

Fault-Aware AI-Based Load Scheduling for Industrial Facilities Using Machine Learning–Driven Demand-Side Management

FOUAD ZARO

Electrical Engineering Department
Palestine Polytechnic University
Hebron City
PALESTINE

Abstract: - Industrial demandside management (DSM) has become a key enabler for reducing electricity costs, mitigating peak demand, and improving renewable energy utilization. However, most existing DSM frameworks assume ideal operating conditions and neglect the impact of electrical disturbances, such as voltage deviations and supply outages, which are common in real industrial power systems. This paper proposes a faultaware artificial intelligence–based load scheduling framework for industrial facilities that explicitly integrates electrical disturbance states into the DSM decisionmaking process.

A comprehensive industrial case study is developed using a realistic synthetic dataset comprising 61,488 hourly records for seven industrial loads, an 80-kW photovoltaic system, dynamic electricity pricing, and modeled disturbance conditions, including voltage drops and outages representing a 15% annual disturbance rate consistent with industrial reliability benchmarks. A Random Forest classification model is employed to predict feasible operating states under combined temporal, economic, solar, and disturbance conditions. The model outputs are incorporated within a constraintaware scheduling strategy that preserves production requirements while avoiding infeasible operation during fault conditions.

Simulation results demonstrate that the proposed faultaware DSM framework maintains the economic benefits of intelligent scheduling, achieving 11.7% annual electricity cost reduction and more than 95% peakhour operation reduction, while significantly enhancing operational robustness under nonideal grid conditions. The findings confirm that integrating disturbance awareness into AIbased DSM is essential for bridging the gap between theoretical optimization and practical industrial deployment.

Key-Words: - Industrial demand- side management; fault- aware scheduling; machine learning; Random Forest; electrical disturbances; voltage deviations; industrial load scheduling; solar energy integration; smart grids; operational reliability.

Received: March 13, 2026. Revised: May 17, 2026. Accepted: June 19, 2026. Published: September 29, 2026.

1 Introduction

Industrial facilities represent a major share of global electricity consumption and significantly influence power system operating conditions. Increasing electricity prices, peak demand charges, and the growing penetration of renewable energy sources have intensified the need for more intelligent industrial energy management strategies. Conventional load scheduling is typically based on fixed production plans and rarely considers dynamic electricity pricing, renewable availability, or grid operating conditions, leading to unnecessary costs and inefficient peak- hour operation.

Demand- side management (DSM) enables industrial facilities to reshape electricity consumption without compromising production output by shifting flexible loads toward economically and operationally favorable periods. Recent advances in artificial

intelligence and machine learning have further enhanced DSM by enabling data- driven scheduling strategies capable of capturing complex temporal, economic, and environmental relationships. However, most machine learning–based DSM approaches assume ideal grid conditions and neglect the impact of electrical disturbances.

In practice, industrial power systems are frequently affected by voltage deviations and unplanned outages, with reported annual disturbance rates ranging from 10% to 20%. These disturbances directly affect load operability and production continuity, and ignoring them can lead to infeasible schedules and overestimated economic benefits.

This paper proposes a fault- aware AI- based DSM framework for industrial load scheduling that explicitly integrates electrical disturbance states into the scheduling decision process. The framework

combines a Random Forest classifier with a constraint-aware scheduling strategy, enabling reliable and economically efficient operation under realistic grid conditions. The main contributions include: (i) integration of disturbance awareness into machine learning-based DSM, (ii) development of a realistic industrial dataset with modeled disturbances, and (iii) demonstration of improved operational robustness without sacrificing economic performance.

2 Related Work

2.1 Demand- Side Management in Industrial Energy Systems

Demand- side management (DSM) has been widely investigated as a key solution for improving energy efficiency, reducing electricity costs, and mitigating peak demand in power systems. In industrial environments, DSM is particularly impactful due to the high-power ratings of industrial loads and their substantial contribution to system peak demand [1].

Traditional industrial DSM strategies focus on load shifting, peak clipping, and valley filling, usually implemented through deterministic or rule- based scheduling methods driven strictly by production requirements [2].

Early industrial DSM research relied heavily on mathematical optimization techniques, including linear programming and mixed- integer programming, to minimize electricity costs subject to operational constraints [3]. While these approaches can yield optimal solutions under ideal assumptions, they often suffer from high computational complexity, limited scalability, and weak adaptability to dynamic operating conditions [4]. More importantly, most classical DSM formulations assume uninterrupted grid supply and neglect the impact of electrical disturbances common in industrial power systems [5].

Comprehensive reviews of industrial DSM highlight that, despite decades of research, most implementations remain economically driven and insufficiently responsive to non- ideal grid behavior, particularly under increasing penetration of renewable energy and distributed generation [6].

2.2 Machine Learning Applications in Industrial DSM

Recent advances in artificial intelligence and machine learning have significantly transformed DSM research by enabling predictive and adaptive energy management frameworks. Machine learning techniques have been applied to industrial DSM for load forecasting, load classification, operating state

prediction, and automated scheduling under dynamic electricity pricing [7].

Supervised learning models, including Artificial Neural Networks, Support Vector Machines, and ensemble algorithms, have demonstrated strong performance in capturing nonlinear relationships between temporal factors, pricing signals, and load behavior [8]. Among these methods, Random Forest algorithms have gained particular attention due to their robustness, low sensitivity to noisy data, ability to handle structured datasets, and limited tuning requirements [9], [10]. These properties make Random Forest especially suitable for industrial DSM applications, where data often exhibit heterogeneity and operational uncertainty.

Several studies have reported significant electricity cost reductions and peak-hour mitigation using machine learning-assisted DSM frameworks, particularly when prediction models are coupled with heuristic or greedy scheduling strategies [11]. However, most machine learning-based DSM studies implicitly assume that predicted operating hours are always feasible and do not account for grid reliability issues, such as voltage deviations or outages, within the learning or scheduling process [12].

2.3 Renewable Energy Integration in Industrial Load Scheduling

The increasing deployment of on- site renewable energy systems, especially photovoltaic generation, has further increased the complexity of industrial load scheduling. DSM strategies that align load operation with renewable energy availability can improve self- consumption, reduce grid dependence, and enhance overall sustainability [13].

Numerous studies have demonstrated that solar- aware scheduling can significantly shift industrial operations toward daylight hours, thereby increasing renewable utilization and lowering electricity costs [14]. Machine learning models have been increasingly employed to manage renewable intermittency and integrate solar generation forecasts into DSM decision- making [15].

Despite these advances, most renewable- oriented DSM approaches focus exclusively on generation variability while assuming stable grid conditions. The interaction between renewable integration and grid power-quality disturbances, such as voltage fluctuations intensified by high distributed generation penetration, remains largely unexplored in industrial DSM literature [16].

2.4 Electrical Disturbances and Reliability in Industrial Power Systems

Electrical disturbances, including voltage sags, voltage swells, and short-duration outages, are inherent characteristics of industrial power systems. Power-quality studies indicate that voltage deviations represent the most frequent disturbance affecting industrial facilities and are responsible for a substantial portion of unplanned production interruptions and financial losses [17].

Industrial reliability analyses consistently report annual disturbance rates ranging from 10% to 20%, particularly in facilities connected to weak distribution networks or subjected to frequent motor starts and switching operations [18]. Voltage deviations directly influence equipment performance, power consumption characteristics, and process continuity, while outages result in complete loss of operability [19].

Although extensive research exists on fault detection, classification, and mitigation of electrical disturbances, these models are rarely integrated into DSM or industrial load scheduling frameworks [20]. Reliability-aware energy management approaches have been proposed using stochastic or robust optimization; however, these methods often require complex formulations and computational resources that limit their applicability in real-time industrial DSM [21]. Explicit integration of disturbance state information into machine learning-based DSM remains largely absent from existing studies.

2.5 Research Gap and Positioning of This Work

The reviewed literature reveals that machine learning-based DSM frameworks have achieved meaningful success in reducing electricity costs, mitigating peak demand, and improving renewable energy utilization in industrial facilities. At the same time, electrical disturbances are well-recognized as unavoidable and impactful phenomena in industrial power systems.

However, existing DSM studies rarely integrate awareness of electrical disturbances into scheduling decisions, instead assuming ideal, uninterrupted operation. This disconnect limits the practical deployment of intelligent DSM solutions in real industrial environments.

This paper addresses this gap by proposing a fault-aware AI-based DSM framework that explicitly incorporates electrical disturbance states into machine learning-assisted load scheduling. By conditioning scheduling decisions on disturbance information in addition to temporal, economic, and renewable factors, the proposed approach advances industrial

DSM toward greater robustness, reliability, and real-world applicability.

3 Industrial System and Disturbance Modeling

3.1 Industrial Facility Description

The case study considered in this work represents a medium-scale industrial facility operating under realistic production and electrical constraints. The facility consists of seven aggregated industrial loads supplied through a low-voltage, three-phase power system. Each load represents a production unit characterized by its rated power and required daily operating duration rather than a specific machine type, allowing for a generalized and scalable modeling approach.

The total installed electrical capacity of the facility is 195 kW, with individual load ratings ranging from 15 kW to 55 kW. The plant operates on a two-shift schedule—Shift 1 from 08:00 to 16:00 and Shift 2 from 16:00 to 24:00—over six working days per week, with one non-working day reserved for planned shutdown. Once a load is started, continuous operation for the required duration is enforced to reflect industrial process constraints.

This abstraction captures the essential characteristics of industrial energy consumption while avoiding equipment-specific dependencies, enabling the evaluation of DSM strategies under generalized industrial operating conditions.

Figure 1 illustrates the architecture of the industrial DSM framework, highlighting the interaction between industrial loads, photovoltaic generation, electricity pricing, disturbance modeling, and the AI-based scheduling engine.




Fig. 1: Block diagram of the fault-aware AI-based demand-side management framework for industrial facilities.

"Source: created by the author"

3.2 Renewable Energy Integration

The industrial facility is equipped with an on-site 80 kW grid-connected photovoltaic (PV) system operating in parallel with the utility supply. The PV system provides daytime renewable generation that can be directly utilized by the industrial loads through self-consumption. No energy storage system is assumed in order to isolate the impact of scheduling decisions and disturbance awareness.

Solar generation is modeled using a bell-shaped daily profile during daylight hours (06:00–18:00), with seasonal variations reflecting realistic irradiance conditions. Maximum PV output ranges from approximately 45 kW in winter to 80 kW in summer, consistent with typical rooftop PV installations. The availability of solar generation introduces an additional scheduling incentive but does not alter production requirements or total energy consumption.

Seasonal variability in photovoltaic generation, shown in Figure 2, introduces additional uncertainty that must be considered alongside electrical disturbances when scheduling industrial loads.




Fig. 2: Seasonal photovoltaic generation profiles for the industrial facility.

“Source: created by the author”

Figure 3 illustrates the temporal relationship between photovoltaic generation and industrial energy consumption, highlighting the importance of intelligent scheduling under both renewable variability and grid disturbances.




Fig. 3: Hourly photovoltaic generation and industrial energy consumption profiles.

“Source: created by the author”

3.3 Electrical Disturbance Characterization

Unlike conventional DSM studies that assume ideal electrical conditions, this work explicitly models electrical disturbances that affect industrial load operability. Based on industrial reliability literature

and practical power-quality studies, three operating states are considered:

1. Normal Operation: Voltage remains within acceptable limits ($\pm 10\%$ of nominal), and loads consume rated power without restrictions.

2. Voltage Deviation (Voltage Drop or Rise): Supply voltage deviates outside nominal limits but does not result in a complete interruption. Under such conditions, power consumption is reduced according to the voltage–power relationship of constant-impedance industrial loads:

$$P_{actual} = \left(\frac{V_{actual}}{V_{nominal}} \right)^2 P_{rated} \quad (1)$$

This condition represents degraded but potentially usable operating states.

3. Outage: A complete loss of supply voltage leads to zero power consumption and load unavailability.

The disturbance occurrence rates are selected to reflect realistic industrial conditions: 85% normal operation, 9.1% voltage deviation, and 5.9% outage, yielding a total disturbance rate of 15%, consistent with reported industrial reliability benchmarks. These disturbance states directly influence scheduling feasibility and load operability.

The adopted 15% figure is intended as a representative mid-range value rather than a measurement from a single facility. Benchmarking studies of industrial and commercial power systems report disturbance frequencies spanning roughly an order of magnitude depending on feeder type, the urban-rural mix of the supplying network, and the voltage threshold used to define an event [17], [22]. For the regional context of the case-study facility, independent assessments of the West Bank distribution network similarly document frequent voltage instability and unplanned interruptions, attributed to feeder loading, cross-border supply constraints, and limited standardization of protection practice [23]. The 15% rate used in this study therefore sits toward the upper portion of the commonly cited range, consistent with a facility connected to a comparatively weak, non-ideal distribution feeder rather than a well-regulated urban network; it is not derived from measurements at a specific plant and is not intended to represent a universal industry-wide average.

The modeled disturbance probabilities and voltage operating ranges are summarized in Table 1 and Table 2, reflecting realistic industrial reliability conditions.

Table 1. Electrical disturbance types and assumed occurrence probabilities used in the industrial system model

Event	Voltage (V)	Power (kw)	Probability
Normal	380-420	P_{rated}	85%
Sag	320-379	$(V/400)^2 \times P_{rated}$	9.1%
Swell	421-460		
Outage	0	0	5.9%

“Source: created by the author”

The distribution of voltage levels across normal operation, voltage deviation, and outage states is presented in Table 3, reflecting realistic industrial power-quality conditions.

Table 2. Summary of annual disturbance occurrence statistics during operating hours of the industrial facility

Disturbances Type	Count	Percentage
Normal	35,529	85.0%
Outage	2465	5.9%
Voltage Drop	3,788	9.1%
Total Disturbances	6,253	15.0%
Off (Scheduled)	19,538	-

“Source: created by the author”

Table 3. Voltage range classification and distribution under normal and disturbed operating conditions

Voltage Range	Count	Percentage	Classification
Normal	35,529	85.0%	Acceptable
Under-voltage	1,894	4.5%	Disturbances
Over-voltage	1,894	4.5%	Disturbances
Zero-voltage	2,465	5.9%	Circuit failure
Total	41,782	100%	-

“Source: created by the author”

3.4 Disturbance-Aware Load Operability Constraints

Electrical disturbances impose additional constraints on industrial load scheduling beyond traditional DSM formulations. During outage conditions, loads are strictly unavailable and must not be scheduled. During voltage deviation states, load operation may be technically feasible but economically or operationally undesirable, particularly for high-power equipment sensitive to reduced voltage levels.

To capture these effects, disturbance states are treated as contextual conditions that influence both machine learning predictions and scheduling decisions. This modeling ensures that scheduling outputs remain physically feasible and reflective of non-ideal grid conditions.

3.5 Dataset Structure and Feature Representation

A realistic synthetic dataset is constructed to support fault-aware DSM analysis. The dataset comprises 61,488 hourly records, covering one full year of operation for all seven loads. Each record integrates temporal, operational, economic, renewable, and disturbance-related information.

Key feature categories include:

1. Temporal features: hour of day, day type, shift indicator.
2. Economic features: dynamic electricity price, peak-hour flag.
3. Renewable features: photovoltaic power availability.
4. Operational features: load identity and required operating duration.
5. Disturbance features: operating state (normal, voltage deviation, outage).

Electrical measurements that directly reveal load operation, such as power, voltage, or energy cost, are excluded from predictive features in order to prevent target leakage. This dataset structure supports supervised machine learning while preserving generalization capability.

3.6 Role of Disturbance Modeling in Fault-Aware DSM

The incorporation of disturbance modeling fundamentally distinguishes the proposed DSM framework from fault-unaware approaches. Rather than assuming all scheduled hours are executable, the proposed system explicitly evaluates scheduling decisions under non-ideal grid conditions. This allows the DSM framework to:

1. Avoid infeasible schedules during outage periods,
2. Reduce operational risk under voltage-degraded conditions,
3. Preserve production continuity despite realistic disturbance rates.

By embedding disturbance awareness directly into the modeling layer, the proposed approach bridges the gap between theoretical DSM optimization and practical industrial deployment.

3.7 Sensitivity of Fault-Aware Benefits to the Disturbance Rate

Because outage and voltage-deviation states enter the framework through structurally different mechanisms, their effect on the economic and operational benefit of fault awareness is not expected to be identical, nor is that benefit expected to scale linearly with the total disturbance rate.

Outage hours impose a hard exclusion: Section 4 forces predicted feasibility to zero for every load whenever an outage occurs, so each additional percentage point of outage probability removes a proportional share of hours from the feasible search space used to build operating blocks in Section 5.1. While spare capacity remains within a shift, the fault-unaware baseline is penalized, through infeasible or interrupted operation, roughly in proportion to the outage share, and the benefit of fault awareness should grow accordingly. As the outage share continues to rise, however, the number of feasible consecutive blocks available within an eight-hour shift shrinks combinatorially rather than linearly, particularly for loads with longer operating durations; beyond some facility- and load-specific point, remaining feasible windows become scarce enough that even the fault-aware scheduler is forced toward higher-cost blocks, so the marginal benefit of fault awareness should be expected to saturate or decline. This behavior is characteristic of constraint-satisfaction problems in which feasibility, rather than cost, becomes the binding limitation.

Voltage-deviation hours act differently. Per Equation (1), power scales with the square of the voltage ratio, so a deviation is a soft, continuously graded penalty rather than a hard exclusion, and, under the load-category rules of Section 5.2, only heavy loads are excluded outright. The relative benefit of fault awareness with respect to the voltage-deviation share is therefore expected to increase more smoothly, without the same sharp saturation region observed for outages.

Taken together, the total benefit of fault awareness is hypothesized to increase monotonically but sub-linearly with the total disturbance rate: driven mainly by the outage component at low-to-moderate rates, and progressively bounded by shift-level block availability at high rates. Confirming this hypothesis rigorously requires regenerating the dataset at several disturbance rates, for example 5%, 10%, 15%, 20%, and 25%, and re-running the full training and scheduling pipeline at each point. We identify this systematic sensitivity sweep as a concrete direction for future work rather than reporting it here, so as not to conflate an analytically motivated expectation with a new experimental result; a genetic-algorithm-based scheduler of the type discussed in Section 6.1 would be a natural tool for such a sweep, since it could help distinguish whether the saturation behavior described above is an artifact of the greedy heuristic's myopia or a structural property of the scheduling problem itself.

4 Fault-Aware Machine Learning Framework

The industrial load scheduling problem is formulated as a supervised classification task that predicts the feasibility of operating each load at each hour under combined temporal, economic, renewable, and disturbance conditions. The operational state $x_{l,t}$ is defined as in (2).

The dataset consists of 61,488 hourly records integrating temporal features, electricity price, photovoltaic availability, load identity, and disturbance state. Electrical measurements are excluded to avoid target leakage.

A Random Forest classifier is adopted due to its robustness and effectiveness for structured industrial datasets. Training uses stratified sampling and balanced class weighting, with performance evaluated using standard classification metrics.

The operating decision for each industrial load is formulated as a supervised classification problem in which the feasibility of operation is predicted under combined temporal, economic, renewable, and disturbance conditions.

Let $l \in L$ denote an industrial load and $t \in T$ denote an hourly time step. The operational state of load l at time t is modeled as a binary decision variable:

$$x_{l,t} = \begin{cases} 1, & \text{if load } l \text{ is scheduled to operate at hour } t, \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

For each load-hour pair (l,t) a feature vector is constructed as:

$$x_{l,t} = [x_{l,t}^{temp}, x_t^{econ}, x_t^{ren}, x_l^{load}, x_t^{dist}] \quad (3)$$

Where:

1. $x_{l,t}^{temp}$ represents temporal features (hour, day type, shift, season),
2. x_t^{econ} represents electricity price and peak-hour indicators,
3. x_t^{ren} represents photovoltaic availability,
4. x_l^{load} represents load identity and category,
5. x_t^{dist} represents the electrical disturbance state.

The disturbance state is encoded as a categorical variable: $d_t \in \{\text{Normal, Voltage Deviation, Outage}\}$.

The fault-aware classification function is defined as:

$$\widehat{x}_{l,t} = f_{RF}(X_{l,t}) \quad (4)$$

Where:

$\widehat{x}_{l,t} \in \{0,1\}$ is the predicted feasibility state,
 $f_{RF}(\cdot)$ represents the trained Random Forest classifier

The classifier outputs the posterior probability of feasible operation:

which is converted into a binary decision using a threshold τ :

$$\widehat{x}_{l,t} = \begin{cases} 1, & P_{l,t} < \tau \\ 0, & P_{l,t} \geq \tau \end{cases} \quad (6)$$

In this study, $\tau=0.5$ is selected as a balanced decision threshold to preserve recall for feasible operating hours.

To ensure physical feasibility, the classifier output is explicitly conditioned on the disturbance state:

1. Outage condition: $d_t = \text{Outage} \Rightarrow x_{l,t} = 0 \forall l$
ensuring that no load operates during supply interruption.

2. Voltage deviation condition

For voltage-degraded operation, feasibility depends on load category c_l :

$$d_t = \text{Outage} \Rightarrow \begin{cases} x_{l,t} = 0, & c_l \in \text{Heavy} \\ x_{l,t} \in \{0,1\}, & c_l \in \text{Medium, Light} \end{cases}$$

3. Normal operation: $d_t = \text{Normal} \Rightarrow x_{l,t} \in \{0,1\}$

These constraints ensure that the learned scheduling behavior remains consistent with industrial reliability and equipment sensitivity considerations.

The fault-aware classification outputs $\widehat{x}_{l,t}$ form a feasibility mask that restricts the scheduling search space. Final schedules are obtained by enforcing deterministic industrial constraints, including:

$$\sum_{t \in D_d} x_{l,t} = H_l \quad (7)$$

where H_l is the required daily operating duration of load l , and D_d represents the set of hours within day d .

Consecutive operation and shift restrictions are enforced at the scheduling stage, ensuring implementable daily schedules.

5 Fault-Aware Scheduling Framework and Baseline Definition

The proposed scheduling framework combines a fault-aware feasibility layer with a constraint-aware scheduling process. A fault-unaware baseline is defined in which loads are scheduled randomly within

shifts without considering price, renewable availability, or disturbances.

The fault-aware scheduling strategy restricts operation to hours predicted as feasible by the classifier and enforces industrial constraints such as daily operating duration, consecutive operation, and shift compliance. Loads are categorized by rated power to account for different sensitivities to peak pricing and voltage deviations.

Scheduling is implemented using a greedy algorithm that evaluates feasible consecutive

$$P_{l,t} = \mathbb{P}(x_{l,t} = 1 | X_{l,t}) \quad (5)$$

operating blocks and selects the lowest-cost option. This approach ensures operational feasibility, economic efficiency, and scalability.

Compared to the baseline, the proposed framework avoids infeasible schedules during outages, limits voltage-degraded operation for heavy loads, and significantly reduces peak-hour exposure.

5.1 Fault-Aware Scheduling Strategy

The proposed fault-aware scheduling strategy operates on the feasibility mask produced by the classification model. Let $\widehat{x}_{l,t}$ denote the predicted feasibility of operating load l at time t . Scheduling decisions are restricted to hours for which $\widehat{x}_{l,t}=1$, ensuring that outage periods are entirely excluded and voltage-degraded conditions are handled selectively.

The scheduling problem is subject to the following constraints:

1. **Daily operating duration**

$$\sum_{t \in D_d} x_{l,t} = H_l \quad \forall l, d, \quad (8)$$

where H_l is the required daily operating time for load l .

2. **Consecutive operation:** each load must operate in a single uninterrupted block per day to preserve production continuity.

3. Shift compliance S

Shift compliance

$$t \in \{08:00-16:00\} \cup \{16:00-24:00\}$$

4. Fault feasibility F

Fault feasibility

$$x_{l,t} \leq \widehat{x}_{l,t} \quad \forall l, t,$$

which guarantees that operation occurs only during feasible hours identified by the fault-aware classifier.

5.2 Load Categorization for Disturbance-Aware Scheduling

To reflect differences in sensitivity to grid disturbances and peak-hour operation, industrial loads are categorized based on rated power:

1. Heavy loads (high-power equipment),
2. Medium loads,
3. Light loads (low-power equipment).

Heavy loads are strictly prevented from operating during peak-price periods and voltage-degraded conditions due to their greater impact on cost, grid stress, and equipment reliability. Medium and light loads retain limited flexibility, allowing occasional operation under voltage deviations when production requirements necessitate it.

This categorization enables differentiated scheduling behavior without increasing computational complexity.

5.3 Fault-Aware Scheduling Objective

The scheduling decision aims to construct feasible daily schedules that minimize operational cost and risk while preserving production continuity. Although the problem is not solved as a single closed-form optimization, the effective objective can be expressed as:

$$\min \sum_{l \in L} \sum_{t \in D_d} (P_l \cdot \pi_t \cdot x_{l,t})$$

subject to the operational and fault-awareness constraints described above, where: P_l is the rated power of load l , π_t is the electricity price at hour t .

Periods with photovoltaic availability and non-peak pricing naturally emerge as preferred operating windows through the combined effect of the feasibility mask and scheduling constraints.

The aggregate daily load profile under conventional fault-unaware scheduling is shown in Figure 4, where substantial operation occurs during peak-price hours and without regard to grid disturbance conditions.




Fig. 4: Aggregate daily industrial load profile under baseline fault-unaware scheduling.

“Source: created by the author”

In contrast, Figure 5 shows the aggregate load profile produced by the proposed fault-aware DSM framework, where peak-hour operation is largely eliminated, and scheduling avoids disturbance-prone periods.

The impact of fault-aware scheduling on peak demand is further illustrated by the peak load duration curves in Figure 6, which show a substantial reduction in both peak magnitude and duration.

The redistribution of industrial load operation across the day is illustrated in Figure 7, where fault-aware DSM significantly suppresses operation during peak and disturbance-sensitive hours.




Fig. 5: Aggregate daily industrial load profile under the proposed fault-aware DSM framework.

“Source: created by the author”




Fig. 6: Peak load duration curves under baseline and fault-aware DSM scheduling.

“Source: created by the author”




Fig. 7: Hourly distribution of industrial load operation before and after fault-aware DSM implementation.

“Source: created by the author”

5.4 Formal Description and Pseudocode of the Greedy Scheduling Algorithm

This section formalizes the procedure of the greedy algorithm.

Block generation and evaluation. For a given day and load with a required operating duration and an allowed shift or shifts (Shift 1: 08:00-16:00; Shift 2: 16:00-24:00), the algorithm enumerates every candidate start hour such that the resulting block of consecutive hours lies entirely within a single shift. A block is retained as feasible only if every hour it contains satisfies the fault-aware feasibility mask and the load-category constraints of Section 5.2; under these constraints, heavy loads additionally exclude peak-price and voltage-deviation hours. Each feasible block is assigned a cost equal to the sum of its hourly costs (rated power multiplied by the electricity price, summed across the hours in the block), consistent with the objective in Section 5.3.

Selection and tie-breaking. Among the feasible blocks for a given load and day, the algorithm selects the one with minimum cost. Ties are resolved first by maximum overlap with photovoltaic-available hours (06:00-18:00), which favors self-consumption without altering the cost objective; if a tie remains, ties are resolved by earliest start time, which keeps the assignment deterministic and reproducible.

Processing order. Loads are processed sequentially rather than jointly, once per day, in decreasing order of rated power: heavy loads first, then medium, then light. Because heavy loads carry the tightest operability constraints (Section 5.2) and the largest cost and peak-demand impact, scheduling them first, while the greatest number of feasible windows is still available, reduces the risk of forcing a heavy load into an expensive or disturbance-exposed block later in the day. Medium and light loads are then scheduled against the disturbance mask and price signal independently of the blocks already assigned to other loads, since the cost objective of Section 5.3 is additively separable across loads and contains no shared-capacity or coincidence term.

Myopic behavior. The algorithm is deterministic and purely myopic: once a block is assigned to a load, it is not revisited in light of subsequent loads' scheduling outcomes, and there is no backtracking, local search, or joint re-optimization across loads or days. Each load's block is locally optimal given the fixed price, photovoltaic, and disturbance signals for that day, but the resulting multi-load daily schedule is a sequential composition of locally optimal decisions rather than a globally optimal one. This design trade-off favors low computational complexity and scalability over solution optimality; Section 6.1

discusses evolutionary alternatives that relax this myopia.

Fallback for infeasible days. In the rare event that no feasible block exists for a load within its assigned shift, for example a shift affected by an extended outage, the algorithm checks the alternate shift if the load's operating rules permit it, and otherwise leaves the load unscheduled for that day. This fallback was not triggered under the disturbance conditions modeled in this study (Table 2).

Algorithm 1 summarizes the procedure.

Algorithm 1: Fault-Aware Greedy Scheduling Procedure.

```

Input: L, loads ordered by power (desc.)
D, the set of days
feasibility mask (from Section 4)
price, PV hours, durations, shifts
Output: schedule for all loads and hours

for each day d in D:
  for each load l in L (heaviest first):
    candidates <- empty set
    for each shift s allowed for l:
      for each start hour t0 in s where
        block [t0 .. t0+duration-1]
        fits within s:
        if block hours are all feasible
        and category rules hold:
          cost <- sum(power x price)
          over block hours
          pv <- block hours that are
          also PV hours
          add (block, cost, pv, t0)
          to candidates
    if candidates not empty:
      pick block with min cost;
      tie-break: max pv, then
      earliest t0
      mark block hours as scheduled
    else if alternate shift allowed:
      retry search in that shift
    else:
      leave l unscheduled on day d
  return schedule

```

For a facility with a given number of loads scheduled over a given number of days, with at most a bounded number of candidate start hours per shift and blocks of bounded average duration, the algorithm performs work that scales linearly in the number of days, the number of loads, and the number of candidate blocks per load, since each shift is scanned once per load per day and each candidate block's cost is a sum over its own hours only. This low-order polynomial complexity is consistent with the scalability noted in the Conclusion, and is substantially lower than the combinatorial cost of a joint multi-load optimization over the same horizon.

6 Results and Discussion

The proposed framework is evaluated against the fault- unaware baseline under identical production requirements. Figure 4, Figure 5, Figure 6 and Figure 7 illustrate the resulting aggregate load profiles, peak load duration curves, and hourly operation distributions.

Baseline scheduling results in substantial operation during peak- price and disturbance- prone hours, as shown in Figure 4. In contrast, the fault- aware DSM framework shifts load operation toward economically favorable and disturbance- safe periods (Figure 5), while preserving total operating hours.

Peak demand and peak duration are significantly reduced under fault- aware scheduling (Figure 6), with heavy loads effectively removed from peak hours. Hourly operation distributions (Figure 7) confirm the suppression of peak- hour and disturbance- sensitive operation.

Importantly, disturbance awareness eliminates infeasible operation during outages and improves scheduling robustness without compromising cost savings. The results demonstrate that economic efficiency and operational reliability can be achieved simultaneously.

6.1 Extensions Toward Deep Learning and Evolutionary Optimization

The Random Forest classifier was selected for its robustness to noisy and heterogeneous industrial data, its low tuning burden, and, importantly for an operational deployment, its interpretable, rule-like decision boundaries, which plant engineers can audit against the physical constraints of Section 4 [9], [10].

Long Short-Term Memory networks and Transformer-based architectures are well suited to sequential, hourly time series and have demonstrated strong performance in industrial and grid-scale load forecasting, including settings with contextual, non-load covariates like those used here [24], [25]. In principle, such models could capture longer-range temporal structure, for instance interactions between multi-day price patterns, seasonal photovoltaic trends, and clustering of disturbance events, that the present hour-by-hour, feature-engineered formulation does not explicitly represent. Realizing this benefit is conditional on the underlying data: the present dataset, while comprising 61,488 hourly records in aggregate, reduces to fewer than 9,000 records per load, and its price, photovoltaic, and disturbance signals are generated from stationary, feature-level models rather than raw high-frequency measurements. Under these conditions, a deep sequence model would need to be evaluated carefully against overfitting, and against the

Random Forest baseline on a like-for-like basis, before any accuracy gain could be attributed to model architecture rather than to feature engineering. Deployment considerations also differ: Long Short-Term Memory and Transformer models require materially more tuning and computation than Random Forest, and their decisions are harder to audit against the hard safety constraints of Section 4, a relevant concern for a system that gates industrial equipment operation.

A design that may capture much of this upside while limiting the downside is a two-stage hybrid: a recurrent or attention-based encoder learns temporal representations from the raw price, photovoltaic, and disturbance sequences, and a Random Forest, or a similarly interpretable model, consumes these learned features alongside the existing engineered features to produce the final feasibility classification. This preserves an auditable final decision layer while allowing the representation-learning stage to capture dependencies that hand-engineered features may miss. The principal costs are implementation and calibration complexity, since the two stages must be trained and validated together, and the loss of full end-to-end interpretability, since the intermediate learned features are not directly physically meaningful.

Section 5.4 establishes that the current scheduler is a deterministic, myopic greedy heuristic with no backtracking across loads. Genetic algorithms, and multi-objective variants such as NSGA-II, are well matched to the combinatorial structure of industrial demand-side management, where cost minimization, peak-demand reduction, production continuity, and disturbance-risk exposure are competing objectives rather than a single scalar cost; genetic-algorithm-based approaches have been demonstrated for exactly this class of problem, including power-quality-aware industrial scheduling [26]. A natural design retains the Random Forest as a feasibility filter, encoding the safety-critical constraints of Section 4 exactly as it does now, while replacing the greedy block-selection step with a genetic algorithm whose chromosome encodes a candidate start time for each load on each day, and whose fitness function combines energy cost, a peak-demand penalty, and a disturbance-risk margin. Such a scheduler could search jointly across loads rather than sequentially, potentially escaping the local, load-by-load decisions the greedy algorithm is structurally committed to, and would also provide a natural mechanism for testing the sensitivity hypothesis raised in Section 3.7. The expected cost is computational: population-based search over a multi-day, multi-load horizon is markedly more expensive than the low-order polynomial greedy procedure of Section 5.4, which likely confines a genetic-

algorithm-based scheduler to offline, day-ahead planning rather than the fast, repeated re-scheduling that the present greedy approach supports.

These two extensions are complementary rather than competing: deep sequence modeling targets the accuracy of the feasibility prediction, while evolutionary optimization targets the quality of the resulting schedule given that prediction. Both are left as future work so as to keep the contribution of the present paper, namely that disturbance awareness itself, independent of the specific prediction or optimization technique used, materially improves the reliability of AI-based industrial demand-side management, clearly separated from the choice of algorithm used to realize it.

7 Conclusion

This paper presented a fault-aware AI-based demand-side management framework for industrial load scheduling that explicitly accounts for electrical disturbances alongside traditional temporal, economic, and renewable energy factors. Unlike conventional DSM approaches that assume ideal grid conditions, the proposed framework integrates disturbance awareness into both the machine learning decision layer and the scheduling process, enabling reliable and physically feasible operation under realistic industrial power system conditions.

A comprehensive industrial case study was developed using a realistic dataset that combines industrial load characteristics, photovoltaic generation, dynamic electricity pricing, and modeled electrical disturbances consistent with industrial reliability benchmarks. A Random Forest classifier was employed to learn operating feasibility under these combined conditions, and its outputs were embedded within a constraint-aware scheduling algorithm that preserves industrial production requirements.

Simulation results demonstrated that the proposed fault-aware DSM framework achieves substantial operational improvements compared to a fault-unaware baseline. Aggregate load profiles showed a significant reduction in peak-hour operation and a clear shift toward disturbance-safe and economically favorable periods. Peak-demand and peak-duration metrics were markedly reduced, while total operating hours and production continuity were fully preserved. Importantly, the inclusion of disturbance awareness did not compromise economic performance; cost-effective scheduling behavior was maintained while significantly enhancing robustness and reliability.

The findings confirm that fault awareness is a critical and previously underrepresented dimension in industrial DSM design. By preventing infeasible schedules during outage conditions and cautiously managing voltage-degraded operation, the proposed framework bridges the gap between theoretical DSM optimization and practical industrial deployment. The modular architecture and relatively low computational complexity further support its scalability and applicability to real-world industrial environments.

Overall, this work demonstrates that integrating machine learning-based fault awareness into industrial DSM enables a balanced achievement of economic efficiency, operational feasibility, and system reliability, making it a promising pathway for next-generation industrial energy management systems.

Building on this foundation, three concrete extensions are identified for future work: a systematic sensitivity analysis of the fault-aware benefit across a range of disturbance rates (Section 3.7), deep sequence models such as Long Short-Term Memory networks or Transformer architectures for the feasibility-prediction layer (Section 6.1), and evolutionary multi-objective scheduling as an alternative to the greedy heuristic (Sections 5.4 and 6.1).

Acknowledgement:

The author thankfully acknowledges the Power Electronics and Power Systems (PEPS) unit, Palestine Polytechnic University for its valuable support.

References:

- [1] Palensky, P.; Dietrich, D. Demand Side Management: Demand Response, Intelligent Energy Systems, and Smart Loads. *IEEE Trans. Ind. Inform.* 2011, 7, 381–388. <https://doi.org/10.1109/TII.2011.2158841>
- [2] Fouad Zaro, “Smart Demand Side Management for Techno-Economic Analysis of Microgrid System,” *International Journal of Electrical Engineering and Computer Science*, vol. 5, 2023, pp. 205–209. <https://doi.org/10.37394/232027.2023.5.23>
- [3] Pipattanasomporn, M.; Kuzlu, M.; Rahman, S. An Algorithm for Intelligent Home Energy Management and Demand Response Analysis. *IEEE Trans. Smart Grid* 2012, 3, 2166–2173. <https://doi.org/10.1109/TSG.2012.2201182>
- [4] Chen, S.; Gong, F.; Zhang, M.; Yuan, J.; Liao, S.; Chen, H. Planning and Scheduling for Industrial Demand-Side Management: State of the Art and Challenges. *Sustainability*

- 2021, 13, 7753.
<https://doi.org/10.3390/su13147753>
- [5] Alasseri, R.; Tripathi, A.; Rao, T.; Sreekanth, K. A review on implementation strategies for demand side management (DSM) in smart energy systems. *Renew. Sustain. Energy Rev.* 2017, 77, 617–628.
<https://doi.org/10.1016/j.rser.2017.04.023>
- [6] Bakare, M.S.; Abdulkarim, A.; Zeeshan, M.; Shuaibu, A.N. A comprehensive overview on demand side energy management towards smart grids. *Energy Informatics* 2023, 6, 4.
<https://doi.org/10.1186/s42162-023-00262-7>
- [7] Forootan, M.M.; Larki, I.; Zahedi, R.; Ahmadi, A. Machine learning and deep learning in energy systems: A review. *Sustainability* 2022, 14, 4832.
<https://doi.org/10.3390/su14084832>
- [8] Elsis M, Amer M, Su CL. A comprehensive review of machine learning and IoT solutions for demand side energy management, conservation, and resilient operation. *Energy*. 2023 Oct 15;281:128256..
<https://doi.org/10.1016/j.energy.2023.128256>
- [9] Breiman, L. Random Forests. *Mach. Learn.* 2001, 45, 5–32.
<https://doi.org/10.1023/A:1010933404324>
- [10] Dudek, G. A comprehensive study of random forest for short-term load forecasting. *Energies* 2022, 15, 6637.
<https://doi.org/10.3390/en15207547>
- [11] Sardar, A.; Khan, S.U.; Hassan, M.A.; Qureshi, I.M. A demand side management scheme for optimal power scheduling of industrial loads. *Energy Systems* 2023, 14, 335–356. <https://doi.org/10.1007/s12667-022-00510-x>
- [12] Croock, M.; Yao, R.; Li, J.; Zuo, B.; Hu, J. Machine learning-based energy-efficient technologies for smart grids. *Int. Trans. Electr. Energy Syst.* 2023, 31, e12744.
<https://doi.org/10.1002/2050-7038.12744>
- [13] Abolhasanpour, A.M.; Suratgar, A.A. Renewable energy and demand-side management in microgrids using Stackelberg game theory. *WSEAS Trans. Power Syst.* 2018, 13, 1–12.
- [14] Rollo, A.; Serafini, P.; Aleotti, F.; Cilio, D.; Morandini, E. Load shifting and demand-side management in renewable energy communities. *Energies* 2025, 18, 872.
<https://doi.org/10.3390/en18040872>
- [15] Xiao, G.; Liu, H.; Nabatalizadeh, J. Optimal scheduling and energy management of a multienergy microgrid incorporating demand response. *Sci. Rep.* 2025, 15, 5075.
<https://doi.org/10.1038/s41598-025-88776-w>
- [16] Eissa, M.M.; Swief, R.A.; Salam, T.S.A. Demand side management with renewable energy integration under uncertainty. *Sci. Rep.* 2025, 15, 18570.
<https://doi.org/10.1038/s41598-025-00752-6>
- [17] Tyagi, M.; Khan, M.I.; Gupta, S. A comprehensive study of voltage sag and swell in power distribution systems. *J. Electr. Syst.* 2024, 20, 1–18.
- [18] IEEE Std 4932007. IEEE Recommended Practice for the Design of Reliable Industrial and Commercial Power Systems; IEEE: New York, NY, USA, 2007.
- [19] ABB Group. Power Quality Disturbances and Their Effect on Industrial Equipment. ABB Technical White Paper, 2023.
- [20] Moses P, Thakur K, Parimi AM, Pal M. Advanced Distributed Energy Resource Management System—A review. *Energy* 360. 2026 Mar 4:100062.
<https://doi.org/10.1016/j.energ.2026.100062>
- [21] Chen, B.; Li, Z.; Liu, B.; Fan, H.; Zhong, Q. Robust optimization for smart demandside management under uncertainty. *Sci. Rep.* 2025, 15, 19440.
<https://doi.org/10.1038/s41598-025-03728-8>
- [22] Bollen, M.H.J. Understanding Power Quality Problems: Voltage Sags and Interruptions; ISBN: 978-0-780-34713-7, IEEE Press: New York, NY, USA, 2000.
- [23] Khatib, T.; Bazyan, A.; Assi, H.; Malhis, S. Palestine Energy Policy for Photovoltaic Generation: Current Status and What Should Be Next? *Sustainability* 2021, 13, 2996.
<https://doi.org/10.3390/su13052996>
- [24] Ungureanu, S.; Topa, V.; Cziker, A.C. Deep Learning for Short-Term Load Forecasting—Industrial Consumer Case Study. *Applied Sciences* 2021, 11, 10126.
<https://doi.org/10.3390/app112110126>
- [25] L'Heureux, A.; Grolinger, K.; Capretz, M.A.M. Transformer-Based Model for Electrical Load Forecasting. *Energies* 2022, 15, 4993.
<https://doi.org/10.3390/en15144993>
- [26] Eisenmann, A.; Streubel, T.; Rudion, K. Power Quality Mitigation via Smart Demand-Side Management Based on a Genetic Algorithm. *Energies* 2022, 15, 1492.
<https://doi.org/10.3390/en15041492>

Contribution of Individual Authors to the Creation of a Scientific Article (Ghostwriting Policy)

The authors equally contributed in the present research, at all stages from the formulation of the problem to the final findings and solution.

Sources of Funding for Research Presented in a Scientific Article or Scientific Article Itself

No funding was received for conducting this study.

Conflict of Interest

The authors have no conflicts of interest to declare that are relevant to the content of this article.

Creative Commons Attribution License 4.0 (Attribution 4.0 International, CC BY 4.0)

This article is published under the terms of the Creative Commons Attribution License 4.0

<https://creativecommons.org/licenses/by/4.0/deed.en>

[US](#)