Technical University of Applied Sciences Rosenheim
University of Applied Sciences Bielefeld

HELIOS: Predictive Spatial Analytics for Solar Energy Grid Integration: Enhancing Reliability and Efficiency

Final Report on a Development Project
funded under the reference number: 39213/01 by the
German Federal Environmental Foundation

Andreas Boschert, Prof. Dr.-Ing. Grit Behrens, Prof. Mike Zehner,
Yassine Ribouh, and Naoufal EL Atmioui

March 2026

Declaration of generative AI and AI-assisted technologies in the writing process.

During the preparation of this work the authors used ChatGPT by OpenAI, DeepL Translator by DeepL SE, and Grammarly by Grammarly Inc. to improve language and readability. All content generated or modified with these tools was subsequently reviewed and edited by the authors, who takes full responsibility for the final content of this manuscript.

Acknowledgements

We would like to express our deepest gratitude to the German Federal Environmental Foundation (Deutsche Bundesstiftung Umwelt, DBU) for its generous financial support and the trust placed in our research project. The funding provided by the DBU was instrumental in enabling the successful completion of this work. It facilitated the development of innovative technical approaches and the generation of valuable scientific insights that would not have been achievable to the same extend without this support.

Our sincere thanks also go to the project partners and collaboration institutions for their unwavering commitment, expertise, and consistently constructive collaboration. The interdisciplinary cooperation and regular, open exchange of ideas proved to be a cornerstone of the project's success, ensuring that the results are not only scientifically robust but also practically applicable. This is of great importance in applied science.

We are equally grateful to all contributors, students, and supporters who played technical, organizational, or scientific roles in the project. Their dedication, creative input, and active engagement enriched the work in countless ways, fostering an environment of collaboration, mutual respect, and constructive dialogue.

To everyone involved – whether through day-to-day collaboration, insightful discussions, or behind-the-scenes support – we extend our heartfelt appreciation. Each contribution, no matter how small, was essential in shaping the project and bringing it to fruition. It is through your collective efforts that we were able to create a productive and inspiring work environment.

On a personal note, we would like to express our sincere gratitude to Dr. Katrin Anneser, our dedicated contact person at the German Federal Environmental Foundation. As the DBU's primary liaison for this project, she provided consistent support, clear guidance, and reliable coordination throughout its duration. Her constructive feedback, prompt responses, and professional expertise ensured smooth communication between the project team and the funding institution. We particularly appreciate her commitment to facilitating an efficient workflow and her role in helping align the project's objectives with the DBU's expectations. Her contributions were instrumental in maintaining a productive partnership, and we are grateful for her professionalism and support.

"*Alone we can do so little; together we can do so much.*" - Helen Keller

Table of Contents

List of Figures

List of Tables

List of Abbreviations

All-Sky Imagers .... *ASI*

Application Programming Interface .... *API*

Clear-Sky Index .... *CSI*

Cloud Motion Vectors .... *CMV*

Diffuse Horizontal Irradiance .... *DHI*

Direct Normal Irradiance .... *DNI*

Endpoint Error .... *EPE*

European Photovoltaic Solar Energy Conference and Exhibition .... *EU PVSEC*

Field Of View .... *FOV*

File Transfer Protocol .... *FTP*

Gated Recurrent Units .... *GRU*

Generative Adversarial Network .... *GAN*

Global Horizontal Irradiance .... *GHI*

HELIOS Data Monitoring System .... *HDMS*

Light Detection and Ranging .... *LIDAR*

Long Short-Term Memory .... *LSTM*

Lucas-Kanade .... *LK*

Masked Autoencoder .... *MAE*

Not-a-Number .... *NaN*

Photovoltaic .... *PV*

Recurrent All-Pairs Field Transforms .... *RAFT*

Research Data Management Plan .... *RDMP*

Self-Supervised Learning .... *SSL*

Spinning Enhanced Visible and Infra-Red Imager .... *SEVIRI*

Technical University of Applied Sciences .... *TUAS*

University of Applied Sciences .... *UAS*

Summary

The »HELIOS« project, funded by the Deutsche Bundesstiftung Umwelt (DBU) and conducted jointly by the Technical University of Applied Sciences Rosenheim and the University of Applied Sciences Bielefeld, developed and validated advanced AI-driven methods for high-resolution short-term solar irradiance forecasting. The overarching goal was to significantly enhance the reliability and efficiency of photovoltaic energy integration into increasingly decentralized power grids, addressing a critical gap in forecast accuracy at minute-scale resolutions.

Classical numerical weather prediction models and satellite-based methods are fundamentally limited for forecast horizons of minutes to a few hours, lacking the temporal and spatial resolution required to capture small-scale cloud dynamics. While satellite instruments such as SEVIRI on Meteosat typically offer spatial resolutions of approximately two kilometres and temporal resolutions of fifteen minutes, All-Sky Imagers enable forecasts at spatial resolutions of around 50 meters and temporal resolutions of up to one minute. The »HELIOS« project systematically exploited this advantage by developing a state-of-the-art measurement infrastructure at a ground-mounted PV installation at the Maierhof site in Buttenwiesen, Bavaria.

The measurement setup comprised three All-Sky Imager, a LIDAR ceilometer, a rotating shadow band pyranometer, high-precision meteorological sensors, and a connection to the site's PV monitoring system via a standardized Application Programming Interface (API) interface. Careful geometric calibration of the Imagers using the OCamCalib toolbox ensured accurate projection of sky images onto a defined coordinate system, forming the basis for reliable cloud analysis. A dedicated Monitoring System was developed to ensure continuous data quality assurance, automated anomaly detection, and structured long-term archiving in line with the project's data management.

Building on this robust data foundation, multiple machine learning and deep learning algorithms were designed, implemented, and rigorously benchmarked under real-world conditions. Key algorithmic developments included a hybrid MAE-DINO self-supervised vision model for accurate cloud detection, a physics-informed Generative Adversarial Network (GAN) -Transformer framework (SkyFusion) for intra-hour sky image forecasting and global horizontal irradiance prediction, and RAFT-based optical flow methods for cloud motion vector estimation. Comparative evaluations demonstrated that RAFT-based approaches achieved subpixel endpoint errors, vastly outperforming classical methods such as Lucas-Kanade and Farnebäck optical flow. The hybrid MAE-DINO cloud detection model similarly demonstrated superior performance across multiple forecast horizons compared to conventional baselines.

The project results were actively disseminated to the scientific community and industry through a range of channels. Contributions were made at international conferences including European Photovoltaic Solar Energy Conference and Exhibition (EU PVSEC) in Vienna, EnviroInfo in Cairo, and a workshop on minute-scale forecasting at DTU Risø Campus in Denmark. A particular highlight was the organization of the inaugural All-Sky Imager Workshop 2025, which brought together researchers and industry representatives to foster exchange and collaboration. A follow-up workshop is planned for October 2026 at the DLR site in Oldenburg. In addition, project findings were integrated into university teaching and advanced through several master's theses and student projects.

Zusammenfassung

Das Projekt »HELIOS«, gefördert von der Deutschen Bundesstiftung Umwelt (DBU) und gemeinsam von der Technischen Hochschule Rosenheim und der Hochschule Bielefeld durchgeführt, hatte zum Ziel, fortschrittliche KI-basierte Methoden für die hochauflösende Kurzfristvorhersage der solaren Einstrahlung zu entwickeln und zu validieren. Das übergeordnete Ziel bestand darin, die Zuverlässigkeit und Effizienz der Einspeisung von Photovoltaikanlagen in zunehmend dezentral organisierte Energieversorgungsnetze deutlich zu verbessern und eine bestehende Lücke in der Vorhersagegenauigkeit im Minutenbereich zu schließen.

Klassische numerische Wettervorhersagemodelle und satellitengestützte Methoden stoßen für Vorhersagehorizonte von wenigen Minuten bis zu einigen Stunden an grundlegende Grenzen, da sie nicht über die räumliche und zeitliche Auflösung verfügen, die zur Erfassung kleinräumiger Wolkendynamiken erforderlich ist. Während Satelliteninstrumente wie SEVIRI auf Meteosat typischerweise räumliche Auflösungen von etwa zwei Kilometern und zeitliche Auflösungen von 15 Minuten bieten, ermöglichen All-Sky Imager Vorhersagen mit räumlichen Auflösungen von rund 50 Metern und zeitlichen Auflösungen von bis zu einer Minute. Das Projekt »HELIOS« nutzte diesen Vorteil systematisch aus, indem am Standort einer Freiflächen-PV-Anlage am Maierhof in Buttenwiesen eine moderne Messinfrastruktur aufgebaut wurde.

Die Messinfrastruktur umfasste drei All-Sky Imager, ein LIDAR-Ceilometer, ein rotierendes Schattenbandpyranometer, hochpräzise meteorologische Sensoren sowie eine Anbindung an das PV-Überwachungssystem der Anlage über eine standardisierte API-Schnittstelle. Eine sorgfältige geometrische Kalibrierung der Imager mithilfe der OCamCalib-Toolbox stellte die genaue Projektion der Himmelsbilder auf ein definiertes Koordinatensystem sicher und bildet damit die Grundlage für eine zuverlässige Wolkenanalyse. Zur kontinuierlichen Qualitätssicherung der erfassten Daten wurde ein eigens entwickeltes Data Monitoring System eingesetzt, das eine automatisierte Anomalieerkennung, strukturierte Langzeitarchivierung und eine konsistente Umsetzung des Forschungsdatenmanagementplans gewährleistete.

Auf Basis dieser Datengrundlage wurden mehre Algorithmen des maschinellen Lernens entwickelt, implementiert und unter realen Bedingungen umfassend evaluiert. Zu den zentralen algorithmischen Entwicklungen zählen ein hybrides, selbstüberwachtes MAE-DINO-Modell zur präzisen Wolkenerkennung, ein physikalisch informiertes GAN-Transformer-Framework (SkyFusion) zur Vorhersage von Himmelsbildsequenzen und der globalen Horizontstrahlung sowie RAFT-basierte optische Flussverfahren zur Schätzung von Wolkenbewegungsvektoren. Vergleichende Evaluierungen zeigten, dass RAFT-basierte Ansätze Sub-Pixel-Genauigkeiten erreichten und klassische Methoden wie Lucas-Kanade und Farnebäck deutlich übertrafen. Auch das hybride MAE-DINO-Modell zur Wolkenerkennung zeigte über mehrere Vorhersagehorizonte hinweg eine überlegene Leistung gegenüber konventionellen Referenzverfahren.

Die Projektergebnisse wurden aktiv in die Wissenschafts- und Fachöffentlichkeit getragen. Beiträge wurden auf internationalen Konferenzen wie der EU PVSEC in Wien, der EnviroInfo in Kairo sowie einem Workshop zur Nowcasting am DTU Risø Campus in Dänemark präsentiert. Ein besonderer Höhepunkt war die Organisation des ersten All-Sky Imager Workshops 2025 an der TH Rosenheim, der Forschende und Industrievertreter zusammenbrachte und den Austausch sowie die Vernetzung innerhalb der Fachgemeinschaft förderte. Eine Fortsetzung ist bereits für Oktober 2026 am DLR-Standort in Oldenburg geplant. Darüber hinaus wurden die Projektergebnisse in der Hochschullehre integriert und im Rahmen mehrerer Masterthesen und Studierendenprojekte weiterentwickelt, wodurch ein nachhaltiger Wissenstransfer sichergestellt wurde.

1. Introduction

The reliable integration of photovoltaic (PV) generation into an increasingly decentralized energy system requires accurate short term forecasts of solar radiation induced power feed in. Classical numerical weather prediction models or satellited based methods are particularly limited for forecast horizons of minutes to a few hours, because they lack both the temporal and the spatial resolution needed to capture small scale cloud dynamics. All-Sky Imagers (ASI) are useful complement to satellite data and offer a promising approach. Their hemispheric sky images make it possible to observe the current cloud situation directly at the site. Compared with satellite based observations methods such as the high resolution visible channel of the Spinning Enhanced Visible and Infra-Red Imager (SEVIRI) instrument on the Meteosat satellites, ASI provide much higher spatial and temporal resolution. While SEVIRI typically offers a spatial resolution of about two kilometers and a temporal resolution of fifteen minutes, ASI enable forecasts with a spatial resolution of around fifty meters and a temporal resolution of up to one minute [SSSB24]. Through this higher level of detail and the analysis of cloud motion in sequences of images, combined with a series of algorithms, fluctuations in global irradiance and the resulting short term PV power production can be predicted on a minute scale.

A key application of short term forecasts is their ability to provide a quantifiable basis for operational decision making. Forecasts of PV generation help to reduce uncertainty and thereby improve grid operation, plant control and market participation. In particular, for the operational management of energy systems such as microgrids, district energy solutions or industrial self-consumption schemes, a reliable irradiance forecast enables optimized scheduling of storage units, controllable loads and conventional generators. Furthermore, improved forecast quality supports more economical participation in the intraday electricity market, because short term fluctuations in PV generation can be anticipated more accurately and marketing strategies can be optimized accordingly. This will make it possible, in the specific context of tenant PV models, to offer higher compensation to customers for the electricity generated locally, which in turn increases the profitability and attractiveness of installing PV systems.

Beyond the technical end economic optimization of plant operation, short term forecasts make an important contribution to the ecological sustainability of the energy system. They support and accelerate the transition toward a renewable energy supply by enabling more efficient use of volatile generation through reliable prediction. At the same time, this leads to decline in conventional energy generation and reduces environmental impacts such as greenhouse gas emissions and resource depletion. As the share of renewable energy in the overall energy mix continues to rise, the ecological footprint of electricity generation decreases significantly. This promotes the development of a more environmentally friendly and more sustainable energy infrastructure. PV generation has a particularly favourable emissions profile. With a net avoidance factor of 683.82 grams of CO2 per kilowatt hour produced [LaMS22], PV contributes significantly to achieving the federal governments climate protection targets, especially with regard to the goal of carbon neutrality. To fully harness this potential, however, precise and reliable forecasts of PV generation are necessary. Only then can security of supply be ensured even at high shares of variable generation. In PV ground mounted system combined with large scale battery storage and predictive energy management, yields can be increased by up to twenty percent [WBOQ00].

Within the HELIOS [BoZe25a] [BoZe25b] project, advanced AI-driven approaches for high-resolution short-term solar irradiance forecasting were developed and rigorously validated to significantly enhance prediction accuracy. To achieve this, a state-of-the-art measurement infrastructure was deployed at ag round-mounted PV system in Buttenwiesen, comprising three ASI and high-precision meteorological sensors. Leveraging this data, multiple AI-based algorithms were designed, implemented, and benchmarked for real-time cloud analysis. The algorithms were tested under real-world conditions, demonstrating substantial improvements in accuracy to conventional methods.

2. System Setup and Data Acquisition

2.1 Installation

2.1.1. PV System Inspection and Technical Planning

A comprehensive on-site inspection of the PV installation at the Maierhof site in Buttenwiesen was conducted on 24 March 2024 to identify optimal locations for measurement and communication equipment. The system, located at GP Joule GmbH's premises and operated by PV Kraftwerk GmbH, is supported by project partner Timeless Planet GmbH.

During the inspection, infrastructural conditions and technical requirements – including power supply, network access, and unobstructed sensor/ASI visibility – were systematically assessed. In collaboration with on-site partners, potential installation positions were evaluated, necessary structural adjustments identified, and initial responsibilities assigned. To support documentation and planning, aerial imagery was captured using a drone (see Figure 1).

Figure 1: Aerial view of the Maierhof PV installation in Buttenwiesen (near Augsburg, Germany), captured via drone during the on-site inspection (24 March 2024) to assess optimal placement of measurement and monitoring equipment.

The assessment provided a critical foundation for detailed planning, ensuring seamless integration of the measurement equipment into the existing PV system. Beyond technical insights, the visit aligned partner expectations, identified logistical challenges early, and enabled the development of a coordinated installation strategy that minimizes operational disruptions while maximizing data quality.

The inspection also uncovered key challenges that required careful consideration during the planning phase. Network infrastructure limitations emerged as a primary concern, with several critical locations lacking stable connections; further complicated by an ongoing transition to fibre-optic cabling. Additionally, ongoing refurbishment work on parts of the PV installation, combined with incomplete or outdated documentation of older system components, introduced uncertainties regarding existing interfaces and operational constraints. Finally, the complex network topology of the site posed additional hurdles for seamless data transmission and secure system integration, making close and continuous coordination with the technical teams essential from the outset. To overcome the network bottlenecks, temporary LTE routers were deployed whilst the fibre-optic upgrade was being completed.

2.1.1 Hardware-Tests

2.1.1.1 LIDAR Ceilometer SkyVue8 (Campbell Scientific)

At the beginning of the commissioning phase on the university campus, the basic device parameters were configurated directly on the terminal module. Subsequently, the communication interface was tested to verify reliable data transmission between the instrument and the data acquisition system.

For the integration of the ceilometer into the DOMINO3 framework, a dedicated device driver was implemented. This driver receives the incoming sensor data, processes it, and stores it in the corresponding database system.

The ceilometer provides two fundamentally different types of data:

- **Scalar value pairs (e.g. cloud base height):** These parameters are stored in the InfluxDB database, which is optimized for time-series data. Its structure allows efficient storage, indexing, and querying of continuously recorded measurements.
- **Backscatter profiles (one dimensional vectors with 5 m resolution up to 8000 m):** Each profile contains 1600 data points per time stamp. Although it would technically be possible to store each data point as an individual time series in InfluxDB, this approach would be inefficient in terms of performance and storage requirements. Instead, the profiles are stored as binary data streams in MariaDB database. Each entry is supplemented with minimal metadata such as timestamp, measurement location, and sensor status.

During the initial test operation on the university campus, conducted prior to relocating the system to the Buttenwiesen measurement site, a recurring malfunction was observed. Shortly after each restart, the ceilometer repeatedly entered a warning state. After several consecutive warnings, the system automatically switched to an alarm state and stopped performing regular measurements. Further analysis identified a defect in the laser module as the underlying cause. After consultation with the manufacturer, the faulty component was replaced, which restored the ceilometer to full operational functionality.

The comprehensive pre-deployment testing ensured reliable system operation prior to installation at the final measurement site. This procedure significantly reduced the risk of data loss and minimized potential downtime during the subsequent measurement campaign.

2.1.2.1.1. Backscatter Dashboard

SkyVue8 provides continuous, vertically resolved measurements of atmospheric backscatter which are highly valuable for cloud monitoring, and atmospheric research. In practice, however, the raw export files generated by the instrument are difficult to process. Profile data are typically stored as compressed hexadecimal strings, metadata formats vary between deployments, and manual decoding is both time-consuming and prone to errors.

To address these challenges, the SkyVue Light Detection and Ranging (LIDAR) Backscatter Dashboard was developed. This web-based system converts raw export files into physically meaningful backscatter profiles, diagnostic visualizations, and printable reports.

The dashboard is designed as a zero-configuration assistant: User simply uploads a data file, after which the system automatically performs several processing steps: it detects the file delimiter, reconstructs or validates header information, identifies the profile encoding format (ASCII or direct hexadecimal), and applies the appropriate conversion to physical units. The processed data are then made available through interactive exploration tools, visualization modules, and automated reporting functions.

The following figures show the central functionalities of the dashboard.

- **Step 1: Data Upload and Ingestion**

  The first step involves uploading of the raw measurement file to the dashboard (Figure 2). The system automatically detects the file delimiter, reconstructs or validates header information, and identifies the profile encoding format (ASCII or direct hexadecimal). Upon the upload, the dashboard performs an initial assessment of the dataset, including:

  - File normalization (e.g., correcting column structures to a standard 12-column format),
  - Memory estimation (approximate file size in MP),
  - Header detection (identifying wheter metadata is present),
  - Timeline extraction (automated detection of the date column and available measurement days).

  This steps ensure that the data is correctly parsed before further processing, reducing the risk of errors due to inconsistent file structures.

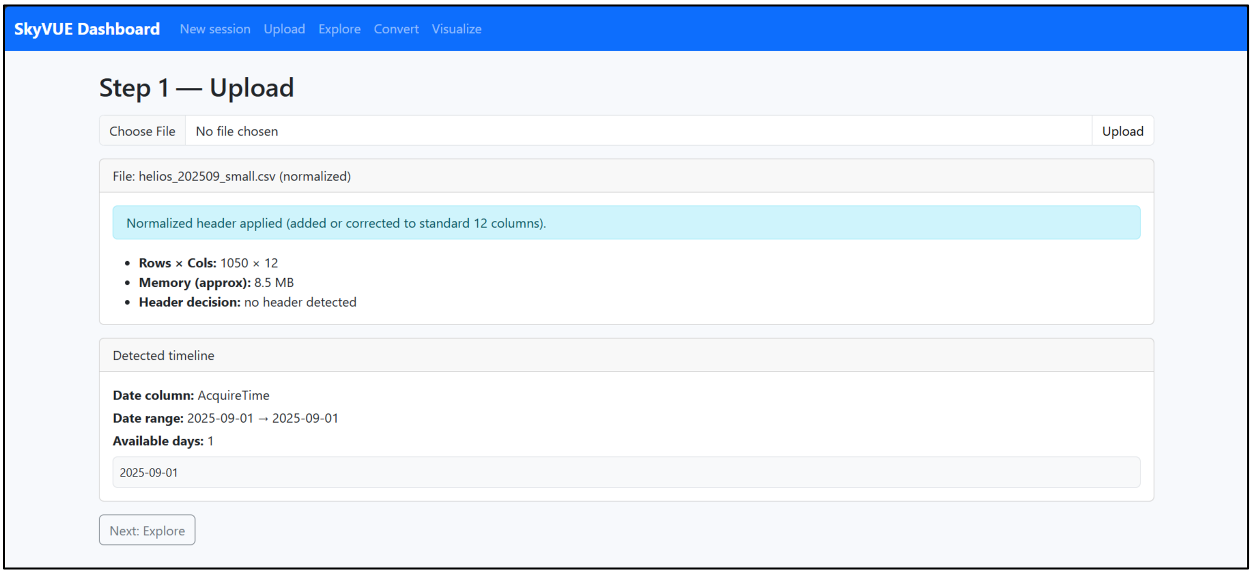

Figure 2: Upload and ingestion of the measurement data. The dashboard automatically detects file properties, such as row count, memory usage, and timeline information.

- **Step 2: Data Exploration and Filtering**

  After successful integration, the dashboard provides an interactive overview of the loaded dataset (Figure 3). Users can:

  - Filter data by date (e.g., selecting specific measurement days),
  - Adjust the number of rows displayed (for performance optimization),
  - Inspect raw data columns (e.g., acquisition time, latitude, longitude, backscatter profiles in hexadecimal format).

  This step is critical for quality control, allowing users to verify the integrity of the dataset before conversion. The dashboard displays key metadata, such as the total number or profiles and altitude bins, ensuring transparency in the data structure.

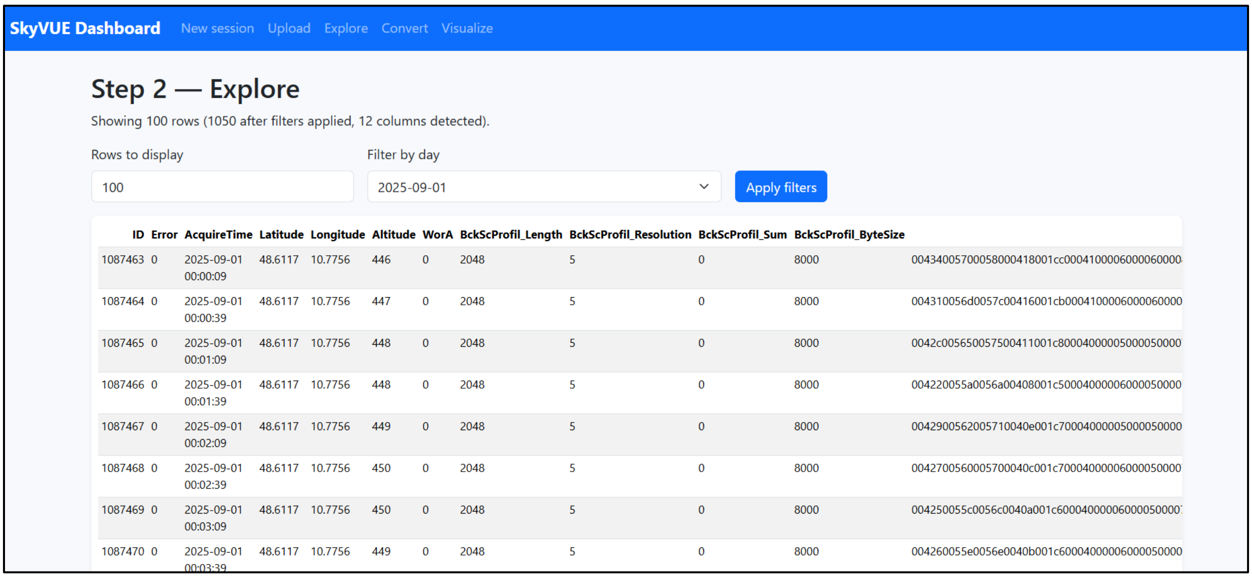

Figure 3: Overview of the loaded data, including filtering options. Users can explore raw data columns and apply temporal filters to refine the dataset.

- **Step 3: Data Conversion**

  In the third step, the dashboard converts the raw backscatter profiles into physically meaningful units (Figure 4). Key features include:

  - Automated format detection (e.g. identifying whether the data is stored in DIRECT or ASCII encoding),
  - Backscatter column selection (with manual override if auto-detection fails),
  - Scaling adjustments (e.g., applying unit conversions or normalization),
  - Output generation (exporting converted profiles, statistics, and altitude references as CSV or TXT files).

  The dashboard also provides a dataset overview, summarizing the total number of profiles, altitude bins, and temporal coverage. A sample profile (e.g., row 0) is displayed for validation, ensuring that the conversion was successful.

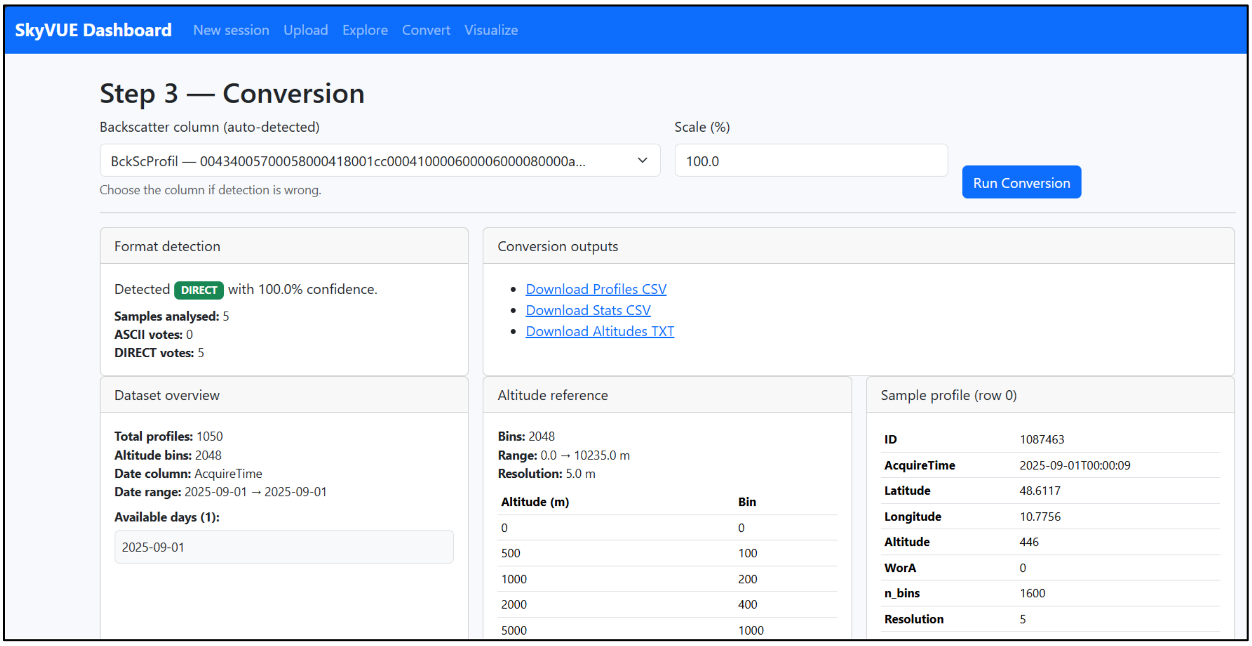

Figure 4: Conversion of raw data into physical units. The dashboard auto-detects the encoding format, applies scaling, and generates downloadable outputs (e.g., CSV files).

- **Step 4: Backscatter Profile Visualization**

  The final step presents the processed data in multiple visualization formats (Figure 6), including:

  1. Vertical Profile Plot – A single backscatter profile as a function of altitude, highlighting atmospheric features (e.g. cloud layers, aerosol concentrations).
  2. Multi-Temporal Comparison – Overlaid profiles from different time steps, enabling temporal analysis of atmospheric changes.
  3. Time-Height Cross Section – A heatmap showing backscatter intensity over time and altitude, ideal for identifying long-term trends or transient phenomena (e.g. cloud formation).

  The visualization is designed to support both exploratory analysis and publication-ready outputs, with customizable axes, colour scales, and annotations.

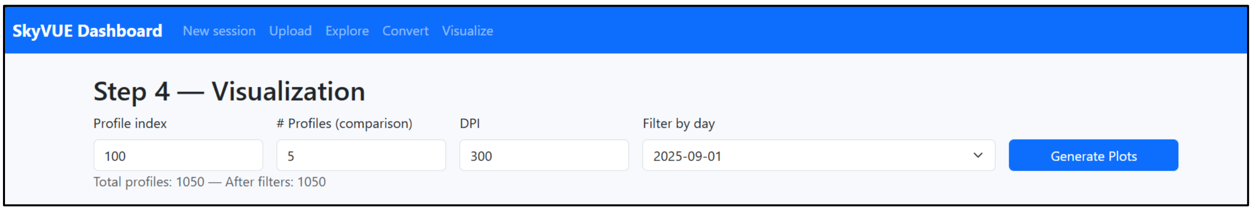

Figure 5: In the fourth step, the visualization settings are configurated.

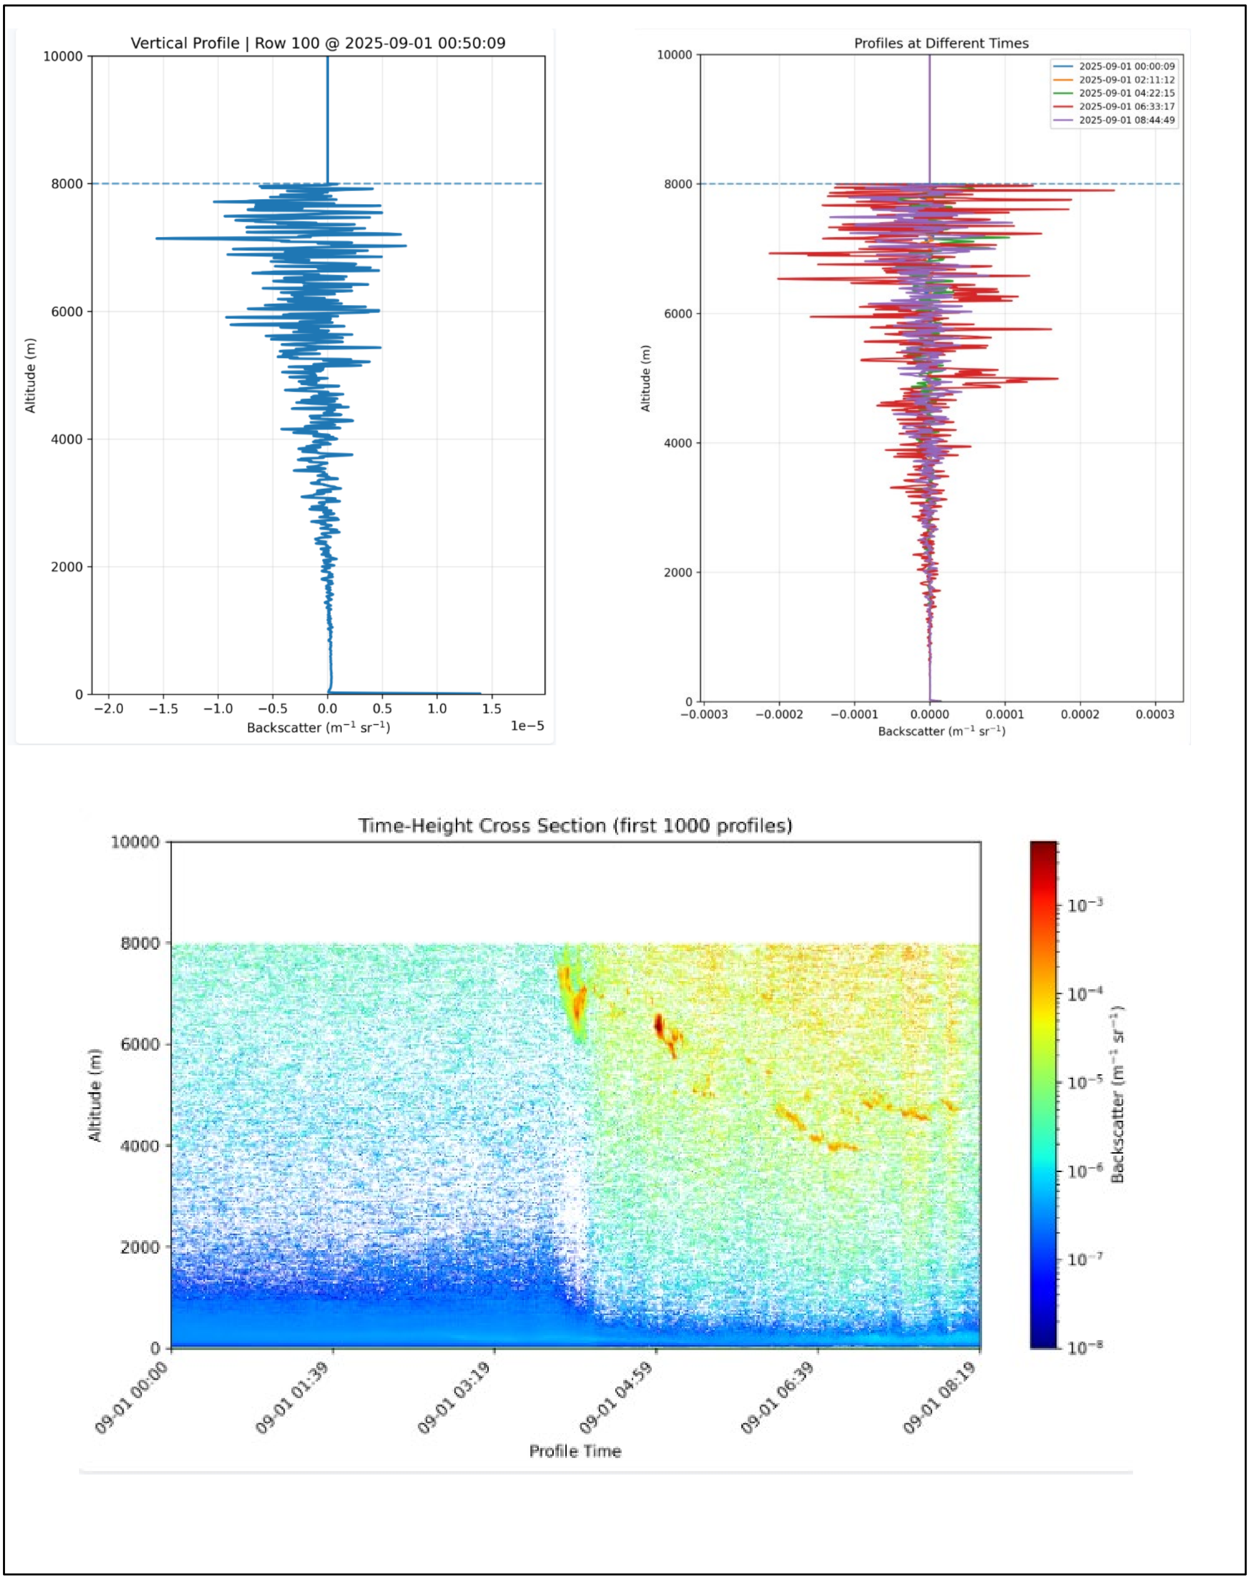

Figure 6: Display of backscatter profiles in different graphs. (Top left) Vertical profile of a single measurement. (Top right) Overlaid profiles at different times. (Bottom) Time-height cross section showing backscatter evolution over the first 1000 profiles.

By streamlining the workflow from raw data to scientific interpretation, the dashboard significantly reduces processing time while improving reproducibility in atmospheric research.

2.1.2.1.2. Software Architecture

The software implementation follows a modular architecture. The Flask-based front end orchestrated several backend modules, each responsible for a specific processing task. These modules are designed to be loosely coupled and can therefore also be reused independently in external scripts or computational notebooks. Table 1 provides an overview of the modules used in the Backscatter Dashboard.

Table 1: Overview of Modules for Backscatter Dashboard.

Module	Responsibility
config.py	Global configuration such as upload and output folders, allowed file extensions, and maximum upload size.
data_reader.py	Robust loading and normalization of raw export files. Detects delimiter and encoding, interprets or constructs headers, and casts columns to appropriate types.
data_inspector.py	Detection of the temporal column and computation of data ranges and per-session timeline summaries.
data_converter.py	Smart detection of the backscatter profile column, automatic classification of the profile encoding (ASCII-encoded versus direct hexadecimal), and decoding into physical backscatter profiles.
data_visualizer.py	Generation of scientific plots: single profiles, temporal comparisons, time-height cross sections, and mean profiles with variability bands.
pdf_generator.py	Assembly of multi-page PDF report that contains tables and plots generated in previous stages.
appy.py	Flask application that defines the web routes, manages the session state, and coordinates calls to all backend modules.

2.1.1.2 Troubleshooting Rotating Shadow Band

During the initial commissioning of the RSB-02, the RS485 interface, operating under the Modbus RTU protocol, was subjected to its first functional test. The primary objective was to establish a stable communication link with the device and successfully read its Modbus registers. However, unexpected complications arose when the master software (Obi, developed by EKO Instruments) failed to receive any response from the RSB-02, despite repeated attempts. To systematically isolate the root cause of the communication failure, a structured diagnostic approach was implemented, comprising the following steps.

2.1.2.2.1. Verification of the USB-to-RS485 Converter

The first diagnostic measure involved validating the functionality of the USB-to-RS485 converters, which serve as the physical interface between the host computer and the RSB-02. Using a known reference sensor (a pre-tested pyranometer or similar Modbus-compatible device), the converter was subjected to the following tests:

- **Driver installation:** Manual installation of the required USB serial drivers was performed to ensure compatibility with the host operating system.
- **Terminal-based communication test:** The converter was connected to a terminal emulator (Hercules SETUP utility) to verify bidirectional data transmission. A successful handshake and register readout confirmed that the converter was fully operational and suitable for subsequent testing.

2.1.2.2.2. Communication Attempts with the RSB-02 via Obi Software

With the converter's functionality confirmed, the next step involved testing the communication link with the RSB-02 using the Obi software. Despite meticulous configuration of the interface parameters – including baud rate, parity, and slave ID – no successful communication could be established. Key observations included:

- **Manual device detection:** Repeated attempts to manually detect the RSB-02 via its Modbus address yielded no response.
- **Autodetect function failure:** The software's built-in Autodetect feature similarly failed to recognize the device.
- **Software-guided troubleshooting:** Following the manufacturer's recommended diagnostic procedures (e.g. adjusting timeout settings, verifying termination resistors) produces no improvements.
- **Physical layer checks:** Swapping the A/B data lines (to rule out polarity issues) and testing alternative cables had no effect on the communication failure.

2.1.2.2.3. Independent Testing of the C-BOX-S2 Communication Module

To further isolate the issue, the C-BOX-S2 communication module, a critical intermediary between the RSB-02 and the host system, was subjected to independent testing. The following steps were undertaken:

- **Alternative device connection:** The C-BOX-S2 was connected to an MS-80SH pyranometer (a known functional Modbus RTU Device) using the Hibi software.
- **Register readout validations:** The pyranometer responded correctly to Modbus queries, and its registers were successfully read, confirming that:
  - The C-BOX-S2's physical interface (RS485 transceiver) was operational.
  - The Modbus protocol implementation in the C-BOX-S2 was functional for at least one device.
- **Firmware update attempt:** Given the successful communication with the pyranometer, a potential firmware incompatibility with the RSB-02 was suspected. An attempt to update the C-BOX-S2's firmware was made but failed due to the device's non-responsiveness.

2.1.2.2.4. Advanced Diagnostic Measures

To eliminate all remaining variables, additional tests were conducted under controlled conditions:

- **Power supply verification:** The supply voltage to the C-BOX-S2 was increased to 24V (the upper limit of its specified operating range) to rule-out insufficient power delivery.
- **Isolation testing:** All connected devices (RSB-02 and pyranometer) were completely disconnected from the C-BOX-S2 to test its standalone functionality. Even in this minimal configuration, the C-BOX-S2 failed to establish communication with the host.
- **Hardware reset attempts:** Multiple power cycles and factory resets of the C-BOX-S2 were performed, with no observable change in behaviour.

2.1.2.2.5. Advanced Diagnostic Measures

The cumulative results of these diagnostic steps strongly suggested that the communication failure was not attributable to external factors. Instead, the evidence pointed toward a hardware of firmware defect within the C-BOX-S2 itself, why the device was replaced.

2.1.2 Commissioning of the Measurement Devices

After the planning phase had been completed, the measurement equipment was installed at the Maierhof site. A total of three ASI systems, each combined with a pyranometer and a weather station, the SkyVue8 lidar ceilometer, and a shadow band pyranometer were set up. The installation work was led by several staff members from the Technical University of Applied Sciences (TUAS) Rosenheim and supported by Timeless Planet GmbH. At the Illemad site, a telehandler was required to install one ASI on the roof of a greenhouse. At the technical building of the Buttenwiesen 1 PV ground mounted system, wall penetrations had to be made for cabling. The SkyVue8 and an additional ASI were mounted on the exterior wall, while the shadow band pyranometer was installed on the roof.

Several additional on-site visits were necessary before all measurement systems were fully operational and delivering reliable data. During the project, the ASI in particular had to be configurated and equipped with current software and firmware versions. Geometric calibration of the ASI systems was also carried out afterward. Due to ongoing construction work at the PV plant and the fibre optic infrastructure that was still being built at the time, the network and power supply of the measurement systems had to be adjusted multiple times and was temporarily implemented with provisional solutions. Only after the gradual optimization of all components stable continuous operation of the measurement equipment could be ensured.

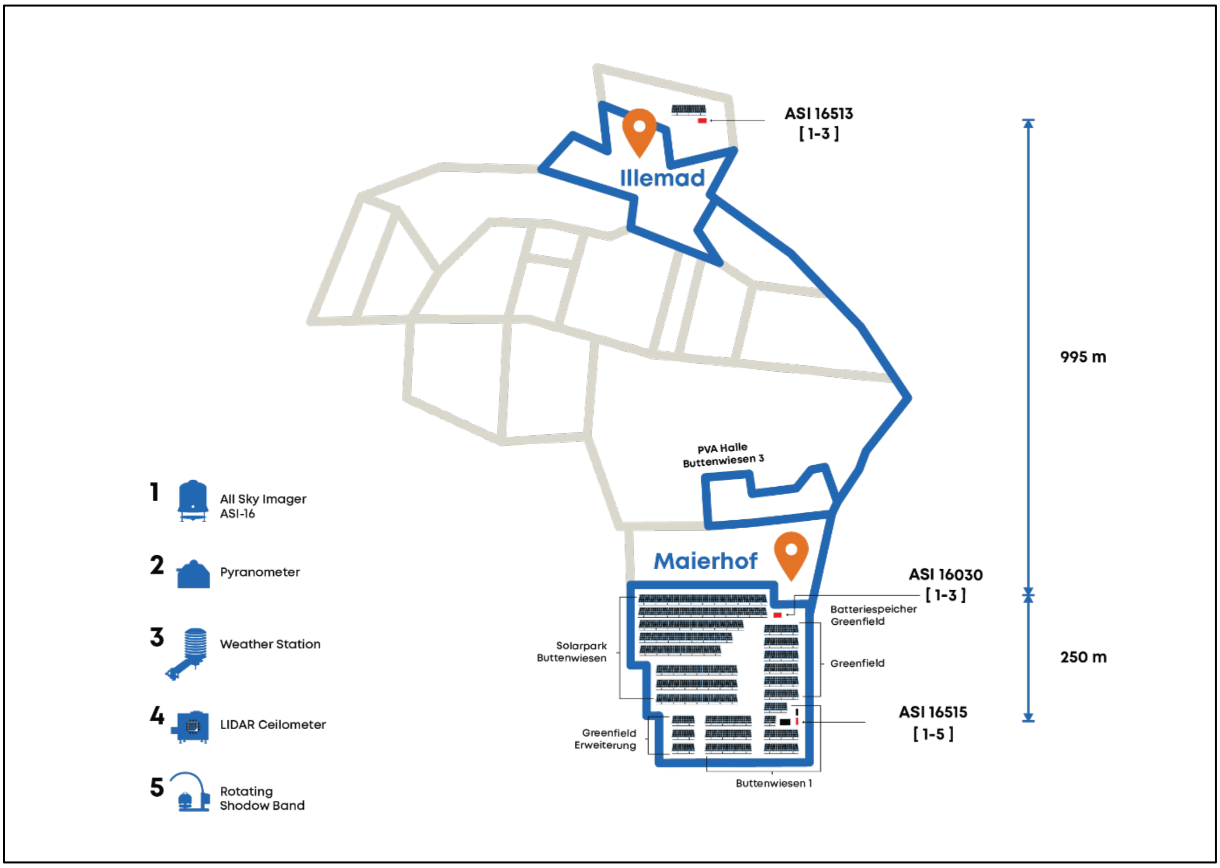

Figure 7: Overview of the Maierhof and Illemad solar parks. Measurement infrastructure and instruments deployment. The map highlights the spatial distribution of measurement locations across the two sites, including key instruments for solar resource assessment and atmospheric monitoring.

Figure 8: Photos of the installation work, on the left the camera installation in Illemad using a telehandler, on the right the installation of the SkyVue8 on the exterior wall of the technical building.

A central component of the measurement setup consists of three ASI from EKO Instruments, used to analyse cloud cover and cloud motion. At the battery storage site of the Greenfield section of the PV plant, an ASI 16/51 is operated together with an MS 80M pyranometer from EKO, ISO 9060:2018, Class A. At the additional sites in Buttenwiesen I and nearby Illemad, two newer ASI 16/55 units are installed, also combined with pyranometers, MS 80S(H), Class A. Temperature and humidity sensors, Sensirion SHT85, are installed at three locations. The sensors are connected to the ASI via Modbus, and the measurement values are written into the metadata, EXIF, of the images.

Figure 9: ASI 16/55 on the roof of a greenhouse in Illemad, the neighbouring village. In the foreground, a pyranometer for measuring global radiation is visible, as well as a ventilated radiation shield on the left, which protects a sensor for measuring air temperature and humidity from direct sunlight.

Figure 10: ASI 16/55 on the roof of a technical container at the Buttenwiesen I site. As in the previous figure, the pyranometer for measuring global radiation is visible here as well. The radiation shield with the temperature and humidity sensor is not visible due to the viewing angle.

To calculate the cloud base height (CBH), using stereoscopic methods, the two ASI 16/55 units were installed at a distance of 1,200 meters from each other. Stereoscopic imaging makes it possible to compute individual heights for detected clouds based on simultaneous observations from both sites. The images are captured every 20 seconds to derive cloud motion vectors (CMV) from consecutive images. Experience has shown that capturing images at 20 second intervals provides an optimal solution for both fast moving and slowly moving clouds.

Figure 11: SkyVue8 lidar ceilometer installation at the Buttenwiesen I site. The device is mounted on the exterior wall of the measurement container to monitor cloud base height and vertical aerosol profiles. To mitigate potential damage from wildlife, the device was installed in an elevated position.

In addition, a SkyVue8 lidar ceilometer from Campbell Scientific was installed. The device records backscatter profiles up to a height of 8 kilometres and provides information on cloud heights and atmospheric layering with a vertical resolution of 5 meters. By comparing these measurements with the results from our stereoscopic method, we can verify the accuracy of cloud heights derived from image data.

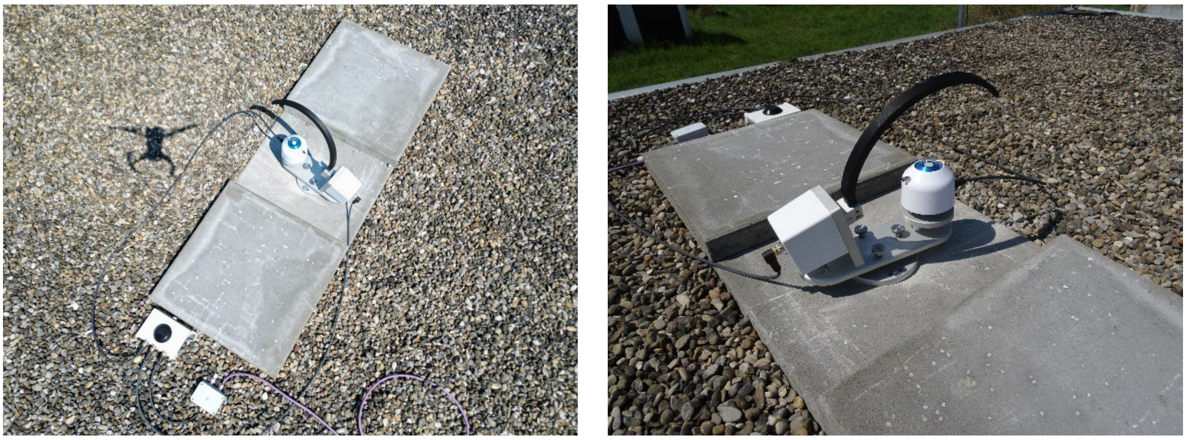

Figure 12: Rotating shadow band (RSB) installation at the Buttenwiesen I site. The system comprises an EKO Instruments MS-80SH pyranometer, a control unit (C-Box) with integrated GPS, and a rotating shadow band for the separate measurement of direct and diffuse solar irradiance.

For the measurement of solar radiation, the setup is complemented by a rotating shadow band (RSB) combined with an MS-80SH pyranometer, Class A, from EKO Instruments. The device measures global horizontal irradiance (GHI), and diffuse horizontal irradiance (DHI), and calculates direct normal irradiance (DNI).

2.2 Geometric Calibration of ASI

For the geometric processing of the images, calibration of the fisheye lenses is required. Due to their wide-angle imaging, the projection geometry deviates significantly from that of an ideal central projection. To model this nonlinear mapping, the Matlab toolbox OCamCalib [ScMS06] [DaMS06] [RuSS08] [Scar08] was used. This method enables the calibration of central omnidirectional cameras using a fourth order polynomial model, which describes the relationship between the distances of an image point from the optical centre and the corresponding incident angle of the incoming light ray.

The calibration determines both the intrinsic and extrinsic camera parameters:

- Intrinsic parameters describe the internal properties of the camera system, including:
  - The coordinates of the image centre,
  - The image dimensions (height and width),
  - The polynomial coefficients that model the nonlinear relationship between the angle of an incident light ray and its radial distance from the image centre. This polynomial coefficients replace the classical focal length and accounts for the fisheye lens's distortion.
- Extrinsic parameters, in contrast, define the spatial position and orientation of the ASI in the world coordinate system.

For calibration, multiple images of a calibration pattern (chessboard) were captured from different viewing angles (Figure 13). OCamCalib automatically detects the chessboard corners and computes the camera parameters, generating a calibration model. Using this model, the functions world2cam and cam2world enable bidirectional projection:

- world2cam maps arbitrary 3D world points onto the 2D image coordinates of the ASI.
- cam2world performs the inverse transformation, reconstructing 3D directions from image points.

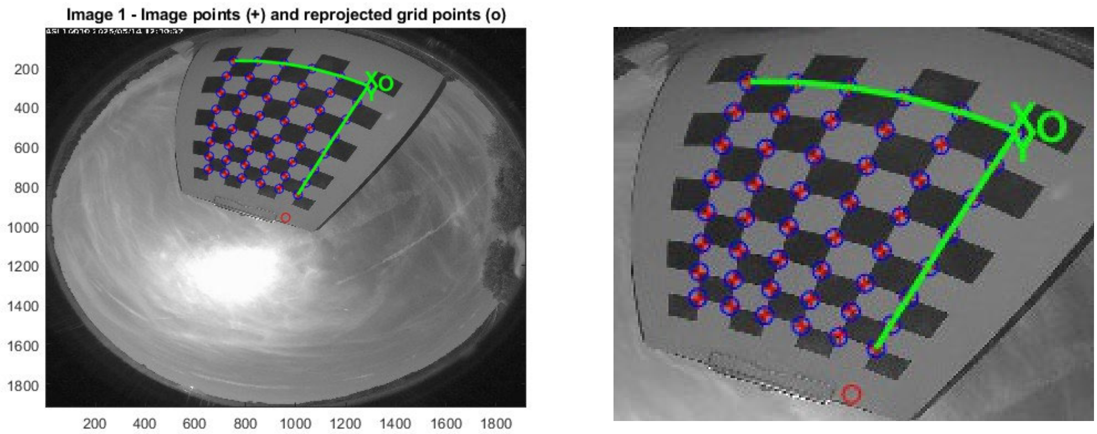

Figure 13: Calibration grid corner detection and projection in All-Sky Images. The image displays the result of automatic corner extraction for the calibration pattern. Pink crosses (+) mark the detected grid corners, while blue circles (o) indicate the corresponding projected grid corners based on the calibrated camera model. The red circle at the edge of the clipboard denotes the determined image centre.

This calibration is essential for accurately interpreting azimuth and elevation angles in the image and serves as the foundation for precise analyses, such as solar position computation.

2.2.1. Calibration Results

Table 2 summarizes the fourth-order polynomial model parameters (affine and distortion coefficients) and reprojection errors (RMSE, max error) obtained via OCamCalib for each ASI model. Calibration was performed using chessboard pattern images captured from multiple viewing angles (Figure 13). Not all images could be analysed, which is why some images from the calibration process were excluded.

Table 2: Calibration Results Using OCamCalib for Two Different ASI Models.

Modell	ASI-16515	ASI-16030
Location	Buttenwiesen-I	Greenfield
Basename (no suffix)	calib_ASI16515_140525_0	calib_ASI16030_14025_0
Number of Images	15	16
Average Reprojection Error Computed for Each Chess-board [Pixels]	001: 4.97 ± 6.07 004: 0.78 ± 0.49 005: 0.68 ± 0.32 006: 0.81 ± 0.45 007: 1.25 ± 0.63 008: 0.52 ± 0.32 011: 1.17 ± 0.76 012: 1.24 ± 0.73 013: 0.88 ± 0.51 016: 0.91 ± 0.46 017: 5.30 ± 9.12 018: 1.16 ± 0.55 019: 1.10 ± 0.61 021: 1.78 ± 0.99 022: 0.99 ± 0.46	001: 2.00 ± 0.90 002: 2.64 ± 1.34 003: 2.34 ± 1.54 004: 2.58 ± 2.04 005: 2.60 ± 1.50 006: 2.39 ± 1.55 007: 3.71 ± 2.05 008: 1.71 ± 0.96 009: 3.36 ± 1.72 010: 5.88 ± 2.68 012: 3.19 ± 1.85 014: 2.50 ± 1.21 016: 2.28 ± 0.94 017: 3.04 ± 1.30 018: 2.95 ± 1.48 020: 3.17 ± 1.96
Excluded Images	002 003 009 010 014 015 020	011 013 015 019
Average Error [Pixels]	1.624871	2.896314
Sum of Squared Errors	9227.032876	9092.574049
ss	1.0e+02 * -6.335456494767021 0 0.000005575211982 -0.000000000840253 0.000000000002138	1.0e+02 * -6.758863193753300 0 0.000006145828378 -0.000000003659945 0.000000000004352
xc **yc**	1.024135787932556e+03 1.023929391484608e+03	9.600452634785977e+02 9.598787665111754e+02
c **d** **e**	1.000355376645339 4.688112427479249e-04 6.188160344287178e-04	1.0014 7.9748e-05 2.2185e-04
Width **Height**	1920 1920	1920 1920

2.2.2. Projection Model and Calibration Validation

The fisheye lens projection is modelled using a fourth-order polynomial equation (Figure 14), which establishes the relationship between the radial distance r of an image point from the optical centre and the incident angle Θ of the corresponding 3D ray. The calibration process generates an ocam_model structure containing all intrinsic parameters.

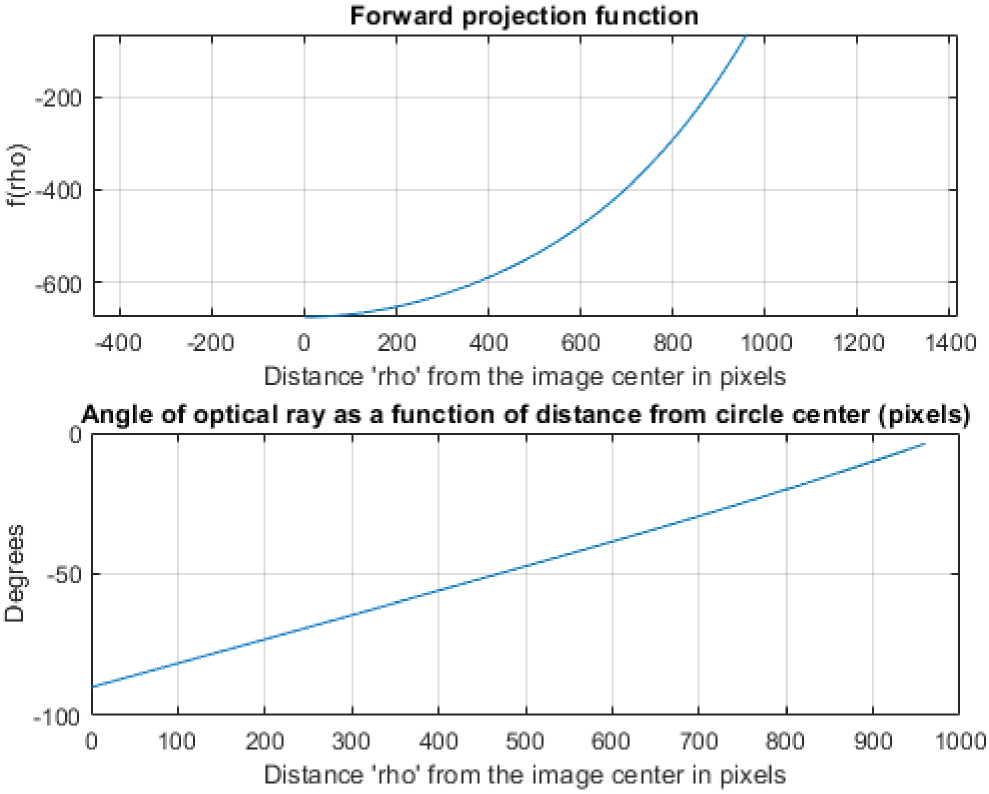

Figure 14: Graph representing function F and plot of angle Θ of the corresponding 3D vector with respect to the horizon (#16030).

The spatial distribution of the chessboard patterns (Figure 15) demonstrates the diversity of poses and orientations captured during calibration. By covering a wide range of viewpoints, from near-field to far-field and varying angles, the procedure ensures robust parameter estimation. The comprehensive sampling of the camera's field of view (FOV) enhances the stability of the derived model, ultimately yielding a highly accurate representation of the system's projection behaviour.

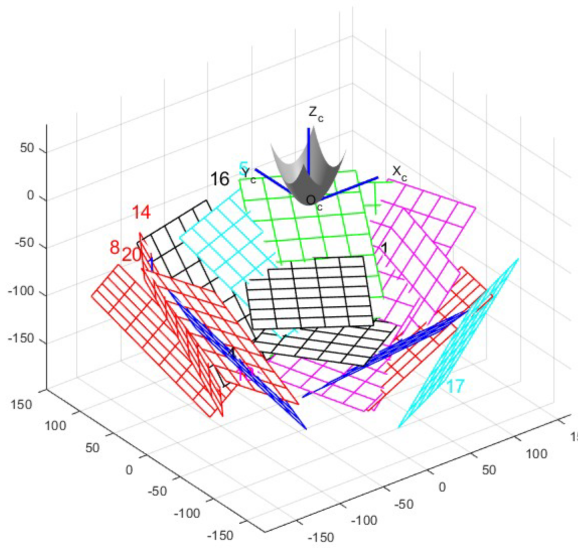

Figure 15: Position of every checkerboard with respect to the reference frame of the ASI #16030 during calibration process.

Figure 16 illustrates the distribution of the reprojection error fore very detected point across all checkerboard images. The colour assignment differentiates the individual images, allowing a clear visual separation of their respective error patterns. The overall spread of the errors provides insight into the consistency of the calibration results and supports the assessment of how well the estimated camera model aligns with the observed image data.

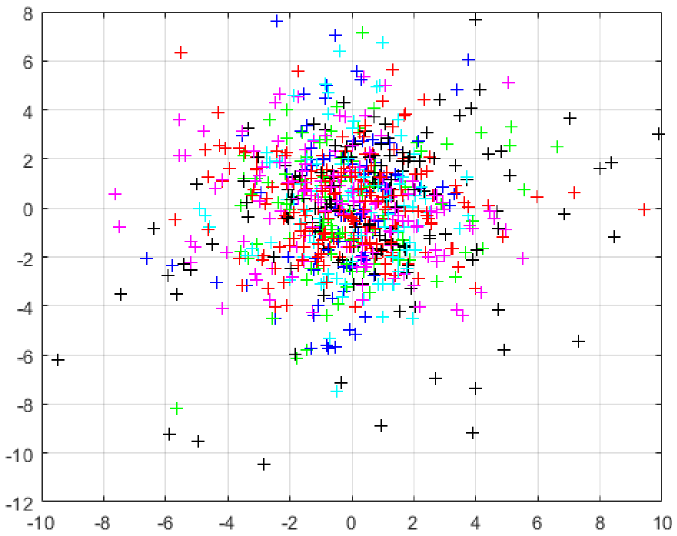

Figure 16: Distribution of the reprojection error of each point for all the checkerboards. Different colours refer to the different images of the checkerboard (see Figure 15).

2.3 API Interface for PV Monitoring

The technical monitoring of the ground mounted PV system at the Maierhof site in Buttenwiesen is using the blue'log data logger in combination with the VCOM (Virtual Control Room) monitoring platform provided by meteocontrol GmbH [Vcom05]. The system continuously records operational data, including electrical output, system status, and relevant environmental parameters, ensuring comprehensive supervision plant performance.

These data are made accessible via a standardized API interface, enabling seamless integration with project-specific algorithms and analytics tools. This facilitates both real-time evaluation and long-term performance analysis, supporting advanced monitoring, diagnostic, and optimization of the PV system.

2.4 Data Connection and Research Data Management

An integral part of this work ins the continuous development and maintenance of a research data management plan (RDMP). This plan ensures the systematic handling of collected measurement data, including quality assurance, traceability, structured archiving, shared use, and long-term availability.

ASI data are stored on an File Transfer Protocol (FTP) server at Bielefeld University of Applied Sciences (UAS). For the two more recent ASI systems (16/55), an additional fallback server is implemented to reduce the risk of data loss. In the event of network disruptions, data loss may still occur, as the CF/SD card (capable of storing more than 8,000 JPG images) is designed only for minute-based acquisitions intervals.

The SmartMonitoringSystem at UAS Bielefeld is responsible for continuous acquisition and storage. The system retrieves data from the PV system via the VCOM API and stores them in a

PostgreSQL database. Through a REST API-based interface, the system enables efficient data access and processing. In addition, a web application supports data acquisition, analysis, and visualization, utilizing standard open-source technologies and Java EE components. The integration of PV production data allows comprehensive monitoring and facilitates the identification of data gaps and inconsistencies.

Measurement data from the RSB and the SkyVue8 ceilometer are stored at the TUAS Rosenheim. SkyVue8 backscatter profiles, represented as one-dimensional vectors, are stored as byte streams (BLOBs) in a MariaDB database, while all other measurement data are recorded as time series in the InfluxDB (MetricX) database.

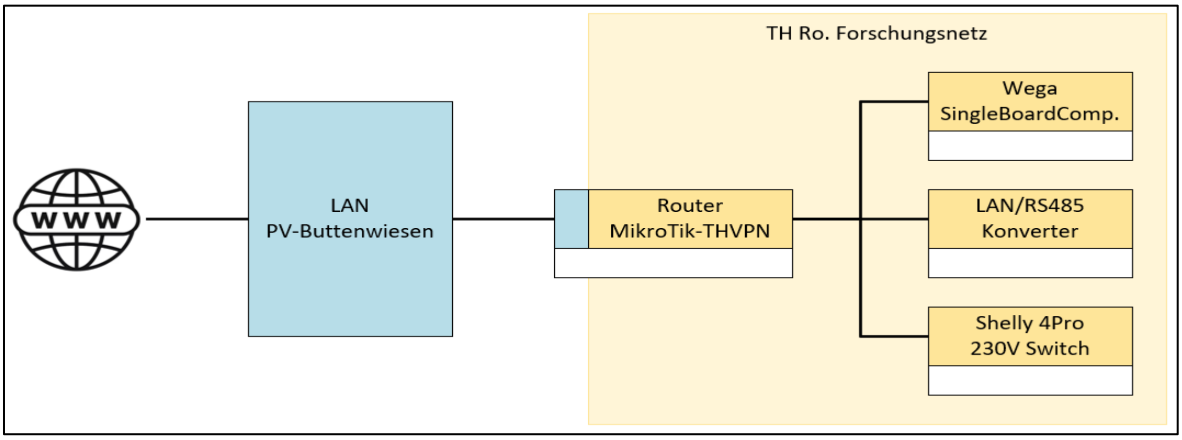

Figure 17: The data from the RSB and SkyVue8 are transmitted from Buttenwiesen via a router and a VPN connection directly into the university research network, where they are stored.

2.4.1 HELIOS Data Monitoring System

The HELIOS Data Monitoring System (HDMS) is a Python-based tool designed to automatically monitor image data stored on FTP servers and detect data quality issues at an early stage. Given that each camera produce image every 15 seconds (approximately 5,760 images per day), manual inspection becomes impractical when operating multiple systems continuously.

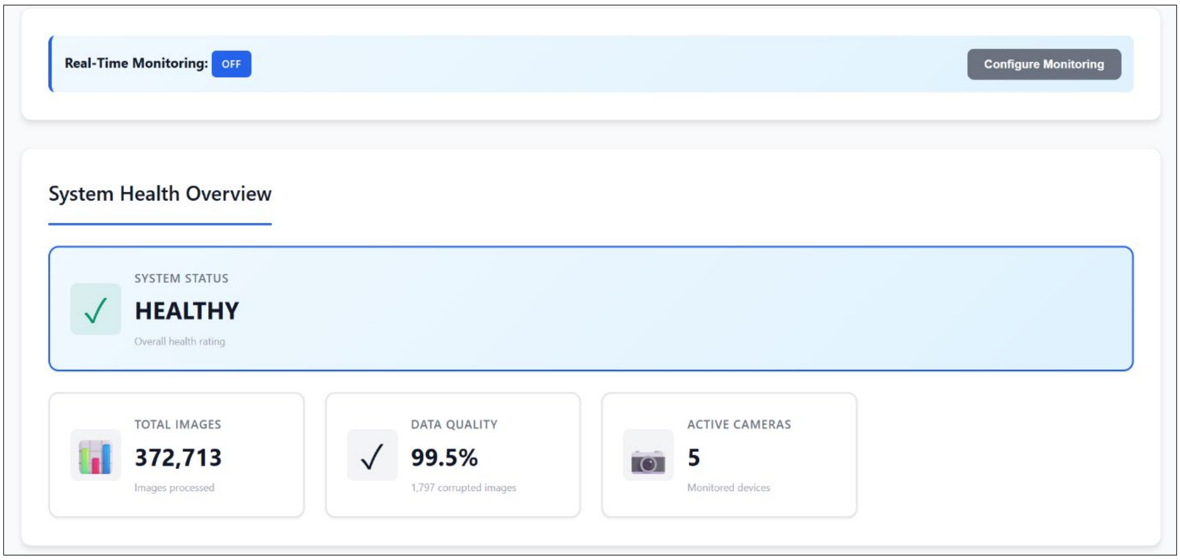

Figure 18: Dashboard overview of HDMS, showing real-time monitoring status, system health, and key data quality metrics.

HDMS addresses this challenge by automating data quality control. It identifies corrupted or incomplete image files, extracts sensor information embedded in image metadata, and logs validation results to enable timely intervention by system operators.

2.4.1.1. Problem and Goals

In long-term camera network operations, typical failure modes include interrupted data transfers resulting in truncated or zero-byte images, missing or incomplete metadata, and sensor malfunctions such as drift or signal freezing. These issues can undermine the reliability of subsequent analyses and may negatively impact machine learning pipelines by introducing erroneous misleading input data.

The HELIOS HDMS addresses these challenges by implementing consistent and reproducible data validation procedure. It provides clear outputs in the form of dashboard and reports, translating detected anomalies into actionable maintenance tasks and thereby supporting efficient system operation and data quality assurance.

2.4.1.2. Architecture Overview

HDMS follows a pipeline in which it scans FTP directories, validates each image, parses EXIF metadata for sensor values, applies rule-based checks, and stores outcomes in SQLite. To support operations, the database separates per-image inspection records from run metadata (for checkpointing) and form aggregated summaries used by the dashboard.

FTP → Scan → Image verify → EXIF parse → Sensor checks → SQLite → Aggregates → Dashboard / Reports

2.4.1.3. Core Checks: Image Integrity and Sensors

Image validation combines fast file-level checks with content-level verification. HDMS flags empty files, detects common transfer truncations, and uses Pillow (PIL) to open and verify images to catch structural corruption. Metadata extraction focuses on EXIF fields commonly used by ASI, including sensor values stored such as UserComment.

Sensor validation mixes plausibility bounds with anomaly detection. Typical checks include temperature, humidity, and irradiance ranges, treatment of known invalid or sentinel values missing, and struck-sensor detection when readings repeat for many consecutive frames. With a 15-second cadence, a threshold on the order of minutes of identical readings is often enough to catch failures without over triggering on natural plateaus.

2.4.1.4. Operations: Scanning Modes and Reporting

HDMS supports full scans for initial backfills, incremental scans for routine monitoring, and resume capability using checkpoints recorded in the database. Incremental scanning keeps daily workloads close to the new data volume per camera, which is the key to scaling beyond a single device.

Results are exposed through a lightweight web dashboard and through generated reports (and exports such as CSV) that summarize corruption rates, data gaps, and sensor anomalies per camera and per day. In practice, operators use these outputs to distinguish network issues from device-side problems and to prioritize maintenance.

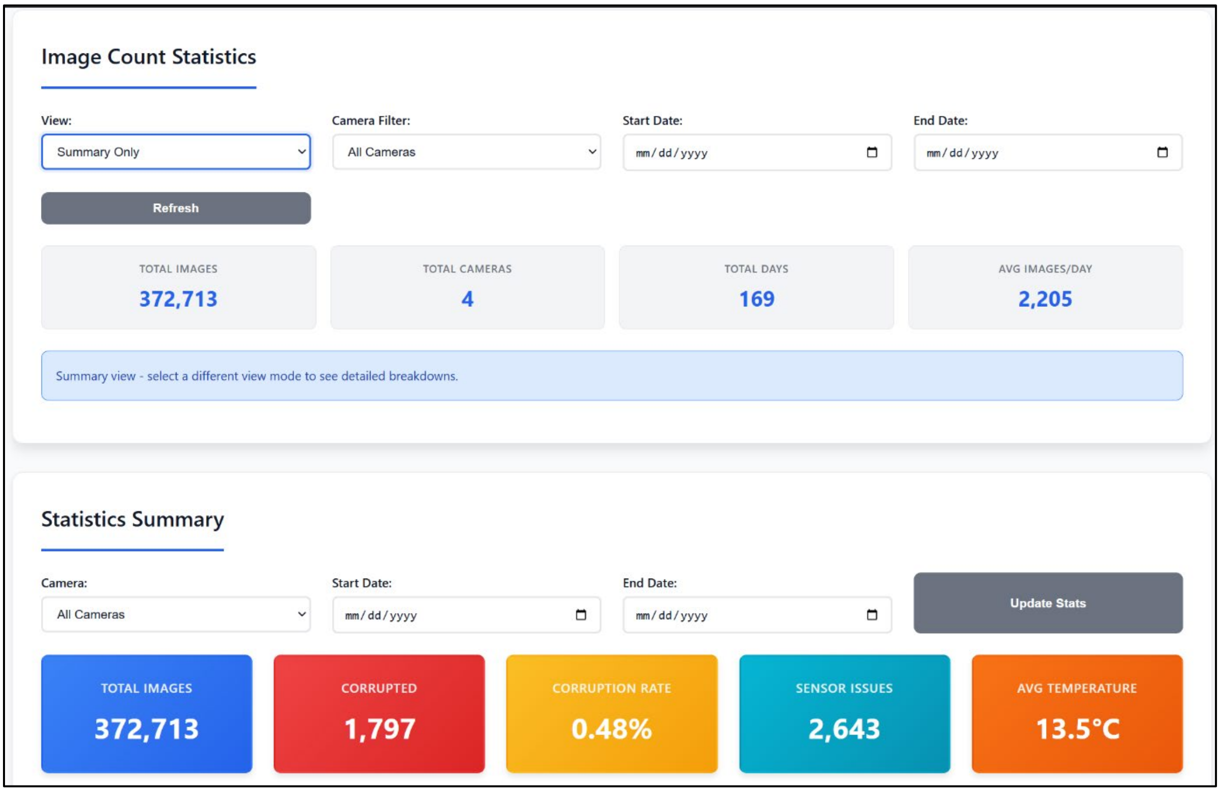

Figure 19: Image statistics dashboard summarizing dataset volume, camera coverage, data quality, and sensor-related metrics over a selected time range.

2.4.1.5. Relevance for Deep Learning

Within deep learning workflows, HDMS mitigates hidden data-quality issues that can otherwise compromise model performance. The exclusion of corrupted or incomplete images reduces training instability and prevents the introduction of spurious artifacts. At the same time, validated metadata minimizes the risk of false corrections when sensor information is used for filtering or as an additional model input.

Furthermore, the use of SQLite-based records ensures full traceability of model inputs, linking each data sample to its acquisition time and validation status. This significantly enhances reproducibility and supports systematic auditing of the data processing pipeline.

3. Analysis of All-Sky Images

3.1. Calculation of Sun Position and Path

This chapter describes the methodology for calculating and projecting the position and trajectory of the sun in All-Sky Images. The objective is to accurately map the true solar position, defined by azimuth and elevation angles, onto the image plane of calibrated fisheye camera.

The workflow comprises several steps, beginning with the astronomical computation of the solar position and concluding with its projection into the camera's image coordinate system.

3.1.1. Fundamentals and Preparatory Steps

The procedure begins with check for leap year conditions to ensure the temporal correctness of subsequent astronomical computations. The solar position, defined by azimuth and elevation angles, is then calculated using the algorithm specified in DIN 5034, which provides the sun's position in a geocentric reference frame suitable for illumination analysis.

All angular quantities are converted from degrees to radians using the Matlab function deg2rad ensuring consistency and compatibility with subsequent numerical operations.

3.1.2. Transformation to Cartesian Direction Vectors

For projection onto an image, the sun is modelled as a direction rather than a spatial point. Due to extremely large distances between Earth and the Sun, its position can be accurately represented by a unit direction vector.

The transformation from spherical to Cartesian coordinates follows the astronomical convention, where the azimuth is measured clockwise from north and the elevation defines the angle above the horizon. Based on this convention, the three-dimensional direction vector is computed as:

```
x = cos(el_rad) * sin(az_rad);
y = cos(el_rad) * cos(az_rad);
z = 1 % unit circle;
```

Solar positions below the horizon are excluded from the projection process by assigning Not-a -Number (NaN) values to the corresponding vector components, ensuring that only physically visible sun position are considered in the subsequent analysis.

3.1.3. Coordinate System Adaption for OCamCalib

OCamCalib adopts a coordinate convention that differs from standard astronomical reference frames. While astronomical systems define the positive z-axis upward (towards the zenith), OCamCalib assumes the positive z-axis to be directed toward the scene (i.e., along the camera's viewing direction).

To reconcile these conventions, the z-component of the direction vector is inverted prior to projection. The resulting three-dimensional point vector is defined as:

```
M = [x
     y
     z * (-1)];
```

The transformation ensures a consistent alignment of the optical axis, such that a direction corresponding to the zenith is correctly projected onto the centre of the image.

3.1.4. Projection onto the Camera Image Plane

Following calibration of the fisheye lens uses the MATTLAB OCamCalib toolbox, the intrinsic projection model is accessed via ocam_model = calib_data.ocam_model (Chapter 2.2). The previously computed direction vectors are projected onto the image plane using the function world2cam, which returns a two-row matrix containing the horizontal and vertical image coordinates.

In MATLAB image representations, the coordinate origin is located in the upper-left corner, whereas OCamCalib defines its origin in the lower-left corner. To ensure consistency between these coordinate systems, the vertical image coordinate is inverted according to:

```
m = world2cam(M, calib_data.ocam_model);
ocam_model = {ss, xc, yc, c, d, e, width, height}
u = m(1)
v = ocam_model.height – m(2)
```

Table 3: Overview of the OCamCalib Parameters.

Variables	Description
ss	Polynomial coefficients describing the projection function
xc yc	Row and column coordinates of the image centre
c d e	Affine transformation parameters
width height	Dimensions of the image

Each parameter contributes to the non-linear mapping between the three-dimensional direction vectors and their two-dimensional coordinates. A correct understanding of these parameters [Scar25] is essential for accurate projection and reprojection.

Prior to visualization, the image itself is flipped horizontally to account for the optical and geometric characteristics of the specific camera system:

```
im_asi = fliplr(im_asi)
```

Two All-Sky images featuring the calibration pattern used to determine the image geometry are shown. The projected sun position, computed from azimuth, elevation and observer location data, is highlighted in red (Figure 20).

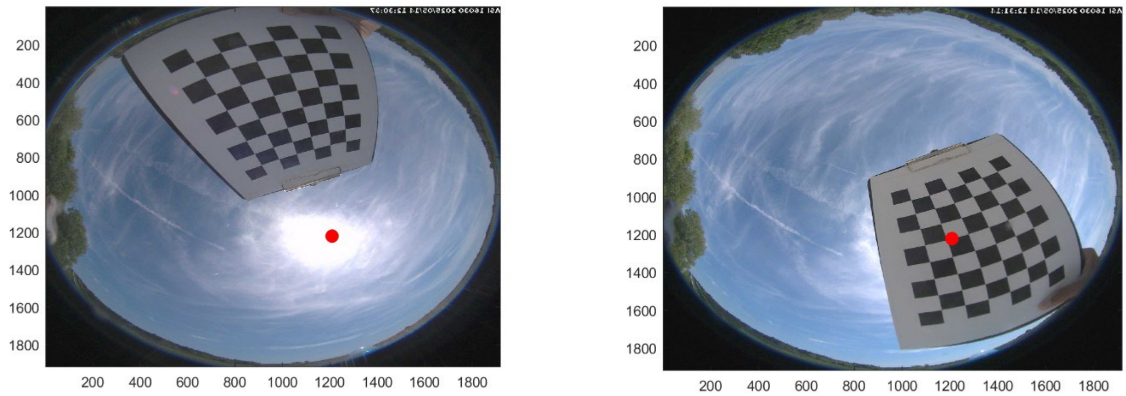

Figure 20: Projection of the solar position onto calibrated images.

3.2. Cloud Detection with Machine Learning

3.2.1.A Hybrid MAE-DINO Approach

Ground based ASI continuously capture hemispheric views of the sky, providing valuable data for solar energy forecasting. However, developing accurate computer vision models for these images faces a critical challenge: manually labelling cloud types and conditions is expensive and time-consuming, resulting in datasets too small to train robust deep learning systems. Additionally, images vary dramatically across locations due to different climates, camera hardware, and lightning conditions, making it difficult for models trained in one location to work well elsewhere. We address these limitations through self-supervised learning, a technique that allows models to learn useful visual features from unlabelled images by solving pretext tasks rather than relying on human annotations. Our approach combines two complementary self-supervised methos: Masked Autoencoders (MAE), which learn by reconstructing randomly hidden portions of images, and DINO (self-distillation), which learns by maintaining consistent representations across different views of the same image. By training a Vision Transformer encoder with both objectives simultaneously, our hybrid model learns to capture both fine-grained cloud textures through reconstruction and high-level semantic patterns through consistency, while remaining robust to illumination changes and geometric distortions inherent to fisheye lenses.

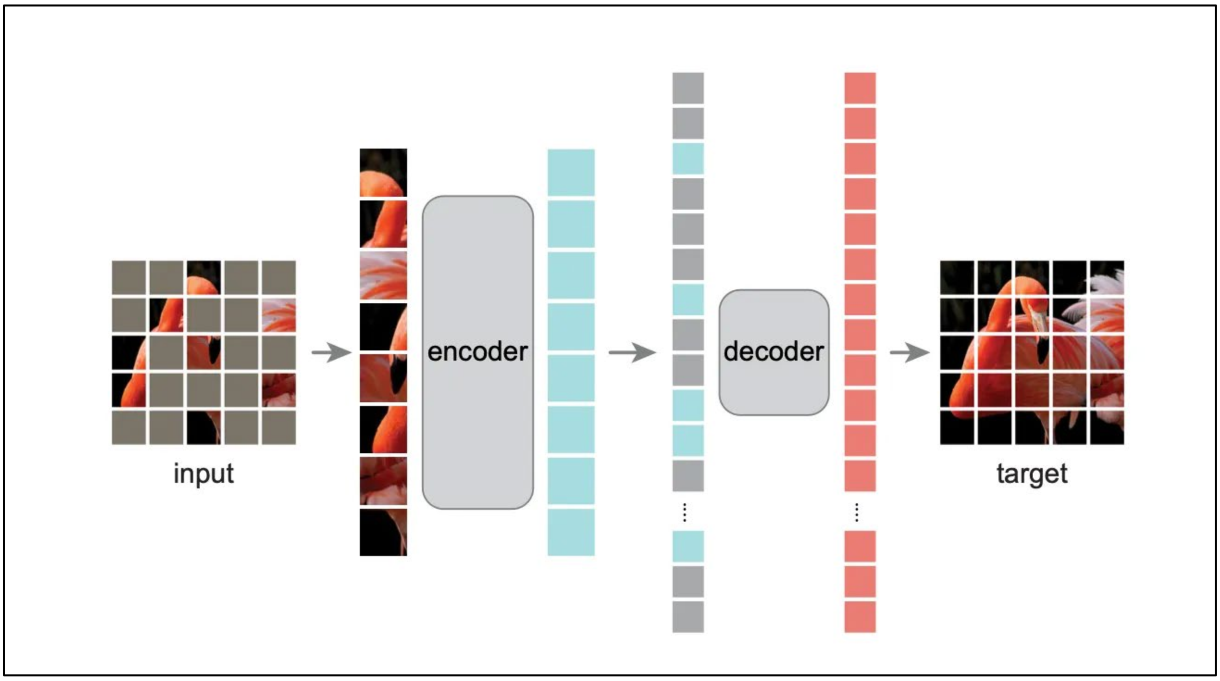

Figure 21: Masked autoencoder architecture where an encoder processes visible patches and a decoder reconstructs the missing image regions.

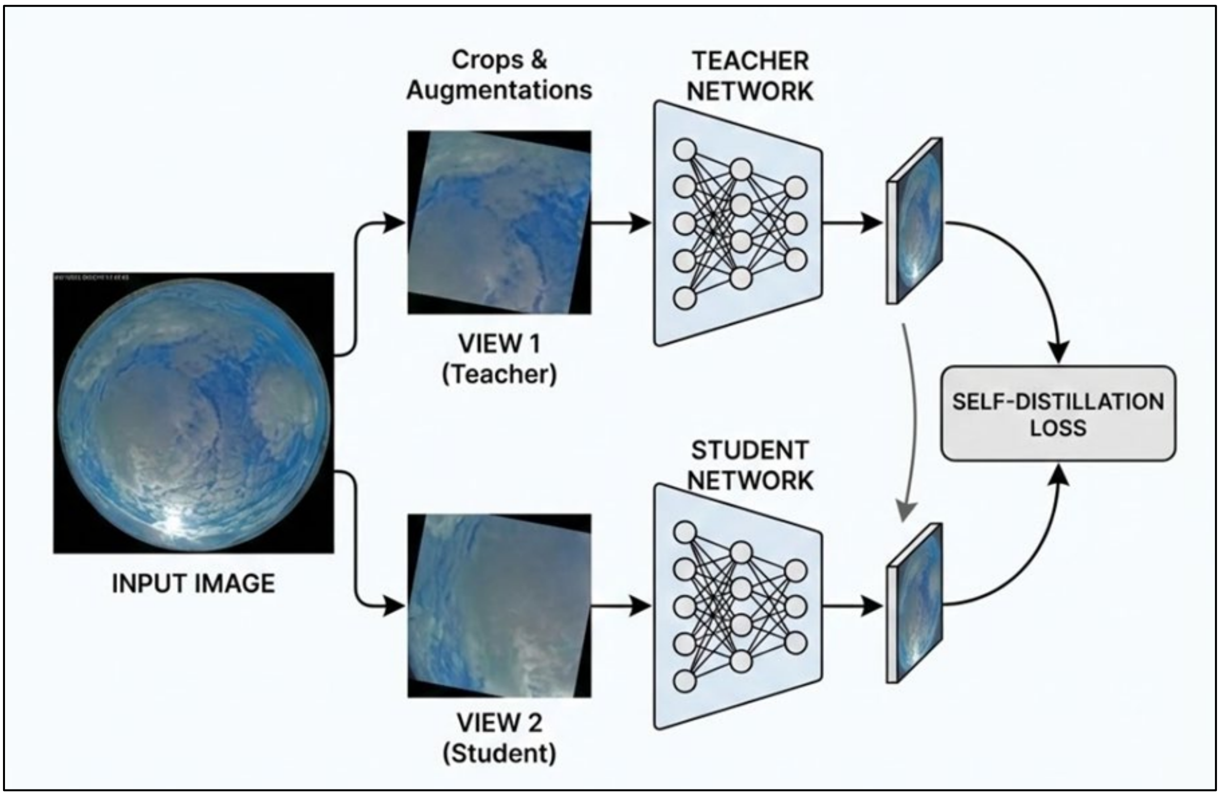

Figure 22: Self-distillation framework where different augmented views of the same image are processed by teacher and student networks, with training driven by a consistency loss.

We pre-trained our model on 150,000 unlabelled images and evaluate it on cloud classification and segmentation tasks. The hybrid approach achieves 89.5 % classification accuracy with minimal task-specific training, performing competitively with models pre-trained on millions of general images from ImageNet. Our results demonstrate that self-supervised learning can effectively extract meaningful representations from unlabelled meteorological imagery, reducing the need for expensive manual annotation while producing models that generalize across different weather conditions and camera deployments. This work represents an important step towards foundation models for atmospheric computer vision that can support operational solar forecasting and weather monitoring with substantially reduced data labelling requirements.

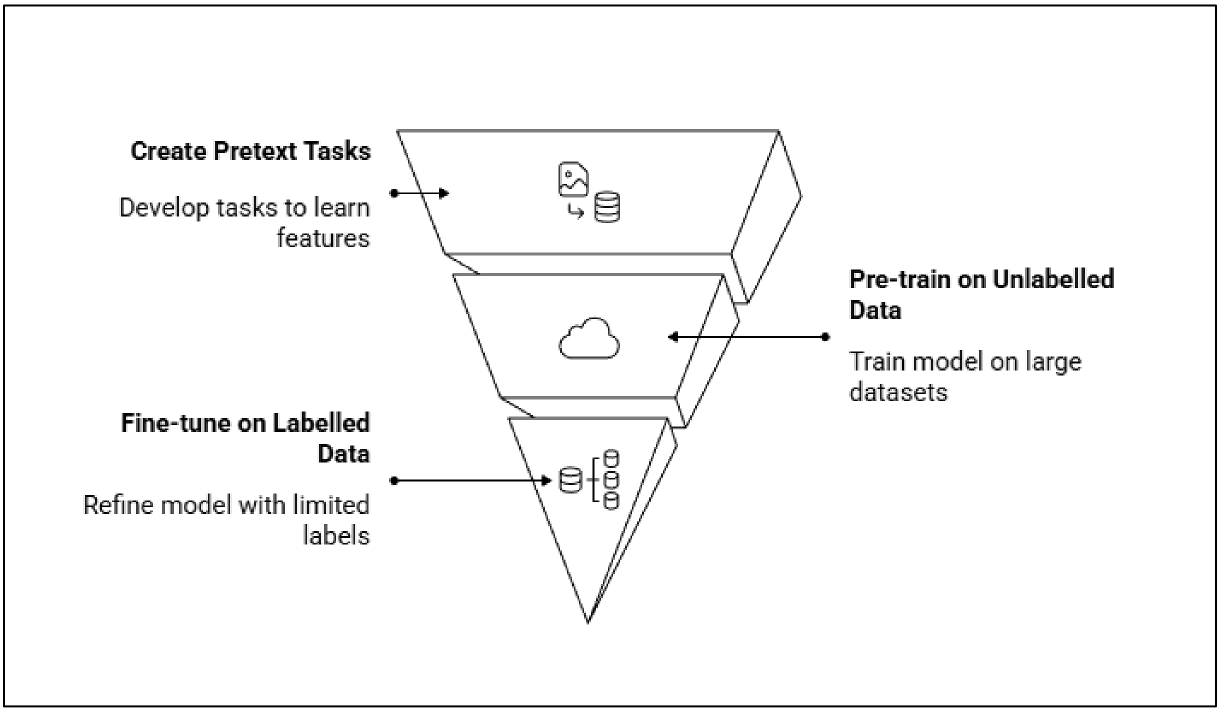

Figure 23: Self-supervised learning workflow: pre-training on unlabelled data with pretext tasks, followed by fine-tuning on limited labelled data.

3.3. Analysis of Solar Irradiance

For the analysis of solar irradiance, a measurement system consisting of three ASI and pyranometer was used. The cameras recorded hemispherical images at 15 seconds intervals, while the rotating shadow band measured global irradiance with one second resolution. The resulting files contain the mean value of all measurements collected with each 60 second period. Data processing and analysis of the radiation measurements were carried out in MATLAB. In addition, the images were visually evaluated by a weather observer from the German Weather Service in order to characterize cloud conditions more precisely and to improve the interpretation of the radiative data. The evaluation of an example data record is shown below.

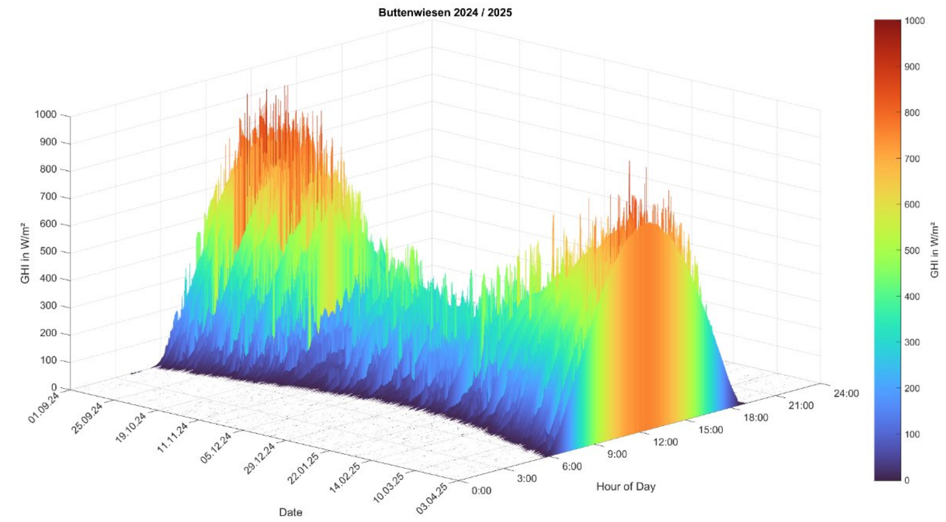

Figure 24: 3D annual overview of measured GHI in Buttenwiesen for the years 2024/25.

3.3.1. Influence of Clouds on Solar Irradiance

The analysis of the irradiance data together with the supplementary cloud observations shows that both irradiance enhancements and strong radiation gradients occur primarily under specific cloud types and synoptic conditions. Days characterized by an unstable atmosphere and convective cloudiness are particularly affected.

Overall, the following relationships were identified:

- Irradiance enhancements occur preferentially in the presence of cumulus clouds and mid-level altocumulus clouds.
- Especially pronounced enhancements develop around midday under broken convective cloud fields, intensified by reflection and scattering effects at cloud edges.
- Strong irradiance gradients are frequently observed during periods of variable convective cloudiness, in particular cumulus and altocumulus.
- The largest gradients result from rapid changes in cloud cover during the midday period and from high cloud velocities.

3.3.2. Sample Analysis for 15 September 2024

On this day, the strongest and most rapid increase in GHI observed over the entire measurement period was recorded. Within a very short time, the value rose extremely steeply, indicating a sudden clearing of the sky following the passage of a dense cloud. A notable feature of this day is a systematic time offset of approximately four minutes between the irradiance measurements and the sky images. This offset was most likely caused by a brief power interruption or an error in the synchronization of the measurement systems. To ensure correct interpretation, the measurement data were temporally adjusted by manually aligning them with the corresponding sky image.

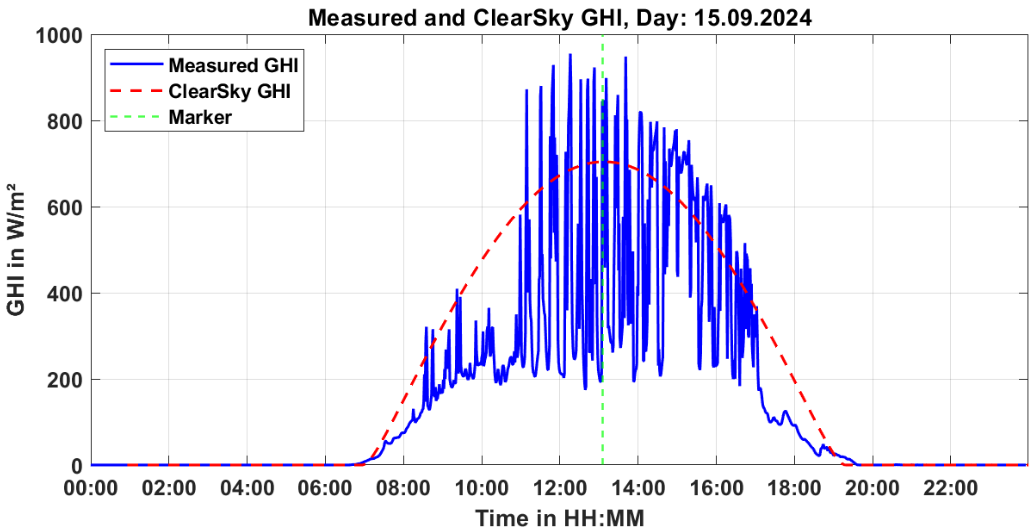

Figure 25: Diurnal evolution of the measured GHI in blue and the calculated clear sky irradiance in red. A pronounced rapid increase is observed between 13:04 and 13:05 local time, during which the irradiance rises from 213.2 W/m² to 844.8 W/m².

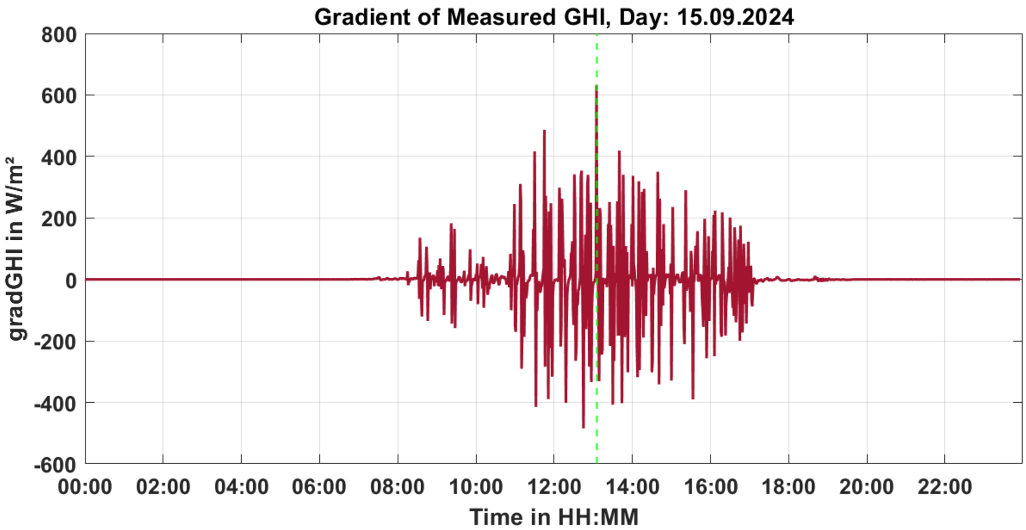

Figure 26: Temporal evolution of the global irradiance gradients. The maximum gradient reaches 631.6 W/m² per minute and is indicated by the green dashed line at approximately 13:05 local time.

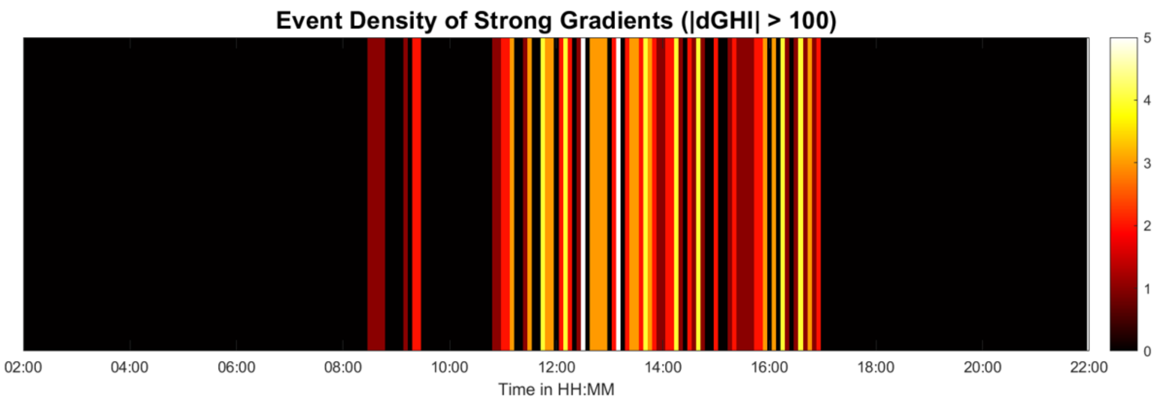

Figure 27: Diurnal distribution of the event density of strong gradients on 15 September 2024.

The figure illustrates the temporal density of strong irradiance gradients, defined as changes greater than 100 W per square meter per minute, throughout the day. Each coloured bar represents a five-minutes interval and indicates how frequently a significant jump in global irradiance occurred within that period. The lighter the bar, the more events were recorded. Bright areas are particularly prominent between 11:00 and 17:00, highlighting periods during which rapid and repeated increases or decreases in irradiance took place. These intervals of high event density point to especially unstable cloud conditions and highly dynamic changes in solar radiation.

Time series of GHI during the rapid increase event: at 13:04, the irradiance was 213.2 W/m^2; at 13:05, it rose sharply to 844.8 W/m^2, corresponding to a maximum positive gradient of 631 W/m^2 per minute, and at 13:06, it slightly decreased to 843.2 $W/^2$.

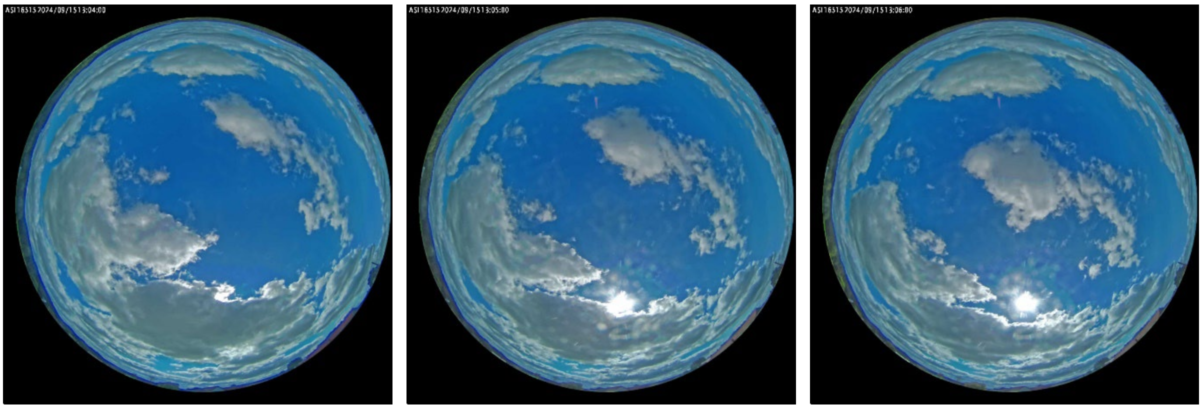

Figure 28: All-Sky Images captured at one-minute intervals from 13:04 to 13:06.

4. Estimation of Cloud Motion Vectors

Within the project, advanced methods for estimating Cloud Motion Vectors (CMV) from ground-based all-sky imagery were systematically investigated and evaluated. Since cloud movement is the primary driver of short-term irradiance variability, reliable estimation of CMVs is a key component in forecasting pipelines.

The analysis focused on three representative approaches for CMV estimation. The classical optical flow methods Luca-Kanade and Farnebäck, as well as the deep learning-based method Recurrent All-Pairs Field Transforms (RAFT). These approaches reflect different trade-offs between computational efficiency, spatial coverage, and accuracy. Lucas-Kanade (LK), as a sparse method, computes motion vectors only at selected feature points such as cloud edges, resulting in low computational cost but limited spatial coverage. Farnebäck, a dense optical flow method, estimates motion for all pixels and is widely used in operational systems due to its robustness and balanced performance. RAFT, a deep learning-based approach, leverages neural network architectures to capture complex motion patterns and deformation, offering the potential for improved accuracy at the cost of higher computational requirements.

A comprehensive dataset of all-sky images was used for training and evaluation, covering multiple seasons and a wide range of cloud conditions. The dataset includes image sequences captures at short time intervals. The images introduce challenges such as varying illumination, and artefacts caused by sun occlusion.

To address these challenges, a dedicated preprocessing pipeline was implemented. This includes circular masking to remove invalid image regions, radiometric normalization to improve robustness against illumination variability, and additional techniques for noise reduction and cloud segmentation. These steps significantly enhance the stability and quality of motion estimation under real-world conditions.

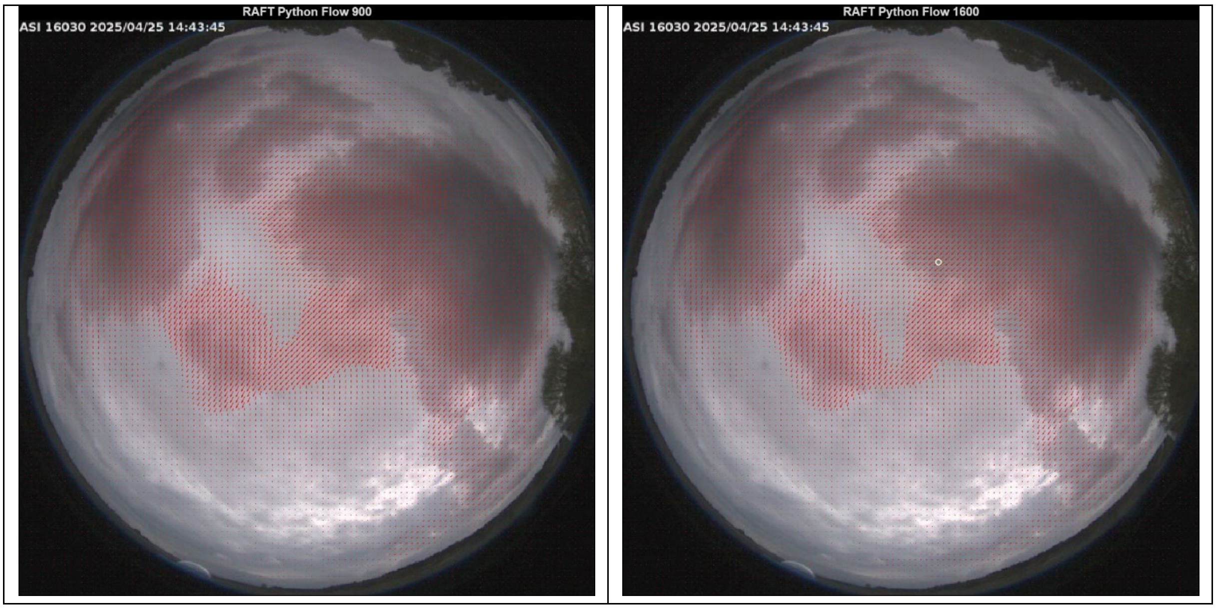

Figure 29: Fine-tuning real cloud motion prediction from 2025-04-25 (RAFT).

The comparative evaluation shows that no single method is universally optimal. Classical approaches such as Farnebäck remain strong baselines due to their efficiency and robustness, particularly for operational use. In contrast, deep learning-based methods such as RAFT demonstrate improved performance in capturing complex cloud dynamics and deformation, especially when adapted to domain-specific data. However, their effectiveness depends on the availability of suitable training data and sufficient computational resources.

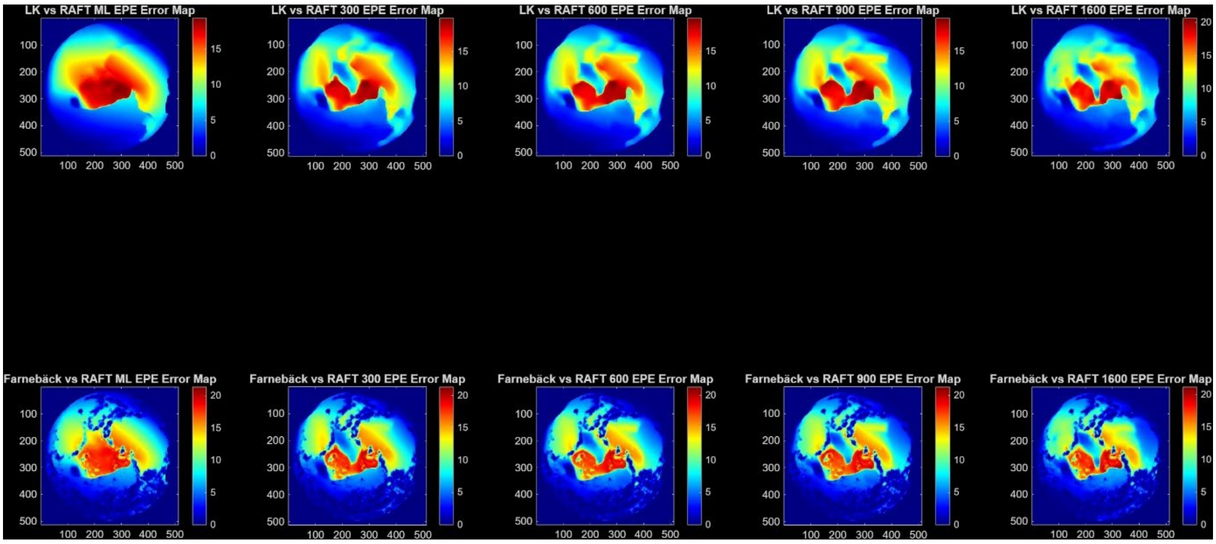

Figure 30: Difference maps cloud motion prediction from 2025-04-25.

4.1. RAFT Training Strategy and Domain Adaption

A key finding in the methodology was the success of a hybrid loss function for RAFT fine-tuning. This approach combined:

- **Supervised Loss:** Using Farnebäck estimates as a "pseudo-ground truth" in regions where the algorithms agreed.
- **Unsupervised Photometric Loss:** Using warping consistency to refine flow in texture less sky regions where classical methods fail.
- **Progressive training stages:** (330, 600, 900, 1600 pairs) revealed that RAFT-900 often provided the best balance of stability and adaption, though RAFT-1600 was uniquely robust in challenging low-light and dawn/dusk conditions.

4.2. Synthetic Flow Evaluation

To establish an accuracy baseline, the methods were tested against synthetic translations (shifts of 5, 10 and 15 pixels) where the ground truth was known.

- **Accuracy:** RAFT variants demonstrated overwhelming superiority, achieving sub-pixel Endpoint Error (EPE) (0.093-0.177 px) and near-zero outliers.
- **Classical Failure:** Both LK and Farnebäck failed significantly on these controlled shifts. LK exhibited 96-100 % outliers, and Farnebäck showed 56-73 % outliers, proving that classical methods struggle with the specific texture and geometric properties of ASI imagery when evaluated against synthetic ground truth.
- **Overfitting:** A minor finding indicated that the most advanced training stage (RAFT-1600) showed a slight EPE increase in texture less regions, suggesting mild overfitting to the set's specific atmospheric conditions.

4.3. Essential Findings

Quantitative and qualitative comparison on real consecutive sky images provides the most critical insights for operational deployment:

- **Coverage Superiority:** RAFT consistently achieved 100 % coverage, providing a coherent motion field across the entire sky. In contrast, LK often covered less than 3-10 %, and Farnebäck coverage fluctuated between 53-90 % depending on cloud density. Dense Refinement: RAFT was found to behave as a "dense, smoothed refinement" of Farnebäck. It generally matched Farnebäck's large-scale motion patterns but selectively departed from them in regions where Farnebäck's polynomial model was violated by brightness-constancy issues or complex deformation.
- **Challenging Conditions:** In "overcast/low-light" scenes, fine-tuned RAFT robustly recovered faint clout motion that was invisible to RAFT or classical baselines, halving the angular error compared to earlier stages.
- **Cirrostratus Detection:** RAFT excelled at predicting motion for very thin cirrostratus clouds, where both Farnebäck and Lucas-Kanade failed to detect any confident motion.
- **Disagreement Patterns:** Disagreement between RAFT and LK was highest in cloud interiors (where LK is sparse), while disagreement with Farnebäck was concentrated at complex cloud edges and uniform low texture regions.

5. SkyFusion: A Physics-Informed GAN-Transformer Framework for Intra-Hour Sky Forecasting and GHI Prediction

5.1 Objective

The developed SkyFusion, a multi-stage framework for 15-minute-ahead sky forecasting and solar irradiance nowcasting. The objective was to predict future cloud evolution from recent all-sky images and then estimate PIRA(t+15) by combining visual forecasts with physical sensor information. Both tasks use 30-second resolution, with 12 input steps and 30 output steps, while the main evaluation focuses on the final t+15 horizon. Image models operate on 128×128 RGB + SunMask sequences, and irradiance models operate on recent physics features such as irradiance, clearness index, variability, and solar geometry.

5.2 Models and Architectures

5.2.1. Pixel-Space Forecasting Path

The main production branch begins with Stage-0: ConvLSTM-UNet (Long Short-Term Memory, LSTM), which serves as the deterministic baseline. Its architecture combines an encoder-decoder with skip connections and a ConvLSTM bottleneck to model spatiotemporal dynamics. The model predicts the future sequence autoregressively, which gives a stable motion baseline and preserves large-scale brightness structure, but it also leads to smoothing and drift at long horizons. This blur at t+15 became the main limitation of the baseline and motivated the following refinement stages.

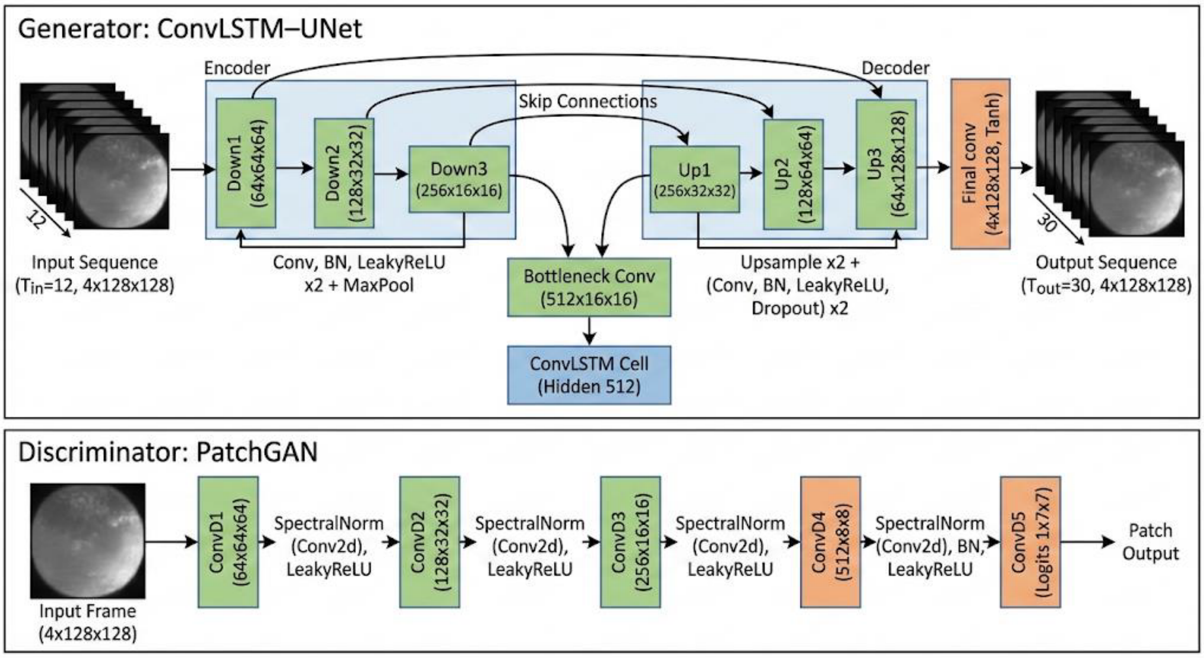

Figure 31: ConvLSTM–UNet + PatchGAN Baseline Architecture.

To reduce the blur of the deterministic baseline, Stage-1 adds a PatchGAN discriminator to the ConvLSTM-UNet generator. A stochastic noise vector is injected at the bottleneck, and the discriminator evaluates local patches to encourage sharper cloud boundaries and more realistic texture. This stage keeps the temporal backbone of Stage-0 but improves visual quality, especially at longer horizons.

Stage-2 extends this approach by introducing physics conditioning. Physical variables, including PIRA history, clear-sky index, solar angles, LIDAR backscatter-derived cloud features, cloud-base height, humidity, and temperature, are encoded with a Gated Recurrent Units (GRU) and injected into the generator through FiLM-style modulation. This improves physical consistency by aligning cloud evolution with irradiance behaviour and atmospheric structure, although it also reduces some sharpness and colour richness.

To restore the texture lost in Stage-2, Stage-3 applies diffusion refinement only at t+15, the most important operational horizon. The diffusion model uses the frozen Stage-2 prediction as a prior and physics embeddings as guidance. In this way, Stage-2 provides the physically consistent structure, while Stage-3 restores finer detail and more natural appearance.

5.2.1.1. Latent-Space Forecasting Path

Alongside the main production branch, the thesis also explores a latent-space approach based on discrete token modeling. A VQ-VAE compresses each sky image into a 32×32 token grid using a 256-entry codebook, transforming the problem from direct pixel prediction into structured token prediction.

On this basis, the FlatSky Transformer predicts the future token grid at t+15 from past token sequences. It uses 12 Transformer blocks, 10 attention heads, and d_model = 640, with positional embeddings over time and space. To address the imbalance between frequent clear-sky tokens and rarer cloud-edge tokens, training uses class-weighted cross-entropy.

After token prediction, a Multi-Scale PhyCell + Spatial Refiner improves motion consistency. The Transformer, VQ-VAE, and frozen RAFT optical flow remain fixed, while the latent physics module and refiner are trained before decoding the latent representation back to RGB. This branch shows the potential of latent discrete forecasting, but remains a proof of concept rather than the main production solution.

5.2.1.2. Irradiance Forecasting and Fusion Models

For PIRA (t+15) prediction, the thesis evaluates three time-series backbones: UNet1D, BiLSTM, and Transformer encoder. These models use recent physical measurements to predict future irradiance as either point estimates or probabilistic quantiles.

Building on these baselines, the final Stage-4 fusion combines the generated t+15 sky image with the physics sequence. Three variants are proposed: Stage-4a deterministic fusion, Stage-4b probabilistic quantile fusion, and Stage-4c hybrid fusion with frozen expert priors from a pre-trained physics LSTM. This final stage moves from standard regression to uncertainty-aware and hybrid prediction.

5.3 Experimental Results

5.3.1. Sky Image Forecasting Results

The first comparison was between Stage-0 ConvLSTM–UNet and Stage-1 GAN refinement. The adversarial stage clearly improved long-horizon forecast quality. Validation error decreased from L1 = 0.1306 and MSE = 0.0443 in Stage-0 to L1 = 0.0608 and MSE = 0.0136 in Stage-1, confirming that PatchGAN training significantly reduced blur and improved local cloud texture.

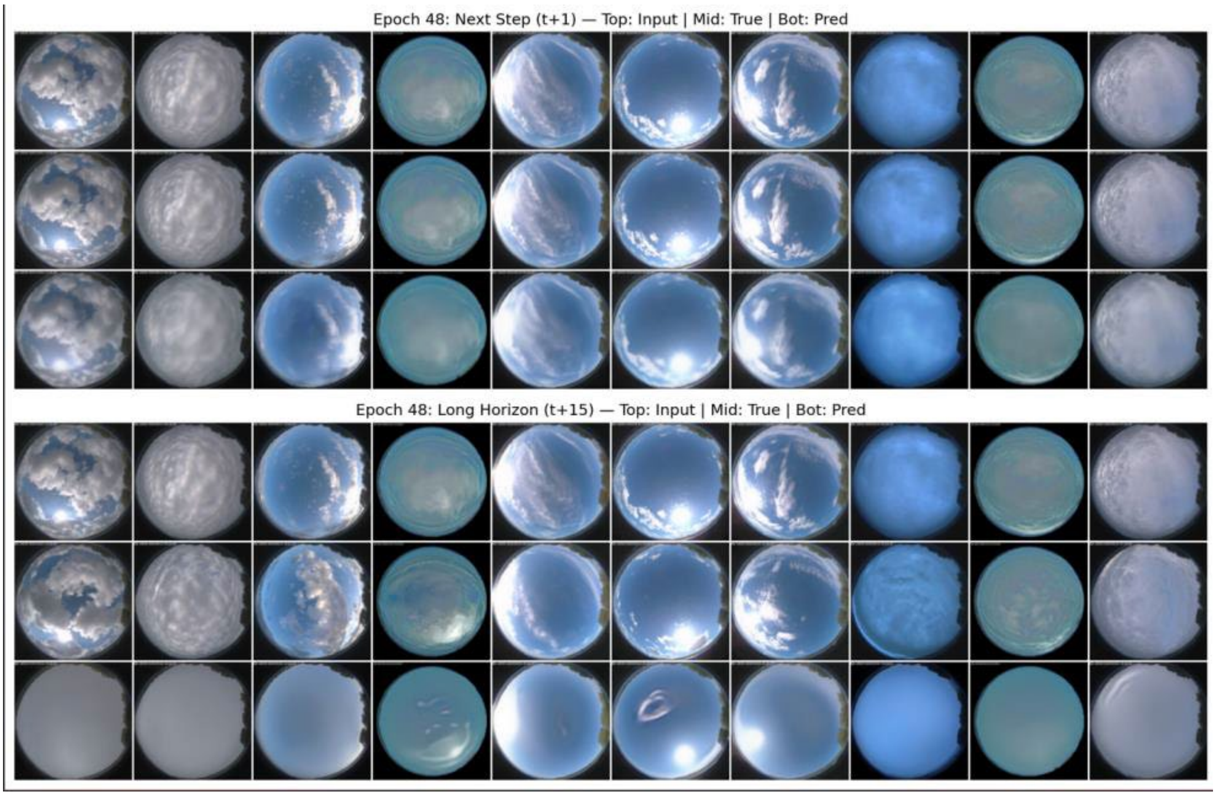

Figure 32: Stage-1 Results

The next improvement came from Stage-2 physics conditioning**.** Compared with the texture-focused Stage-1.6 GAN+VGG model, the physics-conditioned model achieved better pixel-level accuracy, with validation L1 = 0.0494, t+15 L1 = 0.0413, t+15 SSIM = 0.7331, and t+15 PSNR = 23.29 dB. The downside was reduced perceptual sharpness, because the stronger reconstruction emphasis produced more smoothing and lower colour saturation.

To recover that lost texture, Stage-3 diffusion refinement was introduced. The final GAN + Physics + Diffusion model achieved validation/test L1 = 0.1076 / 0.0970**,** MSE = 0.0317 / 0.0255, SSIM = 0.5961 / 0.6065, and PSNR = 19.77 / 20.13 dB. Qualitatively, diffusion restored cloud texture, improved sky colour, and produced more natural-looking outputs while preserving the Stage-2 prior.

Figure 33: Final Result GAN + Physics + Diffusion

For the latent branch, the VQ-VAE achieved SSIM around 0.83–0.85 and 61% codebook utilization (157/256 codes), showing that discrete tokenization preserved atmospheric structure well. The Transformer branch converged to a working set of about 51 codes, which demonstrated proof of concept but also indicated that the dataset was still too limited to fully exploit the full vocabulary.

5.3.1.1. Solar Irradiance Forecasting Results

For deterministic irradiance forecasting on the full dataset, all deep learning models outperformed the Smart Persistence Model (SPM). The best result was obtained by LSTM_Point, with RMSE = 135.47 W/m², MAE = 80.70 W/m², and forecast skill = 26.1%, compared with 183.24 W/m² RMSE and 115.10 W/m² MAE for the SPM baseline.

For probabilistic forecasting, the best baseline was LSTM_Proba, which achieved CRPS = 51.81, median RMSE = 137.47 W/m², median MAE = 76.17 W/m², and forecast skill = 77.5%. These results show that probabilistic forecasting gives a stronger operational basis because it captures uncertainty instead of collapsing all future behaviour into one scalar prediction.

5.3.1.2. Fusion Results

The final step was Stage-4 fusion, where the generated t+15 sky image was fused with the physics sequence. Stage-4a achieved 151.80 / 93.56 RMSE/MAE, Stage-4b improved this to 140.91 / 91.09, and the best result came from Stage-4c hybrid fusion, with 139.10 / 90.72 RMSE/MAE. These results confirm that the strongest performance does not come from image forecasting alone or physics alone, but from a hybrid model that combines visual evidence, temporal physics, and frozen expert priors.

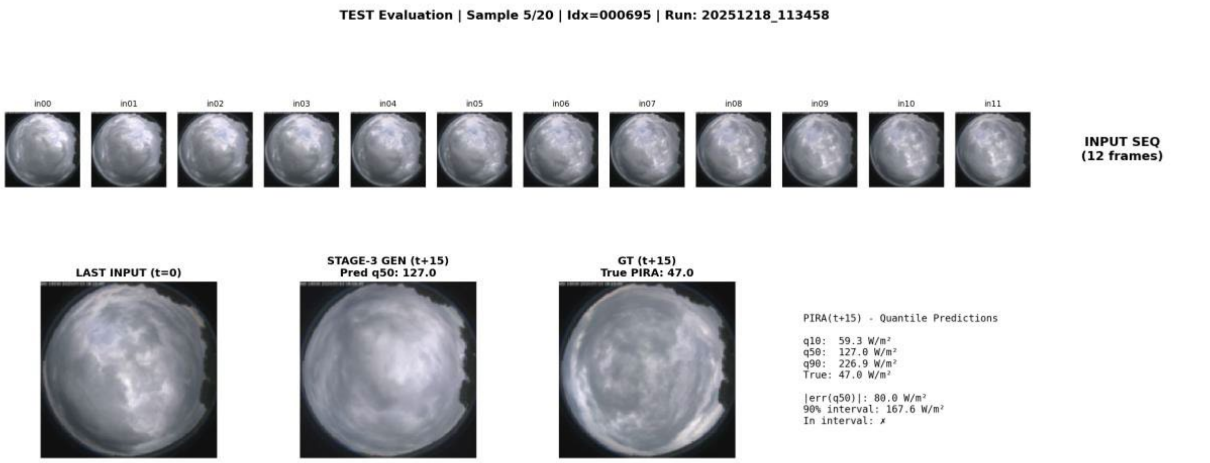

Figure 34: Pre-Trained LSTM Model Samples on Validation Set.

5.4 Conclusion

Overall, the thesis presents a progressive hybrid forecasting framework in which Stage-0 establishes the motion baseline, Stage-1 improves texture realism, Stage-2 adds physical consistency, and Stage-3 restores fine detail at t+15. In parallel, the VQ-VAE → Transformer → PhyCell branch demonstrates the feasibility of latent discrete forecasting. Finally, Stage-4c delivers the best performance within the fusion family, showing that generated sky imagery becomes most effective when combined with physics-based priors in a hybrid uncertainty-aware model. All-Sky Image Foundation Model

6. A Self-Supervised Vision Foundation Model for ASI

6.1 Introduction and Motivation

Ground-based all-sky fisheye cameras capture high-resolution hemispherical cloud observations at sub-minute cadences, enabling real-time cloud monitoring and intra-hour solar irradiance forecasting. However, supervised deep learning on this imagery is constrained by three persistent challenges: scarce and expensive manual annotations, pronounced domain shift across sites and sensors, and overfitting on small site-specific datasets.

Self-Supervised Learning (SSL) sidesteps these constraints by leveraging the large volumes of unlabelled imagery that sky cameras naturally produce. While recent domain-specific efforts have applied contrastive or reconstruction-based SSL to cloud segmentation and classification, no prior work has combined both objectives into a unified pretraining framework for fisheye sky imagery.

This thesis investigates whether a hybrid self-supervised objective — coupling masked autoencoder (MAE) reconstruction with DINO-style self-distillation on a shared Vision Transformer backbone — can produce representations that are both sensitive to fine-grained cloud morphology and invariant to illumination changes, yielding measurable gains on cloud classification and clear-sky index forecasting.

6.2 Data and Problem Formulation

Dataset: After temporal thinning to a 60-second cadence, daytime filtering, and quality control (removing corrupted frames, dome obstructions, and sensor artefacts), the final self-supervised pretraining corpus comprises approximately 143,000 all-sky images pooled from four fisheye cameras, split 90/10 for training and validation.

Downstream Task 1 Cloud Classification: A balanced dataset of 1,400 manually annotated images (350 per class) is split into 1,000/200/200 for train/validation/test. Four meteorologically motivated classes are defined: *clear sky* (≈0–10% coverage), *partial clouds* (≈10–60%), *heavy clouds* (≈60–95%), and *full overcast* (≈95–100%). The task evaluates whether the pretrained encoder captures discriminative cloud morphology under fisheye geometry.

Downstream Task 2 CSI Forecasting: A set of 1,661 sequences (957/207/497 train/val/test) is constructed from 31-frame windows at 1-minute resolution, paired with three future Clear-Sky Index (CSI) targets at horizons of 5, 15, and 30 minutes. CSI normalizes measured irradiance by its clear-sky reference, isolating cloud-driven attenuation from the deterministic solar cycle. Each sequence combines frozen visual embeddings with an 8-dimensional sensor feature vector (irradiance, temperature, humidity, solar geometry, clear-sky reference). The task evaluates whether pretrained representations support continuous, multi-horizon prediction of solar variability beyond persistence assumptions.

6.3 Hybrid MAE-DINO Architecture

The central methodological contribution is a hybrid self-supervised framework that trains a single Vision Transformer (ViT-Base/16, ≈86M parameters) to satisfy two complementary objectives simultaneously: masked image reconstruction (MAE) and self-distillation (DINO). The key intuition is that reconstruction alone encourages sensitivity to local textures and radiometric detail but does not enforce invariance to appearance transformations, while self-distillation promotes view-invariant semantic clustering but may underemphasize fine-grained spatial fidelity. By updating one shared encoder with gradients from both losses, the hybrid model learns representations that are both re-constructible and stable under multi-view perturbations.

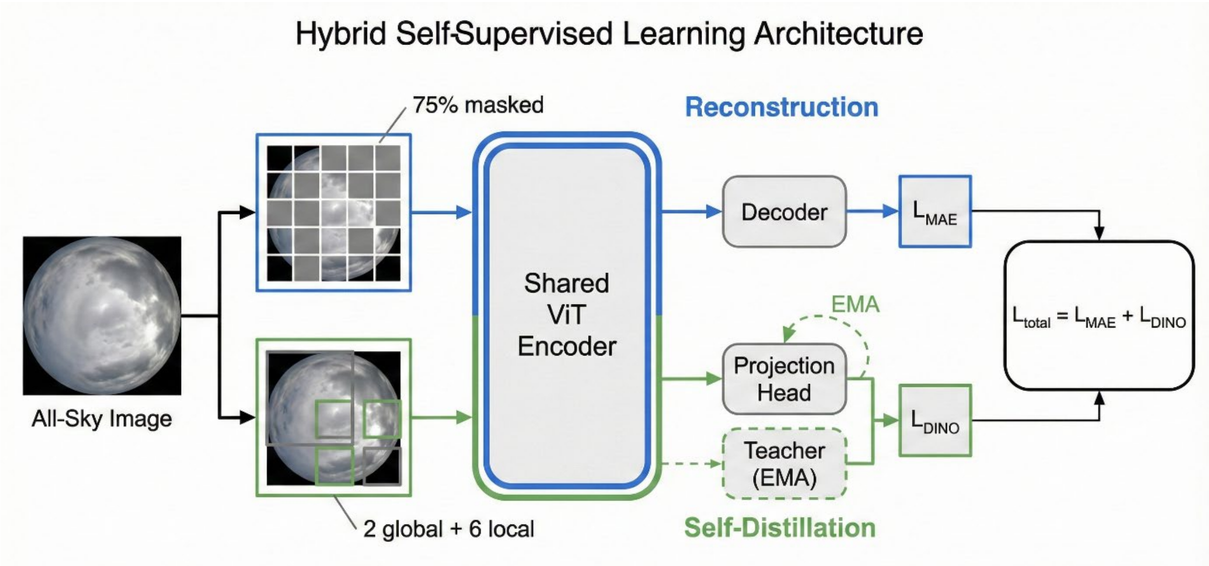

Figure 35: Shared encoder with MAE decoder and DINO projection head

Shared student backbone with dual heads: The ViT-Base encoder (12 layers, 768-dim, 12 heads) serves as a shared feature extractor. Its output feeds two decoupled, lightweight heads: (i) an asymmetric MAE decoder (8 layers, 512-dim, 16 heads) that reconstructs masked patches, anchoring learning in pixel-level cloud structure; and (ii) a DINO projection head (compact MLP) that maps the [CLS] token into a prototype space for self-distillation. The two heads have distinct roles and do not interact — the MAE decoder never participates in distillation, and the projection head never participates in reconstruction — so each objective expresses its inductive bias through a dedicated pathway while the encoder learns shared invariants.

Fisheye-aware patch validity masking: Because fisheye images embed a circular sky region within a rectangular frame, a binary patch validity mask is applied to the reconstruction loss. Only patches whose centres fall inside the precomputed sky circle contribute to the MAE gradient. The reconstruction loss is computed as a mean squared error over patches that are both masked (per the standard 75% random masking) and valid (inside the sky circle), with per-patch normalized targets. This prevents the black border from dominating optimization while leaving the masking strategy unchanged.

Joint objective: The DINO loss aligns student and teacher probability distributions across all cross-view pairs via cross-entropy over learned prototypes. The total training objective combines both terms with equal weighting: the fisheye-aware reconstruction loss plus the DINO self-distillation loss.

6.4 Experimental Setup

Encoder baselines: Four ViT-Base/16 initializations are compared while holding the downstream architecture, data splits, and optimization fixed: (i) random initialization (no pretraining), (ii) supervised ImageNet-1K pretrained weights, (iii) MAE-only self-supervised pretraining on the all-sky corpus, and (iv) the proposed hybrid MAE+DINO pretraining. This spectrum isolates the effect of representation quality from task-head capacity.

Transfer protocols: Two protocols probe different aspects of representation quality. For cloud classification, a partial fine-tuning strategy unfreezes only the last six transformer blocks and normalization layers. For CSI forecasting, a strict frozen-encoder protocol is used: [CLS] embeddings are extracted once and cached, and only a lightweight multimodal regressor (temporal attention pooling over vision embeddings, gated fusion with sensor features, horizon-specific regression heads) is trained on top. This ensures that forecasting improvements are attributable solely to representation quality.

6.5 Results

6.5.1. Cloud Classification

Table 4: Test-set performance across encoder initializations.

Encoder initialization	Accuracy (%)	Macro-F1
Random init	79.5	0.800
ImageNet	91.5	0.914
MAE-only	88.5	0.885
Hybrid	95.5	0.955

A clear ordering emerges. The hybrid encoder achieves 95.5% accuracy and a macro-F1 of 0.955, improving over ImageNet by +4.0 pp in accuracy and +0.041 in F1, and over MAE-only by +7.0 pp and +0.070 respectively. The random baseline confirms that pretraining is essential on this small, labelled set.

The most informative gains appear at the class level. Clear sky and full overcast are reliably classified by all pretrained models, but partial clouds — the most ambiguous regime — shows the sharpest improvement: F1 rises from 0.701 (random) to 0.851 (ImageNet) to 0.848 (MAE-only) and then to 0.949 (hybrid). This suggests that reconstruction alone is insufficient to separate borderline semantic regimes, while adding self-distillation organizes the feature space around meaningful sky-state structure**.**

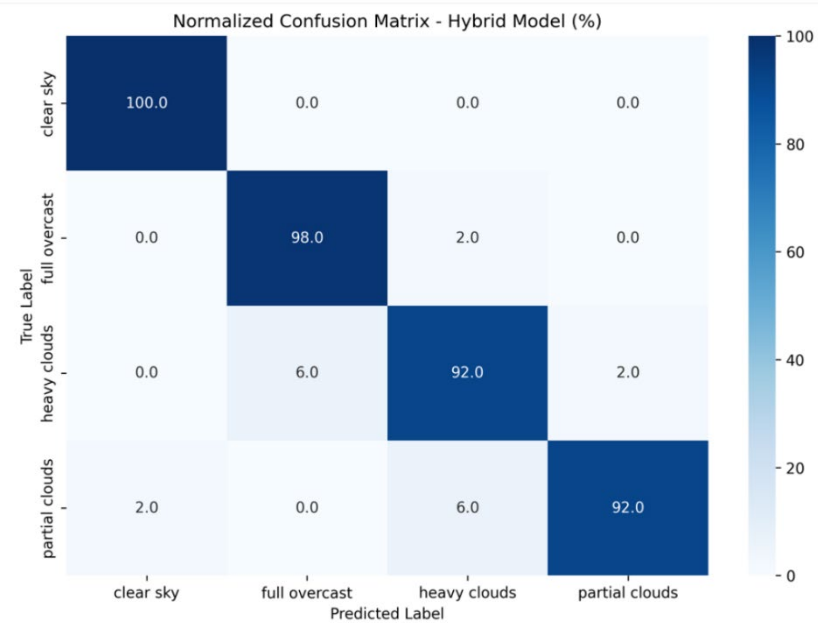

Figure 36: Hybrid model confusion matrix.

Confusion matrix analysis confirms that residual errors concentrate between visually adjacent regimes (partial ↔ heavy, heavy ↔ overcast) rather than between extremes, indicating that the hybrid encoder captures the coarse ordering of cloud coverage and that remaining failures reflect genuinely gradual transitions in sky state.

6.5.2. CSI Forecasting

Table 5: Test-set performance by forecast horizon.

Metric	Horizon	Random	ImageNet	MAE-only	Hybrid
MAE	5 min	0.114	0.113	0.101	**0.100**
	15 min	0.123	0.127	0.116	**0.105**
	30 min	0.155	0.145	0.150	**0.136**
R^2	5 min	0.786	0.799	0.811	**0.819**
	15 min	0.763	0.756	0.768	**0.805**
	30 min	0.651	0.687	0.667	**0.706**

The hybrid encoder achieves the strongest aggregate performance across all horizons, with a mean R^2 of 0.777 and mean MAE of 0.114 — improving over ImageNet (+0.030 R^2, −0.014 MAE) and MAE-only (+0.028 R^2, −0.008 MAE). Skill scores remain positive across all backbones, confirming consistent improvement over smart persistence.

The most striking result occurs at the 15-minute horizon, where the hybrid model reaches R^2 = 0.805 and correlation = 0.908. This horizon sits at a sweet spot for all-sky imagery: long enough that short-lived pixel fluctuations average out, yet short enough that cloud evolution remains partly inferable from the observed sky dome — precisely the regime where representations encoding both localized cloud structure and global sky organization are most valuable.

Performance degrades for all models at 30 minutes (hybrid R^2 drops from 0.805 to 0.706), revealing a fundamental predictability constraint: cloud advection and formation processes increasingly depend on atmospheric dynamics outside the camera's field of view. This degradation is structural rather than an optimization artefact, and representation quality alone cannot overcome it.

6.6 Conclusion

This work provides evidence that combining MAE reconstruction with DINO-style self-distillation forms a robust, domain-adapted pretraining signal for all-sky fisheye imagery. Three key findings emerge:

1. Hybrid pretraining consistently outperforms both ImageNet initialization and reconstruction-only self-supervision across classification (+4.0 pp accuracy over ImageNet, +7.0 pp over MAE-only) and forecasting (+0.030 mean R^2 over ImageNet), demonstrating that neither generic supervised features nor pixel-level reconstruction alone capture the full representational needs of this domain.
2. The largest gains concentrate where ambiguity is highest — the partial-clouds class in classification and the 15-minute horizon in forecasting — suggesting that the hybrid objective specifically improves the encoder's ability to resolve borderline atmospheric states where both local detail and global context matter.
3. Representation quality does not remove fundamental predictability limits. The consistent performance degradation at 30 minutes across all models points to an observability ceiling imposed by the camera's limited field of view and unobserved cloud advection, motivating future work on temporal pretraining, multi-camera fusion, and integration with broader atmospheric context.

7. Dissemination and Education

7.1. All-Sky Imager Workshop 2025

On 25 and 26 February 2025, the first international ASI workshop took place at the Start-Up and Digital Innovation Centre Stellwerk18[1] in Rosenheim (Germany). Initiated within the framework of the project, the workshop attracted considerable interest, with 79 registrations from participants in Germany and abroad. Around 25 participants attended on-site, while additional attendees joined remotely.

Over the course of one and a half days, the workshop featured a diverse program covering a broad range of topics, including hardware systems, low-cost solutions, and available datasets to applications of AI in All-Sky Imaging. Further sessions addressed solar radiation measurement and forecasting, current research gaps and open challenges, representations of ongoing research projects alongside innovative industry solutions.

A key objective of the event was to strengthen collaboration within the research community and to promote exchange between academia and industry. The workshop provided a valuable platform for networking and knowledge sharing across the domains.

Day 1 Program – Tuesday, February 25th 2025

START OF WORKSHOP – DAY I
OPENING SESSION

09:00 – 09:15	**Welcoming Address** *Prof. Dr.-Ing. Grit Behrens, Alexander Kruse* *Prof. Mike Zehner, Andreas Boschert* Bielefeld University of Applied Sciences Technical University of Applied Sciences Rosenheim
09:15 – 09:45	**Best Practice Handbook for the Collection and Use of Solar Resource Data for Solar Energy Applications: Fourth Edition (Online)** *Jan Remund* Meteotest AG – Head of Energy & Climate International Energy Agency: Photovoltaic Power Systems Programm Task 16: Solar Resource for High Penetration and Large Scale Applications
09:45 – 10:00	Networking Break

ALL-SKY IMAGER IN FOCUS: RELIABLE SYSTEMS AND AFFORDABLE HARDWARE
CHAIR: DR. THOMAS SCHMIDT

10:00 – 10:20	**On the design of Eye2Sky camera network subsets for the optimal nowcast of PV production** *André Schaper* German Aerospace Center – Institute of Network Energy Systems
10:20 – 10:40	**Quality control methods and their evaluation in nowcasting of All-Sky Imager network (Eye2Sky)** *Arne Goerlitz* German Aerospace Center – Institute of Network Energy Systems Carl-von-Ossietzky University of Oldenburg
10:40 – 11:00	**Geometric Calibration of All-Sky Cameras Using Sun and Moon Positions: A Comprehensive Analysis (Online)** *Yann Fabel* German Aerospace Center – Solar Research Institute
11:00 – 11:20	**The potential of low-cost cameras for citizen science and data assimilation (Online)** *Dr. Leonhard Scheck* Hans-Ertel-Zentrum für Wetterforschung / LMU München
11:20 – 11:40	**Autonomous All-Sky Cameras for Optimized Control of Photovoltaic Battery Storage Systems** *Paul Matteschk* Wematics FlexCo. / Bonn-Rhein-Sieg University of Applied Sciences
11:40 – 12:40	Lunch Break

2

Day 1 Program – Tuesday, February 25th 2025

PUBLICLY AVAILABLE ALL-SKY IMAGER-DATASETS
CHAIR: ANDREAS BOSCHERT

12:40 – 13:00	**The Eye2Sky all-sky imager network and its data availability** *Dr. Thomas Schmidt* German Aerospace Center – Institute of Network Energy Systems
13:00 – 13:20	**SkyImageNet: A large scale sky image database for solar energy meteorology (Online)** *Dr. Quentin Paletta* European Space Agency
13:20 – 13:30	Networking Break

AI-APPLICATIONS ON SKY IMAGES
CHAIR: PAUL MATTESCHK

13:30 – 13:50	**Synthetic sky image datasets for training AI models** *Max Aragon* Mines Paris PSL
13:50 – 14:10	**A new fast method for simulating all-sky images – Enabling a virtual benchmark suite for nowcasting and other applications?** *Philipp Gregor* Meteorologisches Institut, Ludwig-Maximilians-Universität München
14:10 – 14:30	**Cloud Classification and Solar Irradiance Forecasting Using All-Sky Images** *Yassine Ribouh and Naoufal EL Atmioui* Bielefeld University of Applied Sciences
14:30 – 14:50	**Estimating Global Horizontal Irradiance and Cloud Base Height from All-Sky Images Using a Deep-Learning Approach** *Andreas Boschert* Technical University of Applied Sciences Rosenheim *Alexander Kruse* Bielefeld University of Applied Sciences
14:50 – 15:30	Networking Break

Note: The event will take place in a hybrid format. All presentations will be streamed online, even those held on-site. All timetables in this brochure are in Central European Time (CET).

3

Figure 37: Brochure and conference programme for the 1st All-Sky Imager Workshop 2025.

7.2. All-Sky Imager Workshop 2026 and Long-Term Strategy

Following the strong interest and positive feedback from first ASI workshop, it's continuation is planned for 2026. The second ASI workshop will take place on 5-7 October at the DLR site in Oldenburg, scheduled immediately prior to a meeting of IEA PVPS Task 16.

The workshop series will continue under the patronage of TUAS Rosenheim and UAS Bielefeld, while being hosted at rotating locations in the future. This approach aims to broaden participation and strengthen international engagement within the ASI community.

---

[1] https://www.stellwerk18.de/

The long-term objective is to further consolidate the international network, foster sustained collaboration, and provide recurring platform for the exchange of research results, technological developments, and practical experience in the field of all-sky imaging.

7.3. Publications

As part of the project, particular emphasis was placed on disseminating of the achieved results. A central element of these efforts was the international ASI workshop. In addition, key project outcomes were presented at national and international conferences and workshops, including the PV Symposium and EnviroInfo.

To facilitate knowledge transfer into practice, close collaboration with industry partners was maintained throughout the project. Project-related content was also integrated into university teaching and further explored through project work and academic theses.

The table below provides an overview of the dissemination activities carried out during the project period.

Table 6: Overview of Publications and Presentations at Scientific Conferences.

Conference Seminar	Target Group	Location	Type of Publication	Date	Brief Description
Workshop Minute Scale Forecasting for the Weather Driven Energy System	Science Expert Audience	DTU Risø Campus, Denmark	Presentation	9-12/04/2024	Project Presentation and Exchange on »HELIOS«
41st European Photovoltaic Solar Energy Conference and Exhibition (EU PVSEC)	Science Expert Audience	Vienna, Austria	Scientific Poster	23-27/09/2024	Enhancing Photovoltaic Power Forecasting Accuracy with a Network of All-Sky Imagers and Meteorological Sensors Deployed in a PV Test Field
38th International Conference for Environmental Informatics (EnviroInfo)	Science Expert Audience	Cairo, Egypt	Presentation Journal Lecture Notes of Informatics DOI: 10.18420/en2024_01	12-14/11/2024	Helios: Introducing a Project fore Forecasting Photovoltaic Yields using a Small Network of All-Sky Imagers and Meteorological Sensors
40. PV-Symposium 2025	Science Expert Audience	Bad Staffelstein, Germany	Presentation	11-13/03/2025	Wolken, Daten, Prognosen: Was der 1. Internationale All-Sky Imager

					Workshop 2025 gezeigt hat.
38th International Conference for Environmental Informatics (EnviroInfo)	Science Expert Audience	Potsdam, Germany	Presentation Journal Lecture Notes of Informatics DOI: 10.18420/env2025_1	17-19/ 09/2025	SkyFusion: A Physics-Informed GAN-Transformer Framework for Intra-Hour Sky Forecasting and PV Prediction
			Presentation Journal Lecture Notes of Informatics DOI: 10.18420/env2025_1		A Self-Supervised Vision Foundation Model for All-Sky Imagery

7.4. Presentations

Seminar	**Target Group**	**Location**	**Speaker**	**Date**
Data Science im Spannungsfeld zwischen Prozessoptimierung und Nachhaltigkeit.	Science Expert Audience	Bielefeld	Alexander Kruse	19/09/2024
Scientific Seminar Rosenheimer Technologiezentrum Energie und Gebäude (RoTEG)	Science Expert Audience	Rosenheim	Andreas Boschert	03/06/2025

7.5. Projects Work and Academic Theses

- Modelling and simulation of a ground-mounted PV system using the PV_LIB toolbox in Matlab, Master-Project by Jonathan Huth.
- Cloud Detection using Machine Learning-Techniques, Master-Project by Kumar-Ashish, Konjeti,-Sharath Chandra, Lomada-Dharma Teja Reddy, Mariserla-Varun Kumar, Mohammed-Shaad Sarwar.
- Deep Learning Approaches to Cloud Classification for Solar Energy Optimization, Master-Project by Yassine Ribouh, Naoufal El Atmioui and Alexander Kruse.
- SkyFusion: A Physics-Informed GAN-Transformer Framework for Intra-Hour Sky Forecasting and GHI Prediction, Master-Thesis by Naoufal El Atmioui.
- A Self-Supervised Vision Foundation Model for All-Sky Imagery, Master-Thesis by Yassine Ribouh.
- Estimation of Cloud Motion Vectors from All-Sky Imagers: A comparison of Optical Flow and Machine Learning Approaches for Solar Nowcasting, Master-Thesis by Jonathan Huth.

8. Conclusion and Outlook

8.1. Summary

The »HELIOS« project successfully developed, deployed, and validated a comprehensive AI-driven framework for high-resolution short-term solar irradiance forecasting, making a significant contribution to the reliable integration of photovoltaic energy into decentralized power grids. Over the course of the project, substantial progress was achieved across all core areas, spanning measurement infrastructure, data management, algorithm development, and scientific dissemination.

A state-of-the-art measurement infrastructure was established at Maierhof ground-mounted PV installation in Buttenwiesen, Bavaria. Three All-Sky Imagers, a LIDAR ceilometer, a rotating shadow band pyranometer, and a range of high-precision meteorological sensors was successfully installed, calibrated, and commissioned. The geometric calibration of the ASIs using OCamCalib toolbox (MATLAB) ensured accurate fisheye lens modelling and reliable projection of hemispherical sky images onto a well-defined coordinate system. The PV system's generation data were mace accessible via a standardized API interface. Throughout the measurement campaign, the infrastructure demonstrated robust performance under real-world field conditions, even following initial hardware challenges such as a defective laser module in the ceilometer and communication issues with the rotating shadow band system, both of which were successfully resolved.

To ensure systematic and reproducible data quality, the HDMS was developed as a dedicate pipeline for automated image validation, EXIF metadata extraction, sensor plausibility checks, and structured long-term archiving. The HDMS provided transparent reporting and actionable maintenance alerts, forming an essential backbone of the reliable operation of the measurement network and for the integrity of all downstream analyses.

On the algorithmic side, the project delivered a series of significant advances. For cloud motion estimation, a thorough comparative study of classical optical flow methods and deep learning-based approaches was conducted. RAFT-based models, fine-tuned using a hybrid supervised and unsupervised loss function, demonstrated overwhelming superiority over classical methods such as Lucas-Kanade or Farnebäck, achieving sub-pixel endpoint errors in controlled synthetic evaluations while remaining robust under challenging real-world conditions including low-light and dawn/dusk scenarios. For cloud detection, a hybrid MAE-DINO self-supervised model was developed, combining masked autoencoding for feature learning with DINO-based representation alignment. This approach achieved strong classification performance across multiple cloud types and forecast horizons, outperforming conventional baselines consistently. For sky image forecasting and GHI prediction, the SkyFusion framework was developed, combining a physics-informed GAN with a Transformer-based latent prediction branch and a diffusion-based refinement stage. The final GAN + Physics + Diffusion model achieved a test L1 of 0.0970, MSE of 0.0255, SSIM of 0.6065, and PSNR of 20.12 db, while qualitatively producing natural-looking cloud textures and realistic sky colour transitions. For deterministic irradiance forecasting, LSTM-based point forecasting achieved the best overall performance with an RMSE of 135.47 $W7M^2$, an MAE of 80.70 W/m^2, and a forecast skill of 26.1 % relative to the Smart Persistence Model baseline.

The analysis of solar irradiance data revealed important physical insights relevant to forecasting. Pronounced irradiance enhancement effects were observed under broken convective cloud fields, particularly during midday hours. The largest irradiance gradients were associated with rapid changes in cloud cover during the midday period and with high cloud velocities. These findings underline the importance of high temporal and spatial resolution observations for capturing the full variability of solar resources.

The project's results were actively disseminated through a broad range of scientific and public engagement activities. Contributions were presented at leading international venues including EU PVSEC in Vienna and EnviroInfo in Cairo, as well as at specialized workshops such as the minute-scale forecasting workshop at DTU Risø Campus in Denmark. The organization of the inaugural All-Sky Imager Workshop 2025 represented a particular milestone, bringing together an international community of researchers and industry practitioners and establishing a dedicated recurring forum for the ASI research field. Project outcomes were further embedded in university teaching and explored in depth through multiple master's theses and student project groups at both partner institutions, ensuring broad and sustained knowledge transfer.

In summary, »HELIOS« demonstrated that the combination of a well-designed field measurement infrastructure, rigorous data quality management, and state-of-the-art deep learning methods can substantially improve the accuracy and reliability of short-term solar irradiance forecasting beyond the current state of the art. The project thereby contributes directly to the ecological and economic goals of the energy transition, supporting more efficient use of volatile renewable generation, reducing dependence on conventional backup capacity, and ultimately advancing progress toward a carbon-neutral energy system.

8.2. Outlook

The results achieved within the HELIOS project provide a solid scientific and technical foundation for further research and application work. Several concrete and promising perspectives emerge for the immediate continuation of the project as well as the long-term development of the methods and infrastructure established.

Relocation of the measurement infrastructure to Rosenheim. The next step foresees the relocation of the measurement infrastructure currently operated at the Buttenwiesen site to the campus of the Technical University of Applied Sciences Rosenheim. This includes the three All-Sky Imagers, the LIDAR ceilometer, the rotating shadow band pyranometer, and the associated meteorological sensors. The new site offers the opportunity to operate the measurement technology in a different topographic and climatic environment, while directly applying the experience gained during installation, calibration, and operation throughout the HELIOS project. At the same time, the proximity to the university campus enables an even closer integration of measurement operations, data management, and algorithmic development, facilitating faster iteration cycles and more immediate scientific exploitation of the collected data.

Transfer and validation of algorithms to the new site. With the change of measurement location, the scientifically important question of the transferability of the developed algorithms to new geographic and atmospheric conditions arises. The models developed within HELIOS — including the hybrid MAE-DINO model for cloud segmentation, the SkyFusion framework for irradiance forecasting, and the RAFT-based approaches for cloud motion vector estimation — will be systematically evaluated at the new site and adapted or retrained where necessary. This transfer validation represents an important step toward robust, site-independent forecasting systems and will yield valuable insights into the generalization capabilities of the deployed deep learning architectures across varying environmental conditions.

Integration into a virtual power plant in collaboration with Stadtwerke Rosenheim. A central objective of the follow-up work is real-world testing of the developed short-term forecasts within an operational energy system. In collaboration with Stadtwerke Rosenheim, the integration of solar irradiance and PV generation forecasts into a virtual power plant will be investigated. Virtual power plants aggregate decentralized generation, storage, and consumption units, coordinating them to provide grid services and optimize the feed-in of renewable energy. The incorporation of high-resolution short-term forecasts into such a system offers the potential to increase planning reliability, reduce balancing energy costs, and maximize the self-consumption of locally generated solar energy. Collaboration with Stadtwerke Rosenheim as a regional energy provider enables practical testing under real operational conditions, bridging the gap between research prototype and operational deployment.

Assessment of forecast value and economic relevance. Beyond technical development, future work will systematically investigate and evaluate the concrete benefits of the generated forecasts. Both economic and operational aspects will be considered: To what extent can high-resolution short-term forecasts reduce operating costs, minimize grid interventions, or increase revenues on energy markets? What contribution do improved forecasts make to grid stability and the reduction of balancing energy requirements? And how does forecast quality affect the control of battery storage systems and load flow optimization in distribution networks? Answering these questions requires close collaboration with grid operators and energy suppliers, as well as the development of appropriate evaluation metrics that link technical forecast quality with economic added value. The results of this assessment will help to promote the adoption of AI-based solar forecasting in the energy industry and to transparently demonstrate its relevance for achieving climate protection goals.

9. Literaturverzeichnis

[BoZe25a] BOSCHERT, ANDREAS ; ZEHNER, MIKE: DBU_HELIOS - Predictive Spatial Analytics for Solar Energy Grid Integration: Enhancing Reliability and Efficiency.

[BoZe25b] BOSCHERT, ANDREAS ; ZEHNER, MIKE: Predictive Spatial Analytics for Solar Energy Grid Integration: Enhancing Reliability and Efficiency.

[DaMS06] DAVIDE, SCARAMUZZA ; MARTINELLI, AGOSTINO ; SIEGWART, ROLAND: A Toolbox for Easy Calibrating Omnidirectional Cameras. In: *Proceedings to IEEE International Conference on Intelligent Robots and Systems*. Beijing, 2006

[LaMS22] LAUF, THOMAS ; MEMMLER, MICHAEL ; SCHNEIDER, SVEN: Emissionsbilanz erneuerbarer Energieträger, UMWELTBUNDESAMT (Hrsg.).

[RuSS08] RUFLI, MARTIN ; SCARAMUZZA, DAVIDE ; SIEGWART, ROLAND: Automatic Detection of Checkerboards on Blurred and Distorted Images. In: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. Nice, 2008

[Scar08] SCARAMUZZA, DAVIDE: *Omnidirectional Vision: From Calibration to Robot Motino Estimation*. Zürich, ETH Zürich, 2008

[Scar25] SCARAMUZZA, DAVIDE: *OCamCalib: Omnidirectional Camera Calibration Toolbox for Matlab*. URL https://sites.google.com/site/scarabotix/ocamcalib-omnidirectional-camera-calibration-toolbox-for-matlab. - abgerufen am 2025-12-10

[ScMS06] SCARAMUZZA, DAVIDE ; MARTINELLI, AGOSTINO ; SIEGWART, ROLAND: A Flexible Technique for Accurate Omnidirectional Camera Calibration and Structure from Motion. In: *Proceedings of the Fourth IEEE International Conference on Computer Vision Systems (ICVS 2006)*. New York, 2006

[SSSB24] SCHMIDT, THOMAS ; STÜHRENBERG, JONAS ; SCHELLHORN, MAXIMILIAN ; VON BREMEN, LÜDER ; HAMMER, ANNETTE ; LEZACA, JORGE ; SCHROEDTER-HOMSCHEIDT, MARION: Vergleich der Einstrahlungsvariabilität aus Datenquellen mit unterschiedlicher raum-zeitlicher Auflösung.

[Vcom05] VCOM-API Client.

[WBOQ00] WENIGER, JOHANNES ; BÖHME, NICO ; ORTH, NICO ; QUASCHNING, VOLKER: Sind Solarstromspeicher Klimaschützer? Bd. pv magazine, Nr. September 2019