Energy and AI 25 (2026) 100827

Contents lists available at ScienceDirect

Energy and AI

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

# Price-aware artificial intelligence for market-aligned photovoltaic power nowcasting under quarter-hourly electricity settlement

Saloni Dhingra [a], Sara Carcangiu [b],*, Binh Nam Nguyen [c], Augusto Montisci [b], Giancarlo Storti Gajani [a], Sonia Leva [c], Emanuele Ogliari [c]

[a] Politecnico di Milano, DEIB, Piazza Leonardo da Vinci 32, 20133, Milano, Italy
[b] Università degli Studi di Cagliari, DIEE, Via Marengo 2, 09123, Cagliari, Italy
[c] Politecnico di Milano, DENG, Via Lambruschini 8, 20156, Milano, Italy

## HIGHLIGHTS

- A price-aware framework links photovoltaic nowcasting to imbalance costs.
- Forecast errors are evaluated through a monetary-impact indicator.
- An asymmetric loss embeds market costs into model training.
- Real photovoltaic data are tested over a full unseen year.
- Monetary impact is reduced by up to 54.72% against persistence.

## GRAPHICAL ABSTRACT

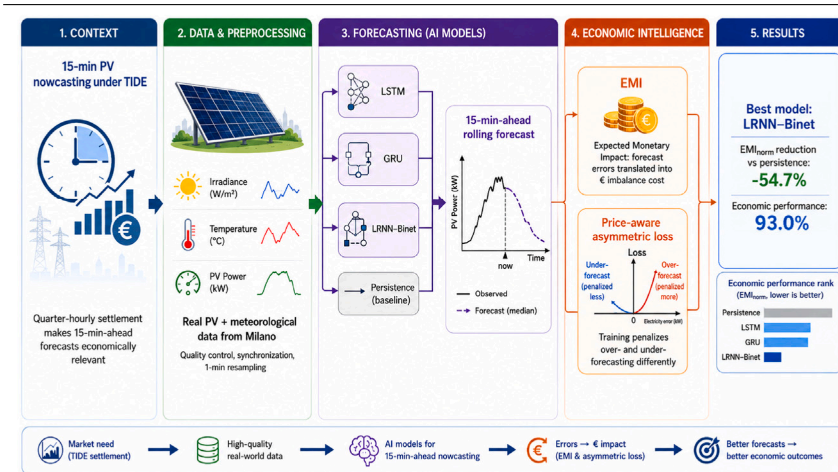


## ARTICLE INFO

*Keywords:*
Photovoltaic power forecasting
Ultra-short-term forecasting
Nowcasting
Expected monetary impact
Asymmetric loss function
Electricity markets

## ABSTRACT

The increasing penetration of distributed photovoltaic generation and the move toward quarter-hourly imbalance settlement, exemplified by the Italian Testo Integrato del Dispacciamento Elettrico reform, make ultra-short-term power forecasting a critical operational and economic task. Although machine-learning forecasters have improved statistical accuracy, they are commonly trained and evaluated with symmetric error indicators that do not reflect the asymmetric monetary effects of over- and under-forecasting. This paper addresses this gap by proposing a decision-oriented framework for fifteen-minute-ahead photovoltaic power nowcasting that embeds market asymmetry in both evaluation and learning. The framework introduces two methodological contributions: an Expected Monetary Impact indicator, which translates forecast deviations into monetary losses at quarter-hour settlement resolution, and a differentiable price-aware asymmetric loss function, which penalizes positive and negative forecast errors according to their market costs. The Italian reform is used as a realistic case study, while the formulation can be transferred to other markets with short settlement intervals and asymmetric imbalance prices. Long Short-Term Memory, Gated Recurrent Unit, and a Locally Recurrent Neural Network trained through the Binet formalism are evaluated as alternative forecasters against persistence under a rolling out-of-sample protocol using real measurements from a photovoltaic facility

* Corresponding author.
*E-mail address:* sara.carcangiu@unica.it (S. Carcangiu).

https://doi.org/10.1016/j.egyai.2026.100827
Received 20 May 2026; Received in revised form 12 June 2026; Accepted 23 June 2026
Available online 30 June 2026

in Milan, with 2023–2024 used for training and 2025 for testing. The learning-based models outperform persistence: Long Short-Term Memory and Gated Recurrent Unit reduce the normalized monetary impact by 18.22% and 19.39%, respectively, while the locally recurrent model achieves a 54.72% reduction. These results show that market-aligned forecasting objectives provide greater operational value than conventional symmetric accuracy indicators.

## 1. Introduction

The increasing penetration of distributed photovoltaic generation is reshaping short-term operation and decision-making in power systems. System operators, aggregators, energy communities, and prosumers are increasingly required to manage generation schedules, flexibility resources, and balancing actions at finer temporal resolution [1]. This need is particularly relevant for ultra-short-term forecasting horizons, typically ranging from a few minutes to one hour, where forecast errors directly affect imbalance costs, dispatch efficiency, reserve activation, and storage operation. In Italy, the Testo Integrato del Dispacciamento Elettrico (TIDE) reform shifts balancing and settlement processes from hourly to quarter-hourly intervals, making fifteen-minute-ahead photovoltaic nowcasting both an operational requirement and an economic opportunity [2].

Although this regulatory reform provides the case study considered in this paper, the underlying problem is broader than the Italian context. Electricity markets are increasingly moving toward shorter settlement intervals, higher renewable penetration, and stronger exposure of distributed energy resources to imbalance prices. This trend is not specific to Italy. Under Regulation (EU) 2019/943, the imbalance settlement period has been harmonized toward fifteen-minute resolution across European scheduling areas, with the explicit aim of better reflecting the real-time value of energy and supporting intraday trading [3]. A comparable, and even finer, shortening of the settlement interval has been adopted outside Europe: the Australian National Electricity Market moved from thirty-minute to five-minute settlement in 2021. The decision-oriented formulation proposed here is therefore relevant well beyond the Italian case, wherever settlement is performed at sub-hourly resolution under asymmetric or time-varying imbalance prices. In such settings, the value of a forecast depends not only on its statistical accuracy, but also on the economic consequences of the resulting scheduling deviations.

Despite significant progress in renewable-energy forecasting, a gap remains between statistical accuracy and operational value. Recent reviews on machine learning for energy applications show that most forecasting studies still focus primarily on conventional statistical performance indicators [4]. In photovoltaic forecasting, indicators such as root mean square error and mean absolute error are widely used because they provide compact and interpretable measures of pointwise accuracy. However, these indicators are symmetric by construction and therefore do not distinguish between over-forecasting and under-forecasting. In electricity markets, this distinction is relevant because positive and negative forecast deviations may be settled at different prices and may trigger different operational actions. As a result, two models with comparable statistical errors may produce substantially different monetary impacts.

This limitation becomes more important as the settlement interval decreases. At quarter-hourly resolution, intra-hour ramps and fast irradiance variations are less likely to be averaged out than in hourly evaluations. Forecasting errors occurring during ramps or economically critical intervals can therefore have a disproportionate effect on the final imbalance cost. Studies on forecast-error costs for photovoltaic prosumers confirm that the economic value of forecasting depends not only on the magnitude of the error, but also on its timing, direction, and market context [5]. This makes quarter-hourly photovoltaic nowcasting a decision-oriented task, in which the timing and sign of the error become as relevant as its magnitude.

The literature on photovoltaic forecasting includes several families of methods. Persistence remains a mandatory reference for ultra-short-term horizons because of its simplicity and strong performance when the forecast lead time is short [6]. Hybrid clear-sky and data-driven methods have been proposed to improve photovoltaic forecasts under changing irradiance conditions [7]. Classical statistical models, such as autoregressive integrated moving average models, have also been used for short-term solar-power prediction [8]. Exponential smoothing provides another lightweight statistical alternative when limited computational complexity is required [9]. Machine-learning methods, including support vector regression, random forests, gradient boosting, and nearest-neighbor approaches, have further improved the ability to capture nonlinear dependencies in photovoltaic power time series [10].

Deep-learning architectures are now widely adopted in photovoltaic and solar irradiance forecasting. Long Short-Term Memory networks are commonly used as recurrent models for time-series forecasting because their gating mechanisms support the learning of temporal dependencies over multiple time scales [11]. Gated Recurrent Unit networks provide a more compact recurrent structure and are often considered a computationally efficient alternative to Long Short-Term Memory models [12]. Convolutional and hybrid convolutional-recurrent models have been proposed to extract temporal or spatio-temporal features from irradiance, sky-image, and photovoltaic data [13]. Transformer-based architectures have also recently been investigated for photovoltaic forecasting, although their effectiveness depends on data availability, regularization strategy, and computational requirements [14]. These studies demonstrate the relevance of machine learning for photovoltaic nowcasting, but they also highlight the need to assess model performance in a way that reflects market operation rather than statistical accuracy alone.

Within this context, recurrent architectures remain particularly relevant for ultra-short-term photovoltaic forecasting because the recent power trajectory contains information on ramps, cloud passages, and short-term irradiance variability. Conventional recurrent models, such as Long Short-Term Memory and Gated Recurrent Unit networks, provide strong and well-established benchmarks for this task. Locally Recurrent Neural Networks offer a different recurrent paradigm, in which temporal dynamics are embedded locally within each neuron through feedback structures. When trained through the Binet formalism, this architecture provides an explicit-gradient mechanism and direct control of the stability [15]. Comparing this Binet-based locally recurrent model with conventional recurrent benchmarks under the same market-aligned evaluation framework allows one to assess whether the local-feedback formulation provides additional operational value for quarter-hourly photovoltaic nowcasting. In addition, all learning-based forecasters are compared against persistence, which represents the standard baseline for ultra-short-term photovoltaic forecasting.

Table 1 summarizes the main research gaps addressed in this work and the corresponding methodological choices adopted in the paper.

Building on these gaps, this paper proposes and validates a decision-oriented artificial intelligence framework for fifteen-minute-ahead photovoltaic power nowcasting under quarter-hourly electricity settlement. The framework is tested on real photovoltaic measurements collected in Milan under realistic operating conditions. Historical data from 2023–2024 are used for training, while the full year 2025 is retained as an unseen chronological test period. All preprocessing steps are based only on past information, in order to avoid data leakage and reproduce a realistic deployment setting. The specific contributions are as follows:

**Table 1**
Research gaps in photovoltaic nowcasting and corresponding contributions of this work.

| Research gap | Representative evidence from the literature | How this paper addresses the gap |
|---|---|---|
| Evaluation granularity is often misaligned with market settlement intervals. | Quarter-hourly balancing and settlement are introduced in the Italian Testo Integrato del Dispacciamento Elettrico reform [2]. | Forecasts are evaluated at the native fifteen-minute cadence, consistently with quarter-hourly dispatch, balancing, and settlement. |
| Conventional error indicators do not account for price asymmetry and operational costs. | Reviews on artificial intelligence for energy systems highlight the dominance of statistical indicators [4], while studies on photovoltaic prosumers quantify the monetary cost of forecast errors [5]. | The Expected Monetary Impact indicator translates signed forecast deviations into monetary losses and separates over- and under-forecasting contributions. |
| Symmetric learning objectives are not aligned with asymmetric imbalance penalties. | Photovoltaic forecasting studies commonly optimize symmetric errors such as root mean square error and mean absolute error [16], although forecast-error costs depend on the sign and market value of the deviation [17]. | A differentiable price-aware asymmetric loss is introduced to embed the different costs of positive and negative errors into model training. |
| Recurrent forecasting models are often compared only through statistical accuracy. | Long Short-Term Memory models are established recurrent benchmarks for renewable time-series forecasting [11]. Gated Recurrent Unit models are commonly used as compact recurrent alternatives [12]. Transformer-based models have also been recently benchmarked for photovoltaic forecasting [14]. | A Binet-based Locally Recurrent Neural Network is compared with conventional recurrent benchmarks using both statistical and economic indicators. |
| The link between model selection and operational economic value remains underexplored. | Recent photovoltaic forecasting studies compare model complexity and accuracy [18], but the monetary value of the resulting forecasting improvement is less frequently assessed [5]. | Long Short-Term Memory, Gated Recurrent Unit, and the Binet-based Locally Recurrent Neural Network are all benchmarked against persistence at quarter-hourly settlement resolution. |

1. A decision-oriented framework is proposed for fifteen-minute-ahead photovoltaic power nowcasting under quarter-hourly electricity settlement. The framework explicitly links forecasting performance to imbalance-cost reduction and is therefore suitable for operators, aggregators, and energy communities that must take short-term decisions under asymmetric market penalties.
2. An Expected Monetary Impact indicator is introduced to translate signed forecast deviations into monetary losses at the native settlement cadence. The indicator includes a normalized form, enabling comparison across systems of different sizes, and a decomposition into over- and under-forecasting components, enabling interpretation of the economic effect of error direction and forecast bias.
3. A differentiable price-aware asymmetric loss function is proposed to incorporate the different costs of over- and under-forecasting directly into model training. This aligns the learning objective with the economic structure of the settlement mechanism, rather than optimizing only conventional symmetric error indicators.
4. A systematic comparison is performed between a Locally Recurrent Neural Network trained through the Binet formalism and two conventional recurrent benchmarks, namely Long Short-Term Memory and Gated Recurrent Unit. All three learning-based forecasters are then evaluated against the persistence baseline, which represents the standard reference for ultra-short-term photovoltaic nowcasting. The comparison is carried out under identical data availability, forecasting horizon, rolling out-of-sample protocol, and statistical and economic evaluation criteria.
5. A full-year rolling out-of-sample validation is carried out on real photovoltaic measurements. The analysis shows how model ranking and practical interpretation change when forecasting performance is assessed in monetary terms, making it possible to distinguish between statistical accuracy and operational value.

The remainder of this paper is organized as follows. Section 2 presents the forecasting framework and the candidate recurrent architectures. Section 3 introduces the Expected Monetary Impact indicator and the price-aware asymmetric loss function. Section 4 describes the dataset, preprocessing procedure, rolling forecasting strategy, and evaluation protocol. Section 5 reports and discusses the results on the Milan case study. Finally, Section 6 summarizes the main findings, practical implications, limitations, and future research directions.

## 2. Methodology

Fig. 1 summarizes the forecasting workflow adopted in this study. Measured photovoltaic and meteorological data are first processed through quality control, synchronization, and one-minute resampling. The resulting time series are then used in a rolling fifteen-minute-ahead forecasting setup, where each model generates a prediction using only information available up to the forecasting time.

The candidate forecasters considered in the comparison are persistence, Long Short-Term Memory, Gated Recurrent Unit, and the Locally Recurrent Neural Network trained through the Binet formalism. Persistence is used as the reference baseline for ultra-short-term photovoltaic nowcasting, while the three learning-based architectures are evaluated under the same data availability, forecasting horizon, rolling out-of-sample protocol, and statistical and economic performance indicators.

The forecasts are assessed at native quarter-hourly settlement resolution using conventional statistical indicators and the proposed economic indicators. In addition, two learning objectives are considered for the neural models: a standard mean-square-error objective and the proposed price-aware asymmetric loss introduced in Section 3. This structure allows the comparison to distinguish between conventional statistical accuracy and operational economic value.

### 2.1. Long Short-Term Memory (LSTM)

Long Short-Term Memory (LSTM) networks extend classical recurrent neural networks by introducing a memory cell and gating mechanisms that stabilize gradient propagation and capture both short- and

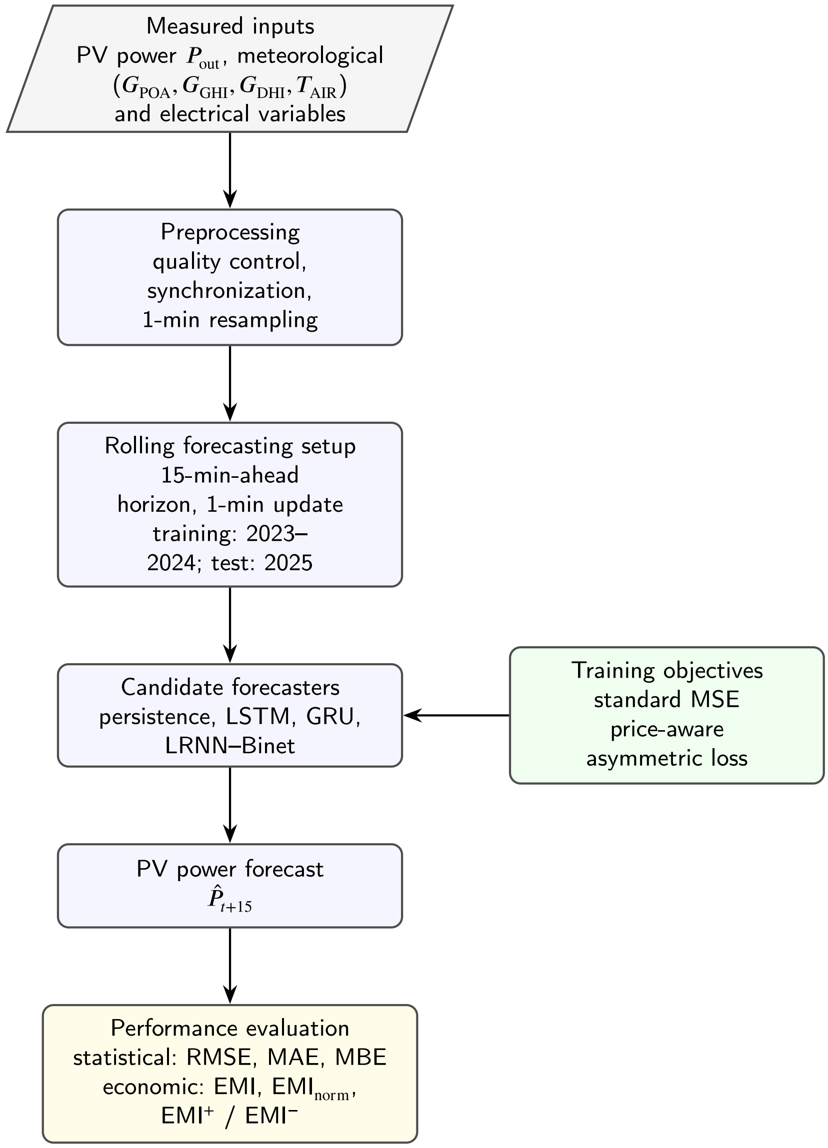


**Fig. 1.** Forecasting workflow adopted in this study. Measured photovoltaic power, meteorological, and electrical variables are processed through quality control, synchronization, and one-minute resampling, and then used in a rolling fifteen-minute-ahead forecasting setup with candidate models (persistence, LSTM, GRU, LRNN–Binet). Forecasts are evaluated at native fifteen-minute resolution through both statistical (RMSE, MAE, MBE) and economic (EMI, $\text{EMI}_{\text{norm}}$, over/under decomposition) indicators. The neural models are trained using either a standard mean-square-error loss or the proposed price-aware asymmetric loss.

long-term temporal dependencies [19,20]. This is particularly suited to quarter-hourly PV forecasting, where fast intra-hour variability coexists with slower diurnal patterns.

Let $x_t \in \mathbb{R}^d$ denote the input at time $t$. Using a context window of length $L$, the LSTM maps $\{x_{t-L+1}, \dots, x_t\}$ to a 15-minute-ahead prediction:

$$\hat{y}_{t+h} = f_\theta(x_{t-L+1:t}), \qquad h = 15 \text{ min}. \tag{1}$$

At each step $\tau$, the hidden state $\mathbf{h}_\tau \in \mathbb{R}^H$ and cell state $\mathbf{c}_\tau \in \mathbb{R}^H$ evolve as:

$$\mathbf{f}_\tau = \sigma(\mathbf{W}_f \mathbf{x}_\tau + \mathbf{U}_f \mathbf{h}_{\tau-1} + \mathbf{b}_f), \tag{2}$$

$$\mathbf{i}_\tau = \sigma(\mathbf{W}_i \mathbf{x}_\tau + \mathbf{U}_i \mathbf{h}_{\tau-1} + \mathbf{b}_i), \tag{3}$$

$$\tilde{\mathbf{c}}_\tau = \tanh(\mathbf{W}_c \mathbf{x}_\tau + \mathbf{U}_c \mathbf{h}_{\tau-1} + \mathbf{b}_c), \tag{4}$$

$$\mathbf{c}_\tau = \mathbf{f}_\tau \odot \mathbf{c}_{\tau-1} + \mathbf{i}_\tau \odot \tilde{\mathbf{c}}_\tau, \tag{5}$$

$$\mathbf{o}_\tau = \sigma(\mathbf{W}_o \mathbf{x}_\tau + \mathbf{U}_o \mathbf{h}_{\tau-1} + \mathbf{b}_o), \tag{6}$$

$$\mathbf{h}_\tau = \mathbf{o}_\tau \odot \tanh(\mathbf{c}_\tau), \tag{7}$$

where $\sigma(\cdot)$ is the logistic sigmoid and $\odot$ denotes elementwise multiplication.

The gating structure regulates information retention, update, and exposure, while the additive cell update ensures stable gradients and enables modeling of both slow trends and rapid PV ramps.

The LSTM is embedded in the same rolling forecasting framework adopted for all models; architecture-specific training details are provided in Section 4.

### 2.2. Gated recurrent unit (GRU)

The Gated Recurrent Unit (GRU) is a simplified recurrent architecture that reduces LSTM complexity by using a single hidden state and fewer gates [21], resulting in lower computational cost.

Using the same input representation, the GRU updates its hidden state $\mathbf{h}_\tau \in \mathbb{R}^H$ as:

$$\mathbf{z}_\tau = \sigma(\mathbf{W}_z \mathbf{x}_\tau + \mathbf{U}_z \mathbf{h}_{\tau-1} + \mathbf{b}_z), \tag{8}$$

$$\mathbf{r}_\tau = \sigma(\mathbf{W}_r \mathbf{x}_\tau + \mathbf{U}_r \mathbf{h}_{\tau-1} + \mathbf{b}_r), \tag{9}$$

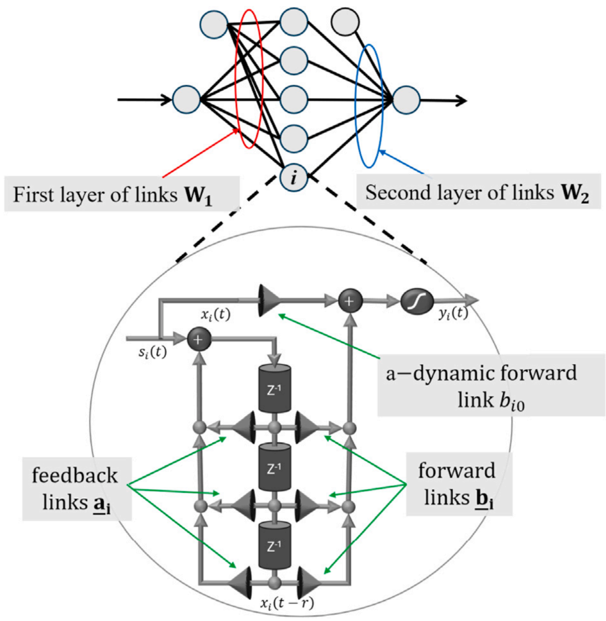


**Fig. 2.** Internal structure of an LRNN neuron with feed-forward FIR coefficients $b_i$, a-dynamic feed-forward link $b_{i0}$, and locally recurrent IIR feedback coefficients $a_i$. Top: global structure with input weights $W_1$, output weights $W_2$, and the $i$th neuron highlighted. Bottom: internal structure of the $i$th neuron.

$$\tilde{\mathbf{h}}_\tau = \tanh(\mathbf{W}_h\mathbf{x}_\tau + \mathbf{U}_h(\mathbf{r}_\tau \odot \mathbf{h}_{\tau-1}) + \mathbf{b}_h), \tag{10}$$

$$\mathbf{h}_\tau = (\mathbf{1} - \mathbf{z}_\tau) \odot \mathbf{h}_{\tau-1} + \mathbf{z}_\tau \odot \tilde{\mathbf{h}}_\tau, \tag{11}$$

where $\mathbf{z}_\tau$ and $\mathbf{r}_\tau$ are the update and reset gates.

This compact formulation efficiently captures short-term temporal dependencies, which dominate ultra-short-term PV forecasting.

As with LSTM, GRU is trained and evaluated under the same rolling framework, with architecture-specific adaptation strategies detailed in Section 4.3.

### 2.3. Locally Recurrent Neural Network (LRNN) with Binet formalism

The Locally Recurrent Neural Network (LRNN) extends the classical multilayer perceptron by introducing a dynamic within each hidden neuron, thus combining static nonlinear mapping with adaptive short-term dynamics.

Unlike gated recurrent architectures, temporal behavior is modeled directly at the neuron level through a compact ARMA structure. Each LRNN neuron behaves as an ARMA-like element composed of a feed-forward finite impulse response (FIR) branch and a recurrent infinite impulse response (IIR) branch. This structure enables the model to capture short- and medium-term temporal dependencies through internal states, while longer-term behavior is encoded in the trainable parameters.

Fig. 2 shows the overall architecture. The upper diagram illustrates the network structure, where input signals are projected through $\mathbf{W}_1$ onto the hidden LRNN layer and the resulting neuron outputs are combined by the output weights $\mathbf{W}_2$. The lower diagram details the internal organization of one LRNN neuron, composed of local feedback loops ($\mathbf{a}_i$), feedforward coefficients ($\mathbf{b}_i$), and an instantaneous dynamic link $b_{i0}$ acting through the nonlinear activation $f(\cdot)$. This mixed FIR/IIR configuration enables the neuron to model both dynamic and static components of the input sequence, forming a compact yet expressive structure suitable for time-series forecasting.

The state evolution of the $i$th neuron is given by:

$$x_i(t) = \mathbf{a}_i^\top \mathbf{x}_i(t-1:t-r) + s_i(t) \tag{12}$$

$$y_i(t) = f\left(b_{i0}\, s_i(t) + \mathbf{b}_i^\top \mathbf{x}_i(t-1:t-r)\right) \tag{13}$$

where $\mathbf{x}_i(t-1:t-r) = [\,x_i(t-1), \dots, x_i(t-r)\,]^\top$ collects the delayed internal states, $\mathbf{a}_i = [\,a_{i1}, \dots, a_{ir}\,]^\top$ and $\mathbf{b}_i = [\,b_{i1}, \dots, b_{ir}\,]^\top$ are the vectors of feedback and feed-forward coefficients, respectively, $b_{i0}$ is the instantaneous feed-forward link, and $f(\cdot)$ is the nonlinear activation (tanh). The parameter $r$ defines the memory depth of the neuron.

Eq. (12) defines a locally recurrent state update whose impulse response can be analytically expanded as follows.

Following [15], the impulse response of each neuron can be written as a linear combination of exponential components:

$$h_i(t) = \sum_{m=1}^{r} C_{im}\, z_{im}^t, \qquad z^r - a_{i1} z^{r-1} - \cdots - a_{ir} = 0 \tag{14}$$

where $z_{im}$ are the roots of the characteristic polynomial. Stability is ensured by enforcing $|z_{im}| < 1$, which directly constrains the local dynamics during training.

Continuous adaptation is achieved through a *sliding window* mechanism that updates the network on recent data. At each step, the model maps a sequence of past samples to a future target shifted by the prediction horizon, so that older samples progressively lose influence as the window advances. After each update, the parameters are frozen and the network produces the next-step forecast in inference mode. Fig. 3 illustrates this concept; the architecture-specific training choices adopted in this study are detailed in Section 4.3.

The gradient of the loss function with respect to the network parameters can be obtained in closed form, yielding the explicit-gradient learning rule proposed in [15]. For each neuron $i$, the gradient of the loss $J$ with respect to its dynamic bases $\mathbf{z}_i = [z_{i1}, \dots, z_{ir}]^T$ reads:

$$\frac{\partial J}{\partial \mathbf{z}_i} = \sum_t [u(t) - d(t)]\, v_i\, f_i'(t) \sum_{j=1}^{r} b_{ij}\, [s_i * z_i^{t-j-1}] \tag{15}$$

where $u(t)$ and $d(t)$ denote the network and target outputs, respectively, $f_i'(t)$ is the derivative of the activation function, and $v_i$ is the $i$th element of the output weight vector $W_2$. This expression generalizes the scalar formulation presented in [15] to the vector representation of the neuron local parameters.

The explicit-gradient update avoids backpropagation through time. Following the *coplanarity paradigm* [22], the output-layer weights $W_2$ are computed analytically as the regression hyperplane fitting the transformed samples in the product space $\{\text{features}\} \times \{\text{output}\}$. Therefore, nonlinear dynamic learning, applied to $\Gamma = \{W_1, \mathbf{a}_i, \mathbf{b}_i, \mathbf{z}_i\}$, is separated from linear output regression. This separation improves convergence while preserving the stability of the local recursions.

The complete training protocol and hyperparameter settings are reported in Section 4.3.

## 3. Economic metric and asymmetric training loss at 15-minute cadence

Under TIDE, positive and negative forecast errors lead to distinct monetary consequences, as outlined in Section 1. To align both evaluation and training with this asymmetry, we propose a twofold modification to the standard framework:

1. A market-consistent **evaluation metric**, the *Expected Monetary Impact* (EMI), which expresses performance directly in euros at the native 15-minute cadence.
2. A differentiable, **asymmetric loss function** used in training, derived from economic principles and designed to emphasize costly forecast errors while preserving smooth optimization.

This dual approach ensures that both the evaluation and the learning phases are aligned with the operational and financial logic of the TIDE market.

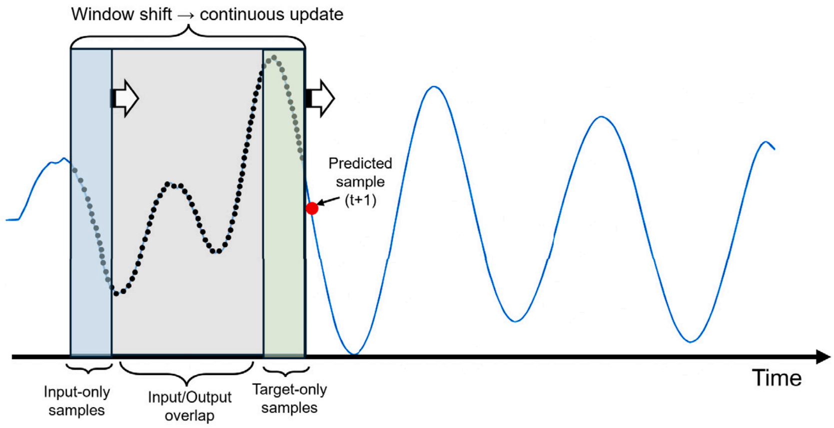


**Fig. 3.** Sliding-window mapping of input and target samples used for LRNN training. Each window maps a sequence of past input samples to a future sample shifted in time; intermediate samples act alternately as input or output in successive updates, while the network predicts the next sample immediately following the window.

### 3.1. Expected monetary impact (EMI) for model evaluation

Let $E_i$ be the measured PV energy (MWh) over the $i$th 15-minute settlement interval, $\hat{E}_i$ the predicted energy, and $e_i = \hat{E}_i - E_i$ the signed forecast error. When only power measurements are available at 1-minute resolution, they are integrated over the settlement interval to obtain $E_i$.

We denote by $c_i^+$ and $c_i^-$ the unit cost [€/MWh] associated with *over-forecasting* ($e_i > 0$) and *under-forecasting* ($e_i < 0$), respectively (the superscript matches the sign of the error that triggers the penalty). For the case study presented here, we assume:

$$c_i^- = 100 \text{ €/MWh}, \qquad c_i^+ = 150 \text{ €/MWh},$$

Under the considered TIDE-aligned remuneration scheme, *over-forecasting* requires purchasing the missing energy at the selling/imbalance price (150 €/MWh), whereas *under-forecasting* implies that the surplus energy is remunerated at a lower price, yielding an effective penalty of 100 €/MWh with respect to the scheduled remuneration. The instantaneous monetary loss at time step $i$ is therefore:

$$S_i = c_i^+ \max(e_i, 0) + c_i^- \max(-e_i, 0) \tag{16}$$

and the total economic impact over a set of $N$ intervals $\mathcal{I}$ is:

$$\text{EMI} = \sum_{i\in\mathcal{I}} S_i, \qquad \text{EMI}_{\text{norm}} = \frac{\sum_{i\in\mathcal{I}} S_i}{\sum_{i\in\mathcal{I}} E_i} \quad [\text{€/MWh}] \tag{17}$$

where $\text{EMI}_{\text{norm}}$ represents the loss per unit of energy produced. A lower EMI indicates a better-performing model not only in terms of accuracy but also in its ability to minimize monetary penalties.

For interpretability, we also report the decomposition of EMI into the contributions associated with positive and negative errors:

$$\begin{cases} \text{EMI} &= \text{EMI}^+ + \text{EMI}^- \\ \text{EMI}^+ &= \sum_{i\in\mathcal{I}} c_i^+ \max(e_i, 0) \\ \text{EMI}^- &= \sum_{i\in\mathcal{I}} c_i^- \max(-e_i, 0) \end{cases} \tag{18}$$

With this convention, $\text{EMI}^+$ collects the cost contribution of positive forecast errors, i.e., over-forecasting events, whereas $\text{EMI}^-$ collects the contribution of negative forecast errors, i.e., under-forecasting events. This formulation allows the evaluation to remain agnostic to the forecasting architecture (GRU, LSTM, or LRNN–Binet) while providing a common economic baseline. It also naturally supports asymmetric weighting: for example, when imbalance prices are time-varying, $c_i^\pm$ can be drawn directly from market data, enabling case-specific sensitivity studies. As a result, model comparisons reflect their real operational value under TIDE rather than purely statistical skill.

### 3.2. Motivation and definition of the asymmetric loss function

While EMI provides an ex-post evaluation metric, the training phase also requires an objective function consistent with market asymmetry. A symmetric loss, such as mean square error or mean absolute error, assigns the same penalty to positive and negative forecast errors. This implicitly assumes that forecasting above or below the actual production has the same operational consequence. In balancing markets, however, the monetary effect of a forecast error depends on both its magnitude and its sign. In the TIDE-aligned scenario considered here, over-forecasting is more expensive than under-forecasting, namely $c_i^+ > c_i^-$. Therefore, the training objective should be differentiable, sign-sensitive, and numerically stable for gradient-based optimization.

For training purposes, the loss is applied to the normalized signed power error. Let $P_{t+h}$ be the measured photovoltaic power at the target time $t+h$, and let $\hat{P}_{t+h}$ be the corresponding forecast issued at time $t$, with $h = 15$ min. The normalized error is defined as

$$\tilde{e}_i = \frac{\hat{P}_{t+h} - P_{t+h}}{P_{\text{nom}}}, \tag{19}$$

where $P_{\text{nom}}$ is the nominal power of the photovoltaic module. This normalization makes the argument of the logistic function dimensionless and is consistent with the scaling adopted in the neural-network training. With this convention, $\tilde{e}_i > 0$ corresponds to over-forecasting, whereas $\tilde{e}_i < 0$ corresponds to under-forecasting.

Let

$$\sigma(x) = \frac{1}{1+\exp(-x)} \tag{20}$$

be the logistic function. The proposed economic loss is

$$\ell_i(\tilde{e}_i) = \exp\left[a_i\sigma(\tilde{e}_i + \delta) + b_i\sigma(-\tilde{e}_i + \gamma)\right] - \exp\left[a_i\sigma(\delta) + b_i\sigma(\gamma)\right]. \tag{21}$$

The coefficients $a_i$ and $b_i$ are the scaled economic weights associated with positive and negative forecast errors, respectively, while $\delta$ and $\gamma$ are shift terms introduced to make the loss minimum occur at $\tilde{e}_i = 0$. The subtraction of the constant term enforces $\ell_i(0) = 0$, so that the loss is referenced to the zero-error condition.

The scaled economic weights are directly linked to the corresponding market costs:

$$a_i = \lambda c_i^+, \qquad b_i = \lambda c_i^-, \tag{22}$$

where $\lambda$ is a normalization factor selected so that the largest scaled weight satisfies

$$\max(a_i, b_i) \le 0.25. \tag{23}$$

This upper bound preserves the cost ratio $c_i^+/c_i^-$, ensures the validity condition $a_i, b_i \le 1/4$ required by the analytical shift-term derivation, and prevents excessive curvature of the exponential loss during optimization. In this sense, the loss is price-aware: increasing the relative cost of one error direction steepens the corresponding branch of the loss.

The shift terms $\delta$ and $\gamma$ are computed by enforcing zero gradient at the origin:

$$\left.\frac{\partial \ell_i}{\partial \tilde{e}_i}\right|_{\tilde{e}_i=0} = 0. \tag{24}$$

This condition gives

$$a_i \sigma'(\delta) = b_i \sigma'(\gamma), \tag{25}$$

where

$$\sigma'(x) = \sigma(x)\left[1 - \sigma(x)\right]. \tag{26}$$

Solving Eq. (25) for the two shift terms yields

$$\exp(-\gamma) = \frac{-(2a_i - 1) + \sqrt{(2a_i - 1)^2 - 4a_i^2}}{2a_i}, \tag{27}$$

and

$$\exp(-\delta) = \frac{-(2b_i - 1) + \sqrt{(2b_i - 1)^2 - 4b_i^2}}{2b_i}. \tag{28}$$

These expressions are valid when $a_i, b_i \le 1/4$. The bound in Eq. (23) therefore guarantees that the square-root terms in Eqs. (27)–(28) are real and that the shift terms are well defined.

In the present case study, the market-cost coefficients are constant over time, with $c_i^+ = 150$ €/MWh and $c_i^- = 100$ €/MWh. Consequently, $\lambda$, $a_i$, $b_i$, $\delta$, and $\gamma$ are computed once and kept fixed throughout training and rolling deployment. The shift terms are therefore not dynamically updated; they only encode the assumed cost asymmetry between over- and under-forecasting. If time-varying imbalance prices are used, the same formulation can be applied with time-indexed values of $c_i^+$ and $c_i^-$, updating the corresponding scaled weights and shift terms accordingly.

The value 0.25 is not a free hyperparameter, but coincides with the analytical limit $a_i, b_i \le 1/4$ required by the shift-term derivation. Above this value, the square-root terms in Eqs. (27)–(28) become complex and the shift terms $\delta$ and $\gamma$ are no longer real. The bound therefore represents the maximum admissible scaling, and any sensitivity analysis can only explore values below it.

Since the scaled weights are proportional to the market costs, namely $a_i = \lambda c_i^+$ and $b_i = \lambda c_i^-$, varying the bound rescales the absolute steepness of the loss while preserving the cost ratio $c_i^+/c_i^-$. Consequently, the direction of the over- and under-forecasting trade-off is fixed by the cost ratio and is unaffected by the scaling factor, whereas its strength is controlled by the selected bound. Smaller values produce a smoother objective closer to a symmetric loss, while values approaching the analytical limit maximize the embedded economic asymmetry. A quantitative assessment of this effect is reported in Section 5.2.

The global objective function for a batch of $N$ samples is then

$$\mathcal{L}(\Theta) = \frac{1}{N}\sum_{i=1}^{N} \ell_i(\tilde{e}_i), \tag{29}$$

where $\Theta$ denotes the trainable parameters of the forecasting model.

Fig. 4 illustrates the resulting asymmetric loss as a function of the normalized prediction error. The minimum occurs at zero error, while the two branches have different steepness according to the market-cost ratio. Increasing $c_i^+$ steepens the positive-error branch, associated with over-forecasting, whereas increasing $c_i^-$ steepens the negative-error branch, associated with under-forecasting.

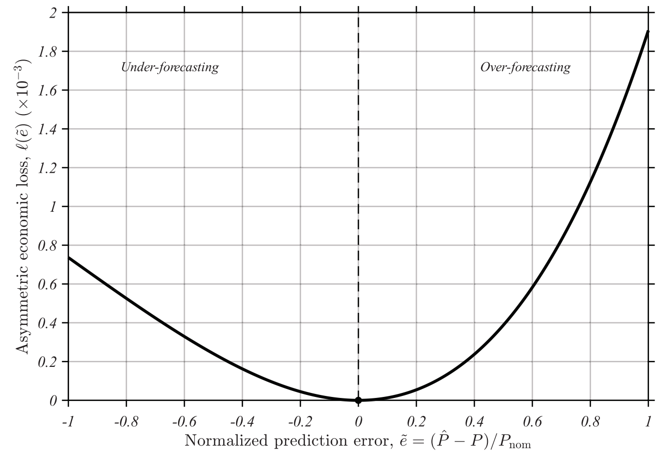


**Fig. 4.** Shape of the proposed asymmetric economic loss versus the normalized prediction error $\tilde{e} = (\hat{P} - P)/P_{\text{nom}}$. The minimum occurs at zero error, while the two branches have different steepness because the loss embeds the asymmetric market costs of over- and under-forecasting.

### 3.3. Interpretation and operational relevance

The proposed metric–loss pair (EMI for evaluation and the price-aware asymmetric loss for training) links forecasting performance to the economic logic of quarter-hour settlement. Rather than optimizing a purely statistical score, the learning process is explicitly shaped by the asymmetric cost of over- and under-forecasting under TIDE-aligned remuneration. As a result, improvements in the forecast are directly interpretable in monetary terms, which is the relevant outcome for scheduling and imbalance settlement at 15-minute granularity.

The formulation generalizes directly to time-varying settlement schemes by using time-indexed coefficients $c_i^{\pm}$ derived from market data, while preserving the same evaluation and training structure.

## 4. Experimental setup

This section describes the datasets, preprocessing steps, and evaluation protocol adopted to ensure a consistent and fair comparison among all forecasting models. Results are reported under a rolling-origin framework aligned with the TIDE market, where forecasts are updated at a 1-minute cadence with a 15-minute lead time, and economic performance is evaluated at quarter-hour settlement intervals as described in Section 3.

### 4.1. Dataset

The proposed framework is validated on the SolarTechLAB (STL) dataset, comprising directly measured photovoltaic power and meteorological observations collected at the SolarTechLAB outdoor facility located at Politecnico di Milano (45°30′10.588″ N, 9°09′23.677″ E) [23].

The photovoltaic module under analysis is a 245 W monocrystalline silicon panel composed of $N_s = 60$ series-connected cells. Its nameplate electrical ratings at Standard Test Conditions, together with the corresponding indoor-flash measurements reported in [24] and the associated relative deviation, are summarized in Table 2. The relative deviation is defined as

$$\varepsilon = \frac{X_{\text{meas}} - X_{\text{DS}}}{X_{\text{DS}}} \cdot 100 \quad [\%], \tag{30}$$

where $X$ denotes any of the electrical quantities listed in Table 2. The mismatch is dominated by an almost 3% drop in $P_{mp,r}$, driven by reductions of comparable magnitude in $V_{mp,r}$ and $I_{sc,r}$, whereas $V_{oc,r}$ and $I_{mp,r}$ remain essentially unchanged. This information is reported

**Table 2**
Electrical ratings of the 245 W monocrystalline silicon photovoltaic module at Standard Test Conditions ($G_r = 1000\,\mathrm{W/m^2}$, $T_{c,r} = 25\,°\mathrm{C}$, AM 1.5): manufacturer datasheet, indoor-flash measurement and relative deviation [24].

| Parameter | Datasheet | Measured | $\varepsilon$ (%) |
|---|---|---|---|
| $V_{oc,r}$ (V) | 37.10 | 37.15 | +0.13 |
| $I_{sc,r}$ (A) | 8.48 | 8.37 | −1.30 |
| $V_{mp,r}$ (V) | 31.30 | 30.40 | −2.88 |
| $I_{mp,r}$ (A) | 7.84 | 7.84 | 0.00 |
| $P_{mp,r}$ (W) | 245.0 | 238.0 | −2.86 |

only to document the physical characteristics of the monitored module and the experimental context of the dataset. It is not used to generate the forecasting target, since the photovoltaic output power is directly measured.

The STL monitoring system captures data at heterogeneous sampling rates. Meteorological variables include global horizontal irradiance ($G_{\mathrm{GHI}}$), diffuse horizontal irradiance ($G_{\mathrm{DHI}}$), plane-of-array irradiance ($G_{\mathrm{POA}}$), and ambient temperature ($T_{\mathrm{AIR}}$). These quantities are recorded every 10 s in order to capture rapid atmospheric transients.

Electrical variables are natively sampled at 1-minute intervals and include nine inverter-side and direct-current-side quantities: alternating-current output power ($P_{\mathrm{out}}$), alternating-current voltage ($V_{\mathrm{out}}$), alternating-current current ($I_{\mathrm{out}}$), line-to-neutral voltages ($V_{\mathrm{L1N}}$ and $V_{\mathrm{L2N}}$), grid frequency ($f_{\mathrm{grid}}$), and direct-current-side power, voltage, and current ($P_{\mathrm{pan}}$, $V_{\mathrm{pan}}$, $I_{\mathrm{pan}}$). To create a unified dataset, the 10 s meteorological observations are first resampled to 1-minute resolution by arithmetic averaging and then synchronized with the electrical time series.

Data quality control is performed in accordance with IEC 61724-1 guidelines, following the algorithmic approach described in [23]. First, an irradiance-based filter excludes records when global irradiance falls outside the range 5–1200 $\mathrm{W/m^2}$, with the lower threshold effectively removing nocturnal and non-productive periods. Second, thermal constraints discard data points when the ambient temperature exceeds 50 °C or the photovoltaic module temperature exceeds 90 °C. Finally, consistency checks remove unreliable measurements, including consecutive identical values and abrupt non-physical variations in power or irradiance, in order to avoid biasing the forecasting models with faulty data.

Models are trained on data from 2023–2024 and evaluated on the full year 2025. The forecasting target is the measured alternating-current output power $P_{\mathrm{out}}$. Two input configurations are considered. In the univariate configuration, only past values of $P_{\mathrm{out}}$ are used. In the multivariate (MISO) configuration, the nine measured electrical variables and the four measured meteorological variables listed above are jointly used as inputs, for a total of 13 measured input variables.

No physically reconstructed photovoltaic power is used either as an input feature or as a forecasting target in the present study. This ensures that the comparison among persistence, Long Short-Term Memory, Gated Recurrent Unit, and the Locally Recurrent Neural Network trained through the Binet formalism is based exclusively on measured time series available in the monitoring dataset. The possible integration of equivalent-circuit photovoltaic models into hybrid or physics-informed forecasting strategies is left for future work.

### 4.2. Data preprocessing

After the quality-control procedure described in Section 4.1, all retained variables are aligned on a common 1-minute time grid. Meteorological measurements originally acquired at 10 s resolution are averaged over each minute and synchronized with the electrical measurements, which are natively available at 1-minute resolution.

The synchronized time series are then used to construct the input–target pairs for the rolling fifteen-minute-ahead forecasting task. At each forecasting instant, only measurements available up to the current time are used as inputs, while the target is the measured photovoltaic output power 15 min ahead. This chronological construction prevents information leakage and ensures consistency with the rolling evaluation protocol described in Section 4.3.

### 4.3. Rolling forecasting strategy

The forecasting framework follows a rolling nowcasting paradigm designed for real-time market participation under the quarter-hourly settlement structure introduced by the Italian Testo Integrato del Dispacciamento Elettrico reform. At each time step $t$, every candidate model generates a prediction $\hat{P}_{t+15}$ using only historical observations available up to time $t$. This corresponds to a fifteen-minute-ahead forecasting horizon with updates every minute, providing a continuously refreshed estimate of short-term photovoltaic generation. In operational terms, this high-frequency update rate allows an aggregator, plant operator, or microgrid controller to anticipate ramps and short-term variability within the quarter-hourly settlement window, thereby supporting corrective actions aimed at reducing asymmetric imbalance penalties.

The comparison is organized at two levels. First, the Binet-based Locally Recurrent Neural Network is compared with two conventional recurrent benchmarks, namely Long Short-Term Memory and Gated Recurrent Unit. Second, all three learning-based forecasters are evaluated against persistence, which represents the standard baseline for ultra-short-term photovoltaic nowcasting. In all cases, the same target variable $P_{\mathrm{out}}$, forecasting horizon, rolling evaluation year, preprocessing procedure, and statistical and economic performance indicators are adopted. This ensures that the comparison is carried out on the same operational task and under the same decision-oriented objective.

The three neural architectures have fundamentally different internal memory and adaptation mechanisms. The Binet-based Locally Recurrent Neural Network combines a short external training window ($L = 25$) with local recursive FIR/IIR dynamics, whose effective memory is shaped by the impulse response of each neuron and by the roots governing stability. By contrast, in Long Short-Term Memory and Gated Recurrent Unit models, the effective temporal memory is mainly controlled by the observed sequence length, the hidden-state dimension, and the learned recurrent representation. For this reason, enforcing identical external memory lengths or identical internal update schemes across architectures would not produce a fair comparison; rather, it would constrain some models to operate outside the regime for which they are designed.

Accordingly, Long Short-Term Memory and Gated Recurrent Unit models are evaluated using architecture-dependent hyperparameters, including the lookback length, while preserving the same forecasting task and evaluation protocol. In particular, the use of a lookback larger than $L = 25$ for the gated recurrent models does not provide privileged future information relative to the Binet-based Locally Recurrent Neural Network; it only defines how much past information is provided to architectures that do not contain the explicit local recursive filter structure available in the latter.

Warm-start and online-adaptation strategies for recurrent forecasters are well established in the literature. Incremental ensemble Long Short-Term Memory models have been proposed to combine transfer learning and incremental learning for continuous model updating [25]. Transfer-learning strategies based on Gated Recurrent Unit architectures have been applied to photovoltaic power forecasting when pre-trained models are adapted to target domains with limited data [26]. Online Long Short-Term Memory models have also been investigated for real-time solar power forecasting with continuous parameter updates as new observations become available [27]. In the present work, these concepts are used to implement a realistic rolling deployment protocol for the conventional recurrent baselines; the methodological contribution remains the market-aligned evaluation and training framework introduced in Section 3.

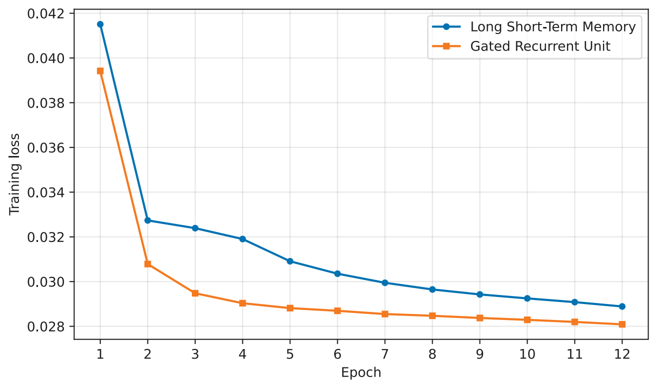


**Fig. 5.** Warm-start training-loss evolution for the Long Short-Term Memory and Gated Recurrent Unit models over the 12 historical-training epochs. Both curves decrease smoothly, indicating stable convergence before the rolling out-of-sample deployment on the 2025 test year.

For the gated recurrent architectures, training is organized in two stages:

1. A warm-start phase is performed on historical data from 2023–2024 in order to initialize the recurrent representation before rolling deployment. Both Long Short-Term Memory and Gated Recurrent Unit models consist of a single recurrent layer with 32 hidden units followed by a linear output head. Training uses the Adam optimizer with learning rate $10^{-3}$, batch size 256, and gradient clipping at 1.0 for 12 epochs. Data normalization maps the interval $[0, P_{\mathrm{nom}}]$ onto $[-1, 1]$, where $P_{\mathrm{nom}}$ is the nominal power of the photovoltaic panel. Since both bounds are physically determined, namely zero generation and the panel nameplate capacity, the normalization is independent of the specific observation period and introduces no information leakage. The same normalization range is adopted for all three neural architectures.
2. During rolling deployment on the 2025 test year, only the output head is updated, corresponding to 33 trainable parameters: 32 weights and one bias. The recurrent backbone remains frozen, with 4 480 parameters for Long Short-Term Memory and 3 360 for Gated Recurrent Unit. This head-only adaptation strategy was adopted because full online adaptation of all recurrent parameters provided only marginal gains while being more prone to drift and instability in the minute-by-minute regime. Restricting adaptation to the final layer preserves the temporal representation learned during warm-start while still allowing local correction of short-term mismatches and slow distributional shifts.

The convergence behavior of the warm-start phase is reported in Fig. 5. The training-loss curves for the Long Short-Term Memory and Gated Recurrent Unit models decrease smoothly over the 12 epochs, with no evidence of instability or divergence before the rolling out-of-sample deployment. This supports the use of the warm-started models as reliable initial conditions for the subsequent online adaptation phase.

The Binet-based Locally Recurrent Neural Network follows a different training paradigm, as described in Section 2.3. The network comprises 5 hidden neurons, each with a memory depth of $r = 14$, corresponding to 15 state variables, a hyperbolic tangent activation function, and a linear output layer whose weights $\mathbf{W}_2$ are computed analytically from the coplanarity condition at every window. Continuous adaptation proceeds through a sliding window of 25 samples. At each step, all local parameters, namely input weights $\mathbf{W}_1$, bypass links $b_{i0}$, FIR coefficients $\mathbf{b}_i$, and dynamic bases $\mathbf{z}_i$, are updated through a single-epoch gradient descent with initial learning rate $\eta = 10^{-2}$. The IIR pole radii are constrained within $|z_{im}| < 0.98$ through a bisection-enforced stability check, and any parameter update that increases the loss on the current window is rejected through backtracking.

The Binet-based Locally Recurrent Neural Network is evaluated in both a univariate configuration, where only past photovoltaic output power is used as input, and a multivariate configuration with 13 input variables. Both configurations share the same training protocol; the only difference is the dimensionality of the input layer.

For the economically trained Long Short-Term Memory and Gated Recurrent Unit models, optimization is initialized from the corresponding warm-started mean-square-error-trained model rather than from random weights. A pre-warming phase of 3 epochs, with learning rate $5 \times 10^{-6}$, is performed on the historical data. This allows the economic fine-tuning to reshape the forecast bias according to the imbalance-cost asymmetry without relearning the entire forecasting task from scratch. During online deployment, the economic models use a wider adaptation context than the mean-square-error-trained ones, with updates every 20 steps instead of every 10 steps, 20 adaptation windows, batch size 20, and learning rate $2 \times 10^{-6}$. This more conservative update configuration reflects the higher sensitivity of asymmetric objectives to noisy gradient updates.

Table 3 summarizes the complete hyperparameter configuration for all three neural architectures.

### 4.4. Evaluation protocol

All models are evaluated under identical conditions regarding data availability, preprocessing, and test horizon.

Since the available measurements are power samples $P_{\mathrm{out}}$ in watts at 1-minute resolution, the settlement energy required by the EMI formulation (Section 3.1) is obtained in MWh as

$$E_i = \frac{\Delta t}{10^6} \sum_{k \in \mathcal{I}_i} P_k \tag{31}$$

where $\Delta t = \frac{1}{60}$ h and $\mathcal{I}_i$ collects the sample indices within the $i$th quarter-hour.

The power prediction error used for the statistical indicators is defined as

$$e_i^P = \hat{P}_{t+h} - P_{t+h}. \tag{32}$$

The corresponding energy error used in the EMI calculation is obtained by multiplying the power error by the settlement duration and converting watt-hours into megawatt-hours. Accordingly, statistical metrics are reported in watts, while EMI and $\mathrm{EMI}_{\mathrm{norm}}$ are computed on the corresponding energy errors in MWh.

As a mandatory benchmark, a persistence model is considered:

$$\hat{P}_{t+h}^{\mathrm{pers}} = P_t, \qquad h = 15 \text{ min}. \tag{33}$$

Statistical performance is assessed using:

$$\mathrm{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^{N} \left(e_i^P\right)^2}, \tag{34}$$

$$\mathrm{MAE} = \frac{1}{N} \sum_{i=1}^{N} \left|e_i^P\right|, \tag{35}$$

$$\mathrm{MBE} = \frac{1}{N} \sum_{i=1}^{N} e_i^P, \tag{36}$$

while economic performance is evaluated through the imbalance cost indicators EMI and $\mathrm{EMI}_{\mathrm{norm}}$ introduced in Section 3.

To improve interpretability, EMI is decomposed into over- and under-forecasting contributions ($\mathrm{EMI}^+$, $\mathrm{EMI}^-$) and also expressed as a

**Table 3**
Hyperparameter configuration for all candidate architectures. LSTM and GRU share the same settings; LRNN–Binet follows a different training paradigm (sliding-window with analytical output-layer regression). The economic loss column refers to the variant described in Section 3.

| Category | Parameter | LSTM/GRU | LRNN–Binet |
|---|---|---|---|
| Architecture | Recurrent/hidden layers | 1 | 1 |
| | Hidden units/neurons | 32 | 5 |
| | Memory depth $r$ | — (implicit) | 14 |
| | Dropout | 0.0 | – |
| | Hidden activation | tanh (built-in) | tansig |
| | Output head | Linear (32, 1) | Analytical ($\mathbf{W}_2$) |
| Input | Lookback window $L$ | 120 (2 h) | 25 (25 min) |
| | Prediction horizon $h$ | 15 min | |
| | Normalization | MinMax $[0, P_{\text{nom}}] \to [-1, 1]$ | |
| Warm-start (2023–2024) | Epochs | 12 | – |
| | Batch size | 256 | – |
| | Learning rate | $10^{-3}$ | – |
| | Optimizer | Adam | – |
| | Gradient clipping | 1.0 | – |
| Online adaptation (MSE) | Scope | Head only (33 params) | All local params |
| | Update frequency | Every 10 steps | Every step |
| | Adaptation windows | 10 | 1 |
| | Batch size | 10 | 25 |
| | Learning rate | $10^{-5}$ | $10^{-2}$ |
| | Max. iterations/window | 1 epoch | 1 |
| Online adaptation (Economic) | Scope | Head only (33 params) | All local params |
| | Update frequency | Every 20 steps | Every step |
| | Adaptation windows | 20 | 1 |
| | Batch size | 20 | 25 |
| | Learning rate | $2 \times 10^{-6}$ | $10^{-2}$ |
| | Loss function | Eq. (21) | |
| Economic pre-warming | Epochs | 3 | – |
| | Learning rate | $5 \times 10^{-6}$ | – |
| Stability | IIR pole constraint | – | $\lvert z_{im}\rvert < 0.98$ (bisection) |
| | Backtracking | – | Reject if loss increases |
| Parameter count | Frozen/fixed | 4 480 (LSTM) / 3 360 (GRU) | $\mathbf{W}_2$ (analytical) |
| | Online-trainable | 33 | $\mathbf{W}_1$, $b_{i0}$, $\mathbf{b}_i$, $\mathbf{z}_i$ |

**Table 4**
Effect of training objective on annual performance over 2025. For each architecture, results are reported for both standard MSE training and the proposed economic loss. Shares are computed with tolerance $\epsilon = 0.1$ W and reported as $\mathbb{P}(e^P > \epsilon)/\mathbb{P}(e^P < -\epsilon)/\mathbb{P}(|e^P| \le \epsilon)$.

| Architecture | Training | RMSE (W) | MAE (W) | MBE (W) | EMI (€) | $\text{EMI}_{\text{norm}}$ (€/MWh) | Perf.(%) | Shares (%) |
|---|---|---|---|---|---|---|---|---|
| LSTM | MSE | 20.98 | 10.89 | −0.146 | 5.505 | 20.271 | 86.49 | 48.76/48.84/2.40 |
| LSTM | Econ. loss | 21.16 | 10.43 | −1.485 | 5.135 | 18.910 | 87.39 | 36.48/55.92/7.61 |
| GRU | MSE | 20.71 | 10.49 | 0.024 | 5.320 | 19.589 | 86.94 | 49.67/48.11/2.22 |
| GRU | Econ. loss | 21.05 | 10.36 | −1.891 | 5.062 | 18.640 | 87.57 | 37.50/55.10/7.40 |
| LRNN–Binet | MSE | 16.75 | 6.11 | −0.087 | 3.089 | 11.373 | 92.42 | 43.16/44.47/12.37 |
| LRNN–Binet | Econ. loss | 14.61 | 5.61 | 0.002 | 2.844 | 10.471 | 93.02 | 42.92/44.50/12.59 |

percentage of the nominal photovoltaic revenue. The nominal revenue is defined as:

$$\text{Revenue} = \sum_{i\in\mathcal{I}} p_i E_i, \tag{37}$$

where $p_i$ is the reference remuneration tariff for photovoltaic energy in interval $i$. In the present case study, a constant value $p_i = 150$ €/MWh is adopted, consistently with the pricing assumptions used in the EMI calculation. The economic performance indicator is then computed as

$$\text{Perf} = 100\left(1 - \frac{\text{EMI}}{\text{Revenue}}\right). \tag{38}$$

Finally, directional error statistics are computed using a tolerance $\epsilon = 0.1$ W:

$$\mathbb{P}(e_i^P > \epsilon), \qquad \mathbb{P}(e_i^P < -\epsilon), \qquad \mathbb{P}(|e_i^P| \le \epsilon), \tag{39}$$

providing insight into the bias and economic asymmetry of forecasting errors.

Overall, this protocol combines rolling 15-minute-ahead forecasting, minute-level updates, and quarter-hourly economic assessment, thereby linking predictive accuracy directly to market performance under the TIDE framework.

## 5. Results

This section presents the results obtained under the experimental framework described in Section 4. The analysis is carried out on the Milan dataset and is reported at yearly, monthly, and daily resolution. The section discusses the effect of the training objective, the comparison among architectures, the temporal evolution of the economic indicators over the year, representative daily case studies, and the impact of exogenous variables through a comparison between univariate and multivariate LRNN–Binet configurations.

### 5.1. Annual performance on the milan dataset

Table 4 reports the annual performance on the Milan dataset for 2025, comparing LSTM, GRU, and LRNN–Binet under both standard

**Table 5**
Sensitivity of the GRU economic model to the upper bound on the scaled economic weights $\max(a_i, b_i)$ (annual results, Milan 2025, $c_i^+ = 150$ €/MWh, $c_i^- = 100$ €/MWh). Directional error shares are reported as $\mathbb{P}(e^P > \epsilon) / \mathbb{P}(e^P < -\epsilon) / \mathbb{P}(|e^P| \leq \epsilon)$ with $\epsilon = 0.1$ W.

| $\max(a_i, b_i)$ | RMSE (W) | MBE (W) | $\mathrm{EMI}_{\mathrm{norm}}$ (€/MWh) | Shares (%) |
|---|---|---|---|---|
| 0.10 | 20.66 | −0.241 | 19.880 | 48.53/48.81/2.66 |
| 0.15 | 20.66 | −0.314 | 19.858 | 48.27/49.07/2.66 |
| 0.20 | 20.67 | −0.426 | 19.832 | 47.84/49.46/2.69 |
| 0.25 (adopted) | 21.05 | −1.891 | 18.640 | 37.50/55.10/7.40 |

MSE and the proposed price-aware asymmetric loss, with persistence as the reference baseline. Several observations emerge.

First, all learning-based models consistently outperform persistence in both statistical and economic terms. This confirms that model-based ultra-short-term forecasting provides a clear operational advantage under the TIDE-aligned framework, particularly when evaluated through economically meaningful indicators.

Second, the gated recurrent architectures (LSTM and GRU) form an intermediate performance group. They exhibit similar statistical accuracy and comparable reductions in imbalance cost, with their economically trained variants achieving the best results within each architecture. Relative to persistence, both models yield a reduction in normalized economic impact on the order of 18%–19%, indicating that the proposed training objective provides a tangible, though moderate, economic benefit.

Third, LRNN–Binet clearly outperforms both LSTM and GRU across all metrics. Its economically trained configuration achieves a reduction in normalized imbalance cost exceeding 50% with respect to persistence. This performance gap is substantial and cannot be interpreted as incremental; rather, it suggests that the LRNN–Binet formulation captures short-term PV dynamics more effectively in the considered forecasting setting.

The economic performance indicator provides an additional perspective. While persistence retains roughly 85% of the nominal revenue, the gated recurrent models increase this value to approximately 87%, and LRNN–Binet further improves it to about 93%. This confirms that the observed statistical improvements translate directly into meaningful monetary gains.

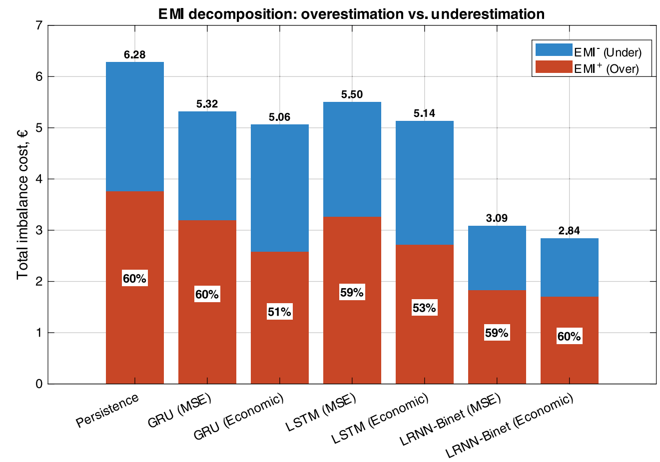


**Fig. 6.** Decomposition of the annual EMI into over-forecasting ($\mathrm{EMI}^+$) and under-forecasting ($\mathrm{EMI}^-$) contributions as defined in Eq. (18). Percentages indicate the $\mathrm{EMI}^+$ share of the total cost for each model.

### 5.2. Effect of the training objective

The comparison in Table 4 clarifies the role of the proposed economic loss. Its impact is consistent at the economic level, but architecture-dependent at the statistical level.

For LSTM and GRU, the economic loss systematically reduces $\mathrm{EMI}_{\mathrm{norm}}$, but does not minimize RMSE. This reflects the objective of the proposed loss, which is not to optimize symmetric pointwise accuracy, but to reduce the most costly forecast deviations under asymmetric settlement pricing. As a result, a slight increase in RMSE is observed alongside a consistent reduction in economic impact.

This behavior is further reflected in the bias term. For both LSTM and GRU, the economically trained models exhibit a more negative MBE compared to their MSE-trained counterparts, indicating a mild tendency to under-forecast. In the considered setting, this shift is economically meaningful because over-forecasting carries a higher penalty than under-forecasting. The directional error shares confirm this interpretation: the economic loss reduces the proportion of positive (over-forecasting) errors and increases the proportion of negative (under-forecasting) ones, effectively steering the forecasts toward a less costly operating regime.

To quantify the role of the upper bound on the scaled economic weights, motivated in Section 3.2, we varied the cap on $\max(a_i, b_i)$ over $\{0.10, 0.15, 0.20, 0.25\}$ while preserving the cost ratio $c_i^+/c_i^-$. The analysis was carried out on the GRU model trained with the economic loss, which is the most sensitive configuration since its economic benefit relies mainly on the bias-redistribution mechanism. The results, summarized in Table 5, show that the effect of the cap is monotonic but strongly nonlinear. For bounds between 0.10 and 0.20, the model behaves almost like a symmetrically trained one: the bias remains close to zero (MBE between −0.24 and −0.43 W), the directional shares stay close to a balanced split (about 48%/49%), and $\mathrm{EMI}_{\mathrm{norm}}$ remains around 19.8 €/MWh, essentially in the mean-square-error regime. The characteristic redistribution of forecast errors toward the less penalized under-forecasting direction (shares shifting to 37.50%/55.10%) and the corresponding reduction of $\mathrm{EMI}_{\mathrm{norm}}$ to 18.64 €/MWh, at the cost of a slight RMSE increase, emerge only as the cap approaches its analytical limit. This confirms that the adopted value of 0.25 is the maximum admissible scaling and the one that effectively activates the price-aware behavior of the loss.

For LRNN–Binet, the effect of the economic loss is both stronger and more favorable. In this case, the asymmetric objective improves not only the economic metric but also the statistical indicators. This indicates that LRNN–Binet is particularly well aligned with the proposed loss formulation and can exploit it without the statistical trade-off observed in gated recurrent architectures. Fig. 6 provides a complementary perspective through the decomposition of the total imbalance cost into over-forecasting ($\mathrm{EMI}^+$) and under-forecasting ($\mathrm{EMI}^-$) components, as defined in Eq. (18).

For persistence and all MSE-trained models, the over-forecasting component accounts for approximately 59%–60% of the total cost, reflecting the higher unit penalty assigned to positive errors ($c^+ > c^-$). When the economic loss is applied to LSTM and GRU, the share of $\mathrm{EMI}^+$ decreases noticeably, confirming that the asymmetric objective effectively redistributes forecast errors toward under-forecasting, which carries a lower marginal cost.

For LRNN–Binet, the over/under ratio remains approximately unchanged across training objectives, while the total cost is substantially reduced. This indicates that its economic advantage arises primarily from a global reduction in forecast error magnitude rather than from a bias shift. In contrast, LSTM and GRU rely more strongly on the bias-redistribution mechanism induced by the asymmetric loss.

**Table 6**
Annual performance comparison against persistence over 2025 (15-minute-ahead forecasting). For each architecture, the best configuration is the one that minimizes $EMI_{norm}$. The calculation formulas for the statistical indicators, economic indicators, performance percentage, and directional shares are provided in Sections 3.1 and 4.4.

| Model | RMSE (W) | MAE (W) | MBE (W) | EMI (€) | $EMI_{norm}$ (€/MWh) | Perf. (%) | $\Delta EMI_{norm}$ (%) |
|---|---|---|---|---|---|---|---|
| Persistence | 24.22 | 12.40 | −0.034 | 6.280 | 23.124 | 84.58 | 0.00 |
| LSTM (best) | 21.16 | 10.43 | −1.485 | 5.135 | 18.910 | 87.39 | −18.22 |
| GRU (best) | 21.05 | 10.36 | −1.891 | 5.062 | 18.640 | 87.57 | −19.39 |
| LRNN–Binet (best) | 14.61 | 5.61 | 0.002 | 2.844 | 10.471 | 93.02 | −54.72 |

### 5.3. Best-model comparison against persistence

To summarize the practical outcome of the yearly analysis, Table 6 compares persistence with the best configuration of each architecture, selected as the one minimizing $EMI_{norm}$ over the full year.

All machine learning architectures outperform persistence. The improvement is modest but clear for LSTM and GRU, and much stronger for LRNN–Binet. GRU is slightly better than LSTM in the present dataset, both statistically and economically, but the difference between the two gated recurrent models remains small. The main gap is instead the one between the gated models and LRNN–Binet.

In particular, LRNN–Binet reduces annual EMI from 6.280 € to 2.844 €, while also achieving the lowest RMSE and MAE. This confirms that the best-performing model is not merely the one with the lowest symmetric error, but the one that best aligns predictive dynamics with the economic structure of the task.

### 5.4. Temporal evolution of the economic impact over 2025

The yearly curves further support the conclusions drawn from the aggregate tables. Fig. 7 reports the cumulative EMI over 2025 for all considered models. Persistence exhibits the steepest growth over the full year, indicating that its forecasting errors accumulate substantial cost in every season. All learning-based models reduce this accumulation, but to different extents.

The LSTM and GRU curves remain relatively close throughout the year, with the economically trained versions consistently below their MSE-trained counterparts. This confirms that the proposed loss function yields a systematic economic advantage even when the difference in statistical metrics is small.

The LRNN–Binet curves remain much lower than all other ones over the entire time horizon. Furthermore, the gap between LRNN–Binet and the gated recurrent models gradually widens over the year, showing that the improvement is not just from a few isolated instances but is continuous and accumulative. In operational terms, this is especially significant because it demonstrates that seemingly modest improvements in short-term prediction can lead to substantial annual savings once the imbalance cost is accumulated over a year.

A complementary view is provided by Fig. 8, which reports monthly $EMI_{norm}$. Seasonal variability is clearly evident, with higher values in certain winter and shoulder-season months, when the combination of lower irradiance and increased relative variability complicates forecasting. However, the relative ranking of the models remains consistent across all months: LRNN–Binet is always the best performer, persistence is always the worst, and LSTM and GRU are in between.

This monthly consistency is important because it shows that the advantage of LRNN–Binet is not dependent on a specific seasonal regime. Likewise, the economically trained versions of LSTM and GRU remain systematically better than their MSE-trained counterparts month after month, confirming that the proposed loss is robust over changing operating conditions.

### 5.5. Representative daily case studies

To complement the yearly and monthly indicators, Fig. 9 reports four representative days selected to span different operating conditions:

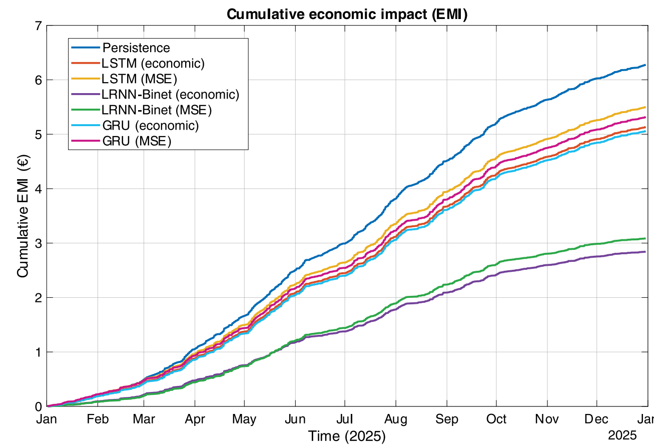


**Fig. 7.** Cumulative economic impact (EMI) over the full year 2025 for persistence, LSTM, GRU, and LRNN–Binet under both MSE and economic training.

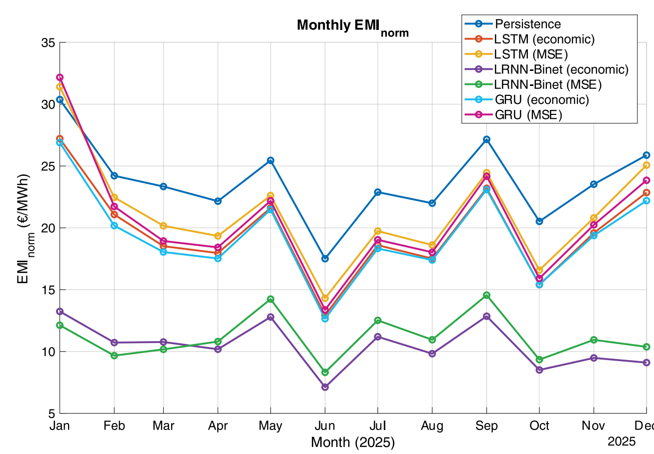


**Fig. 8.** Monthly normalized economic impact $EMI_{norm}$ over 2025 for all considered models on the Milan dataset.

a highly variable day (2025-03-01), a clear-sky day (2025-05-01), a moderately variable day with multiple intra-day ramps (2025-06-04), and a strongly intermittent day with abrupt fluctuations around midday (2025-10-01).

The clear-sky case on 2025-05-01 highlights the main weakness of persistence: a temporal lag during both the morning rise and the evening decay of PV power. This results in consistent cumulative EMI throughout the day. All learned models reduce this lag, but LRNN–Binet provides the closest overall match to the measured profile and the lowest cumulative cost at the end of the day.

On the more variable days, the differences between architectures become clearer. On 2025-03-01 and 2025-06-04, the PV signal presents multiple local ramps and fast fluctuations. In these situations, persistence cannot follow the short-term dynamics and accumulates the highest cost. LSTM and GRU improve tracking of these variations, but still exhibit some smoothing and delayed response around the sharpest

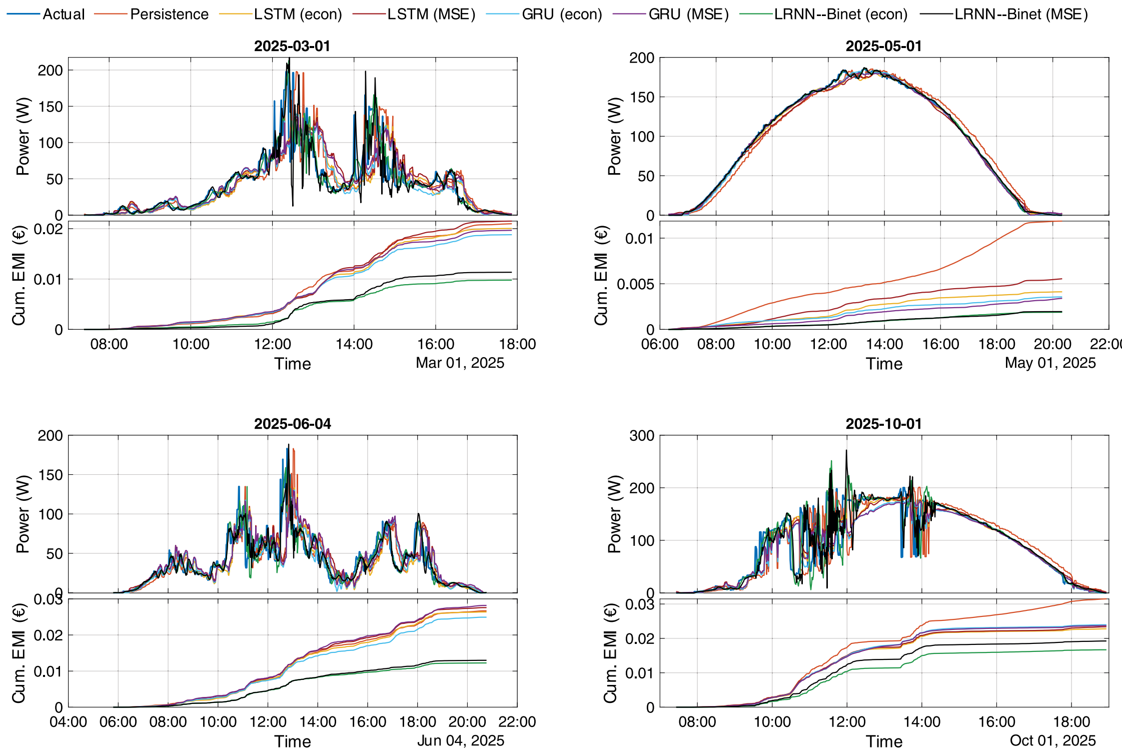


**Fig. 9.** Representative daily case studies for four operating conditions in 2025. Top panels: measured and forecasted PV power. Bottom panels: cumulative EMI over the day.

**Table 7**
Comparison between univariate (SISO) and multivariate (MISO) economically trained LRNN–Binet models over 2025. Shares are computed with tolerance $\epsilon = 0.1$ W and reported as $P(e^P > \epsilon)$ / $P(e^P < -\epsilon)$ / $P(|e^P| \leq \epsilon)$.

| Model | RMSE (W) | MAE (W) | MBE (W) | EMI (€) | $\mathrm{EMI}_{\mathrm{norm}}$ (€/MWh) | Perf. (%) | Shares (%) |
|---|---|---|---|---|---|---|---|
| Persistence | 24.220 | 12.392 | −0.034 | 6.2806 | 23.124 | 84.58 | 47.82/44.71/7.47 |
| LRNN SISO | 14.607 | 5.608 | 0.002 | 2.8441 | 10.471 | 93.02 | 48.63/49.37/2.00 |
| LRNN MISO | 15.186 | 5.707 | −0.551 | 2.8381 | 10.449 | 93.03 | 46.29/50.30/3.42 |

transitions. LRNN–Binet stays more aligned with the measured signal, so its cumulative EMI curve increases more slowly throughout the day.

The day 2025-10-01 is especially informative because it combines irregular early ramps, strong midday variability, and a smoother afternoon decline. Again, all machine learning models outperform persistence, but the ranking remains the same: LRNN–Binet is the best, while GRU and LSTM are in the middle, and persistence is the worst. The comparison between MSE and economically trained models is also consistent with the yearly analysis: for LSTM and GRU, the economic-loss versions do not necessarily appear visually superior at every instant, but they accumulate less economic imbalance by the end of the day. For LRNN–Binet, the economic-loss version is superior both visually and economically.

Overall, the daily plots confirm the global picture emerging from the yearly metrics: persistence suffers from a systematic lag, LSTM and GRU provide substantial improvements but remain close to each other, and LRNN–Binet delivers the strongest and most consistent reduction in imbalance costs across both smooth and highly variable conditions.

### *5.6. Univariate versus multivariate modeling*

An additional question explored in this study is whether including exogenous variables improves forecasting performance at the 15-minute horizon. To address this, LRNN–Binet was evaluated in both a univariate configuration (SISO), using only past PV output, and a multivariate configuration (MISO), incorporating electrical and meteorological variables.

The multivariate (MISO) configuration includes 13 input variables: 9 electrical measurements from the inverter (AC output power, voltage, current, line-to-neutral voltages, grid frequency, and DC-side voltage, current, and power) and 4 meteorological variables (ambient temperature, diffuse horizontal irradiance, global horizontal irradiance, and plane-of-array irradiance). All variables are simultaneously fed to the LRNN–Binet input layer and normalized via mapminmax scaling.

Table 7 reports the comparison. The multivariate model achieves a slightly lower economic impact, with $\mathrm{EMI}_{\mathrm{norm}}$ decreasing from 10.471 to 10.449 €/MWh. However, this improvement is extremely small in absolute terms. At the same time, RMSE increases from 14.61 to 15.19 W and MAE increases from 5.61 to 5.71 W, while the model introduces a more negative bias.

Therefore, the multivariate configuration does not provide a meaningful overall advantage in the present ultra-short-term setting. The small EMI gain is almost negligible, whereas the statistical indicators worsen slightly. This suggests that, at a 15-minute forecasting horizon, most of the predictive information is already embedded in the recent PV power trajectory, and adding electrical and meteorological variables contributes only marginally. In practical terms, the simpler univariate setup is preferable because it delivers nearly the same economic performance with less model and data-management complexity.

This result is important for the paper because it shows that increasing input richness does not automatically improve forecasting quality. In this case, model design and training objectives seem more important than expanding the feature space.

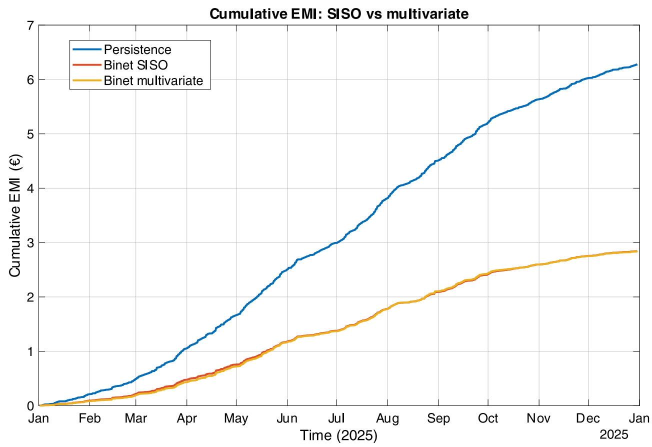


**Fig. 10.** Cumulative EMI over 2025 for persistence, univariate LRNN–Binet (SISO), and multivariate LRNN–Binet (MISO) on the Milan dataset.

The same conclusion is visually confirmed by Fig. 10, where the cumulative EMI curves of the SISO and multivariate LRNN–Binet models are almost indistinguishable over the full year. Their common separation from persistence is large, whereas their mutual difference is negligible. This supports the interpretation that the multivariate extension does not substantially change the forecasting outcome in the considered regime.

## 6. Conclusion

This paper presented a decision-oriented framework for ultra-short-term photovoltaic power nowcasting under quarter-hourly electricity settlement. The central idea is that forecasting should not be assessed only as a statistical prediction problem, but also as an economic decision-support problem. To this end, the proposed framework links forecast errors to their monetary consequences through the Expected Monetary Impact indicator and introduces a differentiable price-aware asymmetric loss function that embeds the different costs of over- and under-forecasting directly into model training.

The experimental analysis was carried out on real photovoltaic measurements from the SolarTechLAB facility in Milan, using historical data from 2023–2024 for training and the full year 2025 as an unseen chronological test period. The results show that all learning-based forecasters outperform persistence, confirming the operational value of data-driven nowcasting at the fifteen-minute horizon. Compared with persistence, the best Long Short-Term Memory and Gated Recurrent Unit configurations reduce the normalized monetary impact by 18.22% and 19.39%, respectively, while the Locally Recurrent Neural Network trained through the Binet formalism achieves a 54.72% reduction. This demonstrates that aligning the learning objective with market costs can produce benefits that are not fully captured by conventional symmetric accuracy indicators.

The comparison also shows that different architectures exploit the economic objective in different ways. Long Short-Term Memory and Gated Recurrent Unit models mainly benefit from a controlled bias shift toward the less penalized error direction. By contrast, the Locally Recurrent Neural Network trained through the Binet formalism reduces the overall forecast error magnitude while preserving a balanced distribution between over- and under-forecasting. This suggests that local recurrent dynamics, analytical output-layer regression, and continuous sliding-window adaptation are particularly effective for capturing short-term photovoltaic variability in a market-aligned forecasting setting.

The analysis of univariate and multivariate input configurations indicates that, at the fifteen-minute horizon considered here, recent photovoltaic output already contains most of the information needed for economically effective forecasting. Adding electrical and meteorological variables produces only a marginal reduction in monetary impact and slightly worsens conventional statistical indicators. This result highlights an important practical implication: increasing input complexity does not automatically improve operational value, and simpler forecasting configurations may be preferable when they provide comparable economic performance with lower data-management requirements.

The proposed framework is not limited to the Italian regulatory case considered in this study. The same metric and training-loss structure can be transferred to other electricity markets characterized by short settlement intervals and asymmetric or time-varying imbalance prices, by simply replacing the cost coefficients $c^+$ and $c^-$ with locally applicable values while preserving the same decision-oriented formulation. This applies, for instance, to the European markets operating under the fifteen-minute imbalance settlement mandated by European regulation, as well as to markets with even shorter intervals, such as the five-minute settlement adopted in the Australian National Electricity Market. This makes the approach suitable for photovoltaic plant operators, aggregators, energy communities, and flexibility providers that need forecasts whose value is measured in terms of imbalance-cost reduction rather than statistical accuracy alone.

The study also has limitations that define the next research steps. The validation is based on a single measured photovoltaic system, so future work should extend the analysis to larger plants, multi-inverter portfolios, and geographically distributed datasets. A further limitation is that the experimental benchmark is restricted to recurrent forecasting architectures and persistence. Future work should extend the comparison to convolutional, hybrid convolutional–recurrent, and transformer-based models under the same market-aligned evaluation protocol. Portfolio aggregation may smooth high-frequency ramps and could modify the relative advantage of different forecasting architectures. A second development will be the integration of time-varying imbalance prices from real market data into both the monetary-impact evaluation and the asymmetric training objective. Finally, future research may investigate hybrid and physics-informed extensions, in which measured time series are combined with physical photovoltaic models without compromising the market-aligned structure of the proposed framework.

## CRediT authorship contribution statement

**Saloni Dhingra:** Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Sara Carcangiu:** Writing – review & editing, Validation, Supervision, Software, Methodology, Formal analysis, Conceptualization. **Binh Nam Nguyen:** Writing – review & editing, Validation, Software, Data curation. **Augusto Montisci:** Writing – review & editing, Software, Project administration, Methodology, Formal analysis, Conceptualization. **Giancarlo Storti Gajani:** Writing – review & editing, Supervision, Methodology. **Sonia Leva:** Writing – review & editing, Supervision, Resources. **Emanuele Ogliari:** Writing – review & editing, Supervision, Project administration, Methodology, Conceptualization.

## Declaration of generative AI and AI-assisted technologies in the manuscript preparation process.

During the preparation of this manuscript, the authors used AI tools (e.g., ChatGPT and Grammarly) to improve language quality, structural clarity, and overall readability. No AI tools were used for data analysis, model development, simulation, result generation, or scientific interpretation. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

## Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Augusto Montisci reports financial support was provided by Fondazione di Sardegna. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

This work was supported by Fondazione di Sardegna, Italy, under the project "Prosumer Resilience-Oriented Management of Enhanced, Thermodynamically Efficient, Unmoving Systems", CUP F83C260003 50007.

## Data availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

## References

[1] Mentens L, Peremans H, Springael J, Nimmegeers P. Flexibility in short-term electricity markets for renewable integration and uncertainty mitigation: A comprehensive review. Smart Energy 2025;18. http://dx.doi.org/10.1016/j.segy.2025.100183.

[2] Terna. The new TIDE regulation explained. 2025, [Accessed 30 October 2025], https://lightbox.terna.it/en/insight/tide-electricity-dispatching.

[3] European Parliament and Council. Regulation (EU) 2019/943 of 5 june 2019 on the internal market for electricity (recast). 2019, Official Journal of the European Union, L 158, Art. 8(4).

[4] Shahverdi N, Saffari A, Amiri B. A systematic review of artificial intelligence and machine learning in energy sustainability: Research topics and trends. Energy Rep 2025;13:5551–78. http://dx.doi.org/10.1016/j.egyr.2025.05.021.

[5] Brusco G, Burgio A, Menniti D, Pinnarelli A, Sorrentino N, Vizza P. Quantification of forecast error costs of photovoltaic prosumers in Italy.. Energies 2017;10(11). http://dx.doi.org/10.3390/en10111754.

[6] Zhang Y, Qin C, Srivastava AK, Jin C, Sharma RK. Data-driven day-ahead PV estimation using autoencoder-LSTM and persistence model. IEEE Trans Ind Appl 2020;56(6):7185–92. http://dx.doi.org/10.1109/TIA.2020.3025742.

[7] Mosaico G, Saviozzi M. A hybrid methodology for the day-ahead PV forecasting exploiting a clear sky model or artificial neural networks. In: IEEE EUROCON 2019 -18th international conference on smart technologies. 2019, p. 1–6. http://dx.doi.org/10.1109/EUROCON.2019.8861551.

[8] Singh C, Garg AR. Enhancing solar power output predictions: Analyzing ARIMA and S-ARIMA models for short-term forecasting. In: 2024 IEEE 11th power India international conference. 2024, p. 1–5. http://dx.doi.org/10.1109/PIICON63519.2024.10995070.

[9] De Falco P, Di Noia LP, Rizzo R. Exponential smoothing model for photovoltaic power forecasting. In: 2021 9th international conference on modern power systems. 2021, p. 1–5. http://dx.doi.org/10.1109/MPS52805.2021.9492538.

[10] Datta L, Rizwan M, Anand P. Comprehensive review and evaluation of machine learning approaches for solar PV output power forecasting. In: 2023 international conference on computational intelligence and sustainable engineering solutions. 2023, p. 794–800. http://dx.doi.org/10.1109/CISES58720.2023.10183451.

[11] Wentz VH, Maciel JN, Gimenez Ledesma JJ, Ando Junior OH. Solar irradiance forecasting to short-term PV power: Accuracy comparison of ANN and LSTM models. Energies 2022;15. http://dx.doi.org/10.3390/en15072457.

[12] Perveen G, Rizwan M, Goel N. Comparison of intelligent modelling techniques for forecasting solar energy and its application in solar PV based energy system. In: IET energy systems integration. vol. 1, 2019, p. 34–51. http://dx.doi.org/10.1049/iet-esi.2018.0011.

[13] Bai M, Chen Y, Zhao X, Liu J, Yu D. Deep attention convlstm-based adaptive fusion of clear-sky physical prior knowledge and multivariable historical information for probabilistic prediction of photovoltaic power. Expert Syst Appl 2022;202. http://dx.doi.org/10.1016/j.eswa.2022.117335.

[14] Suresh V. Benchmarking transformer variants for hour-ahead PV forecasting: PatchTST with adaptive conformal inference. Energies 2025;18(18). http://dx.doi.org/10.3390/en18185000.

[15] Carcangiu S, Montisci A. A training algorithm for locally recurrent neural networks based on the explicit gradient of the loss function. Algorithms 2025;18(104):1–15. http://dx.doi.org/10.3390/a18020104.

[16] Yang B, Zhu T, Cao P, Guo Z, Zeng C, Li D, et al. Classification and summarization of solar irradiance and power forecasting methods: A thorough review. CSEE J Power Energy Syst 2023;9(3):978–95. http://dx.doi.org/10.17775/CSEEJPES.2020.04930.

[17] Jannah N, Gunawan TS, Yusoff SH, Hanifah MSA, Sapihie SNM. Recent advances and future challenges of solar power generation forecasting. IEEE Access 2024;12:168904–24. http://dx.doi.org/10.1109/ACCESS.2024.3496120.

[18] Talukdar A, Hota S, Panda B, Maharana MK, Ravikumar N, Sundar LS, et al. Performance analysis of different machine learning techniques in solar energy forecasting. In: 2025 11th international conference on communication and signal processing. 2025, p. 1654–8. http://dx.doi.org/10.1109/ICCSP64183.2025.11088730.

[19] Hochreiter S, Schmidhuber J. Long short-term memory. Neural Comput 1997;9(8):1735–80. http://dx.doi.org/10.1162/neco.1997.9.8.1735.

[20] Gers F, Schmidhuber J, Cummins F. Learning to forget: continual prediction with LSTM. In: 1999 ninth international conference on artificial neural networks ICANN 99. (conf. publ. no. 470). vol. 2, 1999, p. 850–5 vol.2. http://dx.doi.org/10.1049/cp:19991218.

[21] Dey R, Salem FM. Gate-variants of gated recurrent unit (GRU) neural networks. In: 2017 IEEE 60th international midwest symposium on circuits and systems. 2017, p. 1597–600. http://dx.doi.org/10.1109/MWSCAS.2017.8053243.

[22] Montisci A. A free from local minima algorithm for training regressive MLP neural networks. Adv Artif Intell Mach Learn 2024;4(1):2103–12. http://dx.doi.org/10.54364/AAIML.2024.41120.

[23] Ogliari E, Dolara A, Mazzeo D, Manzolini G, Leva S. Bifacial and monofacial PV systems performance assessment based on iec 61724-1 standard. IEEE J Photovoltaics 2023;13(5):756–63. http://dx.doi.org/10.1109/JPHOTOV.2023.3295869.

[24] Dolara A, Leva S, Manzolini G. Comparison of different physical models for PV power output prediction. Sol Energy 2015;119:83–99. http://dx.doi.org/10.1016/j.solener.2015.06.017.

[25] Wang H, Li M, Yue X. IncLSTM: Incremental ensemble LSTM model towards time series data. Comput Electr Eng 2021;92:107156. http://dx.doi.org/10.1016/j.compeleceng.2021.107156.

[26] Tang Y, Yang K, Zhang S, Zhang Z. Photovoltaic power forecasting: A dual-attention gated recurrent unit framework incorporating weather clustering and transfer learning strategy. Eng Appl Artif Intell 2024;130:107691. http://dx.doi.org/10.1016/j.engappai.2023.107691.

[27] Ullah Z, Shahid SM, Park J, Min B, Kwon S. LSTM-based online learning for real-time solar power forecasting. IEEE Access 2025;13:160456–65. http://dx.doi.org/10.1109/ACCESS.2025.3606910.