photonics

MDPI

Article

An All-Sky Imaging Framework for Cloud-Free Line-of-Sight Assessment in Free-Space Optical Satellite Downlinks

Paul Matteschk [1,2,*], Max Aragon [1,3], Jose Gomez [1,3], Helmut Ribel [4], Marcus Thomas Knopp [4], Niklas Blum [5] and Bijan Nouri [5]

1 Wematics FlexCo, 5020 Salzburg, Austria
2 International Centre for Sustainable Development (IZNE), University of Applied Sciences Bonn-Rhein-Sieg, 53757 Sankt Augustin, Germany
3 Mines Paris, Centre Observation, Impacts, Energy (O.I.E.), PSL University, 06904 Sophia Antipolis, France
4 Responsive Space Cluster Competence Center (RSC^3), Deutsches Zentrum für Luft- und Raumfahrt (DLR), Münchener Straße 20, 82234 Weßling-Oberpfaffenhofen, Germany
5 German Aerospace Center (DLR), Institute of Solar Research, Calle Doctor Carracido 44, 04005 Almería, Spain
* Correspondence: paul.matteschk@wematics.com

Abstract

Free-space optical (FSO) downlinks from satellites enable high data rates but are highly sensitive to cloud-induced attenuation and blockage. We present an integrated all-sky imaging framework for optical ground stations that converts station-local sky observations into direction- and lead-time-dependent cloud-free line-of-sight (CFLOS) decision support along predicted satellite links. The framework combines geometric calibration of hemispheric imagery, two-line element (TLE)-based orbit propagation, stereographic remapping into a common processing domain, short-horizon autoregressive sky-frame prediction using a diffusion-based sequence model, cloud/no-cloud segmentation, and a corridor-based CFLOS decision rule along the projected satellite path. The contribution lies in the operational integration of these components into a unified CFLOS-oriented sensing, prediction, and evaluation chain for optical downlink support. The framework is demonstrated at the German Aerospace Center (DLR) optical ground-station site in Trauen and evaluated using a geometry-controlled reference-track protocol across image sequences acquired at 15 s, 30 s, and 45 s cadence. Under this protocol, nowcasting-based CFLOS decisions outperformed a constant-persistence baseline across all evaluated lead times. At a 90 s lead time, the method achieved an F1-score of 0.857 and a balanced accuracy of 0.865, corresponding to gains of +0.083 and +0.089 over persistence, respectively. Positive performance margins are maintained across the full evaluated range up to a 450 s lead time. These results show that all-sky image sequences can be translated into physically interpretable CFLOS decision support and provide a basis for future network-level site-selection and handover strategies.

Keywords: free-space optical communication; optical ground stations; all-sky imaging; cloud nowcasting; cloud-free line-of-sight

Check for updates

Received: 23 April 2026
Revised: 21 May 2026
Accepted: 22 May 2026
Published: 25 May 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Free-space optical communication (FSOC) uses near-infrared laser links and offers substantially higher data rates than conventional radio-frequency systems. Additional advantages include compact terminals, spectrum-free operation, and highly directional

beams, which reduce spatial spillover and interference with neighbouring links. In space-to-ground applications, however, link availability remains strongly limited by atmospheric propagation effects [1,2].

Among atmospheric impairments, clouds are the dominant limitation for optical downlinks: water droplets strongly attenuate or completely block the beam, leading to recurrent weather-related outages even at otherwise favourable sites [1,3–5]. At the level of individual passes, this creates a short-term, direction-dependent visibility problem along the time-varying line of sight (LOS) to the spacecraft. Existing numerical weather prediction (NWP) and satellite-derived cloud products are generally too coarse in space and time to support minute-scale decisions for individual optical downlinks [6,7].

A common system-level response is optical ground-station networking with geographical diversity. Studies of German and Australasian networks and current European initiatives show that linking sites separated by hundreds of kilometres can improve availability by decorrelating local cloud conditions and enabling coordinated multi-site operation [8–10]. Nevertheless, such networks still require station-local information to determine, for a given contact opportunity, whether cloud-free line of sight (CFLOS) can be maintained along the instantaneous satellite direction [2,11].

An all-sky imager (ASI) is well suited for this purpose. It provides hemispheric observations at a cadence of a few seconds, resolves local cloud structures, and allows cloud motion to be inferred from image sequences, thereby bridging the scale gap between NWP, satellite cloud products, and station-local link conditions [1,6,12]. Multi-camera deployments such as Eye2Sky, historical ASI analyses, and recent site-assessment studies further indicate that all-sky observations can support cloud-motion nowcasting and local availability assessment at candidate station sites [13,14]. Figure 1 gives a conceptual overview of ASI-based station-local cloud monitoring for free-space optical downlinks.

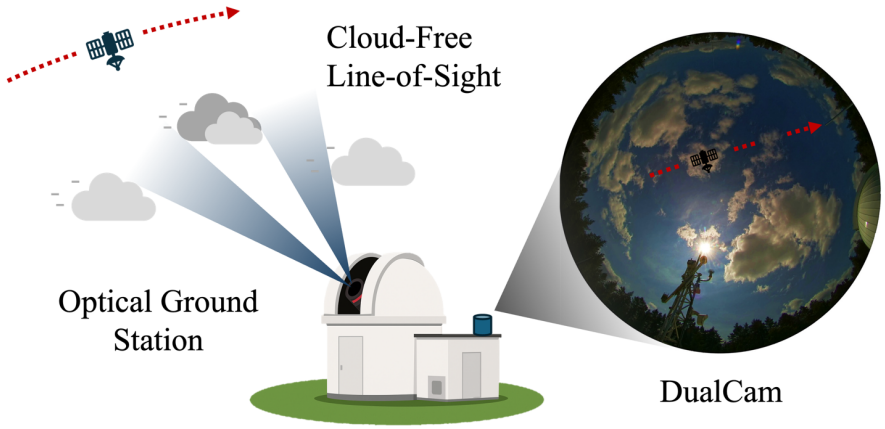


Figure 1. Conceptual overview of ASI-based station-local cloud monitoring for free-space optical satellite downlinks. The dashed red arrows indicate the satellite trajectory.

For scheduling, the relevant question is not only whether the sky is clear at the current time, but whether the LOS is expected to remain clear over the upcoming minutes of a satellite pass. Addressing this requires geometric mapping from image pixels to sky directions, time-resolved satellite-path propagation, and short-horizon prediction of the local cloud field. In the present work, CFLOS is represented by binary cloud masks derived from observed and predicted all-sky frames. Prior work on all-sky-image nowcasting has explored convolutional, recurrent, transformer-based, and hybrid physical–data-driven approaches for short-horizon forecasting of cloud fields and irradiance-relevant sky evolution [15–18], while the broader methodological basis is linked to advances in physics-informed, transformer-based, and diffusion-based video prediction [19–22].

Against this background, we present an integrated all-sky-imaging framework for translating station-local sky observations into CFLOS assessment and decision support along predicted satellite trajectories. CFLOS is evaluated directionally across the sky dome rather than as a single scene-level clear/cloudy label, because blockage risk varies with azimuth and elevation along the satellite path; this is particularly relevant for optical downlinks, where low-elevation segments traverse longer atmospheric slant paths. The framework combines calibrated hemispheric imaging, TLE-based satellite propagation, cloud segmentation based on established sky-image methods [23–27], stereographic reprojection, short-horizon autoregressive sky-frame prediction using a diffusion-based sequence model, and a corridor-based CFLOS decision rule along the projected LOS. While the individual building blocks are largely based on established methods, we are not aware of prior work that integrates them into a unified CFLOS-oriented sensing, prediction, and evaluation chain for optical ground-station support. The main contribution is this integration, together with the explicit assessment of CFLOS along the projected satellite line of sight.

The remainder of the paper is structured as follows. Section 2 describes the ASI hardware, site configuration, and image-sequence dataset. Section 3 presents geometric calibration, and Section 4 describes satellite-pass computation and path projection. Section 5 introduces cloud segmentation and stereographic reprojection. Section 6 presents short-horizon sky-frame nowcasting, and Section 7 formulates the corridor-based CFLOS decision method. Section 8 reports the reference-track evaluation results, while Sections 9 and 10 discuss limitations and summarize the conclusions.

2. Imaging System, Site, and Dataset

This section describes the DualCam imaging system, its installation at the DLR Optical Ground Station (OGS) in Trauen, and the all-sky image-sequence dataset used for nowcasting and CFLOS evaluation.

2.1. Hardware Configuration

The DualCam system comprises two co-mounted imaging modules for visible (VIS) and thermal infrared (TIR) sensing, operated under a common timing and control scheme. Figure 2 illustrates the modular architecture of the DualCam system.

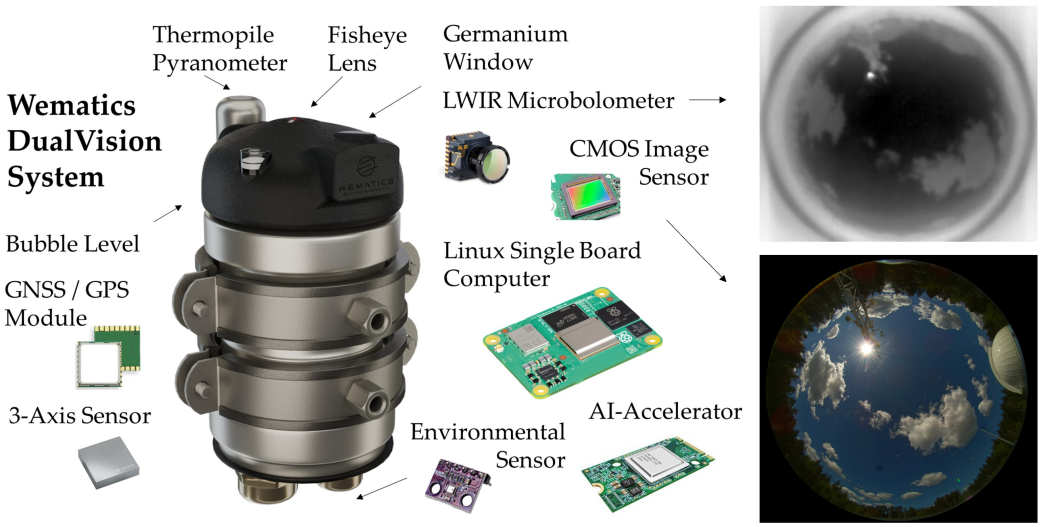


Figure 2. DualCam system architecture with synchronized VIS and TIR modules controlled by a shared onboard computing platform.

The imaging pipeline includes synchronized capture, metadata logging (e.g., temperature, exposure, GPS timing and location information), and storage of both raw and

processed outputs. During the measurement campaigns considered in this study, images were acquired at cadences of 15 s, 30 s, or 45 s, depending on the evaluation setting. For the visible channel, HDR all-sky images were generated from exposure-bracketed acquisitions in order to improve scene readability under high-contrast illumination conditions. Each sensor is geometrically calibrated independently to establish a pixel-to-direction mapping in azimuth and elevation, and subsequent co-registration is achieved by alignment in a common angular reference frame. The geometric calibration procedure is described in detail in Section 3. Static foreground obstructions near the horizon were excluded in downstream processing using a predefined horizon mask. An onboard AI accelerator enables low-latency edge inference and local processing of imaging and sensor metadata. Full technical specifications are provided in Table 1.

Table 1. Summary of hardware characteristics of the DualCam system.

Component	Specification
Compute platform	Quad-core ARM Cortex-A76, 2.4 GHz; AI accelerator: 13 TOPS (tera-operations per second); PoE (IEEE 802.3af)
All-sky camera (VIS)	1/2.3″ CMOS, 4056 × 3040 px, 180° FOV; spectral range: 400–720 nm, IR-blocked; rolling shutter; exposure range: 60 µs to 120 s
Thermal camera (TIR)	Microbolometer, 160 × 120 px, 160° FOV; spectral range: 8–14 µm; uncooled

2.2. Installation at the OGS Trauen Site

The DualCam system was installed at the DLR Optical Ground Station in Trauen, Lower Saxony, Germany [28]. The facility comprises a tracking telescope with largely remote and automated operation and a dedicated control container. The DualCam unit was mounted on a fixed mast within the compound to provide a mostly unobstructed hemispheric view for station-local cloud monitoring. Figure 3 shows the installed system next to the dome.



Figure 3. DualCam system installed at the DLR OGS in Trauen, providing a largely unobstructed hemispheric view for station-local cloud monitoring.

2.3. All-Sky Image Sequence Dataset

The nowcasting dataset comprises 2642 all-sky image sequences acquired between 26 May and 11 September 2025, spanning 86 acquisition days within a 108-day period (temporal coverage: 78.9%). Each sequence contains 23 frames at a spatial resolution of 128 × 128 pixels, yielding 60,766 individual frames in total.

The dataset covers daytime operation between 05:00 and 20:00 and includes a broad range of cloud situations relevant to optical ground-station visibility. For the present study, the sequences were analysed at effective cadences of 15 s, 30 s, and 45 s by temporal subsampling. A temporally stratified train/test split was performed at sequence level using contiguous chronological blocks rather than random frame-level sampling, thereby preserving the seasonal distribution of scenes while reducing temporal leakage between training and testing. The resulting split comprised 2114 sequences (80%) for training and 528 sequences (20%) for testing. The test subset was reserved for evaluation on unseen data. An overview of the temporal coverage and acquisition statistics of the dataset is shown in Figure 4.

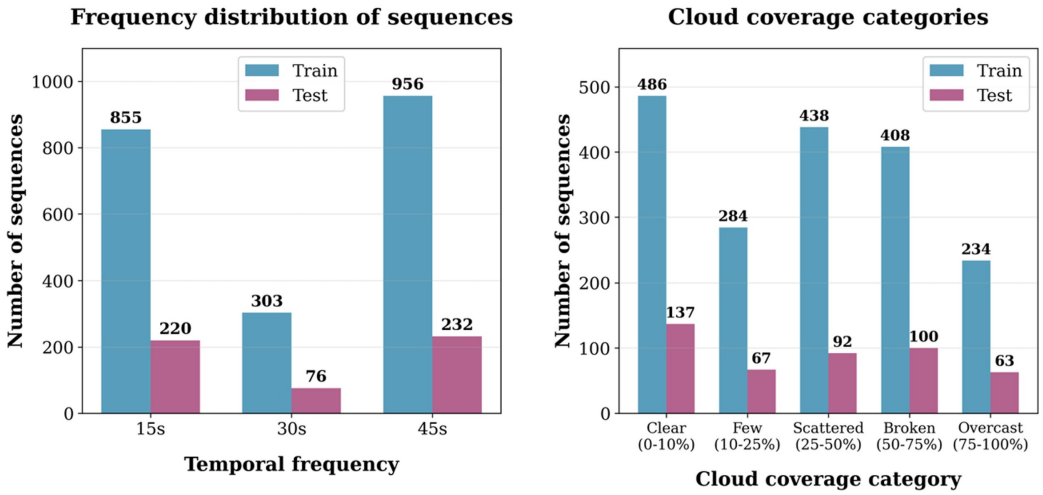


Figure 4. Temporal coverage of the all-sky image-sequence dataset used for nowcasting and CFLOS evaluation.

The dataset should be interpreted as a campaign-based dataset acquired at one optical ground-station site during the available measurement period, rather than as a year-round or site-independent reference dataset.

3. Geometric Calibration

3.1. Definition and Rationale

Geometric calibration establishes a deterministic mapping from image pixels to sky directions expressed as azimuth (Az) and elevation (El). These directions are represented in the local horizon frame (east–north–up, ENU) at the observing site. Each pixel (i, j) is assigned a unique direction by a calibrated mapping $\mathcal{G}$:

$$(\mathrm{Az}(i,j),\, \mathrm{El}(i,j)) = \mathcal{G}(i,j), \tag{1}$$

where, for each pixel (i, j), the mapped azimuth satisfies $\mathrm{Az}(i,j) \in [0°, 360°)$ and the mapped elevation satisfies $\mathrm{El}(i,j) \in [0°, 90°]$. Azimuth increases clockwise from geographic north, and elevation is measured from the horizon (0°) to the zenith (90°). Combined with the cameras' known locations, this mapping enables cloud-field quantification in angular coordinates, projection of predicted satellite paths into the image domain, and fusion of multi-sensor observations on a common angular grid.

3.2. General Calibration Principle

Celestial-object-based calibration relies on associating detected image features of celestial reference objects with the corresponding sky directions predicted from ephemerides for known site coordinates and observation times. A parametric camera model $\mathcal{G}_{\boldsymbol{\Theta}}$, with pa-

rameter vector $\boldsymbol{\Theta}$, is then estimated such that the model-projected positions of the reference directions agree as closely as possible with the detected image positions:

$$\hat{\boldsymbol{\Theta}} = \arg\min_{\boldsymbol{\Theta}} \sum_{k=1}^{N} \left\| \mathbf{u}_k^{\mathrm{det}} - \mathbf{u}_k^{\mathrm{pred}}(\boldsymbol{\Theta}) \right\|^2, \tag{2}$$

where $\mathbf{u}_k^{\mathrm{det}}$ denotes the detected pixel location of celestial reference object k, and $\mathbf{u}_k^{\mathrm{pred}}(\boldsymbol{\Theta})$ denotes the corresponding model-projected pixel location obtained from the ephemeris direction under the camera model. In wide-angle all-sky systems, pixel-space residuals are used here as a practical fitting quantity within the chosen projection model, whereas the physical interpretation of calibration quality is ultimately angular.

3.3. Applied Strategy for VIS and TIR

Distinct calibration routes were applied to the visible (VIS) and thermal infrared (TIR) sensors. The VIS fisheye camera was calibrated exclusively at night from star observations using the ORION solver [29]. The resulting VIS mapping was then validated in two ways: first, by back-projecting the same reference stars used during calibration to assess internal consistency, and second, by an independent daytime check using Sun positions and solar centroids, without refitting the night-time solution. The Sun-based validation followed the Sun-position modelling approach implemented in SUMO [30]. The co-aligned low-resolution TIR camera was likewise calibrated and assessed exclusively from solar centroids using the SuMo methodology.

3.4. Star Identification Support

Stellar identification and naming were supported by the planetarium software Stellarium, version 26.1 [31]. Detections were performed on a frame-wise basis. Isolated and compact stellar footprints with clean morphology and stable apparent motion were retained, whereas close pairs, blends, and tracks exhibiting non-smooth motion, for example due to thin clouds or artefacts, were excluded. Prominent asterisms were used to verify the field orientation prior to inclusion in the star set.

3.5. Calibration Data and Results

For the VIS camera, stellar calibration with ORION was performed on the night of 1–2 July 2025 (22:00–01:30 UTC; 83 frames), anchored to eleven reference stars (Figure 5a). Internal consistency, evaluated from detected star centroids and their ephemeris-projected positions under the fitted model, yielded a mean residual of 0.86°. An independent daytime validation without refitting was then performed using Sun tracks. Across 1531 clear-sky frames acquired on 13 June 2025, the mean angular deviation remained below 1.21°. Sun positions detected and predicted in the VIS images are compared in Figure 5b (top). Relative to the higher accuracies reported for dedicated solar calibration under favourable conditions [30], the larger deviation observed here is likely influenced by the more challenging centroid localisation in HDR all-sky imagery, including glare, blooming, partial saturation, and possible small residual systematic effects.

For the TIR camera, calibration and validation were performed exclusively from solar centroids using the SUMO formulation [30]. Sun detections were obtained on 1 July 2025 (05:51–17:06 UTC; 889 timestamps), of which 219 high-quality detections were retained for parameter estimation and evaluation. The resulting mean angular deviation was 0.27°. Although calibration from multiple days or celestial bodies would generally be preferable [30], the present dataset spans broad solar azimuth and elevation ranges (85 to 281° and 21 to 60°), which constrains the main parameters of the chosen camera model over the angular range most relevant to the downstream analysis.

Overall, both calibrations were sufficiently accurate for the downstream projection, co-registration, and CFLOS analyses carried out in this work. The reported VIS and TIR deviations should not be interpreted as directly comparable sensor-quality indicators, as they arise from different modalities, spatial resolutions, and validation conditions.

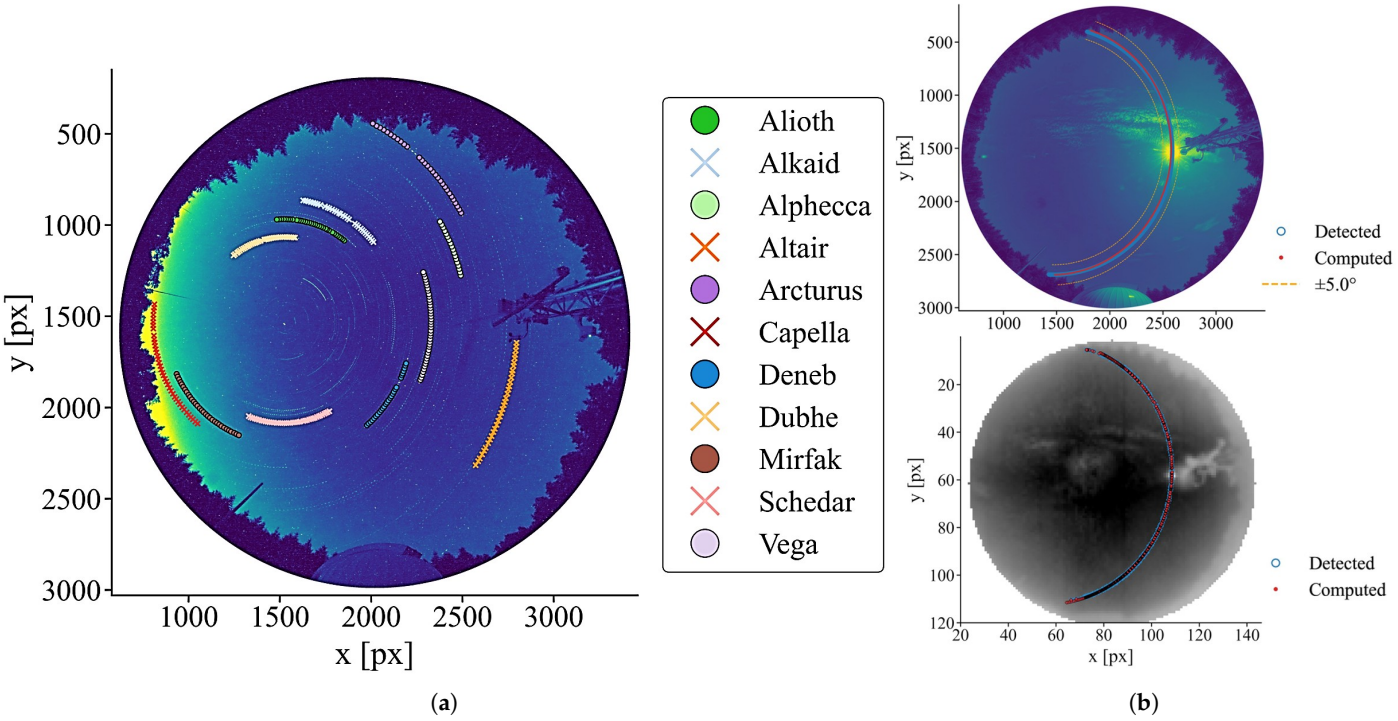


Figure 5. (**a**) VIS star-based calibration based on detected stellar positions and ephemeris-projected reference positions. (**b**) Sun-track validation for VIS (top) and Sun-based calibration and validation for TIR (bottom).

The resulting Az/El LUTs provide a common angular reference for both sensors and thereby enable deterministic VIS–TIR co-registration by sky-direction equivalence. Both modalities are resampled onto a shared angular grid, and pixels mapping to the same $(\mathrm{Az}, \mathrm{El})$ direction are associated by angular nearest-neighbour matching. This co-registered representation is used in Section 5.1 for cloud segmentation and subsequent multi-sensor processing.

4. Satellite Geometry and Path Projection

This section describes the generation of apparent satellite trajectories in azimuth–elevation (Az/El) coordinates for a fixed observing site. Public two-line element (TLE) sets are propagated with the Simplified General Perturbations 4 (SGP4) model to obtain time-resolved line-of-sight (LOS) directions [32]. The resulting Az/El time series provides the geometric input for projecting satellite paths into the all-sky image domain using the calibration from Section 3, thereby forming the basis for the subsequent CFLOS evaluation.

4.1. Inputs, Time Scales, and Reference Frames

Satellite orbits are specified by classical TLE sets (line 1/2) with epoch t_0. The observing site is defined by geodetic latitude, longitude, and ellipsoidal height (φ, λ, h) with respect to the WGS 84 reference ellipsoid.

SGP4 propagation yields satellite state vectors in the True Equator Mean Equinox (TEME) frame. For each evaluation time t_k (UTC), TEME states are transformed to an Earth-fixed frame consistent with the International Terrestrial Reference Frame (ITRF)

using IERS 2010 conventions and Earth orientation parameters (EOP) [33]. The topocentric LOS is then expressed in the local ENU horizon frame at the observing site, using the azimuth–elevation convention defined in Section 3.

4.2. Propagation, Sampling, and Pass Segmentation

For each TLE epoch t_0, satellite states are propagated over a 24 h interval using the SGP4 implementation in Orekit [34]. States are sampled on a uniform time grid with cadence $\Delta t = 10$ s.

At each time step t_k, the apparent azimuth $\alpha(t_k)$ and elevation $\epsilon(t_k)$ are computed for the observing site. Samples are classified as visible if $\epsilon(t_k) \geq \epsilon_{\min}$, where $\epsilon_{\min}$ denotes the minimum elevation threshold imposed by the station mask. Maximal consecutive sequences of visible samples define individual overflight segments.

4.3. Output Representation

All visible samples are exported as a time series of UTC timestamp, azimuth, and elevation. Each row corresponds to one sampled instant within an overflight segment. These trajectories serve as the geometric input for subsequent projection into the all-sky image domain and corridor-based CFLOS evaluation.

Figure 6 illustrates representative Az/El trajectories generated using the described procedure for multiple passes of a low-Earth-orbit (LEO) satellite.

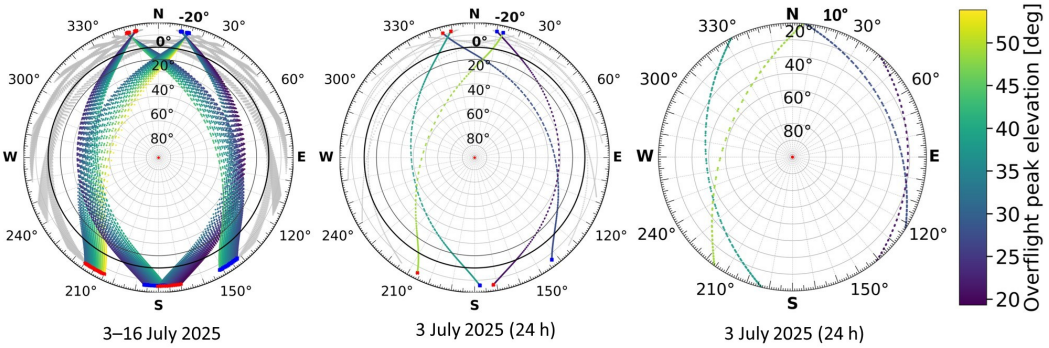


Figure 6. Azimuth–elevation trajectories of the LEO satellite Flying Laptop (NORAD 42831) in radar projection (0° at north, azimuth clockwise, elevation decreasing radially). (**Left**): all passes from 3–16 July 2025 (UTC). (**Center**): passes on 3 July 2025 with start markers in blue and end markers in red. (**Right**): communication-relevant trajectory segments above the minimum elevation threshold of 10°. The trajectory colour encodes the peak elevation of the corresponding pass.

5. Cloud Segmentation and Stereographic Reprojection

This section defines the processing domain for cloud masking, nowcasting, and satellite-path evaluation. VIS and TIR data are co-registered in angular coordinates using the calibration from Section 3. Cloud masks are generated from the co-registered imagery, while the nowcasting and CFLOS evaluation operate in a stereographic KTAN representation that reduces horizon foreshortening and improves the representation of low-elevation cloud structures.

5.1. Cloud-Mask Generation

Using the per-pixel azimuth/elevation lookup tables from Section 3, VIS and TIR data are resampled onto a shared angular grid for direction-consistent fusion in the subsequent refinement stage.

Binary cloud masks $m \in \{0,1\}^{H \times W}$ are first derived in the native fisheye RGB domain, where H and W denote image height and width, respectively, and label 1 denotes cloud. The

primary RGB segmentation model is a compact U-Net encoder–decoder architecture with skip connections [35], trained on 681 manually annotated all-sky image–mask pairs from ASI imagery, including images acquired with other all-sky camera systems. This annotated dataset was split into 440 training, 110 validation, and 131 held-out test samples. Training used binary cross-entropy with logits, the Adam optimizer, light geometric augmentation, and early stopping based on the validation intersection-over-union (IoU). On the held-out test set from this annotated ASI mask dataset, evaluated inside the circular field-of-view mask, the model achieved a cloud-mask IoU of 0.641 and a Dice/F1-score of 0.710; remaining errors were mainly observed near cloud boundaries, glare-affected regions, and optically thin clouds.

For observed frames, the RGB-based masks are refined using the co-registered TIR channel as an auxiliary physical constraint in visually ambiguous situations. The TIR measurements provide additional contrast between clear sky, cloud, and bright visible-channel artefacts, and are used to support mask correction in cases such as glare, bright haze, reduced RGB contrast, or partial saturation. Thus, the observed-frame reference masks are not purely RGB products, but VIS masks with TIR-supported refinement. For predicted frames, no corresponding future TIR observations are available; therefore, masks are derived from the predicted VIS frames alone. In the present implementation, the nowcasting model is applied only to VIS frame sequences; TIR information is used for observed-frame mask refinement but is not itself predicted.

The binary cloud mask is central to the framework because it provides the common decision representation for observed frames, predicted frames, and corridor-based CFLOS evaluation. The reported segmentation metrics document the quality of the mask-generation step used in this processing chain. However, these metrics are not intended as a comprehensive benchmark of cloud-segmentation performance across sensors, cloud regimes, and illumination conditions, since all-sky cloud segmentation remains partly ambiguous at pixel level, especially for optically thin clouds, diffuse boundaries, haze, glare, and low-contrast sky conditions [23–27]. Accordingly, the observed-frame masks are interpreted as segmentation-derived, TIR-supported reference products rather than uncertainty-free reference labels. The present analysis focuses on how this image-derived cloud-state representation is projected into satellite-line-of-sight geometry and translated into corridor-level CFLOS decisions.

5.2. Stereographic Reprojection in the KTAN Domain

Prior to nowcasting, VIS frames and the corresponding cloud masks are reparameterized into a conformal stereographic domain (KTAN; north up, east right). This representation is motivated by the CFLOS task: low-elevation clouds are particularly relevant because the corresponding optical paths traverse long atmospheric slant ranges, yet these regions are strongly compressed in native fisheye imagery. The KTAN mapping, therefore, allocates more effective spatial resolution to low-elevation cloud structures while preserving local angular geometry.

Let el denote elevation, let az denote azimuth, and let

$$\theta = 90^\circ - el$$

denote the zenith angle. The usable angular boundary is defined by $\theta_{\max} = 80^\circ$, corresponding to $el = 10^\circ$ for the employed lens with an effective field of view of approximately 160° in diameter.

Let R denote the radius of the normalized KTAN image domain and let K denote a dimensionless scaling factor controlling the radial stretch of the stereographic mapping. The radial coordinate is defined as

$$r(\theta) = R\,\frac{\tan\left(\frac{K\theta}{2}\right)}{\tan\left(\frac{K\theta_{\max}}{2}\right)}, \qquad 0 \le \theta \le \theta_{\max}, \qquad K = 1.4. \tag{3}$$

Azimuth is then mapped to KTAN image coordinates (x, y) with the image center (x_c, y_c) via

$$x = x_c + r(\theta)\sin(az), \qquad y = y_c - r(\theta)\cos(az), \tag{4}$$

with north oriented upward and east to the right.

The KTAN grid is fixed to 128×128 pixels, and the boundary corresponding to $el = 10^\circ$ is normalized to the rim of the valid KTAN domain, i.e., $r = R$. RGB frames are resampled by bilinear interpolation, whereas binary masks are resampled by nearest-neighbor interpolation. Pixels outside the valid elevation range are excluded from the prediction loss. By reducing horizon foreshortening, the KTAN representation improves the stationarity of apparent cloud motion for short-horizon prediction and stabilizes the geometric overlay of projected line-of-sight paths.

5.3. Satellite-Path Overlay in the Processing Domain

Satellite trajectories propagated in azimuth/elevation (Section 4) are rendered directly in KTAN coordinates to ensure pixelwise consistency with cloud masks. Rendering is restricted to elevations $\geq 10^\circ$, i.e., to the valid angular domain retained after reprojection. Figure 7 compares camera-native fisheye rendering to KTAN rendering and illustrates the resulting co-registered overlay.

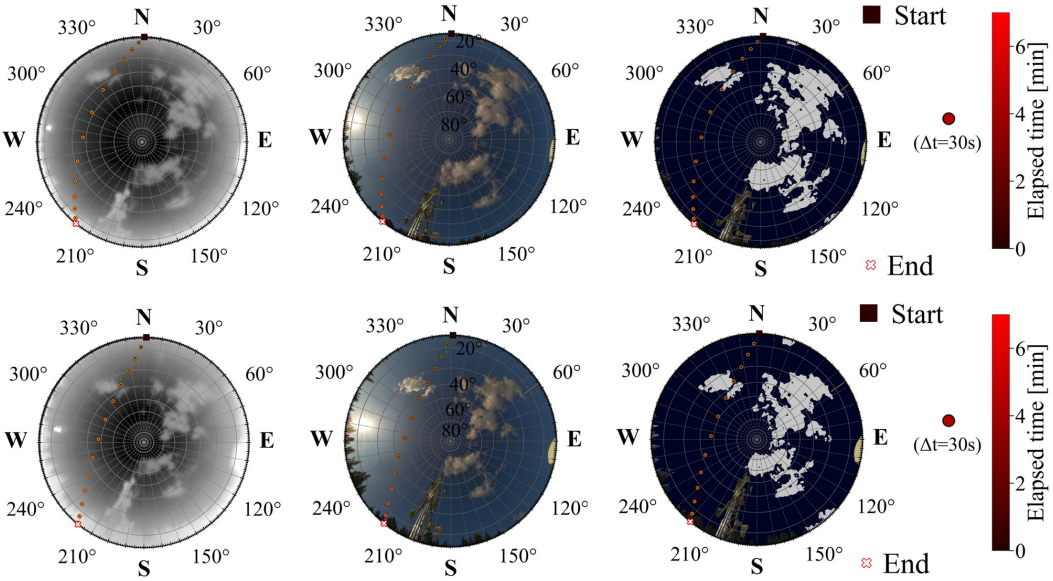


Figure 7. Satellite-track overlay for Flying Laptop (NORAD 42831) during a visible pass on 3 July 2025 from 16:56 to 17:11 UTC. The upper row shows the native fisheye representation and the lower row the KTAN representation. Columns show the TIR image, the RGB image with projected satellite path, and the binary cloud mask. Only trajectory segments above 10° elevation are displayed.

6. Sky-Frame Nowcasting

Operational support for optical downlinks requires short-horizon forecasts of the local cloud field. Here, cloud evolution is formulated as image-sequence prediction in the KTAN domain introduced in Section 5.2. Predicted sky frames are subsequently converted

into binary cloud masks and evaluated for CFLOS feasibility. The dataset is described in Section 2.3.

6.1. Prediction Domain and Target Representation

All frames are represented in the KTAN domain described in Section 5.2, with RGB inputs normalized to $[0, 1]$. The downstream decision variable is the binary cloud mask rather than the RGB image itself. Consequently, predicted RGB frames are converted into cloud masks using the RGB-based segmentation stage, and CFLOS is evaluated exclusively in mask space. For observed frames, masks may additionally be refined using co-registered TIR information as described in Section 5.1. This ensures a consistent decision representation for observed and predicted sky states.

6.2. Autoregressive Diffusion Nowcasting

Future sky frames are generated using SkyDiff, a conditional diffusion-based video prediction model for all-sky image nowcasting. SkyDiff serves as the prediction backbone of the CFLOS pipeline. It operates on RGB image sequences in the 128×128 KTAN domain described in Section 5.2. RGB inputs are normalized to $[0, 1]$ before inference.

At inference time, a fixed-length context window of preceding observed frames is provided. SkyDiff predicts the next RGB sky frame by conditional denoising. Multi-step forecasts are obtained autoregressively: after each step, the generated frame is appended to the context window and the oldest frame is removed. This is repeated until the forecast horizon is reached. The predicted RGB frames are then converted into binary cloud masks using the RGB segmentation stage described in Section 5.1; CFLOS evaluation is performed only on these masks.

For this study, SkyDiff was trained on the training subset described in Section 2.3 and applied without retraining to the reserved test subset. This paper does not introduce a new image-generation model, but evaluates short-horizon all-sky predictions within a CFLOS assessment chain. Architectural choices, diffusion training settings, and standalone image-prediction benchmarks are reported in the companion SkyDiff study by Aragon et al. The inference scheme is shown in Figure 8.

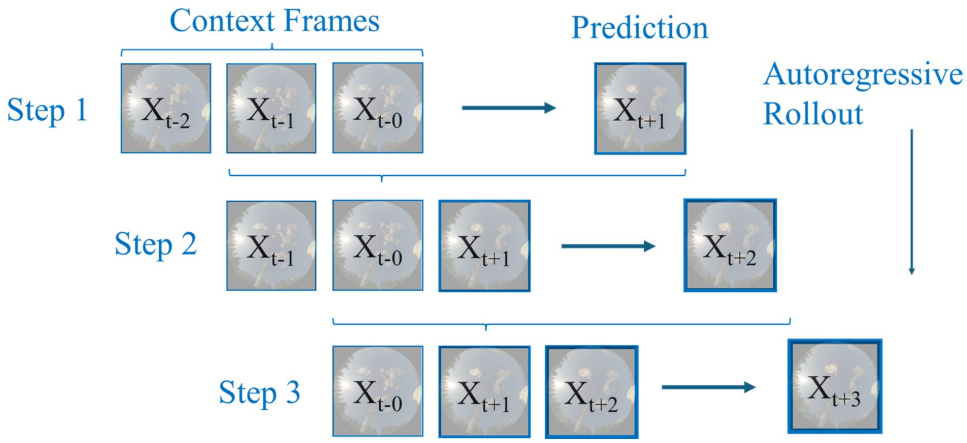


Figure 8. Autoregressive SkyDiff inference scheme. Observed context frames are used to predict the next frame, and generated frames are recursively reused for multi-step prediction.

This approach was chosen because cloud fields exhibit complex, non-rigid motion and deformation. In the present CFLOS framework, the main role of the diffusion model is to preserve cloud morphology sufficiently over short rollout horizons so that predicted frames can be converted into geometrically meaningful cloud masks for line-of-sight evaluation.

7. CFLOS Decision Method

Observed and predicted cloud masks are translated into lead-time-dependent cloud-free line-of-sight (CFLOS) decisions along the satellite trajectory using a corridor-based evaluation rule. The trajectory is discretized into support points, and each interval between two consecutive points is evaluated independently. This yields a binary CFLOS interval sequence that can be compared between observations, model predictions, and a persistence baseline.

7.1. Angular Corridor Definition

The propagated satellite trajectory is represented as a sequence of azimuth–elevation samples, with each direction mapped to a unit vector on the celestial hemisphere. For interval i, the satellite line-of-sight segment Γ_i is defined by spherical-linear interpolation along the great-circle arc between two consecutive support directions.

Using the calibrated pixel-to-direction mapping from Section 3, every valid pixel $\mathbf{x}$ in the projected image domain corresponds to a unit direction vector $\mathbf{u}(\mathbf{x})$. Let Ω_{valid} denote the set of pixels that belong to the evaluable sky domain, i.e., pixels inside the usable field of view after exclusion of invalid regions such as horizon-masked areas or static foreground obstructions.

The angular distance between a pixel direction and the trajectory segment is defined as

$$\delta_i(\mathbf{x}) = \min_{\mathbf{w}\in\Gamma_i} \arccos\left(\mathbf{u}(\mathbf{x})^\top \mathbf{w}\right). \tag{5}$$

The corridor for interval i is then defined as the set of valid pixels whose angular distance to the segment does not exceed a fixed threshold α:

$$\Omega_i = \{\mathbf{x} \in \Omega_{\text{valid}} \mid \delta_i(\mathbf{x}) \leq \alpha\}. \tag{6}$$

In the present evaluation, the corridor half-angle is set to $\alpha = 3^\circ$. This value was chosen as a conservative operational margin based on the geometric calibration accuracy reported in Section 3, where angular deviations remained on the order of 1°. The adopted corridor width, therefore, includes an additional buffer to account for residual projection, centroiding, and discretisation uncertainties in the end-to-end CFLOS evaluation.

7.2. Interval-Based CFLOS Decision Rule

Let t denote the evaluation time associated with a given observed or predicted frame. Furthermore, let $m(\mathbf{x}, t) \in \{0, 1, 2\}$ denote the cloud mask label at pixel $\mathbf{x}$ and time t, where 0 denotes clear sky, 1 denotes cloud blockage, and 2 denotes excluded pixels. Excluded pixels are not considered in the CFLOS decision; this class mainly represents regions outside the evaluable sky domain, for example due to the horizon mask or static foreground obstructions such as trees or infrastructure.

For each corridor Ω_i, only non-excluded pixels are retained. Let $N_i^{\text{valid}}(t)$ denote the number of valid pixels in the corridor at time t, and let $N_i^{\text{blocked}}(t)$ denote the number of blocked pixels among them. The blocked fraction is defined as

$$f_i(t) = \frac{N_i^{\text{blocked}}(t)}{N_i^{\text{valid}}(t)}. \tag{7}$$

The CFLOS state for interval i at time t is then defined as

$$y_i(t) = \begin{cases} -1, & N_i^{\text{valid}}(t) < N_{\text{min}}^{\text{valid}}, \\ 1, & f_i(t) \geq \tau \text{ and } N_i^{\text{blocked}}(t) \geq N_{\text{min}}^{\text{blocked}}, \\ 0, & \text{otherwise}, \end{cases} \tag{8}$$

where 0 denotes *CLEAR*, 1 denotes *BLOCKED*, and -1 denotes that the interval is not evaluable due to insufficient valid pixels. The three cases form a mutually exclusive partition; in particular, the CLEAR state ($y_i = 0$) applies only to evaluable intervals with $N_i^{\text{valid}}(t) \geq N_{\text{min}}^{\text{valid}}$. In the present implementation, fixed thresholds are used, namely $\tau = 0.02$, $N_{\text{min}}^{\text{blocked}} = 3$, and $N_{\text{min}}^{\text{valid}} = 10$. These values were chosen as conservative operational thresholds to suppress isolated mask artefacts and to avoid unstable decisions in sparsely populated corridor segments.

7.3. Reference-Track Evaluation Protocol

To enable aggregation across many cloud scenes under controlled geometric conditions, a geometry-controlled reference-track protocol is used. Instead of synchronizing each image sequence with an individual satellite pass, one representative propagated trajectory is discretized into consecutive line-of-sight intervals and reused across multiple sequences. This isolates the CFLOS decision logic from pass-to-pass geometric variability and allows the evaluation to probe a broad range of cloud situations under fixed corridor geometry. The protocol is, therefore, intended as a controlled evaluation setting for the CFLOS decision chain rather than as a direct emulation of fully time-synchronised operational pass conditions. Accordingly, the reported metrics should be interpreted as trend indicators under the available campaign conditions, not as full proof of general CFLOS prediction performance.

For each sequence, interval-wise CFLOS states are derived along the same fixed reference trajectory from both the observed mask sequence and the model-predicted mask sequence. These state sequences are compared at increasing forecast horizons. As an additional baseline, a constant-persistence forecast is used: the CFLOS state at the first evaluated interval is taken as the predictor for all subsequent intervals within the same sequence. Persistence, therefore, denotes a naive forecast assumption and not a property of the observed CFLOS evolution itself. Sequences analysed at effective cadences of 15 s, 30 s, and 45 s are pooled in the aggregated analysis. Consequently, a horizon defined in interval steps does not correspond to a unique physical lead time, but to a cadence-dependent temporal offset.

7.4. Forecast Horizon and Lead-Time Aggregation

Forecast horizon h is defined in interval steps. For each horizon, evaluation starts at interval $h + 1$, so that predicted and reference CFLOS states are compared only on matched target intervals. Because sequences are analysed at cadences of 15 s, 30 s, and 45 s, the same horizon index corresponds to different physical lead times. Results are, therefore, reported by physical lead time.

Where the same physical lead time occurs in multiple cadence subsets, the corresponding confusion counts are pooled descriptively before metric computation. Only intervals with valid CFLOS labels are retained. The constant-persistence baseline holds the first evaluated CFLOS state fixed for all subsequent target intervals within the same sequence. The pooled metrics are, therefore, interpreted descriptively, since samples at a given physical lead time may originate from different cadence subsets and are not statistically independent.

8. Results and Discussion

This section evaluates the integrated CFLOS processing chain under the geometry-controlled reference-track protocol. The evaluation combines sky-frame prediction, cloud-mask generation, satellite-path projection, and corridor-based CFLOS classification. The objective is to assess whether the nowcasting-based chain provides useful CFLOS information beyond a constant-persistence assumption.

All quantitative results are computed in mask space by applying the same satellite-path corridor and CFLOS decision rule to observed-frame reference masks and predicted-frame masks. The reported values, therefore, describe corridor-level agreement with segmentation-derived reference masks under controlled trajectory geometry. They should be interpreted as descriptive reference-track performance indicators, not as end-to-end measures of optical-link availability during fully time-synchronised satellite passes.

8.1. Representative CFLOS Evaluation Example

Figure 9 illustrates the end-to-end CFLOS evaluation for one 45 s-cadence sequence. Predicted and observed sky frames are segmented into binary cloud masks, and the projected satellite-path corridor is evaluated using the decision rule defined in Section 7. The example shows how image-sequence nowcasting is translated into interval-wise CFLOS decisions along the projected line of sight. Additional examples are provided in Appendix A.

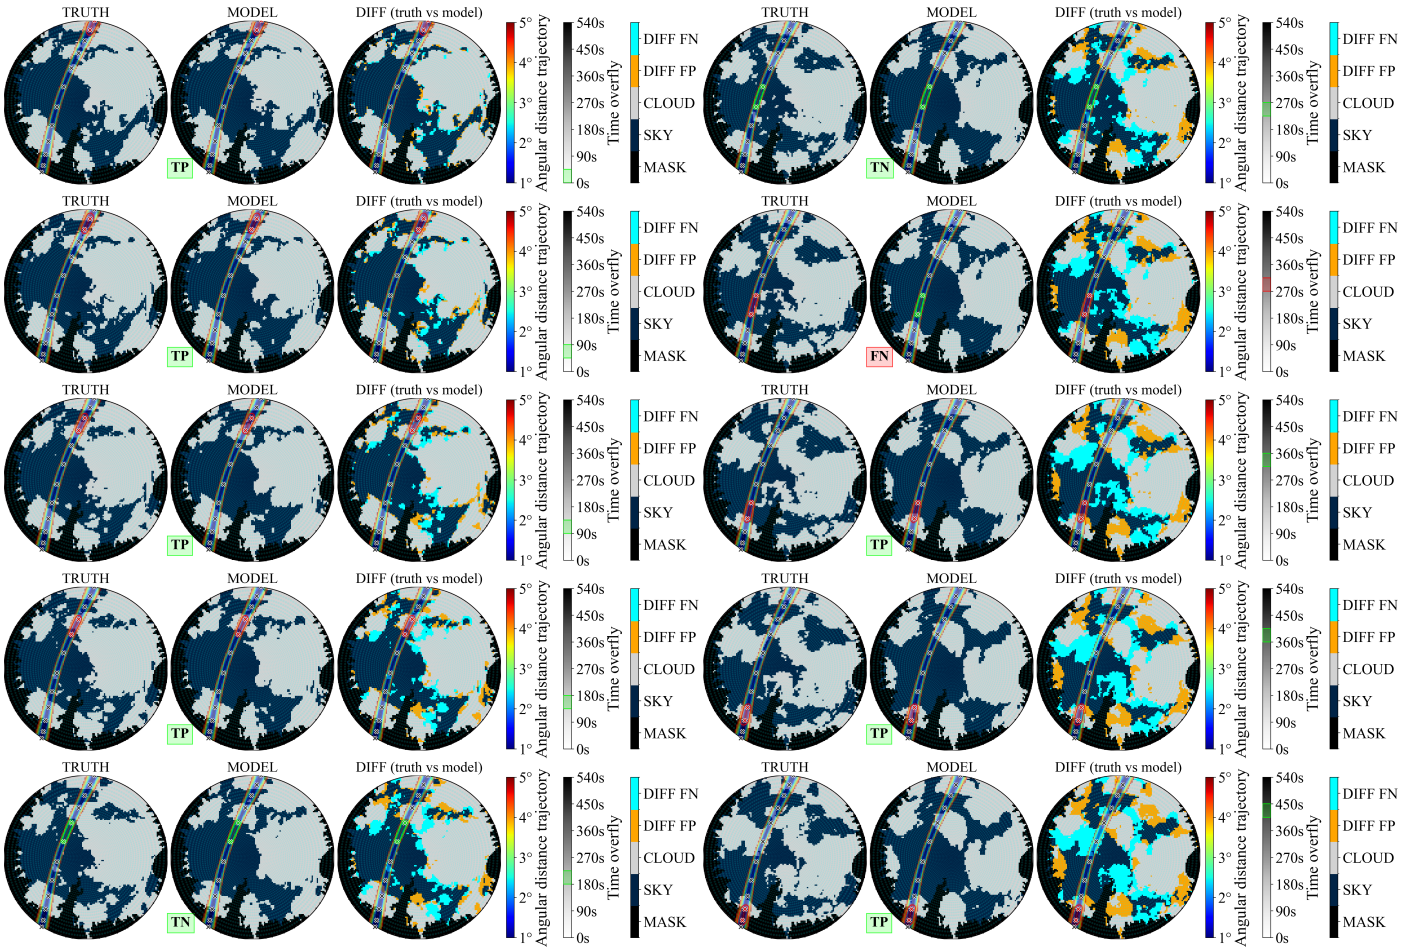


Figure 9. Representative CFLOS evaluation example for a sequence analysed at 45 s cadence. Predicted future sky frames and the matched observed frame are segmented into binary cloud masks and evaluated along the projected satellite-path corridor. The indicated pass time corresponds to forecast lead time.

8.2. Aggregate CFLOS Performance

The aggregate evaluation quantifies how well predicted-mask-based CFLOS decisions reproduce the observed-mask reference decisions at increasing physical lead times. Table 2 summarizes the resulting CFLOS performance. Metrics include F1-score and balanced accuracy, for which higher values are preferable, and the false-negative rate (FNR) and false-positive rate (FPR), for which lower values are preferable. Absolute model performance is complemented by the difference relative to the constant-persistence baseline. The metrics are computed from pooled valid interval-level confusion counts across the contributing cadence subsets.

Table 2. CFLOS performance versus physical lead time. Absolute model performance is reported together with the difference relative to the constant-persistence baseline. Negative ΔFNR and ΔFPR indicate reduced error rates.

Lead Time	Cadences	n	$F1_m$	ΔF1	$BalAcc_m$	ΔBalAcc	FNR_m	ΔFNR	FPR_m	ΔFPR
90 s	15, 30, 45	5220	0.857	0.083	0.865	0.089	0.152	−0.031	0.118	−0.147
180 s	15, 30, 45	2025	0.827	0.070	0.840	0.075	0.185	−0.013	0.135	−0.138
270 s	15, 30, 45	1500	0.816	0.055	0.822	0.061	0.176	−0.035	0.180	−0.105
360 s	30, 45	780	0.804	0.041	0.808	0.047	0.160	−0.030	0.240	−0.080
450 s	30, 45	245	0.790	0.036	0.785	0.038	0.125	−0.035	0.300	−0.060

FNR = unsafe false-clear relative to the observed-frame reference mask (reference blocked → predicted clear). FPR = missed-clear relative to the observed-frame reference mask (reference clear → predicted blocked). Subscript m denotes model. Δ values are computed as model minus persistence. n denotes the number of valid interval-level evaluation samples pooled across all included sequences and cadences for the respective lead time.

The nowcasting-based CFLOS decision shows a consistent positive margin relative to constant persistence across all reported lead times. At 90 s lead time, the model achieves F1 = 0.857 and BalAcc = 0.865, corresponding to gains of ΔF1 = 0.083 and ΔBalAcc = 0.089. Positive margins are maintained up to 450 s, although the magnitude of the improvement decreases with lead time.

The same trend is reflected in the operational error rates. Relative to persistence, the model reduces both the unsafe false-clear rate (FNR) and the missed-clear rate (FPR) at all reported lead times. The largest improvements occur at short lead times, which is consistent with increasing uncertainty from cloud advection and deformation at longer forecast horizons.

8.3. Lead-Time Dependence and Evaluation Context

The decrease in sample count n with increasing lead time results from the finite sequence length and the requirement for valid target intervals. Longer lead times require longer valid sequence support and, therefore, yield fewer eligible samples. In addition, the contributing cadence subsets change at longer lead times, as shown in Table 2.

The reported values should be interpreted in the context of the present reference-track protocol. Interval-level samples are temporally correlated because multiple intervals originate from the same image sequences, and pooling across cadences introduces variations in effective temporal spacing. The results, therefore, describe performance trends within the present binary CFLOS evaluation framework rather than independent estimates of general image-prediction skill.

Overall, the results indicate that short-horizon all-sky nowcasting provides useful additional CFLOS information beyond a stationary persistence assumption. The effect is strongest at short lead times and remains positive over the evaluated range, supporting the use of ASI-based prediction as one component of operational CFLOS assessment and decision support.

9. Limitations and Outlook

The presented framework is subject to several limitations. First, the accuracy of the CFLOS decision depends on the spatial resolution, signal quality, and usable field of view of the all-sky sensor, which may limit detection of small-scale, optically thin, or low-contrast cloud structures. This limitation is particularly relevant near the horizon: cloud structures below or close to the usable lower elevation limit cannot be reliably observed before entering the camera field of view, while low-elevation optical links traverse longer atmospheric slant paths and are, therefore, especially sensitive to such clouds. In addition, horizon regions are strongly compressed in native fisheye imagery. The KTAN reprojection mitigates this foreshortening within the valid angular domain, but it cannot recover information outside the observed field of view. Second, the binary cloud-mask representation does not explicitly distinguish between cloud type, optical thickness, cloud phase, or multi-layer structure. These properties can affect optical attenuation differently; for example, thin cirrus and optically thick low clouds may both be labelled as cloud while having different link impacts. The present CFLOS decision should, therefore, be interpreted as an image-based blockage surrogate rather than as a direct estimate of optical transmittance. Third, segmentation and prediction models may exhibit residual domain shift between training data and operational conditions, for example due to changes in illumination, seasonal cloud regimes, or rare atmospheric morphologies. In the present study, the nowcasting module was applied only to daytime VIS/RGB image sequences; night-time image sequences were not included in the prediction model training or evaluation. Fourth, the current evaluation relies on mask-based CFLOS references derived from image data and is not yet directly coupled to end-to-end link telemetry such as received optical power, fade statistics, or achieved throughput. Finally, the reference-track evaluation protocol intentionally probes the decision logic under controlled geometry rather than representing time-synchronised satellite passes; the reported metrics should, therefore, be interpreted in this methodological context.

Future work will focus on four directions. First, link-to-decision validation is required by pairing predicted CFLOS timelines with operational ground-station logs and satellite pass outcomes. In particular, future validation should relate mask-derived CFLOS states to attenuation-relevant link quantities such as received optical power, fade duration, and link margin. Second, future evaluation should include stronger operational baselines than constant persistence, for example forecasts derived from recent pass history, numerical weather prediction, or satellite-based cloud products. Third, robustness of segmentation and nowcasting models should be improved for challenging observational regimes such as low solar elevation, strong glare, twilight conditions, multi-layer cloud scenes, and optically thin clouds, including explicit uncertainty calibration. Fourth, the nowcasting framework should be extended beyond daytime VIS imagery toward night-time prediction using thermal infrared image sequences. At the system level, the concept will be extended from single-station feasibility assessment to network-level scheduling by combining per-station CFLOS forecasts into coordinated site selection and handover strategies for optical communication networks.

10. Conclusions

This paper presented an integrated all-sky imaging framework for CFLOS-oriented decision support in free-space optical satellite downlinks. By combining station-local hemispheric sensing, cloud-mask generation, short-horizon sky nowcasting, satellite-path projection, and corridor-based CFLOS evaluation, the framework translates image sequences into direction- and lead-time-dependent feasibility information for optical ground-station operation. The main contribution lies in the unified operational formulation of this processing chain rather than in the introduction of a new stand-alone model component.

A central result is that CFLOS should be evaluated along the projected link direction rather than through global cloud statistics. The corridor-based formulation provides a physically interpretable decision variable for scheduling by assessing whether the predicted satellite line of sight remains sufficiently clear within a conservative angular margin.

Under the geometry-controlled reference-track protocol, nowcasting-based CFLOS decisions showed a consistent positive trend relative to a constant-persistence baseline, particularly at short lead times. The model reduced both unsafe false-clear and missed-clear decisions over the evaluated lead-time range, indicating that short-horizon image-sequence prediction can provide useful additional information beyond a stationary persistence assumption within the present mask-based CFLOS framework.

Overall, the results show that all-sky image sequences can be converted into physically interpretable CFLOS guidance within the present reference-track evaluation framework. This provides a practical basis for cloud-aware decision support at single optical ground stations and, more broadly, for future coordinated site-selection and handover strategies in networked optical communication systems.

Author Contributions: Conceptualization, P.M., M.T.K. and B.N.; Methodology, P.M., M.A. and N.B.; Software, P.M., M.A., J.G., H.R. and N.B.; Validation, P.M. and M.A.; Writing—original draft, P.M.; Writing—review and editing, M.A., J.G., H.R., M.T.K., N.B. and B.N. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the ESA Spark Funding Austria initiative, project VISTA, in cooperation with Brimatech Services GmbH. The DLR RSC^3 is supported by the German Federal Ministry of Defense under the technological research and development assignment "Responsive Space Capabilities."

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data and code are available upon request from the corresponding author.

Conflicts of Interest: Authors Paul Matteschk, Max Aragón and Jose Gomez are co-founders of Wematics FlexCo. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest. The authors declare that this study received funding from the cooperation with Brimatech Services GmbH. The funder was not involved in the study design, collection, analysis, interpretation of data, the writing of this article or the decision to submit it for publication.

Appendix A. Additional Nowcasting and CFLOS Evaluation Examples

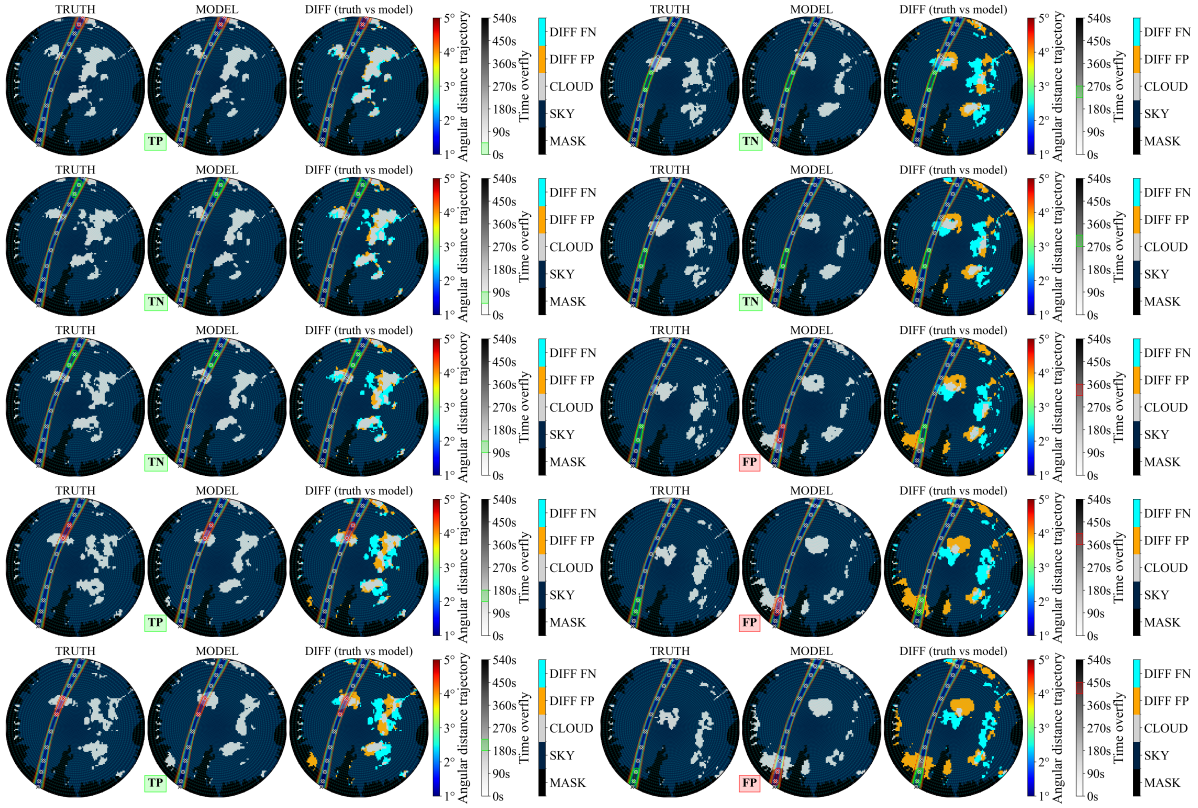


Figure A1. Additional autoregressive prediction rollout with context frames, intermediate predictions, and matched observations. Sequence timestamp: 3 July 2025, 18:54:00 (45 s cadence).

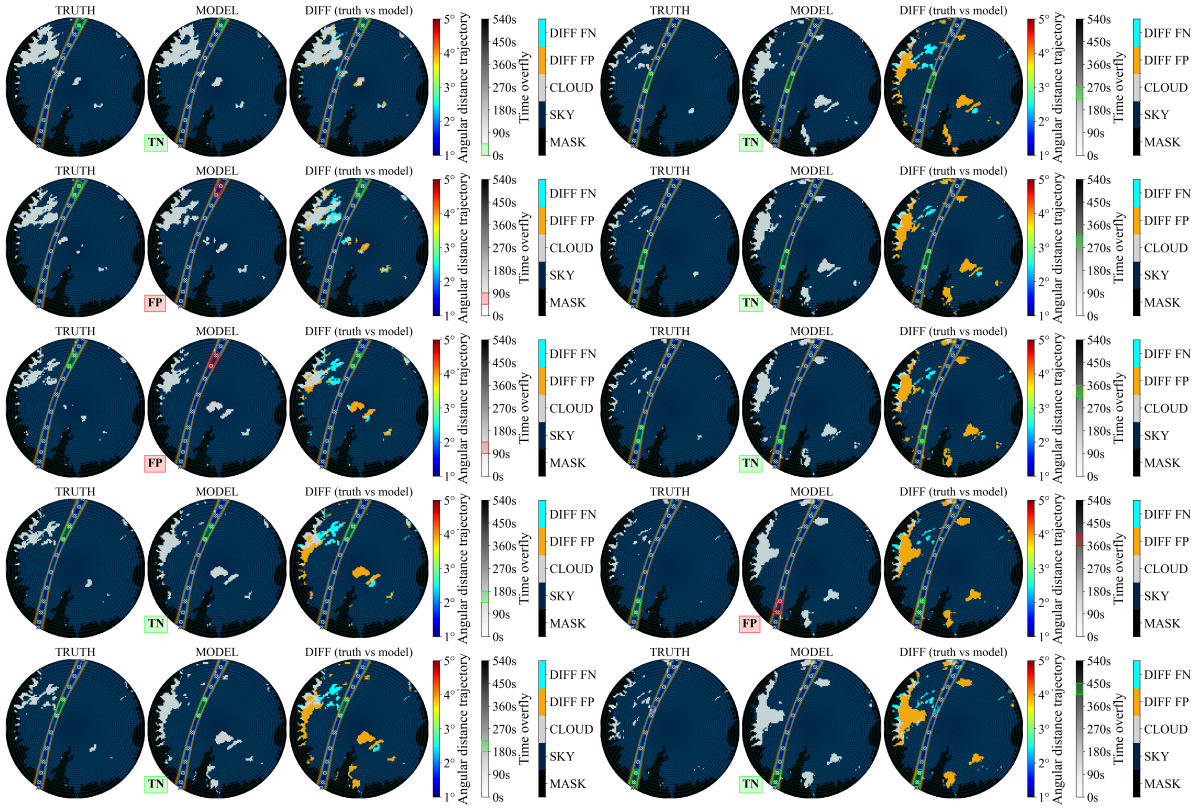


Figure A2. Additional autoregressive prediction rollout with context frames, intermediate predictions, and matched observations. Sequence timestamp: 9 July 2025, 19:21:00 (45 s cadence).

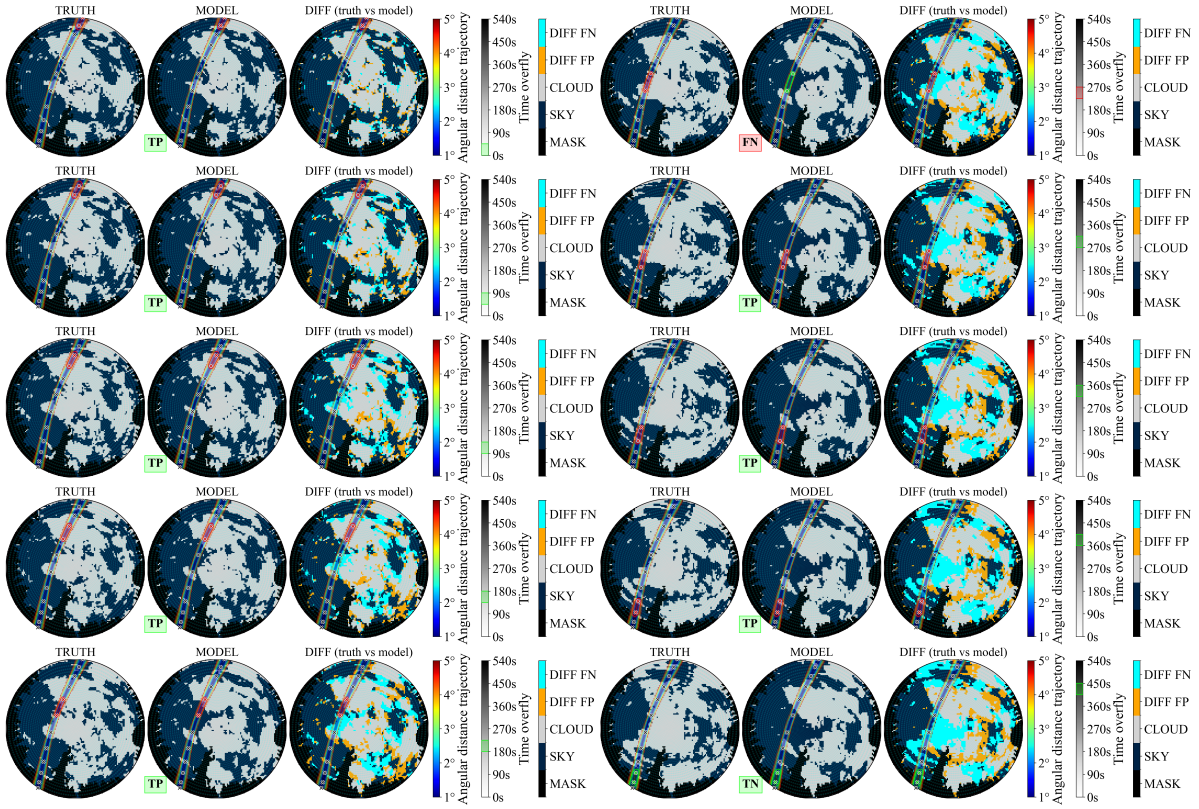


Figure A3. Additional autoregressive prediction rollout with context frames, intermediate predictions, and matched observations. Sequence timestamp: 10 July 2025, 11:27:45 (45 s cadence).

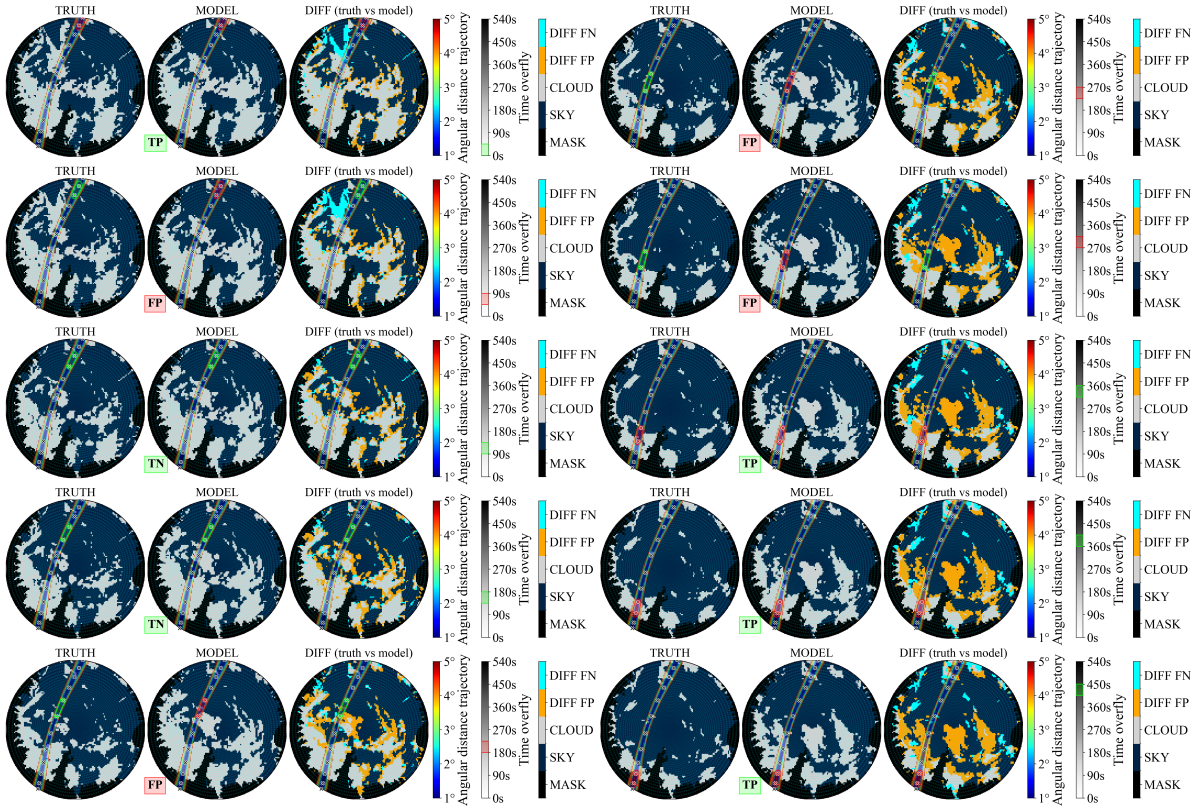


Figure A4. Additional autoregressive prediction rollout with context frames, intermediate predictions, and matched observations. Sequence timestamp: 10 July 2025, 18:32:15 (45 s cadence).

References

1. Fuchs, C.; Moll, F. Ground Station Network Optimization for Space-to-Ground Optical Communication Links. *J. Opt. Commun. Netw.* **2015**, *7*, 1148–1159. [CrossRef]
2. Polnik, M.; Arulselvan, A.; Riccardi, A. Scheduling Space-to-Ground Optical Communication Under Cloud Cover Uncertainty. *IEEE Trans. Aerosp. Electron. Syst.* **2021**, *57*, 2838–2849. [CrossRef]
3. Lyras, N.K.; Kourogiorgas, C.I.; Panagopoulos, A.D. Cloud-Free Line-of-Sight Prediction Modeling for Low Earth Orbit Optical Satellite Networks. *Proc. SPIE* **2019**, *11180*, 111801G. [CrossRef]
4. Poulenard, J.; Cros, S.; Laurens, V. Ground Segment Design for Optical LEO Downlinks. *J. Opt. Commun. Netw.* **2015**, *7*, 325–336. [CrossRef]
5. Bertin, C.; Cros, S.; Saint-Antonin, L.; Schmutz, N. Prediction of Optical Communication Link Availability: Real-Time Observation of Cloud Patterns Using a Ground-Based Thermal Infrared Camera. *Proc. SPIE* **2015**, *9641*, 96410A. [CrossRef]
6. Buntin, S.; Copperwheat, C.M.; Jermak, H.E. Nighttime Cloud Detection, Tracking and Prediction with All-Sky Cameras. *arXiv* **2025**, arXiv:2507.21711. [CrossRef]
7. Fiorino, S.T.; Cronin, A.D.; Sizemore, S.P. Implications of Four-Dimensional Weather Cubes for Improved Cloud-Free Line-of-Sight Estimation. *Opt. Eng.* **2020**, *59*, 081808. [CrossRef]
8. Birch, M.; Beattie, J.R.; Bennet, F.; Rattenbury, N.; Copeland, M.; Travouillon, T.; Ferguson, K.; Cater, J.; Sayat, M. Availability, Outage, and Capacity of Spatially Correlated, Australasian Free-Space Optical Networks. *J. Opt. Commun. Netw.* **2023**, *15*, 415–430. [CrossRef]
9. Rattenbury, N.J.; Ashby, J.; Bennet, F.; Birch, M.; Cater, J.E.; Ferguson, K.; Giggenbach, D.; Grant, K.; Knopp, A.; Knopp, M.T.; et al. Update on the German and Australasian Optical Ground Station Networks. *arXiv* **2024**, arXiv:2402.13282. [CrossRef]
10. Rotherham, E.; Casey, C.; Fernandez Rodriguez, E.; Vidaurre Torrez, K.W.; Mashor, M.; Pike, I. Advancing Free-Space Optical Communication System Architecture: Performance Analysis of Varied Optical Ground Station Network Configurations. *arXiv* **2024**, arXiv:2410.23470. [CrossRef]
11. Le Son, H.; Schwarz, R.T.; Knopp, M.T.; Giggenbach, M.; Nouri, B.; Giggenbach, D.; Knopp, A. Optical Feeder Links to GEO Satellites: Statistical Analysis of Link Availability Using Deep Learning-Based Cloud Segmentation Data. In Proceedings of the 2025 IEEE International Conference on Communications Workshops (ICC Workshops), Montreal, QC, Canada, 8–12 June 2025; pp. 1158–1163. [CrossRef]
12. Aldasoro Marculeta, I. AllSky-Camera System for Monitoring of Optical Satellite Downlinks. Bachelor's Thesis, Universidad Publica de Navarra, Navarra, Spain, 2024.
13. Schmidt, T.; Stührenberg, J.; Blum, N.; Lezaca, J.; Hammer, A.; Wilbert, S.; Nouri, B.; Schroedter-Homscheidt, M.; Heinemann, D.; Vogt, T. Eye2Sky: A Network of All-Sky Imager and Meteorological Measurement Stations for High-Resolution Nowcasting of Solar Irradiance. *Meteorol. Z.* **2025**, *34*, 35–55. [CrossRef]
14. Knopp, M.T.; Giggenbach, D.; Giggenbach, M.; Nouri, B. Assessment of CIEMAT's Plataforma Solar de Almeria as a Ground Station Site for Optical LEO Satellite Downlinks. *Proc. SPIE* **2025**, *13699*, 1067–1080.
15. Sun, Y.; Szocs, G.; Brandt, A.R. Solar PV output prediction from video streams using convolutional neural networks. *Energy Environ. Sci.* **2018**, *11*, 1811–1818. [CrossRef]
16. Paletta, Q.; Arbod, G.; Lasenby, J. Benchmarking of deep learning irradiance forecasting models from sky images: An in-depth analysis. *Sol. Energy* **2021**, *224*, 855–867. [CrossRef]
17. Fabel, Y.; Nouri, B.; Wilbert, S.; Blum, N.; Schnaus, D.; Triebel, R.; Zarzalejo, L.F.; Ugedo, E.; Kowalski, J.; Pitz-Paal, R. Combining Deep Learning and Physical Models: A Benchmark Study on All-Sky Imager-Based Solar Nowcasting Systems. *Sol. RRL* **2024**, *8*, 2300808. [CrossRef]
18. Nie, Y.; Zelikman, E.; Scott, A.; Paletta, Q.; Brandt, A.R. SkyGPT: Probabilistic ultra-short-term solar forecasting using synthetic sky images from physics-constrained VideoGPT. *Adv. Appl. Energy* **2024**, *15*, 100172. [CrossRef]
19. Le Guen, V.; Thome, N. Disentangling physical dynamics from unknown factors for unsupervised video prediction. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 13–19 June 2020; pp. 1112–1121. [CrossRef]
20. Yan, W.; Zhang, Y.; Abbeel, P.; Srinivas, A. VideoGPT: Video generation using VQ-VAE and transformers. *arXiv* **2021**, arXiv:2104.10157. [CrossRef]
21. Peebles, W.; Xie, S. Scalable diffusion models with transformers. *arXiv* **2022**, arXiv:2212.09748. [CrossRef]
22. Blattmann, A.; Dockhorn, T.; Kulal, S.; Mendelevitch, D.; Kilian, M.; Lorenz, D.; Levi, Y.; English, Z.; Voleti, V.; Letts, A.; et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv* **2023**, arXiv:2311.15127. [CrossRef]
23. Fabel, Y.; Nouri, B.; Wilbert, S.; Blum, N.; Triebel, R.; Hasenbalg, M.; Kuhn, P.; Zarzalejo, L.F.; Pitz-Paal, R. Applying self-supervised learning for semantic cloud segmentation of all-sky images. *Atmos. Meas. Tech.* **2022**, *15*, 797–809. [CrossRef]
24. Hasenbalg, M.; Kuhn, P.M.; Wilbert, S.; Nouri, B.; Kazantzidis, A. Benchmarking of six cloud segmentation algorithms for ground-based all-sky imagers. *Sol. Energy* **2020**, *201*, 596–614. [CrossRef]

25. Ye, L.; Cao, Z.; Xiao, Y.; Yang, Z. Supervised fine-grained cloud detection and recognition in whole-sky images. *IEEE Trans. Geosci. Remote. Sens.* **2019**, *57*, 7972–7985. [CrossRef]
26. Ye, L.; Wang, Y.; Cao, Z.; Yang, Z.; Min, H. A self training mechanism with scanty and incompletely annotated samples for learning-based cloud detection in whole sky images. *Earth Space Sci.* **2022**, *9*, e2022EA002220. [CrossRef]
27. Xie, W.; Liu, D.; Yang, M.; Chen, S.; Wang, B.; Wang, Z.; Xia, Y.; Liu, Y.; Wang, Y.; Zhang, C. SegCloud: A novel cloud image segmentation model using a deep convolutional neural network for ground-based all-sky-view camera observation. *Atmos. Meas. Tech.* **2020**, *13*, 1953–1961. [CrossRef]
28. Köhler, A.; Knopp, M.T.; Ohndorf, A. Setup of the Optical Ground Station In Trauen For Optical Free-Space Communication. In Proceedings of the Deutscher Luft- u. Raumfahrtkongress 2022, Dresden, Germany, 27–29 October 2022. [CrossRef]
29. Antuña-Sánchez, J.C.; Román, R.; Bosch, J.L.; Toledano, C.; Mateos, D.; González, R.; Cachorro, V.; de Frutos, Á. ORION software tool for the geometrical calibration of all-sky cameras. *PLoS ONE* **2022**, *17*, e0265959. [CrossRef]
30. Blum, N.; Matteschk, P.; Fabel, Y.; Nouri, B.; Román, R.; Zarzalejo, L.F.; Antuña-Sánchez, J.C.; Wilbert, S. Geometric calibration of all-sky cameras using sun and moon positions: A comprehensive analysis. *Sol. Energy* **2025**, *295*, 113476. [CrossRef]
31. Stellarium Contributors. Stellarium v26.1 Astronomy Software. *Zenodo* **2026**. [CrossRef]
32. Vallado, D.A.; Crawford, P.; Hujsak, R.; Kelso, T.S. Revisiting Spacetrack Report #3. In Proceedings of the AIAA/AAS Astrodynamics Specialist Conference and Exhibit, Keystone, CO, USA, 21–24 August 2006. [CrossRef]
33. Petit, G.; Luzum, B. (Eds.) *IERS Conventions (2010)*; IERS Technical Note; Verlag des Bundesamts für Kartographie und Geodäsie: Frankfurt am Main, Germany, 2010; Volume 36.
34. Maisonobe, L.; Pommier-Maurussane, V.; 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), ESAC, Madrid, Spain, 3–6 May 2010.
35. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Proceedings of the Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015*; Lecture Notes in Computer Science; Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F., Eds.; Springer: Cham, Switzerland, 2015; Volume 9351, pp. 234–241. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.