

Automated OGS Laser Alignment using a Motorized Collimator Mount and Retroreflector Feedback

Mateusz Kucharski* Petros Karafillis Pierre-Alexis Lagadrillière

Department of Communication and Ground Stations,
German Space Operations Center (GSOC), German Aerospace Center (DLR),
Oberpfaffenhofen, Germany

Abstract

Precise laser beam alignment is a prerequisite for reliable operation of telescope-based optical ground stations (OGS) in direct-to-Earth optical communication scenarios. As such systems mature from experimental setups toward operational infrastructure, the limitations of manual alignment procedures become increasingly apparent: they are difficult to standardize, operator-dependent, and offer insufficient traceability across calibration sessions.

Within the GlobeON project, the DLR German Space Operations Centre (GSOC) is developing a software tool to address these shortcomings for the telescope facility in Almería, Spain. The calibration range is a configuration input rather than a design constant: at the operational site the retroreflector stands on Sierra Alhamilla, 11 787 m from the OGS, while the rooftop test arrangement used for the campaign reported here spans 6550 m, and the scan geometry is derived from whichever range is configured. The goal of the calibration procedure is to identify the laser pointing position that maximizes the reflected signal intensity, corresponding to a direct hit of the retroreflector. The tool consolidates this workflow into a single, structured application, replacing ad hoc procedures with a repeatable, documented process.

The software controls a motorized laser mount, steering the beam in two axes across predefined scan patterns to systematically search for the optimal alignment position. A live camera feed allows the operator to monitor the beam and define regions of interest, while automated spot detection analyses each frame through thresholding and centroid estimation, with background subtraction and shape-based filtering. Following a coarse area scan, a local refinement module iteratively explores the neighbourhood of the best candidate position and converges toward the signal maximum. A geometric compensation corrects for near/far field geometry and laser-camera baseline offsets, ensuring the aiming reference accurately reflects far-field target conditions. Each session is logged with images, per-point scan results, and derived coordinates, providing a traceable record for post-session review and comparison.

The tool has been demonstrated end-to-end on a 6550 m outdoor path at the GSOC rooftop test facility, executing a 599-point coarse scan autonomously in 20.8 min and acquiring the target after 38.7 % of that time. This paper presents the system concept, the implementation approach, the results of that validation campaign, and the detector configuration constraints that bound the attainable pointing resolution, which are identified here and carried into the deployment plan.

Keywords: DLR, GSOC, OGS, Optical Communications, GlobeON, Laser alignment

Acronyms / Abbreviations

ADU	Analogue-to-Digital Unit
CCSDS	Consultative Committee for Space Data Systems
CIEMAT	Centre for Energy, Environmental and Technological Research
DTE	Direct-To-Earth
EONN	European Optical Nucleus Network
FOGATA	Free-space Optical Ground Antenna Tabernas
GlobeON	Global Optical Ground Station Network
GSOC	German Space Operations Center
GUI	Graphical User Interface
IFOV	Instantaneous Field of View
LEO	Low Earth Orbit
MOR	Meteorological Optical Range
O3K	Optical On-Off Keying
OGS	Optical Ground Station
ROI	Region of Interest
SNR	Signal-to-Noise Ratio
SOFA	Small Optical Ground Stations Focal-Optics Assembly

1 Introduction

Optical ground stations are the terrestrial endpoint of free-space optical links to airborne and space-based terminals, offering data rates and spectral-allocation freedom that radio-frequency systems cannot match. The capability has been demonstrated over ranges from low Earth orbit to lunar distance [11], and small-satellite optical terminals are now flown as an established payload class [10]. Unlike RF stations, however, an OGS depends on the precise co-alignment of two optical paths: the receive path defined by the telescope line of sight, and the transmit path defined

*Corresponding author: mateusz.kucharski@dlr.de

by the uplink laser beam emerging from a collimator mounted on the telescope structure. Because the uplink divergence is typically on the order of a few hundred microradians, a residual co-alignment error of comparable magnitude translates directly into a loss of received power at the space terminal, and in the worst case into a complete failure to close the link. Static transmitter misalignment is one of the contributions to uplink pointing loss whose effect on received intensity has been quantified for LEO geometries [8], and it is the contribution that ground-segment calibration can remove.

The operational relevance of this problem has grown with the transition of optical ground segments from single experimental installations to coordinated networks. DLR manages the integration of several stations into a German optical ground station network and participates in the European Optical Nucleus Network together with ESA and KSAT [12, 5]. Within this landscape, GSOC is building the Free-space Optical Ground Antenna Tabernas (FOGATA) at the Plataforma Solar de Almería, a research site operated by the Spanish Centre for Energy, Environmental and Technological Research (CIEMAT) [5], as the first terminal of the Global Optical Ground Station Network (GlobeON) [6], using the Small Optical Ground Station Focal-Optics Assembly (SOFA) as communications equipment [3]. Stations operated as network assets rather than as experiments must deliver predictable performance across sessions and operators, which places calibration procedures under requirements they were not originally designed to meet.

Maintaining uplink co-alignment is not a one-time calibration. Thermal gradients across the telescope structure, mechanical settling after slews, and re-installation of the collimator between campaigns all introduce drift, so alignment must be re-established regularly, in practice before each measurement session. It is done against a retroreflector at a known bearing, which returns a fraction of the transmitted beam to the station and provides a direct, end-to-end observable of the transmit path. The range to that target differs between installations, and the procedure is formulated so that it does not depend on any particular value.

In the established procedure, an operator adjusts the collimator mount manually while observing the retroreflected return, iterating until the returned signal is judged to be maximised. This approach is workable but has three practical drawbacks. First, it is *slow*: a full alignment consumes a significant share of the time available in a measurement window, and this cost is paid again after every disturbance. Second, it is *operator-dependent*: the stopping criterion is a subjective judgement, so the achieved alignment quality varies between operators and between runs. Third, it is *poorly documented*: the resulting mount position is recorded, but the search that produced it is not, which makes it impossible to reconstruct after the fact why a given session performed as it did.

This paper presents a software tool that automates the procedure by closing the loop between the motorized collimator mount and the retroreflector return. The tool executes a two-stage search, a coarse scan to acquire the return signal over the full uncertainty region followed by a fine scan to refine the estimate, and determines the optimum from a centroid estimate of the measured intensity distribution rather than from a single peak sample. A geometric compensation step accounts for the difference between the near-field calibration geometry and the far-field operating geometry of an actual space link. Every run is written to a session log that records the raw measurements, the derived estimate, the environmental conditions, and the Git commit hash of the software version used, giving full traceability from a reported result back to the code that produced it.

The contributions of this work are:

- an automated, closed-loop alignment procedure for an OGS uplink collimator using retroreflector feedback, requiring no operator intervention between initiation and result;
- an end-to-end validation of that procedure on a live outdoor path, characterising the scan cadence and its evolution along the pattern, the detection yield, the acquisition timing margin and the internal consistency of the session record;
- an analysis of the two configuration effects that currently bound the attainable pointing resolution, detector saturation and beam-footprint scaling, together with the mitigations they imply;
- a reproducible implementation in which each result is bound to a logged software version and to the environmental conditions under which it was obtained.

Section 2 describes the OGS and the alignment hardware. Section 3 details the scan strategy, centroid estimation, and geometric compensation. Section 4 defines the experimental setup. Section 5 presents the results of the validation campaign, Section 6 discusses their implications and limitations, and Section 7 concludes.

1.1 Relation to Established Methods

The scoring architecture presented in Section 3.2 draws on established techniques from laser beam profilometry and adaptive optics rather than proposing new estimators.

Threshold detection with an intensity-weighted centroid is a simplified form of the second-moment method that underlies ISO 11146 for determining laser beam position and width [14]. The background handling, an exponential moving average model with percentile thresholding, corresponds to standard background suppression in optical trackers and beam

pointing and tracking systems. The weighted combination of peak intensity, spot area and raw score used as a confidence measure serves a role analogous to the SNR proxies employed in Shack–Hartmann wavefront sensors and in FSO spot-tracking systems, where the decision to accept a measurement rests on several correlated image features rather than on amplitude alone [15].

The seeing index, which combines positional jitter and intensity scintillation over a temporal window, corresponds conceptually to the parameters used to characterise atmospheric turbulence in optical communications and astronomy, such as the Rytov variance and the Fried parameter r_0 [16], with the difference that the present system uses a simplified empirical heuristic in place of a rigorous statistical model of propagation through a turbulent medium.

Two further choices are deliberate computational compromises. Trajectory smoothing uses an exponential moving average rather than a Kalman filter, the standard estimator in target tracking; and local optimisation uses a deterministic diminishing-step hill climb rather than a gradient method of the extremum-seeking-control family typical of fibre beam autoalignment systems [17]. Both alternatives are cheaper and deterministic, which matters in a real-time control loop driving stepper actuators, at the cost of the statistical optimality the more advanced approaches offer.

Closer to the present problem, transmitter beam bias has been verified at DLR by recording the intensity distribution of a downlink beam on the ground with the power detectors of three ground stations in parallel and deriving the instantaneous centre of gravity of the beam spot [7]. That work and this one share an objective, determining and removing a static transmitter pointing bias, but differ in direction and in instrument: there the beam originates on the satellite and is sampled by a distributed receiver array, whereas here it originates at the ground station and is sampled by scanning a single co-aligned camera against a cooperative passive target. The two address complementary halves of the same calibration problem.

The contribution of this work is therefore not a new estimator but the integration of established methods into an operational calibration procedure with end-to-end traceability, for a class of ground station making the transition from experimental installation to network asset.

2 System Overview

This section describes the physical calibration infrastructure and the high-level software architecture of the automated OGS laser alignment tool developed at GSOC.

2.1 Calibration Geometry

The calibration arrangement consists of two endpoints: the station, carrying the collimator and a co-aligned camera, and a corner-cube retroreflector at a known approximate bearing. A retroreflector is advantageous as a target because it returns incident light directly toward its source with minimal angular deviation, providing a passive and well-defined alignment reference.

The separation between the two endpoints is a configuration input rather than a fixed property of the system. The scan planner takes the range as a parameter and derives from it the angular extent of the search pattern, the pitch between pointing positions and the expected footprint of the beam in the target plane, so the same procedure applies wherever a co-operative target is available. Two such arrangements are used in this work, both illustrated in Fig. 1: the operational one at the FOGATA telescope facility on the Plataforma Solar de Almería, where the retroreflector is installed on Sierra Alhamilla, a mountain range facing the site, at 11 787 m [1]; and a shorter arrangement at the GSOC rooftop test facility in Oberpfaffenhofen, at 6550 m. The tool has been configured and operated at both. The quantitative results of Section 5 are drawn from the Oberpfaffenhofen arrangement, for which a complete session record was retained.

Ranges of this order are long enough that the alignment obtained is representative of operational conditions in beam divergence and atmospheric path, while remaining short enough for a passive target to return a usable signal. The consequences of operating at one range rather than the other are quantified in Section 6.2.

2.2 Hardware Components

The hardware subsystem at the OGS comprises three primary components. The laser source is coupled to a motorized collimator mount capable of two-axis angular steering. The mount is driven by piezo-inertial actuators commanded over a network interface and provides fine angular positioning with a resolution sufficient to resolve the retroreflector footprint at the calibration distance.

A co-aligned camera is mechanically fixed relative to the laser output aperture. It captures the near-field beam spot and, during calibration runs, records the reflected signal returning from the retroreflector. The fixed geometric relationship between laser and camera introduces a parallax offset that must be compensated in the pointing model, as described in Section 3.9.

2.3 Software Architecture

The calibration tool is implemented as a single structured application integrating five software modules: (i) a hardware control layer for motor commanding




Figure 1: Calibration geometry. The uplink beam emitted from the motorised collimator mount illuminates a corner-cube retroreflector at the configured range D , which returns a fraction of the incident light along the reciprocal path to the co-aligned camera. The range is a parameter of the scan planner, not a fixed property of the system: the same arrangement is used at 6550 m at the Oberpfaffenhofen rooftop facility and at 11787 m against the Sierra Alhamilla retroreflector at Almería. The inset shows what changes with range: at a fixed 5 mrad divergence the illuminated footprint grows from 32.8 m to 58.9 m, while the retroreflector remains an unresolved point within it.

and camera acquisition, (ii) a scan pattern generator for coarse and fine search strategies, (iii) a spot detection and signal estimation module, (iv) a geometric compensation module, and (v) a session logging and data management backend. The inter-module data flow is depicted in Fig. 2.

3 Methodology

The calibration procedure consists of two sequential phases: a coarse area scan that systematically surveys the angular search space, followed by a fine-scan refinement that converges to the global signal maximum. The complete pipeline is illustrated in Fig. 3.

3.1 Coarse Area Scan

The coarse scan phase commands the motorized mount to traverse a predefined two-dimensional pattern about a nominal pointing position derived from the known bearing of the target and the configured range. Three patterns are implemented. A *snake* pattern raster-scans a rectangular region in alternating rows. A *spiral* pattern develops outward from the nominal position along an Archimedean spiral, holding the radial pitch constant so that sampling density is uniform in range from the pattern centre, at the price of a commanded direction that changes at every step. A *hexagonal* pattern traverses concentric hexagonal rings; because the commanded direction changes only at the six vertices of each ring, it reduces the number of direction reversals, which is relevant for piezo-inertial actuators exhibiting direction-dependent hysteresis, at the cost of a radial pitch that is larger toward the vertices than toward the edge midpoints. The spiral and hexagonal patterns are shown in Fig. 4. All three are parameterised by target range, pitch and radial extent.

The spiral was used for the campaign reported in Section 5, and has been the pattern of choice in prac-

tice: developing outward from the nominal position reaches the most probable region of the search space first, so acquisition typically occurs well before the pattern completes. At each pattern point, one or more camera frames are acquired and processed by the detection module. The reflected signal intensity is recorded as a scalar value at each pointing position, forming a two-dimensional signal map of the search space.

The spatial resolution of the coarse scan, set by the angular step size, is chosen to exceed the retroreflector angular extent as seen from the OGS, ensuring that the retroreflector is sampled at least once during the survey. The output of the coarse phase is the pointing position with the highest observed signal, which serves as the seed for the refinement stage.

3.2 Detection and Frame-Level Scoring

The system uses a layered scoring architecture in which each layer addresses a distinct problem: (i) scoring a single detection within one frame, (ii) estimating the reliability of that detection, (iii) assessing signal stability in the time domain, and (iv) aggregating multiple measurements into the single value that represents a scan point. The layers form a processing chain: the raw score feeds the confidence measure, the detection history feeds the seeing index, and the aggregated scan-point value drives the higher-level decisions of the scanner.

The input image I is thresholded relative to its own dynamic range within the region of interest:

$$thr = v_{\min} + \tau(v_{\max} - v_{\min}), \quad (1)$$

where $v_{\min}$ and $v_{\max}$ are the minimum and maximum intensity in the ROI and $\tau \in [0, 1]$ is the relative threshold (default $\tau = 0.6$; the tracking module adapts it within $[0.45, 0.85]$ according to frame contrast, measured as the difference between the 95th and 50th intensity percentiles).




Figure 2: Closed-loop architecture. The scan engine commands angular positions through the motion controller to the piezo-inertial actuators driving the two-axis collimator mount, which steers the collimated beam toward the retroreflector. The co-aligned camera images the return, and the detection and scoring stage of Section 3.2 closes the loop by returning the scan-point score, spot centroid and confidence to the scan engine. The range D to the target enters as a configuration parameter of the scan engine rather than as a property of any block in the loop, which is what allows the same architecture to be used unchanged at both arrangements of Fig. 1. The source wavelength shown is the one used in the reported campaign.

The spot centroid is the intensity-weighted centre of mass of the pixels satisfying the threshold condition:

$$c_x = \frac{\sum_i I_i x_i}{\sum_i I_i}, \quad c_y = \frac{\sum_i I_i y_i}{\sum_i I_i}. \quad (2)$$

The raw score of a single detection is defined as

$$S = v_{\max} + 0.01 A, \quad (3)$$

where A is the number of pixels in the threshold mask. Peak intensity is the dominant term; the area term acts as a tie-breaker between detections of comparable peak intensity. The behaviour of Equation (3) when $v_{\max}$ reaches the sensor ceiling is central to the interpretation of the campaign results and is analysed in Section 5.5.

3.3 Detection Confidence

To limit the influence of false detections (noise, artifacts, specular reflections) on later processing stages, a normalised confidence measure is defined as a weighted sum of three components:

$$C = 0.55 P_n + 0.25 A_n + 0.20 S_n, \quad (4)$$

with each component normalised to $[0, 1]$:

$$P_n = \text{clip}\left(\frac{\text{peak} - v_{\min}}{\max(1, v_{\max} - v_{\min})}, 0, 1\right), \quad (5)$$

$$A_n = \text{clip}\left(\frac{A}{250}, 0, 1\right), \quad S_n = \text{clip}\left(\frac{S}{255}, 0, 1\right). \quad (6)$$

A notable property of Equation (5) is that peak intensity is normalised against the dynamic range of the whole frame rather than the ROI alone, so the confidence measure reflects spot contrast against the global scene background. C gates the transitions of the tracking state machine (searching $\rightarrow$ confirming $\rightarrow$ tracking $\rightarrow$ signal loss); a detection is rejected if C falls below $C_{\min} = 0.35$.




Figure 3: Calibration pipeline. Colour indicates coverage by the validation campaign of Section 5: the acquisition and scoring path was executed in full, while the refinement and geometric compensation stages were not triggered during the reported session.




Figure 4: Two of the three coarse-scan patterns implemented by the tool, plotted by the scan planner in target-plane coordinates with the nominal pointing position marked at the origin; a rectangular snake raster is also available. The spiral trades a constant radial pitch for a continuously varying commanded direction, whereas the hexagonal pattern traverses straight segments at the cost of a radial pitch that varies between vertex and edge. The extents shown here are illustrative and larger than those of the reported session (Table 2), which used the spiral.

3.4 Temporal Signal Quality

To assess signal stability in the time domain, analogous to seeing conditions in optical observation, a rolling history of peak position and intensity is maintained over a window of $N = 24$ frames. Positional jitter is defined as

$$J = \sqrt{\text{Var}(x) + \text{Var}(y)}, \quad (7)$$

and intensity scintillation as the coefficient of variation of the peak,

$$\Sigma = \frac{\sigma(\text{peak})}{\mu(\text{peak})}. \quad (8)$$

After normalisation, $J_n = \min(1, J/5.0)$ and $\Sigma_n = \min(1, \Sigma/0.30)$, the seeing index is

$$Q_{\text{seeing}} = \text{clip}(1 - 0.55 J_n - 0.45 \Sigma_n, 0, 1), \quad (9)$$

where $Q_{\text{seeing}} = 1$ corresponds to a perfectly stable, non-scintillating spot and $Q_{\text{seeing}} = 0$ to strong positional or intensity fluctuation.

3.5 Scan-Point Aggregation

At each scan grid point a series of T temporal frames is acquired. If the number of positive detections in the series meets the temporal voting threshold $T_{\min}$, a mechanism that suppresses isolated false positives, the scan-point score is the mean of the raw scores of the frames in which a signal was detected:

$$S_{\text{point}} = \frac{1}{|F|} \sum_{i \in F} S_i, \quad F = \{i : \text{detection}_i = \text{true}\}. \quad (10)$$

If no consensus is reached, the maximum across all frames of the series is taken and the point is simultaneously flagged as not detected. S_{point} drives the higher-level scanner decisions: selection of the best point so far, the early-stop criterion for acquisition mode ($S_{\text{point}} \geq S_{\text{stop}}$), stagnation detection (no improvement by a defined minimum over a defined number of consecutive points while remaining below the signal threshold), and triggering of local position re-optimisation.

3.6 Local Position Optimisation

When a scan-point score fails to reach the configured target value, a deterministic hill-climbing procedure is initiated around the current position. It samples the raw score S in the four axis directions ($\pm x, \pm y$) with an initial step that is halved whenever no improvement is found, terminating at a minimum step size, an iteration limit, or the target score.

3.7 Detection Gating and Background Handling

A background frame acquired with the laser blocked may be subtracted from the measurement frame to remove ambient illumination and fixed-pattern sensor noise; the background model uses an exponential moving average with percentile thresholding. Connected regions may additionally be evaluated against expected spot morphology (circularity, aspect ratio, minimum area). Optionally, an expected-region gate restricts scoring to a predefined image area, and a region of interest restricts acquisition itself.

The expected-region gate and the ROI restriction were both disabled during the validation campaign

reported in Section 5; the consequences are analysed in Sections 5.5 and 6.3.

3.8 Fine-Scan Refinement

Following the coarse scan, a local refinement module explores the neighbourhood of the best candidate pointing position to locate the signal maximum with higher angular precision. The refinement algorithm performs an iterative search, evaluating candidate positions in a contracting pattern and updating the current best position after each evaluation.

The convergence criterion is defined as a combination of the positional change between successive iterations falling below a threshold and the signal improvement dropping below a minimum increment. The final output of the refinement stage is the motor position corresponding to the estimated signal maximum, expressed in the motor coordinate frame.

3.9 Geometric Compensation

The motor coordinate system defines pointing in terms of the laser aperture axis. However, the retro-reflector return signal is detected by the camera, whose optical axis is displaced from the laser axis by the laser-camera baseline offset. At near-field distances this parallax produces a non-negligible pointing error if uncorrected. Additionally, the near-field beam divergence geometry causes the intensity distribution on the camera to differ from the far-field condition encountered during actual OGS operations.

The geometric compensation module applies an affine correction to the motor-to-camera pointing mapping, derived from a pre-characterized model of the laser-camera geometry. This ensures that the aiming reference output by the calibration procedure corresponds to the correct far-field pointing angle, independent of the calibration distance.




Figure 5: Near-field / far-field geometric compensation. The camera axis is displaced from the laser axis by the baseline b , so the apparent direction of the retroreflected return differs from the true transmit direction by a parallax angle that scales inversely with range. The compensation module removes this term so that the aiming reference produced at the calibration distance is valid for the far-field operating condition.

3.10 Session Logging

Every calibration run is fully documented in a structured log. The log captures timestamp and operator

identification; hardware configuration; software version (Git commit hash); all hyperparameter settings (step size, ROI boundaries, detection thresholds, convergence tolerances); intermediate per-point scan results including motor position, centroid coordinates, and signal intensity; and the final derived calibration result.

Camera frames are archived for each accepted scan point, allowing retrospective review of detection quality. Every quantitative statement in Section 5 was reconstructed from this record after the fact, which is the traceability property that manual alignment does not provide.

4 Experimental Setup

4.1 Test Site and Geometry

The validation campaign was conducted at the GSOC rooftop test facility at the DLR site in Oberpfaffenhofen, Germany. This facility hosts engineering models of station equipment intended for the remote OGS sites and is integrated into the operational network and control environment at DLR, which makes it the natural venue for subsystem-level verification prior to deployment [5].

The tool was configured for a propagation path of 6550 m between the collimator mount and a corner-cube retroreflector, and the scan geometry of Table 2 was generated from that value. The range is 56 % of the one configured for the operational site, which is what makes the pair useful: the same procedure was exercised at two ranges differing by a factor of 1.80, and the quantities that scale with range can therefore be separated from those that do not. The path is near-horizontal and runs at low altitude above built-up terrain for its full length, which places it in the regime of strongest refractive-index structure; a slant path to an orbiting terminal traverses the dense boundary layer over a far shorter fraction of its length. The shorter arrangement is therefore not a benign approximation of the longer one; the range-dependent effects are addressed in Section 6.2.

The campaign was conducted at night under clear conditions. Ambient illumination in the field of view was negligible, which is relevant to the interpretation of both the detection yield and the identification of the return signal in Section 5.

4.2 Hardware Configuration

The hardware configuration is summarised in Table 1. The angular scale factors were determined empirically for this installation and cross-checked against the recorded per-point data.

Table 1: Hardware configuration for the Oberpfaffenhofen validation campaign.

Item	Specification
Source	Fibre-coupled diode, 635 nm, < 5 mW, Class 3R
Beam divergence	5.0 mrad full angle (scan model); 32.8 m spot at 6550 m
Mount	Two-axis kinematic collimator mount
Actuators	Newport 8816 piezo-inertial (picomotor)
Controller	4-axis picomotor controller / driver, Ethernet / USB / RS-485, 12 V
Angular step	0.574 μ rad (X), 0.513 μ rad (Y)
Step at target	3.76 mm (X), 3.36 mm (Y) at 6550 m
Camera	The Imaging Source DMK 33GX265, GigE, monochrome
Sensor	Sony IMX265, 2048 $\times$ 1536 px, 3.45 μ m pitch
Lens	Fujinon HF35XA-5M, $f = 35$ mm
Camera IFOV	98.6 μ rad/px; 0.646 m/px at 6550 m
Field of view	11.6° $\times$ 8.7°
Acquisition	Exposure 5 μ s, gain 0 dB, GenTL / GigE Vision




Figure 6: The validation setup on the GSOC rooftop test facility at Oberpfaffenhofen. The mount carrying the collimator and camera is installed on a mobile bench at the open cabin door, with the control laptop running the calibration GUI alongside; an RF antenna of the co-located ground station is visible outside.




Figure 7: Collimator assembly. Left: the camera with the 35 mm lens and, beneath it, the fibre-fed collimator on the two-axis kinematic mount driven by piezo-inertial actuators. Centre: the four-axis motion controller. Right: the handheld 635 nm Class 3R source, fibre-coupled to the collimator.

4.3 Software Configuration and Hyperparameters

Table 2 lists the scan and detection parameters recorded in the session log. The configuration is reported in full, including the options that were left disabled, because these bear directly on the results interpretation in Sections 5.3 and 5.5.

4.4 Measurement Protocol

The campaign was designed as a functional validation of the coarse-scan stage under field conditions. A single coarse scan was initiated from a nominal pointing position and executed without operator intervention until the pattern completed.

A protocol for the quantitative characterisation has been specified and is stated here so that the campaign is reproducible by others. Each session is initiated from a deliberately offset starting position, displaced from the known retroreflector location by a controlled

Table 2: Scan and detection parameters, session 20260721_223818.

Parameter	Value
Scan profile	<code>retro_11787m_search_v1</code>
Configured range	6550 m
Divergence model	5000 μ rad
Pattern	Archimedean spiral
Coverage at target	-11 to +40 m (X), -37 to +15 m (Y)
Planned points	599
Settle time	150 ms
Snap timeout	1200 ms
Approach policy	disabled (no backlash compensation)
Overshoot	0 steps
ROI	full frame
Expected-region gating	disabled
Image scale calibration	0.037 208 9 mm/px

angular amount, and the tool is executed without intervention until convergence is declared; repetition under nominally identical conditions yields the pointing accuracy distribution and the convergence time statistics. For a human-versus-algorithm comparison, operators perform manual alignment independently using the existing procedure, without access to the automated tool output, with alignment time and final pointing position recorded through the same logging infrastructure to ensure comparability. Executing this protocol requires the detector configuration changes identified in Section 7 and is outside the scope of the present paper, which reports the functional validation only.

4.5 Environmental Conditions

The tool provides for meteorological parameters (ambient temperature, relative humidity, wind speed and direction, and meteorological optical range as a proxy for atmospheric visibility) to be recorded alongside each session from a co-located weather station. This facility was not used during the campaign reported here, so the session record contains no quantitative environmental data. The qualitative conditions were clear and dark, as noted in Section 4.1. Enabling the environmental record is a prerequisite for the comparative campaigns described in Section 7, since performance cannot be attributed to atmospheric conditions that were not measured.

5 Initial Validation Results

This section reports the first end-to-end execution of the tool on a live outdoor path. The campaign validates coarse-scan function rather than providing a statistical characterisation; the repeatability and human-versus-algorithm campaigns are outstanding and are discussed in Section 7.

5.1 Coarse-Scan Execution

A single coarse scan of 599 planned pointing positions was executed over a 6550 m path on 21 July 2026, beginning at 22:38:25 and completing at 22:59:12 local time, for a total duration of 1247s. The mean cadence was 2.08 s per pointing position. Of this, 150 ms is the commanded settle interval and up to 1200 ms the frame acquisition timeout, placing the remainder, approximately 0.7 s per point, in mount slewing and command round-trip.

The scan traversed a spiral pattern spanning approximately 51 m in X and 52 m in Y in the target plane, corresponding to angular excursions of roughly ± 6 mrad about the nominal pointing position. The full pattern executed without operator intervention and without safety-limit violations, confirming that the motion model, the step-to-angle calibration, and the safety envelope behave correctly at operational scale.

The cadence is not uniform across the run. The first 280 pointing positions were completed in 482 s, a mean of 1.72 s per point, while the remaining 319 positions required 765 s, a mean of 2.40 s per point, an increase of 39 %. This follows from the pattern geometry rather than from any change in the acquisition settings: along an outward Archimedean spiral the arc length between consecutive points grows with radius, so the commanded slew distance, and with it the settling interval actually consumed, grows as the scan develops. The practical consequence for scan planning is that the duration of a spiral coarse scan is not linear in the number of points, and an estimate extrapolated from the early portion of a pattern will underestimate the total.

5.2 Acquisition Timing

The highest-scoring position was recorded at point index 280 of the 599 planned positions, 482 s into the 1247 s scan. Acquisition therefore occurred after 46.7 % of the planned points and 38.7 % of the elapsed time, with the remainder of the pattern executed without improvement on that position.

The outward spiral is designed for this. Since the pattern develops from the nominal pointing position outward, the region of the search space carrying the highest prior probability is sampled first, and the expected time to acquisition is shorter than the time to complete the pattern. The early-stop criterion of Section 3.5 exists to exploit this: had a signal threshold been configured for the run, the scan could have terminated at 482 s rather than 1247 s, a reduction of 61 % in the time occupied by the alignment procedure. The saturation described in Section 5.5 means that no meaningful threshold could have been set for this particular session, but the timing establishes the margin available to the mechanism once the detector is operated in its linear range.

5.3 Detection Yield

The spot detection module returned a valid detection at 438 of 599 pointing positions (73.1%); at the remaining 161 positions (26.9%) no candidate satisfied the detection criteria.

Because the campaign was conducted at night with negligible ambient illumination, these non-detections are genuine absences of a detectable return rather than failures of the acquisition chain: with no appreciable background in the field of view, a frame containing no detection is a frame in which no light reached the sensor above the detection threshold. The split reflects the illumination geometry. It is also consistent with the scale relationship between pattern and beam: the pattern spans approximately 51 m in the target plane against an illuminated footprint of 32.8 m, so the commanded excursions amount to roughly 1.6 beam diameters and a substantial fraction of pointing positions keep the retroreflector inside the illuminated spot. A high detected fraction is therefore the expected outcome of this configuration rather than evidence of a large capture region, and the 73.1% figure should not be read as a success rate. It does provide a first empirical bound on the angular extent over which a return is detectable at this range, which can be used to size the pattern in subsequent campaigns.

5.4 Consistency of the Logged Record

The traceability claim of Section 3 can be tested against the record itself. Each logged point carries the commanded actuator position in steps alongside the derived angular offset in microradians. Inverting the ratio of these two quantities recovers the step-to-angle scale factors independently for every point; across the sampled records the recovered values agree with the constants stored in the instrument configuration, 1.741 791 and 1.950 681 steps per microradian, to seven significant figures.

This confirms that the angular quantities reported throughout this paper are reproducible from the raw actuator commands and the stored configuration, and that no undocumented transformation intervenes between commanding and logging. The check covers the integrity of the logging chain, not the physical scale factors themselves, which were determined by a separate calibration procedure. Within that scope it shows that a reported result can be traced back through the record to the commands that produced it.

5.5 Signal Response and Detector Saturation

The scan-point score reached 255.11 at the best-ranked position and remained above 255 across a contiguous sequence of scan positions spanning several metres in the target plane. Because the camera was

operated in 8-bit mode, this indicates that the peak term $v_{\max}$ in Equation (3) reached the sensor ceiling.

The consequence follows directly from the form of the raw score. With $v_{\max}$ clipped at 255, Equation (3) degenerates to $S = 255 + 0.01 A$: the peak term becomes constant and the mask area A , intended only as a tie-breaker, becomes the sole discriminant. Inverting the recorded scores for the four highest-ranked points gives the mask areas listed in Table 3.

Table 3: Decomposition of the highest-ranked scan-point scores under saturation, $S = 255 + 0.01 A$.

Index	S	$v_{\max}$	A [px]
280	255.11	255	11
283	255.09	255	9
282	255.08	255	8
284	255.07	255	7

The entire signal separating the best-ranked point from its neighbours is therefore 4 px of mask area, contributing 0.04 to the score. This is comparable to the frame-to-frame variation of the threshold segmentation itself: a single boundary pixel crossing the threshold of Equation (1) is sufficient to reorder the ranking. The selection of a best point within the saturated plateau is consequently not a meaningful discrimination, and no pointing estimate should be derived from it.

The confidence measure does not provide independent verification in this regime either. When the spot is the brightest object in the scene, $P_n \rightarrow 1$ in Equation (5), and with $A = 11$ the normalised area term of Equation (6) contributes $A_n = 0.044$, giving $C \approx 0.76$. The detection passes the $C_{\min} = 0.35$ gate comfortably, but it does so by construction rather than by evidence.

Two aspects of the measurement are nonetheless informative. First, the non-saturated scan points across the pattern exhibit peak intensities in the range 10 to 20 ADU, so the candidate detection stands more than an order of magnitude above the background level of the scene: the detection is not a marginal noise excursion. Across the repeated trials conducted during the campaign the brightest source in the field was consistently the retroreflector return, identified by its disappearance when the source was interrupted and by its position relative to the known target bearing. That identification rests on repeated operator observation rather than on a logged reference frame, and recording a laser-blocked frame at the best-ranked position as a routine step of the protocol would place it on the same documentary footing as the rest of the session record. Second, the saturation itself is consistent with the configured geometry. At 6550 m with a 5 mrad divergence the illuminated spot is 32.8 m in diameter, substantially larger than the retroreflector aperture, so the return signal varies only weakly over pointing excursions of a few metres. The angular resolution attainable in this configuration is bounded by

the beam footprint rather than by the 0.5 mrad mount step resolution.




Figure 8: Recorded frame `img_000280` at the best-ranked pointing position of session 20260721_223818, displayed at the acquisition contrast. The candidate detection (red) is saturated and exhibits the diffraction pattern of a compact source; three further unresolved sources of comparable morphology (blue) are present in the same field of view. With full-frame scoring and expected-region gating disabled, none of the scoring layers of Section 3.2 distinguishes among them on grounds other than intensity.

5.6 Coverage of the Processing Chain

The campaign exercised the motion model, the scan pattern generator, the frame-level detection and scoring of Section 3.2, and the session logging backend. It did not exercise the fine-scan refinement stage, the geometric compensation of Section 3.9, or the hill-climbing local optimisation, none of which were triggered during the run. The temporal voting parameters of Equation (10) are not recoverable from the retained session record, so the degree of temporal averaging applied per point is unknown.

6 Discussion

6.1 Interpretation of the Validation Campaign

The campaign establishes that the tool operates as an integrated system on a live outdoor path: the mount is commanded correctly in angular units, the scan pattern generator produces a physically realisable trajectory over a multi-kilometre baseline, the acquisition chain sustains a 2 s cadence for over twenty minutes without intervention, and the session logging captures the run in a form that permits this analysis to be performed after the fact.

What the campaign does not establish is pointing accuracy. A single coarse scan without the refinement stage yields one candidate position and no estimate of its uncertainty, and the saturation behaviour of Section 5.5 means that the candidate is drawn from a plateau of indistinguishable positions. A pointing figure derived from this dataset would overstate what

the measurement supports, so none is given. Section 5 makes claims about the execution of the procedure (its cadence, detection behaviour, timing margin and the integrity of its record) and not about the alignment quality it achieves.

6.2 Range Dependence

The procedure itself is range-independent: the scan planner is parameterised by distance, and the control, detection and logging stages make no assumption about its value. What is range-dependent is the physics of the illumination and the return, and this governs how the quantities reported in Section 5 transfer between the 6550 m arrangement analysed there and the 11 787 m operational one, a factor of 1.80 apart. Three effects scale with distance.

Geometrically, the transverse extent of the illuminated spot scales linearly with range: the same 5 mrad divergence produces a 32.8 m footprint at 6550 m and 58.9 m at 11 787 m. Because the planner derives pattern extent and pitch from the configured range, the pattern scales with the footprint automatically and no reconfiguration is required. The angular resolution attainable in the saturated regime, however, is set by footprint size relative to the target, so it degrades in proportion, which argues for reducing divergence rather than for increasing scan density.

Atmospherically, turbulence-induced beam wander and scintillation both increase with path length, so the received signal exhibits larger short-timescale fluctuation at the longer arrangement [16]. The temporal voting and aggregation of Equation (10) exist to absorb this, and the number of frames per point is the parameter to raise when moving to the longer range, at a cost in scan duration that must be traded against the available observation window and that compounds with the larger pattern extent the longer range already implies.

Radiometrically, the two-way loss scales steeply with range for an unresolved retroreflector target, so the factor 1.80 between the two arrangements implies approximately an order of magnitude difference in returned power. Detector settings that saturate at 6550 m are not necessarily those that saturate at 11 787 m, so exposure and attenuation must be re-established per arrangement rather than carried across. Saturation at the shorter range is no guarantee of signal margin at the longer one. A link budget for the calibration path should be derived independently, using the established methodology for optical links of this class [9].

6.3 Limitations

Four limitations apply.

First, the campaign was conducted at 635 nm. For the pointing step this is consistent with operational practice, since the manual procedure also uses a

red seed laser to establish pointing [1], and the visible return is what the station camera detects in both cases. The operational beacon, however, emits at 1590 nm, and beam collimation is adjusted at that wavelength because divergence is wavelength-dependent. Diffraction-limited divergence, retroreflector beam deviation and atmospheric transmission therefore differ between the alignment wavelength and the operating wavelength. The tool addresses the pointing step only; extending it to the collimation step would require the infrared detection path.

Second, detector saturation collapsed the raw score of Equation (3) onto its area tie-breaker, reducing the effective discriminant within the high-signal region to a few pixels of mask area, as quantified in Section 5.5. The scoring function is monotonic in received power only while $v_{\max}$ remains below the sensor ceiling, which constrains the exposure and attenuation settings admissible for any quantitative campaign.

Third, the session was executed with expected-region gating and ROI restriction disabled, leaving the detection stage vulnerable to bright out-of-band sources in the field of view. At an operational site with controlled surroundings this is less severe than at a rooftop facility in a populated area, but spatial gating should be regarded as mandatory rather than optional.

Fourth, backlash compensation was disabled during the campaign, so the reported positions do not account for direction-dependent hysteresis in the piezo-inertial actuators. Characterising and compensating this is a prerequisite for any repeatability claim.

6.4 Generalisability

The architecture makes no assumption specific to either arrangement beyond the presence of a two-axis motorised mount with a serial or network control interface, a co-aligned camera, and a cooperative target at a known approximate bearing. The scan model is parameterised by range and divergence, and the step-to-angle calibration is determined empirically per installation; operating at two ranges differing by a factor of 1.80 required no change to the procedure, only a change to the configured distance. Adapting the tool to another OGS therefore reduces to supplying a device driver, entering the range and re-running the step-to-angle calibration.

7 Conclusion and Future Work

We have presented a software tool that automates uplink laser alignment at an optical ground station by closing the loop between a motorized collimator mount and the return from a retroreflector on a distant calibration target. The tool replaces an operator-dependent manual procedure with a parameterised two-stage search, records every run in a ses-

sion log bound to a software version, and has been demonstrated end-to-end on a 6550 m outdoor path at the GSOC rooftop test facility in Oberpfaffenhofen, executing a 599-point coarse scan autonomously in 20.8 min and reaching its best-scoring position after 38.7 % of that interval. The campaign also established two configuration constraints that bound what the procedure can currently resolve: the scoring function must be kept within the linear range of the detector, and spatial gating must be active for the scoring to discriminate on grounds other than intensity.

The immediate priority is to complete the quantitative characterisation. This requires three changes to the measurement configuration: reduction of the received signal into the linear detector range through attenuation or exposure control, so that the scoring function is monotonic in received power; activation of expected-region gating together with source modulation and frame differencing, to reject ambient sources; and enabling backlash compensation in the motion policy. With these in place, the repeatability protocol of Section 4.4 and the human-versus-algorithm comparison can be executed and reported.

Beyond that, planned work comprises re-characterisation at the operational range in Almería, extension of the refinement stage with temporal averaging to accommodate the increased atmospheric fluctuation expected at the longer path, and integration of the calibration routine into the station monitoring, control and automation system [4] so that alignment can be re-established automatically as part of routine pre-pass preparation.

Acknowledgements

This work was carried out within the GlobeON project at the German Space Operations Center (GSOC) of the German Aerospace Center (DLR). The authors thank the staff of the GSOC rooftop optical test facility at Oberpfaffenhofen for support during the validation campaign.

References

- [1] Lagadrillière, Pierre-Alexis and Karafillis, Petros and Dlubatz, Michael and Knopp, Marcus Thomas and Nouri, Bijan and Zarzalejo, Luis F. (2025) FOGATA — Integrating optical ground stations into common RF ground stations networks. In: *Proceedings of the International Astronautical Congress, IAC*, paper IAC-25-103154. 76th International Astronautical Congress (IAC 2025), 2025-09-29 – 2025-10-03, Sydney, Australia. ISSN 0074-1795.
- [2] Lagadrillière, Pierre-Alexis and Karafillis, Petros and Lucas, Dominic and Falcone, Rossella and Knopp, Marcus Thomas (2025) Latest news from FOGATA — Deployment of the fully automated DLR OGS in Almeria and first results with optical

- DTE communication. In: *18th International Conference on Space Operations (SpaceOps 2025)*, 2025-05-26 – 2025-05-30, Montréal, Canada.
- [3] Knopp, Marcus Thomas and Immerz, Andreas and Giggenbach, Dirk and Köhler, Alexander (2022) The Small Optical Ground Stations Focal-Optics Assembly (SOFA). In: *73rd International Astronautical Congress, IAC 2022*, 2022-09-18 – 2022-09-22, Paris, France. ISSN 0074-1795. <https://dl.iafaastro.directory/event/IAC-2022/paper/68520/>
 - [4] Jain, Himani and Kalkhof, Alexander and Veit, Stefan and Dlubatz, Michael and Knopp, Marcus Thomas (2022) A monitoring, control & automation system for optical communication ground stations. In: *Proceedings of the International Astronautical Congress, IAC*, 2022-09-18 – 2022-09-22, Paris, France. ISSN 0074-1795.
 - [5] Rattenbury, Nicholas and Knopp, Marcus Thomas and Köhler, Alexander et al. (2025) Update on the German and Australasian Optical Ground Station Networks. *International Journal of Satellite Communications and Networking*. doi:10.1002/sat.1564
 - [6] Fürbacher, Anna and Fruth, Thomas and Wiebigke, Armin and Wörle, Maria Th. and Mrowka, Falk and Saucke, Karen and Martín Pimentel, Patricia and Knopp, Marcus (2025) Concept for generic agile, reactive optical link planning. *CEAS Space Journal* 17, pp. 891–900. doi:10.1007/s12567-025-00592-0
 - [7] Giggenbach, Dirk and Karafillis, Petros and Rittershofer, Jonas and Immerz, Andreas and Spörl, Andreas and Gaißer, Steffen and Klinkner, Sabine and Knopp, Marcus Thomas (2023) Transmitter beam bias verification for optical satellite data downlinks with open-loop pointing — the 3-OGS-experiment. In: *Proc. SPIE 12777, International Conference on Space Optics — ICSO 2022*, 127770X. doi:10.1117/12.2689023
 - [8] Carrillo-Flores, Andrea and Giggenbach, Dirk and Knopp, Marcus Thomas and Orsucci, Davide and Shrestha, Amita (2023) Effects of pointing errors on intensity losses in the optical LEO uplink. In: *Proc. SPIE 12777, International Conference on Space Optics — ICSO 2022*, 127775U. doi:10.1117/12.2691100
 - [9] Giggenbach, Dirk and Knopp, Marcus Thomas and Fuchs, Christian (2023) Link budget calculation in optical LEO satellite downlinks with on/off-keying and large signal divergence: a simplified methodology. *International Journal of Satellite Communications and Networking* 41 (5), pp. 460–476. doi:10.1002/sat.1478
 - [10] Fuchs, Christian and Schmidt, Christopher (2019) Update on DLR’s OSIRIS program. In: *Proc. SPIE 11180, International Conference on Space Optics — ICSO 2018*, 111800I. doi:10.1117/12.2535937
 - [11] Boroson, Don M. and Robinson, Bryan S. and Murphy, Daniel V. and Burianek, Dennis A. and Khatri, Farzana and Kovalik, Joseph M. and Sodnik, Zoran and Cornwell, Donald M. (2014) Overview and results of the Lunar Laser Communication Demonstration. In: *Proc. SPIE 8971, Free-Space Laser Communication and Atmospheric Propagation XXVI*, 89710S. doi:10.1117/12.2045508
 - [12] European Space Agency. The European Optical Nucleus Network, ESOC. <https://esoc.esa.int/content/european-optical-nucleus-network> (accessed 15.08.2026).
 - [13] Kongsberg Satellite Services. Optical Networks. <https://www.ksat.no/ground-network-services/new-technologies/optical-comms/> (accessed 15.08.2026).
 - [14] ISO 11146-1:2021 *Lasers and laser-related equipment — Test methods for laser beam widths, divergence angles and beam propagation ratios — Part 1: Stigmatic and simple astigmatic beams*. International Organization for Standardization, Geneva.
 - [15] Hardy, John W. (1998) *Adaptive Optics for Astronomical Telescopes*. Oxford University Press, New York.
 - [16] Andrews, Larry C. and Phillips, Ronald L. (2005) *Laser Beam Propagation through Random Media*, 2nd ed. SPIE Press, Bellingham, WA.
 - [17] Ariyur, Kartik B. and Krstić, Miroslav (2003) *Real-Time Optimization by Extremum-Seeking Control*. Wiley, Hoboken, NJ.