power magnitude
η_{ch}, η_{dis}	Param.	–	Charge / discharge efficiencies
$P_{\max}^{\text{fork}}$	Param.	kW	Maximum forklift charging power
$P_{\text{net}}^{\max}$	Param.	kW	Grid power exchange limit
Market Signals & Penalty Coefficients			
c_t^{buy}, c_t^{sell}	Param.	€/kWh	Grid purchase / feed-in price
γ_{cycle}	Param.	€/kWh	Battery degradation penalty
γ_{fl}	Param.	€/kWh	Forklift unserved energy penalty
γ_{l2}, γ_{l3}	Param.	€/kWh	L2/L3 unserved energy penalties
γ_s	Param.	€/unit	Terminal SOC slack penalty
Decision Variables			
E_t	Var.	kWh	Battery energy level
pt_t^{fork}, pt_t^{l2}, pt_t^{l3}	Var.	kW	Load power setpoints
C_t	Var.	kWh	Cumulative forklift unsatisfied load
s_R	Var.	kWh	Terminal SOC slack
PTO-Specific Variables			
ch_t	Var.	kW	Battery charging power
dis_t	Var.	kW	Battery discharging power
$P_{\text{net,Imp},t}$	Var.	kW	Grid import power
$P_{\text{net,Exp},t}$	Var.	kW	Grid export power
y_t	Binary Var.	–	Charge/discharge exclusivity
z_t	Binary Var.	–	Import/export exclusivity
PAO-Specific Parameters and Variables			
$P_{\text{S,lower},t}$, $P_{\text{S,upper},t}$	Param.	kW	Adaptive battery power bounds
$P_{s,t}$	Var.	kW	Net battery power exchange
$P_{\text{net},t}$	Var.	kW	Net grid active power
$p_{L,t}$	Var.	kW	Convexified battery loss power
$\tilde{P}_{\text{ch},t}$, $\tilde{P}_{\text{dis},t}$	Var.	kW	DCP-safe battery aux. variables
$\tilde{P}_{\text{imp},t}$, $\tilde{P}_{\text{exp},t}$	Var.	kW	DCP-safe grid auxiliary variables
RO-Box Specific Parameters			
Δ_t^f	Param.	kWh	Forklift worst-case energy margin
Δ_t^w	Param.	kW	Net-load grid connection margin
RO-Ellipsoidal Specific Parameters			
Δ_t^w	Param.	kW	Per-timestep net-load tightening margin ($\rho_w \lVert L_w^\top e_t \rVert_2$)
Δ^f	Param.	kWh	Cumulative forklift margin ($\rho_f \lVert L_f^\top \mathbf{1} \rVert_2$)

3.1 PTO Formulation

The PTO layer is a deterministic MILP: binary variables y_t and z_t enforce charge/discharge and import/export exclusivity, respectively.

Objective Function:

$$\min\ C = \sum_{t\in\mathcal{T}} \Bigg[\left(c_t^{\text{buy}} P_{\text{net,Imp},t} - c_t^{\text{sell}} P_{\text{net,Exp},t} \right) \Delta t + \gamma_{\text{cycle}} \left(ch_t + dis_t \right) \Delta t + \gamma_{\text{l2}} U_{\text{l2},t} + \gamma_{\text{l3}} U_{\text{l3},t} \Bigg] + \gamma_{\text{fl}}\, C_{T+1} + \gamma_s\, s_R \tag{1}$$

s.t.: *Battery Dynamics and Exclusivity:*

$$E_1 = E_0, \tag{2}$$

$$E_{T+1} \geq E_r - s_R, \tag{3}$$

$$E_{\min} \leq E_t \leq E_{\max}, \quad \forall t \in \mathcal{T} \tag{4}$$

$$ch_t \leq P_s^{\max}\, y_t, \qquad \forall t \in \mathcal{T} \tag{5}$$

$$dis_t \leq P_s^{\max}\,(1 - y_t), \forall t \in \mathcal{T} \tag{6}$$

$$E_{t+1} = E_t + \Delta t\Big(\eta_{\text{ch}}\, ch_t - \tfrac{1}{\eta_{\text{dis}}}\, dis_t\Big), \quad \forall t \in \mathcal{T} \tag{7}$$

Grid Flow and Power Balance:

$$P_{\text{net,Imp},t} \le P_{\text{net}}^{\max}\, z_t, \qquad \forall t \in \mathcal{T} \tag{8}$$

$$P_{\text{net,Exp},t} \le P_{\text{net}}^{\max}\,(1 - z_t), \forall t \in \mathcal{T} \tag{9}$$

$$\begin{aligned} P_{\text{net,Imp},t} - P_{\text{net,Exp},t} = \frac{\hat{D}_t^{\text{base}}}{\Delta t} - \widehat{PV}_t + ch_t - dis_t \\ + pt_t^{\text{fork}} + pt_t^{\text{l2}} + pt_t^{\text{l3}}, \quad \forall t \in \mathcal{T} \end{aligned} \tag{10}$$

Shiftable Loads:

$$U_{\text{l2},t} = \hat{D}_t^{\text{l2}} - pt_t^{\text{l2}}\,\Delta t, \forall t \in \mathcal{T} \tag{11}$$

$$U_{\text{l3},t} = \hat{D}_t^{\text{l3}} - pt_t^{\text{l3}}\,\Delta t, \forall t \in \mathcal{T} \tag{12}$$

$$pt_t^{\text{fork}} \le P_{\max}^{\text{fork}}, \qquad \forall t \in \mathcal{T} \tag{13}$$

$$C_1 = 0, \quad C_{t+1} = C_t + \hat{D}_t^{\text{fork}} - pt_t^{\text{fork}}\,\Delta t, \quad \forall t \in \mathcal{T} \tag{14}$$

Penalty calibration: Coefficients are scaled relative to the mean purchase price $\bar{c}^{\text{buy}} = \frac{1}{T}\sum_t c_t^{\text{buy}}$ to enforce an operational hierarchy: forklift demand receives a prohibitive penalty $\gamma_{\text{fl}} = 10\,\bar{c}^{\text{buy}}$; non-critical EV loads a moderate $\gamma_{\text{l2}} = \gamma_{\text{l3}} = 2\,\bar{c}^{\text{buy}}$; and terminal SOC violations an aggressive $\gamma_s \ge 10\,\bar{c}^{\text{buy}}$.

3.2 PAO Formulation

To support end-to-end gradient-based training, the MILP is reformulated as a Disciplined Convex Programming (DCP)-compliant convex LP, i.e., binary variables are dropped and BESS operation is represented by a single free variable $P_{s,t}$ (> 0 charging, < 0 discharging), with internal losses handled via a McCormick envelope of the bilinear loss term $p_{L,t}$. Regressor parameters θ are then trained by minimizing a *realized cost*:

$$\mathcal{L}_{\text{PAO}}(\theta) = w\, C_{\text{ops}}(x^*(\hat{y}(\theta))\,, \tilde{y}) + (1-w)\sum_s \ell(\hat{y}_s, \tilde{y}_s) \tag{15}$$

where $x^*(\hat{y}(\theta))$ is the day-ahead schedule produced by the optimization layer on forecast $\hat{y}$, $\tilde{y}$ are the ground-truth realizations, C_{ops} aggregates realized energy cost, load penalties, and cycling cost, and $\ell(\cdot,\cdot)$ is a per-series Huber loss. Gradients $\nabla_\theta \mathcal{L}$ are computed via implicit differentiation with respect to the KKT conditions of the convex layer [6].

Objective Function:

$$\begin{aligned} \min\ C_{\text{PAO}} = \sum_{t\in\mathcal{T}} \Big[& \Big(c_t^{\text{buy}} \tilde{P}_{\text{imp},t} - c_t^{\text{sell}} \tilde{P}_{\text{exp},t}\Big)\Delta t \\ & + \gamma_{\text{cycle}}\Big(\tilde{P}_{\text{ch},t} + \tilde{P}_{\text{dis},t}\Big)\Delta t \\ & + \gamma_{\text{l2}} U_{\text{l2},t} + \gamma_{\text{l3}} U_{\text{l3},t}\Big] \\ & + \gamma_{\text{fl}} C_{T+1} + \gamma_s s_R. \end{aligned} \tag{16}$$

s.t.: *Battery Dynamics and McCormick Relaxation*: Boundary and capacity limits match (2)–(4). Defining loss coefficients $\lambda_c = 1 - \eta_{\text{ch}}$ and $\lambda_d = 1 - \eta_{\text{dis}}$, the McCormick envelope [7] for $p_{L,t}$ is:

$$p_{L,t} \ge \lambda_c\, P_{s,t}, \quad \forall t \in \mathcal{T} \tag{17}$$

$$p_{L,t} \ge -\lambda_d\, P_{s,t}, \quad \forall t \in \mathcal{T} \tag{18}$$

$$p_{L,t} \le \lambda_c P_s^{\max} + a(-P_{s,t} + P_s^{\max})\,, \quad \forall t \in \mathcal{T} \tag{19}$$

where $a = (\lambda_d |P_s^{\min}| - \lambda_c P_s^{\max})/(P_s^{\max} + |P_s^{\min}|)$. SOC dynamics and adaptive bounds:

$$E_{t+1} = E_t + \Delta t\,(P_{s,t} - p_{L,t})\,, \quad \forall t \in \mathcal{T} \tag{20}$$

$$P_{\text{S,lower},t} \le P_{s,t} \le P_{\text{S,upper},t}, \quad \forall t \in \mathcal{T} \tag{21}$$

Power Balance and Grid Envelope:

$$\begin{aligned} P_{\text{net},t} = \frac{\hat{D}_t^{\text{base}}}{\Delta t} - \widehat{PV}_t + P_{s,t} \\ + pt_t^{\text{fork}} + pt_t^{\text{l2}} + pt_t^{\text{l3}}, \quad \forall t \in \mathcal{T} \end{aligned} \tag{22}$$

$$|P_{\text{net},t}| \le P_{\text{net}}^{\max}, \quad \forall t \in \mathcal{T} \tag{23}$$

Shiftable Loads: (11), (12), (13), (14)

DCP-Safe Auxiliary Decomposition:

$$\begin{aligned} P_{s,t} = \tilde{P}_{\text{ch},t} - \tilde{P}_{\text{dis},t}, \quad \tilde{P}_{\text{ch},t}, \tilde{P}_{\text{dis},t} \ge 0 \\ P_{\text{net},t} = \tilde{P}_{\text{imp},t} - \tilde{P}_{\text{exp},t}, \quad \tilde{P}_{\text{imp},t}, \tilde{P}_{\text{exp},t} \ge 0 \end{aligned} \tag{24}$$

3.3 Box-RO Formulation

Let $r_{n,t}^s = |y_{n,t}^s - \hat{y}_{n,t}^s|$ denote absolute forecast residuals for series s. Timesteps are assigned to hourly groups $g(t)$, and group-wise conformal quantiles are computed as:

$$q_k^s = Q_{1-\alpha}\big(\{r_{n,t}^s : g(t) = k\}\big) \tag{25}$$

[8] yielding horizon-indexed margins $\delta_t^s = q_{g(t)}^s$. The net-load margin aggregates demand and PV errors as $\delta_t^w = \delta_t^d/\Delta t + \delta_t^{\text{pv}}$. The box uncertainty set is:

$$\mathcal{U}_{\text{box}} = \big\{r \in \mathbb{R}^T : |r_t| \le \delta_t^w,\ \forall\, t\big\}, \tag{26}$$

and robustness is enforced by setting $\Delta_t^w \equiv \delta_t^w$ and $\Delta_t^f \equiv \delta_t^f$. These margins define the uncertainty set underlying the robust optimization, ensuring worst-case constraint satisfaction in the PAO layer.

Objective Function: (16)

s.t.: *Battery Dynamics:* (2)–(4), (17)–(21)

Power Balance and Robust Grid Connection (22)

$$P_{\text{net},t} \le P_{\text{net}}^{\max} - \Delta_t^w, \quad \forall t \in \mathcal{T} \tag{27}$$

$$P_{\text{net},t} \ge -(P_{\text{net}}^{\max} - \Delta_t^w)\,, \quad \forall t \in \mathcal{T} \tag{28}$$

Shiftable Loads: (13), (11), (12)

$$C_1 = 0, \quad C_{t+1} = C_t + \hat{D}_t^{\text{fork}} + \Delta_t^f - pt_t^{\text{fork}}\,\Delta t, \quad \forall t \in \mathcal{T} \tag{29}$$

Auxiliary Decomposition: (24)

3.4 Ellipsoidal-RO Formulation

To capture temporal correlations, uncertainty is modelled via ellipsoids fitted to joint residual vectors. Let $r_n^w, r_n^f \in \mathbb{R}^T$ be the net-load and forklift *signed* residual vectors for calibration sample n. The sample covariance and its Cholesky factorization are:

$$\hat{\Sigma} = \frac{1}{N-1}\sum_n (r_n - \bar{r})(r_n - \bar{r})^\top = LL^\top, \tag{30}$$

and Mahalanobis scores are computed as:

$$s_n = \|L^{-1} r_n\|_2. \tag{31}$$

and used to set the conformal radius $\rho = Q_{1-\alpha}(\{s_n\}_{n=1}^N)$, ensuring at least $1-\alpha$ empirical coverage [8]. The resulting uncertainty sets are:

$$\begin{aligned}\mathcal{U}_w &= \{r \in \mathbb{R}^T : \|L_w^{-1} r\|_2 \le \rho_w\},\\ \mathcal{U}_f &= \{r \in \mathbb{R}^T : \|L_f^{-1} r\|_2 \le \rho_f\}.\end{aligned} \tag{32}$$

Robust counterparts are linearized via the ellipsoid support function $\max_{\|L^{-1}r\|_2 \le \rho} c^\top r = \rho \|L^\top c\|_2$. For the grid connection, the per-timestep tightening margin is:

$$\Delta_t^w = \rho_w \|L_w^\top e_t\|_2, \tag{33}$$

where e_t is the t-th standard basis vector. For the forklift, the worst-case cumulative deviation is the scalar:

$$\Delta^f = \rho_f \|L_f^\top \mathbf{1}\|_2, \tag{34}$$

distributed uniformly as Δ^f/T per step in the backlog dynamics. These margins augment the PAO layer, as follows:

Objective Function: (16)
s.t.: *Battery Dynamics:* (2)–(4), (17)–(21)
Power Balance and Robust Grid Connection: (22)

$$P_{\text{net},t} \le P_{\text{net}}^{\max} - \Delta_t^w, \quad \forall t \in \mathcal{T} \tag{35}$$

$$P_{\text{net},t} \ge -(P_{\text{net}}^{\max} - \Delta_t^w), \quad \forall t \in \mathcal{T} \tag{36}$$

Shiftable Loads: (13), (11), (12)

$$C_1 = 0, \quad C_{t+1} = C_t + \hat{D}_t^{\text{fork}} + \frac{\Delta^f}{T} - pt_t^{\text{fork}}\,\Delta t, \quad \forall t \in \mathcal{T} \tag{37}$$

Auxiliary Decomposition: (24)

3.5 CL Mechanism

The CL approach treats deployment as a sequence of monthly tasks. Let $\mathcal{D}_k = \{(\psi_n, \tilde{y}_n)\}_{n=1}^{N_k}$ denote the set of input–output day-pairs observed during month k, where ψ_n is the feature vector for day n and $\tilde{y}_n$ the corresponding ground-truth multi-series target. At the beginning of each month, the modelling approach uses the previous checkpoint to forecast and make EMS decisions; after the month is observed, the new samples update the weights, and the resulting checkpoint is stored.

3.5.1 Memory Buffer and FIFO Update

The core challenge in online adaptation is the stability-plasticity trade-off [2], i.e., incorporating new data aggressively improves tracking of current conditions but risks overwriting previously learned behavior, while overly conservative updates fail to capture genuine distributional changes. Following the rehearsal paradigm of [1,2], catastrophic forgetting is mitigated by maintaining a fixed-capacity replay buffer $\mathcal{B}$ of maximum size M_{buf} (in daily sample units). The buffer is operated as a first-in, first-out (FIFO) ring: at the end of month k, the incoming batch $\mathcal{D}_k$ is appended, and the oldest entries are evicted if capacity is exceeded:

$$\mathcal{B}^{(k)} = \text{FIFO-INSERT}\Big(\mathcal{B}^{(k-1)},\ \mathcal{D}_k,\ M_{buf}\Big), \tag{38}$$

where FIFO-INSERT appends all N_k new samples and removes the $\max\big(0,\ |\mathcal{B}^{(k-1)}| + N_k - M_{buf}\big)$ oldest entries, ensuring $|\mathcal{B}^{(k)}| \le M_{buf}$ at all times. The FIFO strategy preserves ordered temporal coverage across all prior months, guaranteeing that the buffer always spans the most recent complete history window [2]. During each weight update, a mini-batch of B_{mem} past day-pairs is drawn uniformly at random from the current buffer:

$$\mathcal{B}_{\text{mini}} \sim \text{Uniform}\Big(\mathcal{B}^{(k-1)},\ B_{mem}\Big), \tag{39}$$

and in addition to the current monthly batch $\mathcal{D}_k$ they form two separate training sets for the gradient step. Balancing α_{cur} and α_{mem}, thus directly controls the plasticity-stability trade-off [2].

3.5.2 PTO CL

In PTO CL, CL is applied to the forecasting models only, since the MILP optimization layer is decoupled from training and is non-differentiable. At month k, the model parameters θ are updated by minimizing a replay-augmented Huber loss [2]:

$$\mathcal{L}_{PTO}^{CL}(\theta) = \alpha_{\text{cur}}\, L_{\text{Huber}}^{\text{cur}}(\theta) + \alpha_{\text{mem}}\, L_{\text{Huber}}^{\text{mem}}(\theta), \tag{40}$$

where the current and memory Huber terms are:

$$L_{\text{Huber}}^{\text{cur}}(\theta) = \frac{1}{N_k} \sum_{n=1}^{N_k} \ell_H(\hat{y}_n(\theta),\, \tilde{y}_n), \tag{41}$$

$$L_{\text{Huber}}^{\text{mem}}(\theta) = \frac{1}{B_{mem}} \sum_{(\psi_b,\, \tilde{y}_b) \in \mathcal{B}_{\text{mini}}} \ell_H(\hat{y}_b(\theta),\, \tilde{y}_b), \tag{42}$$

and ℓ_H denotes the per-sample Huber loss. Operational improvement comes indirectly, since the updated forecasts are passed to the MILP only at evaluation time.

3.5.3 PAO CL

PAO CL follows the same monthly task-and-replay structure, but the update is decision-focused, aligned with the PAO pipeline in Section 3.2. Forecasts are passed through the differentiable PAO layer, the resulting schedule is evaluated against realized trajectories, and the total loss combines prediction error with realized EMS cost [1]:

$$\mathcal{L}_{PAO}^{CL}(\theta) = \alpha_{\text{cur}}\, L_{\text{PAO}}^{\text{cur}}(\theta) + \alpha_{\text{mem}}\, L_{\text{PAO}}^{\text{mem}}(\theta), \tag{43}$$

Unlike PTO CL, the NN weights are updated to optimized the downstream effect on EMS decisions, by minimizing eq. (16).

4 Use case and evaluation metrics

The experimental study considers an industrial EMS with a behind-the-meter BESS, PV generation, base demand, forklift charging demand, and non-critical L2/L3 EV charger loads, operating at 15-minute resolution over a daily scheduling horizon of $T = 96$ intervals. The dataset is characterized by an average base demand of 682.99 kWh ($\sigma = 238.97$ kWh), PV system has a maximum capacity of 2.4 MW, with an average generation of 384.53 kW ($\sigma = 573.47$ kW), and a forklift demand of 79.84 kWh on average ($\sigma = 30.30$ kWh) excluding zeros. The BESS has a nominal capacity of 932 kWh and a 400 kW charge/discharge power limit, while the aggregated forklift charging power is limited to 111 kW. Electricity prices correspond to the local provider tariffs over the same time period as the operational data. The battery cycling cost is derived from the mathematical representation of battery degradation and is applied as a charge/discharge throughput penalty in the EMS objective. Data are split chronologically to reflect a real deployment setting. The first eight months serve as the offline training and calibration period, whereas the remaining 28 months form the test stream. No

Table 3: Aggregated evaluation metrics by period and experiment.

Period	Experiment	Demand MAE (kWh)	PV MAE (kW)	Forklift MAE (kWh)	Opt. Gap (€)	Electricity Cost (€)	Daily unsat. Forklift Energy Penalty (€, $K = 1$)	Daily Unsat. Non-shiftable Loads (€, $K = 0.2$)	Total Cost with Unsat. Loads (€)
Non-drift	PTO Baseline	69.02	124.15	3.13	150.24	6,651.07	65.01	192.14	6,908.22
	PAO Baseline	110.43	97.78	3.02	152.26	6,656.52	63.83	199.56	6,919.91
	PTO CL	56.70	103.85	3.50	147.47	6,651.63	61.73	192.09	6,905.45
	PAO CL, $\alpha = (0.3, 0.7)$	109.03	86.57	3.52	139.10	6,639.23	48.17	193.65	6,881.04
	Box-RO	112.05	98.47	2.97	124.45	6,650.75	24.43	191.50	6,866.68
	Ellipsoidal-RO	112.05	98.47	2.97	**111.61**	6,657.92	4.65	191.50	6,854.07
Forklift drift	PTO Baseline	63.58	102.40	21.03	1,767.68	6,211.00	1,878.98	150.99	8,240.97
	PAO Baseline	105.61	101.69	18.98	1,519.50	6,233.24	1,590.33	156.94	7,980.52
	PTO CL	65.08	92.02	6.89	295.62	6,375.16	242.78	150.96	6,768.91
	PAO CL, $\alpha = (0.3, 0.7)$	103.71	49.06	8.65	290.89	6,367.89	230.63	149.97	6,748.48
	Box-RO	106.66	102.75	19.21	821.66	6,298.20	830.37	150.54	7,279.11
	Ellipsoidal-RO	106.66	102.75	19.21	**119.90**	6,424.80	2.01	150.63	6,577.44
Demand drift	PTO Baseline	151.95	99.40	3.24	121.34	4,955.17	45.95	152.29	5,153.42
	PAO Baseline	173.98	114.38	3.26	124.08	4,633.90	48.12	149.03	4,831.04
	PTO CL	83.86	86.41	3.39	114.21	4,954.50	39.53	152.25	5,146.28
	PAO CL, $\alpha = (0.3, 0.7)$	112.88	33.71	3.57	140.19	4,629.45	69.25	143.30	4,842.00
	Box-RO	166.74	110.97	3.30	**99.33**	4,956.85	6.25	152.50	5,115.59
	Ellipsoidal-RO	166.74	110.97	3.30	100.40	4,964.25	0.05	152.56	5,116.86
Total	PTO Baseline	96.37	105.95	10.86	834.71	5,854.92	837.43	160.20	6,852.55
	PAO Baseline	128.61	104.88	10.35	764.86	5,815.83	752.83	163.99	6,732.65
	PTO CL	70.02	92.53	4.92	199.21	5,925.32	131.56	160.16	6,217.05
	PAO CL, $\alpha = (0.3, 0.7)$	107.84	52.60	5.87	208.54	5,873.28	137.97	157.67	6,168.91
	Box-RO	129.31	104.78	10.06	414.96	5,892.92	364.13	159.94	6,416.99
	Ellipsoidal-RO	129.31	104.78	10.06	**111.16**	5,951.48	1.87	160.00	6,113.35

random splitting is applied, since the objective is to evaluate behavior under a temporal distribution shift. Data drift is assessed using a multivariate Maximum Mean Discrepancy (MMD) detector with a permutation test, where each day is represented as a full intra-day profile and each test month is compared against the training reference distribution at significance level $\alpha_{\text{MMD}} = 0.001$ [9]. Two main drift regimes are identified: a forklift-drift period throughout the second year, linked to a change in charging behaviour, during which the mean monthly MMD p-value reached a value of ≈ 0.00099, representing a 97.0% decrease relative to the remaining test month (≈ 0.0332) and a demand-drift period from mid-third year onward, linked to a shift in base-demand patterns, where the mean p-value likewise reached ≈ 0.00099, corresponding to a 78.1% decrease outside the drift period (≈ 0.0045). Both drift periods are highlighted in Figure 1. Results are reported across four evaluation groups: non-drift, forklift drift, demand drift, and the full 28-month test set. Forecasting accuracy is assessed using per-series MAE for base demand, PV, and forklift demand, while operational performance is evaluated using five daily mean KPIs, namely Optimality Gap (forecast-based C_{ops} minus oracle solve, see Sec. 3.2) , Electricity Cost, Daily Unsatisfied Forklift Energy ($K = 1$€), Unsatisfied Non-critical L2/L3 Loads ($K = 0.2$€), and Total Cost with Unsatisfied Loads.

5 Results

The aggregated results are reported in Table 3, while Figure 1 shows the corresponding monthly evolution of the optimality gap and the total cost with unsatisfied-load penalties. During the non-drift period, all methods perform similarly, with PAO CL, $\delta_{\text{demand}} = 100$, $\delta_{\text{PV}} = 80$, $\delta_{\text{forklift}} = 8$, $\alpha = (0.3, 0.7)$ achieving the largest gap reduction (8.6% relative to the PAO Baseline). This confirms that CL incurs no performance penalty under stable conditions, though its benefit is limited without distributional shift. The forklift-drift period provides the clearest evidence for the value of CL. Both static baselines show a substantial increase in the optimality gap, driven primarily by their failure to satisfy the shifted forklift energy demand. PTO CL reduces the optimality gap by 83.3% relative to PTO Baseline, while PAO CL $\alpha = (0.3, 0.7)$ achieves a reduction of 80.9% relative to PAO Baseline. This improvement is supported by a notably lower Cumulative Unsatisfied Load term for the CL methods and is also visible in Figure 1, where the baseline gap increases sharply during the forklift-drift window.

In the demand-drift period, the relationship between forecasting accuracy and decision quality is less straightforward. PTO CL improves upon the PTO Baseline in terms of optimality gap; however, PAO CL becomes the worst-performing method in this period, with an optimality gap of 140.19, which is also higher than its own PAO Baseline value of 124.08. This occurs despite strong forecasting performance, including the lowest PV MAE in Table 3. This is an important finding of the study: decision-aware CL does not automatically improve downstream decisions under every drift regime, and can destabilize EMS decisions when the drift is not aligned with the learned decision-focused objective. This highlights a limitation of MAE-based evaluation, since average prediction error alone is insufficient to rank EMS configurations, as forecast errors have the greatest impact only when they affect binding constraints, critical loads, or high-cost decision intervals.

Over the full test set, CL substantially improves the performance of both PTO and PAO frameworks. PTO CL reduces the total optimality gap by 76.1% relative to PTO Baseline, while PAO CL, $\alpha = (0.3, 0.7)$ achieves a reduction of 72.7% relative to PAO Baseline. However, the aggregate difference between PTO CL and PAO CL is small, and the demand-drift results show that decision-focused training should not be treated as universally preferable. Total cost rankings are broadly consistent, with minor differences attributable to the interplay of electricity prices and load-violation penalties, as also shown in the monthly cost trajectories of Figure 1.

The robust optimization approaches and Ellipsoidal-RO, in particular, achieve the lowest optimality gaps across the non-

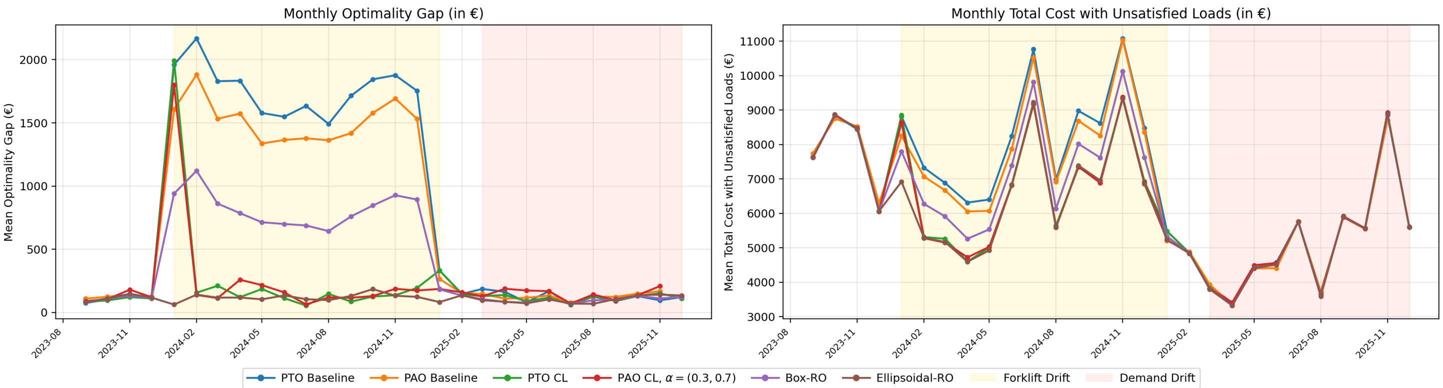

Figure 1: Monthly Optimality Gap and Total Cost with Unsatisfied loads across the evaluated EMS configurations.

drift, forklift-drift, and aggregate evaluation periods. Over the full test set, Ellipsoidal-RO reduces the optimality gap by 86.7% and 85.5% relative to the PTO and PAO baselines, respectively, driven largely by the near-complete elimination of unsatisfied forklift energy. The difference between RO-Box and Ellipsoidal-RO lies solely in the uncertainty set definitions (3.3 and 3.4). While both assume homoscedastic and stationary uncertainty, the ellipsoidal set provides a more descriptive representation of uncertainty. It is important, however, to contextualize these findings within the system's operational constraints. In this setting, the Ellipsoidal-RO prioritises energy sufficiency for forklifts by hedging against worst-case forklift demand, thus inherently sacrificing some arbitrage headroom.

The cost remains marginal because the forklift load represents only 7-10% of base demand and the high base load limits arbitrage potential regardless. Consequently, while the RO approach is inherently conservative, the specific characteristics of this EMS environment allow for enhanced reliability without incurring significant economic penalties during non-drifted periods. Taken together, Table 3 and Figure 1 show that CL improves adaptability to distributional drift, robust optimization improves reliability with respect to the primary decision-quality metric, and decision-focused CL may become unstable when the deployed drift regime differs from the conditions emphasized by its training objective.

6 Conclusion

This study evaluated six EMS modelling approach configurations, namely two static baselines (PTO, PAO), two continual-learning variants (PTO CL, PAO CL), and two robust optimisation approaches (Box-RO, Ellipsoidal-RO), under structured data drift in an industrial facility with behind-the-meter battery storage, PV generation, forklift charging, and non-critical EV loads. The main finding is that CL substantially improves the quality of operational decisions when drift affects critical loads. During the forklift-drift period, PTO CL and PAO CL reduced the optimality gap by 83% and 81%, respectively, relative to their static baselines, and by 76% and 73% over the full 28-month test horizon. Ellipsoidal-RO achieved the best aggregate performance, reducing the optimality gap by 87% relative to the PTO Baseline by nearly completely eliminating Cumulative Unsatisfied Loads, without requiring model updates. PAO CL provided only marginal additional benefit over PTO CL, suggesting that the dominant gain under critical-load drift comes mainly from improved forecasting of the drifted series rather than from the decision-focused training objective. The results also confirm that the MAE metric is an unreliable proxy for EMS performance, reinforcing the need to use downstream decision-quality metrics as the primary evaluation criterion.

Regarding limitations, the experiments were carried out for a single industrial site, so the observed ranking of CL and RO methods may change under different battery-to-load ratios, tariffs, load compositions, or arbitrage opportunities. The MMD drift detector operates at monthly granularity with a very strict, fixed significance threshold, which may miss subtle distribution drifts. In addition, the RO uncertainty margins are calibrated once from the offline residual distribution and remain fixed during deployment, even though residual statistics may drift over time. Finally, the sensitivity of the CL methods to buffer size, replay mini-batch size, and plasticity-stability weights was explored only for selected configurations. Future work should extend the work to additional sites, adaptive or drift-triggered update schedules, dynamic RO recalibration, and broader ablation studies across drift regimes.

References

[1] M. Wu, X. Zhou, S. Li, and H. Shi, "An adaptive continual learning method for nonstationary industrial time series prediction," *IEEE Transactions on Industrial Informatics*, vol. 21, no. 2, pp. 1160–1169, 2025.

[2] T. Zhuo, Z. Cheng, Z. Gao, and M. Kankanhalli, "Continual learning with strong experience replay," *ArXiv*, vol. abs/2305.13622, 2023.

[3] D. Rolnick, A. Ahuja, J. Schwarz, T. P. Lillicrap, and G. Wayne, "Experience replay for continual learning," *CoRR*, vol. abs/1811.11682, 2018.

[4] A. N. Elmachtoub and P. Grigas, "Smart "predict, then optimize"," 2020.

[5] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust Optimization*. Princeton University Press, 2009.

[6] A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and Z. Kolter, "Differentiable convex optimization layers," 2019.

[7] L. Tziovani, L. Hadjidemetriou, and S. Timotheou, "Successive convexification algorithms for optimizing power systems with energy storage models," *IEEE Transactions on Smart Grid*, vol. 15, no. 2, pp. 1807–1820, 2024.

[8] A. N. Angelopoulos and S. Bates, "A gentle introduction to conformal prediction and distribution-free uncertainty quantification," 2022.

[9] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," *Journal of Machine Learning Research*, vol. 13, no. 25, pp. 723–773, 2012.