Energy Reports 14 (2025) 4111–4124

Contents lists available at ScienceDirect

Energy Reports

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

Research paper

Improving sky image-based short-term solar irradiance forecasting through the integration of meteorological data

Martin Ansong [a,b], Albertina M. Amakali [a,c], Thomas N. Nyang'onda [b], Robinson J. Musembi [b], Petja Dobreva [c], Bryce S. Richards [a,d,*]

[a] Institute of Microstructure Technology, Karlsruhe Institute of Technology, Hermann-von-Helmholtz-Platz 1, Eggenstein-Leopoldshafen 76344, Germany
[b] Department of Physics, University of Nairobi, Kenya
[c] Department of Physics, Chemistry and Material Science, University of Namibia, Windhoek, Namibia
[d] Light Technology Institute, Karlsruhe Institute of Technology, Engesserstrasse 13, Karlsruhe 76131, Germany

ARTICLE INFO

Keywords:
Solar irradiance
Sky images
Short-term forecasting
Meteorological data
Deep learning

ABSTRACT

Fluctuations in solar irradiance (SI) cause instability in power generation, challenging grids with high photovoltaic (PV) penetration and impacting power quality. Accurate short-term SI forecasting is crucial for reliable solar energy integration, enabling grid operators to optimise dispatch, manage reserves, and schedule energy storage. It also supports compliance with electricity markets requiring hourly or sub-hourly scheduling. This study presents a hybrid deep learning model combining convolutional neural networks (CNN) with long short-term memory (LSTM) networks to forecast SI 60 min ahead. The model uses images from the Karlsruhe low-cost sky imager (KALiSI) alongside meteorological data from a local weather station. Correlation analysis identifies key meteorological variables associated with global horizontal irradiance, which are integrated with sky images to improve forecasting accuracy. Experimental results show that including relative humidity (RH) and wind direction (WD) yields the lowest average root mean square (RMS) error of 119 W/m^2 and mean absolute error (MA) error of 82 W/m^2. This outperforms image-only and persistence models by 3 % and 42 % in RMS error (9 % and 50 % in MA error), respectively. Adding more meteorological variables does not enhance accuracy and may sometimes degrade performance, emphasising the importance of informed feature selection. At longer horizons (up to 120 min), all models show increased errors, with persistence performing worst (normalised RMS error 60 %, normalised MA error 47 %). In contrast, the image-only and RH+WD models maintain more stable errors (normalised RMS error 34–35 %, normalised MA error < 27 %). These results pave the way for improved forecasting for solar power production.

1. Introduction

1.1. Background

Increasing reliance on fossil fuels, resulting in environmental pollution and greenhouse gas emissions, has led to a surge in interest in the use of renewable energy (RE) sources as a means of protecting the environment on a global scale. Amongst these RE technologies, solar energy technologies, particularly photovoltaic (PV) power generation, stand out as a promising option in the pursuit of environmental protection and meeting energy demands globally (International Energy Agency, 2024). However, the availability of exploitable solar energy is highly dependent on solar irradiance (SI), which exhibits inherent fluctuations and the availability being greatly influenced by various weather conditions (Nikodinoska et al., 2022; Krechowicz et al., 2022). The inherent variability in SI can lead to significant fluctuations in power generation resulting in technical challenges in grids with high PV power integration, leading to impacts on power quality (Pandžić and Capuder, 2024). Thus, to mitigate the effects of SI variability, forecasting techniques are essential, enabling grid operators to more accurately predict solar energy output and improve energy resource planning and management.

SI forecasting can be pursued by a variety of approaches based on the time horizon they address: i) very short-term forecasting, or nowcasting, which covers periods from a few seconds to about 30 mins; ii)

* Corresponding author at: Institute of Microstructure Technology, Karlsruhe Institute of Technology, Hermann-von-Helmholtz-Platz 1, Eggenstein-Leopoldshafen 76344, Germany
E-mail address: bryce.richards@kit.edu (B.S. Richards).

https://doi.org/10.1016/j.egyr.2025.11.027
Received 30 July 2025; Received in revised form 6 October 2025; Accepted 11 November 2025
Available online 18 November 2025

Nomenclature

Abbreviation

ANN	artificial neural network
ATM	atmospheric pressure
CNN	convolutional neural networks
DL	deep learning
FTP	File Transfer Protocol
GHI	global horizontal irradiance
KALiSI	Karlsruhe low-cost sky imager
KIT	Karlsruhe Institute Technology
LSTM	long short-term memory
MA	mean absolute
ML	machine learning
MLP	multilayer perceptron
MS	mean squared
nMA	normalised mean absolute
nRMS	normalised root mean square
NWP	numerical weather prediction
PV	photovoltaic
RH	relative humidity
RMS	root means square error
RF	Random Forest
RI	rain intensity
SI	solar irradiance
T	temperature
VGG	Visual Geometry Group

Symbols

GHI_p	predicted GHI
GHI_m	measured GHI
N	total number of times predictions are performed
t	time
WDsin	sine component of wind direction
WDcos	cosine component of wind direction

short-term forecasting, spanning from over half an hour to 48 h; iii) medium-term forecasting, ranging from more than two days to a week; and iv) long-term forecasting, applicable to periods exceeding a week (Ansong et al., 2024). Additionally, these methods can be grouped according to the type of data or model used – numerical weather prediction (NWP) models, meteorological data-driven approaches, image-based approaches – each with distinct strengths and limitations. The choice of the forecasting model used depends on factors such as the specific application, data availability, accuracy and computational requirements. Combining different methods into a hybrid approach often enhances forecasting accuracy, by taking advantage of various techniques and thus consistently outperforming individual methods (Ansong et al., 2024). NWP models use physical equations describing fluid dynamics of the atmosphere (Martín et al., 2010; Mwigereri et al., 2025) to predict SI over large areas for hours to days. Meteorological data-driven methods rely on statistical models, such as autoregressive integrated moving average, and machine learning (ML) algorithms like artificial neural networks (ANN), random forest or support vector machine (Kumar et al., 2020). Image-based methods use images from satellites or ground-based sky imagers to forecast SI by extracting features such as cloud motion to build physical model or applying ML as well as deep learning (DL) models, to predict SI with or without historical data (Chow et al., 2011; Quesada-Ruiz et al., 2014; Marquez and Coimbra, 2013).

While ML and in particular DL, models can capture complex patterns and offer higher accuracies, these usually require large datasets and high computational power (Nie et al., 2024). Studies suggest that multi-year datasets spanning at least 3 – 5 years are required to capture both seasonal and interannual variability. Indeed, durations above 5 years have been reported beneficial as they enhance robustness and accuracy of DL models for SI forecasting from sky images by incorporating a wider range of weather conditions (Nie et al., 2023; Feng and Zhang, 2020).

Most SI forecasting approaches rely heavily on historical data and are effective for making longer term predictions ranging from hours to several days ahead, but may fail to capture localised and rapid fluctuations in SI caused by dynamic atmospheric conditions. NPW models, for example, typically operate on large areas (up to 10,000 km^2), have coarse special resolution (up to 1 km) and overlook small-scale cloud dynamics. Localised cloud formations can cause significant rapid changes in surface irradiance which such large-scale models typically fail to predict accurately (Samu et al., 2021; Jamal et al., 2018; Yang et al., 2021; Nie et al., 2024). This is especially important for PV power generation considering that a utility-scale solar farm would occupy roughly 1 km^2 of land area.

To address these challenges, the integration of high-resolution ground-based observations into forecast models is being explored, along with ML approaches, to improve SI prediction by capturing complex patterns in atmospheric data. These advanced methods aim to improve performance, particularly in capturing localised and rapid variations in SI. Ground-based sky cameras excel in very short-term predictions, capturing rapid fluctuations at typically 1-min intervals. Enabling forecast of rapidly changing weather conditions, such as cloud movement, which directly affects SI and is therefore crucial for more accurate estimations (Coimbra et al., 2013; West et al., 2014). An advantage of ground-based sky cameras over satellite images for localised analysis is their higher spatial and temporal resolution. Satellite images, with their coarser spatial resolution (typically 2–10 km) and lower temporal frequency (15 min.) with ideal forecast horizon in the range of 30 min to 6 h ahead (Blanc et al., 2017; Lorenz, 2012), will not detect small cloud variations that significantly impact SI. Studies have shown that the normalised root mean square (nRMS) error for satellite image-based forecasts ranges from 20 % to 50 % for a forecast horizon of 15–120 min (Miller et al., 2018; Wang et al., 2019; Yang et al., 2020; Yang et al., 2019). Ground-based sky imagers capture images of the sky at regular intervals (1-5 min or less), creating sequences that illustrate cloud dynamics, enabling the prediction of global horizontal irradiance (GHI) at very short time horizons (up to 30 min). Subsequently, PV power production can be estimated from these GHI predictions using a GHI-to-PV power conversion model (Hendrikx et al., 2024). The combination of sky images and ML models has shown promising results in enhancing the accuracy of SI predictions, particularly for time horizons ranging from a few minutes to about 30 min ahead (Bo et al., 2023; Song et al., 2022). For example, Caldas and Alonso-Suárez (2019) combined sky images with real-time SI measurements to predict 1-min averaged SI up to 10 min ahead with RMS errors of less than 6 W/m^2 for clear skies, 184 – 251 W/m^2 for cloudy skies and 64 – 110 W/m^2 for overcast skies. A long short-term memory (LSTM) model has also been used to predict the cloud cover fraction and, subsequently, the GHI, achieving a nRMS error of 25 – 40 % (Rajagukguk et al., 2021). Dimensionality reduction and ML algorithms (random forest and K-nearest neighbours) were employed for SI prediction up to 4 h ahead, with K-nearest neighbour achieving nRMS errors of 116 –122 (Al-lahham et al., 2020). The SolarNet, a 20-layer convolutional neural network (CNN) using only sky images, produced multi-step GHI forecasts with a 9 % nRMS error and a 25 % forecast skill score, outperforming persistence and six ML models based on six years of data at 10 min resolution. The first three years of data were used for model training, while the following one year and the last two years of data were used for validation and testing respectively (Feng and Zhang, 2020).

Furthermore, hybrid models that combine different DL models capable of accurately predicting very short-term SI can be developed. Hybrid models benefit from the strengths of each approach, and offer more comprehensive and robust predictions (Gala et al., 2016). A hybrid approach that combined a CNN and an ANN has been employed in forecasting GHI. The CNN extracted features from sky images, while the ANN processed time series weather data, including GHI and the sun's angular position. The outputs from these networks were merged and used as the input of another ANN for predicting the GHI up to 20 min in advance, achieving a forecast skill score between 20 % and 45 % (Paletta and Lasenby, 2020). In another approach, a CNN was used to encode features from sky-video frames, which were then processed by a two-tier LSTM network to produce SI forecasts up to 4 h ahead. This CNN-LSTM model performed better than a satellite-based forecast, achieving normalised mean absolute percentage (nMAP) error values of 8 – 47 %, in contrast to the satellite-based forecasts with nMAP error values of 78 – 116 % (Siddiqui et al., 2019). Hybrid DL models combining CNNs and recurrent neural networks like LSTM are promising models for deploying SI forecasts via sky imaging and meteorological data. The advantage of these models is that they can simultaneously analyse sky images and meteorological data. The combination of feature extraction by CNN from the images and the ability of the LSTM to detect long-term dependencies can improve the performance of the implemented SI forecast models. While traditional meteorological data provides a foundation for SI forecasting, its limitations in spatial and temporal resolution necessitate the adoption of more advanced techniques. Incorporating high-resolution image data and advanced modelling approaches is essential for accurately capturing the localised variations in SI caused by rapidly changing atmospheric conditions. On the other hand, sky images alone may not capture cloud information outside the camera's view, making it difficult to predict cloud dynamics beyond the camera's range. One possibility for enhancing SI forecasting performance is via the addition of weather data to sky image-based SI forecasting models. For example, a deep CNN was applied to hemispherical sky images and weather variables for 2 – 20 min SI forecasting. The model predicted the 10-min ahead SI with forecast skill based on RMS error, with nRMS error reaching 40 %. The performance on the 10-min ahead forecast increased by 10 % when past data of the same day was included in the training. These emphasise the benefit of integrating previous same-day data in short-term forecasting, improving forecast skill relative to persistence model – a commonly used benchmark model which assumes the future values will be the same as occurred in the present time (Paletta et al., 2020). Similarly, a model combining a CNN and a multilayer perceptron (MLP) was proposed to predict SI 15 min ahead. This model used sky images, historical GHI data, and other weather variables including temperature (T), relative humidity (RH), wind speed (WS), and wind direction (WD) as inputs (El Alani et al., 2021). The RMS error achieved by CNN-MLP under various sky conditions (clear, mixed and overcast) was only half as large as the RMS error of the persistence model (El Alani et al., 2021). Another approach used sky images and meteorological data including RH and aerosol optical depth in an LSTM network, achieving ~15 % normalised RMS error for 10-min GHI forecasts (Zuo et al., 2022).

Atmospheric conditions such as T, RH, and atmospheric pressure (ATM) play a crucial role in cloud formation and radiation scattering that impact surface irradiance (Mayer et al., 2010). Incorporating these meteorological parameters refines SI predictions, particularly under mixed weather conditions (Demirtas et al., 2012). Wind speed and direction can also impact cloud movement, affecting SI patterns. Including wind data in forecasting models can enhance cloud motion predictions, further improving short-term irradiance estimates (Huang et al., 2013). Recent studies have shown that including these weather data in sky image-based forecasting can results in more reliable very short-term solar predictions with RMS error reductions of up to 60 –70 % compared to persistence models (El Alani et al., 2021) and forecast skill score improvements of ~10 – 25 % depending on forecast horizon and sky conditions (Paletta and Lasenby, 2020).

1.2. Motivation

Conventional forecasting method usually rely on meteorological data alone, which can be insufficient for capturing the localised variations in SI caused by rapidly changing atmospheric conditions (Yang et al., 2014; Chow et al., 2015). Despite recent advancements in use of sky imaging techniques to characterise cloud cover and other atmospheric phenomena, integrating sky images with ML techniques for SI forecasting still presents several challenges (Paletta et al., 2023b). These include the need for high-quality and large-scale image datasets, the development of robust models capable of extracting and interpreting complex image features, and the effective fusion of image-derived information with traditional meteorological data (Nie et al., 2023; Paletta et al., 2023a, 2023b). Most recent studies combining all-sky imagery with meteorological variables target very short-term horizons, usually limited to only up to 30 min ahead (Zuo et al., 2022, 2023; El Alani et al., 2021; Straub et al., 2024; Zhang et al., 2023). In contrast, relatively few investigations have examined horizons beyond 1 h and even pushing the envelope up to 2 h. Furthermore, the performance of SI forecasting based solely on sky images tends to degrade beyond 30-min forecast horizons (Dev et al., 2016; Paletta et al., 2023b), leaving a clear research gap for methods capable of realising accurate $\geq$ 60-minute-ahead predictions. These limitations prompt the need to explore hybrid approaches that integrate meteorological variables, such as T, RH, and WS, with ML techniques to improve predictive accuracy at these horizons. Such multi-modal fusion approaches have the potential to significantly enhance forecasting performance over extended horizons (Diagne et al., 2013; Kumar et al., 2020). One promising methodology involves the deployment of parallel network architectures, where separate networks are used to process the different data inputs. Typically, a CNN is employed to extract spatial features from sky images, while other networks, such as a LSTM, can handle sequential meteorological time-series data. The features extracted by each branch are subsequently concatenated and fed into a fully connected network, which produces the final irradiance forecast (Paletta et al., 2021).

Building on the knowledge from the above, this work uses a CNN inspired by the Visual Geometry Group's (VGG) VGG16 network architecture (Simonyan and Zisserman, 2014) within a hybrid CNN-LSTM framework for 60–120-min ahead SI forecasting. This hybrid CNN-LSTM architecture is proposed, to explore the impact on model performance of different combinations of meteorological data - such as T, RH, WS, and WD – with the images obtained from a custom-built low-cost sky imager. The 60–120 min ahead SI forecasting represents a key short-term prediction horizon, critical for the reliable operation of solar power plants, grid integration, energy trading, and storage optimisation. It bridges the gap between very short-term nowcasting (0–30 min) and longer-term forecasting (several hours to days ahead). This horizon offers sufficient lead time for operational decisions while maintaining relatively high forecast accuracy (Gevorgian and Booth, 2013; Kremer et al., 2021). Within this timeframe, grid operators and energy managers can optimise energy dispatch, manage backup generation, and adjust storage charging or discharging schedules (Pedro and Coimbra, 2012). Additionally, short-term forecasts are integral to regulatory compliance in electricity markets where hourly or sub-hourly scheduling is mandatory (Märkle-Huß et al., 2018; Gevorgian and Booth, 2013). The contributions of this study are as follows:

i) The data from a low-cost, custom-built, sky imager (KALiSI) is integrated with local meteorological variables to develop a hybrid DL model for $\geq$ 60-minute-ahead SI forecasting.
ii) The inclusion of meteorological variables (e.g., T, RH, WS, and WD) is systematically investigated alongside sky images as to how this affects the overall model performance, particularly

beyond the 30-min horizon where image-only models are known to degrade.

iii) A comparative evaluation across multiple forecast horizons (30–120 min) is provided to quantify the benefits of multi-source data fusion on forecasting accuracy.

Accordingly, this work addresses the following research questions:

i) How does a hybrid DL model using only sky image sequences perform (measured by RMS error and MA error) when forecasting SI at 60-min forecast horizon?
ii) How does the inclusion of meteorological variables (e.g., T, RH, WS, and WD) alongside the sky images improve SI forecasting the performance at 60 min horizon, considering the known degradation of image-based models beyond 30 min?
iii) How do selected meteorological variables, when used alongside sky images, affect SI forecasting performance across varying forecast horizons (e.g., from 30 to 120 min)?

2. Materials and methods

2.1. Data collection

Sky images and meteorological data are collected at the Karlsruhe Institute Technology (KIT) Solar Park in Karlsruhe, Germany (49°05'56.2"N 8°26'14.5"E) from March to December 2024 at 1 min intervals from sunrise to sunset. The images were acquired with the Karlsruhe low-cost sky imager (KALiSI), built around the RaspberryPi single board computer equipped with an 8-megapixel camera module and a fish eye lens with 220° viewing angle, as described in more detail in a previous work (Ansong et al., 2025a). This earlier work also demonstrated that the performance of custom-built KALiSI compared favourably against a commercially available sky imager (EKO Instruments, CMS-Schreder, ASI-16, Austria). The GHI is measured with a silicon-based SI sensor (IMT Technology GmbH, Si-V-1.5TC-T, Germany) integrated into KALiSI. The data from the KALiSI is logged directly onto an external storage disk integrated into the system and are subsequently transferred via the internet to a File Transfer Protocol (FTP) server for archiving and remote access. The other meteorological data, including ambient T, RH, WS, WD, ATM and rain intensity (RI), are also collected at 1 min resolution from instruments mounted on the same weather station mast as the KALiSI, as shown in Fig. 1. Details of the instrumentation are presented in Table 1. The sensors are connected to a data logger (OTT Hydromet GmbH, NetDL 1000, Kempten, Germany) for continuous data logging. The recorded data are also transmitted to an FTP server for remote access.

2.2. Data processing

With regard to image data processing, firstly the edges of the images were cropped out as they contained no relevant information about cloud coverage. Secondly, low-quality images, such as those that were blurred or included unwanted elements (e.g. insects or birds) occupying a significant part of the image, were manually removed through visual inspection. Fig. 2 illustrates sample images from the KALiSI with their corresponding pre-processed images. It can be seen that the pre-processing improves the visual quality and uniformity of the sky images collected from KALiSI.

These adjustments play a vital role in boosting the model's ability to learn relevant patterns in the images for SI prediction. While the use of higher-resolution image in training has the potential to improve model performance, it also comes with greater computational requirements. The pixels of the pre-processed images were therefore down-sampled to 64 × 64 pixels and scaled between 0 and 1 by dividing each pixel by 255 to make the training more efficient (Sun et al., 2018). For the weather data, the correlations between the individual parameters and GHI weres determined using the Pearson correlation coefficient. Fig. 3 shows the matrix of Pearson correlation coefficients between the six weather variables and the GHI, revealing a moderate positive correlation with T and a moderate negative correlation with RH.

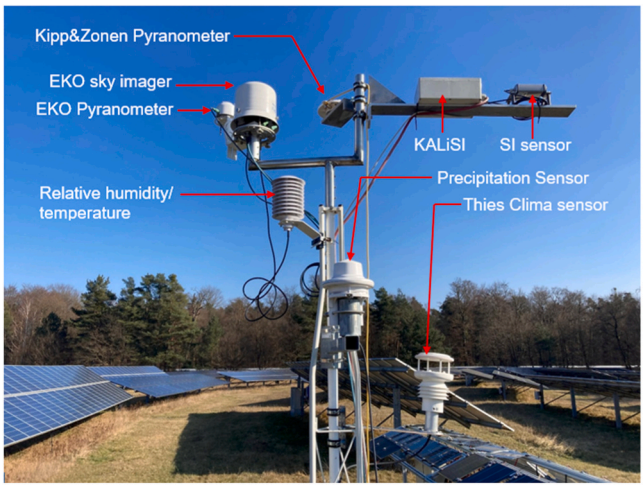

Fig. 1. KALiSI mounted on the weather station mast at the KIT Solar Park together with other sensors including two pyranometers (EKO Instruments, MS-80SH, Japan and Kipp & Zonen, SMP12, The Netherlands) and, the EKO sky imager (EKO Instruments, CMS-Schreder, ASI-16, Austria, 5 MP, 180° field of view). The EKO sky imager, although not used in the current work has been used as a reference in a previous work (Ansong et al., 2025a) which also introduces KALiSI. The arm on which KALiSI is mounted is facing due south and mounted as high as possible to minimise both shading and objects obscuring the images of the hemisphere.

These adjustments play a vital role in boosting the model's ability to learn relevant patterns in the images for SI prediction. While the use of higher-resolution image in training has the potential to improve model performance, it also comes with greater computational requirements. The pixels of the pre-processed images were therefore down-sampled to 64 × 64 pixels and scaled between 0 and 1 by dividing each pixel by 255 to make the training more efficient (Sun et al., 2018). For the weather data, the correlation between the individual parameters was determined and GHI was determined using the Pearson correlation coefficient. Fig. 3 shows the matrix of Pearson correlation coefficients between the six weather variables and the GHI, revealing a moderate positive correlation with T and a moderate negative correlation with RH.

Previous studies suggested variables with a Pearson correlation coefficient greater than 0.2 are considered significant for GHI predictions (Gao et al., 2019; Voyant et al., 2011). However, in this work, combinations of T and RH with the other variables (WS, WD, ATM) with the exception of RI were used with the images. This approach provides a more comprehensive understanding of how multiple factors collectively influence SI predictions, since variables with correlation coefficient below the 0.2 threshold could still have a meaningful indirect impact in combination with the images and other variables with higher coefficient, depending on local climatic conditions. For example, the correlation between WD and GHI is linear and does not capture spatiotemporal dependencies. WD, although weakly correlated in a linear sense, could contribute to critical cloud advection directionality when paired with sky imagery. Local cloud properties, aerosols, and humidity can modulate the amount of SI reaching the surface, thereby affecting its relationship with other weather variables (Hou et al., 2021; Wild et al., 2015). For example, although ATM has very low negative correlation with GHI (-0.04), it has an acceptable negative correlation with temperature (-0.27) and could indirectly influence the prediction of SI when used in combination with temperature. These have motivated the inclusion of ATM, as its indirect relationship with GHI through temperature was expected to enhance predictive skill. Rain intensity

Table 1
List of weather variables and sensor descriptions mounted on the weather mast at the KIT Solar Park. Note, only the instruments marked (*) were used in this work.

Variable	ID	Instrument	Value range	Accuracy
Ambient temperature		Thies Clima sensor US, NHTFB, Germany	−30 − + 80 °C	± 0.3 K@ 25° C, ± 0.1 K above −40…+ 80° C
Temperature/ Relative humidity *	RH, T	EKO, CMS-SChredder, TRH40-STA T/ RH Sensor, Austria	−30 − + 70 °C 0 −100 % rH	±0.2°C (−40…+ 70°C) ±0.1°C (+20…+ 60°C) ±1.5 % @ 25°C, 0…80 % RH) ±2 % (at 25°C, 80…100 % RH)
Relative humidity		Thies Clima sensor US, NHTFB, Germany	0 −100 % rH	±1.8 % of 10… 90 %, ±3.0 % of 0… 10 %
Rain intensity*	RI	Lufft, WS100, radar sensor, Germany	0.01 − 200 mm/h	±0.16 mm or ±10 %
Atmospheric pressure*	ATM	Thies Clima sensor US, NHTFB, Germany	300 − 1100 hPa	±0.25hPa @ + 10…+ 35 ° C, ±1hPa @ −10…+ 60° C
Wind speed *	WS	Thies Clima sensor US, NHTFB, Germany	0.01 − 60 m/s	5 m/s: ±0.3 m/s [rms-mean over 360 °] 5–60 m/s: ±3 % of measured value [rms-mean over 360 °]
Wind direction*	WD	Thies Clima sensor US, NHTFB, Germany	0 − 360°	± 2.0° @ wind speed > 2 m/s
Solar irradiance		Kipp & Zonen, SMP12, The Netherlands	0 − 2000 W/m^2	Zero offset A < ±1 W/m^2 Zero offset B < ±1.5 W/m^2
Solar irradiance		EKO, MS−80SH, Japan	0 − 4000 W/m^2	Zero Offset A < 1 W/m^2 Zero Offset B ± 1 W/m^2

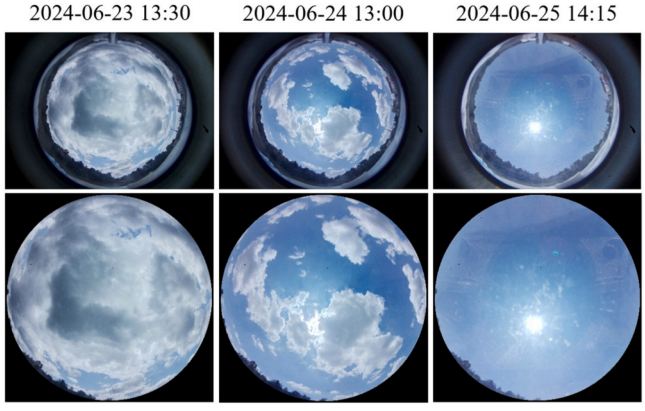

Fig. 2. Images from KALiSI with time stamp (top row) and corresponding pre-processed images (bottom row).

does not have an acceptable correlation with both T or RH and is therefore not used further in this study.

The T, RH, ATM, and WS data were also normalised between 0 and 1 to ensure consistency and comparability across variables. However, due

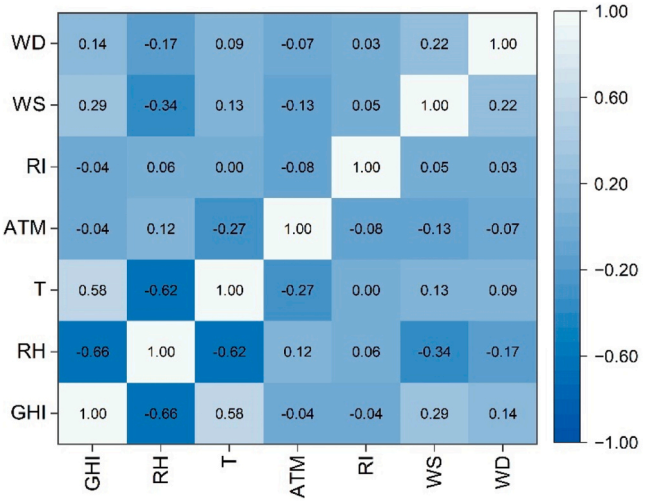

Fig. 3. Pearson correlation coefficient of weather variables: relative humidity (RH), Temperature (T), atmospheric pressure (ATM), Rain intensity (RI), wind speed (WS) and wind direction (WD) with global horizontal irradiance (GHI).

to the cyclic nature of WD data, a direct normalisation approach can lead to incorrect interpretations, as values close to 0° and 360° would be treated as numerically distant despite being nearly identical in direction. To preserve the continuous and circular characteristics of WD (in degrees), the data was encoded into their sine (WDsin) and cosine (WDcos) components (Kuang et al., 2021), resulting in two values for each wind direction observation, ranging between −1 and 1 using Eqs. (1) and (2), before normalisation to between 0 and 1. This transformation ensures that similar wind directions remain numerically close in the transformed space, preventing artificial discontinuities in the dataset. The measured GHI are also rescaled from W/m^2 to kW/m^2 to align with the other normalised data before training. The models' outputs were subsequently scaled back to W/m^2 for error analysis.

$$\mathrm{WDsin} = \sin\left(\frac{\pi}{180} \times WD\right) \tag{1}$$

$$\mathrm{WDcos} = \cos\left(\frac{\pi}{180} \times WD\right) \tag{2}$$

Each sky image was associated with its corresponding GHI measurement and other weather data. Images without corresponding GHI measurement and/or weather variables and GHI or weather data without corresponding sky images were discarded. Sky conditions were categorised using a two-step procedure. Firstly, each day was empirically classified as sunny (mean GHI 490 W/m^2), cloudy (mean GHI ≤ 200 W/m^2), or partly cloudy (mean GHI between 200 and 490 W/m^2) (Ansong et al., 2025b). Secondly, the automatically assigned labels were verified by visual inspection of the sky images and by checking the corresponding daily GHI profiles to ensure consistency with observed cloud cover and irradiance patterns. From this validated data, 9 days comprising of 3 each of clear, cloudy, and partly cloudy days were set aside as the test set (thus representative days for each weather category) and the remainder were used for training and validation of the model. The selection of these days was done to ensure a balanced comparison across varying sky conditions and to allow for a robust assessment of the model's ability to generalise across different sky and weather conditions rather than report large statistical evaluations across years. Comparable sky image-based solar forecasting studies (Hendrikx et al., 2024; Zuo et al., 2022; Sun et al., 2018; Nie et al., 2020) have reported using similar or even smaller test days when balancing sky conditions. The images are stacked into sequences based on the forecasting horizon. For instance, in the case of 60 min ahead forecast, a look-back period of 60 min is used. Given that one image is captured every minute, each training sample for the 60 min horizon consisted of 60 sequential

images. The corresponding meteorological data, recorded at the same frequency, are also stacked in a similar manner. The final dataset used after the processing consisted of observations shown in Table 2.

2.3. Forecasting model

2.3.1. Model architecture

In this hybrid approach, a CNN was combined with an LSTM network. The CNN architecture is based on the model developed by the VGG, popularly known as the VGG16 model (Simonyan and Zisserman, 2014). Fig. 4(a), shows the model consists of five feature learning blocks, each comprising two or three convolutional layers with identical parameters followed by a max-pooling layer. The first block uses 64 filters, the second 128 filters, the third 256 filters, and the fourth and fifth blocks use 512 filters each. All convolutional layers employ a kernel size of 3×3 with 'same' padding. The convolutional layers are followed by a flattening layer, which is used to vectorize the feature maps, followed by a reshape layer that adjusts the output into a format compatible with the LSTM input requirements. The LSTM layer contains 32 units to capture temporal dependencies in the input sequence. After the LSTM, the model includes three fully connected (dense) layers, the first and second with 1024 neurons each, and the final layer with a single neuron, which outputs the predicted SI value. A "ReLU" activation functions are applied to all convolutional and dense layers. To mitigate overfitting, a dropout rate of 0.3 is applied before each dense layer.

A similar VGG16-LSTM model, shown in Fig. 4(b) was built with both sky images and meteorological data as the input. The meteorological data are added to the image sequences to evaluate how they impact forecasting performance, without greatly modifying the underlying model structure. The additional meteorological data are passed through separate dense layers and LSTM network to process the temporal patterns in these variables independently. After processing the inputs, the model concatenates the outputs from the image, and weather data branches, creating a combined feature representation. This concatenated output is then passed through the fully connected layers with drop-outs before going through the single-unit dense layer to generate the GHI prediction output. The models are compared to understand any difference in SI forecasting performance caused by the inclusion of meteorological data in the model.

The convolutional component of the models follows the full VGG16 configuration, implemented from scratch and trained entirely on the present dataset without any pretrained weights. The original number of filters and layer depth are retained because VGG16 has consistently shown strong performance in prior sky-image and cloud-classification studies relevant to solar forecasting (El Alani et al., 2021; Feng and Zhang, 2020). Although the 64×64-pixel input images are smaller than in a typical computer-vision tasks, the lower spatial resolution substantially reduces the overall parameter count and computational cost, which enables the VGG16 to be trained efficiently. No formal architectural ablation was conducted, as the focus is to evaluate the benefit of combining sky images with meteorological features for short-term forecasting beyond 30 min ahead. Consequently, VGG16 serves as a well-validated and robust feature extractor providing a stable baseline for the proposed hybrid CNN–LSTM framework.

2.3.2. Model training

The VGG16-LSTM models are implemented using the Keras TensorFlow DL package and compiled with the Adam optimiser, chosen for its

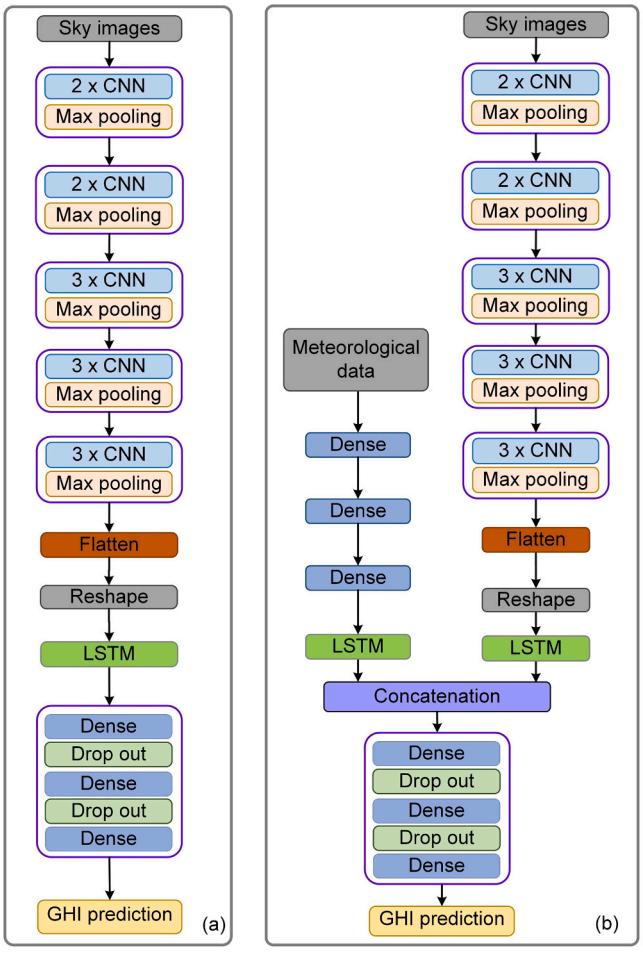

Fig. 4. Overview of the two different forecasting models used. (a) The image-only model – which takes sky images as input. The CNN processes the sky images to extract spatial features and feed into the LSTM (b) The model with meteorological data – which takes both sky images and weather data as input. The CNN processes the sky images to extract spatial features and feed into an LSTM, while the weather data are processed by the separate dense and LSTM layers as additional input. The extracted features from the images and meteorological data are then concatenated and fed into dense layers with drop outs to predict future GHI.

effectiveness in handling complex loss landscapes. The mean squared (MS) error given by Eq. (3), is used as the model's loss function. It measures the difference between the predicted and actual values in the dataset. During training, the model aims to minimise this loss function. A grid search is employed to find the most appropriate hyperparameters, including batch size, learning rate and drop outs, as shown in Table 3, to determine the optimal network architecture. Table 3 lists both the search ranges and the optimal values for the main hyperparameters, thus batch size, learning rate, and dropout rate. The grid search covered batch sizes of 64, 128, 256, and 512, learning rates from 1×10^{-3} to 1×10^{-7}, and dropout rates of 0.2–0.5, with the best configuration found to be a batch size of 256, a learning rate of 1×10^{-6},

Table 2
Training, validation and training samples for each forecast horizon.

Data split	20 min	30 min	60 min	90-min	120-min
Training	159,713	153,738	137,070	134,518	117,447
Validation	15,971	15,373	13,707	13,451	11,744
Testing	7, 720	7516	6916	6303	5673

Table 3
Optimisation hyperparameters.

Hyperparameter	Optimal value	Search space for optimal value
Batch size	256	64,128, 256, 512
Learning rate	1×10^{-6}	1×10^{-3},1×10^{-4}, 1×10^{-5}, 1×10^{-6}, 1×10^{-7}
Dropouts	0.3	0.2, 0.3, 0.4,0.5

and a dropout of 0.3. Early stopping was employed to track validation loss, halting the training process if the validation loss failed to improve for 10 consecutive epochs to avoid overfitting.

$$\text{MSerror} = \frac{1}{N}\sum_{i=1}^{N}\left(y_{p,i} - y_{m,i}\right)^2 \tag{3}$$

where $y_{p,i}$ and $y_{m,i}$ are the respective predicted and measured values, and N is the total number of data points.

2.3.3. Model evaluation

The RMS error and MA error defined in Eqs. (4) and (5), respectively, are used to evaluate the accuracy of model. RMS error and MA error are most widely used metric in SI forecast evaluation. To allow for fair comparison across different datasets of varying scales or different models (Paulescu and Paulescu, 2019), the RMS error and MA error are usually normalised to the mean of the measured data using Eqs. (6) and (7) (Hoff et al., 2013).

$$\text{RMSerror} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(GHI_{p,i} - GHI_{m,i}\right)^2} \tag{4}$$

$$\text{MAerror} = \frac{1}{N}\sum_{i=1}^{N}\left|GHI_{p,i} - GHI_{m,i}\right| \tag{5}$$

$$\text{nRMSerror} = \frac{\sqrt{N}}{\sum_{i=1}^{N} GHI_{m,i}}\sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(GHI_{p,i} - GHI_{m,i}\right)^2} \tag{6}$$

$$\text{nMAerror} = \left(\frac{1}{\sum_{i=1}^{N} GHI_{m,i}}\right)\sum_{i=1}^{N}\left|GHI_{p,i} - GHI_{m,i}\right| \tag{7}$$

where $GHI_{p,i}$ and $GHI_{m,i}$ are the respective predicted and measured GHI values in W/m^2, and N is the total number of times predictions are performed. The benchmark algorithm used here is the persistence model which assumes that the SI value at time $t + 1$, will be the same as the SI value at time t (Diagne et al., 2013).

3. Results and discussion

This section presents and interprets the results obtained from the proposed models for SI forecasting. The forecasting performance of the image-only model at multiple horizons (20, 30, and 60 min) is first reported and compared with persistence model under varying weather conditions. Subsequently, the effect of including meteorological variables alongside sky images within the hybrid framework is examined and compared with other classical models under varying weather conditions, with a discussion of how this multimodal approach influences forecasting accuracy across different time horizons. Finally, an analyses of the broader implications of the results is presented, including the model's performance under different sky conditions.

3.1. Forecast results

The 20-min, 30-min and 60-min ahead GHI forecast is obtained from the output of the VGG16-LSTM image-only model, trained on 10 months of image data. The image-only model's forecast is compared against the persistence forecast to evaluate its performance under different weather conditions, classified as clear, partly cloudy, and cloudy days. The overall performance in terms of RMS error and MA error are summarised in Table 4, while the results of the comparison for the 60 min ahead forecast is shown in Fig. 5. The results of 20-min and 30-min horizon are shown in Sections S1 and S2 respectively, in the Supplementary Information. The VGG16-LSTM forecast consistently outperforms the persistence model across all weather conditions, particularly on clear days. As shown in Fig. 5(a–c). The VGG16-LSTM follows the measured GHI quite well and shows better accuracy. For example, for the 60 min ahead forecast the VGG16-LSTM model achieves an average RMS error and MA error values of 49 W/m^2 and 41 W/m^2, respectively on the clear days, compared to 149 W/m^2, and 139 W/m^2 for the persistence model – thus roughly a factor of three-times better. Under partly cloudy conditions, as depicted in Fig. 5(d–f), again the persistence does not align well with the measured GHI compared with the VGG16-LSTM model which maintains a more matching profile and captures the measured GHI quite well. The average RMS error and MA error values on these days for the 60 min ahead forecast is 137 W/m^2, and 87 W/m^2, respectively, for the VGG16-LSTM model, and 221 W/m^2, and 179 W/m^2, respectively for the persistence model – again, the former showing a big improvement. Similarly on the very cloudy days, illustrated in Fig. 5 (g–i) for the 60 min ahead forecast, the persistence prediction struggles to align with the measured GHI, showing significant deviations. In contrast, the VGG16-LSTM model smooths out most of the variations but still follows the overall trend much better, achieving RMS error, and MA error values of 186 W/m^2, and 142 W/m^2, respectively, compared to 243 W/m^2 and 179 W/m^2 for the persistence model. As expected, the model's performance decreases as weather conditions transition from clear days through partly cloudy days to cloudy days, showing the challenges of forecasting under highly variable conditions.

3.2. Model performance with varying meteorological input combinations

The impact of adding varying combinations of the meteorological parameters, including T, RH, WS, WD and ATM, as shown in Table 5, on the performance of the model for the 60 min ahead forecast was investigated. The evaluation was conducted over the nine test days, that comprise of three representative days each for clear, partly cloudy, and cloudy conditions. Fig. 6 presents boxplots illustrating the distribution of the RMS errors (red plots) and MA errors (blue plots) for the 60-min ahead GHI forecasts across multiple combination of meteorological input based 9 distinct forecast days, comprising 3 days each of clear, partly cloudy, and cloudy conditions. The boxplots give a summary of the spread, median, and outlier behaviour for each configuration that allows comparison beyond mean values. The diamond markers inside boxes represent the mean values and the black horizontal lines represent median values. The detailed numerical results, by sky condition and their overall averages are also provided in Table 6. The forecast results

Table 4
Performance of the models in terms of RMS error (W/m^2) and MA error (W/m^2).

Forecast model	Clear			Partly cloudy			Cloudy			Overall		
	+ 20	+ 30	+ 60	+ 20	+ 30	+ 60	+ 20	+ 30	+ 60	+ 20	+ 30	+ 60
RMS error												
VGG16-LSTM	**29**	**37**	**49**	**123**	**127**	**137**	**158**	**170**	**186**	**103**	**111**	**124**
Persistence	44	67	149	165	181	221	200	228	243	136	159	204
MA error												
VGG16-LSTM	**22**	**29**	**41**	**65**	**72**	**87**	**111**	**125**	**142**	**66**	**75**	**90**
Persistence	40	61	131	94	117	179	128	123	179	87	100	163

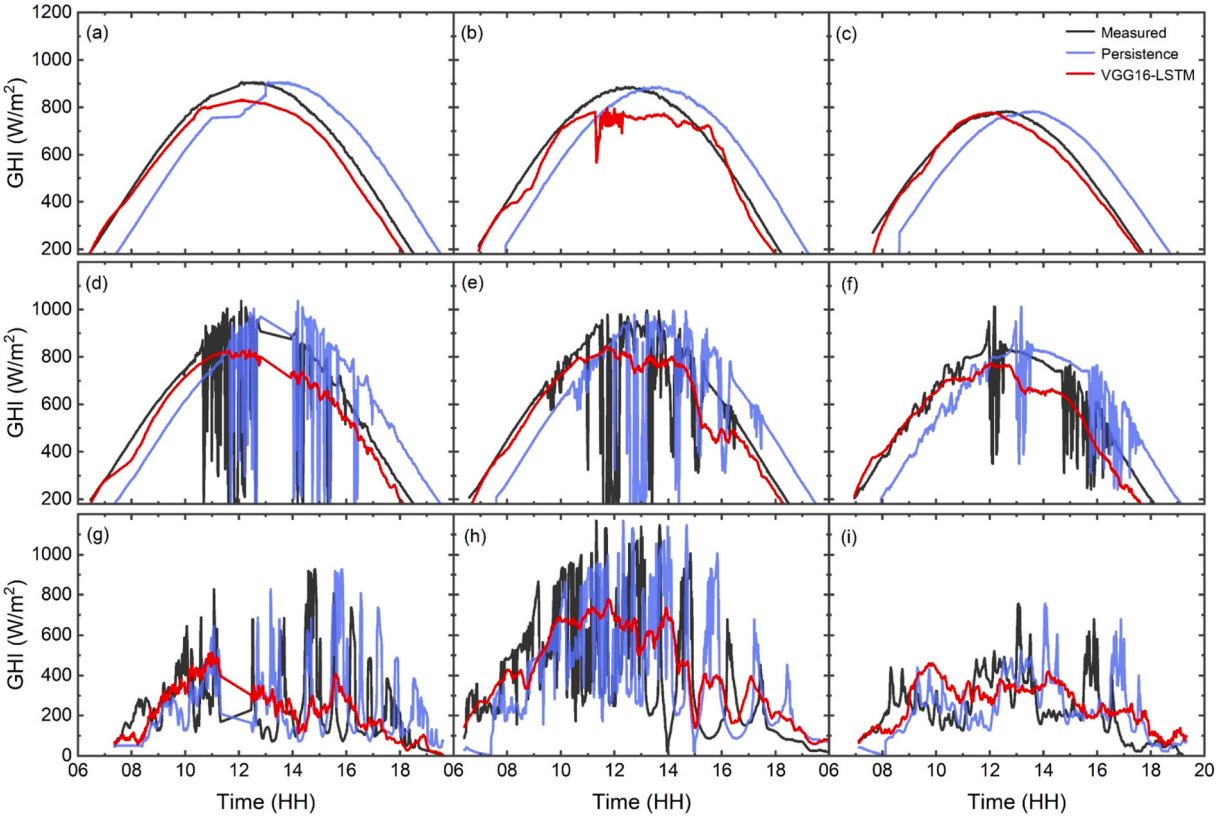

Fig. 5. Comparison between VGG16-LSTM for 60-min ahead forecast (without meteorological data), persistence forecast and the measured GHI; (a)-(c) clear days, (d)-(f) partly cloudy days and (g)-(i) cloudy days.

Table 5
Meteorological variable and combinations used as model input with the images.

Input cases	Parameters combination
1	T
2	RH
3	T + RH
4	T + WS
5	T + WD
6	T + ATM
7	T + RH+WS
8	T + RH+WD
9	T + RH+WS+WD
10	RH+WS
11	RH+WD
12	RH+WS+WD
13	RH +T + WS+WD+ATM

for the various combinations of the meteorological parameters are shown in Section 3 to 15 in the Supplementary Information.

Among all tested configurations, the RH+WD input combination achieved the lowest overall average RMS error (119 W/m^2) and MA error (82 W/m^2), as shown in Table 6. This suggests strong general performance, in terms of average error, despite its simplicity. However, it is important to note that Fig. 6 reveals left-skewed error distributions for all configurations. This implies that average or median values alone may not fully capture robustness, as they can be disproportionately influenced by low-error outliers. Therefore, while RH+WD performs well on average, it does not consistently outperform every alternative across all sky conditions or evaluation metrics, other combinations like RH+WS+WD may offer greater prediction stability, with tighter interquartile ranges. For example, under cloudy or partly cloudy skies, other combinations perform slightly better. The RH+WD input configuration is therefore, considered to be a competitive and efficient configuration, but not universally optimal. The relatively low performance of configurations that include both RH and T may be partly due to multicollinearity. These two variables are often strongly inversely correlated in atmospheric conditions, which can inflate the variance inflation factor (VIF) and introduce instability in the learning process. VIF analysis on the training dataset indicated that RH and T exhibited the highest VIF values (3.7 and 3.2, respectively), which demonstrates moderate multicollinearity. In contrast, WD and WS showed relatively low VIFs (1.6 and 1.9 respectively). These imply that including both RH and T within the same feature set may inflate variance and compromise model stability, whereas WS and WD remain well below commonly accepted multicollinearity thresholds (VIF < 3). This likely explains why RH+WD provides a competitive and computationally efficient feature set without introducing redundant information.

RH+WD achieved the lowest RMS error (33 W/m^2) and MA error (18 W/m^2) under sunny conditions, which is where sky imagery is most informative. Under partly cloudy and cloudy conditions, other combinations such as RH+WS+WD and RH+WS performed slightly better and may offer a better balance of physical relevance and statistical independence. These configurations avoid known collinearity issues (e.g., between RH and T) and includes features that together represent key atmospheric dynamics – moisture content and cloud motion. Their tighter error distribution also suggests greater robustness, making them strong candidates for further consideration. Nonetheless, RH+WD remained highly competitive with RMS error /MA error values of 87/78 W/m^2 and 146/119 W/m^2, respectively. These results suggest that RH and WD effectively capture essential atmospheric drivers of irradiance variability, such as cloud motion and moisture content.

The meteorological relevance of these features is further supported by the Pearson correlation analysis in Fig. 2, where RH showed the strongest negative correlation with GHI (r = -0.66), consistent with its

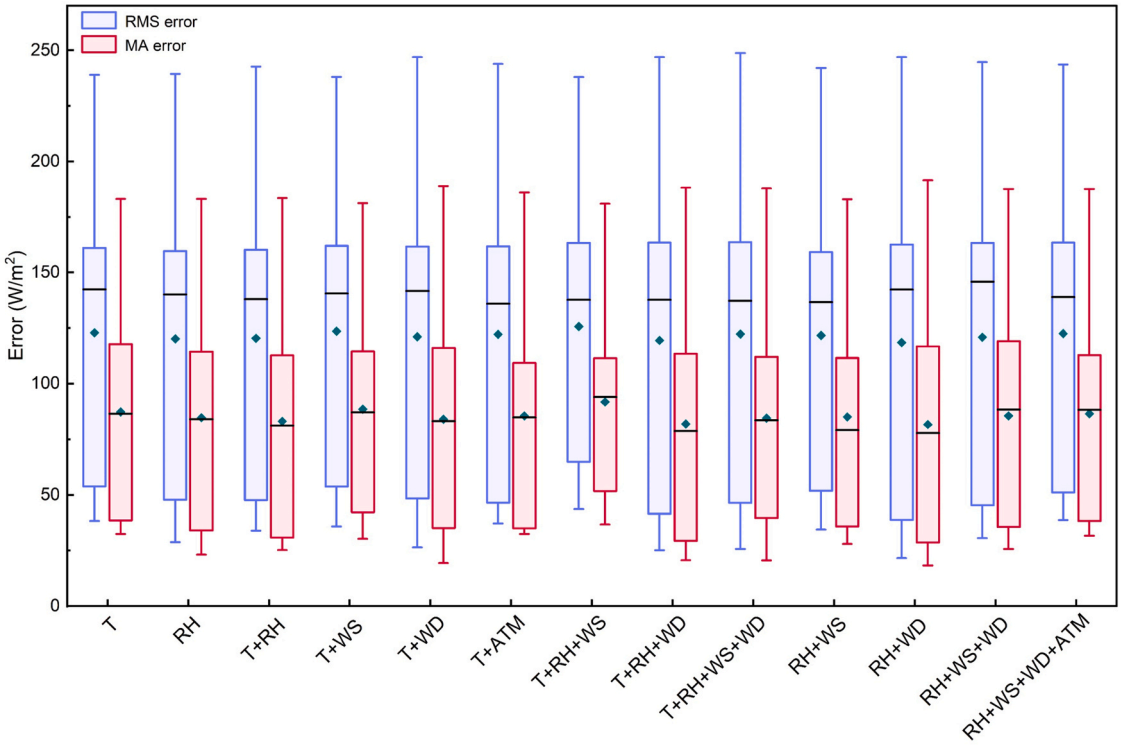

Fig. 6. Boxplot distributions of the RMS error and MA error, for the 60-min-ahead GHI forecasts across various meteorological input combinations of the VGG16-LSTM model. Each box represents 9 data points derived from 9 forecast days: 3 clear, 3 partly cloudy, and 3 cloudy sky conditions. The boxes show the interquartile range, with the black horizontal line indicating the median and the diamond symbol denoting the mean. Whiskers extend to the minimum and maximum values excluding outliers. The RH + WD model configuration achieves the lowest mean and median errors overall, but the distributions reveal noticeable skewness, suggesting variability in performance across different sky conditions.

role in cloud formation and atmospheric turbidity. T exhibited a moderately strong positive correlation (r = 0.58), while WS and WD showed weaker correlations (r = 0.29 and 0.14, respectively). Interestingly, despite its weak linear correlation, WD's combination with RH yielded the lowest overall error values, likely due to the model's ability to learn non-linear and spatiotemporal dependencies from sky images. WD likely improves the model's understanding of cloud movement, enhancing short-term irradiance forecasting. These suggest that RH and WD can be highly complementary features to sky imagery, especially in the given geographic and meteorological context.

Adding more variables such as T, WS, or ATM did not improve performance and in some cases slightly degraded accuracy (e.g., RH+T + WS+WD+ATM yielded RMS error of 123 W/m^2 and MA error of 86 W/m^2). Several factors may explain this: i) feature redundancy, such that variables that correlate well with GHI like RH and T, may provide overlapping information already visible in sky images, so, adding it might introduce redundancy; ii) variable with moderate correlation but may not always be reliable e.g., WS, which may not always move clouds at the camera level or might move thin clouds that do not block irradiance significantly; iii) variable with very low correlation and probably add noise. e.g. ATM, though corelates well with T and could have indirectly influence GHI in combination with T, do not improve the model's understanding of cloud cover or movement; iv) model overfitting, where, larger input dimensionality can reduce generalisation on unseen data, especially with limited training samples per weather type. These observations are consistent with previous studies (Diagne et al., 2013; Yan et al., 2020), which emphasise that selecting a small, physically meaningful subset of inputs can improve both learning efficiency and model interpretability in hybrid DL frameworks. Additionally, recent works show that deep CNN hybrid architectures for very short-term SI forecasts from sky imagery and meteorological variables often yields impressive performance (Zuo et al., 2022; El Alani et al., 2021; Jakoplić et al., 2023). Benchmarking and comparative reviews indicate that VGG-style encoders remain a widely-used and reliable feature extractor for sky imagery. Notable amongst them is SolarNet, a 20-layer deep CNN (Feng et al., 2020) and other VGG-inspired models (El Alani et al., 2021). Probabilistic and ensemble approaches, which combine all-sky imagers, satellite products, and persistence/satellite blending, tend to provide superior robustness for operational applications, particularly when multi-site or multi-sensor fusion is available (Paletta et al., 2023a; López-Cuesta et al., 2023; Straub et al., 2024).

Compared to these studies, the present work demonstrates a favourable accuracy-complexity trade-off. It achieves competitive forecast error for 60-min horizon while using down sampled images (64 × 64 pixels) and only a selective set of usually readily available meteorological variables. This reduces sensor requirements and simplifies deployment, although some limitations remain in single-site validation and multi-season benchmarking. Furthermore, the performance of the proposed hybrid VGG16–LSTM model with the selective RH+WD feature set is consistent with comparable short-term SI forecasts. Recent DL models that fuse sky images with meteorological variables, as stated earlier, often target very-short-term horizons (up to 30 min), and report RMS error values in the range of ∼50–100 W/m^2 over such durations (Shirazi et al., 2024; Chaaraoui et al., 2024; Hendrikx et al., 2024). Models that attempt 60-min ahead forecasting from sky images alone show higher nRMS errors of 15–35 % (Feng et al., 2020; Ansong et al., 2025a). Purely image-based CNN-LSTM approaches generally exhibit even larger errors (Barancsuk et al., 2024), confirming the benefit of carefully selected meteorological predictors for horizons beyond.

In addition to persistence, the RH+WD configuration for the 60-min horizon was compared with two random forest (RF) models, using either all meteorological variables or the selected RH and WD as inputs. Table 7 summarises the results for different sky conditions, Both RF models show higher errors than the VGG16–LSTM model with images

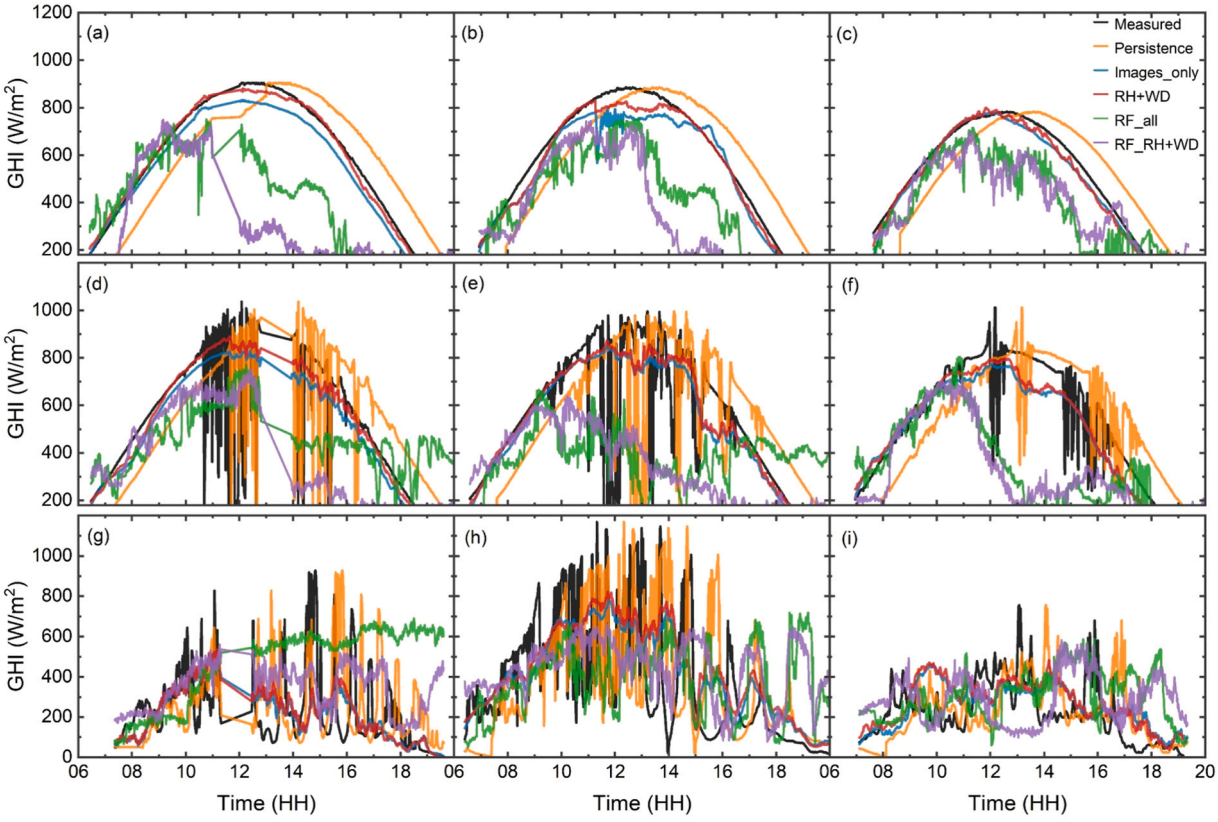

Fig. 7. Forecasting results on selected test days for multiple models, including RH+WD (VGG16–LSTM with images and RH+WD), Image-only (VGG16–LSTM with images only), RF_RH+WD (RF with RH and WD), RF_all (RF with T, RH, WS, and WD), and Persistence compared with measured GHI under: (a–c) clear, (d–f) partly cloudy, and (g–j) cloudy days.

Table 6
RMS (W/m^2) and MA errors (W/m^2) for 60-min ahead GHI forecasts using different combinations of meteorological parameters combined with sky images across three weather conditions. Each row corresponds to one day of the particular sky conditions. The results in bold lowest errors for each sky condition.

Parameter Days	T	RH	T, RH	T, WS	T, WD	T, ATM	T, RH, WS	T, RH, WD	T, RH, WS, WD	RH, WS	RH, WD	RH, WS, WD	RH, T, WS, WD, ATM
RMS error													
Sunny	41	39	34	44	26	44	54	25	26	41	**22**	33	39
	52	46	44	51	48	47	63	41	47	52	**39**	45	47
	38	**29**	35	36	35	37	44	34	46	34	**29**	31	43
Partly	161	**160**	**160**	162	162	162	163	164	164	159	163	163	164
cloudy	156	155	156	156	155	158	157	155	155	**154**	157	157	155
	95	90	93	97	92	97	95	93	97	96	**87**	87	93
Cloudy	180	183	178	184	182	**176**	178	179	182	181	181	181	176
	239	239	243	**238**	247	244	**238**	247	249	242	247	245	244
	142	140	138	141	142	**136**	138	138	137	137	142	146	139
Average	123	120	120	123	121	122	125	120	123	122	**119**	121	122
MA error													
Sunny	32	31	25	37	19	35	46	21	20	31	**18**	29	32
	37	33	30	40	35	33	50	**29**	34	36	**29**	36	36
	33	23	30	30	30	32	37	29	40	28	**24**	26	37
Partly	93	89	86	98	83	89	106	79	84	92	**78**	88	88
cloudy	87	84	81	87	87	85	94	81	84	79	**84**	89	90
	72	68	69	75	69	75	72	69	72	72	**66**	65	69
Cloudy	130	136	129	130	128	125	128	**127**	131	132	129	130	124
	183	183	183	181	189	186	**181**	188	188	183	191	188	188
	118	114	113	115	116	109	**112**	113	**112**	**112**	117	119	113
Average	87	85	83	88	84	86	92	**82**	85	85	**82**	86	86

and (RH+WD). The RF models generally outperform persistence, although persistence performs better than RF in certain regimes. The VGG16–LSTM (RH+WD) achieves the lowest RMS error and MA error across all weather types, which confirms that the inclusion of selected meteorological data with sky images provides more informative representations of cloud dynamics and irradiance than meteorological features alone. Fig. 6 illustrates the forecasts results for the 60 min ahead of the VGG16-LSTM (RH+WD) model compared with the persistent and other RF models.

Table 7
Comparison of results of different models using for the 60-min GHI forecast.

Forecast model	Clear	Partly cloudy	Cloudy	Overall
RMS error (W/m^2)				
VGG16-LSTM (RH+D)	30	136	190	119
RF(RH+WD)	270	286	234	263
RF (all)	203	289	285	259
Persistence	149	221	243	207
MA error (W/m^2)				
VGG16-LSTM (RH+D)	24	76	142	81
RF (RH+WD)	214	219	194	209
RF (all)	171	228	233	211
Persistence	131	179	179	163

3.3. Performance comparison of forecast models across varying forecast horizon

The RH+WD configuration, which demonstrated competitive average performance was selected for further analysis. While it did not consistently exhibit the lowest median error across all weather scenarios, it achieved the lowest overall average RMS error (119 W/m^2) and MA error (82 W/m^2) in Table 6, with especially strong performance under sunny conditions. Its simplicity and strong mean performance make it a suitable candidate for extended evaluation. The associated performance metrics (RMS error and MA error) are summarised in Table 8. Across sky conditions and forecast horizons, the VGG16-LSTM model augmented with RH and WD generally outperforms both baselines. Specifically, for the 60-min horizon, the RH+WD model achieves an overall RMS error of 119 W/m^2, compared to 123 W/m^2 for the image-only model and 204 W/m^2 for persistence. This represents a 3 % improvement over the image-only model and a 42 % improvement over persistence. Similarly, the RH+WD configuration yields the lowest MA error values across most sky conditions and forecast horizons. For 60-min forecasts, the average MA error is 82 W/m^2, compared to 90 W/m^2 for the image-only model and 163 W/m^2 for persistence, equating to 9 % and 50 % improvements, respectively. To evaluate the statistical significance of the performance gains over the image-only model, separate paired tests were conducted on the per-sample absolute errors (to assess MA error differences) and on the per-sample squared errors (to assess RMS error differences). The RH+WD configuration achieved a significantly lower MA error ($p << 0.001$), corresponding to a 9 % reduction relative to the image-only model, whereas the 3 % reduction in RMS error was not statistically significant ($p = 0.14$). This is because RMS error is more sensitive to occasional large errors, which can increase the variability of the squared-error differences across samples, reducing the statistical significance of the improvement even when most errors are consistently smaller.

The forecast performance of RH+WD configuration, was also evaluated for horizons of 90 and 120 min to assess the model's stability beyond the 60 min. As expected, both RMS error and MA error increased steadily with lead time across all sky conditions. The RMS error rose from 119 W/m^2 at 60 min to 125 W/m^2 at 90 min and 127 W/m^2 at 120 min, while MA error increased from 82 W/m^2 to 93 W/m^2 and 92 W/m^2, respectively. This reveals the growing uncertainty in cloud advection and evolution, which reduces the relevance of the most recent sky-image features over longer horizons. Nevertheless, the hybrid VGG16-LSTM maintained a consistent advantage over the image only and persistence baselines even at 120 min. This indicates that adding sky-image data and meteorological (RH + WD) feature set improves performance compared to simple persistence, but the benefit decreases as the lead time increases.

The nRMS error and nMA error, which provide scale-independent evaluations, are presented in Fig. 8 and Table 9 across forecast horizons from 20 to 120 min. As expected, errors generally increase with forecast length. However, the persistence model shows the most significant degradation, reaching 60 % nRMS error and 47 % nMA error at the 120-min horizon. In contrast, both the RH + WD and image-only models maintain more stable performance, with nRMS error around 34–35 % and nMA error below 27 %. The RH+WD configuration slightly outperforms the image-only model, especially in nMA error, suggesting that meteorological features such as RH and WD can enhance the predictive value of sky imagery, particularly for longer horizons. Nonetheless, the relatively small margin between RH+WD and the

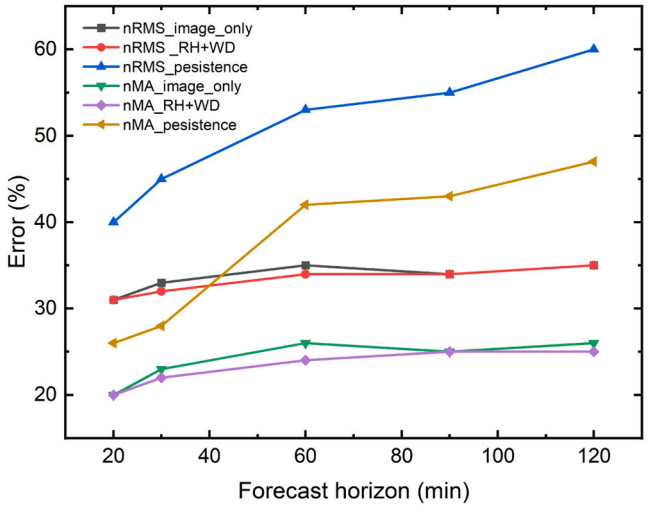

Fig. 8. Comparison of nRMS error and nMA error of the RH+WD model configuration with the image-only and persistence baseline models for different forecast horizons (20 –120 min).

Table 8
Performance of RH+WD model configuration in terms RMS error (W/m^2) and MA error (W/m^2) compared with image-only and persistence (Pers.) baseline models under different weather condition. The results in bold lowest errors in each horizon and sky condition.

Horizon	Clear			Partly cloudy			Cloudy			Overall		
	RH+WD	Image-only	Pers.	RH+WD	Image-only	Pers.	RH+WD	Image- only	Pers.	RH+WD	Image- only	Pers.
RMS error												
20 min	**25**	29	44	**52**	123	165	**113**	158	200	**65**	103	136
30 min	**24**	38	67	**66**	127	181	**125**	170	228	**72**	111	159
60 min	**30**	47	149	**136**	137	221	190	**186**	243	**119**	123	204
90 min	48	**47**	165	135	**134**	248	**192**	**192**	251	125	**124**	221
120 min	**44**	47	314	**141**	142	361	197	**193**	296	**127**	131	324
MA error												
20 min	25	**22**	40	**52**	65	94	113	**111**	128	**66**	**66**	87
30 min	**24**	30	61	**66**	72	117	**125**	**125**	123	**72**	76	100
60 min	**24**	39	131	**76**	87	179	146	**142**	179	**82**	90	163
90 min	41	**39**	134	87	**84**	196	**149**	150	192	93	**91**	174
120 min	**38**	49	143	**87**	97	217	152	**148**	208	**92**	98	189

Table 9
Performance of the RH+WD model configuration in terms nRMS error (%) and nMA error (%) compared with image-only and persistence baseline models under different weather condition. The results in bold lowest errors in each horizon and sky condition.

Horizon	Clear			Partly cloudy			Cloudy			Overall		
	RH+WD	Image-only	Pers.	RH+WD	Image-only	Pers.	RH+WD	Image- only	Pers.	RH+WD	Image- only	Pers.
nRMS error												
20 min	**5**	6	9	26	**25**	33	63	**62**	78	**31**	**31**	40
30 min	**6**	8	13	**26**	**26**	37	**65**	**65**	86	**32**	33	45
60 min	**6**	9	28	**27**	**27**	44	69	**68**	88	**34**	35	53
90 min	**9**	10	31	**26**	**26**	48	**68**	**68**	86	**34**	**34**	55
120 min	**8**	11	35	**27**	**27**	49	69	**68**	96	**35**	**35**	60
nMA error												
20 min	**4**	5	8	**13**	**13**	19	45	**43**	50	**20**	**20**	26
30 min	**5**	6	12	**13**	15	24	**48**	**48**	49	**22**	23	28
60 min	**4**	7	25	**15**	17	35	53	**52**	65	**24**	26	42
90 min	8	**7**	25	17	**16**	38	**52**	**52**	67	**25**	**25**	43
120 min	**7**	9	26	**16**	18	41	53	**52**	72	**25**	26	47

image-only model implies that sky images already encode much of the short-term predictive information, and that selected meteorological inputs can act as useful, but not essential, enhancements in hybrid forecasting architectures.

4. Limitations and outlook

This study establishes that combining sky images with a selective set of meteorological variables, specifically RH and WD, can improve short-term SI forecasting. While the results are promising, several factors limit the scope of the findings. The dataset was collected at a single mid-latitude location (Karlsruhe, Germany) from February to December 2024. Although this period spans late winter, spring–summer, and autumn, it does not represent a complete multi-year record and therefore cannot fully capture inter-annual variability or the full range of global climates. The test set is limited to nine representative days, three each classified as clear, partly cloudy, and cloudy. Although selected to ensure balanced representation of distinct sky conditions it limits statistical power of sky condition analysis, and rare atmospheric events are not captured. The proposed method also depends on co-located sky-imaging hardware and the availability of RH and WD measurements (as can be seen from Fig. 1). Sites lacking these instruments, or employing different calibration standards, may require additional preprocessing or model adaptation. A further limitation arises from the specific choice of RH and WD as meteorological inputs. RH is an indicator of cloud optical depth, while WD provides information on the advection of cloud systems, together, they capture key drivers of SI variability at the study site. However, these relationships are not universal. In tropical or arid regions, cloud formation mechanisms differ, and RH may not represent optical depth in the same way. Likewise, in high-latitude regions, complex wind patterns and low solar elevation angles may diminish the predictive value of WD. Thus, while the RH+WD feature set is physically motivated and effective in the present study, its broader applicability across diverse climatic zones remains to be validated.

Despite these constraints, the findings offer practical guidance and a clear path for future research and deployment. When resources are limited, priority should be given to measurement like RH and WD, while periodic model re-training should be conducted as new local data become available. Future studies still need to investigate alternative CNN architectures to quantify potential trade-offs between accuracy and computational efficiency. Furthermore, future investigations should evaluate the approach on multi-site, multi-year datasets that encompass a wide range of climates, explore transfer-learning and domain-adaptation methods to reduce the need for large site-specific datasets, and integrate additional real-time predictors such as cloud motion, and aerosol optical depth.

5. Conclusions

In this work, a hybrid deep learning model combining a VGG16-like CNN and an LSTM network was developed for 60-min ahead SI forecasting using images from the KALiSI. A correlation analysis using the Pearson coefficient was first conducted to evaluate the relationship between meteorological parameters (T, RH, ATM, RI, WS, and WD and GHI. Based on this analysis, the most relevant features, were selected and used alongside image data as model inputs. To assess the impact of these meteorological variables, various configurations of input combinations were also tested.

The GHI forecasts generated by the hybrid model using only the images from the KALiSI outperformed the persistence forecast. The VGG16-LSTM model achieved an average RMS error of 124 W/m^2 and MA error of 90 W/m^2 at a 60-min forecast horizon. In comparison, the corresponding persistence forecasts yield average RMS error and MA error values of 204 W/m^2 and 163 W/m^2 respectively. The inclusion of meteorological variables into a VGG16-LSTM inputs can marginally enhance forecast accuracy beyond what is achievable using sky images alone. These results show that augmenting the VGG16-LSTM model with meteorological data can improve forecasting performance across all sky conditions and forecast horizons.

The inclusion of RH and WD reduced the overall RMS error from 124 W/m^2 (image-only) to 119 W/m^2, and the overall MA error from 90 W/m^2 to 82 W/m^2, at the 60-min horizon. Compared to the persistence model, which yielded 204 W/m^2 RMS error and 163 W/m^2 MA error. The hybrid model showed substantial improvements (42 % RMS error and 50 % MA error reductions) over persistence. Interestingly, the addition of more meteorological variables beyond RH and WD did not yield further improvements and, in some cases, slightly degraded performance. This suggests that including excessive or less informative features may introduce redundancy or noise, adversely affecting the model's ability to generalise well. Additionally, the forecasting errors generally increased as forecast horizons increased to 120-min for all models. However, the persistence model showed the most significant rise in error, with nRMS and nMA errors reaching 60 % and 47 %, respectively, at 120-min horizon. Conversely, the image-only model and the configuration that included RH and WD exhibited much lower error levels. At the 120-min forecast horizon, the nRMS errors for these models were relatively stable around 34 –35 %, while nMA errors stabilised below 27 %.

These outcomes show the importance of strategic feature selection when designing hybrid deep learning models for SI forecasting. Specifically, meteorological features that are physically relevant and complementary to sky imagery should be prioritised, while redundant or weakly related inputs may be excluded to improve training, avoid overfitting and improve interpretability. In future, the generalisability of this approach would be further validated using larger, more diverse

datasets from multiple geographic locations to test model robustness and transferability.

CRediT authorship contribution statement

Albertina M. Amakali: Writing – review & editing, Visualization, Validation, Software, Methodology, Investigation, Formal analysis. **Martin Ansong:** Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Bryce S. Richards:** Writing – review & editing, Visualization, Supervision, Resources, Project administration, Methodology, Funding acquisition, Formal analysis, Conceptualization. **Petja Dobreva:** Writing – review & editing, Supervision, Resources. **Robinson J. Musembi:** Writing – review & editing, Resources. **Thomas N. Nyang'onda:** Writing – review & editing, Resources.

Declaration of Competing Interest

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

Acknowledgements

The authors acknowledge the Ph.D. scholarship for MA provided by the PASET Regional Scholarships and Innovation Fund (RSIF) and the German Academic Exchange Services (DAAD); the financial support provided by the Helmholtz Association via (i) the Recruitment Initiative Funding to BSR and (ii) the Research Field Energy Program-Materials and Technologies for the Energy Transition-Topic 1 Photovoltaics and Wind Energy (38.01.05). MA and BSR would like to acknowledge Dr Gan Huang (KIT) for his efforts in setting up the weather station and data logging facilities.

Appendix A. Supporting information

Supplementary data associated with this article can be found in the online version at doi:10.1016/j.egyr.2025.11.027.

Data availability

Data will be made available on request.

References

Al-lahham, A., Theeb, O., Elalem, K., Alshawi, T.A., Alshebeili, S.A., 2020. 'Sky imager-based forecast of solar irradiance using machine learning'. Electronics 9, 1700.
Ansong, M., Huang, G., Nyang'onda, T.N., Musembi, R.J., Richards, B.S., 2025a. Very short-term solar irradiance forecasting based on open-source low-cost sky imager and hybrid deep-learning techniques. Sol. Energy 294, 113516.
Ansong, M., Nyang'onda, T.N., Musembi, R.J., Richards, B.S., 2024. "Very short-term solar irradiance forecasting for photovoltaic power integration with the grid: potentials and challenges for Africa.". 2024 IEEE PES/IAS Power 01–05.
Ansong, M., Ogunniyi, E.O., Jiménez, B.P., Richards, B.S., 2025b. Renewable energy powered membrane technology: integration of solar irradiance forecasting for predictive control of photovoltaic-powered brackish water desalination system. Appl. Energy 401, 126651.
Barancsuk, L., Groma, V., Günter, D., Osán, J., Hartmann, B., 2024. 'Estimation of solar irradiance using a neural network based on the combination of sky camera images and meteorological data'. Energies 17, 438.
Blanc, P., J. Remund, and L. Vallance. 2017. '6 - Short-term solar power forecasting based on satellite images.' in George Kariniotakis (ed.), Renewable Energy Forecasting (Woodhead Publishing).
Bo, Y., Xia, Y., Ni, Y., Liu, K., Wei, W., 2023. The ultra-short-term photovoltaic power prediction based on multi-exposure high-resolution total sky images using deep learning. Energy Rep. 9, 123–133.
Caldas, M., Alonso-Suárez, R., 2019. 'Very short-term solar irradiance forecast using all-sky imaging and real-time irradiance measurements'. Renew. Energy 143, 1643–1658.
Chaaraoui, S., Houben, S., Meilinger, S., 2024. Probabilistic end-to-end irradiance forecasting through pre-trained deep learning models using all-sky-images. Adv. Sci. Res. 20, 129–158.
Chow, C.W., Belongie, S., Kleissl, J., 2015. Cloud motion and stability estimation for intra-hour solar forecasting. Sol. Energy 115, 645–655.
Chow, C.W., Urquhart, B., Lave, M., Dominguez, A., Kleissl, J., Shields, J., Washom, B., 2011. Intra-hour forecasting with a total sky imager at the UC San Diego solar energy testbed. Sol. Energy 85, 2881–2893.
Coimbra, C.F.M., J. Kleissl, and R. Marquez. 2013. 'Overview of Solar-Forecasting Methods and a Metric for Accuracy Evaluation.' in Jan Kleissl (ed.), Solar Energy Forecasting and Resource Assessment (Academic Press: Boston).
Demirtas, M., Yesilbudak, M., Sagiroglu, S., Colak, I., 2012. "Prediction of solar radiation using meteorological data. 2012 Int. Conf. Renew. Energy Res. Appl. (ICRERA) 1–4.
Dev, S., Savoy, F.M., Lee, Y.H., Winkler, S., 2016. "Short-term prediction of localized cloud motion using ground-based sky imagers. 2016 IEEE Reg. 10 Conf. (TENCON) 2563–2566.
Diagne, M., David, M., Lauret, P., Boland, J., Schmutz, N., 2013. Review of solar irradiance forecasting methods and a proposition for small-scale insular grids. Renew. Sustain. Energy Rev. 27, 65–76.
El Alani, O., Abraim, M., Ghennioui, H., Ghennioui, A., Ikenbi, I., Dahr, F.-E., 2021. Short term solar irradiance forecasting using sky images based on a hybrid CNN–MLP model. Energy Rep. 7, 888–900.
Feng, C., Zhang, J., 2020. SolarNet: a sky image-based deep convolutional neural network for intra-hour solar forecasting. Sol. Energy 204, 71–78.
Gala, Y., Fernández, Á., Díaz, J., Dorronsoro, J.R., 2016. 'Hybrid machine learning forecasting of solar radiation values'. Neurocomputing 176, 48–59.
Gao, B., Huang, X., Shi, J., Tai, Y., Xiao, R., 2019. Predicting day-ahead solar irradiance through gated recurrent unit using weather forecasting data. J. Renew. Sustain. Energy 11.
Gevorgian, V., and S. Booth. 2013. "Review of PREPA technical requirements for interconnecting wind and solar generation." In.: National Renewable Energy Laboratory.
Hendrikx, N.Y., Barhmi, K., Visser, L.R., de Bruin, T.A., Pó, M., Salah, A.A., van Sark, W. G.J.H.M., 2024. 'All sky imaging-based short-term solar irradiance forecasting with Long Short-Term Memory networks'. Sol. Energy 272, 112463.
Hoff, T.E., Perez, R., Kleissl, J., Renne, D., Stein, J., 2013. Reporting of irradiance modeling relative prediction errors. Prog. Photovolt. Res. Appl. 21, 1514–1519.
Hou, X., Wild, M., Folini, D., Kazadzis, S., Wohland, J., 2021. Climate change impacts on solar power generation and its spatial variability in Europe based on CMIP6. Earth Syst. Dynam 12, 1099–1113.
Huang, H., Xu, J., Peng, Z., Yoo, S., Yu, D., Huang, D., Qin, H., 2013. "Cloud motion estimation for short term solar irradiation prediction. 2013 IEEE Int. Conf. Smart Grid Commun. (SmartGridComm) 696–701.
International Energy Agency, 2024. Renewables 2024. International Energy Agency, Paris. https://www.iea.org/reports/renewables.
Jakoplić, A., Franković, D., Havelka, J., Bulat, H., 2023. Short-Term Photovoltaic Power Plant Output Forecasting Using Sky Images and Deep Learning. Energies 16, 5428.
Jamal, T., Shafiullah, G.M., Carter, C., Ferdous, S.M., Rahman, M., 2018. "Benefits of short-term PV forecasting in a remote area standalone off-grid power supply system. 2018 IEEE Power Energy Soc. Gen. Meet. (PESGM) 1–5.
Krechowicz, M., Krechowicz, A., Lichołai, L., Pawelec, A., Piotrowski, J.Z., Stępień, A., 2022. 'Reduction of the risk of inaccurate prediction of electricity generation from PV farms using machine learning'. Energies 15, 4006.
Kremer, M., Kiesel, R., Paraschiv, F., 2021. An econometric model for intraday electricity trading. Philos. Trans. R. Soc. A Math. Phys. Eng. Sci. 379, 20190624.
Kuang, H., Guo, Q., Li, S., Zhong, H., 2021. Short-term wind power forecasting model based on multi-feature extraction and CNN-LSTM. IOP Conference Series Earth Environmental Science 702, 012019.
Kumar, D.S., Yagli, G.M., Kashyap, M., Srinivasan, D., 2020. Solar irradiance resource and forecasting: a comprehensive review. IET Renew. Power Gener. 14, 1641–1656.
López-Cuesta, M., Aler-Mur, R., Galván-León, I.M., Rodríguez-Benítez, F.J., Pozo-Vázquez, A.D., 2023. Improving solar radiation nowcasts by blending data-driven, satellite-images-based and all-sky-imagers-based models using machine learning techniques. Remote Sens. 15, 2328.
Lorenz, E., Heinemann, D., 2012. Prediction of Solar Irradiance and Photovoltaic Power. In: Sayigh, A. (Ed.). Comprehensive Renewable Energy (Elsevier: Oxford).
Märkle-Huß, J., Feuerriegel, S., Neumann, D., 2018. 'Contract durations in the electricity market: causal impact of 15min trading on the EPEX SPOT market'. Energy Econ. 69, 367–378.
Marquez, R., Coimbra, C.F.M., 2013. Intra-hour DNI forecasting based on cloud tracking image analysis. Sol. Energy 91, 327–336.
Martín, L., Zarzalejo, L.F., Polo, J., Navarro, A., Marchante, R., Cony, M., 2010. Prediction of global solar irradiance based on time series analysis: Application to solar thermal power plants energy production planning. Sol. Energy 84, 1772–1781.
Mayer, O., Becker, G., Giesler, B., Hammer, A., Gottschalg, R., Betts, T.R., Krawczyński, M., Wirth, G., Mayer, B., Weigl, T., Gratzl, T., Weizenbeck, J., Hartmann, M., Zehner, M., 2010. "Systematic analysis of meteorological irradiation effects.". 25th EU PVSEC / WCPEC 5.
Miller, S.D., Rogers, M.A., Haynes, J.M., Sengupta, M., Heidinger, A.K., 2018. Short-term solar irradiance forecasting via satellite/model coupling. Sol. Energy 168, 102–117.
Mwigereri, D.G., Ochieng, F.X., Ouma, J., Mwai, Z., Mwalili, T., 2025. Evaluation of weather research forecasting-solar radiation schemes for solar power forecasting in equatorial Africa. Energy Rep. 13, 4318–4330.
Nie, Y., Li, X., Paletta, Q., Aragon, M., Scott, A., Brandt, A., 2024. Open-source sky image datasets for solar forecasting with deep learning: a comprehensive survey. Renew. Sustain. Energy Rev. 189, 113977.
Nie, Y., Li, X., Scott, A., Sun, Y., Venugopal, V., Brandt, A., 2023. SKIPP'D: a sky images and photovoltaic power generation dataset for short-term solar forecasting. Sol. Energy 255, 171–179.

Nie, Y., Sun, Y., Chen, Y., Orsini, R., Brandt, A., 2020. PV power output prediction from sky images using convolutional neural network: the comparison of sky-condition-specific sub-models and an end-to-end model. J. Renew. Sustain. Energy 12.
Nikodinoska, D., Käso, M., Müsgens, F., 2022. 'Solar and wind power generation forecasts using elastic net in time-varying forecast combinations'. Appl. Energy 306, 117983.
Paletta, Q., Arbod, G., Lasenby, J., 2021. Benchmarking of deep learning irradiance forecasting models from sky images – an in-depth analysis. Sol. Energy 224, 855–867.
Paletta, Q., Arbod, G., Lasenby, J., 2023a. Omnivision forecasting: Combining satellite and sky images for improved deterministic and probabilistic intra-hour solar energy predictions. Appl. Energy 336, 120818.
Paletta, Q., Lasenby, J., 2020. "Convolutional Neural Networks applied to sky images for short-term solar irradiance forecasting. Eur. PV Sol. Energy Conf. Exhib. (EUPVSEC) 2020.
Paletta, Q., Terrén-Serrano, G., Nie, Y., Li, B., Bieker, J., Zhang, W., Dubus, L., Dev, S., Feng, C., 2023b. 'Advances in solar forecasting: Computer vision with deep learning'. Adv. Appl. Energy 11, 100150.
Pandžić, F., Capuder, T., 2024. 'Advances in Short-Term Solar Forecasting: A Review and Benchmark of Machine Learning Methods and Relevant Data Sources'. Energies 17, 97.
Paulescu, M., Paulescu, E., 2019. Short-term forecasting of solar irradiance. Renew. Energy 143, 985–994.
Pedro, H.T.C., Coimbra, C.F.M., 2012. Assessment of forecasting techniques for solar power production with no exogenous inputs. Sol. Energy 86, 2017–2028.
Quesada-Ruiz, S., Chu, Y., Tovar-Pescador, J., Pedro, H.T.C., Coimbra, C.F.M., 2014. 'Cloud-tracking methodology for intra-hour DNI forecasting'. Sol. Energy 102, 267–275.
Rajagukguk, R.A., Kamil, R., Lee, H.-J., 2021. A Deep Learning Model to Forecast Solar Irradiance Using a Sky Camera. Appl. Sci. 11, 5049.
Samu, R., Calais, M., Shafiullah, G.M., Moghbel, M., Shoeb, M.A., Nouri, B., Blum, N., 2021. Applications for solar irradiance nowcasting in the control of microgrids: A review. Renew. Sustain. Energy Rev. 147, 111187.
Shirazi, E., Gordon, I., Reinders, A., Catthoor, F., 2024. Sky images for short-term solar irradiance forecast: a comparative study of linear machine learning models. IEEE J. Photovolt. 14, 691–698.
Siddiqui, T.A., Bharadwaj, S., Kalyanaraman, S., 2019. "A deep learning approach to solar-irradiance forecasting in sky-videos. 2019 IEEE Winter Conf. Appl. Comput. Vis. (WACV) 2166–2174.
Simonyan, K., and A. Zisserman. 2014. 'Very deep convolutional networks for large-scale image recognition', arXiv preprint arXiv:1409.1556.
Song, S., Yang, Z., Goh, H., Huang, Q., Li, G., 2022. A novel sky image-based solar irradiance nowcasting model with convolutional block attention mechanism. Energy Rep. 8, 125–132.
Straub, N., Herzberg, W., Dittmann, A., Lorenz, E., 2024. Blending of a novel all sky imager model with persistence and a satellite based model for high-resolution irradiance nowcasting. Sol. Energy 269, 112319.
Sun, Y., Szűcs, G., Brandt, A.R., 2018. 'Solar PV output prediction from video streams using convolutional neural networks'. Energy Environ. Sci. 11, 1811–1818.
Voyant, C., Muselli, M., Paoli, C., Nivet, M.-L., 2011. Optimization of an artificial neural network dedicated to the multivariate forecasting of daily global radiation. Energy 36, 348–359.
Wang, P., van Westrhenen, R., Meirink, J.F., van der Veen, S., Knap, W., 2019. Surface solar radiation forecasts by advecting cloud physical properties derived from Meteosat Second Generation observations. Sol. Energy 177, 47–58.
West, S.R., Rowe, D., Sayeef, S., Berry, A., 2014. Short-term irradiance forecasting using skycams: Motivation and development. Sol. Energy 110, 188–207.
Wild, M., Folini, D., Henschel, F., Fischer, N., Müller, B., 2015. Projections of long-term changes in solar radiation based on CMIP5 climate models and their influence on energy yields of photovoltaic systems. Sol. Energy 116, 12–24.
Yan, K., Shen, H., Wang, L., Zhou, H., Xu, M., Mo, Y., 2020. Short-term solar irradiance forecasting based on a hybrid deep learning methodology. Information 11(1), 2078-2489.
Yang, L., Gao, X., Hua, J., Wu, P., Li, Z., Jia, D., 2020. Very Short-Term Surface Solar Irradiance Forecasting Based On FengYun-4 Geostationary Satellite. Sensors 20, 2606.
Yang, L., Gao, X., Li, Z., Jia, D., Jiang, J., 2019. Nowcasting of surface solar irradiance using FengYun-4 satellite observations over China. Remote Sens. 11 (17), 1984.
Yang, H., Kurtz, B., Nguyen, D., Urquhart, B., Chow, C.W., Ghonima, M., Kleissl, J., 2014. Solar irradiance forecasting using a ground-based sky imager developed at UC San Diego. Sol. Energy 103, 502–524.
Yang, B., Zhu, T., Cao, P., Guo, Z., Zeng, C., Li, D., Chen, Y., Ye, H., Shao, R., Shu, H., Yu, T., 2021. Classification and summarization of solar irradiance and power forecasting methods: A thorough review. CSEE J. Power Energy Syst. 1–19.
Zhang, L., Wilson, R., Sumner, M., Wu, Y., 2023. Advanced multimodal fusion method for very short-term solar irradiance forecasting using sky images and meteorological data: A gate and transformer mechanism approach. Renew. Energy 216, 118952.
Zuo, H.-M., Qiu, J., Li, F.-F., 2023. Ultra-short-term forecasting of global horizontal irradiance (GHI) integrating all-sky images and historical sequences. J. Renew. Sustain. Energy 15.
Zuo, H.-M., Qiu, J., Jia, Y.-H., Wang, Q., Li, F.-F., 2022. Ten-minute prediction of solar irradiance based on cloud detection and a long short-term memory (LSTM) model. Energy Rep. 8, 5146–5157.