

# Enhanced All-Sky Imaging for Optical Satellite Link Availability: Multi-Site Analysis with Visible and Far-Infrared Cloud Assessment for nowcasting toolchain

Helmut Ribel<sup>a</sup>, Matias Lantschner<sup>b</sup>, Marcus Knopp<sup>a</sup>, Bijan Nouri<sup>c</sup>, Eduardo Saez Martinez<sup>c</sup>, Yann Fabel<sup>c</sup>, and Luis Zarzalecho<sup>d</sup>

<sup>a</sup>(DLR) Responsive Space Cluster Competence Center (RSC3), Trauen, Germany

<sup>b</sup>(DLR) Space Operations and Astronaut Training German Aerospace Center, Oberpfaffenhofen, Germany

<sup>c</sup>(DLR) Institute of Solar Research, Almería, Spain

<sup>d</sup>(CIEMAT) Energy Department – Renewable Energy Division, Almería, Spain

## ABSTRACT

Ground-based optical satellite communication requires accurate short-term knowledge of cloud conditions along the line-of-sight. In this work, we evaluate a nowcasting-based toolchain that uses all-sky imagery to segment clouds in real time and predict cloud evolution several minutes ahead. By combining historical cloud masks with satellite ephemerides, we compute the instantaneous and pass-integrated availability of optical links for a large set of satellite passes. This enables a quantitative assessment of how cloud coverage statistics translate into long-term link availability.

Using multi-year datasets from two geographically distinct sites—Trauen (Germany) and Almería (Spain)—we derive cloud coverage distributions, monthly and diurnal patterns, and long-term temporal trends. Visible all-sky images provided daytime cloud masks. Far-infrared dataset enabled nighttime cloud assessment, extending applicability beyond daylight hours. For each satellite pass, pass-level, monthly, and annual availability statistics were calculated and compared directly with cloud-coverage metrics.

Preliminary results reveal clear climatological differences: Trauen exhibits higher overall cloudiness ( $\approx 75, 87\%$  cloud cover) and pronounced seasonal variability, while Almería shows stable, predominantly clear conditions ( $\approx 34, 6\%$  cloud cover). Satellite-pass availability closely mirrors these patterns, with Trauen yielding lower average availability and stronger seasonal swings. Long-term annual averages of cloud coverage, link availability converge to within 1.5% Trauen’s highly variable periods. Seasonal and site-specific differences in convergence rates are quantified, providing practical guidance for mission scheduling and link-budget planning.

The toolchain reliably processes large historical archives, reproduces expected climatological features, and generates consistent availability statistics. The far-infrared extension demonstrates that nighttime cloud assessment is feasible with modest additional hardware, significantly broadening operational windows. Direct comparison of a Central-European and a Mediterranean site further highlights how regional cloud regimes govern both statistics and link performance.

**Keywords:** free-space optical communication, optical ground stations, all-sky imager, cloud segmentation, cloud nowcasting

## 1. INTRODUCTION

Optical ground-to-space communication has emerged as a pivotal technology for high-rate data transfer between satellites and terrestrial terminals. Compared to traditional radio-frequency links, optical communication offers significantly higher bandwidth, enhanced security, and reduced interference (Khalighi et al.<sup>1</sup>). However, the performance of these systems is fundamentally limited by atmospheric conditions, particularly the presence of

---

Further author information: (Send correspondence to Helmut Ribel)

Helmut Ribel: E-mail: helmut.ribel@dlr.de

Matias Lantschner: E-mail: Matias.Lantschner@dlr.de

clouds along the line of sight (Rotherham et al.<sup>2</sup>). Even thin or partially transparent cloud layers can disrupt or completely block an optical link; consequently, accurate knowledge of cloud dynamics is essential for effective link planning, scheduling, and long-term site assessment. As optical communication systems transition toward operational deployment, the need for reliable, short-term cloud prediction becomes increasingly critical.

All-sky imaging has become a valuable tool for monitoring local cloud conditions at optical ground stations. Modern systems provide continuous hemispherical sky coverage and enable the segmentation of cloud and clear-sky regions. When combined with nowcasting models that predict cloud motion over time, these systems can generate minute-scale forecasts of cloud evolution. Such nowcasts are particularly valuable for satellite optical links, where communication windows are typically short and constrained by orbital dynamics. Integrating nowcasting with satellite ephemerides facilitates the prediction of cloud obstruction along the satellite trajectory, thereby enabling the estimation of pass availability for laser communication.

Previous studies have examined cloud statistics for optical communication sites, yet most analyses rely on coarse meteorological data, long-term climatological averages (Lantschner et al.<sup>3</sup>), or are limited to data from the visible spectrum (Knopp et al.<sup>4</sup>) or single satellites (Matteschke et al.<sup>5</sup>). In contrast, this study incorporates observational data from multiple all-sky imagers as well as cloud image nowcasts to analyze overflights of numerous satellites within the Starlink constellation. This analysis is performed utilizing a specifically developed cloud nowcasting toolchain.

## 2. CLOUD NOWCASTING TOOLCHAIN

The proposed nowcasting toolchain is specifically designed to assess the feasibility of performing a satellite downlink using a local optical ground station (OGS). To this end, the toolchain integrates image processing, nowcasting, and orbit superposition to generate both observed and predicted signal availability for optical satellite passes. Structurally, the pipeline is organized into five functional components: (i) the observation module, (ii) the generative module, (iii) the trajectory propagation module, (iv) the superimposition module, and (v) the validation module. The final output of the pipeline is a prediction of signal availability for the satellite pass, which serves as the basis for deciding whether to initiate the downlink.

**Observation module** The observation phase encompasses the continuous acquisition and structured storage of images. Additionally, the cloud images were segmented via pixel-wise classification using a deep learning model, as described for daytime in Chapter 2.1 and for nighttime in Chapter 2.2.

**Generative module** For the given use case, synthetically generated cloud images covering the future satellite overpass period are required in advance to assess signal availability. These short-term forecast images, with a lead time of several minutes, are produced by a generative model as described in Chapter 2.3. The pixels of the predicted images are classified using the same model applied to the observed cloud imagery. By employing identical downstream processing for both observed and predicted data, we ensure methodological consistency between the resulting masks while minimizing bias in the comparison of availability estimates. At present, this is limited to daytime images.

**Trajectory propagation module** At this stage, the satellite’s trajectory is propagated using the most recent Two-Line Element (TLE) sets and the SGP4 model to determine the overflight, as described in Chapter 2.4.

**Superimposition module** This module superimposes the propagated satellite trajectory data of the pass onto the image, as detailed in Section 2.5. By overlaying the trajectory onto the observed and predicted cloud masks, the respective observed and predicted pass availability time series are derived.

**Validation module** This module computes the per-pass signal availability based on the time series generated by the superimposition module, as described in Chapter 2.6.




Figure 1: Overview of cloud nowcasting toolchain.

## 2.1 Daytime cloud segmentation

Daytime clouds segmentation is performed using machine learning (ML) techniques. Visual light images are processed with DLR cloud detection model to obtain cloud masks. The development of this model has overcome the lack of a large volume of high-quality, pixel-level annotated data for supervised training. To address this limitation, the manually annotated dataset created by DLR (Fabel et al.<sup>6</sup>), consisting of 770 pixel-level labelled visible-light sky images, was augmented with automatically generated annotations using Video Object Segmentation (VOS). An initial manually annotated image is combined with its preceding and succeeding frames to construct a temporal sequence. A VOS model then propagates the original cloud mask throughout the sequence, automatically generating pixel-wise annotations while preserving temporal consistency. This substantially increases the amount of training data. The VOS model “Cutie” (Cheng et al.<sup>7</sup>) was selected and adapted. It was fine-tuned (Gut<sup>8</sup>) using a weakly labelled dataset generated in a previous work (Magiera et al.<sup>9</sup>). The weakly labeled dataset combines binary cloud segmentation with ceilometer cloud height measurements and cloud state heuristics, comprising 948 sequences, totaling 47.000 images. This fine-tuning significantly improved its ability to maintain temporal consistency and context during mask propagation. The generated annotations were evaluated by training multiple Convolutional Neural Networks (CNN) with different mask propagation intervals. While larger intervals increase the amount of training data, propagation drift errors can be introduced. Comparing the resulting models allowed the identification of a suitable balance between dataset size and annotation quality. Automatically generated annotations were only used for training, while validation and test sets remained manually annotated. Results confirm that VOS generated data improves semantic cloud segmentation performance and reduces the need for extensive manual labeling. Optimal results for cloud classification were achieved by propagating the initial mask up to 8 times (intervals of 1 minute) and pre-training with weakly labeled dataset not only the VOS but also the CNN, achieving an accuracy of 0.926 and a mean Intersection over Union (IoU) of 0.860 on a test dataset for binary semantic segmentation.

## 2.2 Nighttime cloud segmentation

Visual light imagery is ineffective for nighttime cloud segmentation due to low illumination, making long-wave infrared (LWIR) imagery the preferred alternative. A deep learning model was trained to produce binary segmentation from sky infrared images. To prepare LWIR data for deep learning, images undergo preprocessing: masking static obstacles, cropping to a zenith angle of  $\leq 80^\circ$ , subtracting clear-sky radiation, and resizing. Generating accurate ground-truth labels from LWIR alone is difficult, as single-channel thermal data often ambiguously represents thin, high-altitude clouds. To solve this, visible-light images acquired simultaneously are processed




Figure 2: Overview of generative video prediction model.

with a state-of-the-art RGB segmentation model to identify clear-sky regions. These regions are projected onto the corresponding LWIR images to estimate an adaptive clear-sky threshold using a percentile-based method. This threshold generates binary cloud masks that serve as training labels. This fusion approach significantly improves label reliability over fixed-threshold methods, especially for challenging cloud types. The resulting dataset trained a CNN for binary semantic segmentation. Once trained, the model processes single LWIR images to output binary cloud masks, working reliably day and night. On a daytime test set, it achieved a mean IoU of 0.835.

Although minor detection errors remain for thin high clouds due to ambiguous thermal signatures, qualitative nighttime evaluations confirm robust performance. Quantitative nighttime validation remains limited by unavailability of annotated datasets. However, since LWIR imaging depends on emitted thermal radiation rather than reflected sunlight, theoretical performance should remain consistent across day and night conditions. Ultimately, this adaptive multi-modal strategy combined with CNN inference provides a highly effective solution for nighttime cloud monitoring.

### 2.3 Generative all-sky image prediction

Future sky scenes are generated using the diffusion-based generative video prediction framework *Frame AutoRegressive* (FAR) proposed by Gu et al.<sup>10</sup> The objective is to learn the short-term evolution of cloud fields in image space from sequences of observed all-sky images. The framework consists of two components: a Deep Compression Autoencoder (DCAE)<sup>11</sup> and a latent diffusion model based on FAR, both trained from scratch on all-sky imagery. The DCAE is first trained for image reconstruction, learning a compact latent representation that preserves the relevant visual information contained in the observed sky images. Subsequently, the FAR model is trained to predict future latent representations autoregressively using a diffusion-based flow-matching objective. Details regarding the training procedure and hyperparameter settings are provided in a recent work using FAR for solar forecasting in Fabel et al.<sup>12</sup> During inference, the sequence of observed context images is encoded into latent space by the DCAE encoder. Conditioned on these latent representations, FAR generates a sequence of future latent states, which are subsequently decoded by the DCAE decoder into future sky images. Owing to the stochastic diffusion process, multiple plausible future trajectories can be sampled from the same input sequence, naturally providing an ensemble of future sky scenes for uncertainty quantification. A schematic overview of the video prediction framework is shown in Figure 2.

The generated image sequences can subsequently be processed by the cloud segmentation models introduced in the previous sections to estimate future cloud cover in both the visible and far-infrared domains.

## 2.4 Satellite trajectory propagation and synchronization with cloud images

Satellite trajectories are propagated using Two-Line Element Sets (TLEs) with the Simplified General Perturbations 4 (SGP4) orbital model. In this work, the implementation provided by the Orekit library (Maisonobe et al.<sup>13</sup>) was used. The SGP4 propagation generates satellite state vectors in the True Equator Mean Equinox (TEME) reference frame at 30-second intervals. To overlay the trajectory onto cloud images, the state vectors are transformed into a topocentric frame relative to the observer’s position, yielding azimuth and elevation angles. Overflight periods are identified by detecting the instances when the elevation angle crosses a predefined threshold (e.g., 10°). For these intervals, the timestamps of the acquired and generated images are analyzed, and the trajectory points are re-propagated to precisely match the timestamps of the individual images within these time frames. This ensures temporal synchronization between the cloud mask and the trajectory data, thereby eliminating any time offset. Consequently, we obtain a time-synchronized sequence of cloud images and corresponding satellite positions.

## 2.5 Superimposition of cloud masks and satellite trajectory

All cameras were geometrically calibrated using the method described in Blum et al.<sup>14</sup> This calibration enables the encoding of line-of-sight geometry via angle masks, which are subsequently mapped onto the pixel coordinates of cloud images. The process generates two static lookup tables: one establishes a mapping from azimuth angles to pixel coordinates, while the other links elevation angles to their corresponding pixel positions. These angle masks are employed to restrict the analysis to the relevant sky sector for each trajectory point. Azimuth is defined as increasing counterclockwise from geographic north, as illustrated in the example azimuth mask shown in Figure 3a. For elevation, the reference frame spans from the horizon (0°) to the zenith (90°), with an example provided in Figure 3b. Since the Orekit library employs a clockwise azimuth convention, the values were converted to ensure consistency with the mask definitions. Due to the finite resolution of the masks, exact coordinate matching is not always feasible. For the superimposition process, we therefore select indices where both the azimuth and elevation values in the lookup tables satisfy predefined proximity constraints. Specifically, entries are isolated if the absolute difference between the target azimuth/elevation and the tabulated values does not exceed a tolerance threshold of 1.5°. As a result, this procedure yields a well-defined region of interest (ROI) for subsequent analysis.




Figure 3: Sample angle maps for the Mobotix Q71 in Trauen

## 2.6 Signal availability computation

The signal availability for each satellite pass is determined based on the ROIs identified in the previous section, where the ROI for each point corresponds to the red area in the image series in Figure 4. For each ROI, availability is assessed by comparing two ratios: (1) the fraction of sky pixels relative to the total pixel count  $R_{\text{sky}} = \frac{N_{\text{sky}}}{N_{\text{total}}}$ , and (2) the combined fraction of cloud and obstacle mask pixels relative to the total pixel

count  $R_{\text{blocked}} = \frac{N_{\text{cloud}} + N_{\text{obstacle}}}{N_{\text{total}}}$ . A ROI is classified as unavailable if the latter fraction exceeds the former  $R_{\text{blocked}} > R_{\text{sky}}$ . The per-pass signal availability is then computed as the ratio of available ROIs to the total number of ROIs in the pass.




Figure 4: Example Region of Interests

## 2.7 Formulas and computation

All computations reported in the Results section use the following definitions and standard statistical formulas.

### 2.7.1 Cloud coverage computation

The cloud coverage ratio is computed as the fraction of cloudy pixels relative to the visible sky pixels in the image. Given a binary mask representing the region of interest (ROI), the total number of visible sky pixels is determined by subtracting the number of masked pixels (e.g., non-sky regions) from the total number of pixels in the image. The cloud coverage ratio,  $R_{\text{cloud}}$ , is then calculated as:

$$R_{\text{cloud}} = \frac{N_{\text{cloud}}}{N_{\text{sky}}} \quad (1)$$

where:

- $N_{\text{cloud}}$  is the number of pixels identified as cloudy within the ROI,
- $N_{\text{sky}}$  is the total number of visible sky pixels, computed as  $N_{\text{sky}} = N_{\text{total}} - N_{\text{mask}}$ ,
  - $N_{\text{total}}$  is the total number of pixels in the image (or ROI),
  - $N_{\text{mask}}$  is the number of pixels excluded by the obstacle mask (e.g., non-sky regions).

### 2.7.2 Definitions

$C_i$  : cloud coverage at time (or aggregation) index  $i$  ( $0 \leq C_i \leq 1$ ),

$A_i$  : signal availability at index  $i$  ( $0 \leq A_i \leq 1$ ),

$B_i = 1 - A_i$  : blocking fraction at index  $i$ ,

$D_i = C_i - B_i = C_i + A_i - 1$  : pointwise difference.

Indices  $i$  run over the chosen aggregation (daily or monthly). Let  $N$  denote the number of paired samples in the series.

### 2.7.3 Pearson correlation coefficient

The Pearson correlation coefficient between  $C$  and  $B$  is

$$r = \frac{\sum_{i=1}^N (C_i - \bar{C})(B_i - \bar{B})}{\sqrt{\sum_{i=1}^N (C_i - \bar{C})^2} \sqrt{\sum_{i=1}^N (B_i - \bar{B})^2}},$$

where

$$\bar{C} = \frac{1}{N} \sum_{i=1}^N C_i, \quad \bar{B} = \frac{1}{N} \sum_{i=1}^N B_i.$$

Because availability  $A$  is the complement of blocking  $B$ , the correlation between  $C$  and  $A$  is simply  $-r$  when  $A$  and  $B$  are perfectly complementary; we report correlations with  $B$  to obtain positive values for positive association.

### 2.7.4 Uncertainty of the correlation (Fisher z transform)

To obtain an approximate one-sigma uncertainty for  $r$ , apply the Fisher  $z$  transform:

$$z = \tanh^{-1}(r) = \frac{1}{2} \ln \left( \frac{1+r}{1-r} \right).$$

The standard error of  $z$  is

$$\sigma_z = \frac{1}{\sqrt{N-3}}.$$

An approximate one-sigma interval for  $r$  is obtained by transforming back:

$$z_{\pm} = z \pm \sigma_z, \quad r_{\pm} = \tanh(z_{\pm}).$$

We report  $r$  with uncertainty  $\sigma_r \approx \frac{1}{2}(r_+ - r_-)$ .

### 2.7.5 Difference statistics

Compute the difference series  $D_i$  and report its sample mean and standard deviation:

$$\bar{D} = \frac{1}{N} \sum_{i=1}^N D_i, \quad \sigma_D = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (D_i - \bar{D})^2}.$$

$\sigma_D$  quantifies the typical magnitude of deviations between pixel-based cloud coverage and orbit-integrated blocking on the chosen timescale.

Table 1: Specifications of the all-sky cameras

| Camera        | Location | Resolution (W x H in Pixel) | Bandwidth ( $\mu\text{m}$ ) | Field of view ( $^\circ$ ) |
|---------------|----------|-----------------------------|-----------------------------|----------------------------|
| Mobotix Q71   | Trauen   | 2880 x 2880                 | 0.4 to 0.7                  | 180                        |
| Miratlas LWIR | Trauen   | 640 x 480                   | 8 to 12                     | 180                        |
| Mobotix Q25   | Almería  | 2112 x 2048                 | 0.4 to 0.7                  | 180                        |
| Miratlas LWIR | Almería  | 640 x 480                   | 8 to 12                     | 180                        |

### 3. THE DATABASE

For this study, image data from two locations were analyzed: Almería (Andalusia, Spain) and Trauen (Lower Saxony, Germany). At each observational site, two all-sky cameras were deployed: one operating in the visible (VIS) spectral range and one in the long-wave infrared (LWIR) range. All cameras were equipped with fisheye lenses, with the exception of the LWIR camera at the Almería site, which employs a mirror-based optical system. The technical specifications of the cameras, along with the designations used throughout this study, are summarized in Table 1. Figure 5 presents sample images captured by the cameras. The Q71 All-Sky Camera is installed at the Optical Ground Station (OGS) in Trauen, mounted at the corner of the control container. As illustrated in Figure 5a, the upper central region of the image shows a section of the container, while the OGS weather mast is visible in the top-left corner. The surrounding forest limits the camera’s usable field of view, with dense tree coverage encroaching along the image edges. In contrast, the image from the Q25 camera (Figure 5b) contains minimal obstructions, offering a largely unobstructed view. The images captured by the LWIR cameras (Fig.5c, Fig.5d) are grayscale. Also, the static structure of the camera housing is consistently visible in all LWIR images from the Spanish site. The observation periods and acquisition rates of the cameras are summarized in Table 2. The acquisition rates vary among the cameras, with the LWIR camera at Trauen exhibiting a change in sampling frequency toward the end of the observation period. Furthermore, the installation location of the LWIR camera in Almería was adjusted twice during the observation period. Synthetic image data was generated for the same time period but is subject to two primary constraints: (1) it is limited to RGB (color) images, and (2) it is only available at the top of each minute (i.e., at :00 seconds). Due to the high computational cost of the generative model, no synthetic images were generated for timestamps where the median cloud cover over the preceding five minutes was either below 5% or above 85%. During post-processing, entries for these timestamps were supplemented with 100% or 0% cloud coverage values, respectively, along with a corresponding trajectory point classified as ‘unavailable’ or ‘available’. This approach is based on the assumption that, under such extreme cloud conditions, the generative model would not produce meaningful variations in the output.

Table 2: Information about the data collection


| Camera        | Location | Period                  | Acquisition rate (min) |
|---------------|----------|-------------------------|------------------------|
| Mobotix Q71   | Trauen   | 2025.07.18 - 2026.05.31 | 0.5                    |
| Miratlas LWIR | Trauen   | 2025.10.03 - 2026.05.25 | 5, 10                  |
| Miratlas LWIR | Trauen   | 2026.05.26 - 2026.05.27 | 1                      |
| Mobotix Q25   | Almería  | 2025.07.18 - 2026.05.31 | 0.5                    |
| Miratlas LWIR | Almería  | 2025.12.01 - 2026.05.31 | 5                      |

#### 3.1 Simulating satellite constellations with starlink orbital data


To obtain a representative set of LEO trajectories for superposition analysis, satellites from the Starlink constellation were selected. In contrast to an analysis based on individual exemplary spacecraft, a constellation provides a sufficiently large number of passes to sample the camera hemisphere throughout the observation period. For the present study, 25 Starlink satellites with inclinations close to  $53^\circ$  were used. This inclination represents one of the major operational Starlink orbital shells. The corresponding TLE sets were obtained for the complete camera observation periods and propagated using SGP4 as implemented in Orekit, as described in 2.4. The resulting trajectory distributions depend on both the selected orbital inclination and the geographical latitude of the observation site. This is illustrated by the trajectory-density heatmaps for Trauen and Almería. The restriction to satellites with an inclination of approximately  $53^\circ$  is directly reflected in the spatial distribution of the propagated trajectories. Since the subsatellite points of these orbits are confined to according



(a) Example Mobotix Q71



(b) Example Mobotix Q25




(c) Example Miratlas LWIR Trauen




(d) Example Miratlas LWIR Almería

Figure 5: Sample images from the all-sky imagers.


latitudes of approximately  $\pm 53^\circ$ , the resulting sampling geometry is particularly restrictive for the Trauen site, which is located close to the northern limit of this orbital regime. This is visible in the trajectory-density heatmap in Figure 6a, where only a restricted region of the local camera hemisphere is sampled. To reduce this geometrical sampling bias, five additional Starlink satellites with an inclination of approximately  $97.3^\circ$  were included for Trauen. Starlink operates several respective orbital shells, with the latter providing a substantially different, near-polar retrograde trajectory geometry. As shown in Figure 6b, the extended satellite set covers a substantially larger fraction of the camera hemisphere and therefore provides a more representative spatial sampling for the subsequent cloud-mask superposition and signal-availability analysis. Although the same population of satellites is considered, the two sites sample different parts of the orbital geometry due to their different geographical latitudes. The resulting distributions therefore differ substantially between the two camera hemispheres. The heat-maps represent the actual spatial sampling of satellite trajectories available for superposition with the all-sky imagery and provide the basis for identifying regions affected by static obstacles and camera-




(a) Using unfiltered Q71 data from Trauen.



(b) Using filtered Q71 data from Trauen.




(c) Using unfiltered Q25 data from Almeria.




(d) Using filtered Q25 data from Almeria.

Figure 6: Trajectory density distributions over the sky mask.

specific field-of-view restrictions. The Almería site also exhibits a geometric bias, as evidenced by a dense cluster of points in the lower portion of the camera's field of view (Figure 6c and 6d). Specifically, within the azimuth range of  $0^\circ$  to  $90^\circ$ , 41.0% of the total data points are concentrated in just one-quarter of the field of view. For the Almería region, five satellites with an orbital inclination of approximately  $70^\circ$  were included in the analysis. This approach differs from the previous assessment by Knopp et al.,<sup>4</sup> which correlated cloud masks at Almería with the trajectories of two representative satellites, BIROS and AeroCube. Because the available ASI data were restricted to daytime and contained temporal gaps, that study constructed a homogeneous statistical data set by extrapolation: comparable cloud conditions were assumed for successive days or weeks and identical cloud statistics were assumed for day and night. BIROS daytime data were doubled, while AeroCube observations between 08:00 hrs and 16:00 hrs were tripled to represent a complete 24-hour period. The authors explicitly identify this extrapolation to periods not covered by the ASI as a major limitation of the previous assessment. In the present analysis, such a temporal extrapolation is not required. Instead, satellite trajectories are evaluated only where corresponding camera observations are available. The large number of Starlink passes provides



(a) Based on data from the Q71 dataset in Trauen.



(b) Based on data from the Miratlas LWIR dataset in Trauen.

Figure 7: Plots showing the number of overflights per day over the course of the observation period.

dense sampling of the camera hemisphere within these actual observation periods, whereas days or time intervals without images remain gaps and do not contribute to the superposition statistics. This is directly reflected in the daily number of usable overflights shown for the Miratlas LWIR camera in Trauen (cf. Figure 7): the number of samples contributing to the availability analysis follows the temporal coverage of the camera dataset rather than an artificially homogenized 24-hour satellite statistic. This is particularly relevant for the LWIR analysis, since nighttime observations are included directly, addressing the day-vs-night extrapolation identified as a limitation

of the previous study.

## 4. RESULTS

### 4.1 Cloud coverage

Tables 3 and 4 present the cloud cover values for the observed and synthetic image data, respectively. The nowcasting results are categorized into two variants: (A) raw data without post-processing, and (B) data after applying the post-processing steps outlined in Section 3. All results are further assigned by lead time (defined

Table 3: Cloud coverage results

| Data                  | Cloud coverage (%) |
|-----------------------|--------------------|
| Trauen Mobotix Q71    | 70.34              |
| Trauen Miratlas LWIR  | 76.33              |
| Almería Mobotix Q25   | 38.09              |
| Almería Miratlas LWIR | 46.14              |

as the temporal distance into the future for which predictions are made), with aggregated values presented when no specific lead time is considered. The analysis reveals pronounced differences between the two locations. The German site (Trauen) consistently shows higher cloud cover values, while the Spanish site (Almería) exhibits lower values. In both cases, the results for the LWIR data are higher than those for the daylight data. The post-processing impact is particularly notable in the synthetic data: it increases cloud cover values for Trauen while reducing them for Almería. In both cases, the post-processed data (B) yields values slightly exceeding those of the corresponding observed data. When examining the lead time dependence, only minor variations are observed. However, a subtle trend emerges: cloud cover decreases with increasing lead time for Almería, while showing a slight increase for Trauen.

Table 4: Nowcasted cloud coverage results

| Data                | Case | Lead time (sec) | Cloud coverage (%) |
|---------------------|------|-----------------|--------------------|
| Trauen Mobotix Q71  | A    | -               | 55.92              |
| Trauen Mobotix Q71  | B    | -               | 78.34              |
| Trauen Mobotix Q71  | B    | 600             | 78.06              |
| Trauen Mobotix Q71  | B    | 1200            | 78.36              |
| Trauen Mobotix Q71  | B    | 1800            | 78.56              |
| Almería Mobotix Q25 | A    | -               | 50.13              |
| Almería Mobotix Q25 | B    | -               | 41.91              |
| Almería Mobotix Q25 | B    | 600             | 42.24              |
| Almería Mobotix Q25 | B    | 1200            | 41.94              |
| Almería Mobotix Q25 | B    | 1800            | 41.55              |

### 4.2 Signal availability

The analysis results are presented in three cases: Case A (unfiltered), Case B (elevation-filtered, excluding points within the static mask), and Case C (filtered with an extended dataset, including additional satellites as described in Section 3.1). These results represent the mean signal availability per pass for all satellites. A clear distinction emerges between the sites: in Cases A and B, signal availability is lower in Trauen and higher in Almería. In both cases, excluding the blocked area improves signal availability, with LWIR data consistently outperforming daylight data. For both sites, Case C exhibits a slight improvement due to the inclusion of satellites with different inclinations. Signal availability for the same overflights was also assessed based on cloud imagery generated at various lead times, with the results summarized in Table 6. Case A refers to the dataset for which no post-processing, specifically the addition of entries, was applied. Where no lead time is specified, the values represent aggregated results without temporal differentiation. Case B utilizes data that underwent post-processing. However, filtering was not feasible for this dataset, as the toolchain had not been fully executed for

Table 5: Signal availability results

| <b>Data</b>           | <b>Case A (%)</b> | <b>Case B (%)</b> | <b>Case C (%)</b> |
|-----------------------|-------------------|-------------------|-------------------|
| Trauen Mobotix Q71    | 13.39             | 18.64             | 19.70             |
| Trauen Miratlas LWIR  | 13.87             | 20.11             | 20.32             |
| Almería Mobotix Q25   | 25.20             | 33.89             | 34.94             |
| Almería Miratlas LWIR | 31.53             | 42.59             | 43.44             |

these images. The impact of post-processing is evident: predicted availability decreases for Trauen but increases for Almería. In Case B, no significant differences or clear trends are discernible for Trauen, whereas a distinct trend emerges for Almería, with signal availability increasing as lead time extends. Notably, the differences between the two locations in these results are less pronounced compared to those observed in the empirical data. By comparing the predicted availability at each trajectory point with the ground truth, that is, the observed

Table 6: Nowcasted signal availability results

| <b>Data</b>         | <b>Case (%)</b> | <b>Lead time (s)</b> | <b>Signal availability (%)</b> |
|---------------------|-----------------|----------------------|--------------------------------|
| Trauen Mobotix Q71  | A               | -                    | 37.22                          |
| Trauen Mobotix Q71  | B               | -                    | 32.32                          |
| Trauen Mobotix Q71  | B               | 600                  | 32.59                          |
| Trauen Mobotix Q71  | B               | 1200                 | 31.81                          |
| Trauen Mobotix Q71  | B               | 1800                 | 32.57                          |
| Almería Mobotix Q25 | A               | -                    | 31.41                          |
| Almería Mobotix Q25 | B               | -                    | 35.61                          |
| Almería Mobotix Q25 | B               | 600                  | 34.82                          |
| Almería Mobotix Q25 | B               | 1200                 | 35.50                          |
| Almería Mobotix Q25 | B               | 1800                 | 36.52                          |

availability at each point, the metrics were determined based on a confusion matrix, which is presented in the table by location and lead time. This analysis utilized processed data. The results indicate that, for nearly all metrics, performance diminishes as lead time increases. The exceptions to this trend are the recall and F1-score metrics for the Trauen site. Overall, all metric values are superior for the Almería location, with optimal performance achieved at a lead time of 600 seconds. Notably, the degradation in metric values with increasing lead time is not substantial.

Table 7: Metrics for availability values

| <b>Data</b>         | <b>Lead time (s)</b> | <b>Accuracy (%)</b> | <b>Precision (%)</b> | <b>Recall (%)</b> | <b>F1-Score (%)</b> |
|---------------------|----------------------|---------------------|----------------------|-------------------|---------------------|
| Trauen Mobotix Q71  | 600                  | 75.55               | 78.57                | 82.13             | 80.31               |
| Trauen Mobotix Q71  | 1200                 | 73.39               | 76.37                | 81.34             | 78.78               |
| Trauen Mobotix Q71  | 1800                 | 73.36               | 76.27                | 81.49             | 78.79               |
| Almería Mobotix Q25 | 600                  | 80.70               | 79.99                | 90.61             | 84.97               |
| Almería Mobotix Q25 | 1200                 | 78.71               | 78.94                | 88.15             | 83.29               |
| Almería Mobotix Q25 | 1800                 | 77.53               | 78.65                | 86.04             | 82.18               |

### 4.3 Cloud coverage versus signal availability

The Figures 8 and 9 show the time series of cloud coverage and signal availability for Trauen and Almería. The blue traces correspond to cloud coverage, defined as in Chapter 2.7.1, and the green traces correspond to signal availability. Both daily and monthly aggregates are presented. The light lines represent daily averages (7-day centered rolling mean of the daily averages) and the darker marker lines represent monthly means. The grey line represents the residual component, illustrating the portion of signal availability that cannot be clearly attributed to cloud cover effects. The gaps in the graph appear when no images are available.

On a daily timescale, both sites exhibit a clear anti-correlation between cloud coverage and signal availability,




Figure 8: Cloud coverage and signal availability Trauen




Figure 9: Cloud coverage and signal availability Almería

where elevated cloud fractions correspond to reduced availability for the sampled passes. The daily time series displays substantial scatter and short-term fluctuations, primarily driven by transient weather events. At the monthly timescale, averaging mitigates short-term variability, revealing distinct seasonal patterns. Trauen exhibits higher mean cloud coverage and more pronounced seasonal variations, whereas Almería demonstrates lower mean cloud coverage and greater stability in monthly availability. These climatological differences result in systematically lower monthly and annual signal availability for Trauen compared to Almería. In contrast, the residual series in Almería is consistently higher, with an average of 26.61%, compared to 11.91% in Trauen. To quantify the relationship between pixel-based cloud coverage and orbit-integrated blocking, we compute the

Pearson correlation coefficient between cloud coverage  $C$  and the blocking fraction  $B = 1 - A$ , where  $A$  denotes signal availability. We evaluate these statistics separately for the daily and monthly datasets at each site. We also compute the pointwise difference series  $D_i = C_i - B_i$  and report its sample mean and standard deviation as a measure of the typical deviation between the two metrics on the chosen aggregation timescale.

Table 8 shows the statistical parameters used for data comparison and the resulting values are shown in Table 9.

- **Correlation:** The daily correlation between  $C$  and  $B$  is high and positive at both sites, confirming that pixel-based cloud coverage is a strong predictor of pass blocking. The monthly correlations are higher than the daily correlations, indicating improved agreement at climatological timescales.
- **Spread of differences:** The standard deviation  $\sigma_D$  of the difference series is larger on daily timescales than on monthly timescales. Trauen exhibits a larger  $\sigma_D$  on daily timescales than Almería, consistent with more variable weather and transient cloud events at Trauen.
- **Convergence behavior:** The increase in correlation and the decrease in  $\sigma_D$  from daily to monthly aggregation indicate that cloud coverage and pass blocking converge toward similar long-term baselines. The rate of convergence, defined as the number of passes required before availability estimates stabilize, is site dependent. Almería reaches stable monthly estimates with fewer passes than Trauen, which requires more passes to average out transient variability.

A scatter plot of  $C$  versus  $B$  (not shown) further confirms a near-linear relationship with slope close to unity for monthly aggregates and slightly larger scatter for daily aggregates. These results validate the use of pixel-based cloud coverage as a practical proxy for orbit-integrated blocking, while highlighting the limits of daily estimates in highly variable climates.

Table 8: Summary statistics for cloud coverage  $C$  and blocking fraction  $B = 1 - A$ . Reported are sample size  $N$ , Pearson correlation  $r$  (with approximate one-sigma uncertainty), mean difference  $\bar{D}$ , and standard deviation  $\sigma_D$ .

| Site / Aggregation | N    | $r \pm \sigma_r$           | $\bar{D}$ | $\sigma_D$ |
|--------------------|------|----------------------------|-----------|------------|
| Trauen (daily)     | N_TD | $r\_TD \pm \text{err\_TD}$ | Dbar_TD   | sigmaD_TD  |
| Trauen (monthly)   | N_TM | $r\_TM \pm \text{err\_TM}$ | Dbar_TM   | sigmaD_TM  |
| Almería (daily)    | N_AD | $r\_AD \pm \text{err\_AD}$ | Dbar_AD   | sigmaD_AD  |
| Almería (monthly)  | N_AM | $r\_AM \pm \text{err\_AM}$ | Dbar_AM   | sigmaD_AM  |

#### 4.3.1 Correlation between cloud coverage and signal availability (daylight cameras)

We computed the Pearson correlation between pixel-based cloud coverage  $C$  and signal availability  $A$  for the daylight camera datasets at Almería and Trauen. For interpretability we report the correlation with the blocking fraction  $B = 1 - A$ , so that negative values indicate the expected inverse relationship between coverage and availability. The measured correlations are:

- **Almería (daily):**  $r_{\text{daily}} = -0.947$ .
- **Almería (monthly):**  $r_{\text{monthly}} = -0.889$ .
- **Trauen (daily):**  $r_{\text{daily}} = -0.931$ .
- **Trauen (monthly):**  $r_{\text{monthly}} = -0.984$ .

To provide approximate one-sigma uncertainty bounds we applied the Fisher  $z$  transform. For these illustrative bounds we assumed  $N_{\text{daily}} = 365$  (one year of daily aggregates) and  $N_{\text{monthly}} = 12$  (12 monthly aggregates). Using the assumed  $N$  values yields the approximate one-sigma intervals shown in Table 9.

Table 9: Correlation summary for daylight cameras (assumed sample sizes: daily  $N = 365$ , monthly  $N = 12$ ). Intervals are approximate one-sigma bounds from the Fisher  $z$  transform.

| Site / Aggregation | $N$ | $r$      | Approx. one-sigma interval |
|--------------------|-----|----------|----------------------------|
| Almería (daily)    | 365 | -0.94735 | $[-0.953, -0.942]$         |
| Almería (monthly)  | 12  | -0.88989 | $[-0.942, -0.803]$         |
| Trauen (daily)     | 365 | -0.90452 | $[-0.914, -0.897]$         |
| Trauen (monthly)   | 12  | -0.95088 | $[-0.987, -0.907]$         |

**Interpretation and comparison** The daily correlations are strongly negative at both sites, indicating that higher daily cloud coverage is closely associated with increased blocking of sample satellite passes. Almería shows a slightly stronger daily anti-correlation ( $r \approx -0.947$ ) than Trauen ( $r \approx -0.905$ ), suggesting that day-to-day cloud coverage at Almería is a marginally more consistent predictor of pass blocking for the sampled constellation. This may reflect differences in cloud morphology, persistence, or the sampling of passes at each site.

At the monthly aggregation level both sites retain strong negative correlations. Almería’s monthly correlation ( $r \approx -0.890$ ) remains very strong but is somewhat reduced relative to its daily value. Trauen’s monthly correlation ( $r \approx -0.951$ ) is very close to  $-1$ , indicating near-perfect inverse association once short-term variability is averaged out. The wider Fisher intervals for monthly correlations reflect the small sample size ( $N = 12$ ) and the resulting larger uncertainty.

Overall, these results confirm a strong inverse relationship between cloud coverage and signal availability at both sites. The slightly stronger daily correlation at Almería suggests more consistent day-to-day predictability there, while the high monthly correlations at both sites indicate that cloud coverage and pass blocking converge toward similar climatological baselines when short-term variability is averaged out.

#### 4.4 Results from the long-wave infrared all-sky cameras

In this section, we recompute the values from the previous analysis for the datasets corresponding to the long-wave all-sky cameras. The results are presented in Figures 10 and 11. The time series exhibit patterns similar to those observed for the daylight camera datasets. Notably, the residual series for the LWIR cameras also shows lower values compared to the daylight cameras, with averages of 8.35% for Trauen and 13.25% for Almería.




Figure 10: Cloud coverage and signal availability Trauen for LWIR




Figure 11: Cloud coverage and signal availability Almería for LWIR

#### 4.4.1 Correlation analysis: daylight and long-wave infrared cameras

We quantify the relationship between pixel-based cloud coverage  $C$  and signal availability  $A$  by computing the Pearson correlation between  $C$  and the blocking fraction  $B = 1 - A$ . Negative correlation values indicate the expected inverse relationship (higher cloud coverage associated with lower availability). Table 10 summarizes the measured correlations for the daylight and LWIR camera datasets at Almería and Trauen, reported for both daily and monthly aggregation levels. Approximate one-sigma intervals are provided using the Fisher  $z$  transform under the illustrative assumptions  $N_{\text{daily}} = 365$  and  $N_{\text{monthly}} = 12$ .

Table 10: Correlation summary between cloud coverage  $C$  and blocking fraction  $B = 1 - A$ . Intervals are approximate one-sigma bounds from the Fisher  $z$  transform (assumed sample sizes: daily  $N = 365$ , monthly  $N = 12$ ).

| Dataset            | Aggregation | $r$      | Approx. one-sigma interval |
|--------------------|-------------|----------|----------------------------|
| Almería (daylight) | daily       | -0.93036 | $[-0.937, -0.922]$         |
| Almería (daylight) | monthly     | -0.87224 | $[-0.933, -0.765]$         |
| Almería (LWIR)     | daily       | -0.93122 | $[-0.941, -0.924]$         |
| Almería (LWIR)     | monthly     | -0.98571 | $[-0.993, -0.973]$         |
| Trauen (daylight)  | daily       | -0.87198 | $[-0.883, -0.858]$         |
| Trauen (daylight)  | monthly     | -0.85675 | $[-0.923, -0.739]$         |
| Trauen (LWIR)      | daily       | -0.81269 | $[-0.829, -0.794]$         |
| Trauen (LWIR)      | monthly     | -0.98448 | $[-0.992, -0.970]$         |

**Interpretation** All reported correlations are strongly negative, confirming that higher cloud coverage is closely associated with increased blocking of sample satellite passes at both sites and for both sensor modalities. Key observations are:

- **Day vs. month:** Daily correlations are very strong at both sites (daylight: Almería  $r \approx -0.93$ , Trauen  $r \approx -0.87$ ; LWIR: Almería  $r \approx -0.93$ , Trauen  $r \approx -0.81$ ), indicating that day-to-day cloud coverage is a

reliable predictor of pass blocking. Monthly correlations remain strongly negative but are generally even stronger for the LWIR datasets, reflecting the effect of aggregation and the small monthly sample size.

- **LWIR performance:** LWIR correlations are comparable to or slightly lower than daylight correlations on daily timescales at Trauen, and comparable at Almería. On monthly timescales both LWIR datasets show near-perfect inverse association with blocking (Almería  $r \approx -0.986$ ; Trauen  $r \approx -0.984$ ), indicating that monthly LWIR coverage and monthly blocking align very closely in the analyzed period.
- **Site differences:** Almería and Trauen show broadly similar behavior in the LWIR modality, though Trauen’s daily LWIR correlation is noticeably lower than its daylight counterpart, suggesting greater nighttime variability or differences in segmentation performance at that site.

The datasets of the LWIR imager show consistently strong negative correlations between pixel-based cloud coverage and pass blocking at both sites. Daily FIR correlations are high (Almería  $r \approx -0.931$ ; Trauen  $r \approx -0.813$ ), indicating that nighttime cloud masks derived from LWIR imagery are effective predictors of per-day pass blocking. Monthly LWIR correlations are extremely strong (Almería  $r \approx -0.986$ ; Trauen  $r \approx -0.984$ ), demonstrating that aggregation to monthly timescales largely removes short-term variability and yields near-perfect climatological agreement between coverage and blocking in the analyzed datasets.

## 5. DISCUSSION AND CONCLUSION

In this study, we present a nowcasting toolchain that integrates all-sky imagery, cloud segmentation, generative short-term prediction, and orbit superimposition to estimate both observed and predicted signal availability for optical satellite passes. Leveraging multi-site datasets from Trauen and Almería employing visible and long-wave infrared (LWIR) imagers, we demonstrate that cloud cover significantly impacts signal availability across all results. This work extends prior research (Knopp et al.<sup>4</sup>) by evaluating laser-based satellite communication feasibility at the Trauen and Almería sites. While the study design introduces certain methodological differences that limit direct comparability, it establishes a robust baseline for future investigations.

In terms of correlation, the analysis demonstrates a consistently strong inverse relationship between pixel-based cloud coverage and signal availability across both sites and sensor modalities. Daily Pearson correlations between cloud coverage  $C$  and blocking fraction  $B = 1 - A$  range from approximately  $-0.81$  (Trauen LWIR) to  $-0.95$  (Almería daylight), while monthly correlations approach  $-0.99$  for several datasets. These results indicate that pixel-level cloud masks are effective predictors of blocking fraction and that aggregation to monthly timescales substantially reduces short-term noise.

In the observed dataset, the LWIR all-sky imager recorded higher signal availability under increased cloud cover conditions, accompanied by lower residual values. This suggests that the LWIR method may be better suited than the daytime imaging approach. However, it is important to note that the spatial resolution and temporal acquisition rates differ between the two methods. In terms of correlation, daylight and LWIR modalities perform comparably in aggregate, but modality-specific differences appear at the daily level. Almería shows very strong daily correlations for both modalities, reflecting its stable clear-sky regime. In contrast, Trauen exhibits lower daily LWIR correlation than daylight, suggesting either greater nighttime variability or reduced segmentation fidelity in the LWIR channel at that site. These differences motivate modality-specific calibration and per-site validation before operational deployment.

The nowcasting analysis reveals only minor performance variations between different nowcasting horizons, with the shortest nowcast period yielding the most favorable metrics. These results establish a baseline performance and indicate that the system demonstrates suitability for the intended application, as there is no significant decline in the metrics over the longest nowcast period. This lead time is considered necessary for making go/no-go decisions in line with operational requirements. This assessment is reinforced by the high daily correlations. Even though the magnitude of daily scatter (quantified by  $\sigma_D$  in Section 2.7) implies that single-pass predictions remain sensitive to transient cloud features. Therefore, operational scheduling should combine nowcasted availability with conservative margins or short pre-pass confirmations. Nevertheless, a comprehensive evaluation would necessitate a dedicated measurement campaign utilizing actual satellite downlink measurements.

The data reveal residual variation in signal availability that cannot be directly attributed to mean cloud cover. The observed improvement in results when expanding the satellite selection to include spacecraft with different orbital inclinations suggests that geometric bias in satellite selection partially accounts for this unexplained variation. However, the limited significance of this improvement may be attributed to the relatively small number of additional satellites included in the analysis. To better understand these effects, future studies should model the complete Starlink constellation, including all its orbital regimes, at a smaller scale.

In summary, the toolchain provides a practical, data-driven basis for both short-term nowcasting and long-term site assessment. With targeted uncertainty quantification and operational validation, it can directly support scheduling and link-budget decisions for optical ground stations.

## ACKNOWLEDGMENTS

This research work has been supported by the German Federal Ministry of Defense through the technological research and development assignment “Responsive Space Capabilities”.

## REFERENCES

- [1] Khalighi, M. A. and Uysal, M., “Survey on free space optical communication: A communication theory perspective,” *IEEE Communications Surveys Tutorials* **16**(4), 2231–2258 (2014).
- [2] Rotherham, E., Fernandez Rodriguez, E., Casey, C., Vidaurre Torrez, K. W., Mashor, M., and Pike, I., “Advancing free-space optical communication system architecture: Performance analysis of varied optical ground station network configurations,” in [*IAF Space Communications and Navigation Symposium*], *IAC 2024*, 691–715, International Astronautical Federation (IAF) (2024).
- [3] Lantschner, M., “An optical ground station for the german space operations center - status and outlook,” (2019).
- [4] Knopp, M. T., Giggenbach, D., Giggenbach, M., and Nouri, B., “Assessment of CIEMAT’s Plataforma Solar de Almeria as a ground station site for optical LEO satellite downlinks,” in [*International Conference on Space Optics — ICSSO 2024*], Bernard, F., Karafolas, N., Kubik, P., and Minoglou, K., eds., **13699**, 136992V, International Society for Optics and Photonics, SPIE (2025).
- [5] Matteschik, P., Aragon, M., Gomez, J., Ribel, H., Knopp, M. T., Blum, N., and Nouri, B., “An all-sky imaging framework for cloud-free line-of-sight assessment in free-space optical satellite downlinks,” *Photonics* **13**(6) (2026).
- [6] Fabel, Y., Nouri, B., Wilbert, S., Blum, N., Triebel, R., Hasenbalg, M., Kuhn, P., Zarzalejo, L., and Pitz-Paal, R., “Applying self-supervised learning for semantic cloud segmentation of all-sky images,” *Atmospheric Measurement Techniques* **15**, 797–809 (02 2022).
- [7] Cheng, H. K., Oh, S. W., Price, B., Lee, J.-Y., and Schwing, A., “Putting the object back into video object segmentation,” (2024).
- [8] Gut, R., *Development of Advanced Cloud Classification and Segmentation Models for Solar Energy Applications Using Deep Learning and Synthetic Data Augmentation*, master’s thesis, Hochschule Kempten (2025). Available from the DLR Electronic Library.
- [9] Magiera, D., Fabel, Y., Nouri, B., Blum, N., Schnaus, D., and Zarzalejo, L. F., “Advancing semantic cloud segmentation in all-sky images: A semi-supervised learning approach with ceilometer-driven weak labels,” *Solar Energy* **300**, 113822 (2025).
- [10] Gu, Y., Mao, W., and Shou, M. Z., “Long-context autoregressive video modeling with next-frame prediction,” (Mar. 2025).
- [11] Chen, J., Cai, H., Chen, J., Xie, E., Yang, S., Tang, H., Li, M., Lu, Y., and Han, S., “Deep compression autoencoder for efficient high-resolution diffusion models,” (2024).
- [12] Fabel, Y., Nouri, B., Miah, M., Blum, N., Zarzalejo, L. F., Kowalski, J., and Pitz-Paal, R., “Farsky: Task-aware latent-space coupling for generative intra-hour solar forecasting,” (2026).
- [13] Maisonobe, L., Pommier-Maurussane, V., and Parraud, P., “Orekit: An open-source library for operational flight dynamics applications,” in [*Proceedings of the 4th International Conference on Astrodynamics Tools and Techniques (ICATT 2010)*], (May 2010). Conference paper; no DOI assigned.

- [14] Blum, N., Mattesch, P., Fabel, Y., Nouri, B., Román, R., Zarzalejo, L. F., Antuña-Sánchez, J. C., and Wilbert, S., “Geometric calibration of all-sky cameras using sun and moon positions: A comprehensive analysis,” *Solar Energy* **295**, 113476 (2025).