

# Accurate and scalable receiver-level flux prediction: A fully data-driven solution

Mathias Kuhl <sup>a</sup>,\* , Max Pargmann <sup>a</sup>, Daniel Maldonado Quinto <sup>a</sup>, Robert Pitz-Paal <sup>a,b</sup>

<sup>a</sup> DLR, Institute of Solar Research, Linder Höhe, Köln, 51147, Germany

<sup>b</sup> RWTH Aachen University, Chair of Solar Technology, Linder Höhe, Köln, 51147, Germany

## ARTICLE INFO

Dataset link: <https://paint-database.org/>, <https://hdl.handle.net/21.11152/474a4b1c-de93-4d4a-b33d-1d32d63baf4b>

### Keywords:

Solar power tower  
Flux density prediction  
Camera-target method  
Heliostat  
Machine learning

## ABSTRACT

Concentrated Solar Technologies (CST) systems, particularly central tower configurations with heliostat fields, play a critical role in the renewable energy landscape. By focusing sunlight from thousands of heliostats onto a central receiver, these systems generate high-temperature heat, which serves as a key resource for dispatchable power generation and industrial processes. Accurate receiver-level flux prediction, which depends on precise heliostat characterization, is essential for optimizing efficiency and operational control. However, existing characterization methods face trade-offs between accuracy and scalability, limiting their practicality for large-scale deployment.

To overcome these limitations, this study introduces a fully data-driven framework that unifies heliostat characterization and flux prediction, leveraging operational data from standard calibration procedures. Expanding upon previous work that employed StyleGAN for beam-characterization-based predictions, this approach advances the methodology to achieve accurate receiver-level flux predictions. While the prior method demonstrated a proof of concept for a unified data-driven approach, it remained constrained to flux predictions on the calibration target itself. This study introduces key innovations, including aim point generalization strategies and a novel receiver projection technique, effectively bridging the gap between beam-characterization-based heliostat characterization and accurate receiver-level flux predictions.

The proposed Transformer-based architecture achieves receiver-level focal spot prediction errors below 12%, exceeding the accuracy of state-of-the-art deflectometry-enhanced ray tracing. By relying exclusively on standard calibration images, the method remains both cost-efficient and scalable, offering a practical solution for large-scale CST applications.

## 1. Introduction

Concentrated Solar Technology (CST) systems, particularly solar power tower plants, are a cornerstone of sustainable energy production, offering the unique advantage of dispatchable power generation and solar fuel production [1,2]. These systems employ extensive heliostat fields to track the sun and concentrate sunlight onto a central receiver, where solar energy is converted into high-temperature thermal energy. This energy plays a critical role in electricity generation, industrial processes, and thermal storage, solidifying CST's importance within the renewable energy landscape.

Achieving an optimal flux density distribution on the receiver surface is paramount for ensuring operational efficiency, safety, and the longevity of the receiver. The spatially distributed flux significantly

influences the receiver's thermal performance, where adherence to thermal thresholds is essential to prevent damage and maintain durability. However, achieving uniform energy absorption under dynamic solar conditions remains a considerable challenge.

### 1.1. Challenges in flux prediction

The flux density on a CST receiver arises from the superposition of individual flux distributions emitted by each heliostat. Accurate control of this distribution necessitates precise prediction of focal spot shapes, which are influenced by heliostat-specific mirror errors. Capturing these errors requires additional measurements to characterize heliostat properties. While deflectometry and photogrammetry are known

\* Corresponding author.

E-mail address: [mathias.kuhl@dlr.de](mailto:mathias.kuhl@dlr.de) (M. Kuhl).

for their high accuracy in this regard, their lack of robustness, high cost, and limited scalability render them impractical for application in large-scale heliostat fields.

Recent advancements have shifted the focus toward leveraging operational data, such as calibration target images, for heliostat characterization and focal spot prediction. These Beam Characterization methods eliminate the need for extensive additional measurements by utilizing existing operational data [3–6]. However, existing approaches are constrained by the limited information available in focal spots derived from calibration images, which makes it challenging to infer heliostat surface details with the precision achievable by more direct measurement methods.

### 1.2. Proposed approach

Building upon the prior work of Kuhl et al. [7], which demonstrated a proof of concept for addressing information loss through a unified purely data-driven framework for heliostat characterization and flux prediction, this study extends and refines the approach to tackle practical implementation challenges. The proposed methodology introduces key advancements:

- **Generalization Across Aim Points:** Two strategies support generalization across varying aimpoints: reparametrization aligns input conditions to enhance learning, while artificial data extends the input range, enabling accurate predictions for unseen aim points.
- **Projection to Receiver Surfaces:** A novel algorithm projects heliostat flux distributions from flat aperture planes onto arbitrary receiver surfaces, including cavity reactors.
- **Transformer-Based Architecture:** The initial StyleGAN architecture is replaced with a transformer-based encoder–decoder structure, offering practical advantages in scalability, robustness, and achieving higher prediction accuracy.

These advancements bridge the gap between calibration target data and receiver-level predictions in a fully data-driven manner. By leveraging a transformer-based architecture, the proposed method achieves robust and accurate flux predictions that surpass the performance of deflectometry-enhanced ray tracing—currently considered the most accurate state-of-the-art approach. Unlike deflectometry, our method remains scalable and practical by relying solely on calibration target images for heliostat characterization. In future applications, the increased predictive accuracy and robustness of the proposed method support more advanced aim point optimization strategies within CST plant operations. Improved control over flux distribution is expected to reduce localized flux peaks, enhance spatial uniformity across the receiver surface, and thereby contribute to higher thermal efficiency and reduced material stress.

## 2. State of the art

### 2.1. Camera-target calibration process

Our approach utilizes images acquired through the camera-target method, first introduced by Stone [8]. This process involves shifting the focal spot of a heliostat onto a Lambertian target located beneath the receiver. A strategically placed camera within the heliostat field captures the light reflected off the target, using the captured image to identify the focal spot's centroid. This information is crucial for adjusting a heliostat alignment model, enabling it to predict the focal spot's location based on the heliostat's actuator motor positions and the sun position. Recognized as a fundamental procedure in power plant operations for calibrating heliostats alignment models [9], its effectiveness has been validated in various studies, including work by Smith and Ho [10], Guangyu and Zhongkun [11] or Pargmann et al. [12].

### 2.2. Predicting heliostat focal spots

Accurately predicting the shape or distribution of heliostat focal spots is critical for optimizing aimpoint distributions and compensating for mirror imperfections such as canting errors, facet misalignments or mirror deformations. Traditional methods for focal spot computation, including Convolution Methods and Monte Carlo Ray Tracing, have been widely used in solar power tower systems.

Convolution methods simulate a heliostat's focal spot by approximating its flux distribution through the superposition of error cones, representing scattering effects from various error sources. These sources are mathematically modeled and aggregated into a composite flux distribution. While computationally efficient, such methods simplify mirror imperfections, limiting accuracy [13,14].

Monte Carlo Ray Tracing offers a more detailed approach by simulating individual light rays interacting with heliostat surfaces. This method can account for detailed surface deformations achieving high accuracy in flux predictions [15–17]. However, it relies on precise surface characterization using techniques like deflectometry or photogrammetry. These heliostat characterization methods, while accurate, are costly, time-intensive, and sensitive to environmental conditions, limiting their scalability for large heliostat fields [18,19].

Beam Characterization methods provide a practical alternative by leveraging the information contained in focal spot measurements rather than direct surface characterization. These methods analyze the reflected focal spot to infer heliostat properties.

Basic shape approximation models use simplified convolution techniques to fit Gaussian or similar distributions to measured focal spots. While computationally efficient, these models often fail to capture detailed distortions caused by mirror imperfections [20].

More advanced methods optimize facet alignment by fitting simulated flux distributions to observed focal spots. Techniques like those proposed by Sánchez-González et al. [21] refine canting and alignment parameters, achieving improved accuracy without requiring detailed physical measurements.

The most recent approaches reconstruct detailed heliostat surface deformations from focal spot measurements. For instance, Pargmann et al. [6] employed differentiable ray tracing to optimize surface parameters based on calibration target images, achieving high accuracy. However, this method is computationally intensive and less effective for heliostats located far from the calibration target due to information loss during beam concentration. Similarly, Martínez-Hernández et al. [22] proposed a movable target system to reduce information loss by positioning the target closer to the heliostat. While effective, this approach requires additional equipment and increases operational complexity, limiting its scalability.

Building on these methods, Kuhl et al. [7] demonstrated that the inherent information loss in beam characterization can be addressed by integrating heliostat characterization and flux prediction into a unified framework. They employed a StyleGAN-based data-driven approach that generalized focal spot predictions using data from the entire heliostat field. While this approach achieved superior accuracy, it was limited in practical application due to its restriction to the calibration target and inability to predict receiver-level flux distributions.

Existing methods for heliostat characterization reveal a fundamental trade-off between accuracy and scalability. Techniques such as deflectometry and photogrammetry offer high accuracy by directly capturing surface deformations, but they require extensive, time-consuming measurements and are highly sensitive to environmental conditions, limiting their feasibility in large-scale deployments. In contrast, beam characterization methods leverage focal spot images for scalable in-field application but may suffer from reduced accuracy due to information loss during beam concentration. For some approaches, their robustness can also diminish for heliostats located farther from the calibration target.

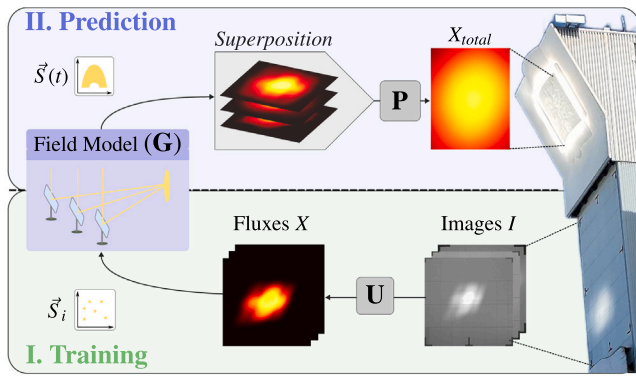


Fig. 1. Overview of the proposed data-driven methodology (Illustration adapted from [7]). Observed target images  $I$ , captured during calibration, are processed into relative flux distributions  $X$  using a UNet (U) as described by Kuhl et al. [23]. These flux distributions  $X$ , obtained under various sun positions  $\vec{S}_i$ , are utilized to train a generalized heliostat model (G). The trained model predicts focal spots for any given sun position  $\vec{S}(t)$  and aim point  $\vec{A}_p$  in the aperture plane of the receiver. These predictions are subsequently projected onto the receiver surface, enabling the aggregation of individual heliostat focal spots into the total flux density  $X_{total}$ .

To address these challenges, this study builds upon the previously introduced StyleGAN-based framework by Kuhl et al. [7], which demonstrated the feasibility of integrating heliostat characterization and flux prediction into a unified data-driven approach. Expanding on this foundation, we introduce a transformer-based architecture with enhanced generalization capabilities. By incorporating strategies for aim point extrapolation and receiver surface projection, the proposed method bridges the gap between calibration target data and accurate receiver-level flux prediction, offering a robust and scalable solution.

### 3. Method

The methodology, depicted in Fig. 1, outlines the progression from raw image acquisition to receiver-level flux prediction. The chapter is structured as follows:

- **Translation of Raw Images to Flux Distributions (U):** Raw calibration images are transformed into relative flux distributions using an image-to-image neural network, as introduced in Kuhl et al. [23].
- **Generalizing Flux Prediction Model (G):** A generalizing model is trained on all focal spot measurements, enabling accurate predictions of focal spot shapes in the aperture plane under varying conditions, including sun positions and aim points.
- **Projection and Superposition on the Receiver (P):** Predicted focal spots are scaled to absolute flux values using an efficiency model, projected from the aperture plane to the receiver surface, and aggregated into total flux distributions across receiver bins.

By introducing generalization across aim point conditions and employing a novel projection method, this approach facilitates fully data-driven receiver-level flux prediction based solely on calibration images. The implementation of each step is elaborated in the subsequent sections.

#### 3.1. Translation of raw images to flux distributions

The initial step involves transforming raw calibration images into precise relative flux distributions. This process employs the image-to-image neural network proposed in Kuhl et al. [23], which facilitates robust, automated flux derivation from calibration images while accounting for background illumination and target reflectivity. The network is based on a UNet architecture and serves as a preprocessing step

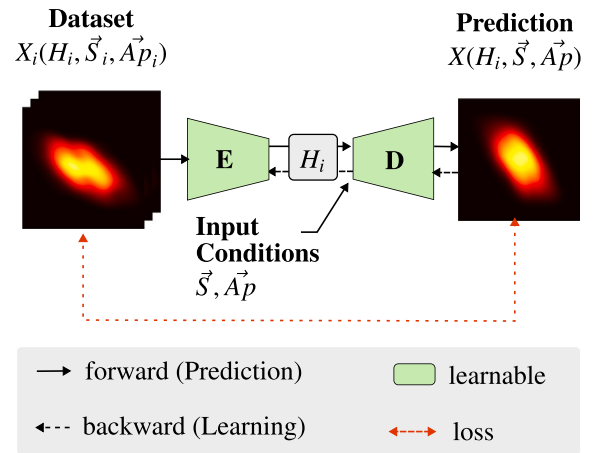


Fig. 2. Encoder-Decoder Method. All calibration data for a heliostat  $H_i$  are processed by the encoder (E) to generate its latent representation  $\vec{H}$ . This representation, encapsulating the heliostat's properties, is combined with input conditions – aim point  $\vec{A}_p$  and sun position  $\vec{S}$  – and fed into the decoder (D). The decoder predicts the focal spot for the specified heliostat under the given conditions.

for raw calibration target images. It is trained on artificially generated data to separate each image into background and spatially resolved relative flux distributions, while accounting for spatial variations in target reflectivity. This design ensures consistent flux extraction across varying illumination conditions and target inhomogeneities. The neural network generates relative flux distributions scaled to a mean value of 0.1. This scaling emphasizes the generalizing model (G) on predicting the shape of the flux distribution, independent of absolute energy values, which are affected by external factors such as DNI and mirror reflectivity. The conversion to absolute flux values is addressed in a subsequent step.

#### 3.2. Generalizing flux prediction model

This fully data-driven methodology advances from individual heliostat characterization to a comprehensive framework for generalizing flux predictions across the entire heliostat field. By incorporating data from multiple heliostats into a single model, the network learns shared patterns across the field, ensuring scalability and robustness. This unified framework integrates heliostat characterization and flux prediction, leveraging neural networks to directly model focal spot variations under diverse conditions such as sun position and aim point (see Fig. 2).

The framework consists of two core components: an encoder, analogous to a Heliostat Characterization Technique, and a decoder, serving as a Flux Prediction Method. The encoder translates focal spot measurements into latent heliostat representations ( $\vec{H}$ ), while the decoder utilizes these representations along with sun position  $\vec{S}$  and aim point  $\vec{A}_p$  to generate focal spot predictions  $X$ . This process is mathematically formalized as:

$$\vec{H}_i = \mathbf{E}(X_i), \quad (1)$$

$$X = \mathbf{D}(\vec{H}_i, \vec{S}, \vec{A}_p), \quad (2)$$

where  $\mathbf{E}$  denotes the encoder and  $\mathbf{D}$  the decoder. This end-to-end pipeline directly optimizes the target value – flux distributions – ensuring high predictive accuracy for focal spot shapes.

Training the model on data from the entire heliostat field allows for robust generalization across different heliostats and conditions. Significant modifications to the transformer network have been implemented to tailor it specifically for focal spot prediction. Further details on these adjustments, including implementation and training procedures, are outlined in Appendix.




Fig. 3. Distribution of input conditions between training and prediction cases, represented as aim point  $\vec{A}_p$  (left) and aim direction  $\vec{A}_d$  (right). Training is performed on targets  $T_{1,2}$  and prediction on target  $T_3$  (see Fig. 5).

### 3.2.1. Advancing encoding techniques for heliostat properties

This study transitions from the single-generator approach of Kuhl et al. [7] to an encoder–decoder architecture, addressing operational inefficiencies while enhancing accuracy and flexibility. In the earlier StyleGAN-based method, heliostat properties ( $\vec{H}$ ) were represented as latent vectors refined iteratively through backpropagation during training. While effective in capturing complex relationships, this approach required retraining or fine-tuning for new calibration data, limiting scalability in field operations.

The encoder–decoder methodology introduces an explicit encoder to compute  $\vec{H}$  directly from calibration images. During training, the encoder learns to encode heliostat properties in a single forward pass. This allows new calibration points to be easily encoded without retraining, significantly enhancing practical applicability in CST operations.

Additionally, the encoder–decoder structure demonstrates superior focal spot prediction accuracy compared to the original StyleGAN architecture, leveraging recent advancements in machine learning.

Moreover, the dedicated encoder enables targeted training strategies, such as encoding heliostat properties from limited data or enhancing aim point extrapolation by strategically separating data ranges for encoding and prediction. This structured approach improves the model’s ability to generalize to previously unseen scenarios.

The transition to this architecture marks a significant advancement, seamlessly integrating the flexibility and efficiency of transformer networks. This approach ensures scalable and accurate heliostat flux predictions, effectively addressing the limitations of the earlier StyleGAN-based approach.

### 3.3. Aim point extrapolation for generalized flux prediction

Accurate receiver-level flux prediction necessitates the ability to generalize beyond the conditions observed during calibration. Since aim points on calibration targets differ significantly from those on the receiver, models trained solely on calibration data face inherent limitations. To overcome this challenge, we introduce two complementary strategies: reparametrization of input conditions and the integration of artificial data.

**Reparametrization approach.** The aim point ( $\vec{A}_p$ ) of a heliostat is reformulated into an aim direction ( $\vec{A}_d$ ), characterized by its azimuth and elevation angles and the vector’s length to the target. This reparametrization establishes a more general and invariant representation of aim conditions, enhancing alignment between training and prediction input spaces. As illustrated in Fig. 3, the reformulation reduces the likelihood of encountering unseen input conditions during prediction by leveraging the natural variations in aim directions present in the training data. Furthermore, it captures the nuanced relationship between changes in aim direction and resulting focal spot shapes, thereby facilitating improved generalization. This approach allows the model to reliably extend its predictions beyond the calibration target, aligning with operational requirements.

**Artificial data for enhanced extrapolation.** Artificial datasets, generated through ray tracing and deflectometric measurements, can be incorporated into the training process when available. These datasets provide highly realistic focal spot samples across an expanded range of input conditions, including those beyond the calibration target. By broadening the coverage of input conditions, artificial data significantly enhances the model’s capability to generalize to unseen scenarios. Moreover, the inclusion of controlled variability mitigates the risk of overfitting to specific calibration target configurations. This strategy also enables knowledge transfer from a small subset of heliostats with deflectometrically measured properties to the entire heliostat field, bolstering the model’s scalability and robustness.

The reparametrization approach provides a robust and invariant representation of aim conditions, while artificial data, when available, extends the model’s ability to extrapolate effectively. Together, these strategies address the challenges posed by the disparity between calibration and receiver-level conditions.

### 3.4. Projection method for receiver-level flux prediction

A purely data-driven approach, trained exclusively on flat calibration data, is inherently limited to predicting flux distributions on flat surfaces. Extending this capability to complex receiver geometries requires a projection technique that maps focal spot predictions from the flat aperture plane to arbitrary receiver surfaces. To address this, we developed a ray tracing-based method that samples rays directly from the predicted flux distribution in the aperture plane.

Each ray is defined by its origin (the center of a pixel in the aperture plane), its energy (calculated from the pixel’s flux value and area), and its direction. While the primary beam direction, defined as the vector connecting the heliostat and aim point, is known, the lack of explicit information about heliostat canting, focusing properties, mirror deformations, and facet orientations in a data-driven approach complicates the precise calculation of ray directions. To overcome this challenge, we integrate a lightweight convolutional neural network (CNN), consisting of only three convolutional layers, into the pipeline, designed to predict ray directions based on the focal spot shape:

$$D = \text{CNN}(X, d), \quad (3)$$

where  $D$  is an array of 3D direction vectors corresponding to the resolution of the input focal spot  $X$ , and  $d$  is the distance between the heliostat and aim point, included as an additional input since it influences the ray directions. The CNN is trained on simulated data generated using a ray tracer and deflectometric measurements from a small set of heliostats.

This approach enables the CNN to accurately learn how rays converge or diverge based on the focal spot shape and distance, facilitating the projection of any focal spot from the aperture plane onto complex receiver geometries. While leveraging standard ray tracing principles, the method significantly reduces computational complexity compared to traditional ray tracing by starting with a pre-predicted flux distribution, treated as a cross-section of the heliostat beam. The relatively short distance between the aperture plane and the receiver surface (e.g., 0–3 m) minimizes the required number of rays, achieving comparable accuracy with only approximately 0.1% of the rays used in conventional ray tracing. For scenarios requiring higher precision or resolution, such as intricate receiver designs, the focal spot image can be upsampled to generate additional rays, further enhancing accuracy while maintaining computational efficiency.

### 3.5. Intensity scaling to absolute flux values

Heliostat characterization and flux prediction methods primarily aim to reconstruct the spatial shape and distribution of the focal spot. Accordingly, the primary focus of this work is to evaluate and improve the spatial accuracy of such relative flux predictions.

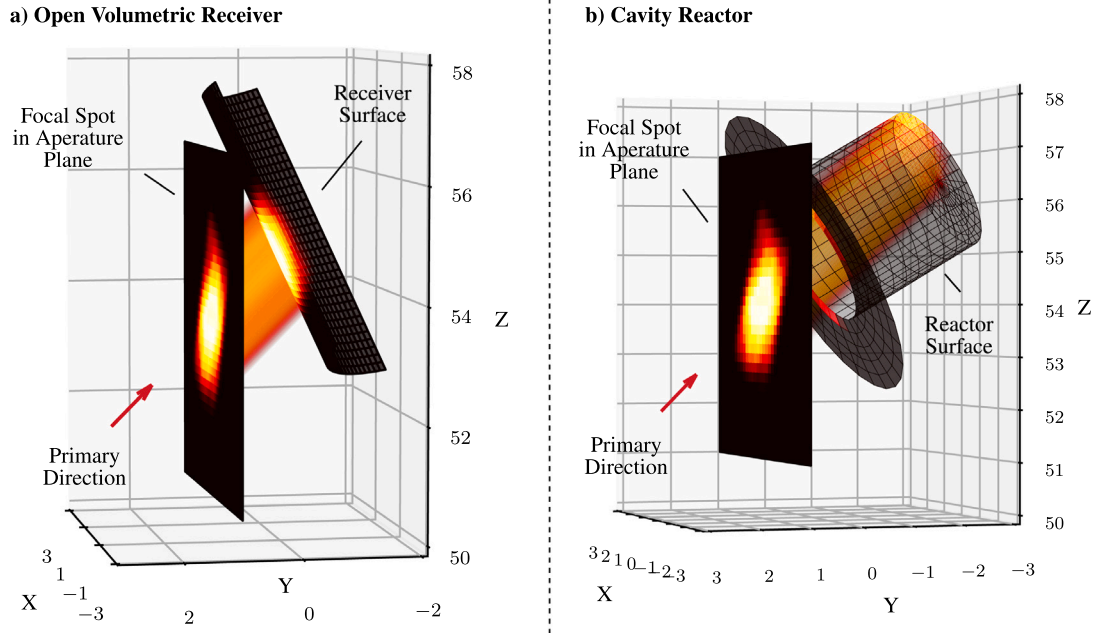


Fig. 4. Projection method: Translating focal spot predictions from the flat aperture plane to complex receiver geometries using a ray tracing-based approach.

However, for practical applications, it is necessary to convert these relative distributions into absolute radiative flux values using an efficiency model. This efficiency model estimates the total energy flux of the heliostat beam by incorporating key parameters such as Direct Normal Irradiance (DNI), mirror reflectivity, and the effective mirror area [23]. The resulting beam power serves as a scaling factor for the normalized flux distribution. Once scaled, the absolute flux predictions can be superposed across the heliostat field and used as input for downstream applications involving receiver performance evaluation and control system integration.

It is important to note that the efficiency model used to convert relative flux predictions into absolute units ( $\text{kW}/\text{m}^2$ ) is independent of the underlying flux prediction method—be it convolutional methods, ray tracing, or the data-driven approach proposed in this work. As it does not influence the spatial shape of the predicted distribution, it is excluded from the core evaluation in this study.

### 3.6. Total error estimation in the data-driven pipeline

A complete end-to-end evaluation – from calibration target images to flux distribution on the receiver surface – is not feasible in our case. This limitation arises from the complex surface structure and non-Lambertian properties of the receiver, which introduce measurement uncertainties ranging from 10%–40% [24]. To ensure a reliable assessment, we therefore decompose the pipeline into its individual components, each of which can be evaluated accurately. We then estimate the overall prediction error using a total error analysis.

The data-driven pipeline consists of multiple interconnected components: the UNet framework (U) for translating raw images to flux values, the Transformer network (G) for predicting focal spots in the aperture plane, and the projection stage (P) for mapping focal spots onto the receiver surface. To evaluate the pipeline's overall reliability, it is crucial to estimate the cumulative error across these stages, accounting for both independent errors and their propagation.

The total error  $E_{\text{total}}$  is expressed as:

$$E_{\text{total}} = \sqrt{E_U^2 + E_G^2 + E_P^2 + 2 \cdot \text{Cov}(E_U, E_G) + 2 \cdot \text{Cov}(E_G, E_P)},$$

where  $E_U$ ,  $E_G$ , and  $E_P$  denote the errors from the UNet, Transformer, and projection stages, respectively. The covariance terms ( $\text{Cov}(E_U, E_G)$ ,

$\text{Cov}(E_G, E_P)$ ) capture the dependencies between adjacent components, ensuring a holistic error assessment.

To compute these covariances, Monte Carlo simulations are employed, introducing realistic perturbations at intermediate stages and quantifying their effects on subsequent predictions. This approach assumes Gaussian error distributions, first-order dependencies, and sequential error propagation, reflecting the modular structure of the pipeline.

While this analysis focuses on the errors intrinsic to the heliostat flux prediction pipeline, external factors such as DNI variability and mirror reflectivity – independent of the prediction method – are excluded but must be considered in a comprehensive system-level evaluation.

### 3.7. Dataset

The foundation of any data-driven approach, particularly in machine learning, lies in the quality and scope of the underlying dataset. For this study, the objective is to predict focal spots  $X$  of various heliostats  $H_i$  under diverse sun positions  $\vec{S}$ . This necessitates a comprehensive dataset encompassing a wide range of heliostat and solar conditions.

**Data collection and processing.** The dataset is sourced from the calibration process conducted at the research power plant Solar Tower Jülich (STJ) [25]. Spanning the years 2022 and 2023, the calibration process produced approximately 200,000 target images from approximately 1000 heliostats. Each image is labeled with its corresponding heliostat ID  $H_i$ , sun position  $\vec{S}$ , and the derived aim point vector  $\vec{A}_p$ . These raw images are processed using the UNet framework [23] to derive relative flux distributions  $X$ .

Each image is cropped to a resolution of  $64 \times 64$  pixels, corresponding to physical dimensions of  $6 \times 6$  meters on the calibration target, centering on the focal spot's intensity. This cropping mitigates tracking errors and ensures that the analysis remains unaffected by such errors.

**Experimental setup.** Fig. 5 illustrates the setup of the STJ, which features two towers and three distinct Lambertian targets ( $T_{1-3}$ ). The availability of multiple targets allows for rigorous evaluation of the model's aim point extrapolation capabilities. Specifically, the unique




Fig. 5. Dataset overview and experimental setup: Left — Illustration of the two towers and three Lambertian targets ( $T_{1-3}$ ) at the Solar Tower Jülich (STJ). Right — Data collection and processing pipeline for the encoder-decoder model.

two-tower configuration enables training the network on focal spots observed at targets  $T_1$  and  $T_2$  from the left tower and testing its predictions on target  $T_3$  of the second tower. This setup ensures the evaluation of extrapolation over distances greater than those to the receiver while maintaining the high measurement accuracy provided by Lambertian targets, which would not be achievable on the receiver surface due to its non-Lambertian properties.

**Heliostat characteristics.** While the proposed method is independent of explicit physical modeling, a brief description of the heliostats at the STJ power plant is provided for context. Each heliostat consists of four non-focusing rectangular facets, each with an area of  $1.92\text{m}^2$ , canted toward the receiver.

**Supplementary simulated data.** In addition to the calibration data, supplementary simulated data is generated using a ray tracer enhanced with deflectometric measurements. Due to the time-intensive nature of deflectometry, this data is limited to 185 heliostats. These simulated datasets serve two critical purposes. First, they provide a benchmark for comparison, allowing direct evaluation against the most accurate state-of-the-art flux prediction method. Second, it is tested whether they can enhance the network’s capabilities by improving extrapolation to aim point conditions not covered in the calibration data. Additionally, the artificial data is utilized to train the CNN responsible for the projection approach, enabling accurate mapping of flux distributions onto complex receiver geometries.

#### 4. Evaluation

This section evaluates the proposed methodology through a step-wise analysis. We assess focal spot prediction accuracy, aim point extrapolation, and receiver projection. A total error analysis combines individual contributions, and robustness is evaluated under limited data and rare conditions. These steps demonstrate the method’s accuracy and practical viability.

**Evaluation metrics.** Accurate evaluation of focal spot predictions is essential to validate the performance of the proposed data-driven methodology. Both visual and quantitative metrics are employed to assess the spatial accuracy of predicted flux distributions relative to




Fig. 6. Comparison of focal spot predictions between measured ground truth, deflectometry-based ray tracing, and the data-driven Transformer approach. Predictions are shown for five heliostats (columns) at varying distances  $d$ . All focal spots are normalized to the same energy flux.

ground truth measurements. Focal spot intensities are visualized using a “hot” colormap, with all focal spots normalized to identical summed pixel intensity, ensuring consistent energy representation. For quantitative evaluation, a pixel-wise loss function  $L$  is defined to measure the discrepancy between predicted ( $X$ ) and ground truth ( $\hat{X}$ ) flux values:

$$L = \frac{\sum_i |X_i - \hat{X}_i|}{\sum_i \hat{X}_i} \quad (4)$$

This loss function directly evaluates spatially resolved flux errors, capturing proportional discrepancies while remaining invariant to image resolution. By normalizing flux distributions, it ensures that predictions are compared on relative terms, unaffected by variations in absolute intensity due to factors like DNI fluctuations or reflectivity changes. This approach provides a robust and precise metric for assessing the accuracy of focal spot predictions in flux modeling, emphasizing spatial distribution over absolute magnitude.

##### 4.1. Focal spot comparison

Fig. 6 visually compares focal spot predictions from the transformer network and deflectometry based raytracing against ground truth flux measurements obtained from the calibration target.

Raytracing enhanced with deflectometry achieves high accuracy in predicting focal spot shapes and orientations, closely aligning with ground truth measurements. Nonetheless, slight size discrepancies are noted for some heliostats ( $H_{2-4}$ ), potentially arising from unaccounted environmental factors or subtle mirror imperfections. Additionally, minor errors in focal spot shape and orientation are observed for heliostats ( $H_{2-5}$ ).

The data-driven Transformer approach demonstrates robust performance in predicting focal spot size, shape and orientation, effectively capturing variations across heliostats and distances. While highly accurate, fine details, such as sharp edges visible in deflectometry-based predictions, are less pronounced ( $H_1$ ). This difference may be attributed to measurement constraints in the training data, such as overexposure, which limit the prominence of these features.

In summary, for very close heliostats, deflectometry excels in resolving finer details for closest heliostats, while data-driven approaches demonstrate greater robustness in accurately predicting focal spot size and orientation across varying conditions.

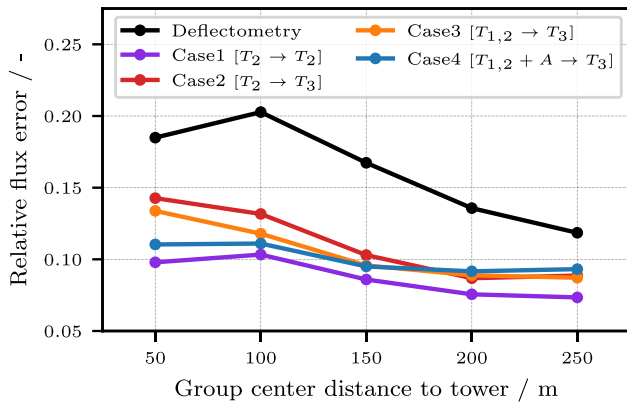


Fig. 7. Accuracy evaluation of the data-driven approach across various training and prediction cases, for the entire field grouped by heliostat distance, compared to deflectometry-based ray tracing.

#### 4.2. Accuracy evaluation for aim point extrapolation

The accuracy of the data-driven approach is evaluated across the entire heliostat field, with loss curves presented in Fig. 7 for varying training and evaluation datasets.

Case 1 (solid violet line) serves as the baseline, with both training and prediction performed on the same calibration target  $T_2$ . Case 2 (dashed violet line) evaluates predictions on  $T_3$  using data exclusively from  $T_2$ , introducing the challenge of aimpoint extrapolation to unseen conditions. Case 3 includes a wider range of calibration data from  $T_1$  and  $T_2$ , improving accuracy for closer heliostats (50–150 m) by leveraging greater aimpoint variability in the training data. Case 4 (dashed blue line) incorporates artificial data, further enhancing extrapolation capabilities for closer heliostats, though this effect diminishes for distant heliostats (> 150 m).

All cases outperform the deflectometry-based prediction method across all distances. The accuracy drop in Case 2 reflects the complexity of predicting unseen aimpoints. Introducing variability in training data (Case 3) and artificial data (Case 4) enhances accuracy for closer heliostats, as predicting aimpoint changes poses greater challenges for these heliostats. For heliostats at shorter distances, aimpoint changes result in larger variations in incidence angle and focal length, increasing the complexity of aimpoint extrapolation.

Notably, even the simplest model trained solely on data from a single target (Case 2) consistently outperforms deflectometry-based ray tracing, demonstrating the robustness and superiority of the data-driven approach across the heliostat field.

#### 4.3. Receiver projection evaluation

To evaluate the accuracy of the proposed projection approach, simulated data generated through ray tracing combined with deflectometry measurements is utilized. While deflectometry measurements may not perfectly align with measured focal spot data, the consistency within the simulation itself ensures robust evaluation across various evaluation planes and receiver surfaces.

**Projection depth analysis.** Projection errors were analyzed as a function of depth along the primary direction between the heliostat and the aimpoint. Focal spots were simulated at parallel planes beginning at the aperture, and the projection method was applied to the focal spot at the aperture plane. The resulting projections were then compared to simulated focal spots at different depths.

The CNN-based projection approach was benchmarked against two alternative assumptions for ray directions in the aperture plane: *Parallel*, where all rays align with the primary direction of the heliostat beam, and *Angled*, where each ray follows the connecting vector between the heliostat and the pixel location.

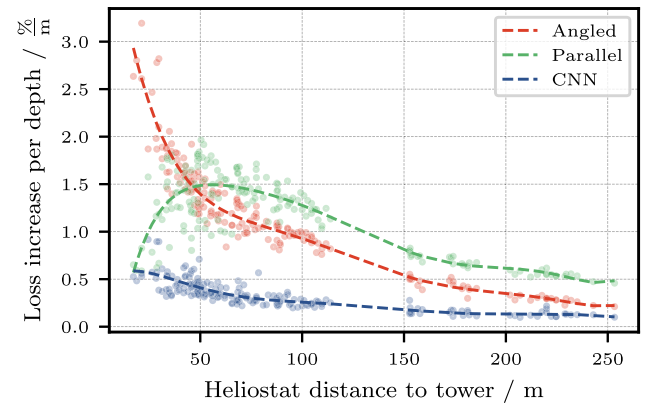


Fig. 8. Scatter plot of simulated heliostats showing the gradient of the depth loss curve for each heliostat, along with trendlines for the different projection methods.

The analysis revealed that projection losses increased nearly linearly with depth in the [0–5] m range for all heliostats. This linearity allows the depth analysis to be characterized by a single gradient, representing the percentage loss increase per meter for each heliostat. Fig. 8 illustrates these gradient values for heliostats located at varying distances from the tower. The *Parallel* method demonstrated superior accuracy for closer heliostats, where the beam remains predominantly parallel. In contrast, the *Angled* method performed better for distant heliostats, as the opening angle of the sun introduces greater beam divergence. The CNN-based projection approach consistently outperformed both methods across all distances, with error gradients ranging from 0.5% per meter for close heliostats to 0.1% per meter for distant ones. Overall, projection accuracy improved for all methods as the heliostat distance increased (> 50 m).

**Receiver surface evaluation.** The CNN-based projection method was further evaluated on the surface of two distinct receiver configurations (as shown in Fig. 4): the open volumetric receiver (OVR) and a cavity receiver. Flux distributions were simulated on both the aperture and receiver surfaces using ray tracing with deflectometry measurements. The projection method was applied to the focal spots at the aperture, and the resulting receiver flux distributions were compared against the simulated ground truth.

For the OVR, the mean deviation between the projected and simulated flux distributions across all heliostats was 2.59%. The cavity receiver exhibited a slightly lower mean deviation of 2.36% (2.82% when accounting for spillage). These deviations align with the per-meter loss gradients observed in the projection depth analysis. Despite the cavity receiver requiring a longer projection distance, its lower error can be attributed to the coarser resolution of its circular mesh, which introduces minor simplifications in the projected flux distribution.

These results validate the CNN-based projection method's ability to accurately project focal spot predictions onto complex receiver geometries, ensuring minimal error and broad applicability across diverse receiver configurations.

#### 4.4. Total error analysis

The total error  $E_{\text{total}}$  provides a comprehensive assessment of the pipeline, combining independent error contributions from each stage with their propagation effects. Fig. 9 summarizes the key findings:

The proposed pipeline consistently achieves lower total error than deflectometry-based methods, even under demanding conditions including far aimpoint extrapolation and receiver projection. The highest error contribution arises from focal spot predictions ( $E_G$ ), reflecting the complexity of this task. In contrast, errors from the UNet ( $E_U$ )

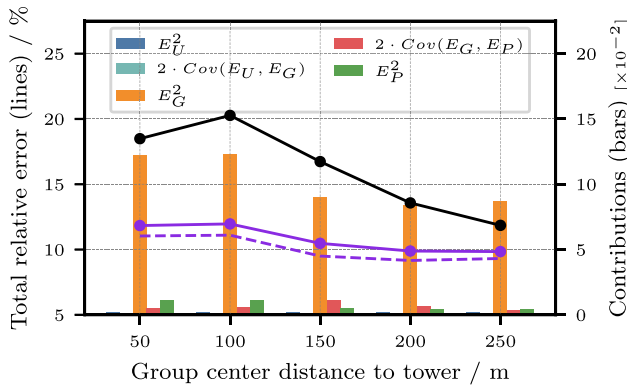


Fig. 9. Total error estimation for the complete pipeline, combining aimpoint extrapolation and receiver projection. The bars represent the individual error contributions, while the violet line indicate the total error of the data-driven pipeline. The dashed line shows the error  $E_G$  solely for comparison. Results are compared to the baseline deflectometry method (black line).

and the projection stage ( $E_P$ ) contribute minimally to the overall error, demonstrating that inaccuracies do not amplify or propagate significantly through the pipeline. The small covariance between stages further confirms the robustness of the error model and validates the assumption of first-order dependencies.

Overall, these results highlight the pipeline's capability to predict highly accurate receiver-level flux distributions, outperforming the most accurate state-of-the-art method -deflectometry-enhanced raytracing.

#### 4.5. Robustness and generalization evaluation

The robustness of the data-driven methodology was analyzed under conditions of scarce sun positions and limited training data. For edge-of-distribution (EOD) cases, representing rare sun positions typically encountered during early mornings or late evenings, the model exhibited slightly elevated loss values due to the underrepresentation of these conditions in the training data. Despite this, the accuracy degradation remained controlled, and these cases were deemed operationally insignificant due to their limited relevance in power plant operations.

The impact of dataset size on model performance was also assessed to ensure scalability to large heliostat fields. Results revealed that as few as five calibration images per heliostat are sufficient for effective field-wide training, enabling high predictive accuracy. For individual heliostats, even a single calibration image sufficed to achieve robust predictions when combined with a comprehensive dataset from the remaining heliostats. This highlights the model's ability to leverage shared information across the heliostat field, ensuring scalability and adaptability with minimal data requirements.

## 5. Discussion

The proposed data-driven pipeline represents a significant advancement in heliostat characterization and flux prediction, addressing critical limitations of traditional methodologies. By employing a transformer-based encoder-decoder architecture, the approach achieves robust generalization across diverse sun positions and aimpoints, ensuring scalability and adaptability. The incorporation of reparametrization strategies and artificial datasets further strengthens the model's extrapolation capabilities, enabling precise predictions for previously unseen aimpoints and conditions.

The projection methodology, supported by a CNN-based ray direction predictor, effectively maps focal spots from flat aperture planes to complex receiver geometries with minimal error. This versatility

ensures the applicability of the pipeline to a range of receiver configurations, which is essential for CST systems. Although the method was validated on a cavity receiver, the heliostat types typically used with such configurations often differ from the non-focusing heliostats investigated in this study. Thus, the accuracy of the projection approach for focusing heliostats remains to be demonstrated and should be a priority for future investigations.

Despite these accomplishments, the pipeline's reliance on high-quality calibration data underscores the critical role of accurate data preprocessing. Errors introduced during calibration may propagate throughout the pipeline, potentially impacting overall performance. Addressing this dependency through enhanced calibration techniques or robust preprocessing workflows is a key area for future improvement.

## 6. Conclusion and outlook

This study introduces a groundbreaking approach to heliostat flux prediction by integrating advanced machine learning techniques with innovative projection methods. The unified pipeline bridges the gap between calibration data and receiver-level flux predictions, delivering superior accuracy, scalability, and robustness. By leveraging transformer-based architectures, reparametrization strategies, and artificial datasets, the method effectively tackles key challenges in aimpoint extrapolation and receiver projection.

Looking ahead, integrating this methodology into operational CSP plants has the potential to significantly enhance efficiency, reduce costs, and improve overall system reliability. Future research should focus on the real-time deployment of the model, incorporating tracking errors and aimpoint optimization algorithms. This integration would allow for a comprehensive investigation into how heliostat focal spot predictions impact the accuracy of the total superposed flux density on the receiver.

The proposed framework provides a robust foundation for advancing CST technologies, with implications that extend to the broader landscape of sustainable energy systems. By addressing longstanding challenges in flux prediction, this work contributes to the ongoing optimization and scalability of renewable energy solutions, aligning with global efforts to achieve a more sustainable and energy-efficient future.

### CRedit authorship contribution statement

**Mathias Kuhl:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology. **Max Pargmann:** Writing – review & editing, Supervision. **Daniel Maldonado Quinto:** Writing – review & editing, Supervision, Funding acquisition. **Robert Pitz-Paal:** Supervision, Funding acquisition.

### Code availability

The code used to develop and evaluate the methods presented in this work is openly available via the **HelioBeamAI** repository maintained by the German Aerospace Center (DLR). The repository can be accessed at <https://github.com/DLR-SF/HelioBeamAI>.

### 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.

### Declaration of generative AI in scientific writing

Statement: During the preparation of this work, the authors used ChatGPT-4, developed by OpenAI, to reformulate and enhance the language for improved clarity and conciseness. After using this tool, the authors thoroughly reviewed and edited the content as needed and take full responsibility for the content of the publication.

## Acknowledgments

We gratefully acknowledge the use of data from the Solar Tower in Jülich, a research power plant operated by the German Aerospace Center (DLR). Additionally, we extend our thanks for the funding provided by the Helmholtz Association (HGF), Germany for the *GANCSTR* project (Grant Number ZT-I-PF-5-069), under which this research was conducted.

## Appendix. Modified transformer architecture

The Transformer architecture, originally introduced by Vaswani et al. [26], was adapted for heliostat flux prediction due to its strengths in processing variable-length inputs and leveraging attention mechanisms for weighting calibration data. These capabilities align well with the task of predicting heliostat focal spots under varying solar and aim point conditions. However, significant modifications were required to tailor the architecture to the challenges of image-based flux prediction.

### A.1. Architecture adaptations

Fig. A.10 illustrates the modified Transformer architecture, which incorporates the following changes:

- **Integration with a Variational Autoencoder (VAE):** A VAE [27] preprocesses focal spot images into compact latent representations suitable for the Transformer encoder. The VAE encoder compresses input images into latent vectors, which are combined with input conditions (e.g., sun position, aim point). The VAE decoder reconstructs flux distributions from the Transformer's predictions, ensuring consistency between input and output spaces.
- **Input Condition Encoding:** Positional encoding, typical for sequential tasks, was replaced by directly concatenating input conditions (e.g., solar angles, aim point vectors) with latent representations. This ensures that input conditions guide the processing and weighting of calibration data, aligning the architecture with the spatial nature of the task.
- **Refinement of Attention Mechanisms:** The encoder retains self-attention to contextualize calibration data and identify relevant features. The decoder uses cross-attention to directly link input conditions with calibration data, while self-attention is excluded to avoid unnecessary dependencies between predicted points, reflecting the independence of focal spot predictions for different conditions.

### A.2. Training procedure

The Transformer architecture processes calibration data in an encoder–decoder framework. Training involves dividing a heliostat's focal spot images into two subsets: one for encoding and another for loss-based evaluation. The procedure is as follows:

1. The encoder processes calibration focal spots along with their corresponding conditions ( $\vec{A}_p, \vec{S}$ ) of one heliostat, generating a latent representation that encapsulates heliostat-specific properties.
2. The decoder utilizes this latent representation in combination with the desired input conditions for prediction (e.g., sun position, aim point) to generate the corresponding focal spot prediction.
3. A loss function is applied to the predicted flux distributions, refining both the Transformer and VAE components to minimize spatial discrepancies and enhance predictive accuracy.

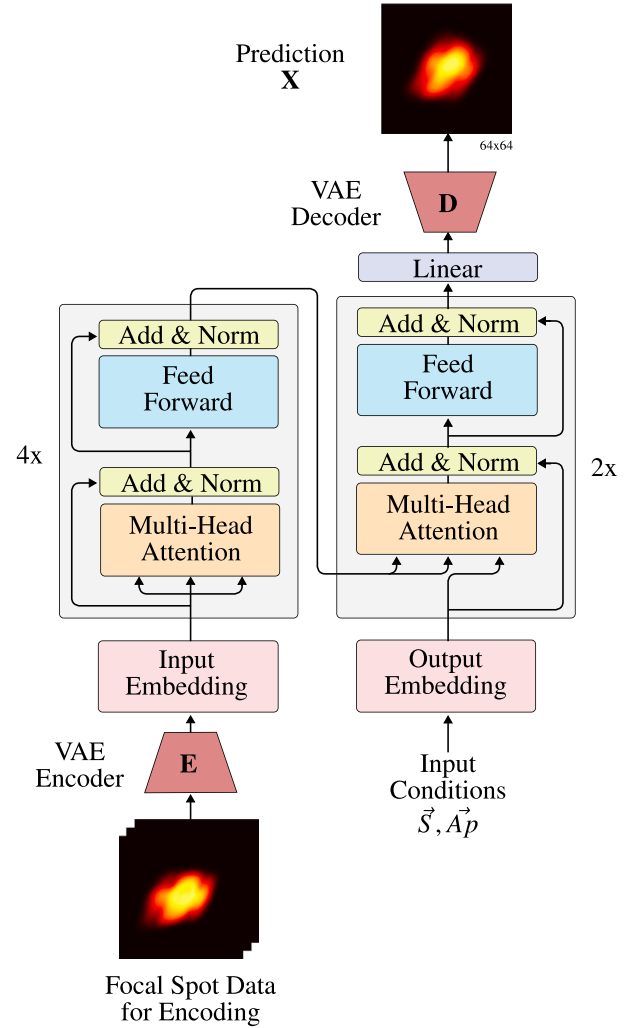


Fig. A.10. Overview of the modified Transformer architecture integrated with a Variational Autoencoder (VAE). The VAE preprocesses focal spot images into latent representations, which are then processed by the Transformer for flux prediction. This architecture enables robust and accurate predictions by leveraging advanced attention mechanisms and task-specific adaptations.

This structured training approach ensures that the model effectively captures heliostat-specific variations while maintaining the flexibility to generalize across different sun positions and aim points.

**Regularization and Model Complexity.** To reduce the risk of overfitting, we incorporate a dropout rate of 0.1 across the Transformer layers and VAE components. Additionally, we use a reduced architecture with only four encoder and two decoder blocks, and a hidden layer dimensionality of 64, which significantly lowers the number of trainable parameters compared to standard Transformer configurations. The model is trained with separate train/test splits per heliostat to ensure generalization across different configurations. No data augmentation was applied, as the model proved robust with the available data volume and structure.

### Data availability

All data used in this work, including focal spot images and the corresponding metadata, are available through the **PAINT Database** [28].

The database can be accessed at <https://paint-database.org/> and is permanently archived under PID: <https://hdl.handle.net/21.11152/474a4b1c-de93-4d4a-b33d-1d32d63baf4b>.

## References

- [1] F. Schöniger, R. Thonig, G. Resch, J. Lilliestam, Making the sun shine at night: comparing the cost of dispatchable concentrating solar power and photovoltaics with storage, *Energy Sources, Part B: Econ. Plan. Policy* 16 (1) (2021) 55–74.
- [2] R. Schäppi, D. Rutz, F. Dähler, A. Muroyama, P. Haueter, J. Lilliestam, A. Patt, P. Furler, A. Steinfeld, Drop-in fuels from sunlight and air, *Nature* 601 (7891) (2022) 63–68.
- [3] F.J. Collado, One-point fitting of the flux density produced by a heliostat, *Sol. Energy* 84 (4) (2010) 673–684.
- [4] M.d.l.R. Rodríguez-Sánchez, C. Leray, A. Toutant, A. Ferriere, G. Olalde, Development of a new method to estimate the incident solar flux on central receivers from deteriorated heliostats, *Renew. Energy* 130 (2019) 182–190.
- [5] A. Sánchez-González, B. Grange, C. Caliot, Computation of canting errors in heliostats by flux map fitting: experimental assessment, *Opt. Express* 28 (26) (2020) 39868–39889.
- [6] M. Pargmann, J. Ebert, D.M. Quinto, et al., Automatic heliostat learning for in-situ concentrating solar power plant optimization with differentiable ray tracing, *Res. Sq.* (2024) URL: <https://doi.org/10.21203/rs.3.rs-2554998/v2>, Preprint (Version 2), under review in *Nature Communications*.
- [7] M. Kuhl, M. Pargmann, M. Cherti, J. Jitsev, D. Maldonado Quinto, R. Pitz-Paal, Flux density distribution forecasting in concentrated solar tower plants: A data-driven approach, *Sol. Energy* 282 (2024) 112894, <http://dx.doi.org/10.1016/j.solener.2024.112894>.
- [8] K.W. Stone, Automatic heliostat track alignment method, 1986, US Patent 4, 564, 275.
- [9] 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.
- [10] E. Smith, C. Ho, Field demonstration of an automated heliostat tracking correction method, *Energy Procedia* 49 (2014) 2201–2210.
- [11] L. Guanyu, C. Zhongkun, Heliostat attitude angle detection method based on BP neural network, in: *MATEC Web of Conferences*, vol. 139, EDP Sciences, 2017, p. 00043.
- [12] M. Pargmann, M. Leibauer, V. Nettelroth, D. Maldonado Quinto, R. Pitz-Paal, Enhancing heliostat calibration on low data by fusing robotic rigid body kinematics with neural networks, *Sol. Energy* 264 (2023) 111962.
- [13] P. Schwarzbözl, R. Pitz-Paal, M. Schmitz, Visual HFLCAL-A software tool for layout and optimisation of heliostat fields, in: *Proceedings*, 2009.
- [14] M. Astolfi, M. Binotti, S. Mazzola, L. Zanellato, G. Manzolini, Heliostat aiming point optimization for external tower receiver, *Sol. Energy* 157 (2017) 1114–1129.
- [15] B. Belhomme, R. Pitz-Paal, P. Schwarzbözl, S. Ulmer, A new fast ray tracing tool for high-precision simulation of heliostat fields, 2009.
- [16] J. Rocca, B. Piaud, C. Coustet, C. Caliot, E. Guillot, G. Flamant, J. Delatorre, SOLFAST, a ray-tracing Monte-Carlo software for solar concentrating facilities, *J. Phys.Conf. Ser.* 369 (1) (2012) 012029.
- [17] T. Wendelin, A. Dobos, A. Lewandowski, Soltrace: a Ray-Tracing Code for Complex Solar Optical Systems, Technical Report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 2013.
- [18] 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.
- [19] A. Bonanos, M. Faka, D. Abate, S. Hermon, M. Blanco, Heliostat surface shape characterization for accurate flux prediction, *Renew. Energy* 142 (2019) 30–40.
- [20] M. Kiera, W. Schiel, Measurement and analysis of heliostat images, 1989.
- [21] A. Sánchez-González, C. Caliot, A. Ferrière, D. Santana, Determination of heliostat canting errors via deterministic optimization, *Sol. Energy* 150 (2017) 136–146.
- [22] A. Martínez-Hernández, R. Conceição, C.-A. Asselineau, M. Romero, J. González-Aguilar, Advanced surface reconstruction method for solar reflective concentrators by flux mapping, *Sol. Energy* 266 (2023) 112162.
- [23] M. Kuhl, M. Pargmann, M. Cherti, J. Jitsev, D. Maldonado Quinto, R. Pitz-Paal, In-situ unet-based heliostat beam characterization method for precise flux calculation using the camera-target method, *Res. Sq.* (2024) URL: <https://doi.org/10.21203/rs.3.rs-3978295/v1>, Preprint (Version 1), under review in *Solar Energy*.
- [24] M. Röger, P. Herrmann, S. Ulmer, M. Ebert, C. Prah, F. Göhring, Techniques to measure solar flux density distribution on large-scale receivers, *J. Sol. Energy Eng.* 136 (3) (2014) 031013.
- [25] S. Pomp, P. Schwarzbözl, G. Koll, F. Göhring, T. Hartz, M. Schmitz, B. Hoffschmidt, The solar tower jülich – first operational experiences and test results, in: *Proceedings of the SolarPACES 2010 Conference*, Perpignan, France, 2010.
- [26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, *Adv. Neural Inf. Process. Syst.* 30 (2017).
- [27] D.P. Kingma, M. Welling, Auto-encoding variational Bayes, 2013, arXiv preprint arXiv:1312.6114.
- [28] K. Phipps, M. Kuhl, M. Weiel, M. Busch, J. Lewen, N. Blumenröhr, D.M. Quinto, C. Debus, F. Göhring, A. Streit, R. Pitz-Paal, M. Götz, M. Pargmann, PAINT database, 2025, <https://paint-database.org/>, PID: <https://hdl.handle.net/21.11152/474a4b1c-de93-4d4a-b33d-1d32d63baf4b>.