

Airborne real-time solar concentrator orientation estimation for heliostat fine calibration

Alexander Schnerring^{a, c, *}, Marc Röger^a, Adrian Winter^{a, 1}, Rafal Broda^{a, c}, Michael Nieslony^a, Niels Algner^a, Benedikt Kölsch^b, Robert Pitz-Paál^{b, c}

^a German Aerospace Center (DLR), Institute of Solar Research, Calle Doctor Carracido 44, 04005, Almería, Spain

^b German Aerospace Center (DLR), Institute of Solar Research, Linder Höhe, 51147, Cologne, Germany

^c RWTH Aachen University, Chair of Solar Technology, Linder Höhe, 51147, Cologne, Germany

HIGHLIGHTS

- Novel airborne method utilizing real-time data analysis enables autonomous acquisition of heliostat fine calibration points without preplanned flight paths.
- A closed-loop feedback controller controls the location of a point light reflection on the concentrator surface, thereby eliminating redundant measurements.
- The method enables targeted recalibration of heliostat outliers in tracking heliostat fields by preparing a fine-calibration surface scan within approximately one minute, without interrupting plant operation.
- The method is systematically tuned and evaluated in a dedicated simulation environment across various system configurations and operating conditions.

ARTICLE INFO

Keywords:

Unmanned aerial vehicle
Heliostat calibration
Condition monitoring
Concentrated solar power
Closed-loop control
Autonomous systems

ABSTRACT

Unmanned aerial vehicle (UAV)-based methods are a promising option for heliostat calibration. However, existing approaches require prior knowledge of the heliostat orientation and rely on predefined flight paths, which can be difficult to determine when the assumed orientations are uncertain. Small orientation errors necessitate redundant flight segments to reliably capture the measurement signal, while largely misaligned heliostats cannot be measured altogether. This raises the question of whether real-time data processing can enable the dynamic adaptation of UAV flight paths. To address this, we present a novel method for the airborne real-time solar concentrator orientation estimation for heliostat fine calibration, termed ARTSCORE-F. Building on the *HelioPoint* measurement principle, the method introduces a closed-loop controller that incorporates the currently observed reflection signal as feedback. For the first time, a heliostat fine calibration point can be acquired autonomously without requiring previous calibration or accurate *a priori* orientation knowledge. The feedback loop actively guides the reflection signal toward defined locations on the concentrator surface while simultaneously controlling the UAV–concentrator distance, allowing online adaptation of the flight path to the actual measurement situation and avoiding redundant measurements. A dedicated simulation environment is used to systematically evaluate the method across varying system configurations and operating conditions. For an exemplary concentrator at the Solar Tower Jülich, the proposed method prepares a final surface scan within approximately one minute, even for arbitrary initial orientations. The method is applicable in tracking heliostat fields and enables targeted recalibration of individual heliostats without interrupting field operation.

1. Introduction

The heliostat field of a concentrating solar power (CSP) tower plant comprises up to several thousand heliostats that concentrate solar radiation onto a central receiver. Each heliostat must continuously track the

angle bisector between the incident solar vector and the vector from the heliostat to the receiver. Various factors can introduce heliostat tracking errors [1], reducing the overall performance of the heliostat field. Accurate calibration is therefore essential to identify and compensate for these errors, ensuring high system efficiency [2].

* Corresponding author at: German Aerospace Center (DLR), Institute of Solar Research, Calle Doctor Carracido 44, 04005, Almería, Spain.

Email address: alexander.schnerring@dlr.de (A. Schnerring).

¹ Now with Hamburg University of Applied Sciences, Department Aeronautical Engineering, Stiftstraße 69, 20999, Hamburg, Germany.

Nomenclature

Acronyms

CSP	concentrating solar power
GPS	global positioning system
ITAE	integral of time-weighted absolute error
ITAESS	integral of time-weighted steady-state-corrected absolute error
LED	light-emitting diode
NIO	Non-Intrusive Optical
NVR	normal vector refinement
NVPS	normal velocity profile shaping
PAE	peak absolute error
PI	proportional-integral
RMS	root mean square
RTK	real-time kinematics
UAV	unmanned aerial vehicle

Coordinate Systems

CCS	concentrator coordinate system
CCSXY	concentrator coordinate system XY-plane
FCS	facet coordinate system
GCS	global coordinate system
ICS	image coordinate system
NCS	navigation coordinate system
OCS	observer coordinate system

Unmanned aerial vehicle (UAV)-based approaches offer advantages over conventional methods for heliostat calibration such as the camera-target method, particularly by enabling flexible high-resolution imaging over large areas as well as the simultaneous measurement of multiple heliostats. Recent research has therefore explored computer vision approaches applied to aerial images captured by UAV-mounted cameras [3–7]. These methods extract measurement signals from recorded images, such as concentrator corners or reflections, and use them to infer quantities of interest such as heliostat orientation.

Among these approaches, the *HelioPoint* method [6] achieves fine calibration accuracy (i.e., an accuracy around 0.3 mrad) by using the position of a light-emitting diode (LED) reflection on the concentrator surface as a measurement signal. However, the method relies on pre-planned flight paths and requires sufficient prior knowledge of the heliostat orientation to ensure visibility of the reflection. In practice, this corresponds to a coarse calibration accuracy on the order of 10 mrad. As a result, the method is primarily suited for the fast initial calibration of a heliostat field during the commissioning phase, where heliostats can be freely controlled and aligned prior to the measurement. Its application in an operating field is limited, as tracking heliostats continuously change their orientation and suitable viewing geometries cannot be guaranteed along pre-planned flight paths.

In this work, we present ARTSCORE-F, a method for the airborne real-time solar concentrator orientation estimation for heliostat fine calibration. This method extends the *HelioPoint* measurement principle toward real-time operation: By incorporating the observed reflection signal into a closed-loop control scheme, ARTSCORE-F enables the dynamic adaptation of the UAV flight path. This allows the targeted positioning of the reflection on the concentrator surface and the control of the UAV–concentrator distance, ensuring reliable signal visibility without requiring accurate prior knowledge of the heliostat orientation.

ARTSCORE-F is particularly suited for use during plant operation as a targeted method for the recalibration of individual heliostats exhibiting tracking errors. For example, such heliostat outliers could first be identified using ARTSCORE-C [7], after which ARTSCORE-F performs a reflection-based measurement to refine their orientation with fine calibration accuracy. This enables efficient recalibration without interrupting normal plant operation and thereby complements existing UAV-based approaches, as outlined in Table 1.

By enabling the dynamic planning of the UAV flight path, the proposed approach opens up new possibilities for the autonomous UAV-based calibration of heliostat fields. This contributes to the broader goal of reducing operational costs in CSP systems and aligns with strategic objectives outlined in the HelioCon Roadmap [8], which emphasizes the development of advanced inspection tools to improve the cost competitiveness of solar tower technology.

A preliminary prototype of the method was first tested in a master's thesis by Winter [9]. The main contributions of the present work are: (i) The extension of the method to larger parts of the measurement process, (ii) a systematic simulation-based tuning and evaluation of the method using a high-fidelity simulation environment representing the heliostats at the Solar Tower Jülich [10], and (iii) an analysis of robustness and operational trade-offs under realistic conditions for the considered field, including the influence of concentrator surface imperfections and modeling approximations.

This paper is structured as follows. Section 2 introduces the notation and conventions and summarizes the mathematical and methodological preliminaries relevant to this work. Section 3 details the proposed system methodology. Section 4 describes the approach used to identify and evaluate suitable system parameters in simulation. Section 5 presents and discusses the simulation results. Finally, Section 6 summarizes the main contributions and outlines directions for future research.

2. Preliminaries

2.1. Conventions and notation

The heliostat field is assumed to consist of one heliostat with two tracking axes, controlling the orientation of a concentrator with M facets. The coordinate system associated with this concentrator is referred to as the concentrator coordinate system (CCS), while each facet m has its own facet coordinate system (FCS), denoted by FCS_m . These coordinate systems are defined with respect to the global coordinate system (GCS), chosen to be a local east-north-up coordinate system on the northern hemisphere in accordance with the *SolarPACES Guideline for Heliostat Performance Testing* [11,12]. Likewise, the pose of a camera moving through the heliostat field is characterized by a coordinate system, which changes over time as the camera navigates the environment. As the acronym CCS is already in use, the camera coordinate system is denoted by the observer coordinate system (OCS). The 2D coordinate system associated with the camera sensor is denoted by the image coordinate system (ICS). Fig. 1 illustrates the coordinate systems used throughout this work.

Scalar quantities are denoted by lowercase letters in regular font, vectors by lowercase letters in bold font, and matrices by uppercase letters in bold font. The individual components of a Cartesian vector $\mathbf{v}$ are written as $(\mathbf{v})_x$, $(\mathbf{v})_y$, and $(\mathbf{v})_z$, corresponding to the directions X, Y, and Z, respectively.


2.2. Coordinate transformations

The transformation between two coordinate systems CS1 and CS2 is described by a rigid body transform and is conveniently summarized in

Table 1

Comparison of selected UAV-based heliostat calibration approaches developed at the German Aerospace Center.

Method	Accuracy	Measurement Principle	Flight Planning	Application
ARTSCORE-C [7]	Coarse	Corner-based	Pre-planned/Dynamic	<i>Commissioning</i> Fast acquisition of coarse calibration points for <i>groups of heliostats</i> <i>Operation</i> Tracking outlier detection
HelioPoint [6]	Fine	Reflection-based	Pre-planned	<i>Commissioning</i> Fast acquisition of fine calibration points for <i>groups of heliostats</i>
ARTSCORE-F (this work)	Fine	Reflection-based	Dynamic	<i>Operation</i> Fast acquisition of fine calibration points for <i>individual heliostats</i>

**Fig. 1.** Overview of the coordinate systems for the case of the heliostats at the Solar Tower Jülich, where each concentrator surface consists of a 2×2 grid of facets, i.e., $M = 4$.

a 4×4 matrix of the form

$$T_{CS2}^{CS1} = \begin{pmatrix} r_{XX} & r_{XY} & r_{XZ} & t_X \\ r_{YX} & r_{YY} & r_{YZ} & t_Y \\ r_{ZX} & r_{ZY} & r_{ZZ} & t_Z \\ 0 & 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} R_{CS2}^{CS1} & t_{CS2}^{CS1} \\ \mathbf{0} & 1 \end{pmatrix}, \quad (1)$$

where t_{CS2}^{CS1} denotes the translation from CS1 to CS2 and R_{CS2}^{CS1} denotes the rotation from CS1 to CS2. The operator $\text{Rot}(\cdot)$ extracts the rotational part from a transformation, while the operator $\text{Trans}(\cdot)$ extracts its translational part, i.e.,

$$\begin{aligned} \text{Rot}(T_{CS2}^{CS1}) &= R_{CS2}^{CS1}, \\ \text{Trans}(T_{CS2}^{CS1}) &= t_{CS2}^{CS1}. \end{aligned} \quad (2)$$

A point $p^{CS2} = (p_X^{CS2}, p_Y^{CS2}, p_Z^{CS2})^T$ expressed with respect to CS2 can be expressed with respect to CS1 in the following way:

$$\begin{pmatrix} p^{CS1} \\ 1 \end{pmatrix} = T_{CS2}^{CS1} \cdot \begin{pmatrix} p^{CS2} \\ 1 \end{pmatrix}. \quad (3)$$

To simplify notation, the transformation of the point p^{CS2} to CS1 is denoted by

$$p^{CS1} = T_{CS2}^{CS1} \cdot p^{CS2} \quad (4)$$

but implicitly applies Eq. (3). Conversely, a point expressed with respect to CS1 can be expressed with respect to CS2 by multiplying with the inverse of the transformation T_{CS2}^{CS1} :

$$p^{CS2} = (T_{CS2}^{CS1})^{-1} \cdot p^{CS1} = T_{CS1}^{CS2} \cdot p^{CS1}. \quad (5)$$

A point cloud P is a collection of N points and can be expressed in the form of a matrix:

$$P^{CS2} = \begin{pmatrix} p_{X1}^{CS2} & \dots & p_{XN}^{CS2} \\ p_{Y1}^{CS2} & \dots & p_{YN}^{CS2} \\ p_{Z1}^{CS2} & \dots & p_{ZN}^{CS2} \end{pmatrix}. \quad (6)$$

Analogously to Eqs. (3) and (4), a point cloud can be expressed with respect to another coordinate system, i.e.,

$$P^{CS1} = T_{CS2}^{CS1} \cdot P^{CS2}. \quad (7)$$

2.3. Kinematic model

The kinematic model K describes the motion of the CCS with respect to the GCS. It is parameterized by the vector k , which contains all kinematic model parameters, such as geometric properties and mounting errors. The parameterization is individually fitted for each heliostat using a set of calibration points. In addition, the model depends on the heliostat origin o^{GCS} . In a forward pass, the kinematic model K^f takes in two tracking angles, τ_1 and τ_2 , and computes the resulting concentrator pose:

$$K = K^f(\tau_1, \tau_2 | o^{GCS}, k) = T_{CCS}^{GCS}. \quad (8)$$

In a backward pass, the kinematic model K^b takes in a desired concentrator normal vector n^{GCS} and computes the concentrator pose that results in this normal vector:

$$K = K^b(n^{GCS} | o^{GCS}, k) = T_{CCS}^{GCS}. \quad (9)$$

2.4. Homography

The planar homography H_{CS1}^{CS2} describes the projective transformation between two two-dimensional coordinate systems CS1 and CS2 up to an arbitrary projective scale factor s [13], i.e.,

$$s \cdot \begin{pmatrix} p_X^{CS2} \\ p_Y^{CS2} \\ 1 \end{pmatrix} = H_{CS1}^{CS2} \cdot \begin{pmatrix} p_X^{CS1} \\ p_Y^{CS1} \\ 1 \end{pmatrix}. \quad (10)$$

In this work, a homography maps image points from the ICS to the concentrator coordinate system XY-plane (CCSXY-plane). The homography matrix H_{ICS}^{CCSXY} is estimated from two corresponding sets of points, P_{corner}^{ICS} and $(P_{corner}^{CCS})_{XY}$. Applying the homography H_{ICS}^{CCSXY} to a point in homogeneous image coordinates yields

$$\begin{pmatrix} p_{X'}^{CCSXY} \\ p_{Y'}^{CCSXY} \\ p_{Z'}^{CCSXY} \end{pmatrix} = H_{ICS}^{CCSXY} \cdot \begin{pmatrix} p_X^{ICS} \\ p_Y^{ICS} \\ 1 \end{pmatrix}. \quad (11)$$

De-homogenization gives the 3D point in the CCSXY-plane:

$$p^{CCSXY} = \frac{1}{p_{Z'}^{CCSXY}} \begin{pmatrix} p_{X'}^{CCSXY} \\ p_{Y'}^{CCSXY} \end{pmatrix}. \quad (12)$$

For brevity, the conversion between Euclidean and homogeneous coordinates is implicitly assumed when applying the homography, i.e.,

$$p^{CCSXY} = H_{ICS}^{CCSXY} \cdot p^{ICS}. \quad (13)$$

2.5. HelioPoint

The proposed system is based on the *HelioPoint* method, whose objective is the acquisition of calibration points with fine calibration accuracy of around 0.3 mrad, using a camera and an LED mounted on a UAV [6].




Fig. 2. Measurement principle of the *HelioPoint* method: A camera observes the reflection of a light source (LED) on the concentrator surface. The reflection position $p_{\text{refl}}^{\text{GCS}}$ in the GCS is obtained by first mapping the reflection position observed in the ICS to the CCS via a homography derived from the observed concentrator corner points $p_{\text{corner}}^{\text{ICS}}$, and subsequently transforming it into the GCS using an estimate of the concentrator pose $T_{\text{CCS}}^{\text{GCS}}$. The local surface normal vector $n_{\text{CCS}}^{\text{GCS}}$ can then be determined as the bisector between the incident light direction, defined by the vector from $p_{\text{refl}}^{\text{GCS}}$ to the light source position $p_{\text{LED}}^{\text{GCS}}$, and the reflected light direction, defined by the vector from $p_{\text{refl}}^{\text{GCS}}$ to the camera position $p_{\text{cam}}^{\text{GCS}}$.

2.5.1. Setup

Fig. 2 illustrates the measurement principle in a simplified 2D view. Both a strong LED and a camera are pointed toward the concentrator surface to be measured. If the camera is positioned close to the reflected light ray of the LED, a reflection of it appears in the ICS. The concentrator normal vector can be determined using

$$n_{\text{CCS}}^{\text{GCS}} = \frac{(p_{\text{cam}}^{\text{GCS}} + p_{\text{LED}}^{\text{GCS}})/2 - p_{\text{refl}}^{\text{GCS}}}{\|(p_{\text{cam}}^{\text{GCS}} + p_{\text{LED}}^{\text{GCS}})/2 - p_{\text{refl}}^{\text{GCS}}\|}, \quad (14)$$

where the vector $p_{\text{cam}}^{\text{GCS}}$ denotes the camera position, the vector $p_{\text{LED}}^{\text{GCS}}$ denotes the LED position, and the vector $p_{\text{refl}}^{\text{GCS}}$ denotes the position of the reflection on the concentrator surface. Note that in practice, reflections may appear as spatially extended regions rather than ideal point features, as discussed in Section 2.5.4. Nevertheless, the geometric relation expressed by Eq. (14) remains valid and can be evaluated for every point belonging to the reflection region. Furthermore, none of the quantities in Eq. (14) are known exactly and must instead be obtained through estimation: The camera position can be estimated using real-time kinematics (RTK) [6,14] or an optical method such as visual simultaneous localization and mapping [15] or ARTSCORE-C [7]. The reflection position $p_{\text{refl}}^{\text{GCS}}$ is estimated by first mapping the detected reflection image coordinates $p_{\text{refl}}^{\text{ICS}}$ to the CCS using a homography. This point can then be transformed to the GCS using a concentrator pose estimate $\hat{T}_{\text{CCS}}^{\text{GCS}}$:

$$\hat{p}_{\text{refl}}^{\text{GCS}} = \hat{T}_{\text{CCS}}^{\text{GCS}} \cdot \hat{p}_{\text{refl}}^{\text{CCS}}. \quad (15)$$

For readability, the hat notation is omitted in the remainder of this work, and all quantities are assumed to represent their respective estimates unless stated otherwise.

2.5.2. Defining the navigation coordinate system

In order to control the reflection position $p_{\text{refl}}^{\text{ICS}}$ with respect to the concentrator corners $p_{\text{corner}}^{\text{ICS}}$, it is convenient to define the navigation coordinate system (NCS), consisting of the navigation vector and the navigation plane: The navigation vector is given as the unit vector pointing from the reflection position to the camera position, i.e.,

$$r^{\text{GCS}} = \frac{p_{\text{cam}}^{\text{GCS}} - p_{\text{refl}}^{\text{GCS}}}{\|p_{\text{cam}}^{\text{GCS}} - p_{\text{refl}}^{\text{GCS}}\|}. \quad (16)$$

The navigation plane is defined as the plane perpendicular to the navigation vector r^{GCS} . While camera movements along the navigation

vector r^{GCS} change the scale of the concentrator surface in the ICS, they do not affect the reflection position relative to the concentrator corners. In contrast, lateral movements in the navigation plane cause the observed reflection to shift relative to the concentrator corners while the concentrator surface size remains approximately constant, given that the camera distance is large compared to the concentrator dimensions.

2.5.3. Simplification of co-located camera and LED

As described in [6], the LED can be mounted directly next to the camera lens. Throughout this work, it is assumed that the camera position $p_{\text{cam}}^{\text{GCS}}$ and the LED position $p_{\text{LED}}^{\text{GCS}}$ are co-located and jointly referred to as the UAV position $p_{\text{UAV}}^{\text{GCS}}$:

$$p_{\text{UAV}}^{\text{GCS}} := p_{\text{LED}}^{\text{GCS}} \approx p_{\text{cam}}^{\text{GCS}}. \quad (17)$$

With this assumption, Eq. (14) simplifies to

$$n_{\text{CCS}}^{\text{GCS}} = \frac{p_{\text{UAV}}^{\text{GCS}} - p_{\text{refl}}^{\text{GCS}}}{\|p_{\text{UAV}}^{\text{GCS}} - p_{\text{refl}}^{\text{GCS}}\|}, \quad (18)$$

which is exactly the definition of the navigation vector r^{GCS} . Consequently, the navigation plane coincides with the concentrator coordinate system XY-plane.

The impact of this approximation can be estimated as follows: For a spatial separation of approximately 40 mm between the camera and LED, the resulting angular deviation in the estimated normal vector is on the order of 1.0 mrad at a distance of 20 m, 0.4 mrad at 50 m, and 0.2 mrad at 100 m, accounting for the reflection geometry. These values represent a systematic bias that becomes relevant for high-accuracy surface characterization. For the purpose of closed-loop navigation considered in this work, this effect is small compared to other modeling approximations, and is therefore not expected to significantly affect the system behavior.

2.5.4. Simplification of the reflection characterization

The surface of a concentrator typically comprises multiple mirror facets whose orientations are individually adjusted (i.e., “canted”) such that the overall concentrator surface approximates an ideal paraboloid. Additionally, the individual facets may be shaped with a slight curvature to better match the local geometry of the paraboloid and improve optical performance. In practice, real concentrator surfaces exhibit deviations from this ideal shape.

As a result of these effects and the finite spatial extent of the light source and optical effects, reflections generally appear as spatially extended regions in the image rather than as ideal point features. The size and shape of these regions depend, among other factors, on the surface quality and the distance between the UAV and the concentrator. Reflections could hence be described as spatial distributions or regions in the recorded images. A perception module could segment the image into reflective and non-reflective regions, for example using blob detection, effectively classifying each pixel as either belonging to a reflection or not. Based on this segmentation, a reflection may be represented by the collection of detected pixels or by fitting a parametric distribution to these pixels. The choice of representation depends on the intended use of the reflection signal.

The concentrator facets at the Solar Tower Jülich are designed to be flat, and experimental data confirm that the observed reflections are consistently compact, even at larger distances. This is illustrated in Fig. 3, where each reflection covers only a small number of pixels. For the scope of the present work, which focuses on the closed-loop navigation of the UAV rather than detailed surface characterization, each reflection is hence represented by a single pixel location per facet. Depending on the chosen interpretation, this location can be understood as the average coordinate of the detected reflection pixels or as the mean of the fitted distribution. Spatial modeling of the reflection signal may become relevant for more extended reflections or when targeting high-accuracy surface characterization, as outlined in Section 3.5.




Fig. 3. Cropped section of an image recorded at the Solar Tower Jülich using a *HelioPoint* measurement setup. Each of the four facets shows a distinct reflection. The reflections appear compact and are therefore represented by single pixel locations rather than spatially extended regions.




Fig. 4. Simplified two-dimensional illustration of the reflection corridors of a concentrator with focal length f , composed of two flat, canted mirror facets. The intersecting corridors define regions in which a camera with an attached LED simultaneously observes multiple LED reflections from different facets.

Assuming the camera and LED to be approximately co-located (see Eq. (17)), each flat facet defines a three-dimensional reflection corridor extending from the facet surface along its normal vector. If the UAV (i.e., camera & LED) is located within such a corridor, the LED reflection can be observed by the camera on the corresponding facet. As illustrated in a simplified 2D view in Fig. 4, a concentrator surface with multiple flat facets defines several distinct reflection regions: When the UAV is positioned in an area where multiple reflection corridors intersect, reflections are simultaneously visible on multiple facets. These regions depend on the facet dimensions, canting configuration, and the concentrator focal length f .

For concentrators comprising M facets, up to M reflections may be observed simultaneously. To represent the reflection signal in a quantity of constant dimension, the collection of observed reflections is described by the $(2 \times M)$ point cloud

$$\mathbf{P}_{\text{refl}}^{\text{ICS}} = \left(\mathbf{p}_{\text{refl},1}^{\text{ICS}}, \dots, \mathbf{p}_{\text{refl},M}^{\text{ICS}} \right). \quad (19)$$

Here, each vector $\mathbf{p}_{\text{refl},m}^{\text{ICS}}$ corresponds to the image coordinates of the reflection associated with facet m . If no reflection is visible on facet m , the m -th column of $\mathbf{P}_{\text{refl}}^{\text{ICS}}$ is set to the void vector $(\text{NaN}, \text{NaN})^T$. This facet-wise representation ensures that at most one coordinate is assigned to each facet, matching the compact reflections observed at the Solar Tower Jülich.

The concentrators at the Solar Tower Jülich consist of a 2×2 grid of facets, indexed by $m \in \{\text{UL}, \text{UR}, \text{LR}, \text{LL}\}$ (i.e., upper left, upper right, lower right, and lower left). Hence, the reflection signal takes the form

$$\mathbf{P}_{\text{refl}}^{\text{ICS}} = \left(\mathbf{p}_{\text{refl},\text{UL}}^{\text{ICS}}, \mathbf{p}_{\text{refl},\text{UR}}^{\text{ICS}}, \mathbf{p}_{\text{refl},\text{LR}}^{\text{ICS}}, \mathbf{p}_{\text{refl},\text{LL}}^{\text{ICS}} \right). \quad (20)$$




Fig. 5. Schematic overview of the four-step system workflow, illustrating how the UAV is navigated to maintain continuous visibility of the reflection signal. For simplicity, only a single facet reflection is shown. To qualitatively indicate facet canting, the spatial extent of the reflection corridor is illustrated as decreasing with increasing distance from the concentrator.

3. System methodology

The presented system builds on the *HelioPoint* measurement setup described in the previous section. Its overall objective is to enable the observation of the LED reflection at specified locations on the concentrator surface from sufficiently large distances between the UAV and the concentrator. As discussed previously, the UAV must be positioned within the concentrator reflection corridor in order for the camera to observe a reflection signal. The location and extent of this corridor depend on the orientation of the concentrator surface, and positioning the UAV in the reflection corridor therefore requires *a priori* knowledge of the concentrator normal vector. In practice, this orientation is only available with limited accuracy. As illustrated in Fig. 5, the resulting uncertainty in the estimated normal vector directly propagates to uncertainty in the predicted location of the reflection corridor and becomes increasingly pronounced with growing distances between the UAV and the concentrator. Hence, the probability of observing a reflection when positioning the UAV without feedback from the reflection signal is significantly reduced for large initial distances between the UAV and the concentrator. To address this challenge, the following approach is proposed:

1. **Stage 1:** Acquire an initial reflection signal by positioning the UAV close to the concentrator, where it is likely that a reflection is visible even with high uncertainty in the initial normal vector estimate.
2. **Stage 2:** Guide the reflection signal toward the concentrator center while keeping the distance between the UAV and the concentrator roughly constant. This reduces the likelihood of losing the reflection when increasing the distance in the subsequent stage.
3. **Stage 3:** Move the UAV further away from the concentrator until a specified target distance is reached.
4. **Stage 4:** Perform a surface scan while keeping the distance between the UAV and the concentrator roughly constant.

Section 3.1 discusses the initial distance between the UAV and the concentrator required to detect the initial reflection signal in Stage 1. Section 3.2 describes the common framework used for Stages 2, 3, and 4, in which the UAV is guided by velocity commands defined in the CCS and subsequently transformed into the GCS. Building on this framework, Section 3.3 proposes an improvement to increase the robustness of Stage 3. While these subsections assume a single two-dimensional reflection signal, Section 3.4 discusses the implications of multi-faceted concentrators, where multiple reflection signals may be observed simultaneously. Section 3.5 discusses the final surface measurement stage,

where a guided *HelioPoint* measurement is enabled, and elaborates on how established results can be transferred and further extended in the context of the proposed framework. [Section 3.6](#) addresses practical challenges arising when transferring the method from simulation to the real world. [Section 3.7](#) concludes this section with remarks on implementation.

3.1. Stage 1

For the reflection to be initially visible, the concentrator orientation must be known with sufficient accuracy. For a rectangular concentrator of dimensions $w_{\text{conc}} \times h_{\text{conc}}$, let

$$l_{\text{conc}} = \min\{w_{\text{conc}}, h_{\text{conc}}\} \quad (21)$$

denote the shorter side length. Assuming a perfectly flat, single-facet concentrator, the extent of the reflection corridor is independent of the UAV–concentrator distance and depends solely on the concentrator dimensions. For the worst-case scenario, the entire orientation error is assumed to act about the axis of this shorter side length. At an initial distance d_{init} between UAV and concentrator, the maximum angular deviation δ of the initial orientation estimate which still ensures visibility of the reflection is hence

$$|\delta| \leq \arctan\left(\frac{l_{\text{conc}}}{2d_{\text{init}}}\right) . \quad (22)$$

Note that for concentrators comprising multiple canted facets, the extent of the reflection corridor depends on the facet canting angles (i.e. on the focal length of the concentrator), with the reflection corridor's extent decreasing with increasing distance. However, during Stage 1 the UAV operates at comparatively short distances from the concentrator. Hence, the assumption of a reflection corridor with approximately constant spatial extent remains a reasonable approximation.

The initial concentrator normal vector estimate can e.g., be obtained from a preceding flight mission, running ARTSCORE-C [7] on images recorded at larger distances between the camera and the heliostat field. Once a potential heliostat outlier is identified in this way, the UAV can be positioned closer to the corresponding concentrator. As described in more detail in [7], a smaller distance between the camera and concentrator also benefits the ARTSCORE-C accuracy. In preliminary Monte-Carlo experiments, an upper bound on the ARTSCORE-C error (and hence a lower bound on the initial distance) is estimated by considering the case where only one image is used to estimate the concentrator’s initial orientation. The results show that for the concentrator dimension at the Solar Tower Jülich ($l_{\text{conc}} = 1.28 \text{ m}$) and the DJI Mavic 3 Enterprise in video mode (image resolution of $3840 \text{ px} \times 2160 \text{ px}$ with a horizontal field of view of 71°), an initial distance of $d_{\text{init}} = 15 \text{ m}$ is sufficiently small to guarantee an orientation estimation error below the limit provided by Eq. (22). However, larger initial distances are achievable by averaging over a collection of images to enhance the initial orientation estimate. This collection can stem from both a flight mission specifically targeted to obtain accurate concentrator orientations (as presented in [7]) or from the flight segment where the camera approaches the concentrator.

In further preliminary experiments with different concentrator geometries, we observed that the required initial distance roughly scales linearly with the concentrator side length l_{conc} . For example, for commercial-scale power plants with $l_{\text{conc}} = 10\text{ m}$, we found an initial distance of $d_{\text{init}} = 60\text{ m}$ to be sufficient.

3.2. Framework for Stages 2, 3, and 4

The Stages 2, 3, and 4 are governed by a controller that operates at discrete, equitemporal timesteps. At each timestep, the controller receives an input signal and, after an inherent processing delay, returns the corresponding control signal. Details on the signal transmission behavior are provided in [Section 3.2.1](#). The remainder of this subsection describes the computations performed within a single timestep. The controller is conveniently implemented in the CCS, where the control signal




Fig. 6. Overview of the framework used to navigate the UAV during Stages 2, 3, and 4. The controller exchanges uplink and downlink signals with the UAV at discrete time steps. For simplicity, time indices are omitted, i.e., all quantities denote their values at a given time step.

is decomposed into two components in the concentrator coordinate system XY-plane and one component along the concentrator normal vector, as described in [Section 3.2.2](#). While the controller operates in the CCS, its input reflection signal is obtained in the ICS, whereas its output signal can only be interpreted in the GCS by the UAV's autopilot software. [Sections 3.2.3](#) and [3.2.4](#) describe how the required coordinate transformations from the ICS to the CCS and from the CCS to the GCS can be estimated, respectively. An overview of the resulting framework for the Stages 2, 3, and 4 is illustrated in [Fig. 6](#).

3.2.1. Signal transmission

In this work, the perception module and control logic are assumed to run on a ground station, communicating with the UAV via a wireless data link. This setup relaxes onboard hardware requirements and simplifies debugging and the implementation of safety mechanisms by providing greater control over which commands are sent to the UAV. However, these advantages come at the cost of additional communication delays. In principle, the same system can be implemented with all data processing and control algorithms running directly on an onboard companion computer. This step is suitable for later development stages once the system's safety and behavior are fully evaluated.

The signal propagation from image capture to availability at the ground station is assumed to take a fixed downlink time t_{down} . Upon receiving the image and the UAV position data, the ground station performs the necessary computation to produce a control command, requiring a processing time t_{comp} . The resulting control command is then transmitted back to the UAV, which takes an additional uplink

time t_{up} . Hence, the total delay between image capture and execution of the corresponding control command on the UAV side is

$$t_{delay} = t_{down} + t_{comp} + t_{up} . \quad (23)$$

The system operates at equitemporal timesteps nt_{sys} , where n indexes the iterations and the system time t_{sys} denotes the time interval between two iterations. For real-time operation, the computation time must therefore satisfy

$$t_{comp} < t_{sys} . \quad (24)$$

This condition ensures that no computational backlog accumulates over successive timesteps. Note that the overall delay t_{delay} may exceed the system time t_{sys} , as the signal may remain in transit for multiple cycles, but the system remains stable as long as each new computation finishes within a single iteration interval.

3.2.2. Control in the CCS

Although the reflection signal is initially detected in the ICS, it is expressed with respect to the concentrator coordinate system XY-plane (CCSXY-plane) using a homography, as described in Section 3.2.3. This design choice allows the controller to operate directly in the CCS in metric units rather than pixels, enabling the guidance of the reflection signal p_{refl}^{CCSXY} toward a desired setpoint s_{refl}^{CCSXY} during Stages 2, 3, and 4. Velocity control is a natural choice for this task: The reflection position is measured relative to the concentrator corners and can therefore be influenced directly through corresponding velocity commands of the UAV. This approach operates independently of the UAV's absolute position and does not require a transformation of the reflection setpoint into a position target. Instead, a velocity control command expressed in the CCSXY-plane produces an immediate and intuitive shift of the reflection on the concentrator surface. In addition, the distance between the UAV and the concentrator can be easily controlled by adding a velocity component along the CCS Z-axis.

In mathematical terms, the controller receives the error signal

$$e_{refl}^{CCSXY}(nt_{sys}) = s_{refl}^{CCSXY}(nt_{sys}) - p_{refl}^{CCSXY}(nt_{sys}) \quad (25)$$

and outputs the velocity command

$$v_{control}^{CCSXY}(nt_{sys}) = C_{\kappa} \left(e_{refl}^{CCSXY}(nt_{sys}) \right) , \quad (26)$$

where the function $C_{\kappa}(\cdot)$ denotes the control law, given the controller parameters κ . In principle, a proportional velocity controller should be sufficient to drive the reflection error to zero, since the commanded velocities are implicitly integrated by the UAV's dynamics. However, tests in simulation reveal that a purely proportional controller leads to a steady displacement of the reflection signal from the setpoint whenever vertical and horizontal motion are commanded simultaneously. The magnitude of this deviation is observed to depend on both the commanded normal velocity and the concentrator elevation angle. A plausible explanation for this coupling effect, together with supporting simulation results, is provided in Appendix B.

To compensate for this steady-state error, an integral term is introduced. The integral state i_{refl}^{CCSXY} is updated at each iteration n according to

$$i_{refl}^{CCSXY}(nt_{sys}) = i_{refl}^{CCSXY}((n-1)t_{sys}) + e_{refl}^{CCSXY}(nt_{sys}) . \quad (27)$$

The resulting proportional-integral (PI) control law is then given by [16]

$$v_{control}^{CCSXY}(nt_{sys}) = k_p \cdot e_{refl}^{CCSXY}(nt_{sys}) + k_i \cdot i_{refl}^{CCSXY}(nt_{sys}) . \quad (28)$$

For brevity, the explicit time dependence will be omitted in the remainder of this work whenever no ambiguity arises. The control parameters

are of the form

$$\kappa = \begin{pmatrix} k_p \\ k_i \end{pmatrix} . \quad (29)$$

While the control law could be extended by a derivative term, this is not pursued in the present work. The considered system shows discontinuities in the reflection feedback signal, as reflections may appear or disappear abruptly with small changes in the UAV pose (see Appendix D for more details). Consequently, the reflection signal can show large variations between successive sampling instants, even for small changes in the UAV position. In practice, derivative action therefore requires filtering to avoid amplification of such effects as well as measurement noise [16], which in turn introduces additional dynamics and complicates both analysis and tuning. For these reasons, a derivative term is not considered here.

Alternatively, the steady-state error could be tackled by means of a feedforward compensation based on a model of the coupling effect [17]. However, this would require prior knowledge or identification of the steady-state behavior, which varies across different UAVs and operating conditions. In contrast, the integral term provides a model-free compensation mechanism that generalizes across different UAVs. It is emphasized that the objective of this work is not to derive an optimal controller, but rather to demonstrate a robust solution for the considered task. To this end, a PI-controller is implemented, suitable parameters are identified and the resulting performance is evaluated experimentally. Further improvements to the controller are possible but lie beyond the scope of this work.

While the control command $v_{control}^{CCSXY}$ lies in the CCSXY-plane, a velocity component $v_{control,Z}^{CCS}$ along the concentrator normal vector can be added in the Z-component:

$$v_{control}^{CCS} = \begin{pmatrix} v_{control}^{CCSXY} \\ v_{control,Z}^{CCS} \end{pmatrix} . \quad (30)$$

In Stages 2 and 4, the normal velocity component $v_{control,Z}^{CCS}$ is set to zero. This way, the reflection position p_{refl}^{CCSXY} can be moved to the desired setpoint s_{refl}^{CCSXY} on the concentrator surface while maintaining an approximately constant distance between the UAV and the concentrator. For Stages 2 and 3, the setpoint is preferably set to the concentrator center

$$s_{refl}^{CCSXY} = (0 \text{ m}, 0 \text{ m})^T . \quad (31)$$

This way, the reflection maintains the maximum distance to each concentrator edge, minimizing the probability that the reflection signal gets lost when increasing the distance between the UAV and the concentrator in the subsequent Stage 3. For Stage 4, the setpoint can be set to any location on the concentrator surface for which a measurement is desired.

In Stage 3, an additional normal component $v_{control,Z}^{CCS} > 0 \text{ m/s}$ serves to increase the distance between the UAV and the concentrator. While the normal movement steers the UAV away from the concentrator along the estimated normal vector, simultaneous control in the CCSXY-plane maintains the reflection position relative to the concentrator corners, compensating for unintended drift of the reflection across the concentrator surface. Such drift can arise due to inaccuracies in the normal vector estimate, introducing a lateral displacement of the reflection position, since a part of the normal velocity component leaks into the CCSXY-plane. In addition, external influences such as wind or the aforementioned drift due to simultaneous vertical and horizontal movement can displace the reflection signal away from the setpoint.

3.2.3. Transformation from ICS to CCS

The controller expects the reflection signal p_{refl}^{CCSXY} defined with respect to the concentrator coordinate system XY-plane. However, all image coordinates of the relevant quantities detected by the perception




Fig. 7. Overview of the signal flow used to estimate the concentrator orientation $\hat{R}_{CCS}^{GCS}$, which is used to transform the controller output signal from the CCS to the GCS. The optional normal vector refinement step may be employed to further improve the accuracy of the estimated orientation.

module are expressed with respect to the ICS. To transform the coordinates from the ICS to the CCSXY-plane, the homography H_{ICS}^{CCSXY} is determined based on the detected concentrator corners P_{corner}^{ICS} and the design concentrator geometry $(P_{corner}^{CCS})_{XY}$. The reflection position p_{ref}^{ICS} is then expressed with respect to the CCSXY-plane via

$$p_{ref}^{CCSXY} = H_{ICS}^{CCSXY} \cdot p_{ref}^{ICS}. \quad (32)$$

3.2.4. Transformation from CCS to GCS

The transformation from the CCS to the GCS requires an estimate of the concentrator orientation $\hat{R}_{CCS}^{GCS}$, which comprises three rotational degrees of freedom. Two degrees of freedom are determined by the concentrator normal vector, while the third corresponds to a rotation about the concentrator optical axis. The complete estimation procedure, including the combination and refinement of partial estimates, is detailed below. The corresponding signal flow is shown in Fig. 7.

Normal vector estimation. In a first step, the concentrator normal vector is approximated using the known heliostat position o^{GCS} and the UAV position p_{UAV}^{GCS} , i.e.,

$$\hat{n}_{CCS}^{GCS} = \frac{p_{UAV}^{GCS} - o^{GCS}}{\|p_{UAV}^{GCS} - o^{GCS}\|}. \quad (33)$$

This approximation assumes that, as long as a reflection signal is visible, the vector from the heliostat position to the UAV provides a sufficiently accurate estimate of the concentrator normal vector.

This vector is then used to construct a rotation matrix whose Z-axis is aligned with the concentrator normal and whose X-axis lies parallel to the ground plane:

$$\hat{R}_{normal}^{GCS} = \text{ROTATIONFROMNORMAL}(\hat{n}_{CCS}^{GCS}). \quad (34)$$

Further details on the construction of this rotation matrix are provided in Appendix C.

Rotation about the optical axis. Depending on the heliostat kinematic system, the concentrator surface may be additionally rotated about its optical axis by an angle ψ , defined as the angle at which the line described by the concentrator's lower edge intersects the ground plane. This remaining rotation can be estimated from the detected concentrator corner points, i.e.,

$$\hat{\psi} = \text{ESTIMATE}_{\psi}(P_{corner}^{ICS}), \quad (35)$$

described in more detail in Appendix C.

Concentrator orientation estimate. Combining the two steps yields the concentrator orientation estimate

$$\hat{R}_{CCS}^{GCS} = \hat{R}_{normal}^{GCS} \cdot R_Z(\hat{\psi}), \quad (36)$$

where the matrix $R_Z(\cdot)$ denotes a rotation about the Z-axis. This orientation estimate does not require any information from the reflection signal. It relies solely on the known heliostat position, the estimated UAV position and quantities derived from the recorded image.

Normal vector refinement. A more accurate concentrator orientation estimate can be obtained by iteratively refining the normal vector using the reflection position expressed in the GCS. As described in Eq. (15), this position is obtained by transforming the observed reflection signal to the GCS using a concentrator pose estimate.

While Eq. (36) provides the rotational component of this pose, the translational component can be estimated using the known heliostat kinematic system and the current normal vector estimate in a backward pass (see Eq. (9)), as illustrated in Fig. 7. This iterative refinement procedure is referred to as normal vector refinement (NVR) in this work and can be repeated until convergence of the normal vector estimate. Empirical results indicate that two iterations are sufficient to achieve convergence.

The application of this refinement is optional, as it introduces additional computational effort and requires more *a priori* knowledge. Its impact on system performance is therefore evaluated explicitly in the evaluation studies presented later in this work.

3.3. Normal velocity profile shaping in Stage 3

When the UAV is commanded with an instantaneous velocity step in the normal direction during Stage 3, large displacements of the reflection position are observed due to the transient response associated with the required acceleration. As long as the reflection signal remains visible, the controller can compensate for these displacements after a short time. However, strong transient displacements may lead to a loss of the reflection signal, particularly for large commanded normal velocities $v_{control,Z}^{CCS}$.

To mitigate this effect, normal velocity profile shaping (NVPS) is applied: Instead of commanding the maximum normal velocity instantaneously, the commanded velocity is shaped using the cubic function

$$s(x) = 3x^2 - 2x^3, \quad \text{for } 0 \leq x \leq 1, \quad (37)$$

which yields a smooth velocity profile with continuous acceleration and deceleration phases. More precisely, the commanded normal velocity at the discrete time steps $t = n t_{sys}$ is defined as

$$v_{control,Z}^{CCS}(t) = \begin{cases} 0 \text{ m/s}, & t < 0 \text{ s}, \\ v_{Z,max}^{CCS} \cdot s\left(\frac{t}{t_{rise}}\right), & 0 \leq t < t_{rise}, \\ v_{Z,max}^{CCS}, & t_{rise} \leq t < t_{rise} + t_{hold}, \\ v_{Z,max}^{CCS} \left[1 - s\left(\frac{t - (t_{rise} + t_{hold})}{t_{rise}}\right)\right], & t_{rise} + t_{hold} \leq t < 2t_{rise} + t_{hold}, \\ 0 \text{ m/s}, & t \geq 2t_{rise} + t_{hold}. \end{cases} \quad (38)$$




Fig. 8. Exemplary normal velocity profile according to Eq. (38) for $v_{z,max}^{CCS} = 3$ m/s, $t_{rise} = 10$ s, and $t_{hold} = 15$ s, illustrating the smooth acceleration, constant-velocity, and deceleration phases introduced by the normal velocity profile shaping.

Fig. 8 shows an exemplary normal velocity profile. When combining normal velocity profile shaping with the previously described PI-controller without further modification, an additional effect was observed: After the UAV comes to a stop, the integral action accumulated due to the preceding non-zero normal velocity remains active. As a result, even if the reflection signal is already approximately aligned with its setpoint, the controller produces a second transient response. This additional response does not improve alignment but increases the risk of losing the reflection signal.

To mitigate this behavior, the integral gain is modulated depending on the current normal velocity component: The smooth profile described in Eq. (38) is normalized to take values between zero and one and is used as a weighting factor for the integral term. The influence of the integral term is thus increased during phases with non-zero normal motion, where it is required to compensate for coupling effects, and reduced once the normal velocity component returns to zero, where integral action is no longer necessary. Since the UAV does not track the commanded velocity instantaneously, the weighting profile is slightly shifted in time to account for the UAV's delayed response. This gain-scheduling approach proved sufficient to suppress the undesired transient in the considered experiments. Further refinements of the controller design are left for future work.

The distance traveled during Stage 3 in the normal direction is given by the area under $v_{control,Z}^{CCS}(t)$. With

$$\int_0^1 s(x) dx = \frac{1}{2}, \quad (39)$$

and assuming that the UAV approximately follows the normal velocity command, the traveled distance is

$$d_{S3} \approx \int_0^{2t_{rise}+t_{hold}} v_{control,Z}^{CCS}(t) dt = v_{z,max}^{CCS} \cdot (t_{rise} + t_{hold}). \quad (40)$$

This work uses two different operating modes:

1. **Fixed Time:** Choose t_{rise} and t_{hold} , and the resulting distance is $d_{S3} \approx v_{z,max}^{CCS} \cdot (t_{rise} + t_{hold})$.
2. **Fixed Distance:** For a desired distance d_{S3} , choose t_{rise} and compute

$$t_{hold} = \frac{d_{S3}}{v_{z,max}^{CCS}} - t_{rise}, \quad t_{hold} \geq 0. \quad (41)$$

3.4. Considerations for multi-faceted concentrators

The system presented in the previous subsection is formulated for a single reflection signal p_{ref}^{CCSXY} as an input. However, as described




Fig. 9. Simplified two-dimensional view of reflection corridors for a concentrator composed of two flat, canted mirror facets. The figure illustrates two effects: (i) For UAV–concentrator distances below $2f$ (points 1 and 2), the reflection signal can be used to reliably guide the UAV toward the desired setpoint (cyan arrow), whereas larger distances may result in the loss of the reflection signal (point 3). (ii) Assuming symmetric canting and perfectly planar facet surfaces, averaging the facet normal vectors yields an accurate estimate of the overall concentrator normal (point 2). If only a non-symmetric subset of reflections is available (such as at points 1 and 3) the estimated normal vector becomes biased.

in Section 2.5.4, a multi-faceted concentrator may produce multiple reflection signals simultaneously. This subsection analyzes how such multi-reflection signals can be reduced to a single reflection signal. More precisely, it is shown that averaging all visible reflections constitutes a suitable strategy for Stage 3, provided that the distance between the UAV and the concentrator remains within a certain operating range defined by the concentrator focal length. The reflection signal is used in two components of the system:

1. Within the CCS-controller, the reflection signal determines the velocity command that guides the UAV.
2. When applying the normal vector refinement, the reflection signal is required to transform velocity commands from the CCS to the GCS.

The effect of reflection averaging in both components is illustrated in Fig. 9 and discussed in what follows.

3.4.1. Implications for the CCS-controller

The controller computes velocity commands such that the reflection signal is driven toward a specified setpoint. In Stage 2 and 3, this setpoint corresponds to the concentrator center. As illustrated in Fig. 9, this approach is effective for UAV positions closer than $2f$ (points 1 and 2): Although the averaged reflection signal decreases in magnitude with increasing distance (shorter blue arrow for point 2 than for point 1), the resulting velocity command remains qualitatively correct. However, for UAV positions beyond $2f$ (point 3), this behavior no longer holds: Despite the UAV being positioned in the upper half of the concentrator surface, the reflection signal appears on the lower half of the surface, causing the velocity command to drive the UAV further upward. As a result, the reflection signal may be lost entirely in the next iteration.

This behavior illustrates that, under the assumption of canted but approximately flat facets, the effective operating range of the proposed system is bounded by

$$0 \text{ m} \leq d \leq 2f. \quad (42)$$

Further details on the mapping between the UAV position and the reflection signal are discussed in Appendix D, also considering additional surface slope errors that are present in real scenarios.




Fig. 10. Simulated quantities in the ICS for a concentrator as installed at the Solar Tower Jülich. Each concentrator consists of $M = 4$ canted facets. If multiple reflections are observed, the combined reflection signal $P^{\text{ICS}}_{\text{refl}}$ is obtained by averaging all visible reflections. Due to the kinematic configuration, concentrators at the Solar Tower Jülich may appear rotated with respect to the ICS. The rotation angle ψ can be estimated directly from the corner point image coordinates $P^{\text{ICS}}_{\text{corner}}$, as described in more detail in [Appendix C](#).

3.4.2. Implications for the normal vector refinement-based transformation

The second use of the reflection signal arises when applying the normal vector refinement, which requires the reflection signal to transform the velocity command from the CCS to the GCS. To navigate the UAV away from the concentrator while maintaining visibility of multiple reflection signals, the UAV must remain within the intersection region of the corresponding reflection corridors. A navigation strategy that treats all facets symmetrically is best formulated with respect to the concentrator normal vector.

Assuming symmetric canting, this normal vector can be obtained by the vector from the averaged reflection signal to the UAV position (point 2 in [Fig. 9](#)): Averaging the individual reflection signals is equivalent to computing and averaging the corresponding facet-wise normal vectors. Even when only a subset of facets show reflections, the concentrator normal vector can still be approximated by the vector from this subset of reflections to the UAV position: As shown for points 1 and 3 in [Fig. 9](#), the estimated navigation plane remains approximately correct, since at sufficiently large distances the estimated concentrator normal vector is only weakly misaligned with the true concentrator normal vector.

For Stage 4, it may be advantageous to focus on a single facet m by considering only the reflection signal associated with this facet, $P^{\text{ICS}}_{\text{refl},m}$, while disregarding all other observed reflections. In this case, the facet corners can be used to express the setpoint with respect to the XY-plane of the corresponding facet coordinate system FCS_m . All relevant quantities in the ICS are illustrated in [Fig. 10](#) for a concentrator representative of those at the Solar Tower Jülich.

3.5. Stage 4

Stage 4 aims at accurately measuring the concentrator normal vectors over a sufficiently large portion of the concentrator surface. Conceptually, this stage is addressed by the *HelioPoint* method [6]. Hence, the present work does not implement or evaluate Stage 4 in detail. Instead, the focus lies on establishing the navigation framework for the Stages 2, 3, and 4. Nevertheless, the present subsection outlines how Stage 4 could be realized within this framework and summarizes relevant insights from the *HelioPoint* method, in particular regarding the required surface coverage under the assumption of point-like reflection measurements.

The navigation during Stage 4 can continue to employ the approximation [Eq. \(17\)](#) of a co-located camera and LED, while the actual *HelioPoint* measurement in Stage 4 can utilize a more accurate geometric model to improve the measurement accuracy by explicitly accounting for the spatial separation between the camera and the light source, as described by Krauth et al. [6].

An important question for Stage 4 is what constitutes sufficient coverage of the concentrator surface. Prior work by Krauth et al. [6] represents the reflections as point-like features and indicates that on the order of 40–60 evenly distributed sampling points per concentrator are sufficient to achieve high calibration accuracy for the considered heliostats at the Solar Tower Jülich.

If reflections are instead treated as spatially extended regions, the definition of sufficient surface coverage may change, as individual observations could provide information about multiple surface locations simultaneously: For example, if a reflection is described as a collection of pixels, a normal vector could be estimated for each of these pixels instead of assigning a single normal to the reflection centroid. Determining how such region-wise estimates can be consistently aggregated across multiple observations remains an open question for future work.

The proposed method may also provide additional information about facet alignment. In the original *HelioPoint* approach, measurements are performed along predefined flight paths, which do not guarantee uniform coverage of all facets, particularly at larger distances. In extreme cases, misaligned facets may not show reflections at all. In contrast, the closed-loop nature of the proposed method enables the UAV to adapt its trajectory based on the observed reflection signals. For example, the loss of reflections on individual facets during motion may indicate facet canting errors. The adaptive capabilities of the proposed method could therefore be used to ensure sufficient coverage of all facets and to identify inconsistencies in their alignment. While such strategies are not explicitly implemented in the present work, they represent a promising direction for future research.

3.6. Practical considerations

While the navigation framework of ARTSCORE-F is evaluated in a simulation environment as part of this work, its assessment in a real heliostat field lies beyond the scope of the present work. The simulation allows for a controlled analysis of the method, assuming that all required quantities (namely the reflection positions, the concentrator corner positions, and the UAV position) are directly and reliably available. In practice, these quantities must be obtained from recorded image data and onboard sensing, and their accurate estimation is a key requirement for real-world deployment. The purpose of this subsection is to discuss practical challenges that could arise when transferring ARTSCORE-F from simulation to real-world operation, and to outline how such challenges may be addressed.

Detection of the concentrator corner points is not expected to pose a major challenge, as robust AI-based perception modules have already been developed for this task. In particular, the approach proposed by Broda et al. [18] has been successfully applied in ARTSCORE-C [7]. Similarly, the UAV position can be estimated using hardware-based methods such as RTK, which have demonstrated positioning uncertainties on the order of a few centimeters [6]. Alternatively, optical methods such as ARTSCORE-C can be employed in scenarios where no RTK correction signal is available.

The remaining challenges to ensure the robustness of the proposed method are therefore the reliable detection of the reflection signals, and the ability to handle cases where these signals are lost. In real-world applications, such losses may occur for different reasons, which can be broadly categorized into two types.

- (i) *Perception-related losses* occur when a reflection is physically visible given the geometric configuration but is not detected by the perception module. This may be caused by low contrast

of the reflection, for example due to unfavorable lighting conditions or local surface slope deviations. Soiling can further impair reflection visibility by reducing intensity and contrast. However, practical soiling tests indicate that the reflection remains visible for moderately soiled mirrors, and that a substantial degree of soiling would be required before the reflection becomes undetectable.

- (ii) *Geometry-related losses* occur when the UAV is not located within the reflection corridor of a facet, such that the reflection is not visible due to the geometric configuration. This may result from transient deviations during UAV motion (e.g., acceleration or deceleration) or from errors in the estimated normal vector leading to incorrect positioning of the UAV. External disturbances, in particular wind, can further contribute to deviations from the desired flight path. Modern UAV autopilot systems are generally capable of compensating for moderate and slowly varying wind disturbances. Remaining drifts of the reflection position can be compensated by the navigation controller through the adaptation of the UAV trajectory. However, strong wind gusts may induce rapid UAV displacements that can lead to a loss of reflection visibility.

The system behavior depends on the number of visible reflections: As long as at least one reflection remains visible, the controller can continue to operate. While the estimated normal vector and the resulting CCS may change slightly, leading to small changes in the intended flight path, the system remains stable. More critical situations arise when all reflection signals are lost. To handle such cases, it is beneficial to maintain a history of detected reflection signals together with the corresponding UAV positions. This information can be used to identify the likely cause of the reflection loss and to select an appropriate recovery strategy.

If reflections are observed to drift toward the concentrator edges before disappearing, this indicates a geometry-related loss. In this case, the direction of the drift provides information about how the UAV position must be adjusted to restore visibility, for example by moving in the opposite direction of the observed drift.

If no clear drift is observed, the loss may be perception-related or caused by an unfavorable combination of effects. In such cases, it may be beneficial to continue the motion based on predicted reflection positions derived from past observations, allowing the reflection to reappear. If this strategy fails, a robust fallback is to return the UAV to the last position at which reflections were reliably observed.

While geometry-related losses are captured in the simulation studies, the present work assumes idealized perception, i.e., a reflection is detected whenever it is physically visible and its position is computed without additional noise. Future work could incorporate a perception noise model accounting for effects such as false positives, false negatives, and measurement noise. Tests with the current AI-based perception model on image data recorded at the Solar Tower Jülich indicate low error rates. Occasional perception-related losses are therefore expected to be rare and could likely be mitigated using predictions based on the history of previously detected reflections, as discussed above.

Similarly, other sources of measurement noise, such as UAV positioning or reflection localization uncertainty, are not explicitly modeled. These effects are expected to be small compared to the dominant geometric approximations considered in this work, in particular the omission of the normal vector refinement (see Section 3.2.4). A more detailed analysis of these effects is left for future work.

3.7. Implementation

The proposed method is implemented in PYTHON. Coordinate transformations rely on the packages NUMPY [19] and SCIPY [20] for efficient vectorized computation. Computer vision tasks are performed using the OPENCV-PYTHON package, a PYTHON wrapper for the open-source computer vision library OPENCV [21]. The simulation environment runs a

software-in-the-loop instance of ARDUPILOT [22] and therefore expects flight commands in the MAVLINK protocol [23]. The proposed method utilizes PYMAVLINK to encode velocity control commands in the required format.

4. Tuning and evaluation methodology

The system architecture described in Section 3 relies on a set of parameters, referred to in the following as the *system configuration*. The system behavior for a given system configuration depends on the conditions in which it operates. For example, a larger concentrator surface permits greater deviations of the reflection signal from the concentrator center before the signal is lost, compared to a smaller surface. The set of parameters characterizing the conditions under which the system operates is referred to as the *boundary conditions*.

The process of finding a suitable system configuration that performs well under a relevant set of boundary conditions is referred to as *tuning*. Rather than assuming an analytical model of the system behavior, the tuning process is conducted using the simulation environment proposed in [10], which serves as a high-fidelity representation of the physical environment and enables systematic evaluation of the system behavior under controlled and repeatable conditions. Within this context, tuning is treated as an empirical exploration of system configurations under varying boundary conditions. Two main challenges arise in this approach:

1. The combined space of possible system configurations and boundary conditions is high-dimensional, making exhaustive simulation infeasible.
2. The tuning should account for general physical principles and establish limitations and trade-offs that are independent of a particular heliostat type or concentrator properties. At the same time, the resulting system configuration should perform robustly across different heliostat kinematic systems, concentrator orientations, dimensions, and reflective properties.

To address the first challenge, the exploration of the parameter space is restricted in two ways: (i) Certain parameters are fixed throughout this work because they are either physically constrained or symmetric (such that evaluating a single representative case is sufficient). (ii) Only a limited subset of parameters is varied jointly during tuning, focusing on parameters that are expected to interact based on physical considerations and dependencies identified during the system design phase. While this selection relies on engineering judgment, it enables a tractable exploration of the most influential dimensions of the parameter space.

To address the second challenge, the tuning procedure considers only boundary conditions that influence system behavior in a consistent manner across different scenarios, i.e., boundary conditions that reflect general effects rather than heliostat- or heliostat field-specific properties. For example, the choice of the integral controller gains has to consider different concentrator elevation angles, since the coupling of horizontal and vertical movement introduces an elevation angle-dependent velocity drift. For this reason, the elevation angle is considered already during tuning. On the other hand, some properties such as the concentrator dimensions may vary from field to field. Some properties even vary across different heliostats in the same field, such as the reflective properties. Based on this distinction, the tuning process identifies a robust system configuration by varying only those boundary conditions that capture the essential physics common to a large class of real-world scenarios. The remaining boundary conditions are idealized during tuning: The underlying hypothesis is that configurations obtained under these idealized conditions remain robust to moderate deviations from the idealized setting. This hypothesis is tested in the *evaluation* phase by testing the tuned configuration under altered boundary conditions that more closely reflect realistic environments.

The remainder of this section is structured as follows. Section 4.1 describes how the system behavior can be simulated and its performance

Table 2

Overview of the parameters defining the system configuration and the boundary conditions throughout the tuning and evaluation phases. Some parameters remain fixed throughout this work, while others are systematically varied during tuning and/or evaluation.

Variable	Description	Varied During	Comment	Type
$s_{\text{refl}}^{\text{CCSXY}}$	Reflection Setpoint	Fixed to $(0 \text{ m}, 0 \text{ m})^T$	Stages 2, 3, 4	System Configuration
k_p	Proportional Controller Gain	Tuning		
k_i	Integral Controller Gain	Tuning		
$v_{z,\text{max}}^{\text{CCS}}$	Maximum Normal Velocity	Tuning + Evaluation	Normal Velocity Profile Shaping (Stage 3)	
t_{rise}	Sweep Time	Tuning + Evaluation		
t_{hold}	Hold Time	Tuning + Evaluation		
NVR	Normal Vector Refinement	Evaluation	<i>Applied or Not Applied</i>	
t_{sys}	System Time	Fixed to 0.5 s	Signal Transmission	Boundary Conditions
t_{delay}	Delay Time	Fixed to 0.6 s		
$p_{\text{UAV,init}}^{\text{CCS}}$	UAV Initial Position	Tuning	Only relevant in Stage 2	
o^{GCS}	Heliostat Position	Fixed to $(0 \text{ m}, 0 \text{ m}, 1.5 \text{ m})^T$	Concentrator Pose Characterization	
az	Azimuth Angle	Fixed to 0°		
el	Elevation Angle	Tuning + Evaluation		
ψ	In-Plane Rotation	Evaluation		
S	Concentrator Surface Characterization	Evaluation	Idealized During Tuning	

quantified, given a set of parameters for both system configuration and boundary conditions. Sections 4.2 and 4.3 describe the parameters defining the system configuration and the boundary conditions, and indicate which ones are fixed and which are varied during tuning and evaluation. Table 2 summarizes both subsections in a single table to provide a clear overview of all relevant parameters and their respective treatment. Sections 4.4 and 4.5 describe the tuning studies and the evaluation studies in detail, respectively.

4.1. Experimental framework

4.1.1. Simulation runs and studies

The system is tested within the simulation environment presented in [10], which simulates the quantities observed by a camera that moves in a virtual heliostat field, as well as the camera pose as a function of UAV control commands: At each discrete timestep, ARTSCORE-F receives the simulated input signals and outputs a velocity command. This command is passed to the simulation environment with a simulated delay (see Section 4.3.1), which updates the UAV state and recomputes the input signal. The resulting closed-loop system can be executed either in fixed-time mode or in fixed-distance mode, and the sequence of signals and system states generated during this execution constitutes the outcome of one *simulation run*. In contrast, a *study* in this work refers to a set of simulation runs in which one or more parameters are systematically varied while all others are held constant.

For each simulation run in a study, the closed-loop system produces a time-discrete sequence of reflection signals. Depending on the concentrator surface model, one or multiple facets may be simulated such that one or multiple reflection signals may be visible at a given timestep. For the case of a single-facet concentrator, let

$$p_{\text{refl}}^{\text{CCSXY}}(n t_{\text{sys}}) \in \mathbb{R}^2, \quad n = 1, \dots, N_{\text{study}}, \quad (43)$$

denote the 2D reflection signal at discrete timestep n for a study with a total of N_{study} timesteps. For the case of a multi-faceted concentrator, let

$$P_{\text{refl}}^{\text{CCSXY}}(n t_{\text{sys}}) \in \mathbb{R}^{2 \times M}, \quad n = 1, \dots, N_{\text{study}}, \quad (44)$$

denote the 2D reflection signal at discrete timestep n for a study with a total of N_{study} timesteps and a concentrator with M facets. The collection of all 2D reflection signals for a given study is referred to as the *reflection trajectory*.

4.1.2. Error metrics

Throughout this work, the system performance is evaluated using discrete-time metrics based on the reflection trajectory. As defined

in Eq. (25), the reflection error is the difference between setpoint the $s_{\text{refl}}^{\text{CCSXY}}$ and the reflection signal $p_{\text{refl}}^{\text{CCSXY}}$. Since the setpoint is chosen as $s_{\text{refl}}^{\text{CCSXY}} = (0 \text{ m}, 0 \text{ m})^T$ throughout this work (see Eq. (31)), the reflection signal and reflection error differ only in sign and are otherwise identical. Consequently, the terms *reflection signal* and *reflection error* are used interchangeably in the following. Several discrete-time metrics are used to quantify the performance of a given system configuration within given boundary conditions during both the tuning and evaluation phases.

ITAE (integral of time-weighted absolute error). The integral of time-weighted absolute error (ITAE) introduces a time-weighting factor to the absolute reflection [16],

$$\text{ITAE} = t_{\text{sys}} \cdot \sum_{n=1}^{N_{\text{study}}} (n t_{\text{sys}}) \cdot \|p_{\text{refl}}^{\text{CCSXY}}(n t_{\text{sys}})\|. \quad (45)$$

By penalizing reflection errors more strongly at later time instants, this metric reduces the influence of initial transients and emphasizes long-lasting deviations, making it suitable for assessing steady-state behavior.

ITAESS (integral of time-weighted steady-state-corrected absolute error). While the ITAE penalizes persistent offsets, it can be desirable to detect oscillatory behavior independently of steady-state reflections. To isolate such deviations, the steady-state reflection signal is subtracted, i.e.,

$$\text{ITAESS} = t_{\text{sys}} \cdot \sum_{n=1}^{N_{\text{study}}} (n t_{\text{sys}}) \cdot \|p_{\text{refl}}^{\text{CCSXY}}(n t_{\text{sys}}) - p_{\text{refl,SS}}^{\text{CCSXY}}\|, \quad (46)$$

where $p_{\text{refl,SS}}^{\text{CCSXY}}$ is obtained by averaging over the last N'_{study} time steps:

$$p_{\text{refl,SS}}^{\text{CCSXY}} = \frac{1}{N'_{\text{study}}} \cdot \sum_{k=N_{\text{study}}-N'_{\text{study}}+1}^{N_{\text{study}}} p_{\text{refl}}^{\text{CCSXY}}(k t_{\text{sys}}). \quad (47)$$

This metric emphasizes sustained oscillations around the steady-state reflection signal.

PAE (peak absolute error). The PAE captures the maximum reflection signal magnitude along the reflection trajectory,

$$\text{PAE} = \max_{n=1, \dots, N_{\text{study}}} \|p_{\text{refl}}^{\text{CCSXY}}(n t_{\text{sys}})\|. \quad (48)$$

This metric is sensitive to peaks in the reflection trajectory and is hence used in studies where the risk of losing the reflection signal is assessed.

Settling time. The settling time is defined as the first time instant after which the reflection signal magnitude remains below a threshold ϵ [16],

$$t_{\text{settling}}(\epsilon) := \min \left\{ n t_{\text{sys}} \geq 0 \text{ s} : \forall n' t_{\text{sys}} \geq n t_{\text{sys}}, \left\| p_{\text{refl}}^{\text{CCSXY}}(n' t_{\text{sys}}) \right\| \leq \epsilon \right\}. \quad (49)$$

Throughout this work, the threshold is fixed to $\epsilon = 0.1$ m. For brevity, the explicit dependence on ϵ is therefore omitted.

Note that this metric is used exclusively to tune and evaluate the system for Stage 3. Since the UAV requires a minimum deterministic time

$$t_{\text{arrival}} = 2 \cdot t_{\text{rise}} + t_{\text{hold}} \quad (50)$$

to complete Stage 3 by definition, the settling time metric is modified: After reaching the target position and commanding a normal velocity speed of zero, the simulation is continued for a specified duration to allow the controller to align the reflection signal properly. Consequently, the settling time is not measured from the start of Stage 3, but from the time instant following t_{arrival} . The effective settling time used during tuning and evaluation can be expressed as

$$t_{\text{S3}} = t_{\text{arrival}} + t_{\text{settling}}^{(\text{post-arrival})}, \quad (51)$$

where $t_{\text{settling}}^{(\text{post-arrival})}$ denotes the settling time determined with respect to a time origin shifted to t_{arrival} . This adjusted definition is implicitly used in all tuning and evaluation results.

4.2. System configuration

The system configuration is characterized by seven parameters: the reflection setpoint $s_{\text{refl}}^{\text{CCSXY}}$, the controller gains (k_p, k_i) , the normal velocity profile shaping parameters $v_{\text{Z,max}}^{\text{CCS}}$, t_{rise} & t_{hold} , and whether the normal vector refinement is applied.

As described in Eq. (31), the reflection setpoint is chosen to coincide with the concentrator center in Stages 2 and 3 in order to minimize the risk of losing the reflection signal. While Stage 4 requires the reflection setpoint to be positioned in different locations on the concentrator surface to perform localized measurements, the analysis of this stage is beyond the scope of this work. The reflection setpoint is hence kept fixed at the concentrator center throughout both tuning and evaluation.

The controller gains determine the closed-loop behavior during Stages 2, 3, and 4, ensuring that the reflection signal is guided toward and maintained close to the desired setpoint. They are determined during tuning with the objective of achieving robust performance in both Stage 2 and 3 across a range of normal velocities. The gains identified during tuning are subsequently held fixed during evaluation.

The normal velocity profile shaping parameters $v_{\text{Z,max}}^{\text{CCS}}$, t_{rise} , and t_{hold} are only relevant in Stage 3, where motion normal to the concentrator surface is introduced. As described in Eq. (41), these three parameters determine the total distance traveled by the UAV during Stage 3. Their values define a trade-off between how quickly Stage 3 can be completed and how safely it can be executed, in the sense that faster completion increases the risk of losing the reflection signal. To evaluate this trade-off under realistic operating conditions, the normal velocity profile shaping parameters are also varied during evaluation.

An evaluation of the normal vector refinement's impact on the system behavior is conducted only during the evaluation phase. During the tuning phase, an idealized model of the concentrator surface is assumed, along with exact knowledge of the normal vector, as detailed in the subsequent subsection.

4.3. Boundary conditions

4.3.1. Signal transmission characterization

In simulation, the perception module is replaced by geometric computations, where all processing runs on a single machine. Consequently,




Fig. 11. Illustration of the timeline used to emulate communication and processing delays of a real system in simulation. Control commands are executed after a fixed offset t_{off} relative to the time of image capture. The upper case shows a total delay $t_{\text{delay}} = t_{\text{up}} + t_{\text{comp}} + t_{\text{down}} < t_{\text{sys}}$, where the control command derived from an image captured in a given cycle is still executed within the same cycle. The lower case shows $t_{\text{delay}} = t_{\text{up}} + t_{\text{comp}} + t_{\text{down}} > t_{\text{sys}}$, where command execution moves to the subsequent cycle.

only a negligible computation and communication delay is present, since the computations required for simulation execute orders of magnitude faster than a real perception module and no physical communication link is involved. To realistically emulate the behavior of the actual system, in which the perception and control logic are subject to a total communication and processing delay time t_{delay} , this delay is modeled using a queue-based mechanism: At the beginning of each clock cycle, the relevant input quantities (i.e., the reflection signal $p_{\text{refl}}^{\text{ICS}}$, the concentrator corner points $p_{\text{corner}}^{\text{ICS}}$ and the UAV position $p_{\text{UAV}}^{\text{GCS}}$) are computed and fed to the controller. The output control command is pushed into a queue together with its timestamp. After a time offset of $t_{\text{off}} = t_{\text{delay}} \bmod t_{\text{sys}}$ within the same cycle, the newest command whose timestamp exceeds the delay threshold t_{delay} is executed. Note that if $t_{\text{delay}} < t_{\text{sys}}$, the offset simplifies to $t_{\text{off}} = t_{\text{delay}}$, meaning that the command derived from an image captured in a given cycle is still executed within that same cycle. Fig. 11 illustrates the resulting simulation timeline.

For this work, the signal transmission parameters are fixed to

$$t_{\text{sys}} = 0.5 \text{ s} \quad \text{and} \quad t_{\text{delay}} = 0.6 \text{ s}. \quad (52)$$

These values are informed by the typical times the perception module requires to process a given image, as well as the delay present in a typical wireless communication channel between the ground station and UAV. These values remain fixed throughout all studies in both tuning and evaluation.

4.3.2. Concentrator characterization

The simulation of the system input requires a model of the concentrator, characterized by its pose and surface description. In this work, a distinction is made between tuning and evaluation, which differ in the level of modeling fidelity applied to the concentrator.

Idealized concentrator pose. For tuning, the concentrator pose is simplified in order to decouple controller parameter selection from heliostat-specific kinematics and geometric details. In this mode, the concentrator pose is not obtained from a specific kinematic model. Instead, the concentrator plane is constructed by rotating the global XY-plane about the X-axis by the elevation angle el and translating it to intersect the heliostat origin o^{GCS} . The azimuth angle is fixed to the representative value

$$az = 0^\circ, \quad (53)$$

since the UAV can be aligned through its yaw angle to face the concentrator for any azimuth angle. Overall, the concentrator pose is obtained as

$$T_{\text{CCS}}^{\text{GCS}} = \begin{pmatrix} R_X(el) & o^{\text{GCS}} \\ \mathbf{0} & 1 \end{pmatrix} \quad (54)$$

Table 3

Overview of the system configuration and boundary condition parameters considered during the tuning studies, indicating for each study whether parameters are swept, fixed, or adopted from preceding tuning results.

Tuning Study	k_p	k_I	$v_{Z,max}^{CCS}$ [m/s]	t_{sweep} [s]	t_{hold} [s]	$p_{UAV,init}^{CCS}$ (id)	el [°]
1	0.1, ..., 0.8	0	0	0	Fixed Time: $t_{hold} = 25$ s	1, 2, 3, 4	45
2	0.1, ..., 0.8	0	1, 3, 5	0	Fixed Time: $t_{hold} = 50$ s	Concentrator Center	45
3	Best of Tuning Studies 1, 2 Value: 0.4	0, 0.01, ..., 0.1	1, 3, 5	0	Fixed Time: $t_{hold} = 50$ s	Concentrator Center	45
4	Best of Tuning Studies 1, 2 Value: 0.4	Best of Tuning Study 3 Value: 0.04	1, 3, 5	0	Fixed Time: $t_{hold} = 50$ s	Concentrator Center	0, 15, ..., 90
5	Best of Tuning Studies 1, 2 Value: 0.4	Best of Tuning Study 3 Value: 0.04	Fixed Distance: $d_{S3} \approx 100$ m NVPS Configurations C1–C23 (see Table 4)			Concentrator Center	Worst of Tuning Study 4 Value: 60

and is defined by the chosen elevation angle and the heliostat origin. This approach avoids propagating a desired normal vector through the full kinematic model and results only in a small positional shift of the plane compared to the use of a kinematic model, without affecting the behavior relevant for tuning.

Idealized concentrator surface. The concentrator surface used during tuning is idealized to further simplify the simulation: the surface is assumed to be perfectly planar and infinitely extended, ensuring that a reflection signal is always available regardless of the UAV position relative to the CCS. Under these assumptions, the reflection position corresponds to the projection of the camera position onto the CCSXY-plane

$$(p_{UAV}^{CCS})_{XY} = (T_{GCS}^{CCS} \cdot p_{UAV}^{GCS})_{XY}, \quad (55)$$

where the transformation T_{GCS}^{CCS} is obtained as the inverse of T_{CCS}^{GCS} .

Realistic concentrator pose. For evaluation, the concentrator pose is determined using the full kinematic model for the Solar Tower Jülich, which computes the transformation T_{CCS}^{GCS} from a desired concentrator normal vector, the heliostat origin o^{GCS} , and the kinematic parametrization k . The concentrator orientation can be fully described by the Euler angles azimuth angle az , elevation angle el , and in-plane rotation angle ψ , i.e.,

$$R_{CCS}^{GCS} = R_Z(az) \cdot R_X(el) \cdot R_Z(\psi), \quad (56)$$

A south-facing normal vector (i.e., azimuth $az = 0^\circ$) leads to the lower concentrator edge being parallel to the ground plane (i.e., $\psi = 0^\circ$) for nominal kinematic parameters. However, other azimuth directions result in nonzero in-plane rotations. To reproduce this effect while keeping $az = 0^\circ$, the kinematic parameter corresponding to a rotation of the heliostat pylon around the Z-axis is varied, allowing the simulation of effective in-plane rotations during evaluation.

Realistic concentrator surface. For evaluation, the concentrator's reflective properties are modeled as described in [10]: The surface is discretized using a collection of small tiles, where each tile is assumed to be a perfectly flat mirror surface. The tiles are oriented by considering the facet canting and a slope deviation map, measured using the QDEC-H system [24]. As a result, the reflection signal during evaluation is shifted by surface errors and can disappear once the UAV leaves the reflection corridor.

4.4. Tuning studies

Considering the fixed parameters and using the idealized concentrator pose and reflection model, the following parameters remain for

tuning: (i) controller gains: k_p, k_I , (ii) normal velocity profile shaping parameters: $v_{Z,max}^{CCS}, t_{rise}, t_{hold}$, (iii) concentrator pose characterization (elevation angle): el and (iv) UAV initial position: $p_{UAV,init}^{CCS}$. To further reduce the space of parameter combinations, only subsets of parameters are varied jointly. This yields a total number of five studies: Tuning Studies 1–4 focus on steady-state behavior during Stage 2 and Stage 3 and are run in fixed-time mode, whereas Tuning Study 5 is run in fixed-distance mode, as it explores the trade-off between the time required to complete Stage 3 and reflection signal displacements for different normal velocity profile shaping parameters. For consistency across all studies, the reported normal velocity parameter is always $v_{Z,max}^{CCS}$. In Tuning Studies 1–4, normal velocity profile shaping is not applied, corresponding to the case of a constant commanded normal velocity, i.e.,

$$v_{control,Z}^{CCS}(n t_{sys}) = v_{Z,max}^{CCS} \quad \text{for all } n t_{sys}. \quad (57)$$

In this case, the maximum normal velocity $v_{Z,max}^{CCS}$ directly corresponds to the commanded normal velocity throughout the entire study. In contrast, Tuning Study 5 applies normal velocity profile shaping, such that $v_{Z,max}^{CCS}$ denotes the maximum commanded normal velocity, while the temporal shaping is additionally defined by t_{rise} and t_{hold} .

The elevation angle is set to $el = 45^\circ$ in Tuning Studies 1, 2, and 3, where the coupling effect between horizontal and vertical movement is expected to be maximized, thus representing the most challenging scenario for the system. The identified system configuration parameters are confirmed in Tuning Study 4 for different elevation angles. An overview of the five studies for system tuning is provided in Table 3. The individual studies are described in more detail in what follows.

4.4.1. Tuning Study 1

Tuning Study 1 covers Stages 2 and 4, in which the UAV moves parallel to the concentrator surface. To identify an appropriate proportional controller gain, k_p is swept while the normal velocity profile shaping parameters (i.e., the maximum normal velocity $v_{Z,max}^{CCS}$ and the rise time t_{rise}) are set to zero. Since there is no coupling between horizontal and vertical motion, no velocity drift occurs and the integral gain k_I is set to zero. To ensure a robust choice of the proportional gain k_p for different starting positions (corresponding to different initial concentrator orientation estimates), Tuning Study 1 considers different initial UAV positions $p_{UAV,init}^{CCS}$, as shown in Fig. 12. The initial normal distance is fixed to $d_{init} = 15$ m.

4.4.2. Tuning Study 2

Tuning Study 2 repeats the sweep of the proportional gain k_p under the conditions of Stage 3, where the UAV moves away from the concentrator. In addition, the normal velocity $v_{Z,max}^{CCS}$ is swept to ensure that the choice of proportional gains is robust for different non-zero normal




Fig. 12. Illustration of the four starting points used in Tuning Study 1 (points 1–4), chosen such that they span all four quadrants of a concentrator at the Solar Tower Jülich, with varying lateral distances from the concentrator center (point c).

velocities. Again, the integral gain k_i is set to zero, since its analysis is the purpose of the subsequent study. Stage 3 is preceded by Stage 2, which aligns the reflection signal with the concentrator center. Hence, the initial UAV position $p_{UAV,init}^{CCS}$ is set to the concentrator center for this and all subsequent tuning studies.

4.4.3. Tuning Study 3

Tuning Study 3 addresses the steady-state reflection error arising from non-zero normal velocities in Stage 3: The integral gain k_i is swept for different non-zero normal velocities $v_{z,max}^{CCS}$. The proportional gain is fixed at $k_p = 0.4$, as selected from the preceding tuning studies (see Section 5.1), and the initial UAV position $p_{UAV,init}^{CCS}$ is again set to the concentrator center due to the alignment performed in Stage 2.

4.4.4. Tuning Study 4

The preceding studies identify the proportional gain k_p and the integral gain k_i for different normal velocities $v_{control,Z}^{CCS}$, assuming an elevation angle of $e_l = 45^\circ$. To confirm that the identified system configuration generalizes well to other elevation angles, the elevation e_l is swept in Tuning Study 4 for different normal velocities.

4.4.5. Tuning Study 5

The Tuning Studies 2, 3 and 4 explore the system behavior during Stage 3, with a focus on steady-state behavior in a fixed-time operating mode. While this approach serves to identify a robust choice of controller gains, the results obtained with different normal velocities cannot be directly compared, since higher normal velocities naturally lead to larger distances from the concentrator at the final timestep. In practical scenarios, the objective is to reach a specified distance between the concentrator and UAV, determined by a desired measurement accuracy in Stage 4, as derived in Appendix A.

A key disadvantage of flying at high normal velocities is the increased acceleration and deceleration required to reach and maintain such velocities over a finite distance: Strong changes in velocity induce transient behavior, causing potentially large displacements of the reflection signal p_{refl}^{CCSXY} with respect to the desired setpoint s_{refl}^{CCSXY} . Consequently, there exist practical limits on how rapidly the UAV can be accelerated and decelerated without risking the loss of the reflection signal. Hence, the maximum achievable normal velocity $v_{control,Z}^{CCS}$ (and hence the minimum achievable time to terminate Stage 3) is constrained by both the desired final distance and the allowable transient displacement of the reflection signal.

Tuning Study 5 investigates this trade-off by fixing the target distance and analyzing how different normal velocity profile shaping

parameter settings affect the reflection signal magnitudes. Specifically, the maximum normal velocity $v_{z,max}^{CCS}$ (which determines the measurement speed) and the rise time t_{rise} (which determines the smoothness) are varied. This results in a set of 23 normal velocity profile shaping (NVPS) configurations, denoted as C1–C23, which are listed in Table 4.

The hold time t_{hold} is selected such that Eq. (41) yields the desired distance $d_{S3} = 100$ m. With an initial distance of $d_{init} = 15$ m, the UAV reaches a distance of $d_{S4} = 115$ m upon completion of Stage 3. This value is chosen based on two considerations: (i) As discussed in Appendix A, the measurement accuracy does not improve substantially beyond distances between 100–120 m, assuming a representative camera interior orientation and realistic noise levels in the individual system components. (ii) As explained in Section 3.4.1, the operating range of the system is limited to twice the focal length of the measured concentrator. For the Solar Tower Jülich, the smallest value of twice the focal length is 123.9 m. Consequently, all concentrators at the Solar Tower Jülich can be measured using this configuration.

Note that in practice, the final distance may vary since the UAV does not instantaneously reach the commanded velocity, causing the traveled distance to deviate from the anticipated value. The controller gains k_p and k_i are chosen as identified in the tuning studies.

4.5. Evaluation studies

The evaluation studies investigate how well the system configurations identified under idealized assumptions generalize to conditions specific to the Solar Tower Jülich. For increased simulation fidelity, the concentrator pose is computed using the kinematic model, and the reflection model uses a discretized slope deviation map instead of assuming a perfectly flat mirror surface.

Across all evaluation studies, the proportional gain k_p and integral gain k_i are fixed to the values determined in the tuning studies, namely $k_p = 0.4$ and $k_i = 0.04$ (see Section 5.1). As in Tuning Study 5, the normal velocity profile shaping parameters are swept jointly to achieve different trade-offs between reflection signal displacement and convergence time. To limit the total number of simulation runs, only the subset of the normal velocity profile shaping (NVPS) configurations C1, C6, C11, C16, and C21 is considered during evaluation. This subset is selected to represent characteristic operating regimes across the normal velocity profile shaping parameter space and is used consistently in all evaluation studies.

4.5.1. Evaluation Study 1

Evaluation Study 1 investigates the influence of the concentrator surface S on the system performance. With the exception of the purely flat reference case, all surface configurations are based on one exemplary concentrator for which a deflectometric measurement is available. The concentrator is selected based on both focal length and optical quality. Evaluating the slope deviation maps from 19 concentrators yields an average root mean squared (RMS) slope error of approximately 1.5 mrad, with a maximum RMS slope error of 1.8 mrad. A concentrator with a comparatively high slope error is therefore chosen to represent a near worst-case scenario. At the same time, its focal length $f = 164.3$ m lies close to both the mean focal length $f_{mean} = 165.4$ m and the median focal length $f_{median} = 159.8$ m across all concentrators at the Solar Tower Jülich. Based on this concentrator, four surface configurations are defined:

1. S^F (Flat): Reference case assuming a perfectly flat mirror surface, resulting in a single reflection signal.
2. S^{CF} (Canted Flat): Each of the four facets is ideally canted to match the focal length of the exemplary concentrator. All facets are assumed to be perfectly flat.

Table 4

Summary of all normal velocity profile shaping (NVPS) configurations introduced in Tuning Study 5, characterized by different combinations of the maximum normal velocity $v_{z,max}^{CCS}$, the rise time t_{rise} , and the hold time t_{hold} . The arrival time $t_{arrival} = 2 \cdot t_{rise} + t_{hold}$ specifies the time at which the UAV no longer moves away from the concentrator. Due to the discretization of the system time in steps of $t_{sys} = 0.5$ s, the commanded distance slightly deviates from 100 m as a result of rounding effects. A subset of these configurations (C1, C6, C11, C16, and C21) is subsequently used throughout the evaluation studies.

NVPS Configuration	$v_{z,max}^{CCS}$ [m/s]	t_{rise} [s]	t_{hold} [s]	$t_{arrival}$ [s]	d_{S3} [m]
C1	1.5	20	46.5	86.5	99.75
C2	1.5	15	51.5	81.5	99.75
C3	1.5	10	56.5	76.5	99.75
C4	1.5	5	61.5	71.5	99.75
C5	1.5	0	66.5	66.5	99.75
C6	2	20	30.0	70.0	100.0
C7	2	15	35.0	65.0	100.0
C8	2	10	40.0	60.0	100.0
C9	2	5	45.0	55.0	100.0
C10	2	0	50.0	50.0	100.0
C11	3	20	13.5	53.5	100.5
C12	3	15	18.5	48.5	100.5
C13	3	10	23.5	43.5	100.5
C14	3	5	28.5	38.5	100.5
C15	3	0	33.5	33.5	100.5
C16	5	20	0.0	40.0	100.0
C17	5	15	5.0	35.0	100.0
C18	5	10	10.0	30.0	100.0
C19	5	5	15.0	25.0	100.0
C20	5	0	20.0	20.0	100.0
C21	7	10	4.5	24.5	101.5
C22	7	5	9.5	19.5	101.5
C23	7	0	14.5	14.5	101.5

3. S^M (Measured): Each facet is ideally canted as for the surface S^{CF} . Measured slope deviation maps obtained using the QDEC-H system [24] are superimposed onto this ideal reference surface, i.e., slope errors are defined with respect to the ideally canted geometry.
4. S^{MA} (Measured Amplified): Identical to the surface S^M , but with the measured slope deviations amplified by a factor of two.

The RMS of the discretized slope error is 1.73 mrad for the surface S^M . For the surface S^{MA} , amplification by a factor of two results in an RMS slope error of 3.46 mrad. All remaining boundary conditions are defined as summarized in Table 5.

4.5.2. Evaluation Study 2

While the normal vector refinement can be applied using the (measured or design) kinematic parameters of the measured heliostat, operating the control scheme without the normal vector refinement simplifies the overall implementation and reduces its dependence on properties specific to a particular heliostat type or heliostat field. Evaluation Study 2 therefore examines whether the application of the normal vector refinement has a significant impact on the system behavior during Stage 3, or whether a simplified approach without normal vector refinement provides comparable performance. Both estimation schemes are assessed under identical boundary conditions, as summarized in Table 5.

4.5.3. Evaluation Study 3

Evaluation Study 3 explores the sensitivity of the system behavior in Stage 3 with respect to the in-plane rotation ψ of the concentrator. While the control algorithm is formulated independently of this angle, changes in the in-plane rotation ψ affect the spatial distribution of reflections on the concentrator surface and potentially influence the observed

reflection trajectories. To analyze this effect, a set of four in-plane rotation angles, $\psi = 0^\circ, 30^\circ, 60^\circ, 90^\circ$, is simulated. The remaining boundary conditions are chosen as described in Table 5.

4.5.4. Evaluation Study 4

All previous evaluation studies are conducted at the elevation angle identified as the most challenging case in Tuning Study 4. This worst-case scenario is determined under idealized simulation assumptions and without the application of normal velocity profile shaping, where it induces the strongest transient system response. Evaluation Study 4 examines whether the system behavior observed under these idealized conditions transfers to more realistic settings. Specifically, the study evaluates the system performance for a range of elevation angles while shaping the normal velocity profile and maintaining the higher-fidelity simulation model used in the preceding evaluation studies. The objective of this study is to verify that the control strategy remains robust across different elevation angles, using a model of higher fidelity.

5. Results and discussion

As described in the previous section, each study yields a collection of reflection trajectories, where every trajectory corresponds to a specific system configuration and set of boundary conditions. Presenting all reflection trajectories in the time domain would be impractical and hinder interpretability. Therefore, the results are visualized in the metric domain: For each trajectory, one or more metrics are computed (see Section 4.1.2). The resulting metric values are plotted as a function of the system parameters and boundary conditions varied in the respective study to enable a systematic comparison across all investigated parameter combinations. To still convey an intuition of the system behavior, a representative subset of reflection trajectories is shown for each study. The remaining reflection trajectories are provided in Appendix E.

5.1. Tuning studies

5.1.1. Tuning Study 1

To provide an intuitive impression of the system behavior, Fig. 13 shows the X-component of the reflection trajectories for the initial position $p_{UAV,init,1}^{CCS}$ and proportional gains $k_p = 0.2, 0.4, 0.6, 0.8$. While only the X-component is shown, the Y-component exhibits similar behavior. Lower gains result in slower convergence, whereas higher gains cause oscillatory behavior.

For a systematic evaluation across all parameter combinations, the integral of the time-weighted absolute error (ITAE) is computed for each trajectory. The resulting metric values are shown in Fig. 14 as a function of the proportional gain k_p for all four initial positions. Initial positions located farther from the concentrator center lead to larger ITAE values. However, the gain $k_p = 0.4$ consistently yields the lowest ITAE across all initial positions. Consequently, $k_p = 0.4$ is selected as the optimal proportional gain for the subsequent study.

5.1.2. Tuning Study 2

To illustrate the system behavior in the presence of a normal velocity component, Fig. 15 shows the Y-component of the reflection trajectories for a normal velocity of $v_{z,max}^{CCS} = 5$ m/s and the proportional gains $k_p = 0.2, 0.4, 0.6, 0.8$. The X-component is not presented, as it remains close to zero across all simulations. Consistent with Tuning Study 1, gains above $k_p = 0.4$ cause increasingly oscillatory behavior, whereas smaller gains result in slower convergence. A steady-state reflection error is observed, whose magnitude decreases with increasing proportional gains k_p . As discussed in Appendix B, this effect is likely caused by the coupling of horizontal and vertical motion.

For this tuning study, the integral of the time-weighted steady-state corrected absolute error (ITAESS) is evaluated for each trajectory to distinguish persistent oscillations from steady-state effects. The resulting

Table 5

Overview of the system configuration and boundary condition parameters considered during the evaluation studies, indicating for each study whether parameters are swept, fixed, or adopted from preceding tuning results.

Evaluation Study	Concentrator Surface	$v_{Z,max}^{CCS}$ [m/s]	t_{sweep} [s]	t_{hold} [s]	NVR	ψ [°]	el [°]
1	S^{CF}, S^F, S^M, S^{MA}	Fixed Distance: $d_{S3} \approx 100$ m NVPS Configurations C1, C6, C11, C16, C21 (see Table 4)			Applied	0	Worst of Tuning Study 4 Value: 60
2	S^M	Fixed Distance: $d_{S3} \approx 100$ m NVPS Configurations C1, C6, C11, C16, C21 (see Table 4)			Applied/ Not Applied	0	Worst of Tuning Study 4 Value: 60
3	S^M	Fixed Distance: $d_{S3} \approx 100$ m NVPS Configurations C1, C6, C11, C16, C21 (see Table 4)			Not Applied	0, 30, 60, 90	Worst of Tuning Study 4 Value: 60
4	S^M	Fixed Distance: $d_{S3} \approx 100$ m NVPS Configurations C1, C6, C11, C16, C21 (see Table 4)			Not Applied	0	0, 15, ..., 90




Fig. 13. X-component of the reflection trajectories for the initial position $p_{UAV,init,1}^{CCS}$ and the proportional gains $k_p = 0.2, 0.4, 0.6, 0.8$. The plot illustrates the influence of k_p on convergence speed and oscillatory behavior in Stage 2.




Fig. 14. Integral of the time-weighted absolute error (ITAE) as a function of the proportional gain k_p for all four initial positions considered in Tuning Study 1. Each point corresponds to one simulation run.

metric values are shown in Fig. 16 as a function of the proportional gain k_p for the normal velocities $v_{Z,max}^{CCS} = 1$ m/s, 3 m/s, 5 m/s.

Larger normal velocities lead to increased overall error levels. The proportional gain that minimizes the integral of time-weighted steady-state-corrected absolute error (ITAE_{SS}) depends on the normal velocity and lies approximately between $k_p = 0.5$ and $k_p = 0.7$. However, this dependence is also influenced by transient behavior, particularly the larger initial errors associated with smaller gains. Since the systematic optimization of transient response is addressed separately in Tuning Study 5, and to maintain consistency with Tuning Study 1, the proportional gain $k_p = 0.4$ is retained as a robust choice across all considered normal velocities.




Fig. 15. Y-component of the reflection trajectories for the normal velocity $v_{Z,max}^{CCS} = 5$ m/s and the proportional gains $k_p = 0.2, 0.4, 0.6, 0.8$. The figure illustrates the trade-off between convergence speed, oscillatory behavior, and steady-state error.

5.1.3. Tuning Study 3

To illustrate the influence of the integral gain, Fig. 17 shows the Y-component of the reflection trajectories for a normal velocity of $v_{Z,max}^{CCS} = 5$ m/s and integral gains $k_I = 0.0, 0.02, 0.03, 0.05$. Again, the X-component is not shown, as it remains close to zero across all simulations.

As demonstrated in the previous tuning study, a steady-state error persists for the integral gain $k_I = 0$. Introducing a non-zero integral gain eliminates this offset and leads to the convergence of the reflection error signal to zero. The convergence characteristics depend on the magnitude of the gain k_I : Smaller values yield slower but well-damped responses, whereas larger values increase the convergence rate at the expense of overshoot.

For this tuning study, the ITAE is evaluated for each trajectory. The resulting metric values are shown in Fig. 18 as a function of the integral gain k_I for the normal velocities $v_{Z,max}^{CCS} = 1$ m/s, 3 m/s, 5 m/s. Again, larger normal velocities lead to increased overall error levels. The ITAE is maximal for the gain $k_I = 0$ due to the steady-state offset. As the integral gain increases, the ITAE initially decreases due to improved steady-state performance, but rises again beyond approximately $k_I = 0.05$ as overshoot becomes dominant.

Although the integral gain that minimizes the ITAE depends on the normal velocity, $k_I = 0.04$ consistently yields low metric values across all tested normal velocities. It is therefore selected as a robust gain for the subsequent study.

5.1.4. Tuning Study 4

To illustrate the influence of the elevation angle, Fig. 19 shows the Y-component of the reflection trajectories for a normal velocity




Fig. 16. Integral of the time-weighted steady-state corrected absolute error (ITAESS) as a function of the proportional gain k_p for the normal velocities $v_{z,\max}^{\text{CCS}} = 1 \text{ m/s}$, 3 m/s , 5 m/s considered in Tuning Study 2. Each point represents one simulation run.




Fig. 17. Y-component of the reflection trajectories for the normal velocity $v_{z,\max}^{\text{CCS}} = 5 \text{ m/s}$ and the integral gains $k_I = 0.0, 0.03, 0.06, 0.09$. The plot illustrates the elimination of the steady-state error through integral action and the influence of k_I on convergence speed and overshoot.




Fig. 18. Integral of the time-weighted absolute error (ITAE) as a function of the integral gain k_I for the normal velocities $v_{z,\max}^{\text{CCS}} = 1 \text{ m/s}$, 3 m/s , 5 m/s considered in Tuning Study 3. Each point corresponds to one simulation run.

of $v_{z,\max}^{\text{CCS}} = 5 \text{ m/s}$ and elevation angles $el = 0^\circ, 30^\circ, 45^\circ, 60^\circ$. All configurations show comparable convergence times, reaching steady-state within 10–20 s. However, the transient amplitudes vary with the elevation angle. The largest peak response occurs for the elevation $el = 60^\circ$, whereas the elevation $el = 0^\circ$ yields the smallest transient amplitude.

For a systematic evaluation, the peak absolute error (PAE) is computed for all combinations of elevation angles and normal velocities. The resulting metric values are shown in Fig. 20. The largest errors occur for intermediate elevation angles, approximately between 15° and 75° . In contrast, the extreme configurations $el = 0^\circ$ (concentrator pointing toward the horizon) and $el = 90^\circ$ (concentrator pointing toward the




Fig. 19. Y-component of the reflection trajectories for the normal velocity $v_{z,\max}^{\text{CCS}} = 5 \text{ m/s}$ and the elevation angles $el = 0^\circ, 30^\circ, 45^\circ, 60^\circ$. The plot illustrates the influence of the elevation angle on the transient amplitude while maintaining comparable convergence times.




Fig. 20. Peak absolute error (PAE) as a function of the elevation angle el for all considered normal velocities in Tuning Study 4. Each point corresponds to one simulation run.

sky) exhibit the smallest errors. Based on these observations, the subsequent analysis focuses on the most demanding configuration, namely the elevation angle $el = 60^\circ$ at a normal velocity of $v_{z,\max}^{\text{CCS}} = 5 \text{ m/s}$.

5.1.5. Tuning Study 5

To characterize the trade-off between measurement speed and the risk of reflection loss, each reflection trajectory is evaluated using two metrics: (i) the convergence time t_{s3} in Stage 3 (see Eq. (51)) and (ii) the peak absolute error (PAE) (see Eq. (48)). Large peak absolute errors correspond to trajectories where the reflection signal moves far away from the concentrator center at some point during Stage 3. Consequently, the PAE serves as an indicator of the likelihood of losing the reflection signal under real-world conditions. It should be noted that the PAE becomes meaningful for real-world scenarios only in combination with the concentrator dimensions: For larger concentrators, higher PAE values are less critical than for smaller concentrators, as the available margin before reflection loss increases with concentrator size. Furthermore, slope errors contribute to the loss of the reflection signal, which are not captured in the tuning studies but are added in the evaluation studies.

The upper plot of Fig. 21 shows the sign-inverted Y-component of the reflection trajectories over time for the three exemplary normal velocity profile shaping (NVPS) configurations C16, C18, and C20 listed in Table 4. These configurations correspond to a normal velocity of $v_{z,\max}^{\text{CCS}} = 5 \text{ m/s}$ with sweep times $t_{\text{rise}} = 20 \text{ s}$, 10 s , and 0 s , respectively. The sign is inverted to facilitate comparison with the metric-domain representation shown in Fig. 22. The plot shows that longer sweep times t_{rise} reduce the maximum reflection signal magnitude, but come at the cost of increased convergence time. To quantify this trade-off, each error trajectory is mapped to the metric tuple (t_{s3}, PAE) , representing the error trajectory in




Fig. 21. Sign-inverted Y-component of the reflection trajectories for the three exemplary normal velocity profile shaping (NVPS) configurations C16, C18, and C20, corresponding to the maximum normal velocities $v_{Z,max}^{CCS} = 5$ m/s with sweep times $t_{rise} = 20$ s, 10 s, and 0 s. The upper plot shows the trajectories over time, including their associated metric tuples (t_{S3} , PAE). The lower plot shows the same trajectories as a function of traveled distance.




Fig. 22. Peak absolute error (PAE) plotted over the convergence time t_{S3} for all normal velocity profile shaping (NVPS) configurations from Tuning Study 5. Each point corresponds to one simulated reflection trajectory. Pareto-optimal NVPS configurations are shown with full opacity, while dominated NVPS configurations are shown with reduced opacity. The three exemplary NVPS configurations from Fig. 21 are marked in matching colors.

a two-dimensional trade-off space between convergence time and peak absolute error, as also indicated in the upper plot of Fig. 21.

The lower plot of Fig. 21 depicts the same reflection trajectory Y-component as a function of the traveled distance, illustrating how the reflection displacements are distributed along the trajectory.

Fig. 22 visualizes the metric tuples for all simulated NVPS configurations of Tuning Study 5 in the (t_{S3} , PAE)-plane. C23 is omitted for readability, as it yields a large error of PAE = 2.31 m. The three metric tuples corresponding to the exemplary error trajectories shown in Fig. 21 are marked using the marker edge colors in the metric-domain plot in Fig. 22.

Each of the 23 simulated trajectories corresponds to a single point in the metric plane. The distribution of these points reveals a trade-off:




Fig. 23. Peak absolute error (PAE) plotted over the convergence time t_{S3} for the normal velocity profile shaping (NVPS) configurations C1, C6, C11, C16, and C21 for the four concentrator surfaces S^F , S^{CF} , S^M , and S^{MA} .

Larger convergence times t_{S3} correlate with lower peak absolute error. Hence, slower convergence generally reduces the magnitude of reflection displacements. Points for which no other NVPS configuration yields simultaneous improvements in both metrics form the *Pareto front* [25]. These Pareto-optimal configurations represent the best achievable trade-offs between convergence time and reflection displacement. All dominated configurations are shown with reduced opacity in Fig. 22.

5.2. Evaluation studies

The results of Tuning Study 5 form the basis for the evaluation studies, as they enable the realization of different performance trade-offs by selecting appropriate Pareto-optimal normal velocity profile shaping (NVPS) configurations. To limit the overall number of simulations, only a subset of configurations is considered. Specifically, for each normal velocity $v_{Z,max}^{CCS}$, the NVPS configuration with the largest sweep time t_{rise} on the Pareto front is selected. This yields the five NVPS configurations C1, C6, C11, C16, and C21, consistently considered in all evaluation studies in order to illustrate the achievable trade-offs between the convergence time t_{S3} and the peak absolute error (PAE).

5.2.1. Evaluation Study 1

Fig. 23 shows the resulting metric tuples in the (t_{S3} , PAE)-plane for the four simulated surfaces S^F , S^{CF} , S^M , and S^{MA} , evaluated using the NVPS configurations C1, C6, C11, C16, and C21.

For the slower configurations C1, C6, and C11, all concentrator models lead to similar system behavior, with ideal canting resulting in slightly higher PAEs. For the fastest configuration C21, clear differences emerge: The idealized surfaces S^F and S^{CF} yield lower PAEs, whereas the surfaces incorporating realistic slope errors, S^M and S^{MA} , result in increased PAEs.

The corresponding reflection trajectories for configuration C21 are shown in Fig. 24, illustrating how the peak absolute error and the convergence times are determined from the time signals.

The differences observed between the concentrator surfaces can be explained by the surface-dependent relation between the UAV position and the resulting reflection signal: Since the UAV motion is controlled based on the reflection signal, variations in this relation directly influence the closed-loop system behavior. For the ideal flat surface S^F , the reflection signal is given by the projection of the UAV position onto the concentrator coordinate system XY-plane (CCSXY-plane). In contrast, multiple canted facets (modeled by S^{CF}) introduce discontinuities due to reflections appearing on or disappearing from individual facets. Additional slope errors (modeled by the surfaces S^M and S^{MA}) alter the sensitivity of the reflection signal with respect to the UAV position due to local surface deformations introducing nonlinear distortions.




Fig. 24. Reflection trajectories in Stage 3 obtained with normal velocity profile shaping configuration C21 for the four concentrator surfaces S^F , S^{CF} , S^M , and S^{MA} . The trajectories illustrate the resulting peak absolute errors (PAEs) and convergence times t_{S3} .




Fig. 25. Peak absolute error (PAE) plotted over the convergence time t_{S3} for the normal velocity profile shaping configurations C1, C6, C11, C16, and C21, with and without the application of the normal vector refinement (NVR).

For the slower NVPS configurations, the UAV position remains close to the concentrator center, where the mapping between position and reflection signal is approximately the identity: In this region, reflections are visible on all four facets and, due to the symmetric canting, the averaged reflection signal closely approximates the projection of the UAV position onto the CCSXY-plane. Consequently, the system behavior is nearly identical across all surfaces. For the fastest NVPS configuration, stronger transients move the UAV into regions where nonlinearities and discontinuities become significant. This results in larger reflection errors and hence increased PAE values for surfaces implementing slope errors.

A detailed geometric analysis of the surface-dependent position-to-reflection mapping is provided in [Appendix D](#).

5.2.2. Evaluation Study 2

[Fig. 25](#) compares the metric tuples in the (t_{S3}, PAE) -plane for system configurations with and without the application of the normal vector refinement (NVR). While both approaches show similar behavior, the application of this refinement consistently results in slightly lower peak absolute error, particularly for faster NVPS configurations.




Fig. 26. Upper plots: Reflection trajectories in Stage 3 obtained with the normal velocity profile shaping configuration C21, comparing cases with and without the application of the normal vector refinement (NVR). Lower plot: Corresponding geodesic orientation error $\delta(\hat{\mathbf{R}}_{\text{CCS}}^{\text{GCS}}, \mathbf{R}_{\text{CCS}}^{\text{GCS}})$ for both approaches.

The purpose of the normal vector refinement is to improve the estimation of the concentrator orientation during each iteration of the control loop. To quantify the improvement achieved by the normal vector refinement, the geodesic distance

$$\delta(\hat{\mathbf{R}}_{\text{CCS}}^{\text{GCS}}, \mathbf{R}_{\text{CCS}}^{\text{GCS}}) = \arccos\left(\frac{\text{Tr}\left(\left(\hat{\mathbf{R}}_{\text{CCS}}^{\text{GCS}}\right)^T \cdot \mathbf{R}_{\text{CCS}}^{\text{GCS}} - 1\right)}{2}\right) \quad (58)$$

is computed for both approaches, with and without normal vector refinement application for the exemplary NVPS configuration C21. The result is shown in the lower plot of [Fig. 26](#), while the upper plots show the corresponding Y-components of the reflection trajectories.

The application of the normal vector refinement consistently reduces the orientation estimation error. Nevertheless, estimation errors persist whenever reflections are not visible on all facets. This behavior is consistent with the description in [Section 3.4.2](#): When applying the refinement, the concentrator normal vector is estimated from the vector connecting the UAV position to the averaged reflection signal. If reflections from all facets are visible, the biases introduced by symmetric canting cancel out, resulting in an accurate estimate of the concentrator normal. However, if only a subset of reflections is available, the averaged reflection position becomes biased toward the corresponding facet orientations, leading to orientation errors on the order of the canting angle (approximately 3 mrad).

Without normal vector refinement, the concentrator normal is instead estimated based on the heliostat origin. As a result, the estimation error does not depend on the reflection signal but rather on the distance from the projection of the UAV position onto the CCSXY-plane to the heliostat position, as well as on the distance between the UAV and the concentrator. This approach yields smoother error curves but larger overall estimation errors.

Although the normal vector refinement improves the concentrator orientation estimation and leads to reduced peak absolute error, the




Fig. 27. Peak absolute error (PAE) plotted over the convergence time t_{s3} for the normal velocity profile shaping configurations C1, C6, C11, C16, and C21 under four simulated in-plane rotations, $\psi = 0^\circ, 30^\circ, 60^\circ, 90^\circ$.

improvements remain moderate. The orientation error without normal vector refinement remains below 11 mrad. Evaluation Study 2 therefore indicates that Stage 3 is robust with respect to inaccuracies in the concentrator normal estimation. Since the application of the normal vector refinement requires *a priori* knowledge of the concentrator geometry and increases computational complexity, it is not considered necessary for Stage 3 under the tested operating conditions. Given that the initial UAV position is sufficiently close to the optical axis and at an adequate distance from the concentrator, the simpler estimation scheme is sufficient.

Beyond the comparison of estimation schemes, Evaluation Study 2 demonstrates that the system operates reliably despite the approximation of the concentrator normal vector. In particular, it tolerates discrepancies between the true reflection location and the assumed reflection location (i.e., the heliostat origin) of up to approximately one meter. This finding implies that other sources of uncertainty are comparatively less critical during Stage 3. For instance, RTK positioning errors typically remain below 5 cm, and the reflection detection accuracy of the perception module is expected to be of a similar magnitude. Even if such errors propagate into the normal vector computation, they remain significantly smaller than the error introduced by the assumption that reflections always originate from the heliostat position.

5.2.3. Evaluation Study 3

Fig. 27 shows the metric tuples for the four simulated in-plane rotations, $\psi = 0^\circ, 30^\circ, 60^\circ, 90^\circ$, evaluated using the NVPS configurations C1, C6, C11, C16, and C21. The overall system behavior remains mostly unchanged across all rotation angles.

The corresponding reflection trajectories for the exemplary NVPS configuration C21 are depicted in Fig. 28, showing both the X- and Y-components of the reflection signal. Depending on the in-plane rotation angle, the reflection displacement caused by the UAV motion is distributed differently between the CCS X-axis and CCS Y-axis. For example, for the rotation $\psi = 0^\circ$, the displacement occurs primarily in the Y-component, while the X-component remains close to zero. For the rotation $\psi = 90^\circ$, this behavior is reversed, and the entire displacement is observed in the X-component. Intermediate rotation angles lead to a corresponding redistribution of the displacement between both axes.

Since the PAE measures the peak magnitude of the reflection trajectory, it is largely insensitive to how this displacement is distributed between the X- and Y-components. Consequently, the metric values are approximately invariant with respect to the rotation angle ψ . Remaining variations can be attributed to variations in the concentrator slope errors across the surface, as different in-plane rotations cause different




Fig. 28. Reflection trajectories in Stage 3 (X- and Y-components) obtained with the normal velocity profile shaping configuration C21 for the four simulated in-plane rotation angles, $\psi = 0^\circ, 30^\circ, 60^\circ, 90^\circ$.




Fig. 29. Peak absolute error (PAE) plotted over the convergence time t_{s3} for the normal velocity profile shaping configurations C1, C6, C11, C16, and C21 under seven simulated elevation angles, $el = 0^\circ, 15^\circ, \dots, 90^\circ$.

regions of the concentrator to reflect the incoming light back to the camera sensor. Overall, the results demonstrate that the system performance in Stage 3 is robust with respect to in-plane rotations of the concentrator.

5.2.4. Evaluation Study 4

Fig. 29 shows the metric tuples for the seven simulated elevation angles $el = 0^\circ, 15^\circ, \dots, 90^\circ$, evaluated using the NVPS configurations C1, C6, C11, C16, and C21.

A symmetry with respect to the elevation $el = 45^\circ$ can be observed. The metric curves for $el = 0^\circ$ closely match those for $el = 90^\circ$, the curves $el = 15^\circ$ closely match those for $el = 75^\circ$ and the curves $el = 30^\circ$ closely match those for $el = 60^\circ$. Both the largest peak absolute error values and the longest convergence times t_{s3} occur at $el = 45^\circ$.

The corresponding reflection trajectories for the exemplary NVPS configuration C21 are shown in Fig. 30. The plots also explain why the elevation angles $el = 0^\circ$ and $el = 90^\circ$ yield shorter convergence times:




Fig. 30. Reflection trajectories in Stage 3 obtained with the normal velocity profile shaping configuration C21 for the seven simulated elevation angles, $el = 0^\circ, 15^\circ, \dots, 90^\circ$. The subplots are arranged to emphasize the symmetry around the elevation $el = 45^\circ$.

For these elevations, the reflection displacement remains small throughout Stage 3, such that the reflection magnitude already lies within the settling corridor defined by $\varepsilon = 0.1$ m (see Eq. (51)) immediately after the arrival time $t_{\text{arrival}}^{(C21)} = 24.5$ s. In contrast, larger reflection displacements occur for elevation angles closer to 45° , requiring additional time for the controller to realign the reflection signal within the required tolerance band. The arrangement of the subplots emphasizes the symmetry of this behavior, illustrating that elevation-angle pairs with the same distance from 45° produce nearly identical reflection trajectories.

This symmetry indicates that the dominant factor influencing the peak absolute error remains the reflection signal velocity drift, which is maximized at $el = 45^\circ$. As shown empirically in Appendix B, this drift is symmetric about $el = 45^\circ$ and proportional to the normal velocity component. While the integral action mitigates this effect, it does not eliminate it entirely.

For qualitative illustration, a 3D visualization of an exemplary simulated flight is additionally provided as supplementary video material (see Appendix F). The depicted scenario shows the execution of Stage 3 for the elevation angle $el = 45^\circ$ using NVPS configuration C16, recorded as part of Evaluation Study 4, and preceded by Stage 2 to initially align the reflection signal with the concentrator center. The video illustrates the system behavior in the GCS, ICS, and CCS, the relation between the reflection signals and the resulting velocity control commands, and the intended reduction of the normal vector measurement uncertainty at increased distances, based on the uncertainty model derived in Appendix A for a DJI Zenmuse P1 camera.

5.2.5. Evaluation summary and operational implications

Across all evaluation studies, normal velocity profile shaping (NVPS) configuration C21 consistently achieves the shortest completion times for Stage 3, remaining below 30 s in all considered scenarios. However, this configuration also results in the largest peak absolute errors (PAEs). For the measured surface S^M , the maximum observed PAE is approximately 1.2 m, which is still smaller than the shorter side length of a concentrator at the Solar Tower Jülich. Hence, even under

the worst-case tested conditions and without applying the normal vector refinement, the reflection signal is not lost in simulation.

If a larger safety margin is desired, a slower NVPS configuration can be selected. In particular, C16 consistently limits the peak absolute error to below 0.5 m across all tested system configurations and boundary conditions. This increased robustness comes at the expense of approximately 10 s of additional measurement time.

Based on these results, Stage 3 can be safely completed within 40 s using the NVPS configuration C16 for concentrators at the Solar Tower Jülich. Considering the convergence time of Stage 2 of 10–15 s, the combined execution time for Stages 2 and 3 remains well below one minute.

6. Conclusion and outlook

In this work, we present ARTSCORE-F, a novel method for airborne real-time solar concentrator orientation estimation for heliostat fine calibration, building on the *HelioPoint* measurement principle. Rather than relying on predefined flight paths, the method feeds the observed reflection signal back into a closed-loop controller, enabling the adaptive guidance of the reflected signal toward specified locations on the concentrator surface as well as the control of the distance between the UAV and the concentrator.

The method can be applied in tracking heliostat fields, enabling the targeted recalibration of individual heliostats (such as heliostat outliers) while maintaining uninterrupted field operation. However, operation in a tracking heliostat field necessitates additional safety measures, in particular with respect to concentrated solar flux, which may damage or destroy the UAV. Future work should therefore address appropriate safety measures, such as exclusion zones, trajectory constraints, or coordination with plant operation.

The proposed system was prototyped and systematically tuned and evaluated in simulation. A suitable system configuration was identified under idealized conditions and subsequently assessed under increasingly realistic scenarios representative of the heliostat field at the Solar Tower Jülich. These scenarios included measured slope deviation maps as well as varying in-plane rotation and elevation angles. Across all investigated evaluation scenarios, reflection signal visibility was maintained throughout the measurement process. The combined execution time for guiding the reflection signal toward the concentrator center (Stage 2) and increasing the UAV–concentrator distance to reduce normal vector measurement uncertainty for the subsequent surface scan (Stage 3) remained below one minute. These values exclude the time required for the initial reflection signal acquisition (Stage 1) and for the final concentrator surface scan (Stage 4).

Overall, presented results indicate that the proposed method enables the controlled acquisition of *HelioPoint* calibration points, whose accuracy has been reported to be suitable for fine calibration of heliostats at the Solar Tower Jülich [6]. By adapting the flight plan online based on the observed reflection signals, the proposed method reduces the need for predefined scan patterns, avoids redundant measurements, and increases robustness against concentrator orientation uncertainty and surface slope errors. In this sense, the approach represents a step toward autonomous condition monitoring systems in solar tower power plants, where adaptive flight route planning can improve operational efficiency.

Stage 4 was not investigated quantitatively in the present work and should therefore be analyzed in more detail in future studies. Further improvements to the control strategy such as feed-forward control are expected to enhance the system performance. In addition, the applicability of the proposed method should be examined for different concentrator geometries, including varying numbers of facets, alternative facet layouts, and different surface characteristics.

Future work should also focus on the assessment of the method in a real solar tower plant using a real UAV system. This requires the further development of reliable real-time signal processing, autonomous

flight safety mechanisms, and integration into existing field operation procedures. Although communication delays and processing latencies were explicitly considered in the simulation studies, real-world deployment will depend on sufficiently fast and reliable communication under field conditions, which may be affected by interference, bandwidth limitations, and infrastructure constraints. Overall, the realization of a fully operational autonomous UAV-based condition monitoring system is expected to involve numerous practical challenges beyond those addressed in simulation, including hardware integration, synchronization of distributed components, fail-safe behavior, and certification aspects.

CRediT authorship contribution statement

Alexander Schnerring: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Marc Röger:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization. **Adrian Winter:** Writing – review & editing, Visualization, Validation, Software, Methodology, Investigation. **Rafal Broda:** Writing – review & editing, Software, Methodology, Conceptualization. **Michael Nieslony:** Writing – review & editing, Supervision, Methodology, Conceptualization. **Niels Algnier:** Writing – review & editing, Resources, Methodology. **Benedikt Kölsch:** Writing – review & editing, Supervision. **Robert Pitz-Paal:** Supervision, Funding acquisition.

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

During the preparation of this work, the authors used ChatGPT (GPT 5, OpenAI) to improve the readability and language of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Declaration of competing interest

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

Acknowledgements

This research was funded in part by the [European Union](#) under Grant Agreement No. [101234781](#) (*Sun-DT* project) and in part by the Ministry of Economic Affairs, Industry, Climate Action and Energy of the State of North Rhine-Westphalia under the Funding ID 005-2108-0068 (*Shine* project). The authors gratefully acknowledge this financial support. We also thank CSP Services for the prior collaborative work in the *HelioPoint-II* project (funded by the [German Federal Ministry for Economic Affairs and Climate Action](#), Contract No. [3028684](#)) on which this work builds. Furthermore, we thank our colleagues Oliver Kaufhold, Moritz Wirger, Marcel Sibum, Max Pargmann, Felix Göhring, and Peter Schwarzbözl at the Solar Tower Jülich.

Appendix A. Measurement uncertainty derivation

This appendix derives the measurement uncertainty of the local concentrator surface normal vector associated with the observed reflection point in a *HelioPoint* measurement setup. The derivation expresses this uncertainty as a function of uncertainties in the UAV position and reflection position, as well as the UAV–concentrator distance and the camera parameters. For simplicity, the camera and LED positions are assumed to coincide and are jointly referred to as the UAV position, as discussed in [Section 2.5.3](#). This assumption is introduced to simplify the uncertainty derivation by treating both positions as a single point. The actual normal vector measurement can explicitly account for their physical separation. Since this separation is small relative to the UAV–concentrator

distance, the approximation is expected to have a negligible influence on the resulting uncertainty model.

Let the measured UAV position be

$$\hat{p}_{\text{UAV}}^{\text{GCS}} = p_{\text{UAV}}^{\text{GCS}} + e_{\text{UAV}}^{\text{GCS}}, \quad (\text{A.1})$$

where $p_{\text{UAV}}^{\text{GCS}}$ denotes the true UAV position and $e_{\text{UAV}}^{\text{GCS}}$ denotes the error in the UAV position. Analogously, let the measured reflection position be

$$\hat{p}_{\text{refl}}^{\text{GCS}} = p_{\text{refl}}^{\text{GCS}} + e_{\text{refl}}^{\text{GCS}}, \quad (\text{A.2})$$

where $p_{\text{refl}}^{\text{GCS}}$ denotes the true reflection position and $e_{\text{refl}}^{\text{GCS}}$ denotes the error in the reflection position. The true connecting vector is then given as

$$b^{\text{GCS}} = p_{\text{UAV}}^{\text{GCS}} - p_{\text{refl}}^{\text{GCS}}, \quad (\text{A.3})$$

while the measured connecting vector is given as

$$\hat{b}^{\text{GCS}} = \hat{p}_{\text{UAV}}^{\text{GCS}} - \hat{p}_{\text{refl}}^{\text{GCS}} = b^{\text{GCS}} + \underbrace{e_{\text{UAV}}^{\text{GCS}} - e_{\text{refl}}^{\text{GCS}}}_{e_b}. \quad (\text{A.4})$$

For the angular error of the estimated vector, only the projected error

$$e_{\perp} = P \cdot e_b \quad (\text{A.5})$$

is relevant, where the projection matrix

$$P = I_3 - nn^T \quad (\text{A.6})$$

projects vectors expressed with respect to the GCS onto the plane defined by the normal $n = b^{\text{GCS}} / \|b^{\text{GCS}}\|$. For large distances d , the angular error is given as

$$e_{\text{ang}} \approx \frac{\|e_{\perp}\|}{\|b^{\text{GCS}}\|} = \frac{\|e_{\perp}\|}{d}. \quad (\text{A.7})$$

Let $\Sigma_{\perp}$ denote the covariance of the projected error $e_{\perp}$. Then the RMS projection error is given as [\[26\]](#)

$$\text{RMSE}_{\perp} = \sqrt{\mathbb{E}[e_{\perp}^T e_{\perp}]} = \sqrt{\text{Tr}(\Sigma_{\perp})}, \quad (\text{A.8})$$

where $\text{Tr}(\cdot)$ denotes the trace. The RMS projection error translates to the RMS angular error as

$$\text{RMSE}_{\text{ang}} \approx \frac{\text{RMSE}_{\perp}}{d} = \frac{\sqrt{\text{Tr}(\Sigma_{\perp})}}{d}. \quad (\text{A.9})$$

In what follows, the term $\text{Tr}(\Sigma_{\perp})$ is identified. In total, four factors are assumed to contribute to $e_{\perp}$. One factor originates from the UAV position uncertainty, while the remaining three are associated with determining the reflection position in the GCS.

1. *UAV Position Uncertainty:* The UAV position error $e_{\text{UAV}}^{\text{GCS}}$ is assumed to be zero-mean with an isotropic covariance

$$\Sigma_{\text{UAV}} = \sigma_{\text{UAV}}^2 I_3. \quad (\text{A.10})$$

As described in [Section 2.5.1](#), the reflection position is obtained by mapping the image coordinates $p_{\text{refl}}^{\text{ICS}}$ to the CCS and then transforming them to the GCS using the estimated concentrator pose $\hat{T}_{\text{CCS}}^{\text{GCS}}$. Three distinct uncertainty sources contribute to the reflection position error $e_{\text{refl}}^{\text{GCS}}$.

2. **Perception Module:** The image coordinates $p_{\text{refl}}^{\text{ICS}}$ can only be determined with some uncertainty. When this estimate is mapped into 3D space, the resulting uncertainty lies entirely within the plane orthogonal to b^{GCS} , since the corresponding 2D image coordinates are mapped onto this plane. Therefore, this uncertainty is described in two-dimensional space assuming zero-mean and isotropic covariance

$$\Sigma_{\text{perc}}^{\text{ICS}} = (\sigma_{\text{perc}}^{\text{ICS}})^2 \mathbf{I}_2. \quad (\text{A.11})$$

Table A.1

Uncertainty values used in the angular uncertainty model Eq. (A.19).

Parameter	Value	Reference/Description
σ_{UAV}	7 mm	[6]
σ_{origin}	7 mm	[10]
σ_{kp}	5 mm	Assumption
$\sigma_{\text{perc}}^{\text{ICS}}$	1 px	Conservative Estimate

Table A.2

Camera parameters used in the angular uncertainty model Eq. (A.19).

Parameter	DJI Mavic 3 Enterprise (Video)	DJI Zenmuse P1 (Photo)
s_{mm}	17.3 mm	35.9 mm
f_{mm}	11.7 mm	35 mm
s_{px}	3840 px	8192 px




Fig. A.1. Normal vector measurement uncertainty as a function of the UAV-concentrator distance d , computed using Eq. (A.19) and inserting the parameters listed in Tables A.1 and A.2 for two different cameras. The dashed lines indicate the asymptotic accuracy limits of the uncertainty model.

Note that $\sigma_{\text{perc}}^{\text{ICS}}$ is expressed in pixel units, while the corresponding value in the CCS, obtained via homography, is expressed in meters. These two units are related by the ground sampling distance, which depends on the distance between the camera and the concentrator d , the camera focal length f_{mm} and the camera sensor size in metric units s_{mm} and pixel units s_{px} :

$$\text{GSD} = \frac{s_{\text{mm}}}{f_{\text{mm}} \cdot s_{\text{px}}} \cdot d. \quad (\text{A.12})$$

Overall, the perception covariance is

$$\Sigma_{\text{perc}} = \text{GSD}^2 \cdot (\sigma_{\text{perc}}^{\text{ICS}})^2 \mathbf{I}_2. \quad (\text{A.13})$$

3. **Heliostat origin uncertainty:** Errors in the heliostat origin o^{GCS} are assumed to be zero-mean with an isotropic covariance

$$\Sigma_{\text{origin}} = \sigma_{\text{origin}}^2 \mathbf{I}_3. \quad (\text{A.14})$$

4. **Kinematic parametrization uncertainty:** Using a noisy kinematic parametrization $\hat{k}$ or (in case no estimate is available) the design kinematic parametrization k in the kinematic backward pass Eq. (9) introduces additional translational uncertainty due to lever-arm effects. The associated error $e_{\text{kp}}^{\text{GCS}}$ is assumed to be zero-mean with the isotropic covariance

$$\Sigma_{\text{kp}} = \sigma_{\text{kp}}^2 \mathbf{I}_3. \quad (\text{A.15})$$

Assuming independence, each of these four covariance terms contributes additively. Since the covariance matrices Σ_{UAV} , Σ_{origin} and Σ_{kp} characterize noise in three-dimensional space, their contribution to $\text{Tr}(\Sigma_{\perp})$ is determined by applying the projection matrix P . For an isotropic covariance $\Sigma = \sigma^2 \cdot \mathbf{I}_3$, the covariance of the projection is given by [27]

$$\Sigma_{\perp} = P \cdot \Sigma \cdot P^T = \sigma^2 P, \quad (\text{A.16})$$

with trace

$$\text{Tr}(\Sigma_{\perp}) = 2\sigma^2. \quad (\text{A.17})$$

On the other hand, Σ_{perc} is already orthogonal to b^{GCS} and does not need to be projected. Hence, the overall uncertainty orthogonal to b^{GCS} is

$$\begin{aligned} \text{Tr}(\Sigma_{\perp}) &= \text{Tr}(P \cdot (\Sigma_{\text{UAV}} + \Sigma_{\text{origin}} + \Sigma_{\text{kp}}) \cdot P^T) + \text{Tr}(\Sigma_{\text{perc}}) \\ &= 2 \cdot \left(\sigma_{\text{UAV}}^2 + \sigma_{\text{origin}}^2 + \sigma_{\text{kp}}^2 + \left(\frac{\sigma_{\text{perc}}^{\text{ICS}} \cdot s_{\text{mm}}}{f_{\text{mm}} \cdot s_{\text{px}}} \right)^2 \cdot d^2 \right). \end{aligned} \quad (\text{A.18})$$

Inserting into Eq. (A.9) yields

$$\text{RMSE}_{\text{ang}} = \frac{\sqrt{2 \cdot \left(\sigma_{\text{UAV}}^2 + \sigma_{\text{origin}}^2 + \sigma_{\text{kp}}^2 + \left(\frac{\sigma_{\text{perc}}^{\text{ICS}} \cdot s_{\text{mm}}}{f_{\text{mm}} \cdot s_{\text{px}}} \right)^2 \cdot d^2 \right)}}{d}. \quad (\text{A.19})$$

Assuming the components of the projected error $e_{\perp}$ to be independent, zero-mean, normally distributed random variables with identical variance, its magnitude follows a Rayleigh distribution with cumulative distribution function

$$P(e_{\text{ang}} < r) = 1 - \exp\left(-\frac{r^2}{2\sigma^2}\right), \quad (\text{A.20})$$

where the scaling parameter is given by $\sigma = \text{RMSE}_{\text{ang}}/\sqrt{2}$. Thus, the probability that the angular error remains below a radius r is

$$P(e_{\text{ang}} < r) = 1 - \exp\left(-\frac{r^2}{\text{RMSE}_{\text{ang}}^2}\right). \quad (\text{A.21})$$

Conversely, for a desired confidence level p , the corresponding angular confidence radius is

$$r_p = \text{RMSE}_{\text{ang}} \cdot \sqrt{-\ln(1-p)}. \quad (\text{A.22})$$

This yields, for example,

$$\begin{aligned} r_{0.95} &\approx 1.73 \cdot \text{RMSE}_{\text{ang}}, \\ r_{0.99} &\approx 2.15 \cdot \text{RMSE}_{\text{ang}}. \end{aligned} \quad (\text{A.23})$$

Figure Fig. A.1 shows the resulting RMSE_{ang} as a function of the UAV-concentrator distance d , obtained by inserting the parameter values listed in Tables A.1 and A.2 into Eq. (A.19).

The reflection detection uncertainty $\sigma_{\text{perc}}^{\text{ICS}}$ is set to 1 px, which denotes a standard deviation and should therefore be interpreted as a statistical

measure rather than a strict bound. This value represents a conservative estimate informed by experimental evaluations of the current AI-based perception module on image data recorded at the Solar Tower Jülich. Under controlled conditions, smaller detection uncertainties were observed. The chosen value includes a safety margin to account for varying observation distances, changing environmental conditions, and possible generalization effects beyond the evaluated dataset.

While Eq. (A.19) predicts a finite asymptotic limit for $d \rightarrow \infty$, this limit is of a purely theoretical nature: Increasing the UAV–concentrator distance reduces the ground sampling distance such that the concentrator eventually occupies only a few pixels, preventing reliable concentrator corner point and reflection detection. Moreover, the reflected intensity decreases with increasing distance, further complicating the reflection detection. Similar challenges may arise for concentrator surfaces with stronger curvature or surface imperfections, where reflections appear more spatially extended and less well-defined (as discussed in Section 2.5.4). In addition, Appendix D establishes the operating range $0 \text{ m} \leq d \leq 2f$, which is required to ensure a monotonic variation of the reflection signal with the UAV position under the assumption of ideally canted, flat facets. It should therefore be emphasized that Eq. (A.19) represents a simplified geometric uncertainty model and provides only a first-order performance estimate. The expression captures the dominant error sources but does not account for additional practical effects. Furthermore, the uncertainty parameters themselves may vary between different heliostat fields: The achievable UAV positioning accuracy depends on factors such as the proximity of an RTK base station or the performance of alternative positioning methods. Consequently, the results shown in Fig. A.1 should be interpreted as a performance estimate rather than a strict guarantee.

Appendix B. Reflection signal drift and steady state

This appendix provides a physically motivated hypothesis for the origin of the reflection signal drift observed in simulation. The objective is not to claim a definitive causal explanation, but to demonstrate that a simple geometric thrust-coupling model predicts the reflection signal drift's dependence on both elevation angle and normal velocity. However, tests with a real UAV platform are required to validate the hypothesis.

B.1. Reflection drift

When the commanded velocity contains a vertical component (i.e., normal to the ground plane), the total thrust must change to balance gravity and achieve the desired climb rate. If the vehicle is simultaneously commanded a horizontal motion (i.e., parallel to the ground plane), the thrust axis is slightly tilted with respect to the global vertical. The vehicle therefore accelerates along a slightly tilted thrust axis, inducing a horizontal velocity component along the tilt direction in addition to the commanded horizontal motion.

To quantify this effect, let el denote the concentrator elevation, defined as the angle between the ground plane and the desired flight direction. An incremental change in total thrust ΔT under a tilted thrust axis produces a horizontal force component proportional to $\Delta T \cdot \cos(el)$, driving the unintended horizontal motion of the UAV. However, the controller observes only the projection of this motion onto the concentrator coordinate system XY-plane (CCSXY-plane), which scales with $\sin(el)$. Consequently, the CCSXY-plane reflection drift velocity scales as

$$\|v_{\text{drift}}^{\text{CCSXY}}\| \propto v_{\text{control,Z}}^{\text{CCS}} \cdot \sin(el) \cdot \cos(el) = \frac{1}{2} \cdot v_{\text{control,Z}}^{\text{CCS}} \cdot \sin(2el). \quad (\text{B.1})$$

This simple geometric argument therefore predicts (i) a linear dependence on the commanded normal velocity $v_{\text{control,Z}}^{\text{CCS}}$ and (ii) a $\sin(2el)$ elevation dependence with a maximum at the elevation angle $el = 45^\circ$.




Fig. B.1. Steady-state reflection signal in the concentrator coordinate system XY-plane (CCSXY-plane) for different elevation angles and commanded normal velocities. The curves closely match the predicted $v_{\text{control,Z}}^{\text{CCS}} \cdot \sin(2el)$ dependency.

B.2. Steady-state reflection signal under proportional control

Consider an exclusively proportional controller that commands the CCSXY-plane velocity according to

$$v_{\text{control}}^{\text{CCSXY}} = k_p \cdot (s_{\text{refl}}^{\text{CCSXY}} - p_{\text{refl}}^{\text{CCSXY}}), \quad (\text{B.2})$$

where $s_{\text{refl}}^{\text{CCSXY}}$ denotes the reflection setpoint and $p_{\text{refl}}^{\text{CCSXY}}$ denotes the measured reflection position. With the reflection setpoint set to zero, the control law simplifies to

$$v_{\text{control}}^{\text{CCSXY}} = -k_p \cdot p_{\text{refl}}^{\text{CCSXY}}. \quad (\text{B.3})$$

In an equilibrium, the reflection signal remains constant since the commanded velocity compensates the reflection drift velocity, i.e.,

$$v_{\text{control}}^{\text{CCSXY}} = -k_p \cdot p_{\text{refl,SS}}^{\text{CCSXY}} = -v_{\text{drift}}^{\text{CCSXY}}. \quad (\text{B.4})$$

Substituting the predicted disturbance yields

$$\|p_{\text{refl,SS}}^{\text{CCSXY}}\| \propto \frac{c}{k_p} \cdot v_{\text{control,Z}}^{\text{CCS}} \cdot \sin(2el), \quad (\text{B.5})$$

where the constant scalar c accounts for the UAV's dimensions, mass, and effective velocity and acceleration bandwidths. If the UAV moves away from the concentrator, the CCSXY-plane reflection drift is directed toward decreasing elevation angles, as is the steady-state reflection signal induced by the proportional controller.

Another study is conducted to confirm the dependency predicted by this explanation model, where the elevation angle el is swept for the different normal velocities $v_{\text{control,Z}}^{\text{CCS}} = 1 \text{ m/s}, 3 \text{ m/s}, 5 \text{ m/s}$. The proportional gain is fixed to $k_p = 0.4$, as determined in the tuning studies. Fig. B.1 shows the steady-state reflection signal (computed as described in Eq. (47)) as a function of el and $v_{\text{control,Z}}^{\text{CCS}}$. The results show the predicted linear scaling with $v_{\text{control,Z}}^{\text{CCS}}$ and the characteristic $\sin(2el)$ elevation dependence. Note that for the depicted study, the CCS X-axis is aligned parallel to the ground. The drift direction therefore projects entirely onto the CCS Y-axis, which points above the horizon. Consequently, only the Y-component of the steady-state reflection signal is visible in the figure, and its value appears with a negative sign. While the proportional gain is fixed for this study, Fig. 15 demonstrates that varying k_p scales the steady-state reflection signal proportionally to $1/k_p$, in agreement with the relation above.

Appendix C. Further details on the transformation from CCS to GCS

This appendix provides further details on the estimation of the concentrator orientation estimate $\hat{R}_{\text{CCS}}^{\text{GCS}}$ used to transform the controller velocity commands from CCS to GCS.

C.1. Construction of rotation matrix from a normal vector

Two degrees of freedom of the rotation $\hat{\mathbf{R}}_{\text{CCS}}^{\text{GCS}}$ can be derived using the normal vector estimate $\hat{\mathbf{n}}_{\text{CCS}}^{\text{GCS}}$, as computed by Eq. (33). For this, let the vectors $\mathbf{x}'$, $\mathbf{y}'$ and $\mathbf{z}'$ be the basis vectors of the estimated rotation, i.e.,

$$\hat{\mathbf{R}}_{\text{normal}}^{\text{GCS}} = (\mathbf{x}' \ \mathbf{y}' \ \mathbf{z}') \quad (\text{C.1})$$

The rotated Z-axis is chosen as

$$\mathbf{z}' = \hat{\mathbf{n}}_{\text{CCS}}^{\text{GCS}}. \quad (\text{C.2})$$

The rotated X-axis $\mathbf{x}'$ is chosen parallel to the ground plane. Since $\mathbf{x}'$, $\mathbf{y}'$ and $\mathbf{z}'$ are orthonormal, it can be concluded that

$$\mathbf{x}' = \frac{1}{\sqrt{(\hat{\mathbf{n}}_{\text{CCS}}^{\text{GCS}})^2_{\text{X}} + (\hat{\mathbf{n}}_{\text{CCS}}^{\text{GCS}})^2_{\text{Y}}}} \begin{pmatrix} (\hat{\mathbf{n}}_{\text{CCS}}^{\text{GCS}})_{\text{Y}} \\ -(\hat{\mathbf{n}}_{\text{CCS}}^{\text{GCS}})_{\text{X}} \\ 0 \end{pmatrix}. \quad (\text{C.3})$$

The remaining basis vector $\mathbf{y}'$ is given by

$$\mathbf{y}' = \hat{\mathbf{n}}_{\text{CCS}}^{\text{GCS}} \times \mathbf{x}'. \quad (\text{C.4})$$

C.2. Estimation of in-plane rotation ψ

The remaining in-plane rotation angle ψ can be estimated from the detected concentrator corner points

$$\mathbf{P}_{\text{corner}}^{\text{ICS}} = (\mathbf{p}_{\text{corner,UL}}^{\text{ICS}}, \mathbf{p}_{\text{corner,UR}}^{\text{ICS}}, \mathbf{p}_{\text{corner,LR}}^{\text{ICS}}, \mathbf{p}_{\text{corner,LL}}^{\text{ICS}}). \quad (\text{C.5})$$

A basic estimate can be obtained from the angle between the bottom concentrator edge

$$\mathbf{l}_{\text{bottom}} = \mathbf{p}_{\text{corner,LR}}^{\text{ICS}} - \mathbf{p}_{\text{corner,LL}}^{\text{ICS}} \quad (\text{C.6})$$

with the ICS X-axis:

$$\hat{\psi}_{\text{bottom}} = \arctan2((\mathbf{l}_{\text{bottom}})_{\text{Y}}, (\mathbf{l}_{\text{bottom}})_{\text{X}}). \quad (\text{C.7})$$

To improve robustness against noise in the detected corner positions, the same procedure can be applied to all four concentrator edges. Note that, depending on the considered edge, permutations of the arguments of $\arctan2(\cdot, \cdot)$ and sign adjustments are required to obtain the angle $\hat{\psi}_{(\cdot)}$ with respect to the correct ICS axis: Horizontal edges must be measured with respect to the ICS X-axis, whereas vertical edges must be measured with respect to the ICS Y-axis.

The final estimate is computed as the average of the individual edge-based estimates:

$$\hat{\psi} = \frac{\hat{\psi}_{\text{bottom}} + \hat{\psi}_{\text{right}} + \hat{\psi}_{\text{top}} + \hat{\psi}_{\text{left}}}{4}. \quad (\text{C.8})$$

This estimation assumes the camera orientation about its optical axis is known. In typical operation, this camera roll angle is approximately zero due to gimbal stabilization aligning the camera sensor's lower edge with the ground plane. If the camera roll angle is non-zero but known, it can be compensated for by subtracting it from the estimated $\hat{\psi}$.

Appendix D. Surface-dependent mapping between UAV position and reflection signal

This appendix describes the relation between the UAV position and the reflection signal as a mapping from 3D space to 2D space in order to (i) provide intuition for the theoretically achievable operating range and (ii) explain the emergence of surface-dependent closed-loop behavior.

D.1. Mapping formulation

A concentrator surface $\mathcal{S}$ defines a mapping from the UAV position expressed with respect to the CCS to the observed reflection signal:

$$\mathcal{S}(\mathbf{p}_{\text{UAV}}^{\text{CCS}}) = \mathbf{p}_{\text{refl}}^{\text{CCSXY}}. \quad (\text{D.1})$$

Strictly speaking, the reflection signal is measured in the ICS and mapped to the concentrator coordinate system XY-plane (CCSXY-plane) via a homography (see Section 3.2.3), which is omitted here for simplicity.

For multi-faceted concentrators, multiple reflections may be visible simultaneously. The signal used for control is obtained by averaging over all visible facets $\mathcal{M}_{\text{vis}}$ i.e.,

$$\mathbf{p}_{\text{refl}}^{\text{ICS}} = \sum_{m \in \mathcal{M}_{\text{vis}}} \mathbf{p}_{\text{refl},m}^{\text{ICS}}. \quad (\text{D.2})$$

The resulting mapping exhibits discontinuities due to facet visibility changes and is generally nonlinear due to surface slope errors. Since the controller operates on $\mathbf{p}_{\text{refl}}^{\text{CCSXY}}$, these properties directly influence the closed-loop system behavior.

To visualize the mapping, UAV positions are sampled in simulation and the corresponding reflection signals are computed. To reduce dimensionality, the X-component is fixed at $(\mathbf{p}_{\text{UAV}}^{\text{CCS}})_{\text{X}} = 0$ m, and the dependency between the Y-components is analyzed for different distances $d = (\mathbf{p}_{\text{UAV}}^{\text{CCS}})_{\text{Z}}$.

D.2. Operating range of $\mathcal{S}^{\text{CF}}$

The mapping between Y-components for the surface $\mathcal{S}^{\text{CF}}$ with ideally canted, perfectly flat facets is shown in Fig. D.1 for $d = 0.5f$, f , $2f$, $2.5f$. Each curve corresponds to a trajectory along one of the vertical lines indicated in Fig. 9, as indicated by the colored boxes in matching colors. The identity mapping (1-curve) corresponding to the perfectly flat surface $\mathcal{S}^{\text{F}}$ is shown for reference.

Within the range $0 \leq d \leq 2f$, the mapping remains monotonic, i.e., the reflection signal increases with the UAV position. Beyond this range (i.e., for $d = 2.5f$), the mapping becomes non-injective: UAV positions in one half of the concentrator may map to the reflection signals in the opposite half. This illustrates the fundamental operating limit.

D.3. Behavior in the evaluated range

The evaluation studies are conducted for an exemplary concentrator with focal length $f = 164.3$ m, while the starting point and endpoint of each trajectory lie at the distances $d = 15$ m and $d = 115$ m, respectively. This corresponds to a normalized range of approximately $0.1f$ to $0.7f$ for the evaluated concentrator. While this range is specific to the evaluated concentrator, it is representative, as this exemplary concentrator closely matches the mean and median focal length (see Section 4.5.1).

The Y-component mappings within this range are shown in Fig. D.2 for the surface $\mathcal{S}^{\text{CF}}$ and in Fig. D.3 for the surface $\mathcal{S}^{\text{M}}$, each for four exemplary distances in the evaluation range. Near the concentrator center, both surfaces approximately exhibit an identity mapping, as all facets contribute and averaging stabilizes the signal. This explains the similar system behavior observed for small-displacement trajectories (normal velocity profile shaping configurations C1, C6, C11, C16).

With increasing displacement, nonlinear distortions become more pronounced, especially at larger distances. For the surface $\mathcal{S}^{\text{M}}$, slope errors amplify these effects. In contrast, the surface $\mathcal{S}^{\text{CF}}$ exhibits smaller regions of identity due to the earlier loss of facet visibility, while deviations outside these regions remain comparatively moderate. Consequently, for large displacements (configuration C21), the surface $\mathcal{S}^{\text{CF}}$ behaves closer to $\mathcal{S}^{\text{F}}$, whereas differences are more pronounced for smaller displacements.




Fig. D.1. Mapping between the Y-component of the UAV position and the Y-component of the reflection signal (both facet-wise and averaged) for the surface S^{CF} , i.e., canted and perfectly flat facets. The UAV position is held constant at $(p_{\text{UAV}}^{\text{CCS}})_x = 0$ m. Each subplot corresponds to a different UAV–concentrator distance $d = (p_{\text{UAV}}^{\text{CCS}})_z$, expressed as a multiple of the concentrator focal length f .




Fig. D.2. Mapping between the Y-component of the UAV position and the Y-component of the reflection signal (both facet-wise and averaged) for the surface S^{CF} , i.e., canted and perfectly flat facets. The UAV position is held constant at $(p_{\text{UAV}}^{\text{CCS}})_x = 0$ m. Each subplot corresponds to a different UAV–concentrator distance $d = (p_{\text{UAV}}^{\text{CCS}})_z$.




Fig. D.3. Mapping between the Y-component of the UAV position and the Y-component of the reflection signal (both facet-wise and averaged) for the surface S^{M} , i.e., canted facets with measured slope errors. The UAV position is held constant at $(p_{\text{UAV}}^{\text{CCS}})_x = 0$ m. Each subplot corresponds to a different UAV–concentrator distance $d = (p_{\text{UAV}}^{\text{CCS}})_z$.

D.4. Implications for closed-loop behavior

If the UAV operates in regions where the mapping is smooth and approximately linear, the closed-loop behavior is largely independent of the surface realization.

In contrast, trajectories that traverse regions with strong non-linearities or facet-visibility transitions amplify surface-dependent effects. This explains the increased reflection displacement observed for fast trajectories and for surfaces with slope errors in the evaluation studies.

Appendix E. Supplementary reflection trajectories

This appendix presents additional reflection trajectories that are not included in the main manuscript in order to maintain clarity and readability. The figures below show an extended set of simulated reflection trajectories for the parameter combinations considered in the tuning and evaluation studies. While almost all reflection trajectories are included, a small number are omitted in cases where their inclusion would clutter the plots. This is done in cases where symmetry arguments allow the presentation of only a representative subset, or in cases where the system behavior can be deduced from neighboring parameter values and does not show any unexpected effects. The presented results provide a comprehensive overview of the system behavior across the configurations analyzed during tuning (Figs. E.1–E.5) and evaluation (Figs. E.6–E.9).




Fig. E.1. Tuning Study 1: X-components of the reflection trajectories for all starting points and proportional gains $k_P = 0.2, 0.4, 0.6, 0.8$.




Fig. E.3. Tuning Study 3: Y-components of the reflection trajectories for all simulated normal velocities and integral gains $k_I = 0, 0.02, 0.03, 0.05$.




Fig. E.2. Tuning Study 2: Y-components of the reflection trajectories for all simulated normal velocities and proportional gains $k_P = 0.2, 0.4, 0.6, 0.8$.




Fig. E.4. Tuning Study 4: Y-components of the reflection trajectories for all simulated normal velocities and the elevation angles $el = 0^\circ, 30^\circ, 45^\circ, 60^\circ$.




Fig. E.5. Tuning Study 5: Y-components of the reflection trajectories for all simulated normal velocities $v_{\text{control,Z}}^{\text{CCS}}$ and rise times t_{rise} .




Fig. E.7. Evaluation Study 2: Y-components of the reflection trajectories for the normal velocity profile shaping configurations C1, C6, C11, C16 and C21 with and without the application of the normal vector refinement (NVR).




Fig. E.6. Evaluation Study 1: Y-components of the reflection trajectories for the normal velocity profile shaping configurations C1, C6, C11, C16 and C21 and for all concentrator surfaces.




Fig. E.8. Evaluation Study 3: X- and Y-components of the reflection trajectories for the normal velocity profile shaping configurations C1, C6, C11, C16 and C21 for the in-plane rotations $\psi = 0^\circ, 30^\circ, 60^\circ, 90^\circ$.




Fig. E.9. Evaluation Study 4: Y-components of the reflection trajectories for the normal velocity profile shaping configurations C1, C6, C11, C16 and C21 and for the elevation angles $el = 0^\circ, 15^\circ, 30^\circ, 45^\circ$.

Appendix F. Supplementary data

Supplementary data for this article can be found online at doi:10.1016/j.solener.2026.114823.

Data availability

Data will be made available upon request.

References

- [1] S.A. Jones, K.W. Stone, Analysis of solar two heliostat tracking error sources, Sandia National Laboratories (SNL), Albuquerque, NM, and Livermore, CA (United States), 1999, <https://www.osti.gov/biblio/3312>.
- [2] J.C. Sattler, M. Röger, P. Schwarzbözl, R. Buck, A. Macke, C. Raeder, J. Götsche, Review of heliostat calibration and tracking control methods, *Sol. Energy* 207 (2020) 110–132, <https://doi.org/10.1016/j.solener.2020.06.030>.
- [3] R.A. Mitchell, G. Zhu, A non-intrusive optical (NIO) approach to characterize heliostats in utility-scale power tower plants: methodology and in-situ evaluation, *Sol. Energy* 209 (2020) 431–445, <https://doi.org/10.1016/j.solener.2020.09.004>.
- [4] J. Yellowhair, P.A. Apostolopoulos, D.E. Small, D. Novick, M. Mann, Development of an aerial imaging system for heliostat canting assessments, *AIP Conf. Proc.* 2445 (1) (2022) 120024, <https://doi.org/10.1063/5.0087057>.
- [5] W. Jessen, M. Röger, C. Prah, R. Pitz-Paal, A two-stage method for measuring the heliostat offset, *AIP Conf. Proc.* 2445 (1) (2022) 070005, <https://doi.org/10.1063/5.0087036>.
- [6] J.J. Krauth, C. Happich, N. Algner, R. Broda, A. Kämpgen, A. Schnerring, S. Ulmer, M. Röger, HelioPoint – a fast airborne calibration method for heliostat fields, *J. Sol. Energy Eng.* 146 (6) (2024) 061005, <https://doi.org/10.1115/1.4065868>.
- [7] A. Schnerring, R. Broda, M. Nieslony, N. Algner, E. Saez Martinez, M. Röger, S. Kallio, R. Pitz-Paal, Airborne real-time solar concentrator orientation estimation for heliostat coarse calibration, *Sol. Energy* 310 (2026) 114495, <https://doi.org/10.1016/j.solener.2026.114495>.
- [8] G. Zhu, C. Augustine, R. Mitchell, M. Muller, P. Kurup, A. Zolan, S. Yellapantula, R. Brost, K. Armijo, J. Sment, R. Schaller, M. Gordon, M. Collins, J. Coventry, J. Pye, M. Cholette, G. Picotti, M. Arjomandi, M. Emes, D. Potter, M. Rae, HelioCon: a roadmap for advanced heliostat technologies for concentrating solar power, *Sol. Energy* 264 (2023) 111917, <https://doi.org/10.1016/j.solener.2023.111917>.
- [9] A. Winter, Enhancement of a Simulation Environment for UAV-Based Condition Monitoring Systems of Heliostat Fields (Master's thesis), Technical University of Hamburg, 2024, <https://elib.dlr.de/206394/>.
- [10] A. Schnerring, R. Broda, A. Winter, M. Nieslony, J.J. Krauth, M. Röger, S. Kallio, R. Pitz-Paal, A simulation environment for UAV-based real-time condition monitoring of solar tower power plants, *Sol. Energy* 300 (2025) 113803, <https://doi.org/10.1016/j.solener.2025.113803>.
- [11] M. Röger, T. Schlichting, K. Blume, Guidelines for heliostat testing, In G. Zhu, M. Röger, Z. Wang (Eds.), *Advances in Solar Energy: Heliostat Systems Design, Implementation, and Operation*, vol. 12671, International Society for Optics and Photonics, 2023, <https://doi.org/10.1117/12.2682675>, p. 1267106.
- [12] M. Röger, et al., SolarPACES guideline for heliostat performance testing, version 1.0, 2023, <https://doi.org/10.5281/zenodo.15365475>.
- [13] R. Hartley, A. Zisserman, *Multiple View Geometry in Computer Vision*, 2 ed, Cambridge University Press, 2004, <https://doi.org/10.1017/CBO9780511811685>.
- [14] S.J. Lombard, W.J. Smit, Practical challenges to calibrate a heliostat with a multi-copter, *AIP Conf. Proc.* 2445 (1) (2022) 050005, <https://doi.org/10.1063/5.0085695>.
- [15] A. Macario Barros, M. Michel, Y. Moline, G. Corre, F. Carrel, A comprehensive survey of visual SLAM algorithms, *Robotics* 11 (1) (2022), <https://doi.org/10.3390/robotics11010024>.
- [16] D. Seborg, T. Edgar, D. Mellichamp, F. Doyle, *Process Dynamics and Control*, Wiley, 2016, <https://books.google.de/books?id=ZZVFEAAAQBAJ>.
- [17] K. Åström, R. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, second ed, Princeton University Press, 2021, <https://books.google.de/books?id=qZODEAAAQBAJ>.
- [18] R. Broda, A. Schnerring, D. Schnaus, M. Nieslony, J.J. Krauth, M. Röger, S. Kallio, R. Triebel, R. Pitz-Paal, Bridging the sim2real gap: training deep neural networks for heliostat detection with purely synthetic data, *Sol. Energy* 300 (2025) 113728, <https://doi.org/10.1016/j.solener.2025.113728>.
- [19] C.R. Harris, K.J. Millman, S.J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N.J. Smith, R. Kern, M. Picus, S. Hoyer, M.H. van Kerkwijk, M. Brett, A. Haldane, J.F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, T.E. Oliphant, Array programming with NumPy, *Nature* 585 (7825) (2020) 357–362, <https://doi.org/10.1038/s41586-020-2649-2>.
- [20] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S.J. van der Walt, M. Brett, J. Wilson, K.J. Millman, N. Mayorov, A.R.J. Nelson, E. Jones, R. Kern, E. Larson, C.J. Carey, Í. Polat, Y. Feng, E.W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimman, I. Henriksen, E.A. Quintero, C.R. Harris, A.M. Archibald, A.H. Ribeiro, F. Pedregosa, P. van Mulbregt, SciPy 1.0 Contributors, SciPy 1.0: fundamental algorithms for scientific computing in Python, *Nat. Methods* 17 (2020) 261–272, <https://doi.org/10.1038/s41592-019-0686-2>.
- [21] G. Bradski, *The OpenCV library*, Dr. Dobb's J. Softw. Tools (2000) 2236121.
- [22] ArduPilot Team, ArduPilot, <https://ardupilot.org/>.
- [23] A. Koubâa, A. Allouch, M. Alajlan, Y. Javed, A. Belghith, M. Khalgui, Micro air vehicle link (MAVLink) in a nutshell: a survey, *IEEE Access* 7 (2019) 87658–87680, <https://doi.org/10.1109/ACCESS.2019.2924410>.
- [24] S. Ulmer, T. März, C. Prah, W. Reinalter, B. Belhomme, Automated high resolution measurement of heliostat slope errors, *Sol. Energy* 85 (4) (2011) 681–687, <https://doi.org/10.1016/j.solener.2010.01.010>.
- [25] K. Deb, Multi-objective optimization using evolutionary algorithms: an introduction, 2011, <https://www.egr.msu.edu/kdeb/papers/k2011003.pdf>.
- [26] G.A.F. Seber, A. Lee, *Linear Regression Analysis*, John Wiley and Sons, 2003.
- [27] S.M. Kay, *Fundamentals of statistical signal processing*, in: Volume I: Estimation Theory, Prentice Hall PTR, Upper Saddle River, NJ, 1993.