Energy and AI 24 (2026) 100763

Contents lists available at ScienceDirect

Energy and AI

journal homepage: www.sciencedirect.com/journal/energy-and-ai

A data-driven digital twin approach for real-time flux density prediction in solar power towers using ray-level deep learning corrections

Sergio Díaz-Alonso *, Christian Raeder, Allen Charly, Bernhard Hoffschmidt

German Aerospace Center (DLR), Institute of Solar Research, Im Langenbroich 13, 52428 Jülich, Germany

HIGHLIGHTS

- Ray directions compressed 99.99 % enabling near-real-time inference at 4.6 ms.
- Power and direction CNN corrections validated at 80.9 % cavity flux accuracy.
- 5-fold cross-validation confirms robustness with fold std below ±0.005.
- 13 × 13 grid resolution verified as optimal via Kolmogorov-Smirnov sensitivity.
- Aggregated training across 931 conditions generalizes without heliostat calibration.

GRAPHICAL ABSTRACT

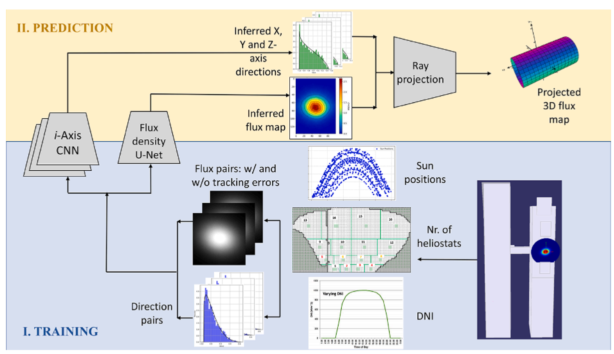

ARTICLE INFO

Keywords:
Concentrating solar technologies (CST)
Deep learning
Convolutional neural networks (CNN)
Flux density predictions
Real-time simulation

ABSTRACT

Concentrating solar technologies offer significant potential for sustainable and dispatchable heat, power and fuel production, especially through solar power tower systems. However, accurate flux density measurement at the receiver remains a critical barrier to the commercial deployment of this technology. Current direct and indirect methods are disruptive, expensive and inapplicable to cavity receivers. In order to overcome these challenges, this work proposes a computational data-driven method that couples four convolutional neural networks operating sequentially with a Monte Carlo ray-tracing simulator to correct both the power and directional components of simulated rays at the receiver aperture plane. Ray direction data are compressed from over 100×10^6 parameters to 2028 values through statistical distribution fitting, enabling near-real-time inference on standard computing hardware. The power correction achieves accuracies up to 95.0 % at the aperture plane, while direction predictions exceed 80.0 % Pearson correlation for the three ray direction components across 63 diverse meteorological conditions spanning direct normal irradiance values from 100.0 to 900.0 W/m^2. The corrected rays are subsequently projected into the three-dimensional cavity receiver geometry, yielding an integrated flux prediction accuracy of 80.9 % compared to experimental camera-based measurements at the Solarturm Jülich facility. An aggregated training strategy across 931 meteorological conditions and multiple heliostat configurations ensures generalization without cost-extensive per-heliostat calibration. These results demonstrate that ray-level deep learning correction provides a universal, non-disruptive, and industrially scalable solution for flux density prediction in both open and cavity solar receivers.

* Corresponding author.
E-mail address: sergio.diazalonso@dlr.de (S. Díaz-Alonso).

https://doi.org/10.1016/j.egyai.2026.100763

Available online 23 April 2026

1. Introduction

Concentrating solar technologies (CST) play a key role in the energy transition to renewable sources. These systems stand out due to their versatile outputs, including low and high-temperature heat as detailed in the thermal analysis by [1], electric power generation [2], and the production of solar fuels through thermochemical cycles [3]. Furthermore, they enable dispatchable power production [4], especially when integrated with thermal energy storage systems [5]. Recent studies have also demonstrated their effectiveness when hybridized with photovoltaic technologies [6]. Among the resulting designs developed over the last 50 years – from Fresnel collectors to parabolic trough systems – solar power towers (SPT) show especially promising features, including high solar concentration at the receiver aimpoints and lower O&M costs in the heliostat field.

Nevertheless, solar power towers also present significant challenges related to their intrinsic complexity. For example, flux density distribution at the receiver is essential and critical to reach high operational efficiency, safety performance, and long-term durability of concentrating solar power (CSP) plants. This distribution, representing the incoming solar energy along the receiver surface, remains an open challenge in SPT development. Operation becomes increasingly difficult when weather is irregular or when the irradiated surface is geometrically complex, such as cavity or curved receivers – precisely the configurations that offer the highest thermal performance. This fact underlines the necessity of a universal flux density prediction method applicable to diverse geometries.

Further problems arise when the measurement methods must scale to industrial plants with thousands of heliostats and strict economic targets. Existing methods rely heavily on precise heliostat aiming, which is highly sensitive to industrial-level inaccuracies. Camera-target measurement campaigns require operational disruptions, which directly reduce plant revenue since output is proportional to duty time. No current method combines non-disruptive operation, applicability to cavity receiver geometries, and scalability to commercial heliostat fields — a gap that motivates the present work.

In response to these challenges, usually related to measurement methods, several prediction and simulation virtual environments have been deployed. The classic codes from 70 s (HELIOS, MIRVAL, or DELSOL) were usable for small and simple facilities and they were based on cone optics and convolution methods in order to predict focal spot characteristics. However, as noted in the comparative study of optical simulation tools by García et al. [7], inexperienced users may face high computational costs due to the division of heliostat facets into numerous differential elements. Subsequent software developments have focused on Monte-Carlo ray-tracing to improve efficiency, including SOLTRACE tool [8], the open-source TONATIUH platform [9] or the computationally efficient code STRAL [10]. These codes restrict ray-tracing in complex surfaces, which presents a new problem for the new generation of receivers.

Addressing these challenges, this work presents a purely data-driven method for universal flux density predictions based on near-real-time AI-enhanced simulations. Unlike, existing approaches that operate on flux maps at the individual heliostat level, the proposed method corrects both the power (X) and directional components (d_x, d_y, d_z) of simulated rays at the aperture plane through four parallel convolutional neural networks (CNN). Then, it projects the corrected rays towards any receiver geometry – including cavity receivers, not accessible to cameras. The key contributions of this work are threefold. First, a ray-level parametric reduction compresses over 100×10^6 directional parameters into 2028 values through statistical distribution fitting, making real-time correction computationally feasible on standard hardware. Second, an aggregated training strategy generalizes across 931 diverse meteorological conditions and 17 heliostat configurations without per-heliostat calibration, overcoming a fundamental scalability limitation of existing approaches. Third, aperture-plane correction accuracy is transferred into three-dimensional cavity receiver geometries where no direct measurement method currently exists, filling a critical gap in the state of the art. The complete process is depicted in Fig. 1.

2. State of the art

2.1. Flux density measurement

Flux density measurement (FDM) methods are mostly divided into direct and indirect techniques, and these methodologies are usually supported by ray-tracing simulations. Röger et al. provide a comprehensive division in ref [11] and scarce advances have been introduced since then in this field. In general terms, the FDM methods included in these classes are designed to obtain the flux distribution at the receiver aperture plane.

On the one hand, direct methods rely on the usage of flux gauges to determine the exact superposed flux in different locations and interpolated values between those points. Different configurations have been tested throughout the years. For example, Osuna et al. [12] demonstrated the use of randomly distributed static radiometers scattered across the receiver area. King and Arvizu [13] subsequently tested structured arrays of static flux gauges. Scanning configurations have also been investigated using two distinct approaches. In the first, sensors are deployed on moving bars or surfaces that sweep across the irradiated area while the solar beam remains stationary, as demonstrated by Ballestrín and Monterreal [14] using a hybrid radiometric scanning system, by Elsayed et al. [15] for flux mapping on a plane receiver, and by Yogev et al. [16] through a high-heat solar flux scanner. In the second approach, Pacheco et al. [17] positioned sensors on fixed auxiliary structures located outside the receiver aperture, where the solar beam is deflected perpendicularly across the stationary bar to achieve spatial scanning without mechanical movement of the sensors themselves.

On the other hand, indirect methods employ cameras to detect the reflected irradiation from the central receiver, as demonstrated by Offergeld et al. [18], or from auxiliary Lambertian targets. In the latter case, the target configuration can be fixed, such as the HERMES I and II systems described by von Tobel et al. [19] for high resolution analysis, or mobile, such as the FMAS presented by Thelen et al. [20] and the PROHERMES system described by Ulmer et al. [21]. In general, these camera-based methods also require a complementary radiometer to scale the relative pixel intensity distributions into heat flux values, as Rodríguez-Sanchez et al. [22] noted.

Direct methods are limited by the complexity of water-cooling circuits, flux gauge operation and maintenance costs, and the capacity to interpolate between radiometer locations. The SOTA indirect methods present other limitations regarding their general use in industrial-scale plants. Operation disruptions, target reflectivity degradation, and dependency on receiver geometry especially handicap these methods. Despite decades of development, no existing experimental flux measurement method combines non-disruptive operation, applicability to cavity receiver geometries, and scalability to commercial heliostat fields of several thousand mirrors — a gap that motivates the virtual prediction approach presented in this work.

2.2. Heliostat beams predictions

Another predominant approach is the superposition of individual heliostats' focal spots. Most representative examples found in the literature rely on ray-tracing simulation techniques, as reviewed by Garcia et al. [7]. Among these computing techniques, it is possible to find Convolution Methods and Monte Carlo ray-tracing. The first type involves the superposition and convolution of optical error cones from all heliostats, represented by their mathematical simplified functions [23]. In the second case, the Monte Carlo method is used to generate pseudo-random ray trajectories in order to solve the deterministic problem of heat flux density simulation. This probabilistic approach has been widely adopted in the field of CSP: Belhomme et al. implemented it in the STRAL code [10], Roccia et al. developed SOLFAST ray-tracing software [24], Wendelin leveraged it for SolTrace development [25], and Leary and Hankins utilized it for MIRVAL code [26]. However, these codes usually require surface simplifications or expensive and time-consuming additional measurements like deflectometry [27] or photogrammetry [28].

In order to address purely computing methods' lack of robustness, a new strategy based on measuring target images was developed by Collado [29]. It is theoretically possible to reconstruct heliostat surfaces (with their respective errors) from these individual fluxes, as proposed in the flux reconstruction models developed by Sánchez-González et al. [30] and in the surface error estimation methods by Zhu et al. [31]. However, these techniques still relied on simplistic heliostat models and the most complex distortions within the focal spots were not captured, although canting errors were fitted to the images. All of these methods, as well as the technique introduced by Rodríguez-Sánchez et al. [32] are time-consuming and usually dependent on climatic conditions. It is rather impractical to incorporate these spot scanning routines in real power-plant operation and, if possible, the accuracy would have to be questioned.

Considering the shortcomings of the presented methods, analyzing superposed fluxes can reduce prediction times drastically, avoid operation disruptions, and enhance cost competitiveness. Furthermore, a holistic approach helps to neglect the individual effects of each heliostat in the superposed flux. In practical operation, a single heliostat facet canting or slope error effect is negligible compared to the superposed values of flux at the receiver. However, computing the slope and canting of all heliostats individually requires high computing power and complex measurement methodologies. While Monte Carlo ray-tracing methods offer geometric flexibility, they remain computationally intensive and dependent on accurate mathematical models of optical errors that require expensive calibration campaigns and do not capture transient atmospheric phenomena during actual flux measurement. These constraints motivate data-driven correction approaches that learn realistic ray behaviour directly from measurements.

2.3. Deep learning methods for flux predictions

The increasing challenges in concentrating solar power operation have motivated the application of deep learning to flux prediction and heliostat characterization tasks. These novel techniques typically leverage neural networks trained on ground truths acquired through the camera-target method [33], which is often part of the standard heliostat calibration routines and therefore represents a non-disruptive and cost-efficient source of data. The resulting methods broadly fall into two categories: image-to-image translation approaches that directly infer real-world flux distributions from ray-tracing simulations, and surface reconstruction approaches that first recover heliostat geometry before computing flux through ray-tracing.

In the first category, M. Kuhl et al. [34] proposed an image-to-image methodology based on the usage of convolutional neural networks (U-Net) to directly infer the individual heliostat flux distributions, given a simulated heliostat flux map and a background target image captured under operational conditions. This approach provides non-disruptive flux estimates but requires individual training data for each heliostat, limiting its scalability to large commercial fields with thousands of mirrors. As an extension, M. Kuhl et al. [35] proposed a generative adversarial networks (GANs) framework to forecast flux distributions from heliostat characteristics and flux data, reducing the dependence on real-time target images. While this approach improves scalability, it remains sensitive to variable atmospheric conditions and cloud covers.

In the second category, Pargmann et al. [36] introduced a differentiable ray-tracing framework that backpropagates irradiance gradients to recover high resolution surface and mechanical representations of individual heliostats through NURBS, enabling automatic in-situ metrology without dedicated calibration campaigns. Lewen et al. [37] complemented this by combining deflectometric measurements with inverse deep learning ray tracing to reconstruct virtual heliostat surface models. Both approaches represent significant methodological advances but require the costly superposition of individual heliostat flux maps and assume full visibility of each heliostat light spot on the measurement target.

All existing deep learning approaches to solar flux prediction share

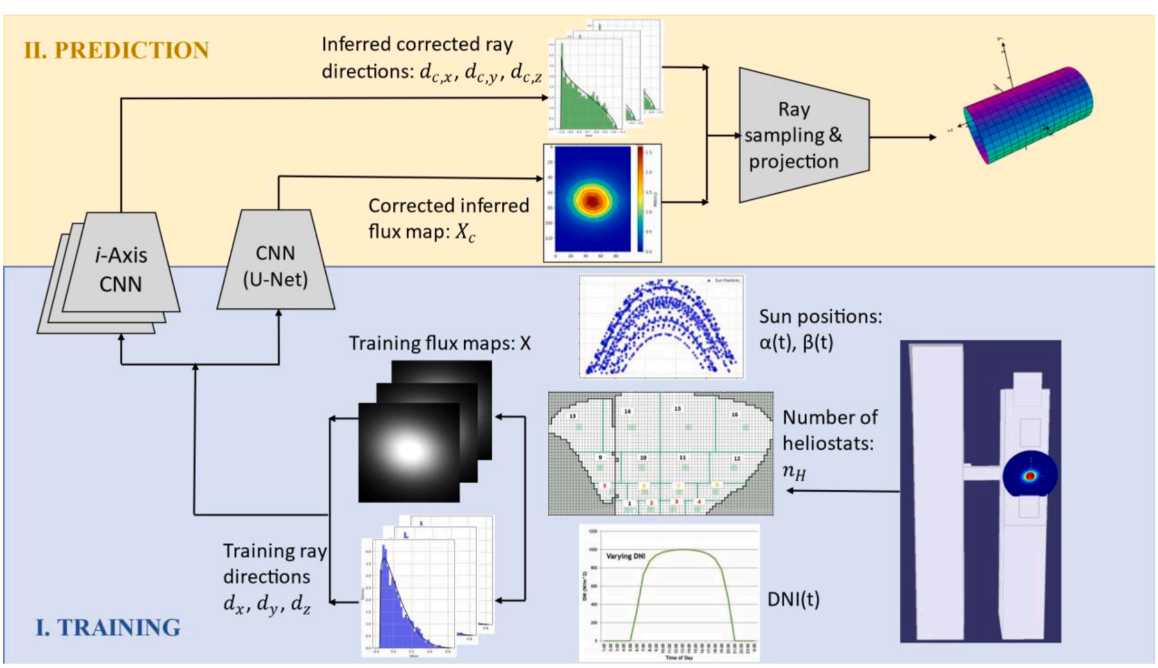

Fig. 1. Training and inference pipeline of the ray-level data-driven correction method. Ideal Monte Carlo ray-tracing inputs (sun position, DNI, heliostat count) feed four parallel CNNs that correct ray power (U-Net) and directional components (X, Y, Z axes). Corrected rays are subsequently projected onto the three-dimensional receiver geometry to produce the final flux map.

three fundamental limitations that the present work addresses. First, they are restricted to flat or near-flat receiver geometries where the heliostat light spot is fully visible on the measurement target, excluding cavity receivers where direct optical access is unavailable. Second, they require either per-heliostat training data or full light spot visibility, both of which impose practical constraints on industrial deployment at the scale of commercial plants. Third, none operate at the ray level — corrections are applied to integrated flux maps rather than to the underlying ray populations, preventing the projection of corrections into complex three-dimensional receiver geometries. This work addresses all three limitations through a ray-level correction framework that operates at the aperture plane and projects corrected rays into arbitrary receiver geometries, trained on aggregated flux data across diverse meteorological conditions without per-heliostat calibration requirements.

2.4. Relevant machine learning advances, architectural foundations and parameter reduction strategies

Beyond the solar energy domain, the broader machine learning literature offers architectural innovations and optimization strategies that inform the design choices made in this work. Several multimodal feature fusion networks have demonstrated that combining complementary information sources through encoder-decoder architectures substantially improves structured prediction accuracy. Zhou et al. [38] introduced MFFENet, combining RGB and thermal modalities through multi-scale feature fusion with spatial attention for improved semantic segmentation. Dong et al. [39] proposed EGFNet, embedding prior edge maps into cross-modal feature representations with multitask deep supervision for RGB-thermal scene parsing. PGDENet [40] addresses weak modality representations through progressive top-down guided fusion with channel and spatial correlation enhancement. ECFFNet [41] introduces bilateral foreground-background reversal fusion with multi-level feature aggregation for salient object detection under low-light conditions. Wang et al. [42] presented LSNet, combining large-kernel static perception with small-kernel dynamic aggregation inspired by the peripheral-foveal mechanism of human vision. MMSMCNet [43] fuses modal memory sharing with a three-stage morphological decoding unit extracting complementary contour, skeleton, and semantic features. Zhou et al. [44] constructed IRFR-Net, integrating gated attention fusion with weighted atrous spatial pyramid pooling for multi-scale RGB-D feature extraction. Zhou et al. [45] proposed WaveNet, applying discrete wavelet transform-based cross-modality fusion with knowledge distillation from a transformer teacher to a lightweight Wave-MLP student, achieving state-of-the-art RGB-thermal detection accuracy at reduced computational cost. CCAFNet [46] combines RGB and depth modalities through cross-modal attention for safety-critical scene understanding.

Optimization-enhanced deep learning has also shown consistent benefits in energy prediction tasks. El-Kenawy et al. [47] proposed an adaptive grey wolf-dipper throated optimization algorithm for wind speed forecasting, achieving a root mean square error of 0.0035 on the Global Energy Forecasting Competition dataset. Takieldeen et al. [48] separately introduced the Dipper Throated Optimization algorithm for feature selection, outperforming particle swarm optimization, whale optimization, grey wolf optimizer, and genetic algorithm on standard benchmarks. Khafaga et al. [49] applied Al-Biruni Earth radius optimization to fine-tune deep convolutional networks for monkeypox classification, and Alhussan et al. [50] proposed a binary waterwheel plant optimization algorithm demonstrating that metaheuristic search substantially reduces computational complexity in high-dimensional datasets' feature selection. Finally, [51] proposed a Greylag Goose Optimization-enhanced hybrid model for smart city electricity load forecasting, achieving a mean squared error of 0.0021 and an R^2 of 0.9999.

While none of these approaches are directly applicable to ray-level solar flux correction, their core principles of multi-scale feature aggregation, spatial attention, cross-modal complementarity, and optimization-driven parameter reduction informed the architectural choices and parametric reduction strategy proposed in this work — specifically the compression of 103.2×10^6 ray direction parameters into 2028 statistically representative values that preserves predictive accuracy while enabling near-real-time inference.

2.5. Knowledge gaps and positioning of this work

Table 1 summarizes the capabilities and limitations of the most closely related deep learning methods for solar flux prediction in comparison with the approach proposed in this work. Three specific gaps motivate the present contribution: the absence of ray-level correction methods in the existing literature, the lack of flux prediction capability for cavity and complex receiver geometries, and the dependence of existing data-driven methods on per-heliostat training data or disruptive measurement campaigns.

This comparison establishes the proposed method as the only published approach integrating all five functional capabilities within a single framework. Its novelty is therefore systemic and task-centric rather than strictly architectural: the neural components utilize established building blocks selected for their empirical reliability, integrated into a physically-informed pipeline that solves a high-dimensional problem no existing method addresses.

Reported accuracy figures are task-specific and not directly comparable across literature due to heterogeneous receiver geometries and distinct prediction targets (e.g., individual vs. superposed flux, aperture-plane vs. 3D receiver-surface). Nevertheless, this model demonstrates superior performance in comparable metrics and uniquely accommodates universal geometries. In terms of planar flux, it exceeds the precision reported in [35] and [37]. Furthermore, as Lewen et al. outperformed Pargmann et al., our methodology—which surpasses Lewen et al.—effectively achieves a new performance benchmark over [36]. While the accuracy from Kuhl et al. [34] is significant, its application framework remains restricted to flat targets, lacking the 3D cavity projection and real-time flexibility central to this work.

3. Methodology

This work generalizes the methodology presented in ref [52], originally developed for flat surfaces. The objective is to accurately predict flux distribution across any receiver geometry by employing different data-driven models that efficiently correct ray generation in Monte-Carlo simulations. To achieve this broader application, the authors propose a discrete correction based on the systematic comparison of directions and power levels between ideal ray-tracing simulations and equivalent results from data-driven models trained on realistic measurements/simulations. These comparisons are performed at the receiver aperture plane to leverage SOTA methods for generating accurate ground truths. Subsequently, the corrected rays are projected toward any desired target geometry. This procedure enables the transfer of flat-surface prediction accuracy to complex geometries, while facilitating performance monitoring without the need for internal cavity sensors, where harsh conditions typically demand intensive O&M.

3.1. Ray-level correction

The correction module is sequentially coupled with the legacy ray-tracer, utilizing ideal simulations as inputs. In this context, "ideal" refers to a simulation devoid of mathematical model, such as slope errors or circumsolar ratio. The only realistic features included in the ideal environment are the actual heliostat surfaces, which were explicitly characterized via deflectometry [27]. The neural network responsible for the power correction (U-Net+AG) receives a flux map from the first component. Similarly, the neural networks tasked with correcting X, Y and Z ray directions are provided with a ray file containing the position, direction, and power of each individual ray intercepted at the receiver

Table 1
Systematic capability comparison of published data-driven solar flux prediction methods, highlighting the three knowledge gaps addressed by this work.

Feature	Kuhl et al. [34]	Kuhl et al. [35]	Pargmann et al. [36]	Lewen et al. [37]	This work
Input flux	Individual heliostat	Individual heliostat	Individual heliostat	Yes	Aggregated
Heliostats per inference	1	1	1	1	All
Receiver geometry	Flat target	Flat target	Flat target	Flat target	Flat target + cavity
Training data required	Per-heliostat images	Per-heliostat + historical	Per-heliostat deflectometry	Per-heliostat deflectometry	Aggregated flux maps (flat)
Disruption	No	No	No	No	No
Near-real-time inference	No	No	No	No	Yes
Validated metric	Per-heliostat flux uncertainty	Flux distribution error	Heliostat surface error	Flux + heliostat surface error	Integrated power + flux (flat and cavity)
Reported accuracy	~4.5 %	89 %	2 mm	92 % (flux) + 0.14 mm	95 % (flat) + 80.9 % (cavity)

aperture plane. This algorithm is trained using superposed fluxes (i.e. from multiple heliostats), which necessitates the generation of an extensive number of rays to ensure final accuracy. Consequently, 34.4×10^6 rays were used for each flux prediction composing the datasets detailed in Section 3.2, as well as images of 256 × 256 px and 1 channel (grayscale). Considering that 15,928 images and 8379 ray files are processed during training, computational cost remains the primary constraint. Furthermore, the authors aim to implement a real-time correction based on the trained neural networks; thus, addressing the dimensional complexity is crucial. To this end, a dimensional reduction based on Region of Interest (ROI) discretization has been carried out, as depicted in Fig. 2.

To perform ray-level correction, each of the 34.4×10^6 ray interceptions is firstly localized within the ROI, allowing each ray to be assigned its corresponding correction features. The initial step in this dimensional reduction involves discretizing the potential ray locations – reducing them from the original 256 × 256 pixel resolution to a 13 × 13 grid consisting then of 169 elements. This discretization is applied to both subsequent correction stages, which are executed element-wise for all rays located within those units. While the grid can be increased to capture steeper flux gradients, it remains a sensitive parameter regarding computational overhead. The grid resolution sensitivity was assessed using the Kolmogorov-Smirnov D-statistic on 13.05.2022 at, which reaches global minima at 13 × 13 across all three axes simultaneously (X: 0.3026, Y: 0.0990, Z: 0.0758), representing improvements of 1.6 %, 14.9 %, and 2.2 % over the 7 × 7 configuration and 0.9 %, 1.4 %, and 2.4 % over the 10 × 10 configuration, confirming better statistical fittings and superior capture of steep directional gradients within the ROI. Additionally, the 13 × 13 grid remains 33 % and 40 % faster than the 17 × 17 and 20 × 20 configurations respectively, making it the optimal trade-off between spatial resolution, fitting accuracy, and computational feasibility.

3.1.1. Power correction

The digital rectification presented herein builds upon the deep-learning-based flux prediction for flat surfaces presented in ref [52]. When provided with an ideal simulation, the trained U-Net architecture generates a realistic flux distribution, that accounts for the mathematical models of tracking error and CSR. This algorithm processes the neural network's input – an ideal 256 × 256 px, single-channel flux map (F^i) – and its corresponding output, a realistic 256 × 256 px single-channel flux map (F^r). Subsequently, the average value for each grid elements is computed for both cases ($\overline{F^i}$ and $\overline{F^r}$), as defined in Eqs. (1)-(2).

$$\overline{F^i_{i,j}} = \frac{1}{|B_{i,j}|} \sum_{(m,n)\in B_{i,j}} F^i_{m,n} \tag{1}$$

$$\overline{F^r_{i,j}} = \frac{1}{|B_{i,j}|} \sum_{(m,n)\in B_{i,j}} F^r_{m,n} \tag{2}$$

With $\overline{F^i}$, $\overline{F^r} \in \mathbb{R}^{256\times256}$;$F^i$, $F^r \in \mathbb{R}^{13\times13}$; $B_{i,j} \subseteq \{1,\ldots\ldots,256\} \times \{1,\ldots\ldots,256\}$ the set of indices in the (i, j)-th element of the images F^i or

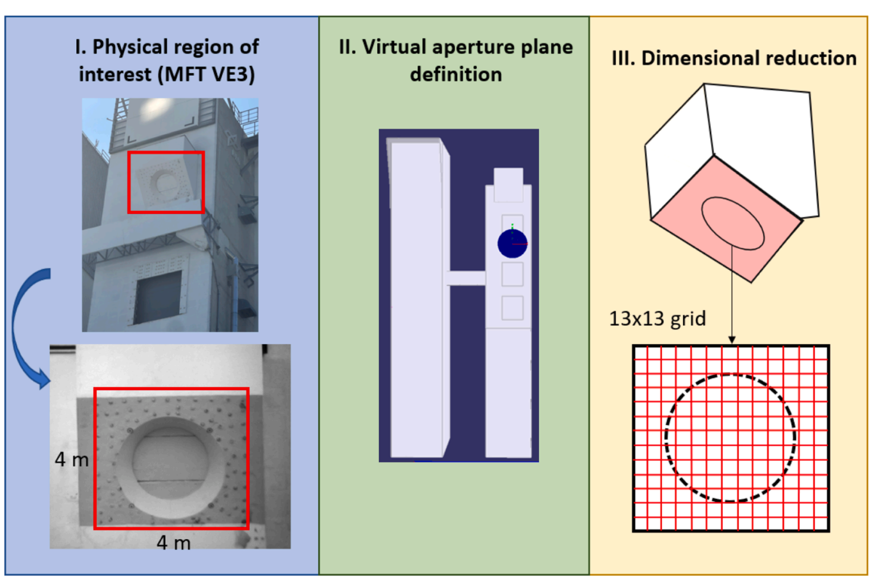

Fig. 2. Dimensional reduction procedure applied to the correction stages. Discretization process involving: (I) physical region of interest at the MFT VE3 aperture plane, (II) virtual aperture plane definition, and (III) reduction from 256 × 256 pixel resolution to a 13 × 13 spatial grid (169 elements).

F^r and $|B_{i,j}|$ the number of pixels in that element. When this downsampling is performed, the average values of both cases are divided element-wise, which gives a correction matrix composed of 13 × 13 correction factors (A^{corr}), as shown in Eq. (3).

$$A^{corr} = \frac{\overline{F^i_{i,j}}}{\overline{F^r_{i,j}}} \tag{3}$$

The subsequent stage involves classifying the 34.4 × 10^6 rays across the grid domain. During this process, the ray file (a table of 34.4 × 10^6 rows and 9 columns in ".txt" format) is parsed to extract the first two columns, which represent the horizontal and vertical coordinates of the ray interceptions. These coordinates are systemically compared against the grid boundaries; subsequently, the ray IDs (corresponding to their row indices) are sorted and stored in a nested list of 169 elements based on their spatial allocation. To optimize this classification, the authors divide the ray file into 10 chunks (3.4 × 10^6 rays each) to enable parallel processing. The simultaneous utilization of multiple cores reduced the computing time by three orders of magnitude on a high-performance workstation equipped with a 13th Gen Intel Core i9–13950HX processor (24 cores, 32 threads).

Once the spatial classification is complete, the 7th column of the ray file, that comprehends the power carried by each beam, is also extracted (p^{init}). A^{corr} is flattened to an array of 169 elements (α^{corr}) and the ray powers (identified by the element where the beam is located) are multiplied element-wise by its corresponding correction factor, as depicted in Eqs. (4)-(5). Thus, the result of this phase is an array of 34.4 × 10^6 elements with the tuned powers (p^{corr}).

$$\boldsymbol{\alpha}^{corr} = \boldsymbol{vec}(A^{corr}) \tag{4}$$

$$\boldsymbol{p}^{corr} = \left\{ \left\{ \boldsymbol{\alpha}_i^{corr} \cdot \boldsymbol{x}_{ij}^{init} \right\}_{j=1}^{N_i} \right\}_{i=1}^{169} \tag{5}$$

Where $\alpha^{corr} = \{\alpha_1^{corr}, \ldots\ldots, \alpha_{169}^{corr}\} \subset \mathbb{R}$ is the flattened correction array; $p^{init} = \{p_1^{init}, \ldots\ldots, p_{169}^{init}\}$ with $p_i^{init} = \left\{x_{i1}^{init}, \ldots\ldots, x_{iN_i}^{init}\right\} \subset \mathbb{R}$ is the original nested power array and p^{corr} is the analog corrected power array. The conceptual process is summarized in Fig. 3 below.

3.1.2. Direction correction

As previously noted, the authors identified the necessity of correcting not only the scalar power, but also the vector direction of the simulated rays using truthful sources. Indeed, these rays undergo directional deviations analogous to those observed in power distribution as a result of tracking errors and CSR mathematical modeling. In this case, the dimensional complexity is significantly higher due to the vectorial nature of the direction magnitude. Considering the initial ray classification into 169 grid elements, each of the three directional components for the 34.4 × 10^6 rays (totaling 103.2 × 10^6 parameters) must be adjusted using a specific triplet of correction factors selected from the 169 available sets.

This purely data-driven method is designed for near-real-time execution on commercial laptops; therefore, the parametric complexity must be mitigated through an additional reduction process performed prior to computation.

3.1.2.1. Parametric reduction. This mathematical procedure is based on the parallel statistical distribution fitting of ray directions within each of the grid elements. Following a visual inspection of representative samples, it was observed that ray directions consistently exhibit similar patterns across different elements within the ROI and across various timestamps. Consequently, the X axis data is fitted with a double step piecewise density function (Eq. (6)):

$$f_x(x) \begin{cases} 0 \text{ if } x < x_1 \\ a_1 \text{ if } x_1 < x < x_2 \\ a_2 \text{ if } x > x_2 \end{cases} \tag{6}$$

In the equation above, x_1 and x_2 represent the thresholds where the two steps are located, respectively. Moreover, a_1 and a_2 denote the amplitude of these steps. Meanwhile, the Y and Z axes are characterized by generalized Beta distributions (Eqs. (7)-(8)):

$$f_y(x;\alpha,\beta,a,s) = \frac{1}{s \cdot B(\alpha,\beta)} \left(\frac{x - - a}{s}\right)^{\alpha-1} \cdot \left(1 - \frac{x - - a}{s}\right)^{\beta-1} \sim f_z(x;\alpha,\beta,a,s) \tag{7}$$

Where α, $\beta > 0$ are the shape parameters, a is the location parameter, $s > 0$ is the scale parameter and $B(\alpha, \beta)$ is the standard Beta function:

$$\boldsymbol{B}(\boldsymbol{\alpha},\boldsymbol{\beta}) = \int_0^1 \boldsymbol{t}^{\alpha-1}(1-\boldsymbol{t})^{\beta-1} \boldsymbol{dt} \tag{8}$$

These fittings, compared to their corresponding histograms for three scattered elements along the mentioned ROI, are provided in Fig. 4 at a specific timestamp.

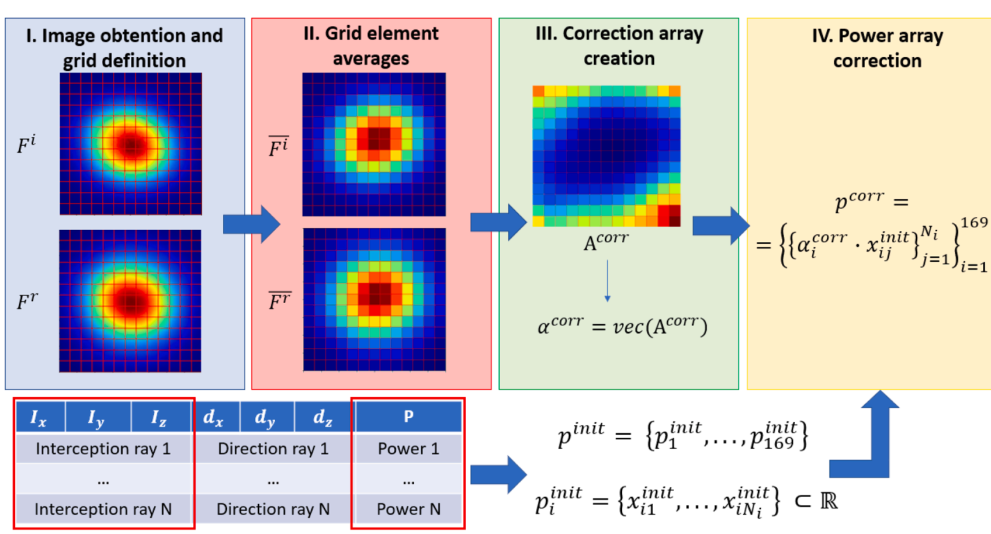

Fig. 3. Power correction complete flow diagram and ray-scaling logic. (I) U-Net generates a realistic flux map $\overline{F_r}$ from an ideal input $\overline{F_i}$, (II) element-wise averages are computed on the 13 × 13 grid, (III) the correction matrix $\boldsymbol{A}^{corr}$ is derived as the element-wise ratio, (IV) individual ray powers ($\boldsymbol{p}^{corr}$) are scaled by their correction vector.

This mathematical procedure is performed for each of the 169 grid elements across the three axes, covering both the inputs and the ground truth labels for the machine learning algorithms detailed hereinafter. As defined in the governing equations, each density function is characterized by four parameters. Consequently, the total number of parameters involved in the directional correction is reduced from 103.2×10^6 to just 2028 (4 parameters per 3 distributions per 169 mesh elements) – a reduction of > 99.9 %. Thus, the machine learning models are trained to correct the distribution parameters rather than individual ray directional components.

The result of this approach, combined with grid-based downsampling, not only ensures the feasibility of training an AI model – which would otherwise be unattainable given the computational constraints of processing 100×10^6 input parameters – but also enables near-real-time correction using the trained architecture.

3.1.2.2. Machine-learning models for direction prediction. Machine learning models for multi-axis direction distributions prediction are designed to incorporate the effects of tracking errors into directional, analogous to the power correction stage. These models process the vertically stacked statistical data mentioned in Section 3.1.2.1. to predict the realistic ray directions from ideal Monte-Carlo ray-tracing simulations. Prediction accuracy varies across axes due to the presence of outlying values or irregular spatial series particularly within Y-axis.

Data preprocessing involves comprehensive smoothing methods tailored to each parameter: "Alpha" and "x_1" parameters utilize LOWESS

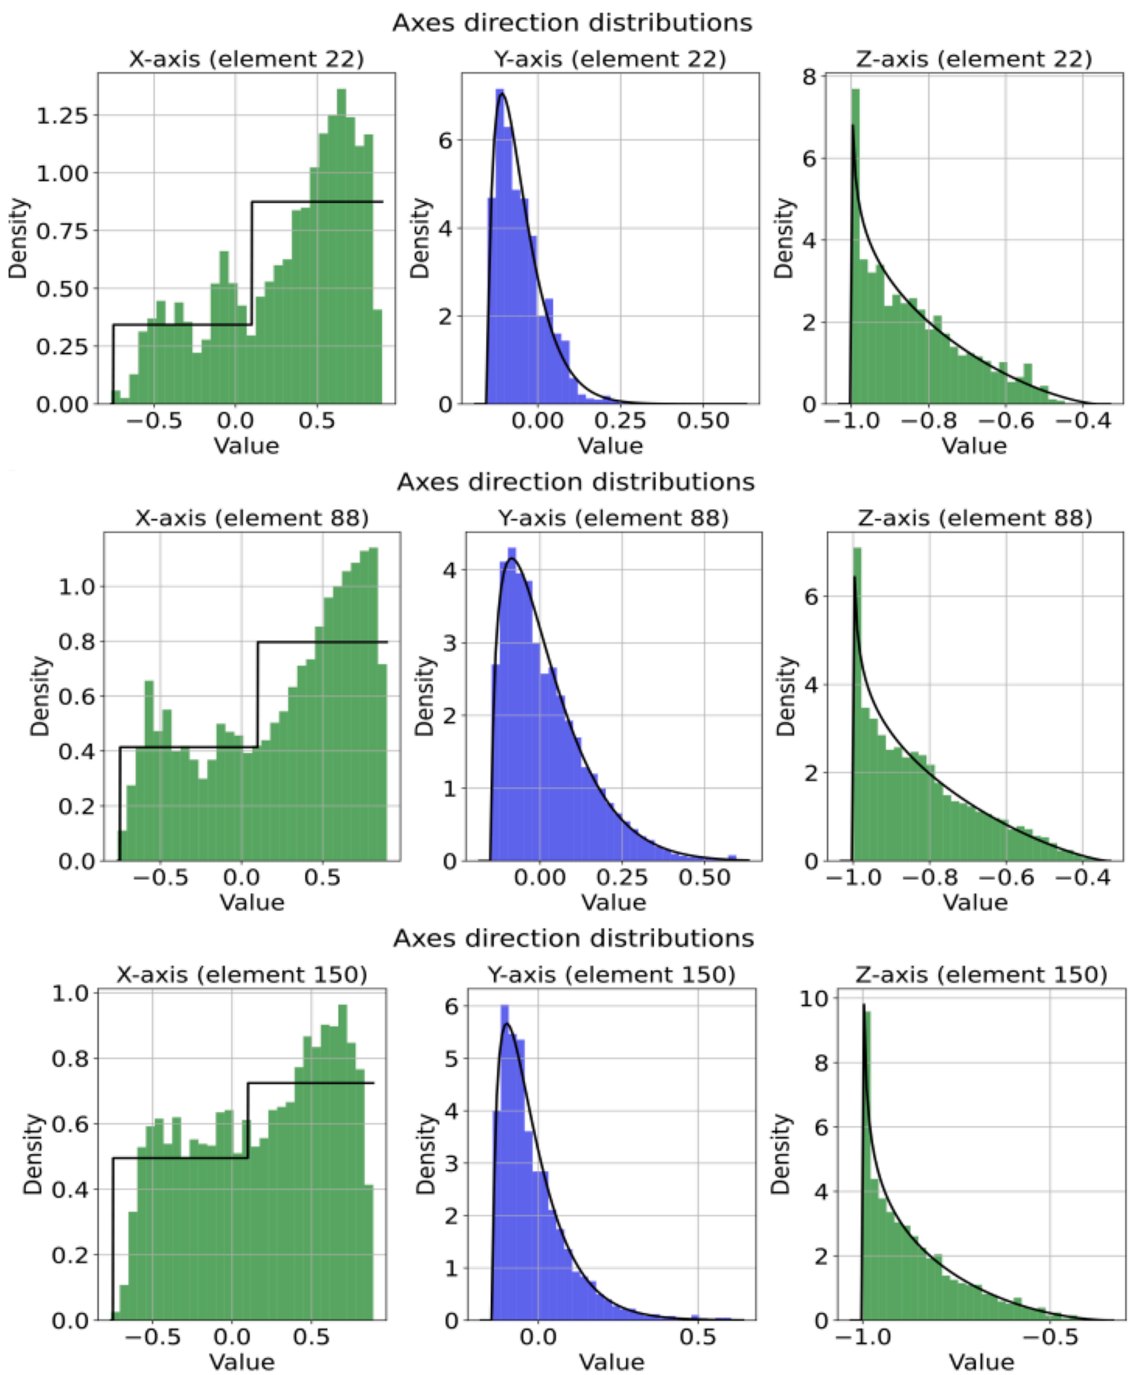

Fig. 4. Statistical distribution fitting for ray directional components. Comparison between original Monte Carlo data and the fitted parametric PDF: double step function for the azimuthal axis (X) and generalized Beta distributions for Y and Z-axes (13.05.2022 at 11:30 AM).

with Gaussian Kernel components for 1D serial pattern conservation [53]. “Beta” and “x_2” parameters use an adaptive Gaussian smoothing method with bilateral filtering [54]. “Location” and “a_1″ parameters apply spline-based method with edge preservation [55], while “State” and “a_2” parameters employ a two-step cleaning processing [56] combined with extreme threshold removal and exponential moving average (EMA) smoothing [57]. These methods effectively mitigate parameter-specific noise constraints.

Six machine learning architectures were implemented and compared for the multi-axis ray direction prediction, each utilizing different computational patterns for capturing the spatio-temporal relationships within the dataset. The aim is to find the optimal fit for the overall methodology.

XGBoost [58] has an ensemble learning method using gradient-boosted decision trees. It processes the tabular data through feature significance, handles non-linear relationships using decision tree making it a suitable model for the heterogenous regression tasks.

Feed Forward Neural Network (FFNN) implements a multi-layered perceptron featuring dense layers, ReLU activations, and dropout regularization. This model processes flattened input features, mapping non-linear complex relationships between ideal and realistic ray data via backpropagation.

CNN Spatial Ray Predictor uses convolutional neural networks to model relationships from the solar field. By treating ray data as 2-Dimensional spatial arrangements instead of temporal patterns, the architecture employs 1D filters with attention mechanisms to extract spatial features and capture connectivity across the statistical array.

Autoencoder features an encoder-decoder structure for feature learning and dimension reduction. The model reshapes 1D input into a 2D format, compresses data into a latent space, and reconstructs the output through normalization and dropout regularization.

Bayesian Autoencoder is a probabilistic extension of the standard autoencoder framework utilizing Monte Carlo Dropout. The model uses dense layers to generate multiple predictions, enabling the estimation of output mean and variance.

WaveNet [59] represents a deep generative model which uses dilated convolutions along with residual connections. Although originally designed for audio generation, this temporal model is able to capture long-range dependencies through increasing dilations across multiple convolutional layers. The model processes the input sequence through gated activation functions (tanh and sigmoid), enabling successful modeling of temporal patterns of ray data.

The implementation progressed from adapted traditional standard ensemble methods (XGBoost)and standard deep-learning algorithms (FFNN) to complex spatial models (CNN) and self-tailored architectures such as Auto-Encoders or Bayesian Auto-Encoder with uncertainty quantification. Additionally, an intricate LSTM-type CNN (WaveNet) has also been tested against our self-tailored approaches. Detailed architectures for all models are provided in Fig. A 2.

3.2. Datasets

Datasets constitute the foundation upon which these data-driven corrections are built. On the one hand, the ray-level power correction is based on the flux prediction method presented in ref [52]. It takes the input and output images from that procedure to generate the correction matrix. This dataset is thoroughly described and available in [60].

On the other hand, the direction correction dataset follows the same underlying logic, encompassing the same 931 meteorological conditions (defined by DNIs and sun positions) as the power dataset. Likewise, it also comprises nine different defocused heliostat configurations. The training inputs consist of rays generated from simulations devoid of tracking errors, while the labels correspond to simulations with modelled tracking errors. Thus, the training is considered successful when the tracking error models are accurately predicted by the architecture.

While the previous dataset consisted of 15,928 “.jpg” images, this one comprises 25,137 “.txt” files. Each file contains the 4 governing parameters for each of the 169 distributions of a specific axis at a determined timestep, stacked as a column array. This structure results in 3 separated files per timestep (X, Y and Z axes) labelled by date and axis. For each of the 931 timesteps, the files consist of 676 components, where every 4 components describe the distribution fitting of a single grid element. Then, the same conditions are repeated across 9 different groups of focused heliostats, covering the majority of the CSP plant operating conditions [61].

3.3. Geometrical ray projection towards the receiver

By implementing the machine-learning-based ray-level correction described before, the authors ensure the rectification of rays at the aperture plane of any given receiver. At this stage, the algorithm stores the interception location of each ray – coordinates and grid element –, its rectified scalar power, and the 4 parameters defining the directional distributions for each mesh element across the three axes. However, the incident flux must ultimately be evaluated at the actual receiver surface (e.g. cylindrical or hemispherical cavity). Therefore, the specific direction for each power-corrected ray has to be derived, enabling their subsequent projection onto the complex geometry of the cavity, as Fig. 5 depicts.

To re-parameterize the problem, it is necessary to generate new directional values following the corrected distributions across the three axes. The algorithm generates three arrays for each grid element (X, Y and Z axes) composed of pseudo-random values sampled from their respective corrected probability distributions. These arrays’ lengths are determined by the specific number of rays intercepted within each mesh element, ensuring that a sufficient number of random values is generated for all existing rays. Afterwards, the three pseudo-random components are concatenated into a single vector which is then assigned to the corresponding ray. As a result, a set of fully rectified rays is obtained element-wise ($\overrightarrow{r'_n}$) from the original ray file ($\overrightarrow{r_n}$), as illustrated in right. Finally, the tuned rays are projected toward the contiguous receiver surface following the new director vector (left).

4. Results

The performance of the proposed methodology is assessed in this chapter by evaluating each of its components separately. First, the effect of the dimensional reduction is examined through a goodness-of-fit analysis for the three axes across the 169 mesh elements, using the previously mentioned timestamp (13.05.2022, 11:30 AM). Then, the ability of the neural networks to predict the effect of tracking errors at the aperture plane is demonstrated, based on ref [61] assessment framework, 5-fold validation and learning curve analysis. Finally, the universal AI-enhanced 3D flux prediction is exemplified for a real cylindrical geometry that reproduces HEHTRES receiver, installed at the DLR facilities in Jülich, by comparing experimental measurements and the corresponding AI-enhanced simulated cases.

4.1. Statistical direction fittings

The authors analyze the goodness-of-fit for each element’s directional data fittings across all three axes. With this aim, three statistical criteria are used: the Kolmogorov-Smirnov D-value, the normalized Anderson-Darling (A^2) metric, and the Cramér-von Mises D-value. Due to the high density of rays in certain mesh elements – exceeding a million rays within the central items – the usage of p-values was rendered unfeasible, as even negligible deviations lead to statistical significance at such scales.

The D-value from the Kolmogorov-Smirnov test provides the maximum deviation between the actual cumulative distribution

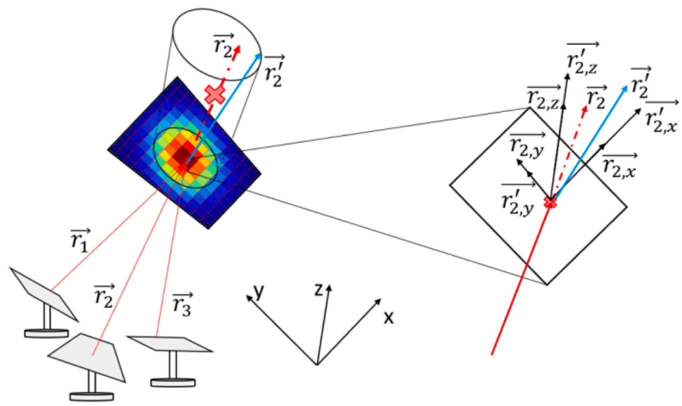

Fig. 5. Stochastic re-parameterization and 3D ray projection. (Right) Detailed picture of the original director vector ($\vec{r_2}$) rectification via stochastic re-sampling from fitted distributions at the aperture plane ($\vec{r_2'}$). (Left) Geometric projection of corrected rays onto the internal 3D surface of HEHTRES receiver.

function (CDF) and the theoretical CDF obtained from the fitting. Fig. 6 illustrates these values for the 169 mesh elements at the timestamp presented in Fig. 4.

The Cramer-von Mises D-value represents the mean square deviation between the theoretical (fitted) and realistic CDFs. Unlike the Kolmogorov-Smirnov test, this metric accounts for the cumulative differences across the entire domain rather than focusing solely on the peak deviation. These values are displayed in Fig. 7 for the whole mesh to facilitate a direct comparison with the Kolmogorov-Smirnov results.

Finally, the Anderson-Darling metric (A^2) quantifies the distance between the fitting and the sample along the whole domain, while assigning greater weight to the distribution tails and considering the sample size. In the case of asymmetric CDFs, this metric (Eq. (9)) provides critical additional information, as well as a more demanding goodness-of-fit criterion.

$$\boldsymbol{A}^2 = -\boldsymbol{n} - \frac{1}{\boldsymbol{n}} \cdot \sum_{i=1}^{\boldsymbol{n}} \left[(2 \cdot \boldsymbol{i} - 1) \cdot (\boldsymbol{ln}\, \boldsymbol{F}(\boldsymbol{x_i}) + \boldsymbol{ln}(1 - \boldsymbol{F}(\boldsymbol{x_{n+1-i}}))) \right] \tag{9}$$

Where n represents the sample size, x_i is the sorted data and $F(x)$, the CDF of the theoretical distribution.

Considering the conceptual alignment between the Anderson-Darling and Cramér-von Mises analyses, the authors correlated both methods (Fig. 9) to identify the outlying mesh elements exhibiting significant deviation between the sample data and the fitted distribution. With this aim, the five elements with the highest A^2 values and the five with the highest CVM D-Values were explicitly labeled in the correlation plots provided below. This dual-metric approach ensures that both tail-end discrepancies and cumulative distribution misalignments are accounted for during the validation of the fitting process.

Fig. 6 demonstrates that deviations in the X-axis are consistently higher than those in Y and Z-axes. Whereas the double-step-function fitting presents deviations up to 35.0 % in outlying elements, the two beta-distribution fittings show maximum deviations between 12.0 % and 14.0 %. Specifically, for the X-axis, only 14 samples deviate by >25.0 %, while the remaining elements yield accurate fittings with D-values between 10.0 % and 25.0 %. In the case of Y and Z-axes, the average peak inaccuracies remain below the 5.0 % (with typical deviations below 10.0 %), which shows consistent behavior during the dimensional reduction. Notably, the highest deviations are located in regions distant from the high flux density area (ROI center). Considering that the aimpoint is centered within the ROI (element 84), the items from 60 to 125 are the most critical for the fitting; consequently, outliers outside this range may be considered negligible.

The analyses presented in Figs. 7 and 8 illustrate the extent to which the fitted data deviate from the sample across the entire domain. In the case of Y and Z axes, the behavior is consistent and stable. In the Y axis, only eight elements present outlying deviations, with only two located within the central region (elements 60 and 125). In the Z-axis, the variability is slightly higher, but <10.0 % of the elements are outliers (with only 6 samples within the central region). The metrics for these two axes are indeed within the same order of magnitude, and the difference between the A^2 and CVM D-Value is only one order of magnitude owing to the moderate skewness of the Beta-function. Regarding X-axis, the metrics remain correlated, but the high skewness causes A^2 to be two orders of magnitudes larger than CVM D-value. Furthermore, the values here appear rather chaotic, with the lowest deviations located at the boundaries; this is due to the inherent ability of a double step to fit scattered data. For example, a single ray would behave as a step with infinite amplitude. Thus, rather than isolated outliers, there is a continuous, albeit acceptable, deviation across the axis.

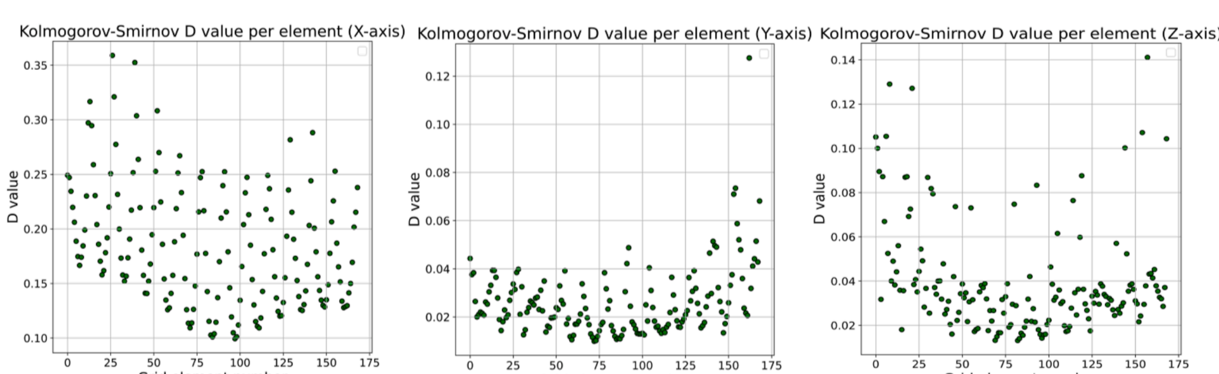

Fig. 6. Kolmogorov-Smirnov D-value for each of the 169 mesh elements in their three axes.

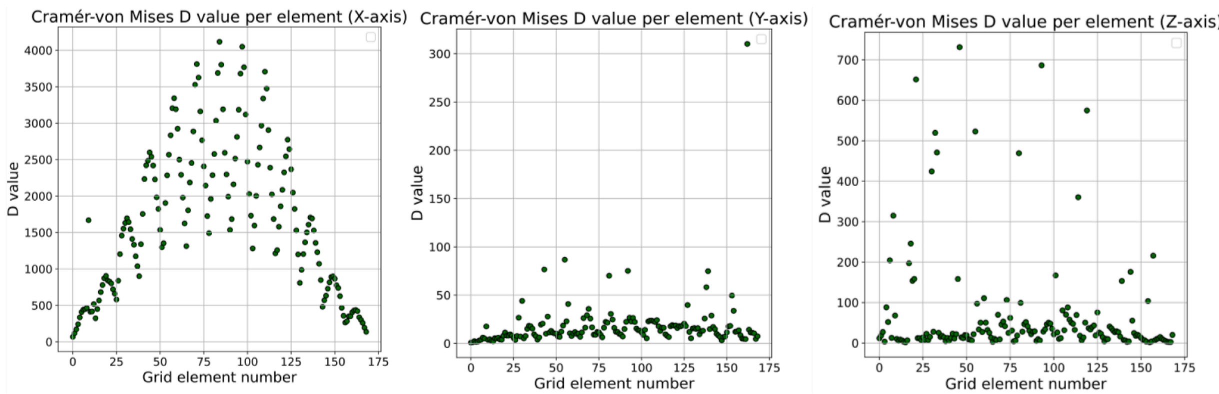

Fig. 7. Cramér-von Mises D-value for each of the169 mesh elements in their three axes.

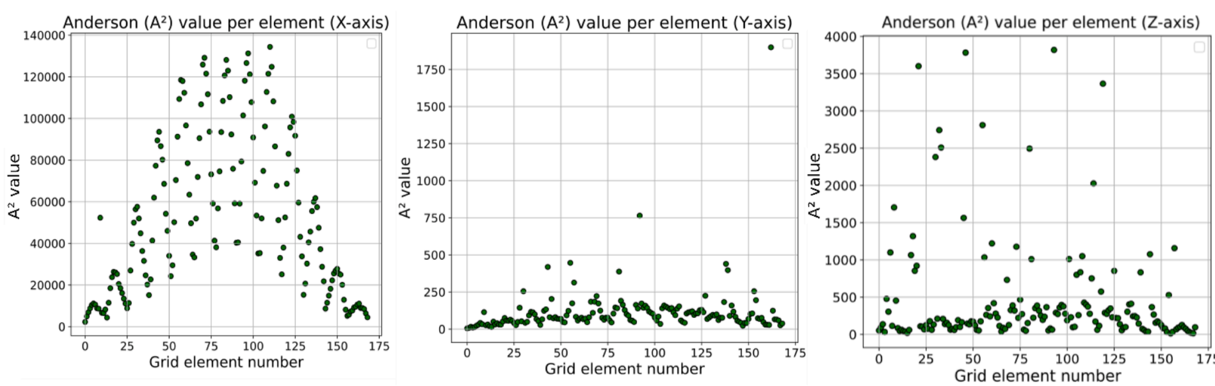

Fig. 8. Anderson-Darling (A^2) results for each of the 169 mesh elements in their three axes.

Fig. 9 confirms the correlation between A^2 and CVM D-value, highlighting the dispersed distribution of the X-axis deviations, compared to the other directions. These results indicate that the fitting accuracy for the Y and Z-fittings is outstanding across the whole domain, particularly in highly irradiated areas. While the X-axis re-parameterization shows low peak deviations, the trend remains consistent across the whole domain. Therefore, it is possible to estimate that the average inaccuracies will fall within the same range as the peak values (10.0–25.0 %), which still represents a remarkably robust fitting. To validate these results, Fig. A 1 provides sanity checks for three representative elements: 22 (lower-left corner), 85 (center), and 145 (upper-right corner).

4.2. *Computational performance and model efficiency*

Training time, inference latency, and parameter count for all six architectures are reported in Tables 2 and 3, measured on an NVIDIA RTX 2060 GPU (CUDA 11.8) under sequential single-file inference conditions to simulate real-time deployment.

Training times span two orders of magnitude. XGBoost requires the longest total training (650.3 min) due to its 676 independent regressors, each comprising 200 gradient-boosted trees — a cost that scales with output dimensionality rather than dataset size, and that is reflected in its 68.9 million decision tree nodes, which are not directly comparable to neural network weight counts. WaveNet occupies an intermediate tier

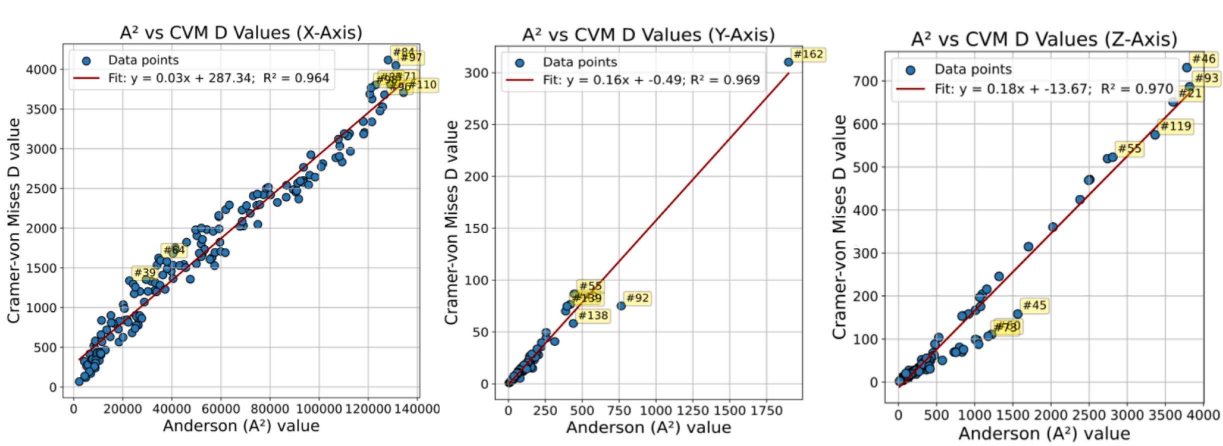

Fig. 9. Cross correlation analysis between Cramer-von Mises D-value and Anderson-Darling (A^2) metrics correlated for the three axes. Labeled points highlight the top five outliers of each metric, primarily located in peripheral regions within the ROI.

(38.5 min total), while the CNN (6.6 min), FFNN (7.9 min), Auto-Encoder (13.6 min), and Bayesian Auto-Encoder (7.6 min) all complete training within 14 min across all three axes. Since training is a one-time offline investment, inference latency is the primary operational constraint for deployment.

Inference times reveal critical differences in deployment suitability. XGBoost exhibits by far the highest total per-file latency (2009.3 ms across three axes), incompatible with near-real-time monitoring requirements and effectively disqualifying it from operational deployment despite its competitive prediction accuracy. WaveNet also shows elevated total latency (53.0 ms). The CNN achieves the lowest confirmed sequential inference time (4.6 ms total, 1.5–1.6 ms per axis), followed by the Auto-Encoder (13.6 ms). It should be noted that FFNN and Bayesian Auto-Encoder inference times (0.04 ms and 0.06 ms total respectively) reflect batch GPU processing rather than true sequential single-file inference and are therefore underestimated; these figures are not directly comparable to the remaining models.

4.3. Ray-level direction prediction assessment

The detailed evaluation of the six machine learning models provided distinct performance insights, identifying the optimal models for multi-axis directional prediction. The authors chose the Pearson correlation coefficient to assess model behavior across the datasets (Tables 4–6).

X-axis predictions yielded the highest accuracies across all ML models, likely due to the relatively simplistic azimuthal orientation prediction. The CNN Spatial Ray predictor achieved the best overall results among the deterministic networks by effectively modeling spatial relationship, although WaveNet also delivered competitive prediction results using its temporal dependency capturing. High overall accuracies across all parameter predictions identify X-direction as the most predictable axis. While simple robust models such as Feed Forward Neural Network successfully predicted x_1 and a_2 parameters, x_2 and a_1 prediction accuracies were only moderate. XGBoost failed to predict x_1, while Bayesian Auto-Encoder performed exceptionally well in x_1 and a_1 and comparatively moderate in x_2 and a_2 performance.

Y-axis data provided the greatest complexities due to the elevation angle dependencies. Bayesian Auto-Encoder provided the second-highest overall accuracy and outperformed the other models in "Location" prediction by a wide gap; this indicates that statistical uncertainty quantification is dominant for such directionally complex dataset. However, the CNN spatial predictor achieved the highest prediction ability for 3 out of 4 parameters, making it the most robust overall option. While XGBoost and WaveNet performed consistently above 80.0 % accuracy, neither reached top-tier performance. Finally, the Auto-Encoder and FFNN models exhibited moderate learning capabilities (below 85.0 % accuracy), making them basic but less precise benchmarks (Table 5).

Z-axis predictions required intermediate computational complexities to balance spatial dependencies with temporal variations. Bayesian Auto-Encoder showed exceptional alpha performance, while XGBoost achieved balanced beta and location predictions. FFNN demonstrated strong location predictions, proving that traditional methods are biased toward certain parameters. The Auto-Encoder model's scale parameter declined dramatically, which highlights the criticality of model selection for such direction prediction tasks. Once more, the CNN spatial predictor shows the strongest average performance – all accuracies over 80.0 % and an overall average above 85.0 %.

The CNN was selected as the final architecture based on three criteria evaluated simultaneously: lowest total training time among spatially-aware architectures at 6.6 min; lowest confirmed real-time inference

Table 2
Axis-wise training times, total training times and parameter count for the six tested architectures.

ML Model	X-axis (min)	Y-axis (min)	Z-axis (min)	Total (min)	Parameter count (-)
XGBoost	223.8	277.8	148.7	650.3	68,900,000*
Wavenet	20.2	9.8	8.5	38.5	494,245
CNN Spatial Ray Predictor	3.3	1.3	2.0	6.6	1023,568
Bayesian Auto-Encoder	2.3	4.5	0.8	7.6	187,332
Auto-Encoder	4.5	5.1	4.0	13.6	437,772
Feed Forward Neural Network	3.2	2.7	2.0	7.9	563,044

* Parameters from decision trees are not directly comparable to NN weights.

Table 3
Average axis-wise inference times for the 63 samples analyzed in Section 4.3.

ML Model	X-axis (ms)	Y-axis (ms)	Z-axis (ms)	Total (ms)
XGBoost	693.1	623.3	692.9	2009.3
Wavenet	17.0	16.6	19.4	53
CNN Spatial Ray Predictor	1.6	1.5	1.5	4.6
Bayesian Auto-Encoder	0.02	0.02	0.02	0.06**
Auto-Encoder	4.5	5.1	4.0	13.6
Feed Forward Neural Network	0.02	0.01	0.01	0.04**

** FFNN and Bayesian AE inference times reflect batch GPU processing and are underestimated; per-file sequential timing was not available for these runs.

Table 4
Comparative performance metrics of six machine learning architectures for X-axis prediction. Mean Pearson correlation coefficients are reported for the distribution parameters in validation phase.

X-Direction Pearson correlations					
ML Model	x_1	x_2	a_1	a_2	Average
XGBoost	0.733	0.910	0.943	**0.980**	0.891
WaveNet	0.766	0.925	0.926	0.950	0.891
CNN Spatial Ray Predictor	0.835	0.924	**0.965**	0.950	0.918
Bayesian Auto-Encoder	**0.932**	**0.945**	0.946	0.951	**0.943**
Auto-Encoder	0.850	0.860	0.795	0.850	0.838
Feed Forward Neural Network	0.761	0.930	0.897	0.889	0.869

Table 5
Comparative performance metrics of six machine learning architectures for Y-axis prediction. Mean Pearson correlation coefficients are reported for the distribution parameters in validation phase.

Y-Direction Pearson correlations					
ML Model	Alpha	Beta	Location	Scale	Average
XGBoost	0.901	0.806	0.853	0.815	0.843
WaveNet	0.918	0.807	0.860	0.821	0.851
CNN Spatial Ray Predictor	**0.920**	**0.814**	0.853	**0.830**	**0.854**
Bayesian Auto-Encoder	0.905	0.793	**0.953**	0.758	0.852
Auto-Encoder	0.817	0.778	0.908	0.749	0.813
Feed Forward Neural Network	0.910	0.789	0.768	0.820	0.821

Table 6
Comparative performance metrics of six machine learning architectures for Z-axis prediction. Mean Pearson correlation coefficients are reported for the distribution parameters in validation phase.

Z-Direction Pearson correlations					
ML Model	Alpha	Beta	Location	Scale	Average
XGBoost	0.879	**0.894**	**0.894**	0.773	**0.860**
Wavenet	0.857	0.854	0.866	0.760	0.834
CNN Spatial Ray Predictor	0.866	0.881	0.878	0.815	**0.860**
Bayesian Auto-Encoder	**0.907**	0.804	0.705	0.793	0.802
Auto-Encoder	0.842	0.811	0.718	0.509	0.720
Feed Forward Neural Network	0.855	0.866	0.888	**0.831**	**0.860**

latency at 4.6 ms total; and the best overall performance in terms of validation accuracy. With 1023,568 parameters, it also represents a compact architecture relative to WaveNet (494,245 parameters but 53 ms latency) and XGBoost. No competing architecture matches the CNN across all three criteria simultaneously: XGBoost is unsuitable for real-time inference regardless of its accuracy; WaveNet incurs an eightfold inference penalty relative to the CNN without compensating accuracy gains; the Bayesian Auto-Encoder, despite providing valuable 99 % confidence interval estimates aligned with the stochastic nature of Monte Carlo simulations, was excluded from final deployment due to its batch-processing dependency being incompatible with sequential near-real-time execution. Although computing times appeared low, this model caused a significant computational overhead. Consequently, the CNNs were examined under 63 diverse meteorological conditions to test their inference abilities. For the sake of generalization, 9 different DNI levels have been tested – from 100.0 to 900.0 W/m^2 – and 7 timesteps per DNI (22–03–2022@09:00:00, 22–03–2022@13:00:00, and 22–03–2022@ 18:00:00 close to spring equinox; 22–06–2022@09:00:00, 22–06–2022@13:00:00, and 22–06–2022@18:00:00 close to summer solstice; and 01–08–2022@15:00:00, reaching the late high-irradiation season) with randomized focused heliostats. The timestamps were selected to cover the widest range of sun positions – height and azimuth – as well as DNI values, following the generalization principles commented from ref [52]. The authors use Pearson correlation coefficient (X-axis) and Spearman rank (Y and Z axis) to quantify the performance of these models across all 63 cases for the 4 parameters. Results are grouped by DNI in boxplots, as displayed in Figs. 10 and 11, and model robustness was assessed through analysis of variance (ANOVA) and confidence interval (CI) studies, as displayed in Fig. 11.

The authors considered 80.0 % Pearson correlation as the foundational performance target for this novel prediction method. Regarding X-axis inference, this aim was surpassed by a wide margin for 3 out of 4 variables. Average correlations ranging from 81.0 % (for x_1 parameter) to 92.7 % (for a_2) were reached across the entire dataset. In addition, the results show a robust and consistent behavior across the 9 DNIs and all parameters, with nearly every individual sample and all categorical averages exceeding the 80.0 % correlation. The only exception was caused by the 7 timesteps analyzed at 500.0 W/m^2, where the poor performance at 22–06–2022@13:00 (anomalous DNI for this timestep) and 21–03–2022@09:00 (edge case) hindered performance. These findings, along with the confidence intervals for each variable are detailed in Fig. 11.

The confidence interval analysis yields consistent results across all variables. For x_1, the CIs indicate a ± 15.0 % variance in Pearson correlations, with narrower spans at lower DNIs. Regarding performance against timestamp, results remain comparatively stable across most timesteps, with the exception of 22–03–22@18:00 (an edge case due to low solar altitude) and 21–06–2022 (weighed down by the anomalous case of 500.0 W/m^2 at 13:00). Beyond these outliers, the prediction performs robustly within a ± 15.0 % range for all other timesteps. Apart from that, the ANOVA shows that the robustness of x_1 prediction is not dependent of DNI – F = 0.812 and p > 0.050 – while showing a slight

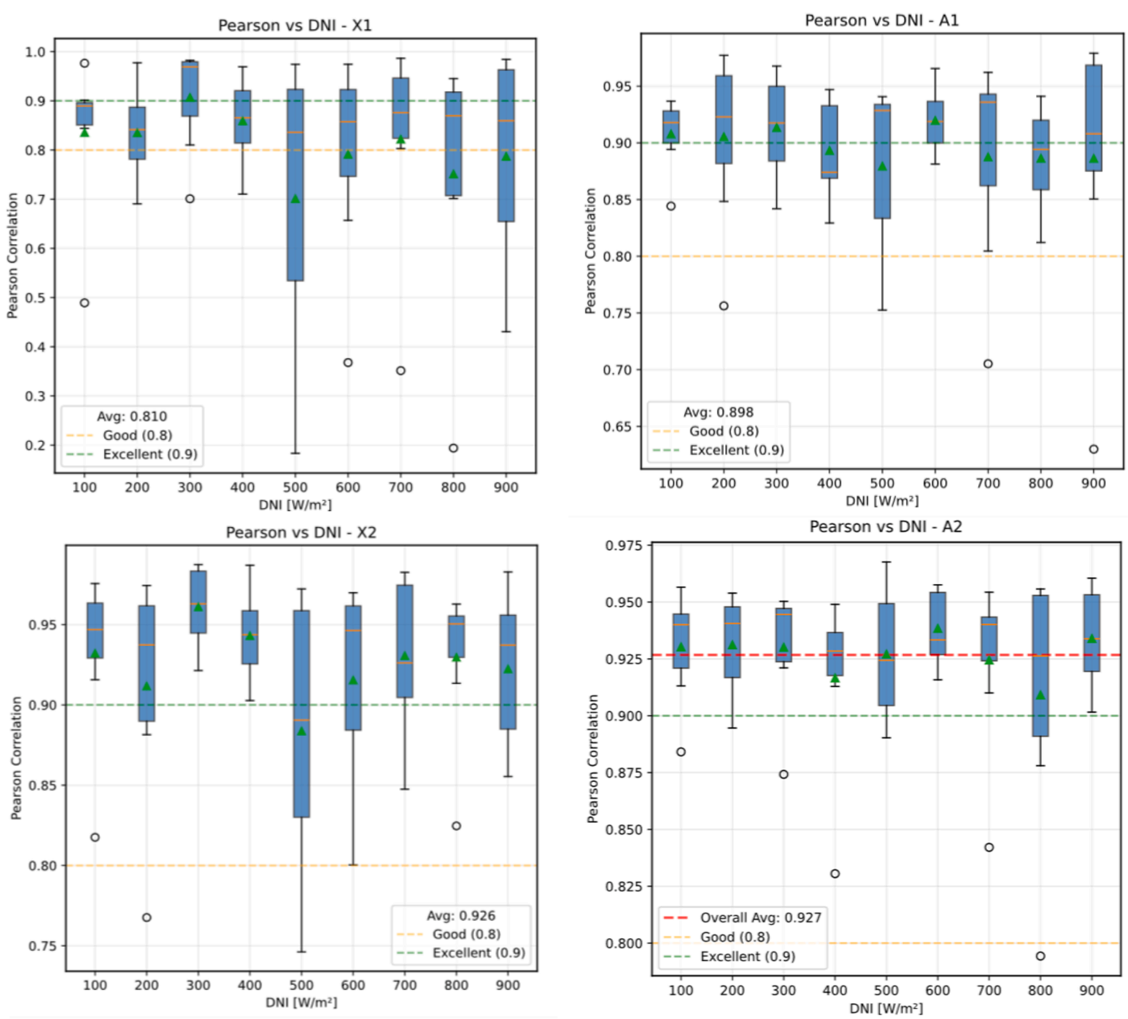

Fig. 10. Boxplot of X-axis parameters' Pearson accuracies against DNI for the 63 samples.

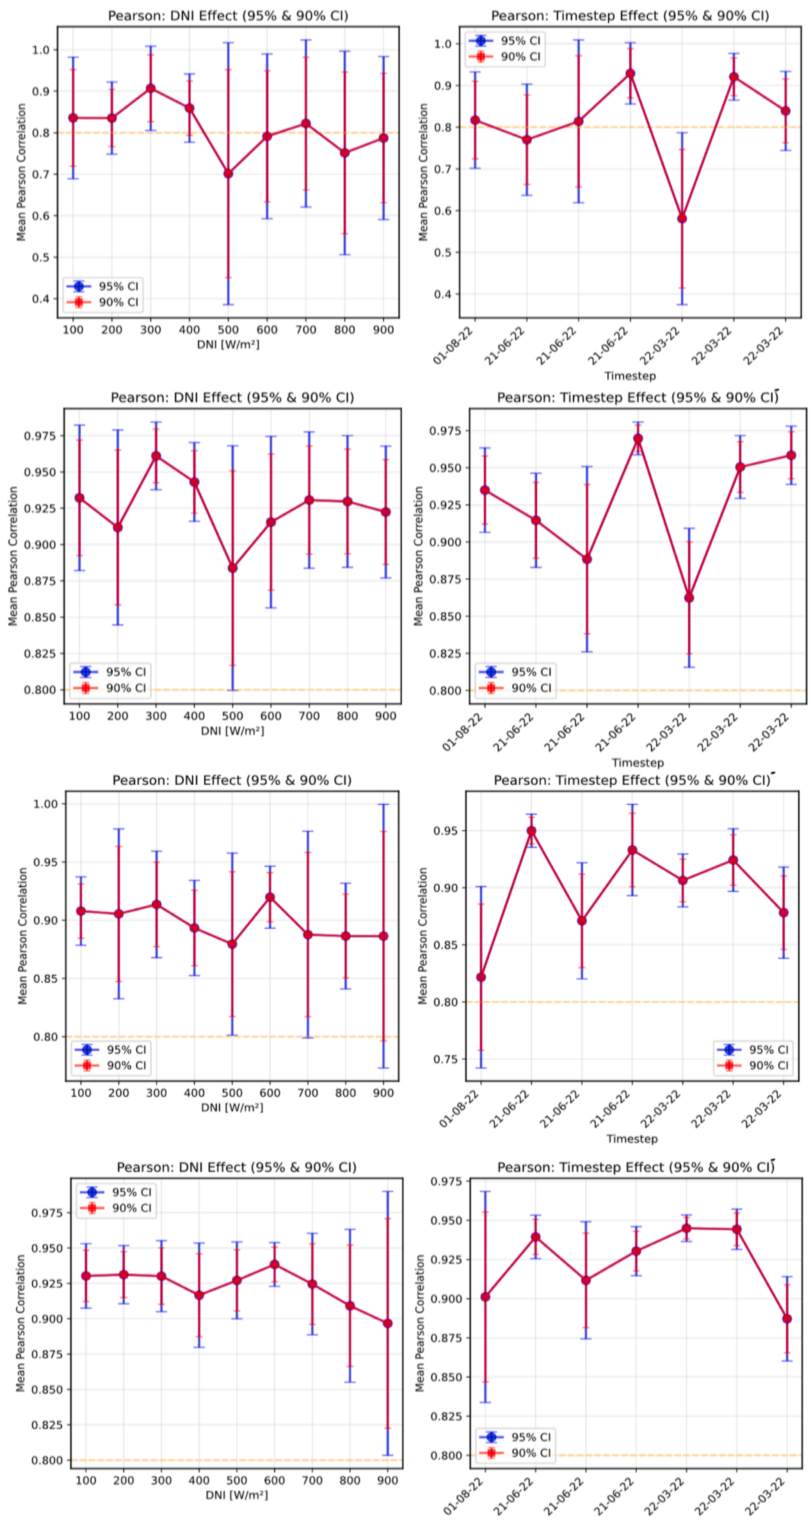

Fig. 11. Average Pearson correlation per DNI (left column) and timestep (right column) with 90 % and 95 % CI. First row: x_1; second row: x_2; third row: a_1; fourth row: a_2.

dependence on the timestep due to the aforementioned outlying cases – F = 3.834 and p < 0.050.

The same conclusions apply to the rest of variables: x_2, a_1 and a_2 prediction accuracies are independent of DNI but slightly influenced by the timestep. In these three cases, p-values are near the 0.05 threshold, showing a marginal dependence primarily caused by edge cases (22–03–22@09:00, 22–03–22@18:00, and 21–06–2022 at low DNIs). Notably, the CI ranges decrease to ±5.0 % for a_2 and ±10.0 % for $x_2$and a_1. As a conclusion, the model behaves accurately (above 80.0 %) and robustly for the X-axis – with almost no functional dependencies on the representative variables.

While Pearson correlation was utilized for the X-axis correlation due to its linear nature, the authors chose Spearman's rank for the Y and Z-axes. These axes employ Beta distribution fittings characterized by strong non-linearities, which result in non-linear distribution relationships. As displayed in Fig. 12 for the Y-axis, the four average correlation values exceed again the 80.0 % target, which ensures a high degree of correlation. In fact, 3 out of 4 parameters surpass that benchmark by >5.0 %. In this axis, the model robustness is even higher than in the X-axis, showing fewer than 5 anomalous values per parameter.

Additionally, the DNI-wise behaviour here is notably more stable; unlike the 500.0 W/m^2 case observed for the X-axis, variations here remain consistently low across all DNIs. Average correlations exceed 80.0 % for all parameters, with the sole exception of the Beta parameter during 01–08–2022 at 15:00 timestamp. For this specific scenario at 900.0 W/m^2, the Spearman's rank correlations yield a negative result for Alpha, alongside diminished values for Beta and Location parameters. Considering that Spearman's rank assesses the monotonic relationship between variables, a negative alpha value shows an inaccurate prediction (below 70.0 %) that consistently follows an opposite trend relative to the ground truth.

Furthermore, the low correlations for Beta and Location highlight an inability to properly predict that specific anomalous case. However, given the consistently high performance at the same timestamp for DNI levels ranging from 100.0 to 800.0 W/m^2, the authors attribute this behavior to artifacts present within the original ray file obtained for 01–08–2022 at 15:00 and 900.0 W/m^2 scenario, rather than a systemic failure of the parametric prediction. This phenomenon is deeply analysed at Fig. 13.

The Y-axis ANOVA yielded even more robust results than the X-axis. As shown in Fig. 13, none of the parameters presented a dependence on the DNI variable ($p > 0.050$), supporting the model's equivalent robustness across different irradiance levels. In fact, all 4 variables present uniform 95 % CIs around ±7.5 %, with the only observable trend being a slightly growing CI for the location and scale parameters at high DNIs (> 700 W/m^2). At 900 W/m^2 results were influenced by the 01–08–2022 anomaly. However, the other two DNIs were affected by higher probability concentrations around smaller Y values, which makes the Beta function more difficult to map due to higher frequency and magnitude of outliers. In fact, F-value is higher than 1 for Alpha and Scale, which means that noise is significant in these 63 samples even without dependence. Given these results, refining the dataset to exclude extreme edge cases would substantially enhance future inference performance.

On the other hand, this analysis also shows the robustness dependence against the timestamp for the four parameters. Although there are not observable biases, the 22–03–22 and the 21–06–22 show a consistently lower performance than the other timestamps. However, their 95 % CI remain in the range of ±7.5 % (or even ±5.0 % for alpha or scale) although the mean prediction accuracies decrease, which shows robustness for the edge cases. This behaviour is supported by ANOVA, that provides p-values below 0.050 for all the variables, penalising edge cases. It is also worth mentioning that the confidence interval at 01–08–2022 is increased because of the noisy sample at 900.0 W/m^2. Finally, the equivalent analysis is shown in Fig. 14 for Z-axis.

Fig. 14 indicates that Z-axis generalization was partially successful, but certain challenges remain. Beta parameter exceeded the 80.0 % average correlation target in 7 out of 9 DNI levels. Z-axis ray distribution

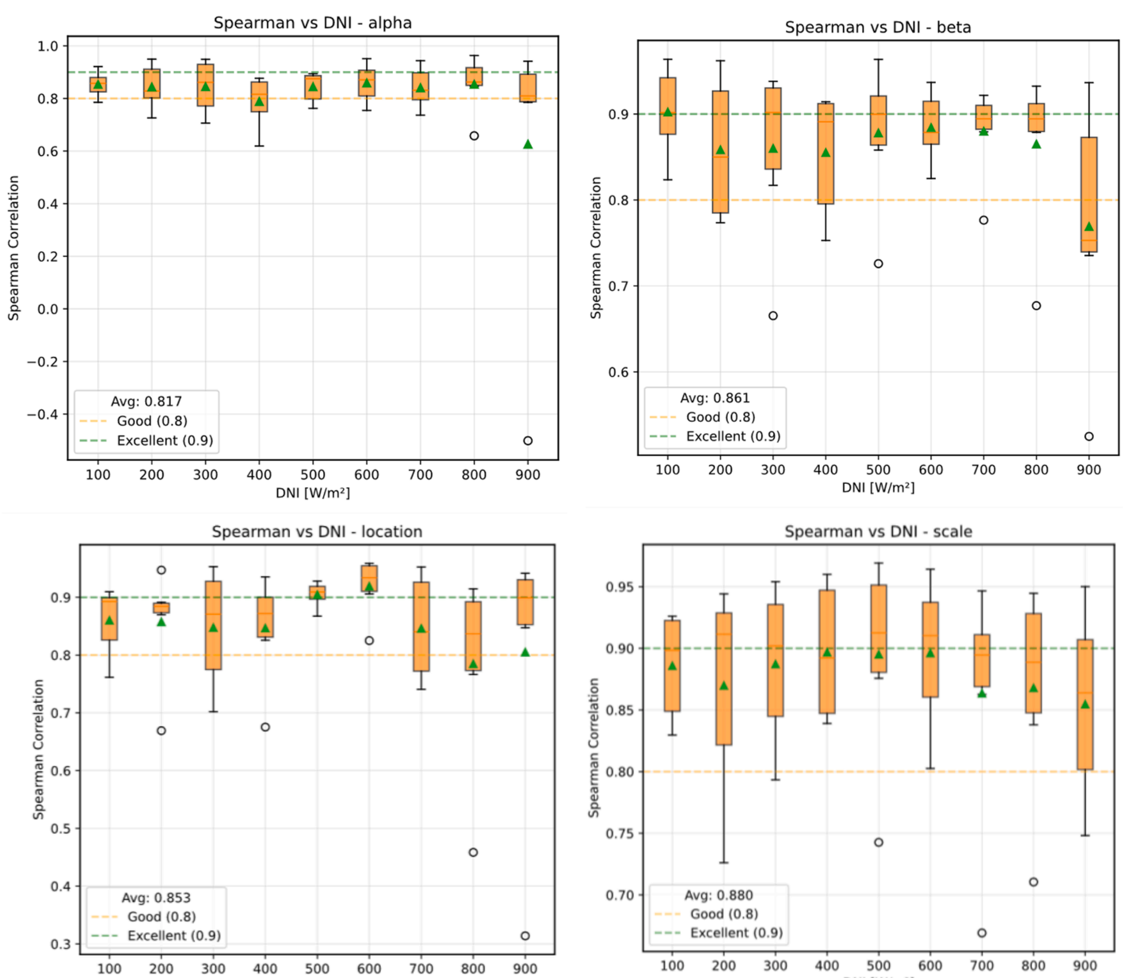

Fig. 12. Boxplot of Y-axis parameters' Pearson accuracies against DNI for the 63 samples.

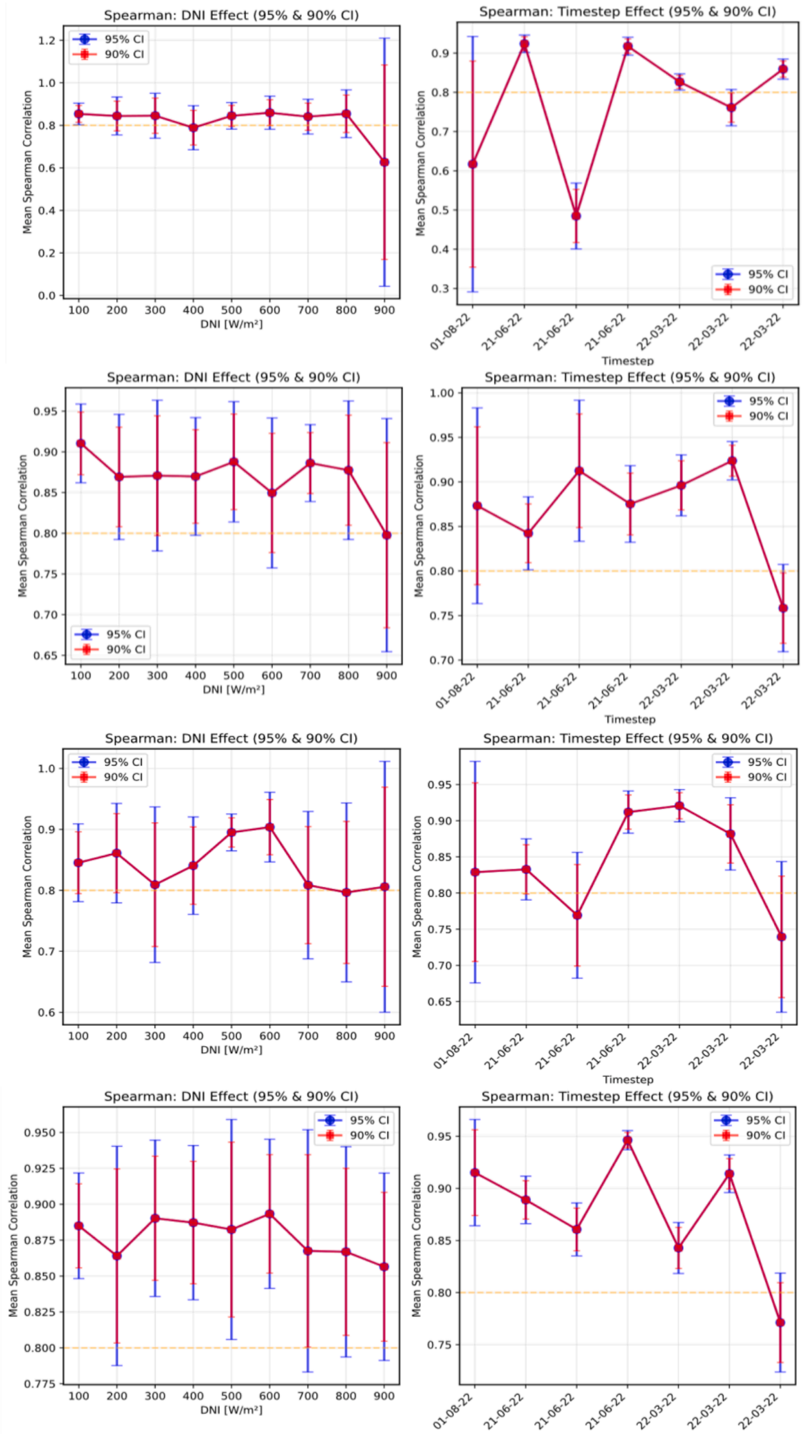

Fig. 13. Average Spearman rank correlation per DNI (left column) and timestep (right column) with 90 % and 95 % CI. First row: Alpha; second row: Beta; third row: Location; fourth row: Scale.

also shows anomalous values for 01–08–2022 at 15:00 and 900.0 W/m^2 scenario, so this outlier is not considered a failed simulation. The alpha and location inferences showed acceptable values (over 70.0 %), with some of the DNI averages reaching 80.0 %. This represents a solid first step for this predictive approach. However, scale prediction is considered incomplete, with average values below 70.0 % and weaker robustness – high IQR and varying correlations. This phenomenon, attributed to the high obliquity of rays at specific timesteps, is thoroughly studied in Fig. 15.

While the model demonstrated high robustness for X and Y predictions against DNI, the Z-axis exhibited a measurable dependence on irradiance levels. Alpha and beta p-values (against DNI) are 0.033 and 0.014, respectively, which indicates significant sensitivities for these edge cases. In fact, it is possible to observe a certain Beta prediction's accuracy loss alongside an expansion of the CI at higher DNIs. However, the average 90 % CI remain within ±7.5 % (at all DNIs except 900 W/m^2), with average prediction values between 70.0 % (for Alpha) and 85.0 % (for Beta).

Conversely, location parameter showed a robust behaviour at 72.3 ± 5.2 % (90 % CI) without dependencies on DNI, which demonstrates

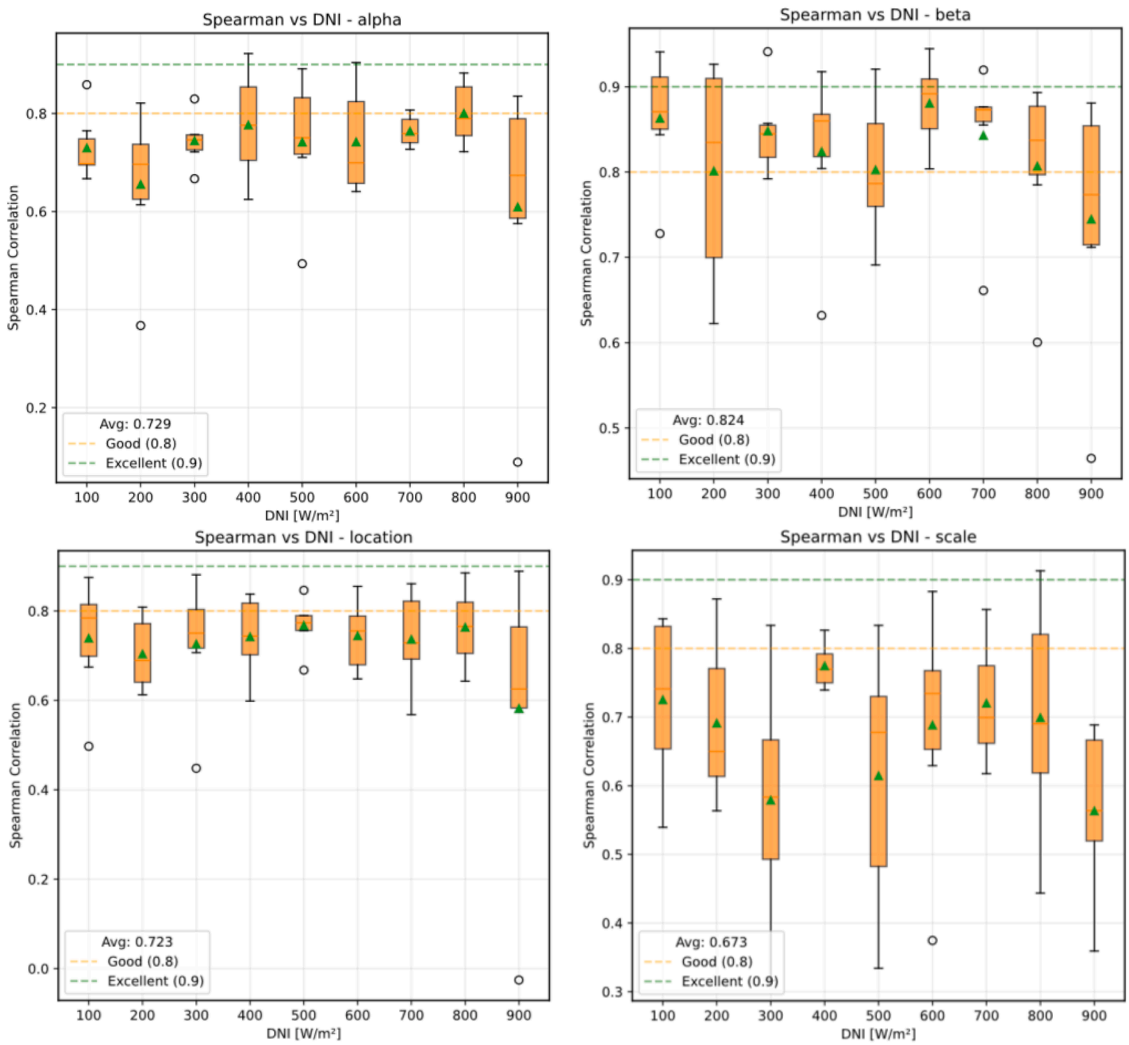

Fig. 14. Boxplot of Z-axis parameters' Pearson accuracies against DNI for the 63 samples.

possibility to improve but a certainly solid first step toward ray-direction prediction. Finally, scale parameter showed the highest DNI sensitivity and a ± 10.0 % (90 % CI) variance, with average accuracy barely reaching 70.0 %.

Furthermore, all parameters exhibited a certain dependence on the timestamp. Despite this, the confidence intervals showcase a solid behavior, with averages over 70.0 % for all parameters except scale. This robustness was primarily hindered by the edge case on 22–03–22 at 9:00 (when the rays are extremely oblique). To sum up, the Z-axis prediction generalization against edge cases has been partially successful for 2 out of 4 parameters (alpha and location), has exceeded the target for beta, and needs to be further improved for scale.

4.4. Model robustness and generalization assessment

Three complementary analyses address overfitting risks and generalization bounds: 5-fold cross-validation, learning curve analysis, and a review of embedded regularization mechanisms.

Fig. 16 presents the 5-fold cross-validation results. Standard deviations are consistently small (±0.001–0.005), confirming performance stability across data partitions and providing direct evidence against overfitting. All parameters exceed the 0.80 reliable threshold across all axes, with overall correlations of 0.964 ± 0.001 (X-Pearson), 0.888 ± 0.002 (Y-Spearman), and 0.867 ± 0.004 (Z-Spearman). The highest variance is observed for Z-axis Scale (±0.004) and Y-axis Beta (±0.002) — precisely the parameters showing weaker absolute performance in the 63-condition generalization test — confirming their underperformance reflects genuine physical complexity rather than a training artifact.

Fig. 17 presents the learning curves across 20 %–100 % training fractions. All axes plateau before 80 %, with negligible gains between 80 % and 100 % (ΔOverall = +0.0014 on Y, +0.0001 on Z), confirming the model is not data-starved and that performance ceilings on Scale and Beta parameters are not attributable to insufficient data. Notably, the Y-axis Overall Spearman correlation reaches 0.9070 at 80 % and only 0.9056 at 100 %, while the Z-axis reaches 0.9272 at 80 % and 0.9222 at 100 % — both showing effective convergence well before the full dataset is used. Future improvements should therefore target architectural modifications rather than data collection.

Regarding regularization, the CNN incorporates dropout (rate 0.1) after each convolutional layer, input batch normalization, and a two-head self-attention module that enforces spatially distributed representations. The Bayesian Auto-Encoder's superiority on Y-axis Location prediction (Pearson 0.953) confirms that uncertainty-aware training benefits the most physically complex targets, motivating lightweight uncertainty quantification in future CNN iterations. Together, these analyses confirm that weaker generalization on Z-axis and Y-axis scale parameters reflects physical limitations at edge sun elevation angles rather than modelling deficiencies.

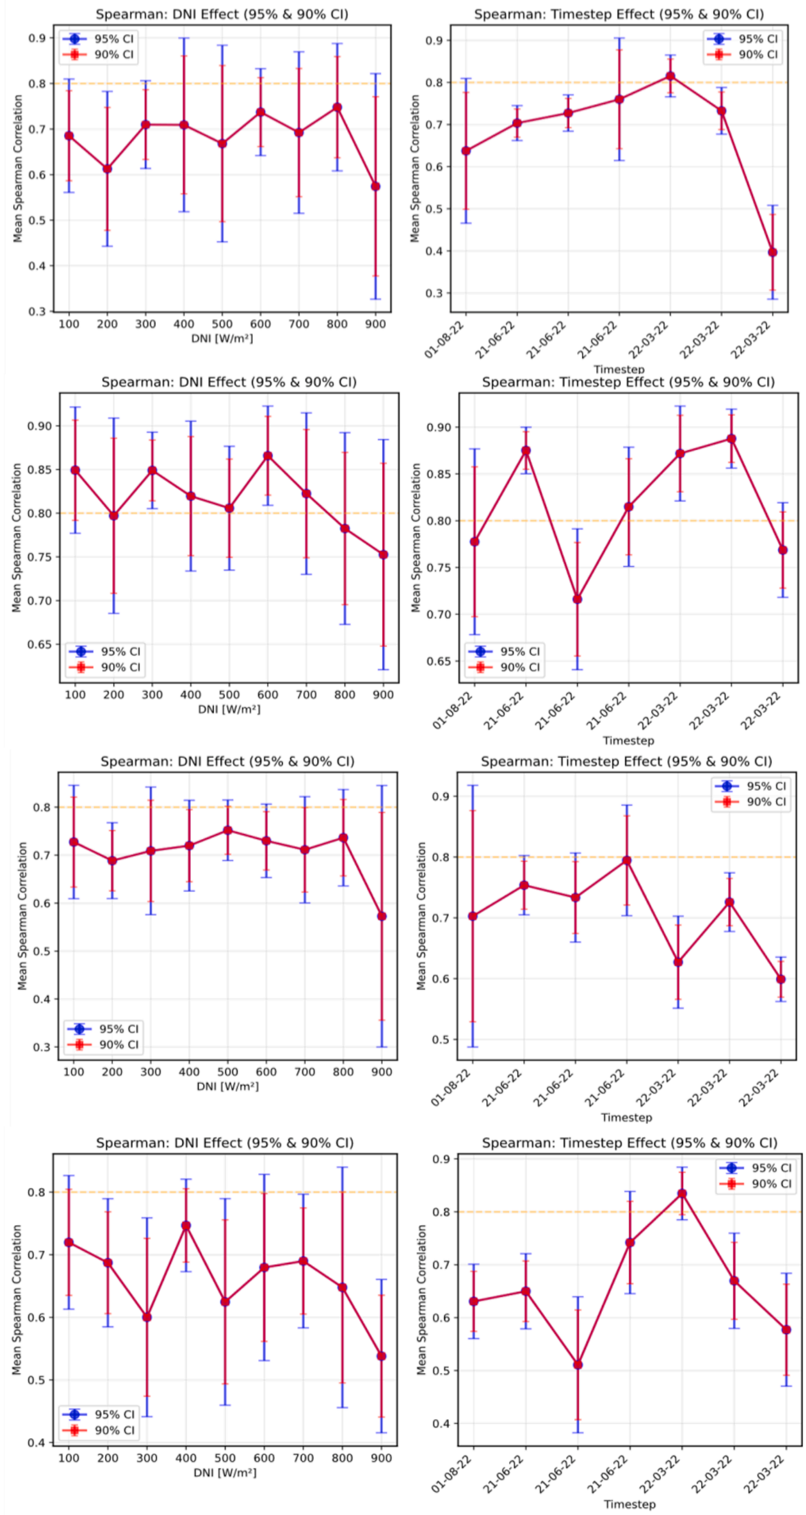

Fig. 15. Average Spearman rank correlation per DNI (left column) and timestep (right column) with 90 % and 95 % CI. First row: Alpha; second row: Beta; third row: Location; fourth row: Scale.

4.5. Cavity receiver predicted flux density distribution

To assess the global performance of this integrated solution, the authors have compared the final simulated flux density results across the 3D cavity receiver surface against camera-based measurements at the aperture plane of the real receiver (depicted in Fig. 2). The experiments were conducted at the third level (40.8 m high) of the *Multifokus Turm* (MFT) at the DLR premises in Jülich, utilizing a white Lambertian target at the aperture plane of a CentRec® receiver (HEHTRES). All the meteorological variables were measured by the meteorological station located at *Solarturm Jülich* heliostat field, and the flux images were captured by a Baumer TXG14 camera (100.0 mm lens) located 88.4 m from the receiver.

The digital twin methodology – comprising a virtual STJ ray-tracing

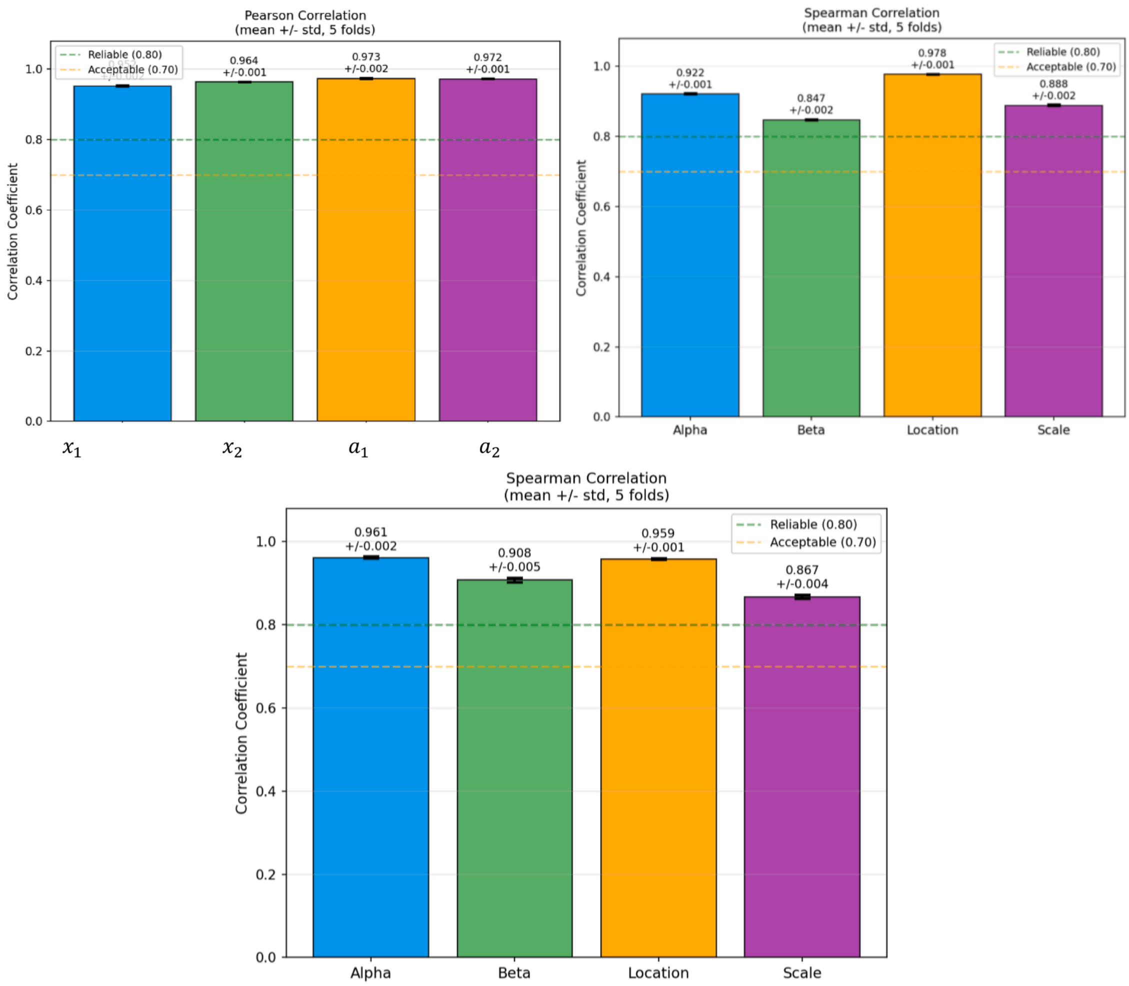

Fig. 16. (Upper left) 5-fold Pearson correlation results for X-axis. (Upper right) 5-fold Sperman correlation results for Y-axis. (Lower center) 5-fold Spearman correlation results for Z-axis.

environment, four convolutional neural networks trained on simulated data, and a ray projector inside of the cavity – replicates the meteorological and technical conditions of the real experiment. The comparison details are summed up in the Table 7 below.

First, the flux density distribution was measured in the white Lambertian target through the usage of the optical camera. The result is a greyscale 12-bit image that needs to be scaled against a radiometer value. To this end, the light beam was centered on the lower left radiometer, and the peak flux density value was used to re-scale the picture. The calibration process consists of assigning the peak flux value to the peak greyscale value and then interpolating the rest of grey values given the linear behavior of both the radiometer and the camera. The receiver aperture is then cropped and its flux values are extracted as the only relevant ones.

Once the flux map is obtained, the two horizontal black stripes (check Fig. 2) resulting from the manufacturing process need to be rectified. The flux values are interpolated linearly interpolated in this region – between the flux values below and above the breaches – after removing the original values. The whole flux map processing method is detailed in Fig. 18.

Then, the digital twin simulation is performed with the corresponding input parameters and without tracking errors. The ray files and the flux map are provided with the tracking error effects through the usage of the different neural networks in the aperture plane, and the rays are projected towards the receiver 3D shape. The corrected power and ray interception distributions along the aperture plane are depicted in Fig. 19.

The receiver flux distribution, resulting from the projection of the previously mentioned rays, is depicted in Fig. 20. Both the measured and simulated powers are calculated by integrating both fluxes. If the method is accurate, the simulated integrated power at the receiver surface and the integrated measured power at the aperture plane must be the same.

The integrated flux density within the cavity indicates a total thermal heat of 577.2 kW, against 714.9 kW integrated from the camera-based aperture-plane measurement, demonstrating an accuracy of 80.9 % in terms of power.

This experiment constitutes direct experimental validation of the complete integrated method under real operating conditions. The digital twin was deployed using live meteorological measurements without any site-specific retraining, demonstrating its applicability as a non-disruptive monitoring tool. From an operational standpoint, 80.9 % integrated power accuracy represents a meaningful first achievement for cavity receivers, where no prior computational or experimental method provides any prediction capability. For plant operators, this translates into the ability to estimate receiver thermal load distribution during normal operation without interrupting the heliostat field, enabling proactive aiming strategy adjustments and reducing the risk of localized overheating — a primary cause of receiver degradation in commercial plants. To further consolidate these results, a serial experimental campaign comprising repeated measurements across diverse operating conditions is planned at the Solarturm Jülich facility during summer 2026, which will provide statistically robust validation of the method's consistency and temporal stability under real plant operation.

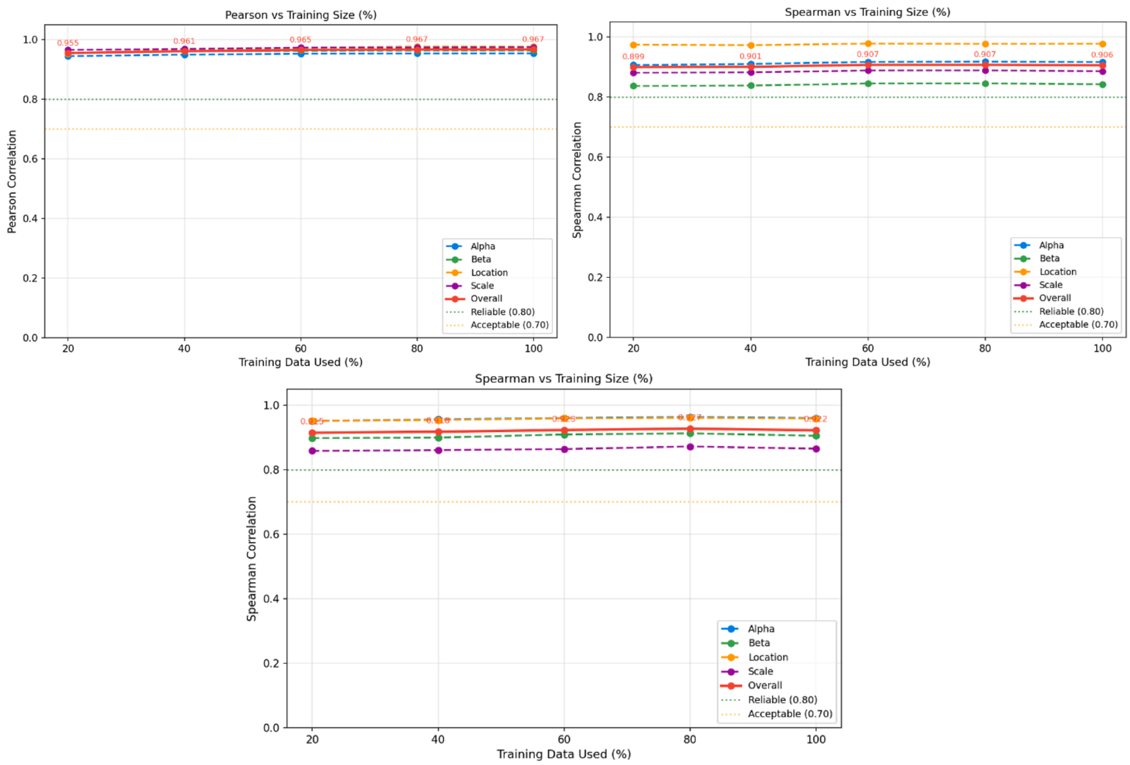

Fig. 17. Learning curves for 20 %–40 %–60 %–80 %–100 % of the dataset. (Upper left) Pearson correlation for X-axis training; (upper right) Spearman's rank correlation for Y-axis; (lower) Sperman's rank correlation for Z-axis.

Table 7
Meteorological and technical conditions of the comparative experiment.

Parameter	Value
Location (latitude [deg], longitude [deg.])	50.9, 6.7
Date/time	13th May 2025, 14:17:48
DNI (W/m^2)	948.9
Focused heliostats	576
Average mirror reflectivity (-)	0.813
Peak measured flux (kW/m^2)	571.0
Cavity height (m)	40.8
Cavity diameter (m)	1.5
Cavity depth (m)	2.6
Receiver tilt angle (deg)	15
Image resolution (px x px)	816 × 703
Camera distance (m)	88.4

5. Conclusion and outlook

This work presents a universal, virtual, and non-disruptive flux density prediction method for concentrating solar power central receivers, based on the near-real-time correction of Monte-Carlo ray-tracing simulations through for convolutional neural networks working in parallel. This procedure corrects both the power and directional components of simulated rays at the receiver aperture plane, achieving correction accuracies up to 95 % and ray direction prediction correlations exceeding 80 % across a wide range of solar irradiance levels and sun positions. A statistical parametric reduction compressing over 100 million ray direction parameters into 2028 values enables near-real-time inference on standard computing hardware. The corrected rays are subsequently projected into three-dimensional receiver geometries, yielding an integrated cavity flux prediction accuracy of 80.9 % against experimental measurements. These results demonstrate that the method successfully bridges the accuracy gap between computationally efficient ray-tracing simulations and disruptive camera-based measurement techniques, while extending flux prediction capability to cavity receivers where direct measurement is physically infeasible.

However, this method still presents some limitations that must be acknowledged. First, ray direction sources remain inherently tied to ray-tracing simulations, meaning that systematic inaccuracies in

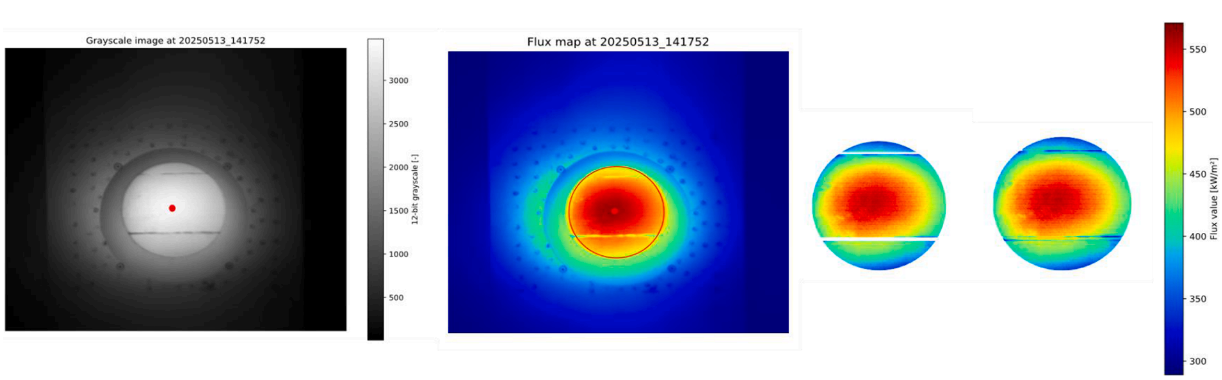

Fig. 18. Flux density measurement obtention process from a 12-px greyscale image.

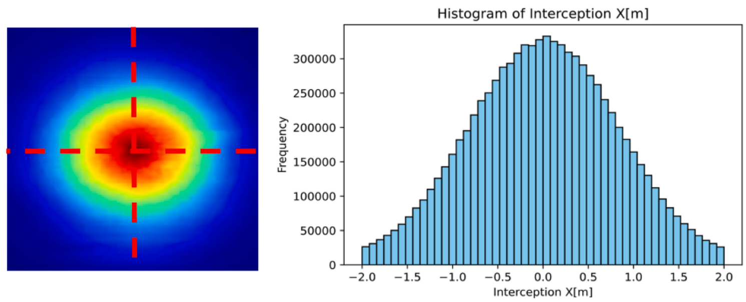

Fig. 19. Aperture-plane outputs after CNN correction: (left) corrected flux density map at the receiver aperture plane on 13 May 2025 at 14:17 UTC with 576 focused heliostats and DNI = 948.9 W/m^2; (center) histogram of X-axis ray interception positions showing the spatial distribution across the aperture; (right) corresponding Y-axis interception distribution prior to 3D projection.

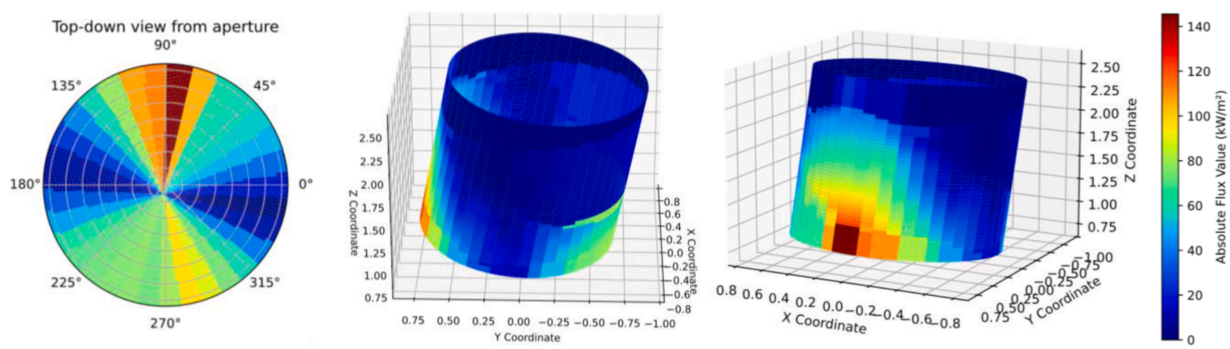

Fig. 20. Simulated three-dimensional flux density distribution inside the HEHTRES cylindrical cavity receiver (diameter 1.5 m, depth 2.6 m) resulting from projection of CNN-corrected rays: (left) top-down view showing receiver bottom end cap perpendicularly; (center) Y-Z plane view; (right) isometric perspective. Integrated simulated power: 577.2 kW vs. 714.9 kW measured (80.9 % accuracy).

circumsolar ratio modelling and similar optical phenomena are not entirely eliminated, though the neural networks partially compensate through learned extrapolation. Second, tracking error effects on ray interception positions at the aperture plane are not fully corrected, although their practical impact is mitigated by the small magnitude of typical tracking errors and the stochastic ray projection strategy applied during cavity simulation. Third, physical phenomena occurring inside the cavity — such as retroreflection and internal scattering — cannot be corrected through aperture-plane measurements alone, and quantification of these effects through spillage monitoring or aperture-plane simulation strategies remains an open challenge. Finally, the Z-axis ray direction scale parameter prediction requires further improvement for edge irradiation conditions, such as very low sun elevation angles, though performance under typical operating conditions is satisfactory.

From an operational and responsible deployment perspective, several considerations are relevant for industrial adoption of this method. Model performance may drift over time as heliostat surfaces degrade, mirror reflectivity decreases, or atmospheric conditions shift beyond the training distribution, necessitating periodic revalidation against experimental flux measurements. The computational cost of the training phase is a one-time investment, while near-real-time inference makes continuous on-duty monitoring viable without disrupting plant operation. Integration with existing plant control and supervisory systems would benefit from standardized data interfaces and operator-facing uncertainty estimates, which probabilistic neural network architectures can provide in principle. Although the method is validated at a single Central European facility, the training pipeline is site-agnostic and requires only a ray-tracing plant model and a flux measurement dataset, both of which are obtainable at any concentrating solar power installation worldwide. Validation across climatically diverse sites, including high-irradiance desert environments representative of the majority of global concentrating solar power capacity, is planned as part of ongoing international research efforts.

CRediT authorship contribution statement

Sergio Díaz-Alonso: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Funding acquisition, Data curation, Conceptualization. **Christian Raeder:** Writing – review & editing, Project administration. **Allen Charly:** Writing – review & editing, Investigation. **Bernhard Hoffschmidt:** Resources.

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

Authors acknowledge the funding of European Union for this project (Horizon MSCA doctoral network, project number 101072537).

Appendix I

Fig. A1, Fig. A2

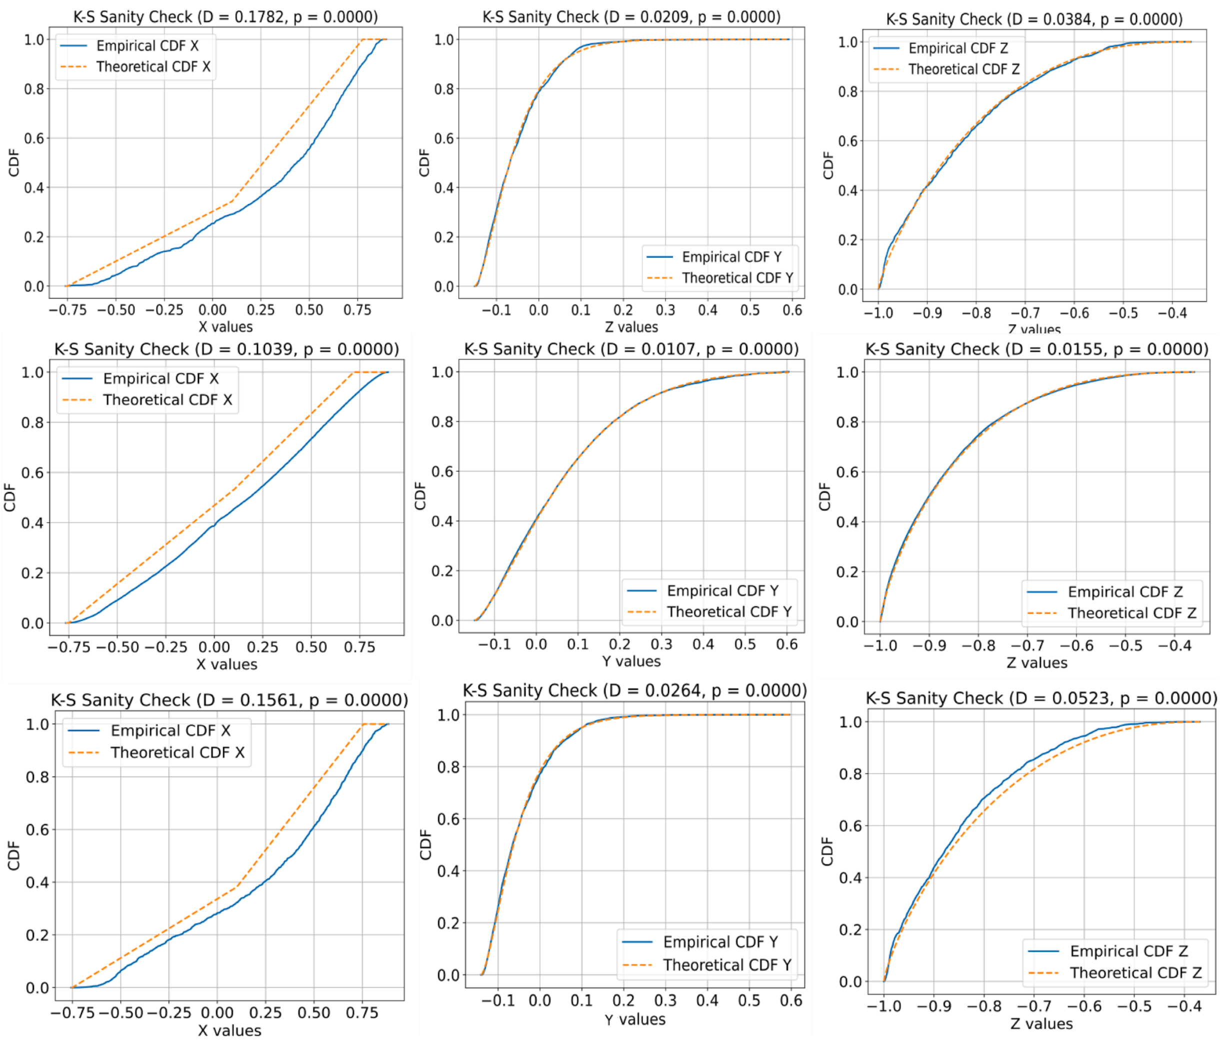

Fig. A1. Sanity check for Kolmogorov-Smirnov test in 3 different mesh elements and in their 3 axes: 22, located at the ROI's lower left corner (upper row); 85, located at the ROI's centre (2nd row); and 145, located at the ROI's upper right corner (lower row). These 9 graphs represent the systematic comparison among empirical and theorical CDFs in order to compute the greatest difference: the Kolmogorov-Smirnov D-value.

Appendix II

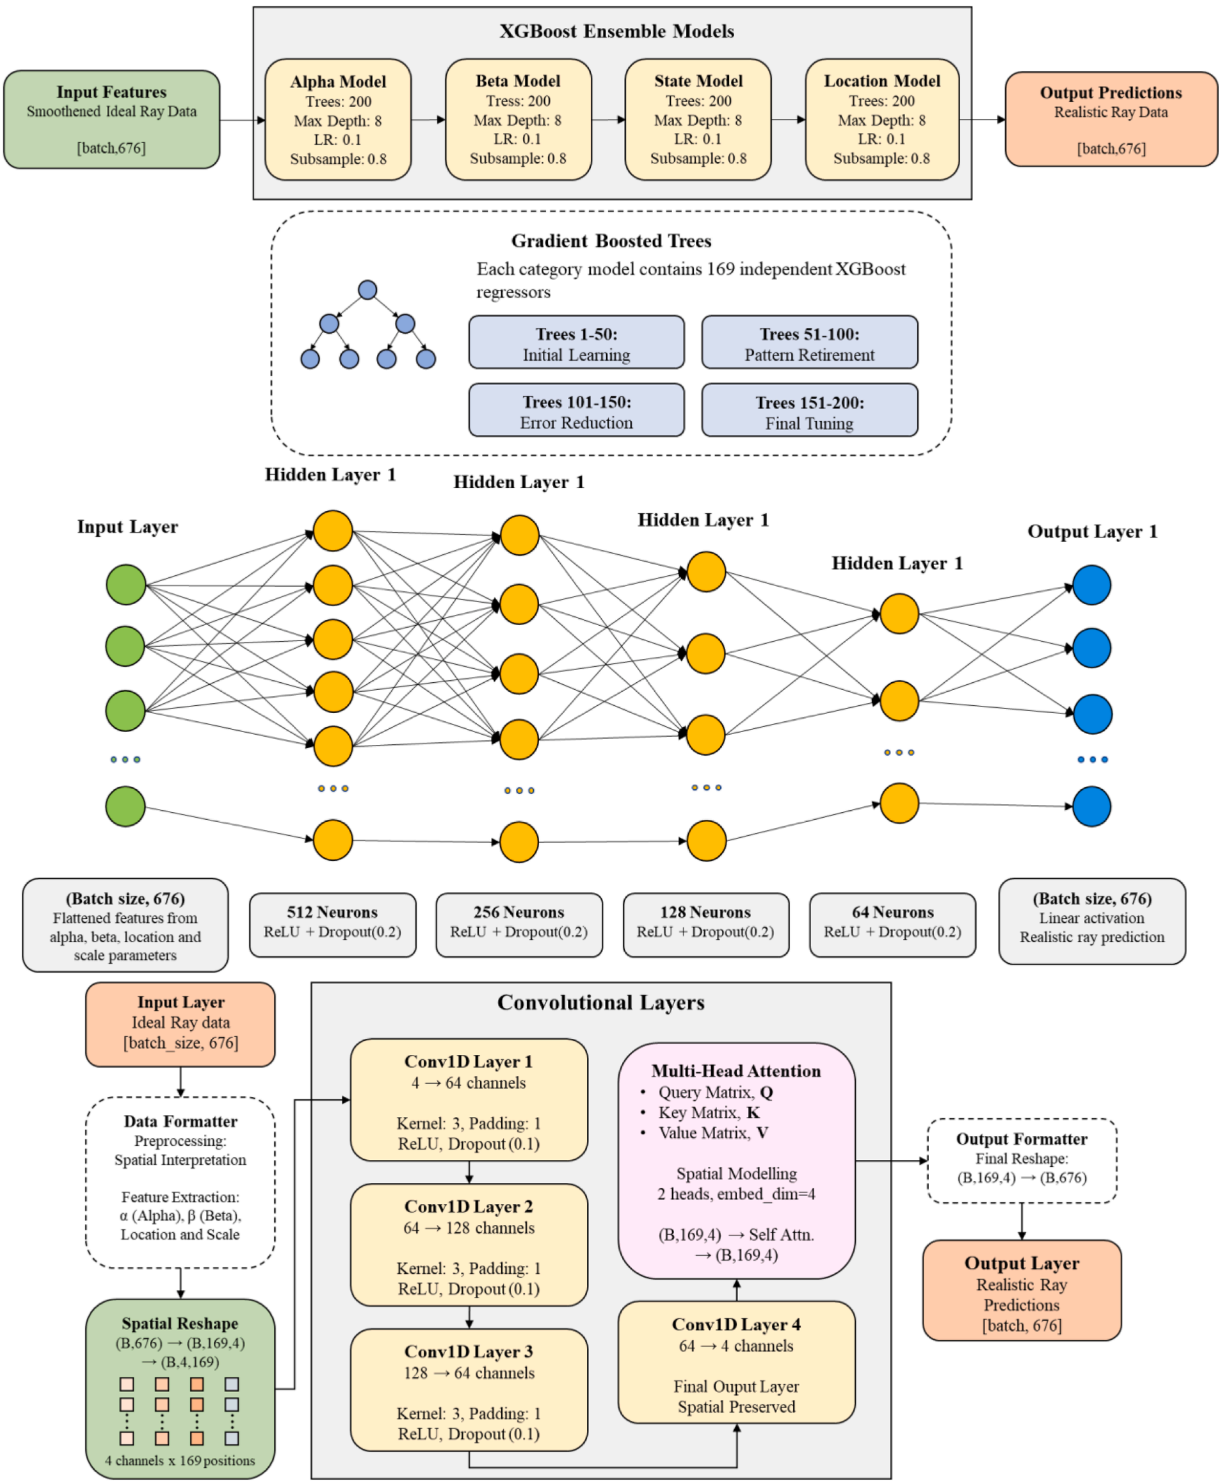

Fig. A2. Summary of the applied machine learning models for ray direction correction: from standard decision trees to self-tailored deep-learning models. I) XGBoost, II) FFNN, III) CNN spatial ray predictor, IV) Auto-Encoder, V) Bayesian Auto-Encoder, VI) WaveNet.

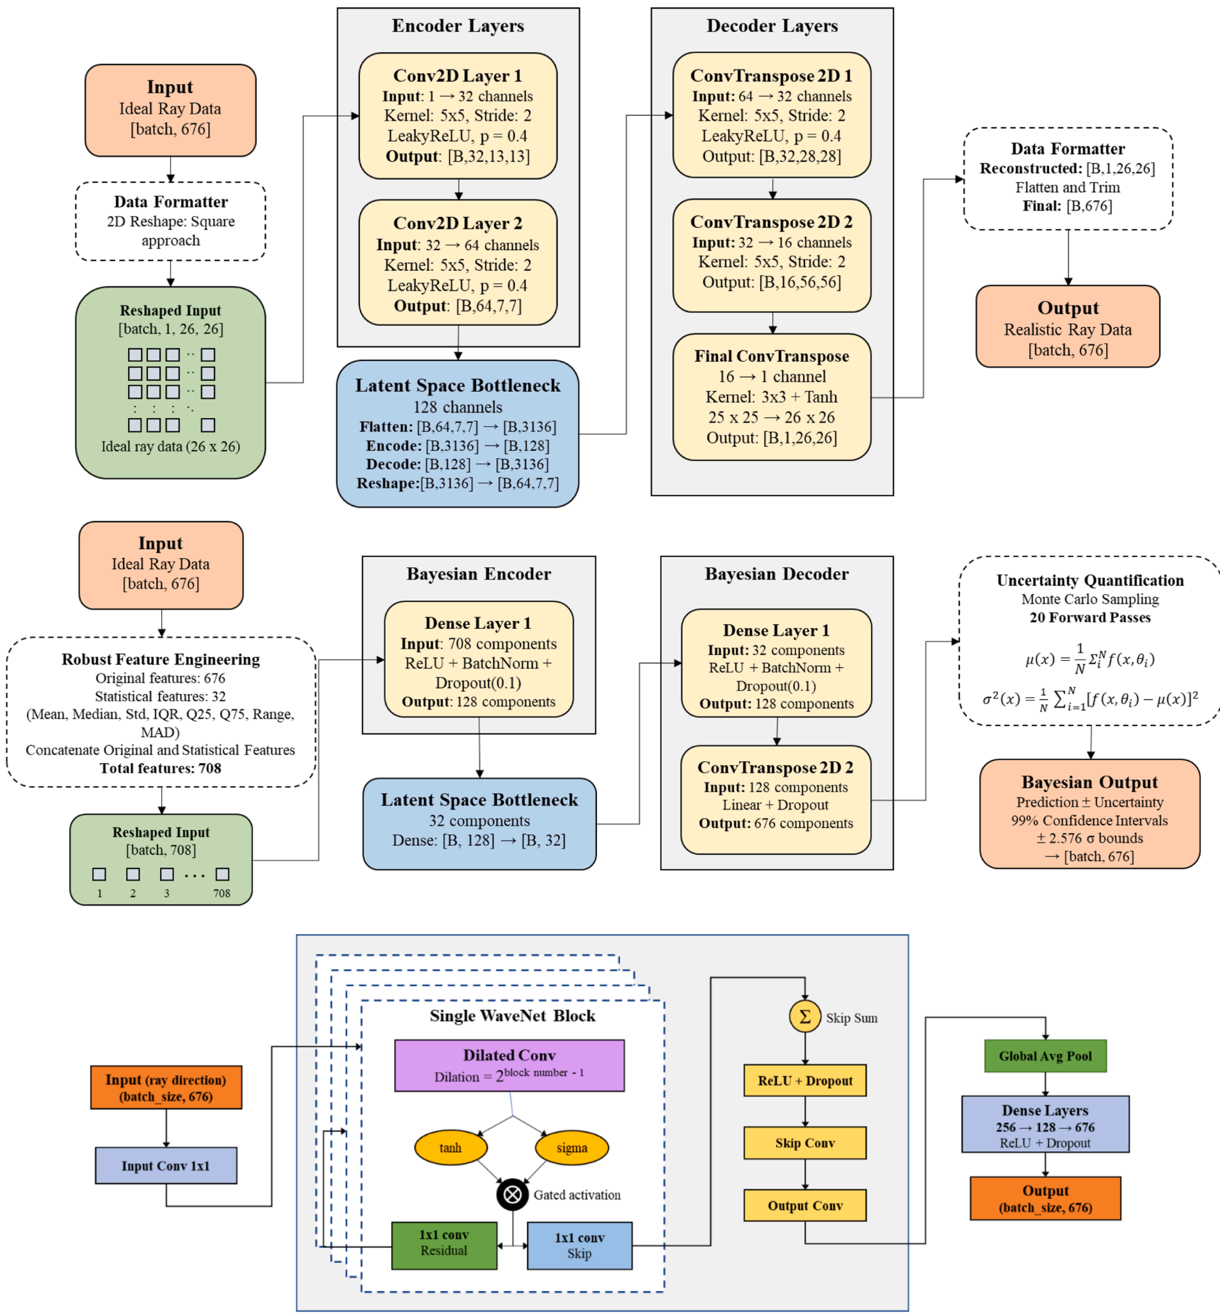

Fig. A2. (*continued*).

Data availability

Software can be made available on request with individual study of the case. It will be fully released after physical-world neural network training. Semi-controlled conditions dataset made available on Zenodo repository: 10.5281/zenodo.13941464 and 10.5281/zenodo.15740381.

References

[1] Famiglietti A, Lecuona A. Energetic and economic analysis of a novel concentrating solar air heater using linear Fresnel collector for industrial process heat. Renew Energy 2024;236:121362. https://doi.org/10.1016/j.renene.2024.121362.

[2] Collado FJ, Guallar J. A review of optimized design layouts for solar power tower plants with campo code. Renew Sustain Energy Rev 2013;20:142–54. https://doi.org/10.1016/j.rser.2012.11.076.

[3] Schappi R, et al. Drop-in fuels from sunlight and air. Nature 2022;601(7891):63–8. https://doi.org/10.1038/s41586-021-04174-y.

[4] Schöniger F, Thonig R, Resch G, Lilliestam J. Making the sun shine at night: comparing the cost of dispatchable concentrating solar power and photovoltaics with storage. Energy Sources B: Econ Plan Policy 2021;16(1):55–74. https://doi.org/10.1080/15567249.2020.1843565.

[5] Martín-Alcántara A, Serrano-Aguilera JJ, Sánchez Sánchez de Puerta A, Aranda-Hidalgo JL. Decarbonizing an industrial process through a combined, high-temperature CSP and sensible heat storage. J Energy Storage 2024;97:112727. https://doi.org/10.1016/j.est.2024.112727.

[6] Sumayli H, et al. Integrated CSP-PV hybrid solar power plant for two cities in Saudi Arabia. Case Stud Therm Eng 2023;44. https://doi.org/10.1016/j.csite.2023.102835.

[7] Garcia P, Ferriere A, Bezian J-J. Codes for solar flux calculation dedicated to central receiver system applications: a comparative review. Sol Energy 2008;82(3):189–97. https://doi.org/10.1016/j.solener.2007.08.004.

[8] NREL. SolTRACE: a new optical modeling tool for concentrating solar optics. presented at the. In: ISEC 2003: International Solar Energy Conference; 2003.

[9] Blanco MJ, Mutuberria A, Garcia P, Gastesi R, Martin V. Preliminary validation of Tonatiuh. presented at the. In: SolarPACES 2009 International Conference; 2009.

[10] Belhomme B, Pitz-Paal R, Schwarzbözl P, Ulmer S. A new fast ray tracing tool for high-precision simulation of heliostat fields. J Sol Energy Eng 2009;131. https://doi.org/10.1115/1.3139139.

[11] Röger M, Herrmann P, Ulmer S, Ebert M, Prahl C, Göhring F. Techniques to measure solar flux density distribution on large-scale receivers. J Sol Energy Eng 2014;136(3). https://doi.org/10.1115/1.4027261.

[12] Osuna R, Morillo R, Jiménez JM, Fernández-Quero V. Control and operation strategies in PS10 Solar plant. In: 13th SolarPACES; 2006. p. 1–5. Jun. 20–23.

[13] King DL, Arvizu DE. Heliostat characterization at the Central Receiver Test Facility. J Sol Energy Eng 1981;103(2):82–8. https://doi.org/10.1115/1.3266229.

[14] Ballestrín J, Monterreal R. Hybrid heat flux measurement system for solar central receiver evaluation. Energy 2004;29(5–6):915–24. https://doi.org/10.1016/s0360-5442(03)00196-8.

[15] Elsayed MM, Fathalah KA, Al-Rabghi OM. Measurements of solar flux density distribution on a plane receiver due to a flat heliostat. Sol Energy 1995;54(6):403–11. https://doi.org/10.1016/0038-092X(95)00010-O.

[16] Yogev O, Gleckman P, Rozler M. High-heat solar flux scanner. In: SolarPACES 2009; 2009. Sept. 15–18.

[17] Pacheco JE, Houser RM, Neumann A. Concepts to measure flux and temperature for external central receivers. presented at the. In: Joint Solar Engineering Conference; 1994. 1993–12–01.

[18] Offergeld M, Röger M, Stadler H, Gorzalka P, Hoffschmidt B. Flux density measurement for industrial-scale solar power towers using the reflection off the absorber. presented at the. In: SOLARPACES 2018: International Conference on Concentrating Solar Power and Chemical Energy Systems; 2019.

[19] G. von Tobel, C. Schelders, and M. Real, "Concentrated solar-flux measurements at the IEA-SSPS solar-central-receiver power plant, Tabernas - Almeria (Spain). Final report. Technical report No. 2/82," 1982.

[20] Thelen M, Raeder C, Willsch C, Dibowski G. A high-resolution optical measurement system for rapid acquisition of radiation flux density maps. In: SolarPACES 20161850. AIP Publishing; 2017. https://doi.org/10.1063/1.4984534. 2016.

[21] Ulmer S, Lüpfert E, Pfänder M, Buck R. Calibration corrections of solar tower flux density measurements. Energy 2004;29(5–6):925–33. https://doi.org/10.1016/s0360-5442(03)00197-x.

[22] Rodríguez-Sánchez MR, Sánchez-González A, Santana D. Field-receiver model validation against Solar two tests. Renew Sustain Energy Rev 2019;110:43–52. https://doi.org/10.1016/j.rser.2019.04.054.

[23] P. Schwarzbözl, R. Pitz-Paal, and M. Schmitz, "Visual HFLCAL - a software tool for layout and optimisation of heliostat fields," 2009.

[24] Roccia JP, et al. SOLFAST, a Ray-Tracing Monte-Carlo software for solar concentrating facilities. J Phys: Conf Ser 2012;369(1):012029. https://doi.org/10.1088/1742-6596/369/1/012029.

[25] Wendelin TJ. SolTRACE: a new optical modeling tool for concentrating solar optics. Sol Energy 2003:253–60.

[26] P.L. Leary and J.D. Hankins, "User's guide for MIRVAL: a computer code for comparing designs of heliostat-receiver optics for central receiver solar power plants," United States, 01 1979. [Online]. Available: https://www.osti.gov/biblio/6371450 https://www.osti.gov/servlets/purl/6371450.

[27] Ulmer S, März T, Prahl C, Reinalter W, Belhomme B. Automated high resolution measurement of heliostat slope errors. Sol Energy 2011;85(4):681–7. https://doi.org/10.1016/j.solener.2010.01.010.

[28] M. Röger, C. Prahl, and S. Ulmer, *Fast determination of heliostat shape and orientation by edge detection and photogrammetry*. 2008.

[29] Collado FJ. One-point fitting of the flux density produced by a heliostat. Sol Energy 2010;84(4):673–84. https://doi.org/10.1016/j.solener.2010.01.019.

[30] Sánchez-González A, Grange B, Caliot C. Computation of canting errors in heliostats by flux map fitting: experimental assessment. Opt Express 2020;28(26):39868–89. https://doi.org/10.1364/OE.412116.

[31] Zhu R, Ni D, Yang T, Yang J, Chen J, Xiao G. Heliostat field aiming strategy optimization with post-installation calibration. Appl Therm Eng 2022;202. https://doi.org/10.1016/j.applthermaleng.2021.117720.

[32] Rodríguez-Sánchez MR, Leray C, Toutant A, Ferriere A, Olalde G. Development of a new method to estimate the incident solar flux on central receivers from deteriorated heliostats. Renew Energy 2019;130:182–90. https://doi.org/10.1016/j.renene.2018.06.056.

[33] Stone KW. Automatic heliostat track alignment method. Google Patents; 1986.

[34] Kuhl M, Pargmann M, Cherti M, Jitsev J, Quinto DMaldonado, Pitz-Paal R. In-situ UNet-based heliostat beam characterization method for precise flux calculation using the camera-target method. Sol Energy 2024;279:112811. https://doi.org/10.1016/j.solener.2024.112811.

[35] Kuhl M, Pargmann M, Cherti M, Jitsev J, Quinto DMaldonado, Pitz-Paal R. Flux density distribution forecasting in concentrated solar tower plants: a data-driven approach. Sol Energy 2024;282:112894. https://doi.org/10.1016/j.solener.2024.112894.

[36] Pargmann M, Ebert J, Götz M, Maldonado Quinto D, Pitz-Paal R, Kesselheim S. Automatic heliostat learning for in situ concentrating solar power plant metrology with differentiable ray tracing. Nat Commun 2024;15(1):6997. https://doi.org/10.1038/s41467-024-51019-z.

[37] Lewen J, Pargmann M, Cherti M, Jitsev J, Pitz-Paal R, Quinto DM. Inverse deep learning raytracing for heliostat surface prediction. Sol Energy 2025;289:113312. https://doi.org/10.1016/j.solener.2025.113312.

[38] Zhou W, Lin X, Lei J, Yu L, Hwang JN. MFFENet: multiscale feature fusion and enhancement network for RGB–Thermal urban road scene parsing. IEEE Trans Multimed 2022;24:2526–38. https://doi.org/10.1109/TMM.2021.3086618.

[39] Dong S, Zhou W, Xu C, Yan W. EGFNet: edge-aware guidance fusion network for RGB–Thermal urban scene parsing. IEEE Trans Intell Transp Syst 2024;25(1):657–69. https://doi.org/10.1109/TITS.2023.3306368.

[40] Zhou W, Yang E, Lei J, Wan J, Yu L. PGDENet: progressive guided fusion and depth enhancement network for RGB-D indoor scene parsing. Trans Multi 2023;25:3483–94. https://doi.org/10.1109/tmm.2022.3161852.

[41] Zhou W, Guo Q, Lei J, Yu L, Hwang JN. ECFFNet: effective and consistent feature fusion network for RGB-T salient object detection. IEEE Trans Circuits Syst Video Technol 2022;32(3):1224–35. https://doi.org/10.1109/TCSVT.2021.3077058.

[42] Wang A, Chen H, Lin Z, Han J, Ding G. Lsnet: see large, focus small. In: Proceedings of the computer vision and pattern recognition conference; 2025. p. 9718–29.

[43] Zhou W, Zhang H, Yan W, Lin W. MMSMCNet: modal memory sharing and morphological complementary networks for RGB-T urban scene semantic segmentation. IEEE Trans Cir Sys Video Technol 2023;33(12):7096–108. https://doi.org/10.1109/tcsvt.2023.3275314.

[44] Zhou W, Guo Q, Lei J, Yu L, Hwang J-N. IRFR-net: interactive recursive feature-reshaping network for detecting salient objects in RGB-D images. IEEE Trans Neural Netw Learn Syst 2021.

[45] Zhou W, Sun F, Jiang Q, Cong R, Hwang JN. WaveNet: wavelet network with knowledge distillation for RGB-T salient object detection. IEEE Trans Image Process 2023;32:3027–39. https://doi.org/10.1109/TIP.2023.3275538.

[46] Zhou W, Zhu Y, Lei J, Wan J, Yu L. CCAFNet: crossflow and Cross-scale adaptive fusion network for detecting salient objects in RGB-D images. IEEE Trans Multimed 2022;24:2192–204. https://doi.org/10.1109/TMM.2021.3077767.

[47] El-kenawy E-SM, et al. Feature selection in wind speed forecasting systems based on meta-heuristic optimization. PLOS ONE 2023;18(2):e0278491. https://doi.org/10.1371/journal.pone.0278491.

[48] Takieldeen AE, El-Kenawy E-SM, Hadwan M, Zaki RM. Dipper throated optimization algorithm for unconstrained function and feature selection. Comput Mater Contin 2022;72:1465–81.

[49] Khafaga DS, et al. An Al-Biruni Earth radius optimization-based deep convolutional neural network for classifying monkeypox disease. Diagnostics 2022;12(11):2892 [Online]. Available: https://www.mdpi.com/2075-4418/12/11/2892.

[50] Alhussan AA, et al. A binary waterwheel plant optimization algorithm for feature selection. IEEE Access 2023;11:94227–51. https://doi.org/10.1109/ACCESS.2023.3312022.

[51] El-Kenawy E-SM, Ibrahim A, Alhussan AA, Khafaga DS, Ahmed AEM, Eid MM. Smart City electricity load forecasting using Greylag Goose Optimization-enhanced time series analysis. Arab J Sci Eng 2025. https://doi.org/10.1007/s13369-025-10647-3. 2025/09/25.

[52] Sergio DA, Christian R, Bernhard H. Concentrating solar power (CSP) plant data-driven digital twin: a novel method for flux density prediction. Results Eng 2025;28:107096. https://doi.org/10.1016/j.rineng.2025.107096.

[53] Dai Y, Wang Y, Leng M, Yang X, Zhou Q. LOWESS smoothing and Random Forest based GRU model: a short-term photovoltaic power generation forecasting method. Energy 2022;256:124661.

[54] Paris S, Kornprobst P, Tumblin J, Durand F. Bilateral filtering: theory and applications. Found Trends® Comput Graph Vis 2009;4(1):1–73.

[55] Breaz N, Aldea M. On the smoothing spline regression models. Acta Univ Apulensis Math-Inform 2008;15:33–51.
[56] Centofanti F, Lepore A, Menafoglio A, Palumbo B, Vantini S. Adaptive smoothing spline estimator for the function-on-function linear regression model. Comput Stat 2023;38(1):191–216.
[57] Hunter JS. The exponentially weighted moving average. J Qual Technol 1986;18 (4):203–10.
[58] Chen T, Guestrin C. XGBoost: A Scalable Tree Boosting System. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. New York, NY, USA; 2016. p. 785–94. https://doi.org/10.1145/2939672.2939785.
[59] A. Van Den Oord et al., "Wavenet: a generative model for raw audio," *arXiv preprint arXiv:1609.03499,* vol. 12, p. 1, 2016.
[60] Díaz Alonso S. Concentrating solar power (CSP) plants AI-training dataset for flux density measurements. Zenodo 2025. https://doi.org/10.5281/zenodo.15173386.
[61] Díaz Alonso S, Raeder C. Ray-level dataset for deep-learning direction correction in Concentrating Solar Power (CSP) plants. Zenodo 2025. https://doi.org/10.5281/zenodo.15740381.