IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 30 June 2026, accepted 17 July 2026, date of publication 14 August 2026, date of current version 21 August 2026.

Digital Object Identifier 10.1109/ACCESS.2026.3724048

RESEARCH ARTICLE

An Explainable Agentic Architecture for Multi-Site Photovoltaic Power Forecasting

LEÏLA SAKLI, (Member, IEEE), SEIFEDDINE BEN ELGHALI, AND YACINE MERRAD
CNRS, LIS, Aix Marseille Université, 13397 Marseille, France

Corresponding authors: Leïla Sakli (leila.sakli@lis-lab.fr) and Seifeddine Ben Elghali (seifeddine.benelghali@lis-lab.fr)

This work was supported by European Union's Horizon Europe Research and Innovation Program under Grant 101096946 (FlexCHESS) and Grant 101096836 (MASTERPIECE).

ABSTRACT Photovoltaic (PV) power forecasting is critical for grid stability and renewable energy integration. However, existing approaches prioritize predictive accuracy while neglecting operational requirements: physical validity under noisy data, transparent decision-making, and scalable multi-site deployment. This paper proposes an explainable architecture for multi-site photovoltaic forecasting that extends a validated hybrid ensemble pipeline through system-level design. Forecasting is structured as an explicit reasoning process comprising perception, inference, validation, explanation, and action stages. Each photovoltaic site uses an autonomous Forecasting Unit (FU), deployed as a microservice for scalable and consistent operation across sites. Explainability is achieved through explicit model orchestration, constraint-based validation, and traceable fallback mechanisms. Validation on real-world PV systems demonstrates that the architecture maintains competitive forecasting accuracy while ensuring physical validity and transparent model behavior. All forecasts are operationally applicable with traceable decision-making processes. These results highlight the relevance of architectural organization and an agentic-by-design (deterministic reasoning loop) architecture, based on explicit validation and orchestration mechanisms, for deploying trustworthy photovoltaic forecasting systems in industrial environments.

INDEX TERMS Photovoltaic power forecasting, explainable artificial intelligence, agent-based systems, hybrid ensemble learning, physical constraint validation, microservices architecture, multi-site energy systems.

I. INTRODUCTION

Over the past decade, solar energy systems have been increasingly integrated into modern power grids. This growth is driven by sustainability targets and lower installation costs [1]. As the capacity for renewable energy continues to proliferate, the significance of PV power forecasting has emerged as an essential element of contemporary grid operations [2], [3]. Although machine learning frameworks have considerably increased the reliability of short-term forecasting [4], [5], their practical application within operational frameworks uncovers challenges that cannot be overcome exclusively through algorithmic refinement. In practice, forecasting systems must simultaneously ensure physical validity under imperfect data conditions, provide transparent and trustworthy decision-making processes, and operate autonomously across geographically distributed sites [6], [7].

Physical validity under real-world conditions remains a fundamental challenge. Data collected from operational PV systems are frequently affected by sensor noise, communication failures, and transient anomalies [8], [9], [10], [11], [12]. When trained on historical data and deployed on corrupted inputs, machine learning models may generate physically implausible predictions, such as non-zero power production at night or outputs exceeding nominal capacity [12]. Post-processing techniques can correct obvious violations, but when employed in isolation, they act reactively and may obscure systematic model failures rather than preventing them.

Transparency and operational trust are also vital components for the successful implementation in real-world contexts. Grid operators necessitate more than mere precision

The associate editor coordinating the review of this manuscript and approving it for publication was Yifan Zhou.

in forecasting. They require a clear understanding of the methodologies employed in generating predictions, especially when such forecasts influence pivotal operational or market decisions [6]. Recent research underscores the inherent trade-off between the accuracy of forecasts and their interpretability in the domain of photovoltaic power forecasting, highlighting that exceedingly precise black-box machine learning models may impede the establishment of operational trust [13]. In response, most existing solutions address transparency primarily at the model level, without making the forecasting process itself explicit or traceable during operation.

Autonomous operation across distributed sites introduces a third, often overlooked challenge. Large photovoltaic parks are spread across multiple locations, and each site faces its own data-quality issues and local environmental conditions [14]. Many forecasting pipelines rely on centralized coordination, which can create a single point of failure. Others operate as isolated black boxes and provide little insight to operators. In both cases, scalability and transparency remain limited for large-scale deployment.

Recent works tackle parts of these problems. However, these requirements are rarely combined within a single operational framework. Physics-informed models incorporate physical constraints during training but typically lack explicit validation at inference time [15], [16]. Ensemble methods improve robustness but do not enforce physical coherence explicitly [17], [18]. Explainable AI techniques remain largely post-hoc [13], and distributed system architectures focus on deployment scalability rather than structured forecasting reasoning [19], [20].

To address these limitations, we propose an agentic architecture that reframes PV forecasting as a process of autonomous reasoning rather than a black-box input/output transformation. Here, "agentic" denotes a deterministic reasoning loop instead of a learning-based agent, while "autonomous" refers to operational independence, meaning that each FU runs without central coordination, and not to autonomous decision-making or learning. In this architecture, each PV installation is managed by an FU that executes five reasoning stages: perception, inference, validation, explanation, and action. Physical constraints are validated within this loop, and FUs operate independently as microservices without centralized coordination.

This work makes the following contributions. First, we propose an agentic multi-site architecture for photovoltaic power forecasting that integrates established forecasting models, implemented as autonomous Forecasting Units executing a five-stage reasoning loop and designed for traceable operation. Second, we embed constraint validation and fallback mechanisms within the runtime loop to ensure physically valid forecasts under degraded data conditions. Third, we deploy the architecture as microservices and validate it on operational multi-site PV systems, demonstrating competitive accuracy alongside end-to-end traceability.

The remainder of this paper is organized as follows. Section II reviews related work and positions our contribution within the existing literature. Section III presents the forecasting pipeline. Section IV introduces the reasoning-oriented forecasting methodology and its built-in explainability. Section V describes the system architecture and microservices deployment. Section VI reports the experimental evaluation. Section VII discusses results and limitations, and Section VIII concludes the paper and outlines directions for future research.

II. BACKGROUND ON HYBRID FORECASTING AND AGENTIC SYSTEMS

Photovoltaic power forecasting has been extensively studied across multiple research dimensions, including data-driven predictive modeling, physical constraint integration, explainability mechanisms, and system-level deployment architectures. Substantial progress has been achieved in each dimension independently. However, integrating these advances into coherent, operational forecasting systems remains limited. This section reviews representative advances and positions the gap addressed by our agentic architecture.

A. ACCURACY-FOCUSED FORECASTING METHODS

Data-driven machine learning approaches are widely used for PV power forecasting, with deep learning models demonstrating strong performance in capturing nonlinear relationships and temporal dependencies. A thorough examination [4] contrasts multiple machine learning methods and underscores the superior performance of recurrent neural networks (RNNs) and long short-term memory (LSTM) models in representing temporal dynamics of PV energy production. More recent studies increasingly explore attention-based mechanisms and transformer-inspired architectures to capture longer-range temporal dependencies and complex multivariate interactions in PV forecasting [6], [7].

Ensemble methods have emerged as a robust approach to improve forecasting accuracy and reliability [21]. Reference [17] proposes a stacking framework that combines extreme learning machines (ELM) and convolutional neural networks (CNN) as base learners with a Bayesian Ridge meta-learner, incorporating Monte Carlo Dropout for uncertainty quantification. Their results show that ensemble aggregation can improve both accuracy and reliability, although physical constraint validation and architectural explainability are not explicitly addressed. Similarly, [9] conducts a systematic mapping study of forecasting models in solar and wind energy, confirming the prevalence of ensemble approaches while noting limited attention to operational reliability and validation mechanisms.

These approaches often deliver strong predictive performance. On the other hand, they usually frame forecasting as a black-box mapping from meteorological inputs to power outputs, with limited attention to physical plausibility and transparent reasoning.

B. PHYSICS-INFORMED AND CONSTRAINT-AWARE MODELING

Recognizing the limitations of purely data-driven forecasting approaches, recent research has increasingly explored the integration of physical principles and domain knowledge into machine learning models for power systems. Surveys on physics-informed learning in power systems emphasize that conventional deep learning models may suffer from physically infeasible or inconsistent predictions, limited generalizability, and low interpretability, motivating the integration of physics-informed rules into data-driven methods [15]. In this context, physics-informed neural networks (PINNs) explicitly embed physical laws and constraints into the learning process, enabling improved physical consistency and robustness in energy applications [22]. In parallel, recent probabilistic forecasting approaches incorporate physical insight implicitly through problem formulation and uncertainty modeling, improving forecast reliability without enforcing explicit physical constraints [23].

Beyond model-centric constraint integration, several works have investigated constraint-aware data processing techniques to address data quality issues in operational energy systems. Adaptive imputation frameworks for SCADA data leverage spatiotemporal correlations and feature dependencies to reconstruct missing measurements under diverse missing-data patterns, improving the reliability of downstream forecasting models [10].

Recent photovoltaic forecasting systems combine physical models with data-driven predictors and apply domain-specific constraints, such as capacity limits and daylight conditions, to enhance physical plausibility and robustness in operational settings [16], [21].

However, across physics-informed modeling and constraint-aware data processing, constraint handling is predominantly embedded within training procedures or applied as offline or post-processing mechanisms. Constraints are therefore enforced reactively rather than through explicit validation stages operating at runtime. As a result, existing approaches do not provide proactive mechanisms for systematically assessing the physical validity of forecasts before their release for operational use.

C. EXPLAINABILITY IN ENERGY FORECASTING SYSTEMS

The need for transparent and interpretable forecasting systems has driven significant research in explainable artificial intelligence (XAI) for energy applications. Post-hoc explanation techniques have been widely applied to solar forecasting models to provide insights into model behavior. In particular, SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations) have been used to derive feature importance rankings and local explanations for individual predictions [24]. Similarly, SHAP values have been employed alongside optimized forecasting models to enhance interpretability for stakeholders [25].

Recent work has explored attention-based interpretability mechanisms. MATNet, a multi-level fusion and self-attention transformer, is proposed in [11] and provides interpretability through attention weight visualization, enabling users to understand which temporal features and spatial relationships drive predictions. The GPT-Agent approach in [26] introduces an LLM-based agent to interpret weather descriptions and provide natural language explanations for forecasting decisions, demonstrating the potential of agentic components for enhancing transparency.

While these XAI techniques successfully provide post-hoc explanations for model outputs, they fundamentally differ from explainability-by-design. Post-hoc methods explain what a black-box model has decided; they do not make the decision-making process itself transparent through architectural design. Furthermore, they typically ignore key elements of operational forecasting systems, such as constraint validation, fallback mechanisms, and model orchestration.

D. MULTI-AGENT SYSTEMS AND DISTRIBUTED ARCHITECTURES

Multi-agent systems have been explored in renewable energy contexts primarily for coordination, control, and distributed learning. Intelligent agent orchestration has been proposed for climate-resilient energy systems, with a focus on autonomous management of renewable integration and distributed resources through coordinated agent behaviors [19]. In a similar vein, a comprehensive framework for AI in renewable energy management is presented in [27], discussing multi-agent coordination patterns, constraint-aware decision methods, and assurance artifacts for trustworthy AI systems. While conceptually rich, these works mainly provide design patterns and architectural guidelines rather than implemented forecasting systems with operational agentic reasoning.

In the context of distributed learning, multi-agent federated learning has been applied to regional photovoltaic power prediction, where each site operates as an agent that trains local models while preserving data privacy [28]. This approach demonstrates effective multi-agent coordination for distributed training, but does not frame forecasting itself as an explicit agentic reasoning process with well-defined perception, inference, validation, explanation, and action stages.

This distinction is critical. Existing multi-agent approaches employ agents for coordination, control, or distributed learning, but not for structuring the forecasting reasoning process itself. To the best of our knowledge, no prior work has architected photovoltaic forecasting as an explicit agentic process in which each autonomous unit implements staged reasoning with built-in validation and fallback mechanisms.

E. MICROSERVICES AND DEPLOYMENT ARCHITECTURES

Service-oriented architectures have been widely recommended to support scalable deployment of forecasting systems in operational energy environments. For example, the integration of forecasting services within a local energy trading platform is demonstrated in [29], where forecasting profiles are coupled with blockchain-based marketplace mechanisms using microservices patterns. In a related direction, a systematic review of AI-enabled digital twins for smart grids is presented in [30], emphasizing closed-loop validation, semantic interoperability, and container-orchestrated microservices deployment for field systems.

Practical deployment challenges have also been addressed in photovoltaic forecasting contexts. AutoPV, proposed in [31], introduces an automated forecasting framework based on an ensemble of pre-trained models designed to operate under limited information, such as missing metadata and cold-start scenarios. While this work explicitly considers deployment constraints, it does not implement autonomous forecasting agents with explicit reasoning stages.

A recent line of work structures PV forecast delivery as an explicit, constraint-aware process qualifying a baseline forecast against physical and operational constraints at runtime and emitting machine-readable evidence (triggered checks, reason codes, and correction events) for auditability [32]. This highlights the value of treating validation and explanation as first-class runtime artifacts. However, it remains confined to a single forecast stream and does not address operation across geographically distributed installations or per-site model behavior. This paper extends and generalizes that single-site wrapper into a multi-site, microservice-based agentic architecture.

Operational robustness has been further explored through predictive analytics for distributed photovoltaic systems. The framework presented in [18] emphasizes real-time adaptability and heterogeneous data fusion to improve grid resilience. However, it does not incorporate agentic reasoning or explicit runtime validation stages within the forecasting process.

Beyond PV forecasting, operational energy AI increasingly emphasizes resilience and explainability at the system level rather than predictive accuracy alone. For instance, Kırat et al. [33] combine deep learning with mixed-integer optimization for battery energy storage operation, highlighting the role of validation and decision logic in deployment. Similarly, AI-based operational frameworks for hydrogen refueling infrastructure stress resilience-oriented design principles for critical energy services [34]. These systems show that forecasting is not an end in itself but rather a component of a larger decision-making architecture that must balance multiple objectives under operational constraints.

Current deployment architectures either propose microservices patterns conceptually without demonstrating operational forecasting agents, or implement forecasting services as conventional pipelines without agentic reasoning. The gap lies in the absence of autonomous Forecasting Units that combine microservices scalability with agentic reasoning, runtime validation, and autonomous fallback mechanisms.

F. MACHINE LEARNING AND ENSEMBLE METHODS FOR PV FORECASTING

New studies are now delving into blended strategies that integrate several paradigms. In [17], ensemble learning has been integrated with uncertainty quantification and SHAP-based explainability to address multiple objectives within a unified framework. Nevertheless, this methodology does not encompass explicit constraint validation or agentic reasoning phases within the forecasting mechanism.

Research into hybrid renewable energy forecasting systems [35] looks into merging diverse data sources and modeling strategies within containerized deployment frameworks, confronting challenges of scalability and interoperability.

These systems illustrate the viability of intricate integration but lack the explicit reasoning architecture and autonomous decision-making capabilities that are posited in our research.

G. RESEARCH GAP AND POSITIONING

Table 1 summarizes the capabilities of representative recent works across five key dimensions: agentic reasoning architecture, proactive constraint validation, explainability-by-design, autonomous microservices deployment, and data-quality fallback mechanisms. As shown, existing approaches address individual dimensions or pairs of dimensions, but no prior work integrates all five within a unified system.

The limitations of existing approaches can be grouped into several categories, as summarized below.

- *High-accuracy models without operational reliability:* Machine learning and ensemble methods [4], [11], [17] optimize for predictive accuracy but lack explicit constraint validation and transparent reasoning mechanisms required for operational deployment.
- *Constraint-aware models without explicit validation stages:* Physics-informed approaches [22], [23] embed physical constraints during training but do not implement runtime validation as a distinct and inspectable reasoning stage.
- *Post-hoc explainability without architectural transparency:* Explainable AI techniques [24], [25] provide explanations of model outputs but do not make the forecasting process itself transparent through architectural design.
- *Multi-agent coordination without agentic forecasting reasoning:* Multi-agent systems [19], [27], [28] employ agents for coordination or distributed learning rather than for structuring the forecasting reasoning process itself.
- *Deployment patterns without operational agents:* Microservices architectures [29], [30] recommend scalable deployment patterns but do not demonstrate autonomous forecasting agents with explicit reasoning capabilities.
- *Partial integrations without unified systems:* Some works combine two or three of these dimensions [10], [27], [30] but do not achieve end-to-end

TABLE 1. Comparison of related work across key dimensions. A work is considered to satisfy a given dimension only when the corresponding capability is explicitly implemented as an operational system component, rather than implicitly addressed at the model or conceptual level.

Work	Year	Agentic Reasoning	Proactive Constraints	Explainability by Design	Autonomous Microservices	Data Fallback
Gaboitaolelwe et al. [4]	2023	not addressed	not addressed	not addressed	not addressed	not addressed
Hassan et al. [17]	2025	not addressed	not addressed	SHAP	not addressed	uncertainty
Kuzlu et al. [24]	2020	not addressed	not addressed	SHAP/LIME	not addressed	not addressed
Baelongandi et al. [22]	2024	not addressed	training	architecture	not addressed	not addressed
Simeunovic et al. [23]	2024	not addressed	training	architecture	not addressed	not addressed
Yang et al. [10]	2025	not addressed	dynamic	not addressed	not addressed	imputation
Yan et al. [26]	2024	interpretation	not addressed	LLM	not addressed	not addressed
Li et al. [28]	2024	coordination	not addressed	not addressed	distributed	not addressed
Ali et al. [29]	2023	not addressed	not addressed	not addressed	services	not addressed
Ahmed and Khan [30]	2025	not addressed	closed-loop	patterns	patterns	not addressed
Pasupuleti [27]	2024	coordination	optimization	artifacts	not addressed	not addressed
Li et al. [18]	2025	not addressed	not addressed	not addressed	distributed	resilience
This Work	**2025**	**fully implemented**	**fully implemented**	**fully implemented**	**fully implemented**	**fully implemented**

integration of all five key characteristics within a single operational framework.

This paper addresses this gap by proposing an explainable agentic architecture that integrates these five dimensions into an operational multi-site PV forecasting system.

III. FORECASTING PIPELINE

A. OVERVIEW OF THE FORECASTING PIPELINE

The forecasting pipeline adopted in this work follows a layered and modular design, intended to provide reliable short-term photovoltaic power predictions under heterogeneous and potentially degraded data conditions. As shown in Fig. 1, rather than introducing novel predictive models, the pipeline combines complementary forecasting approaches within a unified and operationally validated structure [16]. Three complementary predictors are employed: (i) a physics-based predictor grounded in physical principles, (ii) an LSTM-based sequence model designed to capture temporal dependencies, and (iii) an XGBoost-based tree ensemble regressor modeling nonlinear interactions between meteorological variables and recent production lags.

Each predictor captures a specific aspect of photovoltaic generation behavior. Their outputs are combined through constrained ensemble aggregation and refined using context-aware mechanisms. A final validation stage enforces physical and operational consistency before forecasts are delivered to downstream systems.

The pipeline operates at an hourly resolution and produces forecasts over a 24-hour horizon.

B. PHYSICS-BASED COMPONENT

The physics-based component provides a physically coherent estimate of photovoltaic power production grounded in solar geometry and irradiance relationships. Global horizontal irradiance (GHI), measured on a horizontal plane, is decomposed into its diffuse horizontal component (DHI) and the direct normal irradiance (DNI), i.e., beam irradiance on a plane normal to the sun rays, as:

$$\mathrm{GHI} = \mathrm{DHI} + \mathrm{DNI} \cdot \cos(\theta_z), \tag{1}$$

where θ_z denotes the solar zenith angle.

The irradiance is then projected onto the plane of the photovoltaic array (POA) using a transposition model:

$$\mathrm{POA} = \mathrm{POA}_{\mathrm{direct}} + \mathrm{POA}_{\mathrm{diffuse}} + \mathrm{POA}_{\mathrm{reflected}}, \tag{2}$$

where $\mathrm{POA}_{\mathrm{direct}}$ denotes the beam component incident on the tilted plane, $\mathrm{POA}_{\mathrm{diffuse}}$ the sky-diffuse contribution, and $\mathrm{POA}_{\mathrm{reflected}}$ the ground-reflected term (often linked to surface albedo).

Rather than modeling internal electrical or thermal subcomponents, the resulting irradiance signal is mapped to a power estimate through calibrated site-specific coefficients that capture aggregate system behavior. This abstraction preserves physical consistency while remaining robust under partial or imperfect data availability.

To ensure operational continuity under degraded input conditions, the physics-based component incorporates an explicit fallback mechanism. When critical inputs are missing, inconsistent, or deemed unreliable, the pipeline switches to a simplified physically grounded estimator:

$$P_{\mathrm{fallback}} = P_{\mathrm{nameplate}} \cdot \frac{\mathrm{POA}}{1000} \cdot [1 - 0.004\,(T_{\mathrm{module}} - 25)] \cdot \eta_{\mathrm{system}}. \tag{3}$$

Here, P_{fallback} denotes the fallback power estimate, $P_{\mathrm{nameplate}}$ is the rated PV capacity at standard test conditions (STC), and POA is the plane-of-array irradiance incident on the module surface. The constant 1000 is the STC reference irradiance used to normalize POA. T_{module} is the PV module temperature, and 25 °C is the STC reference temperature; the term $[1 - 0.004\,(T_{\mathrm{module}} - 25)]$ implements a linear temperature derating using a temperature coefficient of 0.004 °C^{-1}. Finally, η_{system} is a dimensionless aggregated efficiency factor capturing average system-level losses. The activation of this fallback mechanism is deterministic and governed by explicit validity checks on the inputs, ensuring that fallback-based predictions remain physically plausible.

Fallback activation is explicitly reported as part of the forecasting output, allowing downstream components

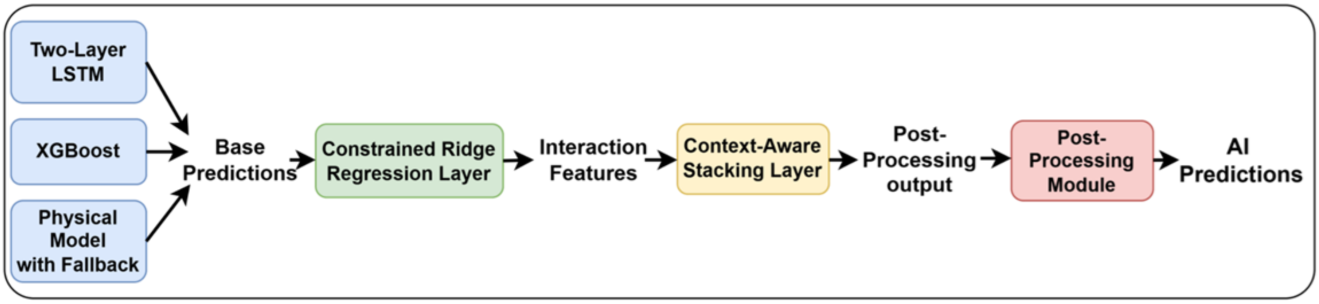

FIGURE 1. Foundational hybrid forecasting pipeline.

to distinguish between nominal and fallback-based predictions and contributing to operational transparency and auditability.

C. SEQUENCE-BASED LEARNING MODEL

Temporal dependencies inherent in photovoltaic generation are primarily addressed through a sequence-based learning framework based on recurrent neural networks. This component focuses on modeling short-term temporal dynamics and diurnal patterns and is trained on sliding windows of historical data. It leverages meteorological variables, temporal indicators, and recent production measurements to capture sequential behavior under varying operating conditions. Temporal features are encoded using cyclical representations to preserve periodicity, such as:

$$\text{hour}_{\text{sin}} = \sin\left(\frac{2\pi \cdot \text{hour}}{24}\right), \quad \text{day}_{\text{cos}} = \cos\left(\frac{2\pi \cdot \text{day}}{365.25}\right). \tag{4}$$

Model training relies on a weighted mean squared error objective. Given N samples, the loss is defined as:

$$\mathcal{L} = \frac{1}{N}\sum_{i=1}^{N} w_i (y_i - \hat{y}_i)^2, \tag{5}$$

where y_i denotes the observed PV power for sample i, $\hat{y}_i$ the corresponding prediction. The weights w_i are assigned during low-production periods to limit sensitivity to near-zero outputs.

D. TREE-BASED REGRESSION MODEL

A tree-based regression model is employed to capture nonlinear interactions between meteorological conditions, temporal variables, and recent photovoltaic power history. The feature set includes smoothed meteorological indicators, daylight flags, interaction terms, and lagged power values.

The model is trained using a regularized objective function:

$$\mathcal{L}(\phi) = \sum_i l(y_i, \hat{y}_i) + \sum_k \Omega(f_k), \tag{6}$$

where regularization terms control model complexity and support generalization under heterogeneous operating conditions.

E. ENSEMBLE AGGREGATION AND CONTEXT-AWARE REFINEMENT

Outputs from the individual predictors are combined using a constrained weighted aggregation scheme:

$$\hat{y}_{\text{ens}} = \sum_m w_m p_m, \tag{7}$$

subject to:

$$w_m \geq -0.5, \quad \sum_m w_m > 0. \tag{8}$$

These constraints prevent dominance by a single predictor while allowing limited corrective contributions to compensate for systematic biases. A subsequent refinement stage incorporates contextual information to adapt aggregation behavior across temporal and meteorological conditions without explicit data segmentation.

F. VALIDATION AND CONSISTENCY ENFORCEMENT

To ensure operational usability, aggregated forecasts are subjected to a final validation stage enforcing physical and operational constraints. Predicted power values are bounded by nominal system capacity according to:

$$\hat{y}_{\text{final}} = \min(\hat{y}_{\text{ens}}, \alpha \cdot P_{\text{nameplate}}), \tag{9}$$

where α is a safety factor.

Day–night consistency is enforced by setting predicted power to zero when solar elevation is non-positive. Model-level outliers are constrained using statistical bounds derived from historical behavior.

Through this combination of physical modeling, data-driven learning, constrained aggregation, explicit fallback handling, and deterministic validation rules, the forecasting pipeline provides a stable and inspectable prediction backbone. However, it does not, by itself, define how predictions are validated, contextualized, or operationally exposed within an industrial environment.

IV. REASONING-ORIENTED FORECASTING AND BUILT-IN EXPLAINABILITY

This section introduces the reasoning-oriented methodology adopted in this work, which structures an existing forecasting pipeline into an explicit, inspectable, and deterministic reasoning process.

The workflow corresponds to a functional decomposition of the forecasting pipeline into five reasoning stages. Each stage relies on established classes of methods commonly used in time-series forecasting, physical modeling, and constraint-based validation, together forming a coherent reasoning structure suitable for industrial deployment and directly mappable to software services within a microservices-based architecture (Fig. 2).

A. SCOPE AND MOTIVATION

In operational PV systems, forecasts are consumed as part of broader monitoring, planning, and energy management workflows, where downstream use depends on assessing forecast validity and propagating reliability indicators to subsequent systems.

In this work, reasoning-oriented forecasting refers to the explicit structuring of forecasting activities into well-defined stages governed by physical and operational constraints. This perspective enables transparent handling of forecast generation, validation, and delivery within deterministic and auditable industrial processes, without relying on post-hoc explanation techniques.

B. FORECASTING UNIT ARCHITECTURE

The system is organized around a site-level forecasting unit that encapsulates the forecasting pipeline together with its associated reasoning stages. Each photovoltaic installation operates an independent instance of this unit, which processes local data streams while following an identical reasoning structure across sites.

This organization ensures consistent forecasting behavior, validation logic, and explainability across heterogeneous installations. It also allows site-specific data characteristics to be handled locally. The FU serves as a self-contained, deployable building block that integrates predictive computation, validation mechanisms, and explanatory outputs.

C. PERCEPTION

At the perception stage, each installation contributes its historical power measurements, meteorological variables, and temporal context, which are assembled into the feature representation used downstream.

Time variables receive a cyclical encoding so that diurnal and seasonal periodicity is retained, and the inputs undergo integrity controls covering missing values, inconsistent timestamps, and the delimitation of nighttime intervals. The feature set produced in this manner constitutes the input to inference.

D. INFERENCE

Inference executes, for the installation under consideration, the hybrid ensemble pipeline of Section III. The physics-based estimator, the sequence model, and the tree-based regressor are combined through constrained aggregation, after which a stacking mechanism reweights each contribution according to the temporal context, the site-specific characteristics, and the prevailing data quality. What emerges is a preliminary power forecast, informed jointly by physical reasoning and data-driven temporal patterns, that has not yet been subjected to operational validation.

E. VALIDATION

Although validation follows inference in the execution order, it is architecturally proactive: it constitutes a distinct, explicitly formalized stage that every forecast traverses before release. Each inferred output is examined against the physical and operational constraints relevant to photovoltaic systems, with deterministic rules enforcing domain consistency through three controls: zero production outside daylight, an upper bound tied to nominal capacity, and the flagging of anomalous values. Whenever a constraint is breached, the corrective action and an accompanying validation indicator are recorded, and these indicators travel with the forecast to downstream components so that corrected or constrained predictions remain transparent in operation.

F. EXPLANATION

Explainability in the proposed framework is achieved by architectural design through explicit separation and inspection of reasoning stages. The objective is not to explain how a specific model computes its output but to make the forecasting process itself explicit, auditable, and traceable through well-defined reasoning stages and deterministic validation mechanisms.

In particular, the system explicitly reports:

- the activation of validation rules,
- the application of corrective constraints,
- and the triggering of fallback estimators under degraded input conditions.

In addition to predictive accuracy, the system enables explicit monitoring of operational reasoning behavior. During the evaluation period, validation rules and fallback mechanisms were automatically stored as part of the forecasting process. These records allow operators to quantify how often forecasts are corrected or generated under degraded input conditions, providing an additional layer of operational insight beyond traditional accuracy metrics. For example, whenever a fallback mechanism is activated, a structured explanation is stored together with the forecast, explicitly indicating the reason for this activation (e.g., missing or inconsistent meteorological inputs).

G. ACTION

At the action stage, the validated forecast is released to its consumers, accompanied by the indicators and explanations produced upstream. Recipients may include monitoring dashboards, energy management systems, or decision-support tools. Because each forecast carries its contextual and validation information, this stage enables informed operational use

Algorithm 1 Reasoning-Oriented Forecasting Process

Require: PV site s, horizon H, reference time t
Ensure: Forecast $\hat{P}$, validation status $\mathcal{V}$, explanation $\mathcal{E}$

```
1:  Perception: acquire data and construct feature set. X
2:  Inference: generate raw forecast P̂_raw from X
3:  Validation: check P̂_raw against operational constraints
4:  if constraints satisfied then
5:      P̂ ← P̂_raw; V ← compliant
6:  else
7:      apply correction or fallback; P̂ ← corrected; V ← corrected
8:  end if
9:  Explanation: construct explanation artifact E
10: Action: deliver P̂, V, E and persist decision artifacts
11: return P̂, V, E
```

while keeping predictive analytics distinct from the decision-making that builds upon it.

H. MULTI-SITE REASONING CONSISTENCY

In multi-site deployments, each photovoltaic installation operates an independent forecasting unit instance. While training and inference are performed on site-specific data, the forecasting pipeline, reasoning stages, and validation logic remain identical across installations.

This organization supports scalable deployment across geographically distributed assets. It preserves consistency in forecasting behavior, explainability, and operational semantics across sites.

While the workflow illustrated within each forecasting unit in Fig. 2 provides a functional decomposition of the reasoning-oriented forecasting process, Algorithm 1 explicitly formalizes the deterministic operational loop executed by each forecasting unit. This formalization specifies the sequential execution of perception, context-aware inference, validation, explanation, and action stages, together with the handling of constraint violations and fallback conditions in an auditable manner.

V. SYSTEM ARCHITECTURE AND MICROSERVICES DEPLOYMENT

This section describes the system architecture supporting the proposed agentic forecasting methodology as shown in Fig. 2. The objective is to operationalize the forecasting and reasoning pipeline described in Sections III and IV within a scalable and modular software architecture suitable for industrial photovoltaic deployments.

A. ARCHITECTURAL OVERVIEW

The proposed system follows a cloud-centric, microservices-based architecture designed to support multi-site photovoltaic forecasting. Within each site, multiple photovoltaic installations may be present. Each installation is associated with an independent autonomous Forecasting Unit, which encapsulates the forecasting pipeline and reasoning stages described previously. All FUs share the same forecasting structure and configuration. They operate autonomously on site-specific data streams.

A centralized cloud infrastructure provides shared services, including data storage, orchestration, and monitoring. Site-level autonomy is preserved for forecasting and validation. This design enables scalable deployment across geographically distributed assets without introducing inter-site dependencies.

B. DATA MANAGEMENT AND INGESTION

Operational data generated at each photovoltaic site are ingested into a centralized data lake hosted in the cloud. The data lake stores historical power measurements, meteorological inputs, and contextual information required for forecasting. Each FU retrieves the relevant data corresponding to its site and installation, ensuring data isolation and consistency.

Data ingestion and preprocessing are implemented as dedicated microservices, allowing independent scaling and maintenance. This separation supports robustness against data quality issues and facilitates integration with heterogeneous data sources.

C. AUTONOMOUS FORECASTING UNIT AS A MICROSERVICE

Each FU is deployed as an independent microservice responsible for executing the complete agentic reasoning loop, including perception, inference, validation, explanation, and action. The FU exposes well-defined interfaces for receiving input data and delivering forecasts together with associated indicators.

By encapsulating forecasting logic within a microservice, the system enables flexible deployment, fault isolation, and independent lifecycle management. This design also allows multiple FUs to run concurrently for different installations using the same underlying models.

D. MULTI-SITE DEPLOYMENT STRATEGY

In a multi-site context, the same forecasting and reasoning pipeline is instantiated independently for each photovoltaic installation. Training and inference are performed on a per-site basis, reflecting local characteristics and data distributions, even though the model architecture remains identical across sites.

This strategy avoids cross-site interference and preserves consistency in forecasting behavior and explainability. It also facilitates incremental onboarding of new sites without requiring system-wide retraining or reconfiguration. Table 2 reports the predictive performance of the forecasting pipeline for two photovoltaic sites located in different geographical regions. The same hybrid ensemble model and training procedure are applied independently at each installation, and performance is evaluated using the standard accuracy metrics adopted in prior work. This comparison aims to assess the

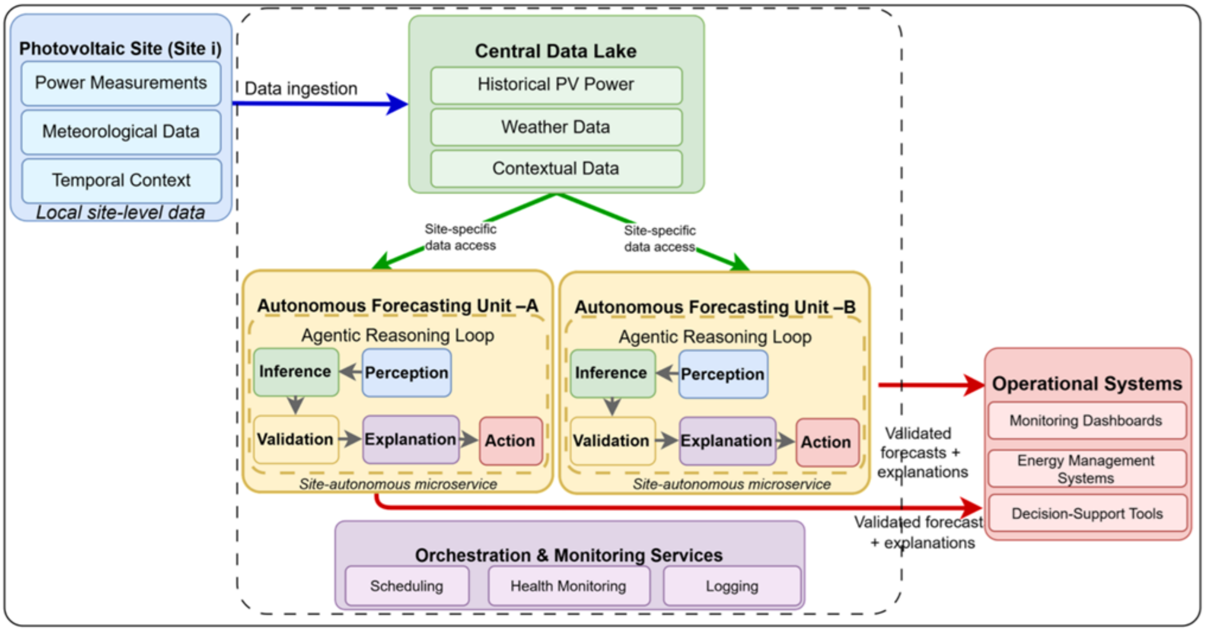

FIGURE 2. System architecture and microservices deployment.

robustness and consistency of the forecasting pipeline under heterogeneous climatic and operational conditions rather than to derive site-specific optimizations.

E. INTEGRATION WITH OPERATIONAL SYSTEMS

Validated forecasts and their associated indicators are delivered by the FUs to downstream operational systems, such as monitoring dashboards and energy management platforms. The microservices architecture supports asynchronous communication, ensuring that forecasting services remain decoupled from consumption and visualization layers.

This integration enables operational stakeholders to access forecasts enriched with validation and explanation artifacts, supporting informed decision-making without requiring direct interaction with predictive models.

F. ARCHITECTURAL CONSIDERATIONS

The proposed architecture emphasizes modularity, scalability, and transparency. By separating data ingestion, forecasting, reasoning, and delivery into distinct services, the system supports maintainability and extensibility. Moreover, the explicit alignment between architectural components and reasoning stages facilitates traceability and auditability in industrial environments.

VI. EXPERIMENTAL EVALUATION

This section presents the experimental evaluation conducted to assess the proposed explainable forecasting architecture in a multi-site photovoltaic context. The objective of the evaluation is twofold: first, to verify that the proposed system preserves the predictive performance of the underlying forecasting pipeline when deployed across multiple sites; second, to illustrate the operational benefits introduced by the explicit reasoning structure and built-in explainability mechanisms.

A. EXPERIMENTAL SETUP

The experimental evaluation is performed on real-world industrial photovoltaic installations located in different geographical regions. It includes two installations in Gökçeada, Turkey (TR-1 and TR-2), and one installation in Formentera, Spain (ES-1).

Each installation operates independently and is associated with its own autonomous FU, as described in Section IV. While the architecture and reasoning workflow are identical across all installations, model training and inference are conducted on a per-installation basis using local historical data. This setup reflects realistic operational conditions and ensures that observed differences in performance are attributable to installation-specific characteristics rather than architectural variations.

B. DATASETS AND SITES

The datasets consist of historical photovoltaic power measurements and corresponding meteorological variables collected from industrial photovoltaic systems. The selected installations exhibit heterogeneous climatic conditions, seasonal patterns, and operating regimes.

Data preprocessing follows the perception stage described in Section IV and includes temporal alignment, handling of missing values, and cyclical encoding of time-related features. For each installation, the dataset is divided into training and evaluation periods using a chronological split to reflect realistic forecasting scenarios and avoid information leakage.

TABLE 2. Predictive performance of the different forecasting algorithms for each photovoltaic installation.

Installation	Model	RMSE	MAE	MAPE (%)	Corr.
ES-1	Physics-based	0.57	0.18	155.0	0.86
ES-1	LSTM	0.45	0.15	181.5	0.90
ES-1	XGB	0.35	0.10	74.6	0.95
ES-1	Ensemble	**0.33**	**0.095**	**92.0**	**0.96**
TR-1	Physics-based	0.28	0.14	37.4	0.92
TR-1	LSTM	0.50	0.27	89.5	0.75
TR-1	XGB	0.10	0.050	84.4	0.96
TR-1	Ensemble	**0.09**	**0.045**	**101.0**	**0.97**
TR-2	Physics-based	0.20	0.10	40.1	0.95
TR-2	LSTM	0.38	0.21	154.0	0.86
TR-2	XGB	0.24	0.12	84.4	0.93
TR-2	Ensemble	**0.19**	**0.095**	**102.0**	**0.96**

C. EVALUATION METRICS

Forecasting performance is evaluated using a set of complementary accuracy metrics, including the root mean square error (RMSE), the mean absolute error (MAE), the mean absolute percentage error (MAPE), and the Pearson correlation coefficient. Together, these metrics provide a comprehensive view of forecasting behavior by capturing not only error magnitude and relative deviation but also robustness across installations with different nominal capacities and the degree of temporal agreement between predicted and observed power outputs.

It is noted that MAPE may exhibit high values during near-zero production periods, where percentage-based errors are known to be unstable; RMSE and correlation are therefore considered primary indicators of predictive behavior.

All metrics are computed independently for each installation to preserve installation-level variability and avoid aggregation effects that could mask site-specific behavior.

Table 2 reports the predictive performance of the physics-based model, the data-driven predictors, and the ensemble approach for each installation separately. By avoiding site-level aggregation, this evaluation protocol preserves installation-specific behavior and highlights both inter-site and intra-site variability. Across all three installations, the ensemble model exhibits the most stable overall performance on average, confirming its ability to leverage the complementary strengths of individual predictors under heterogeneous operating conditions.

D. MODEL SELECTION AND OPERATIONAL BEHAVIOR

Beyond aggregated accuracy metrics, the proposed system explicitly compares the outputs of multiple forecasting algorithms at runtime, including physics-based models, data-driven predictors, and the ensemble approach. This comparison enables the identification of the best-performing model for each installation over the evaluation horizon.

This selection is reported for analysis purposes and is not intended to define a universally optimal forecasting strategy. Instead, it illustrates how the agentic system exposes model-level performance information to support operational assessment. For each installation and time period, the system provides both a robust ensemble forecast and an indication of the best-performing individual model, allowing operational stakeholders to choose between robustness-oriented ensemble predictions and context-specific high-performance models. Table 3 summarizes the selected algorithm for each installation together with its associated performance metrics.

TABLE 3. Best performing forecasting algorithm for each photovoltaic installation.

Installation	Selected model	MAPE (%)	Corr.
ES-1	Ensemble	92.0	0.96
TR-1	XGB	84.4	0.96
TR-2	Physics-based	40.1	0.95

The results show that while the ensemble model is selected for multiple installations, alternative predictors may outperform it under specific operating conditions. This behavior highlights the relevance of exposing model selection outcomes in an agent-oriented forecasting system, enabling interpretability and traceability without modifying the underlying predictive models.

The operational behavior of the architecture is further reflected in the activation rates of its validation and fallback mechanisms over the evaluation horizon. Validation rules were triggered for 9.4% of forecasts, enforcing physical consistency and filtering anomalous predictions. The physics-based fallback estimator was activated for 0.75% of forecasts, corresponding to situations in which real production data were unavailable.

E. CASE STUDIES AND QUALITATIVE ANALYSIS

To complement the quantitative evaluation, representative case studies are presented to qualitatively analyze the behavior of the forecasting system under different operating conditions. These examples are illustrative rather than statistically exhaustive and are intended to highlight temporal dynamics, model-specific biases, and ensemble behavior in realistic scenarios.

Figure 3 presents a representative two-day forecasting example for a Turkish photovoltaic installation, illustrating the behavior of the different predictive models and the ensemble approach. The physics-based model, the data-driven predictors (XGBoost and LSTM), and the ensemble forecast are shown together to highlight their respective temporal dynamics.

During daytime periods, noticeable differences can be observed among individual models, particularly around peak production, reflecting their distinct modeling assumptions and sensitivities. In contrast, the ensemble forecast exhibits a more stable and balanced behavior, lying between individual predictions.

Importantly, the absence of positive power predictions during nighttime periods confirms the effective application of validation rules enforcing physical consistency. By exposing both individual model outputs and the ensemble forecast, the system enables transparent assessment of model behavior and supports informed operational decision-making.

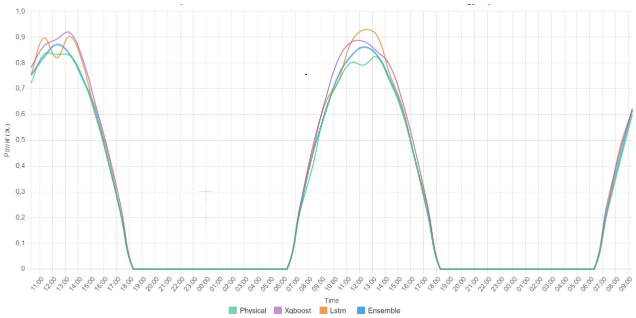

FIGURE 3. Measured and forecasted photovoltaic power for Turkish site TR-2.

VII. DISCUSSION AND LIMITATIONS

The experimental evaluation confirms that the proposed explainable forecasting architecture can successfully embed a validated hybrid forecasting pipeline within a multi-installation operational context. This work's primary value is in structuring PV forecasting as an explicit and inspectable process, supporting transparency, traceability, and integration into operations.

The results highlight that no single forecasting algorithm consistently outperforms others across all installations and operating conditions. This observation reinforces the relevance of exposing model-level behavior and performance variability, rather than enforcing a unique forecasting strategy. By delivering both a reliable combined forecast and clear identification of the top-performing individual models, the system enables well-informed operational decisions while maintaining human supervision.

Beyond predictive accuracy, the system demonstrates reasoning-oriented operational behavior through deterministic validation mechanisms and explicit communication of forecast status. The enforcement of physical consistency rules and the activation of fallback estimators under degraded data conditions are not treated as post-processing steps but as explicit and traceable architectural actions reported to downstream systems together with their underlying causes. This design choice directly addresses industrial requirements for auditability and trust in forecasting systems.

Several limitations of the proposed approach should be acknowledged. First, the reasoning mechanisms implemented in this work are deterministic and rule-based and do not involve adaptive learning or autonomous decision-making. Second, while we report aggregate activation rates for validation and fallback, the current evaluation does not provide a detailed quantitative analysis of per-rule and site-wise activation patterns, nor their marginal impact under heterogeneous operating conditions. The evaluation likewise does not characterize computational overhead, runtime performance, or the behavior of the architecture under a large number of concurrently deployed Forecasting Units. Future work will complement these results with dedicated per-rule and site-wise analyses and with a quantitative characterization of runtime performance and deployment scalability. Finally, while the experimental results are illustrated using a subset of representative installations, the proposed architecture was evaluated in a multi-site operational context. The selected sites are presented to illustrate heterogeneous operating conditions rather than to exhaustively characterize large-scale photovoltaic portfolios.

As a direction for future work, a natural extension of the proposed architecture is to integrate probabilistic uncertainty metrics alongside point forecasts. Forecasting Units could output prediction intervals (e.g., via quantile regression, conformal prediction, or ensemble-based uncertainty estimates) while preserving the deterministic reasoning loop. These uncertainty estimates can be logged and reported in the explanation stage to support risk-aware operation, for instance, by triggering conservative actions or fallbacks when uncertainty is high. Future evaluations could complement RMSE/MAE with calibration and coverage metrics to assess operational reliability under uncertainty.

VIII. CONCLUSION

This paper proposed an agentic architectural framework for multi-site photovoltaic power forecasting that addresses operational requirements often overlooked by accuracy-driven approaches. The proposed architecture structures forecasting as an explicit reasoning process composed of perception, inference, validation, explanation, and action stages.

The motivation for this work arises from the observation that high predictive accuracy alone is insufficient for industrial deployment. Operational forecasting systems must ensure physical validity under imperfect data conditions, provide transparent and traceable decision processes, and operate autonomously across geographically distributed installations. These requirements cannot be reliably addressed through algorithmic optimization alone.

The proposed architecture addresses these challenges through explicit architectural design choices. Physical constraints are enforced as validation rules within the reasoning process, ensuring that forecasts remain physically coherent before operational use. Explainability is achieved by design through the explicit separation of reasoning stages and the systematic reporting of validation actions and fallback activations. Each photovoltaic installation is managed by an autonomous Forecasting Unit deployed as a microservice, enabling decentralized operation while preserving consistency across sites.

Experimental evaluation on operational photovoltaic systems showed that this architectural approach preserves the predictive performance of a validated hybrid forecasting pipeline while ensuring physical consistency and traceable decision paths. By orchestrating physical and data-driven models within a transparent ensemble framework, the system supports robust forecasting under heterogeneous data quality conditions without relying on a single predictive model.

Several limitations remain. As noted, the reasoning remains deterministic and rule-based, and the computational cost of the staged process has not been characterized.

Overall, this work highlights architectural design as a first-class concern in operational forecasting systems. Treating forecasting as an explicit and inspectable reasoning process provides a practical foundation for trustworthy deployment in industrial photovoltaic applications. The architectural principles presented here may be extended to other renewable energy forecasting contexts where reliability, transparency, and autonomous operation are essential.

REFERENCES

[1] C. J. Ejiyi, D. Cai, D. Thomas, S. Obiora, E. Osei-Mensah, C. Acen, F. O. Eze, F. Sam, Q. Zhang, and O. O. Bamisile, "Comprehensive review of artificial intelligence applications in renewable energy systems: Current implementations and emerging trends," *J. Big Data*, vol. 12, no. 1, p. 169, Jul. 2025, doi: 10.1186/s40537-025-01178-7.

[2] A. Mystakidis, P. Koukaras, N. Tsalikidis, D. Ioannidis, and C. Tjortjis, "Energy forecasting: A comprehensive review of techniques and technologies," *Energies*, vol. 17, no. 7, p. 1662, Mar. 2024, doi: 10.3390/en17071662.

[3] K. Barhmi, C. Heynen, S. Golroodbari, and W. van Sark, "A review of solar forecasting techniques and the role of artificial intelligence," *Solar*, vol. 4, no. 1, pp. 99–135, Feb. 2024, doi: 10.3390/solar4010005.

[4] J. Gaboitaolelwe, A. M. Zungeru, A. Yahya, C. K. Lebekwe, D. N. Vinod, and A. O. Salau, "Machine learning based solar photovoltaic power forecasting: A review and comparison," *IEEE Access*, vol. 11, pp. 40820–40845, 2023, doi: 10.1109/ACCESS.2023.3270041.

[5] M. Gupta, A. Arya, U. Varshney, J. Mittal, and A. Tomar, "A review of PV power forecasting using machine learning techniques," *Prog. Eng. Sci.*, vol. 2, no. 1, Mar. 2025, Art. no. 100058, doi: 10.1016/j.pes.2025.100058.

[6] M. Husein, E. J. Gago, B. Hasan, and M. C. Pegalajar, "Towards energy efficiency: A comprehensive review of deep learning-based photovoltaic power forecasting strategies," *Heliyon*, vol. 10, no. 13, Jul. 2024, Art. no. e33419, doi: 10.1016/j.heliyon.2024.e33419.

[7] X. Huang, T. Zhao, B. Huang, Z. Zhang, and M. Yue, "Advancing energy system optimization via data-centric task-oriented forecasting: An application in PV-battery operation," *Appl. Energy*, vol. 378, Jan. 2025, Art. no. 124753, doi: 10.1016/j.apenergy.2024.124753.

[8] U. Hijjawi, S. Lakshminarayana, T. Xu, G. P. M. Fierro, and M. Rahman, "A review of automated solar photovoltaic defect detection systems: Approaches, challenges, and future orientations," *Sol. Energy*, vol. 266, Dec. 2023, Art. no. 112186, doi: 10.1016/j.solener.2023.112186.

[9] W. Castillo-Rojas and J. Pastén Salinas, "Forecasting models applied in solar photovoltaic and wind energy: A systematic mapping study," *IEEE Access*, vol. 12, pp. 151092–151111, 2024, doi: 10.1109/ACCESS.2024.3471073.

[10] L. Yang, Z. Huang, X. Mo, and T. Luo, "Enhanced GAIN-based missing data imputation for a wind energy farm SCADA system," *Electronics*, vol. 14, no. 8, p. 1590, Apr. 2025, doi: 10.3390/electronics14081590.

[11] M. Tortora, F. Conte, G. Natrella, and P. Soda, "MATNet: Multi-level fusion and self-attention transformer-based model for multivariate multi-step day-ahead PV generation forecasting," 2023, *arXiv:2306.10356*.

[12] N. Rahimi, S. Park, W. Choi, B. Oh, S. Kim, Y.-H. Cho, S. Ahn, C. Chong, D. Kim, C. Jin, and D. Lee, "A comprehensive review on ensemble solar power forecasting algorithms," *J. Electr. Eng. Technol.*, vol. 18, no. 2, pp. 719–733, Jan. 2023, doi: 10.1007/s42835-023-01378-2.

[13] Z. Wang, H. Liu, S. Wu, N. Liu, X. Liu, Y. Hu, and Y. Fu, "Explainable time-varying directional representations for photovoltaic power generation forecasting," *J. Cleaner Prod.*, vol. 468, Aug. 2024, Art. no. 143056, doi: 10.1016/j.jclepro.2024.143056.

[14] I. Ioannou, S. Javaid, Y. Tan, and V. Vassiliou, "Autonomous reinforcement learning for intelligent and sustainable autonomous microgrid energy management," *Electronics*, vol. 14, no. 13, p. 2691, Jul. 2025, doi: 10.3390/electronics14132691.

[15] B. Huang and J. Wang, "Applications of physics-informed neural networks in power systems—A review," *IEEE Trans. Power Syst.*, vol. 38, no. 1, pp. 572–588, Jan. 2023, doi: 10.1109/TPWRS.2022.3162473.

[16] L. Sakli and S. B. Elghali, "PV forecasting with constrained ensemble learning and context-aware stacking," in *Proc. IEEE 37th Int. Conf. Tools Artif. Intell. (ICTAI)*, Athens, Greece, Nov. 2025, pp. 932–937, doi: 10.1109/ICTAI66417.2025.00136.

[17] I. Hassan, I. Alhamrouni, N. H. Azhan, S. Mekhilef, M. Seyedmahmoudian, S. I. Majid, and A. Stojcevski, "A predictive stacking framework for energy output in photovoltaic system with uncertainty estimates," *IEEE Access*, vol. 13, pp. 154046–154068, 2025, doi: 10.1109/ACCESS.2025.3604038.

[18] Z. Li, S. Liu, and G. Liu, "Improving power grid disaster resilience through predictive analytics based on data-driven photovoltaic power plant and distribution network power forecasting," *IEEE Access*, vol. 13, pp. 1–18, 2025, doi: 10.1109/ACCESS.2025.3611104.

[19] K. Kravari, M. Roussi, K. Ladomenou, A. Thysiadou, and M. Chalaris, "Intelligent agent orchestration for climate-resilient energy systems: Autonomous management of renewable integration and distributed resources," in *Proc. E3S Web Conf.*, vol. 638, 2025, p. 02009, doi: 10.1051/e3sconf/202563802009.

[20] S. Boudoudouh, M. Ouassaid, and M. Maâroufi, "Multi agent system in a distributed energy management of a multi sources system with a hybrid storage," in *Proc. 3rd Int. Renew. Sustain. Energy Conf. (IRSEC)*, Marrakech, Morocco, Dec. 2015, pp. 1–6, doi: 10.1109/IRSEC.2015.7454933.

[21] W. Fekih Hassen and M. Challouf, "Long short-term renewable energy sources prediction for grid-management systems based on stacking ensemble model," *Energies*, vol. 17, no. 13, p. 3145, Jun. 2024, doi: 10.3390/en17133145.

[22] A. T. Baelongandi, W. V. Kambale, A. M. Nzung, M. N. Bampende, D. K. Kadurha, and K. Kyamakya, "Towards a physics informed neural network approach for modeling and prediction of renewable energy production: Use case of solar cells," in *Proc. IEEE PES/IAS PowerAfrica*, Johannesburg, South Africa, Oct. 2024, pp. 01–05, doi: 10.1109/PowerAfrica61624.2024.10759416.

[23] J. Simeunovic, B. Schubnel, P.-J. Alet, P. Frossard, and R. Carrillo, "Physics informed graph neural networks for multi-site solar forecasting," in *Proc. 41st Eur. Photovoltaic Sol. Energy Conf. Exhib. (EU PVSEC)*, 2024, pp. 1–14, doi: 10.4229/EUPVSEC2024/4CV.1.25.

[24] M. Kuzlu, U. Cali, V. Sharma, and O. Guler, "Gaining insight into solar photovoltaic power generation forecasting utilizing explainable artificial intelligence tools," *IEEE Access*, vol. 8, pp. 187814–187823, 2020, doi: 10.1109/ACCESS.2020.3031477.

[25] R. M. Rizk-Allah, L. M. Abouelmagd, A. Darwish, V. Snasel, and A. E. Hassanien, "Explainable AI and optimized solar power generation forecasting model based on environmental conditions," *PLoS ONE*, vol. 19, no. 10, Oct. 2024, Art. no. e0308002, doi: 10.1371/journal.pone.0308002.

[26] Z. Yan, D. Zhenyuan, Y. Xu, and Z. Zhou, "Probabilistic PV power forecasting by a multi-modal method using GPT-agent to interpret weather conditions," in *Proc. IEEE 19th Conf. Ind. Electron. Appl. (ICIEA)*, Aug. 2024, pp. 1–6, doi: 10.1109/ICIEA61579.2024.10665072.

[27] M. K. Pasupuleti, *AI in Renewable Energy Management: Tools for Optimizing Sustainable Power Systems*. National Education Services, Oct. 2024, doi: 10.62311/nesx/rb978-81-979057-2-8.

[28] L. Ding, D. Wu, C. Zhang, Z. Chen, L. Ma, Z. Li, D. Liu, B. Han, and D. Ge, "Regional photovoltaic power prediction based on multi-agent federated learning and deep learning," in *Proc. 3rd Int. Conf. Energy, Power Electr. Technol. (ICEPET)*, May 2024, pp. 16–21, doi: 10.1109/ICEPET61938.2024.10627618.

[29] L. Ali, M. I. Azim, N. B. Ojha, J. Peters, V. Bhandari, A. Menon, J. Green, and S. M. Muyeen, "Integrating forecasting service and Gen2 blockchain into a local energy trading platform to promote sustainability goals," *IEEE Access*, vol. 12, pp. 2941–2964, 2024, doi: 10.1109/ACCESS.2023.3347432.

[30] M. Ahmed and M. R. Khan, "Artificial intelligence-enabled digital twins for energy efficiency in smart grids," *Rev. Appl. Sci. Technol.*, vol. 4, no. 2, pp. 580–615, Jun. 2025, doi: 10.63125/12kp9w74.

[31] S. Meisenbacher, B. Heidrich, T. Martin, R. Mikut, and V. Hagenmeyer, "AutoPV: Automated photovoltaic forecasts with limited information using an ensemble of pre-trained models," in *Proc. 14th ACM Int. Conf. Future Energy Syst.*, Orlando, FL, USA, Jun. 2023, pp. 386–414, doi: 10.1145/3575813.3597348.

[32] L. Sakli and S. Ben Elghali, "Towards an agentic and explainable runtime architecture for PV forecast delivery," in *Proc. IEEE Int. Conf. Complex Syst. Intell. Comput. (ICCSIC)*, Mar. 2026, pp. 279–285, doi: 10.1109/ICCSIC69301.2026.11521689.

[33] O. Kirat, A. Çiçek, and T. Yerlikaya, "A new artificial intelligence-based system for optimal electricity arbitrage of a second-life battery station in day-ahead markets," *Appl. Sci.*, vol. 14, no. 21, p. 10032, Nov. 2024.

[34] O. Kirat and A. Çiçek, "Artificial intelligence based optimal operation of green hydrogen refueling," in *Proc. 32nd Telecommun. Forum (TELFOR)*, Belgrade, Serbia, Nov. 2024, pp. 1–4, doi: 10.1109/TELFOR63250.2024.10819143.

[35] M. Mansouri, K. Attouri, and S. S. Refaat, "Multimodal learning techniques for time series forecasting in renewable energy systems: A comprehensive survey," *IEEE Access*, vol. 13, pp. 151970–151992, 2025, doi: 10.1109/ACCESS.2025.3602914.

LEÏLA SAKLI (Member, IEEE) received the Engineering degree in computer science and the M.Sc. degree in control and signal processing from the University of Tunis El-Manar, Tunisia, and the Ph.D. degree in computer science from Aix-Marseille Université, France, in 2016.

She is currently a Postdoctoral Researcher with the Laboratoire d'Informatique et des Systémes (LIS, CNRS 7020), Aix-Marseille Université. Her recent work focuses on explainable forecasting and resilient supervision for smart grids and photovoltaic energy systems. She has contributed to several European Horizon projects, including VPP4Islands, MASTERPIECE, OMRES, and FlexCHESS. Her research interests include data-driven modeling, AI-based forecasting, time-series analysis, risk and vulnerability assessment, decision-support systems, and multi-agent approaches for complex distributed systems.

SEIFEDDINE BEN ELGHALI received the Engineering degree in electrical engineering from the National School of Engineers of Tunis, Tunisia, in 2005, the M.Sc. degree in science and technology from ENSIP, France, in 2006, the Ph.D. degree in electrical engineering from the University of Western Brittany, France, in 2009, and the Habilitation to Supervise Research (H.D.R.) degree from Aix-Marseille University, France, in 2023.

He is currently an Associate Professor with the Faculty of Sciences, Aix-Marseille University, and a member of the Laboratoire d'Informatique et des Systémes (LIS, CNRS 7020). His research interests include modeling, control, and optimization of multi-source energy systems, virtual power plants, energy storage systems, and renewable energy integration. He has coordinated several European research projects and is the author or co-author of numerous journals and conference publications in these areas.

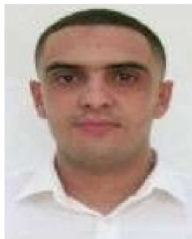

YACINE MERRAD received the Ph.D. degree in engineering from International Islamic University Malaysia, in 2022. He is currently a Postdoctoral Researcher with the Laboratoire d'Informatique et des Systémes (LIS, CNRS 7020), Aix-Marseille Université, France. He has contributed to several European Horizon projects, including VPP4Islands, MASTERPIECE, and FlexCHESS. His research interests include decentralized smart-grid management, multi-agent, convex, and stochastic optimization for energy systems, blockchain-based electricity markets, privacy-preserving energy trading, and the IoT-enabled energy monitoring.