

Inverse Deep Learning Raytracing for heliostat surface prediction

Jan Lewen^{a,b,*}, Max Pargmann^a, Mehdi Cherti^b, Jenia Jitsev^b, Robert Pitz-Paal^a, Daniel Maldonado Quinto^a

^a German Aerospace Center (DLR), Institute of Solar Research, Linder Höhe, D-51147 Köln, Germany

^b Research Center Jülich, Jülich Supercomputing Centre, Wilhelm-Johnen-Straße, 52428 Jülich, Germany

ARTICLE INFO

Keywords:

CSP
Deep learning
Generative model
StyleGAN
Sim2real transfer
Heliostat surface
Mirror error
Raytracing
Applied artificial intelligence

ABSTRACT

Concentrating Solar Power (CSP) plants play a crucial role in the global transition toward sustainable energy. A key factor in ensuring the safe and efficient operation of CSP plants is the distribution of concentrated flux density on the receiver. However, the non-ideal flux density generated by individual heliostats can undermine the safety and efficiency of the power plant. The flux density from each heliostat is influenced by its precise surface, which includes factors such as canting and mirror errors. Accurately measuring these surfaces for a large number of heliostats in operation is a formidable challenge. Consequently, control systems often rely on the assumption of ideal surface conditions, which compromises both safety and operational efficiency. In this study, we introduce inverse Deep Learning Raytracing (*iDLR*), an innovative method designed to predict heliostat surfaces based solely on target images obtained during heliostat calibration. Our simulation-based investigation reveals that the flux density distribution of a single heliostat contains sufficient information to enable deep learning models to accurately predict the underlying surface with deflectometry-like precision in most cases, achieving a median Mean Absolute Error of approximately 0.14 mm). When integrating the *iDLR* surface predictions into a ray-tracing environment to compute flux densities, our method achieves an accuracy of 92%, surpassing the performance of the ideal heliostat assumption by 25%. Additionally, we assess the limitations of this method, particularly in relation to surface prediction accuracy and resultant flux density predictions. Furthermore, we present an innovative and efficient heliostat surface model based on NURBS. This approach achieves a highly compact representation, requiring only 256 parameters to define the surface—a reduction of 99.97% in the amount of parameter and a 99.91% in memory usage. This efficient model enables resource-effective deep learning for heliostat surface predictions, positioning it as a promising state-of-the-art solution for heliostat surface parameterization. Our findings demonstrate that *iDLR* has significant potential to optimize CSP plant operations, enhancing overall efficiency and increasing the energy output of power plants.

1. Introduction

In order to enhance the economic viability of solar tower power plants, it is crucial to enhance their efficiency. A viable approach for achieving this is through the optimization of the flux density distribution on the receiver. In commercial plants this can be done by optimal aim point control. e.g. using the ant-colony optimization meta-heuristic [1,2] or deep learning models [3] to calculate a flux density for each heliostat using raytracing and optimize in a next step the flux density distribution on the receiver as a superposition of those. Moreover, using knowledge about each heliostat's flux density distribution and its position on the receiver, even higher efficiency gains can be achieved. Even simple assumptions can increase the energy output by up to 20% [4]. Commercial power plant control systems commonly use the Vant–Hull Algorithm [5,6] for aim point optimization. However, in

most cases, flux density predictions are made either through raytracing of an ideal heliostat or by adopting simplified flux density assumptions. This can be explained by the fact that exact information about the heliostats' surfaces is not easily accessible. Inherent mirror errors cause the individuality of each heliostat's flux density distribution.

The mirror error consists of roughness, slope, and canting errors. The roughness arises from sub-micrometer flaws on the reflective surface. Slope error measures the deviations of the mirror surface from its ideal shape. Canting error, on the other hand, reflects misalignment among mirror facets. Among these errors, slope and canting are particularly crucial for the flux density distribution [7–9]. A precise heliostat representation, including its mirror error, can be loaded in a raytracing environment to predict the flux density of the heliostat under certain environmental conditions accurately [10–12].

* Corresponding author at: German Aerospace Center (DLR), Institute of Solar Research, Linder Höhe, D-51147 Köln, Germany.
E-mail address: jan.lewen@dlr.de (J. Lewen).

List of Abbreviations

CSP	Concentrating Solar Power
STJ	Solar Tower Jülich
DLR	Deutsches Zentrum für Luft- und Raumfahrt
iDLR	inverse Deep Learning Raytracing
GAN	Generative Adversarial Network
GPU	Graphics Processing Unit
NURBS	Non-Uniform Rational B-Spline
MAE	Mean Absolute Error
Q1/3	Quartile 1/3
IQR	Interquartile Range
CAD	Computer Aided Design
ACC	Accuracy
SOTA	State of the Art
CSR	Circumsolar Ratio

The most prominent and state-of-the-art (SOTA) method for obtaining the surface shape of heliostats is the deflectometry measurement [11,13]. This method involves capturing camera images of stripe patterns with diverse frequencies projected onto a Lambertian target and subsequently reflected from the heliostat. However, this approach faces obstacles in the heliostat field because of factors such as dew, wind, and the fact that the measurement has to be conducted at night with long exposure times. Alternatively, some methods use a laser [14,15], while others employ photogrammetry [16–20] or flux mapping [21]. However, up for today these measurements are unreliable or cost intensive. For a complete review, the reader is referred to [9].

As a result, considerable efforts are made to extract additional information about each heliostat from more readily available data sources. Zhu et al. [4] have developed a post-installation calibration procedure to find four geometrical parameter per facet by optimization, using only target images of the corresponding heliostat. Even this simple heliostat model achieved the above mentioned efficiency gains. These results underline the potential of including an accurate flux density prediction in the aim point optimization. A key drawback of this method is the simple and specific heliostat model, assuming only canting and focusing error. For example, the heliostats at Solar Tower Jülich (STJ), Germany are not focused and the important surface features that influences the flux density besides the canting error, are the waviness and the bending of the facets at the edges and corners. Those surface features are not predictable by the approach of Zhu et al. [4]. The latest method for enhancing the precision of flux density predictions is pioneered by Martínez-Hernández et al. [21]. They have devised a technique to reconstruct the surface of heliostats using focal spots during daylight hours. However, this reconstruction capability decreases when the distance between the heliostat and target exceeds 6 m. To address this limitation, additional measures such as implementing a moving target are necessary.

Pargmann et al. [12,22] introduced a versatile approach for heliostat calibration and surface reconstruction using differentiable raytracing. This optimization-based method leverages target images for both tasks, providing a cost-effective and fully automated solution that requires no additional hardware. However, the approach faces challenges in accurately predicting heliostat surfaces due to the inherently underdetermined and ill-posed nature of the problem. This highlights the need for a reliable and cost-effective method capable of predicting heliostat surfaces during routine power plant operations.

In addition to their differentiable raytracing formulation, Pargmann et al. [12] introduced a NURBS-based heliostat model. This model enables the optimization of NURBS parameters directly from target images, offering a more compact and efficient representation compared

to traditional normal vector point cloud representations [11]. The primary advantage of NURBS parameterization is its low dimensionality, which makes it particularly suitable for machine learning applications. However, their study did not involve fitting NURBS surfaces to normal vectors derived from deflectometry measurements. Consequently, the resulting NURBS surfaces often yielded suboptimal heliostat surface predictions due to the previously discussed underdetermination. Unfortunately, the traditional normal vector point cloud representations, which can have up to one million parameters for heliostats at the STJ, make deep learning methods exceedingly challenging due to their significant computational demands and difficult trainability. Hence, it is still desirable to find a way to effectively parameterize heliostat surfaces with a low dimensional model, like a NURBS representation.

Generative deep learning models have emerged as a prominent area of research within computer science in recent years, characterized by groundbreaking advancements enabled by architectures such as StyleGAN [23], Large Language Models [24,25], and Diffusion Models [26]. The term “generative” refers to the capacity of these models to synthesize new instances that belong to the learned data domain. For instance, the StyleGAN model proposed by Karras et al. [23] can produce highly realistic synthetic images of human faces. In many scenarios, generating synthetic data necessitates the inclusion of additional input to control the output of the model, a framework known as conditional generative modeling. A famous example is the stable diffusion architecture [26], which takes natural language as input to generate corresponding images. Generative models acquire domain-specific knowledge during training, enabling them to give a good output even for underdetermined tasks. This property is particularly advantageous in addressing *Inverse Problems*. Solving an inverse problem involves deducing causal factors from observed effects, typically in the opposite direction of physics. These problems are inherently underdetermined and ill-posed. Deep learning has been extensively applied to solve inverse problems in fields such as medicine [27,28], physics [29,30], and seismology [31,32], owing to its ability to learn prior knowledge about the target domain. For an in-depth review, see [33].

In this work, we propose a conditional generative deep learning framework designed to predict heliostat surfaces from target images. The model is trained on synthetic data generated using a raytracing simulation, solving the inverse problem of raytracing. Our results demonstrate that the deep learning model achieves surface predictions with accuracy comparable to deflectometry methods, while offering significant cost advantages. Additionally, we decouple the NURBS-based heliostat representation of Pargmann et al. [12] from the differentiable raytracing process and introduce a new NURBS fitting routine allowing us to parameterize all deflectometry measurements efficiently as NURBS. The NURBS-based representation for heliostat surfaces achieves a drastic reduction in parameters and in memory usage compared to the conventional normal vector point cloud representation. These improvements establish the proposed representation as a likely new SOTA for heliostat surface modeling in the digital twins of CSP plants.

2. Workflow of inverse deep learning raytracing

The primary measurements executed at CSP plants are target images employed for calibration with the *Camera-Target Method* [34], offering a direct observation of a heliostat’s flux density. The flux density distribution is caused by the precise heliostat surface characteristics and known parameters such as sun and heliostat positions, and hence the flux density contains information about the surface. However, predicting heliostat surfaces based on their flux density is highly challenging because it is an underdetermined problem. Factors like canting of the facets and the overlap of reflected and scattered rays on deformed surfaces make it particularly challenging.

This study aims to use deep learning models to predict the heliostat surface using the heliostat’s target images as input. This approach

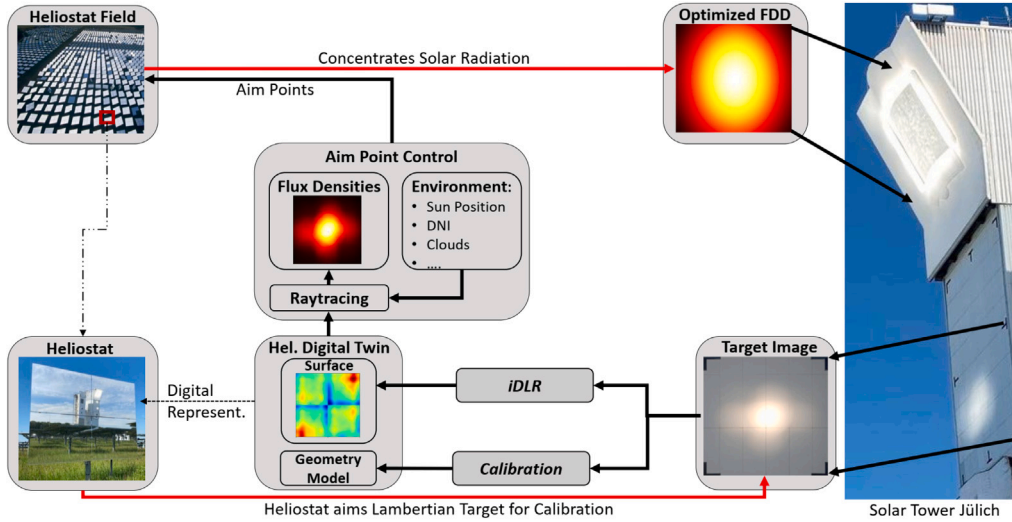


Fig. 1. Inverse Deep Learning raytracing (iDLR) embedded in the regular power plant operation. The heliostats are sequentially and fully automatically focused on a Lambertian target and target images are taken for the calibration (Camera-Target Method). Those target images contain information about the heliostat surface. Leveraging a deep learning model, this information is extracted and utilized to predict the heliostat surface, without the necessity of introducing new hardware or executing measurement which are not done yet during regular power plant operation. Consequently, the existing power plant systems and software can now operate seamlessly with an improved heliostat digital twin. This enhancement enables the power plant to operate more efficiently by achieving a more optimal flux density distribution (FDD) on the receiver.

is the inverse direction to physical raytracing and hence we term the method *inverse Deep Learning raytracing (iDLR)*. Fig. 1 shows the integration of iDLR in the routine power plant operation. Deep learning models have the capability to learn typical surface deformations constrained by material properties and mechanical construction of the mirror surface. This unique ability enables them to predict heliostat surfaces even within the underdetermined regime, leveraging the surface information learned during training. With this capability iDLR should perform better in this underdetermined problem than comparable physical algorithms like differentiable raytracing [12].

Before training a deep learning model, the question arises which data can be used. A purely real data approach would involve collecting real data pairs of heliostat surfaces and target images. However, there are significant drawbacks to this approach. First, the high costs associated with surface measurement for all existing methods. Second, the method should be applicable to commercial power plants that do not have a surface measurement setup. Finally, the substantial data demand for training generative deep learning models may pose challenges due to the limited number of heliostats in most fields.

To circumvent these limitations, our methodology aims to train on a semi-artificial dataset. The first step is to gather real surface measurements of the heliostat model used at the power plant where it is intended to implement iDLR. These measurements can be taken either at the power plant implementing iDLR or at any other power plant utilizing the same heliostat model, or even at the manufacturer's facility. Subsequently, a model is trained using a dataset comprising augmented real surface data and simulated flux density data obtained through raytracing. This model is presented in this work.

In the subsequent phase, the simulated model can be applied to real-world target images using deep learning sim-to-real transfer techniques, to bridge the gap between simulated flux densities and target images [35–38].

2.1. Comprehensive NURBS heliostat model

We employ a heliostat model proposed by Pargmann et al. [12], which integrates a traditional geometric model with an innovative spline approach to parameterize the reflecting surface of the heliostat. Current SOTA methods represent heliostat surfaces as a point cloud of normal vectors [11], providing precise representation but at the cost of a high number of parameters. For instance, the deflectometry measurement

used at the Solar Tower Jülich measures approximately 80,000 normal vectors per facet (1.6 m × 1.25 m; slope deviation with mean Root Mean Square of 1.36 mrad). This presents a challenge for training deep learning models due to the high computational resource demands and the potential decrease in model precision from the large number of free parameters. Pargmann et al. [12] employ a differentiable Non-Uniform Rational B-Spline (NURBS) within a differentiable raytracing routine (code published at [39]) as a trainable parameter to fit a heliostat surface to a given flux density. NURBS are the industrial standard to represent complex geometries with few parameters in Computer-Aided Design (CAD) [40]. A NURBS surface $S(u, v)$ is defined as:

$$S(u, v) = \frac{\sum_{i=0}^n \sum_{j=0}^m N_{i,p}(u) M_{j,q}(v) w_{i,j} \mathbf{P}_{i,j}}{\sum_{i=0}^n \sum_{j=0}^m N_{i,p}(u) M_{j,q}(v) w_{i,j}}, \quad (1)$$

with:

- n and m are the number of control point in the u - and v -directions, respectively.
- $\mathbf{P}_{i,j}$ are the control points that shape the surface.
- $w_{i,j}$ are the weights associated with each control point.
- $N_{i,p}(u)$ and $M_{j,q}(v)$ are the B-spline basis functions of degree p and q in the u - and v -directions, respectively.

NURBS elegantly enable the smooth and continuous parameterization of a wide range of geometries with minimal parameters, while preserving high accuracy. This is accomplished through the weighted superposition of piecewise-defined B-splines, which are constructed to meet continuity requirements along their boundaries.

Before fitting the NURBS parameter against the normal vectors from the deflectometry measurement through a gradient-based optimization process we adapt the NURBS slightly. First, to each facet of the heliostat, one NURBS surface is assigned. Hence, the heliostat with four facets is described by 4 NURBS. Then, the number of control points (n and m) must be chosen. Since a network architecture benefits from having the number of control points as a power of two, an investigation was conducted to determine the smallest power of two that yields sufficient results for surface parameterization. It was found that 8 control points are sufficient for the parameterization ($n = 8$ and $m = 8$). Next, we reduce the number of NURBS parameters by fixing the xy -position of each control point on a predefined grid:

$$\mathbf{P}_{i,j} = (P_{i,j}^x, P_{i,j}^y, P_{i,j}^z), \quad \text{where } P_{i,j}^x = \text{constant}, P_{i,j}^y = \text{constant}$$

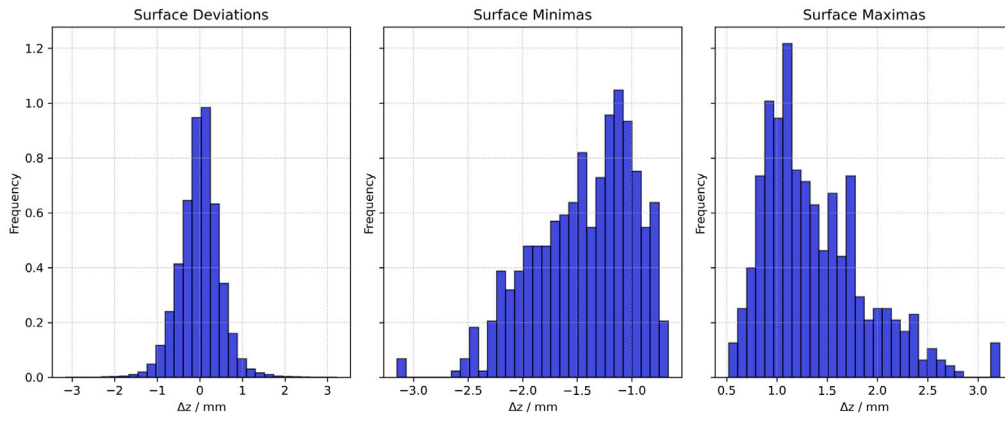


Fig. 2. Histograms representing the distribution of z-control point deviations in heliostat surface data, parameterized by NURBS. The left panel shows the overall distribution of surface deviations across all z-control points from all measured heliostat surfaces. The middle panel displays the distribution of the strongest negative deviations (surface minima) for each heliostat, while the right panel shows the strongest positive deviations (surface maxima). The views in the middle and right panel highlight outliers, providing clearer insight into the extremes of the heliostat surfaces.

This results in $P_{i,j}^z$ being the only control points that must be fitted. Finally, the NURBS weights $w_{i,j}$ are set to zero:

$$w_{i,j} = 0, \quad \forall i, j$$

This leads to only $4 \times 8 \times 8$ z-control points $P_{i,j}^z$ as NURBS parameter for the whole heliostat. The spline degree is set to three. The z-control points are then fitted using the differentiable formulation of Pargmann et al. [12] and gradient descent against the normal vector cloud of the deflectometry measurement. The optimization is done for 3000 epochs, by an Adam optimizer with betas of (0.9,0.999), eps of $1e-8$ and a weight decay of 0.1. The learning rate is $1e-5$, whereby an exponential decay learning rate scheduler is used with gamma of 0.995. As a result, the z-control points of the NURBS surface act as a comprehensive representation of the heliostat surface, reducing the number of free parameters from nearly 1 million to just 256—a reduction of 99.97%. This significantly enhances the efficiency of training machine learning models that involve heliostat surfaces. Furthermore, the NURBS representation is drastically more memory and computationally efficient than the normal point cloud. A single deflectometry measurement’s normal vectors occupy 8.6 MB, while the NURBS representation requires only 7 KB, a reduction of 99.91%.

2.2. Data acquisition and augmentation

The utilized dataset comprises stripe pattern deflectometry measurements conducted at the STJ, encompassing a total of 458 heliostat surfaces. A split into training data (428), validation (32) and testing data (32) was made, ensuring the exclusion of surfaces from validation and testing heliostats in the model training. For each heliostat, the z-control points of the NURBS were computed as detailed in Section 2.1 and act as surface representation. As there is a constant unknown offset for the surface of each facet, the mean value of the surface from each facet was set to zero. Fig. 2 shows the statistics of the data set. While the dataset comprises an unusually extensive collection of deflectometry data, it remains relatively small for effectively training a neural network. To address this limitation, two types of data augmentation were applied. Firstly, the heliostat surface was rotated by 180 degrees. Secondly, the weighted average between two randomly selected surfaces was calculated using $z_{augm} = \alpha * z_1 + (1 - \alpha) * z_2$ with $\alpha \in (0, 1)$. This approach ensures training on a diverse artificial dataset, avoiding the use of surfaces with unrealistic features given the physical constraints of the material and mechanical construction. In total, around 160,000 artificial heliostat surfaces were generated.

For testing and validation the 32 heliostats were placed in total 4544 times over the field, with a distance of up to 300 m to the tower

and with a maximum azimuth to the sides of 45 degrees (see Fig. 3). These heliostat positions were chosen as they match those at the STJ. Subsequently, environmental parameters such as sun position, heliostat aim point, and position were randomly drawn in accordance with the geographical position and heliostat field of the STJ for every simulated heliostat. For the flux density simulation, sun positions were uniformly sampled from the sun positions under which a target image was taken at the STJ for calibration (see Fig. 3). Whereby the sun shape of Buie et al. [41] was used with a Circumsolar Ratio (CSR) value of 5%. Finally, eight flux densities to different sun positions were simulated for each heliostat surface using the ray tracer developed by Pargmann et al. [12]. Those flux densities were centered around the center of mass, cropped onto a size of 4 m x 4 m and finally normalized that the sum of the discrete points of the flux density equals to 100.

Finally, the training, validation and test set are formed by the simulated flux densities, the sun position and the heliostat position as input parameter and the z-control points of the NURBS as the surface representation the model will predict.

2.3. Model and training

The nature of the problem suggests the adoption of an Encoder–Decoder architecture (see Fig. 4). The encoder is tasked with fusing two distinct input data streams: the eight flux densities, parameterized as an $8 \times 64 \times 64$ tensor, and scalar data containing the corresponding sun positions and the position of the heliostat in the field. The Encoder comprises five identical blocks. Each block begins with a convolutional layer that processes the output from the previous block or, in the case of the first block, the input flux densities (eight in total). Additionally, scalar input values are incorporated into the visual data stream using an affine transformation (A) followed by *weight demodulation*, as described by Karras et al. [42]. This process is repeated within the same block, and before downsampling, a residual connection is applied. The residual connection, based on the method introduced by He et al. [43], adds the output of the first layer to the output of the second layer. Throughout the encoder, the number of channels remains constant at 256. The final output of the encoder is a latent representation of the heliostat surface with dimensions 3×32 , using the extended $w+$ latent space of the StyleGAN2 architecture, as proposed by Richardson et al. [44]. For the Decoder, we utilize the StyleGAN2 Generator [42], which has demonstrated exceptional success in generating domain-specific data. The decoder begins with a starting block, which consists of a trainable tensor, followed by a convolutional layer that receive inputs through an affine transformation (A) and *weight demodulation* from the $w+$ latent space. This is followed by three style blocks. Each style block

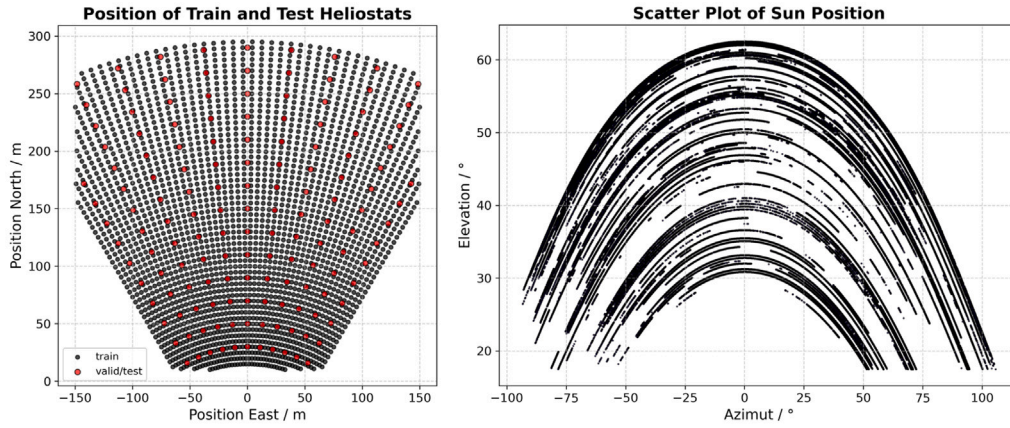


Fig. 3. Left Panel: Simulated heliostat field with black dots indicating training heliostat positions and red dots marking validation and testing heliostat positions. A total of 32 different heliostat surfaces were used for validation and testing and placed on every position marked by a red dot, resulting in a dataset of 4544 simulated heliostats. Right Panel: Scatter plot of sun positions during target image captures used for flux density simulations. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

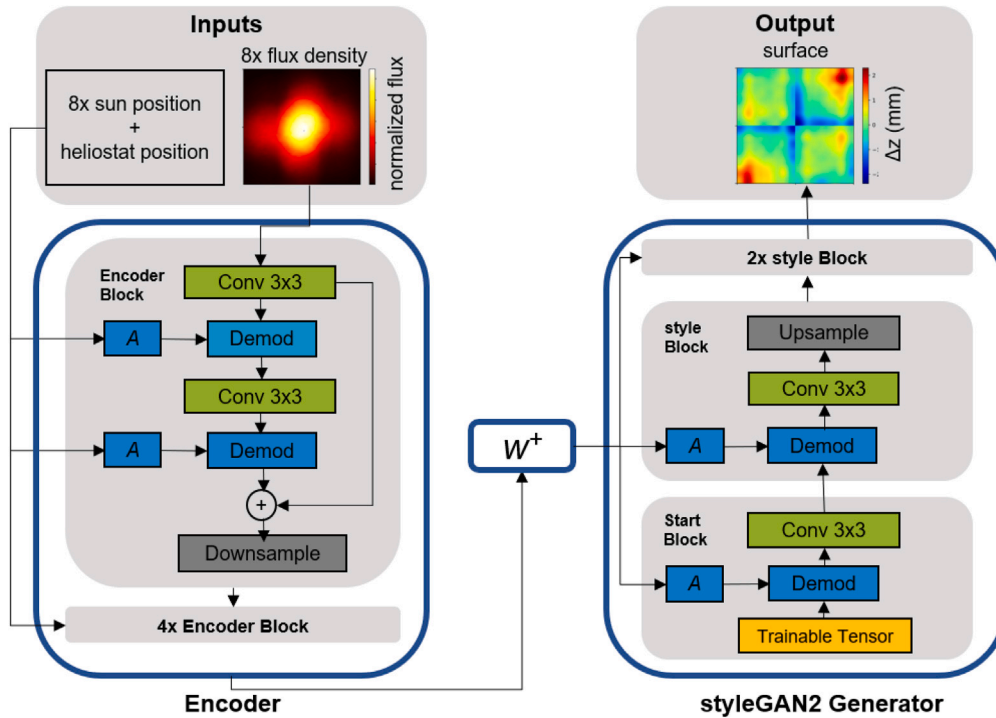


Fig. 4. The employed model integrates up to 8 flux densities and their corresponding sun and heliostat positions as inputs. These scalar inputs are processed in each Encoder Block, where they first undergo an affine transformation A , followed by weight demodulation, effectively mapping them into the image data stream. This stream is subsequently refined through convolutional neural layers. The Encoder consists of 5 sequential Encoder Blocks, which output the latent space representation $w+$. This $w+$ latent space serves as the input to the styleGAN2 generator. The generator begins with a trainable tensor in the Start Block and continues with 3 style Blocks. Each style Block receives inputs both from the preceding block and the $w+$ latent space, transformed through an affine transformation and weight demodulation. These are then fed into convolutional layers, followed by an upsampling step. The generator's output is the Cartesian representation of the z -control points, which parameterize the heliostat surface.

takes input from the preceding block as well as from the $w+$ latent space, again processed through an affine transformation (A) and weight demodulation. Within each style block, a convolutional layer is applied before an upsampling operation. The final style block in the decoder generates a Cartesian representation of the heliostat surface's z -control points. Two additional operations are applied after every convolutional layer and affine transformation in the network, though these are not explicitly depicted in Fig. 4. First, a dropout layer is introduced with a dropout probability of 0.2. Second, a rectified linear unit (ReLU) with a negative slope of 0.1 is applied. The total number of trainable parameters in the network is to 3.5 million.

The model underwent training utilizing the semi-artificial dataset introduced in Chapter 2.2. By using also less than eight input flux densities and sun positions during the training we ensure that the model can predict surfaces for any number of input flux densities up to eight. The objective of the training is the minimization of the *mean absolute error* (MAE) between the predicted z -control points $P_{i,j}^{z,pred}$ and those from the training set $P_{i,j}^{z,GT}$:

$$MAE = \frac{1}{N} \sum_{i=1}^{n=8} \sum_{j=1}^{m=8} \left| P_{i,j}^{z,pred} - P_{i,j}^{z,GT} \right| \quad (2)$$

Where:

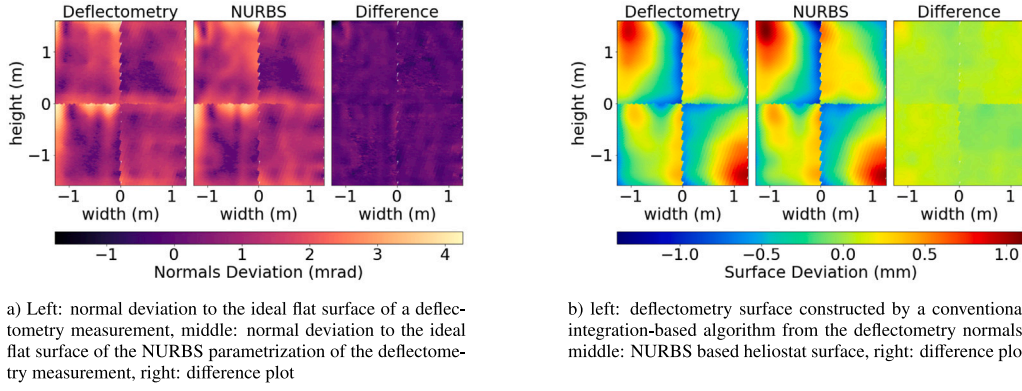


Fig. 5. Illustration of the NURBS fitting procedure for an individual heliostat. The left panel depicts the normal vectors resulting from the NURBS parameterization, while the right panel displays cartesian representations of the surface.

- n and m are the number of control points.
- $N = n \times m$ is the total number of control points.

The model is trained for a total of 50 epochs. Training is stopped at this point as the validation loss has converged, indicating no further improvements can be achieved with additional training.

3. NURBS for surface representation: Quality and computational efficiency

To assess the quality of the NURBS parameterization, we compare it to the deflectometry normal vector point cloud. Fig. 5 shows an example heliostat surface represented using NURBS. On the left, Fig. 5(a) compares the normal vectors from deflectometry measurements with those from the NURBS surface, demonstrating minimal differences and underscoring the accurate parameterization of the heliostats' normal vectors using the $4 \times 8 \times 8$ NURBS z-control points. In Fig. 5(b), the surface deviation of the heliostat from the ideal assumption is depicted, with a conventional integration-based algorithm estimation on the left and the NURBS parameterization on the right. The strong alignment between both underscores the robust parameterization achieved with the NURBS parameters. To quantify the quality of the NURBS parameterization, we calculate the angular deviation between the normal vectors from the NURBS spline and the deflectometry measurements. 81% of the NURBS normal vectors have an angular deviation of zero mrad from the original deflectometry normal vectors within machine accuracy. The remaining 19% form an extremely skewed tail, with the largest outlier being 3mrad. While the majority of the normal vectors are fitted perfectly, the outliers have mainly three causes: the NURBS normals must be smoother than the original, sharp edges in the contour lines are rounded, and at the boundary of the facets, the NURBS parameterization becomes slightly less accurate due to boundary effects.

Next, we compare the flux density predictions based on the normal vector cloud and the NURBS spline loaded into the ray tracer proposed by Pargmann et al. [12]. Approximately 2.3 million rays are simulated per heliostat, using the CSR sun shape model proposed by Buie et al. [41], with a CSR value of 5%. Fig. 6 shows on the left (Fig. 6(a)), the flux density prediction from deflectometry normals and NURBS for the heliostat with the largest outlying normal vector deviation (10.2mrad) and the smallest (2mrad). In both cases, the flux densities match strongly with nearly no visible difference, highlighting the robust surface parameterization of the NURBS. To quantify the quality of the flux density prediction, the following accuracy metric is calculated:

$$\text{ACC}_{\text{GT,pred}} = \frac{\sum |\phi_{\text{GT}} - \phi_{\text{pred}}|}{\sum |\phi_{\text{GT}}|} \quad (3)$$

As raytracing involves Monte-Carlo Sampling, slight variations in flux densities are expected when raytracing the same scene twice. To

account for this, the flux density for the normal vector cloud was calculated twice $\phi_{\text{Normals1/2}}$, and the loss of accuracy caused by the NURBS parameterization is calculated as:


$$\text{loss} = 1 - (\text{ACC}_{\text{NURBS, Normals1}} + (1 - \text{ACC}_{\text{Normals1, Normals2}})) \quad (4)$$

Fig. 6(b) illustrates the accuracy loss in flux density predictions attributable to the NURBS parameterization for 184 deflectometry measurements conducted at STJ. These 184 measurements constitute the subset with more than 99% of their mirror surface measured successfully (see Section 2.2 for more details on data). The distribution centers around its mean $\text{loss}_{\text{NURBS}} = 2.2$, indicating a significant but relatively small loss due to the NURBS.

The substantial parameter reduction by 99.97% (from more than one million to 256) is essential for the implementation of iDLR. Although it is technically feasible to train generative models with over a million pixels, the computational expense is prohibitive; for example, the styleGAN architecture requires 41 days and 4 h of training on a Tesla V100 GPU for 1024×1024 pixels [23]. In contrast, training the iDLR model to predict the NURBS surface takes just 12 h on a single NVIDIA A100 GPU. Furthermore, training such models is highly challenging due to the large number of trainable parameters, which complicates regularization. Additionally, the parameter reduction significantly enhances memory efficiency, decreasing the memory requirement by 99.91% from 8.6MB to 7kB per heliostat. This improvement is critical for two main reasons. First, deep learning models often necessitate hundreds of thousands of training samples, making a memory-efficient data representation imperative. Second, large concentrated solar power (CSP) plants may comprise up to 200,000 heliostats, necessitating an efficient surface representation for real-time flux density predictions during operation, which includes heliostat digital twins with surface information.

The deviations between the NURBS normals and the original normals have resulted in a minor flux density accuracy loss of 2.2%. This loss is particularly minimal when contrasted with the accuracy of the ideal heliostat assumption at the STJ, which is $\text{ACC}_{\text{ideal}} = 67\%$. Consequently, the potential gain in flux density prediction accuracy substantially outweighs the loss introduced by the NURBS parameterization. It is feasible to further reduce the mrad and flux density loss by using additional NURBS control points. However, the accuracy achieved with our current parameterization is considered sufficient for our purposes.

The proposed NURBS parameterization represents a significant advancement over the current SOTA in heliostat surface parameterization, specifically the normal vector cloud. Key advantages include enhanced memory and computational efficiency, which are crucial for large power plants employing a heliostat digital twin that incorporates surface information for real-time flux density prediction. Moreover, the NURBS method allows for a reduced number of free parameters in



a) Simulated flux densities for the heliostat with the smallest outlying fitted normal vector (upper row) and the largest outlying fitted normal vector (lower row) are shown. The left column displays the flux density prediction using the normal vector cloud from the deflectometry measurement, while the middle column presents the flux density resulting from the NURBS parameterization. The right column illustrates the difference between these two flux density predictions.

b) The accuracy loss in flux density prediction that can be attributed to the NURBS parameterization. The distribution is centred around the median accuracy loss of 2.2%. Hence there is a significant, but relatively small loss due to the NURBS parameterization.

Fig. 6. Flux density prediction accuracy under the NURBS representation. Two example flux densities are shown on the left and the histogram with the loss in accuracy on the right.




Fig. 7. (a): The presented surface predictions are randomly selected. The top row illustrates the ground truth obtained through deflectometry, while the bottom row depicts the predictions generated by iDLR. b: The provided flux density predictions correspond to the heliostats shown above. The top row presents the ground truth, derived through raytracing deflectometry surfaces, while the middle row showcases predictions generated by raytracing the iDLR predictions. As a point of reference, the bottom row displays flux density predictions derived from the ideal heliostat assumption.

machine learning models and provides a differentiable formulation. Consequently, all training losses and validation metrics in this work are based on the NURBS parameterization. Overall, our findings suggest that NURBS will become the new SOTA for heliostat surface parameterization.

4. Results

4.1. Surface prediction

Fig. 7 shows the results of the iDLR prediction. In Fig. 7a, ten randomly selected surface predictions by iDLR are shown, whereby seven simulated flux densities were used as an input. The mean absolute error per control-point (MAE) between the predicted surface and the deflectometry-derived surface has a median of 0.14 mm with (Minimum/Quantile 1/Quantile 3/Maximum) of (0.07/0.12/0.17/0.7) mm. For comparison the typical surface deviation falls within the range

of -2 mm to 2 mm , but exhibiting considerable variability among heliostats. As illustrated in Fig. 7a, a strong alignment is evident between the predicted surface and the deflectometry ground truth surface. Notably, the model demonstrates accurate prediction of heliostat waviness, representing surface features with the highest spatial frequency that is representable with the NURBS spline. Fig. 8 shows an analysis of iDLR performance. On the left, Fig. 8(a) shows the performance in dependency to the heliostats distance to the tower (a few outliers are not displayed). The surface predictions become significantly better at larger distances up to the considered maximum distance of 300 m. However, no analogous effect is identified concerning the azimuth of the heliostat position. Fig. 8(b) illustrates the model's surface reconstruction performance in relation to the number of input flux densities. Notably, a substantial enhancement in precision is observed with an increased number of input flux densities. However, given that the model was trained with a maximum capacity for processing eight flux densities, the potential impact of utilizing an even larger number of

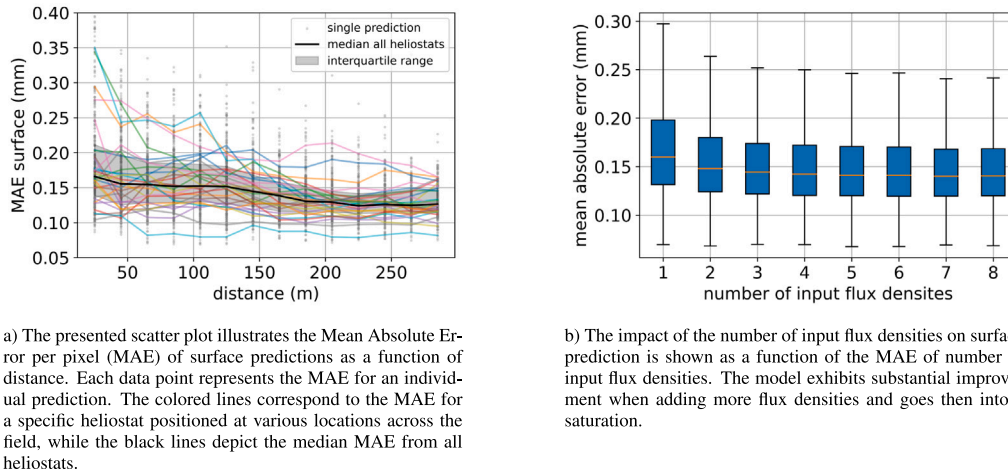


Fig. 8. MAE of iDLR surface predictions in dependency on the distance to the tower (left) and the number of input flux densities (right).

input flux densities on the results remains unclear as the saturation of the curve could be at a smaller MAE.

4.2. Flux density prediction

In the subsequent phase, the predicted surfaces serve as a basis for forecasting flux densities of heliostats using raytracing. Randomly selected flux density predictions are visualized in Fig. 7b and compared to the, deflectometry based flux density prediction and the ideal heliostat assumption. A high match between the flux densities predicted from iDLR and deflectometry surfaces is visible across all shown heliostats. Remarkably, this does not only hold for the heliostats with accurate surface prediction (compare Fig. 7a) but also for those with a poor surface reconstruction (the third heliostat from the left). This phenomenon can be attributed to the inherent underdetermination of the problem, as surface deformations that significantly change the facets mean normal vector (especially canting errors) have a stronger effect on the flux density distribution for heliostats at larger distances. To quantify the quality of the flux density predictions, the mean accuracy is computed as described in Section 2.1. The median accuracy of the flux density prediction of the model is 0.92 with (Min/Q1/Q3/Max) of (0.43/0.90/0.94/0.98), proving a very accurate and reliable flux density prediction using the surfaces predicted with the model for most predictions. Generally, even fine details can be predicted using the method. However, some small deviations from especially inhomogeneous and exceptional flux densities are possible. For comparison, the ideal heliostat assumption of a flat surface without deformations achieves a median accuracy of only 0.67 with (Min/Q1/Q3/Max) of (0.37/0.59/0.74/0.9). This is not only on median 0.25 (−0.08/0.18/0.33/0.55) less accurate than the iDLR enhanced raytracing prediction, but also the interquartile range (IQR) of the ideal heliostat prediction ($IQR_{ideal} = 0.15$) is more than three times as high as those from iDLR ($IQR_{iDLR} = 0.04$). This makes the iDLR enhanced flux density prediction more precise and reliable than the ideal heliostat assumption. Over the whole test data set consisting of 32 simulated heliostats on 142 positions, the ideal heliostat flux density was more accurate than the iDLR prediction for only four heliostats, showing a failure rate of iDLR based flux density prediction around 1‰ in the presented simulated heliostat field.

Fig. 9 compares deflectometry and iDLR enhanced flux density forecasts with the corresponding target image taken at the STJ. The heliostat is positioned at 4.4 m west and 25 m north of the tower. Three source flux densities are used to predict the heliostats surface. This iDLR surface prediction is used to predict three flux densities under new sun position. The simulated flux densities correspond to target images that were taken for calibration purpose at the STJ (right side). The high

match between raytracing deflectometric measured heliostats and the target images taken during the calibration measurement shows that the information necessary to predict the surfaces are in the target images as well as in the raytracing simulation, emphasizing the possibility to transfer the simulative model to real data by machine learning techniques. Moreover, Fig. 9 illustrates a spatial extrapolation flux density prediction and its comparison to a real target image. The STJ incorporates a secondary target positioned 18 m west and 15 m than the main target. The high correspondence between the iDLR flux density prediction for spatial extrapolation and the deflectometry flux density prediction and the real target image, underscores the capability of the predicted iDLR surface to not only forecast flux densities on the source target but also on another target plane. This is important as receivers are situated in a different plane than the source targets and hence the flux density prediction on the receiver is usually a spatial extrapolation.

5. Discussion

The results obtained through simulation underscore the significant ability of iDLR. The method demonstrates its capability to predict heliostat surfaces with high accuracy using simulated flux densities, despite the inherently undetermined and ill-posed nature of the inverse problem. While the majority of the reconstructed surfaces closely align with the deflectometry ground truth (median MAE of 0.14 mm with (Min/Q1/Q3/Max) of (0.07/0.12/0.17/0.7) mm), it is important to acknowledge a potential drawback: there exists a risk of less accurate surface predictions compared to deflectometry measurements for a minority of heliostats. This phenomenon is particularly evident for heliostats with limited input flux density data and those located close to the tower, where surface predictions may deviate from deflectometry ground truth. These discrepancy arises from canting errors and surface deformations, which alter the mean normal vector of the facets. For heliostats positioned farther from the tower, the impact of these factors on the flux density is more pronounced. As a result, the flux density conserves more information about canting errors and surface deformations, which alter the mean normal vector, at greater distances, making these deviations easier to predict. However, this relationship does not hold for surface errors with higher spatial frequency, such as waviness. With higher distances to the tower, the features in the flux density caused by waviness diminishes and the flux density becomes more smooth. Notably, the model demonstrates the ability to predict higher frequency heliostat features even for heliostats at greater distances, indicating the utilization of learned static material constraints to make accurate predictions rather than relying solely on conserved information in the flux density.




Fig. 9. Comparison of simulated flux densities arising from raytracing deflectometry and iDLR surfaces with target images obtained at the Solar Tower Jülich. The top row displays three simulated flux densities (left) corresponding to three target images (right). These flux densities were used to predict the heliostat surfaces (not shown). Subsequently, flux density predictions for three new sun positions were generated with a ray tracer (middle row, left). The three flux densities on the right are generated under the ideal heliostat assumption. The bottom row presents the ground truth flux densities along with their corresponding target images. Notably, the red-marked target images were captured from an additional target positioned 18 m west and 15 m higher than the source target.

In contrast to the singular existing alternative for heliostat surface prediction from target images, the differentiable ray tracer proposed by Pargmann et al. [12], our findings demonstrate substantial improvement. While the differentiable ray tracer can only predict the surface for one singular heliostat with pronounced, artificially introduced surface deformations, iDLR is able to predict a wide range of measured, realistic heliostat surfaces. The problem of the differentiable ray tracer lies in the underdetermined nature of the problem as it does not acquire any preknowledge about the physical constraints of the surface deformations. Our approach successfully navigates this challenge by learning the physical constraints of the heliostat surfaces, thereby narrowing the solution space and enabling precise predictions despite the inherent underdetermination and ill-posed characteristics of the problem. Furthermore, iDLR predictions come at real time while the predictions of differentiable raytracing is depended on the number of epochs and number of input flux densities, resulting in computation times in the range of a few minutes up to some hours per heliostat.

On the other hand, we identified three drawbacks of iDLR. First, while the model provides real-time predictions, its training process is notably time-intensive. The execution of raytracing and subsequent flux density processing requires approximately 0.1 s per iteration. As a result, generating the training dataset demanded approximately 45 h of computation on a single GPU (NVIDIA A100). Additionally, training and testing the model on the same GPU required an additional 12 h. Despite this initial time investment, the eventual computational costs are anticipated to be outweighed by iDLR's real-time inference capabilities, making it superior to optimization-based algorithms like differentiable raytracing over even short periods of use. Second, to achieve accurate surface predictions, the model requires training on realistic heliostat surfaces, which can be obtained through measurement or simulation. While surface measurement entails associated costs, it is important to note that each power plant implementing iDLR does not require its own measurement setup. Instead, a single setup per heliostat model, located either at a power plant or on the manufacturer's side, is sufficient. Alternatively, finite-element simulation of heliostats can be used to predict potential deformations caused by construction-induced stress, eliminating the need for surface measurements. However, simple surface simulations that do not account for physical factors, such as adding Gaussian deformations to an ideal heliostat, were found to be insufficient. Models trained on such simplified data struggled to predict real heliostat surfaces based on simulated flux densities, as they failed to learn the physically constrained heliostat surface space. Lastly, the model may face challenges in accurately predicting rare or unique surface deformations that deviate significantly from those encountered

during training. These out-of-distribution predictions present a notable challenge, necessitating ongoing refinement and adaptation of the iDLR framework.

The minimum requisite quantity of measured heliostat surfaces for model training is contingent upon two primary factors. Firstly, it relies on the diversity present in surface deformations, as the iDLR model necessitates comprehending a broader surface space, thereby requiring a larger dataset for effective learning. Secondly, it is influenced by the generalizability of a distinct set of deflectometry measurements. Biases often manifest within datasets, such as an overrepresentation of deflectometry measurements from heliostats situated in close proximity to the tower, owing to the ease of application in these areas. Given the likely existence of correlations between mirror error and heliostat position (e.g. larger canting errors occurring at closer positions), a biased dataset may diminish the predictive accuracy of the model, even when training set sizes remain constant.

When comparing to existing methods for surface measurements such as deflectometry [11], laser scanning, photogrammetry [14,16,18] or flux mapping [21], iDLR stands out as a software-only approach for surface prediction. It relies on hardware already employed in commercial power plants and eliminates the need for active human involvement during execution, making it considerably more cost-effective. However, as mentioned above, this is traded against the risk of having a less accurate surface prediction for a minority of the heliostats.

The flux density predictions obtained through raytracing the iDLR prediction significantly outperform the ideal heliostat assumption (accuracy of iDLR is 0.92 compared to accuracy of ideal heliostat assumption of 0.67). This highlights the potential benefits of replacing the ideal heliostat model with the iDLR prediction in the operation of a solar power plant as aim-point optimization strategies reliant on raytracing a heliostat model, as seen in works like [2–4,45], can experience enhancements when the ideal heliostat assumption is substituted with the iDLR surface predictions. Notably, Zhu et al. [4] demonstrated that their post-installation calibration procedure “can lead to up to a 20% increase in utilized solar power”. However, their method is only applicable to focused heliostats, whereas our approach is adaptable to all heliostat models. A comparable efficiency gain, exceeding the accuracy achieved by Zhu et al. [4], could potentially be realized with iDLR. Moreover, this is achieved with a generic heliostat model, that can, unlike the model of Zhu et al. [4] parameterize all heliostat models precisely. Nevertheless, the precise efficiency gain remains to be determined in future research.

In addition to enhance aim point optimization, achieving reduced tracking errors and enhanced heliostat calibration is desirable. Current

calibration methods often use the flux density's center of mass as input to a rigid-body regression model [22,34,46], assuming the center of mass aligns with the desired aim point. However, due to irregularities in flux densities, significant discrepancies can arise between the desired aim point and the center of mass, introducing modeling errors that compromise the accuracy of the rigid-body model. This challenge could be tackled by integrating the approach of Pargmann et al. [22], which employs a combination of differentiable raytracing and deep learning for calibration, with iDLR. In particular, iDLR's surface prediction can be initialized within the differentiable raytracing framework. The calibration process could then be further refined by treating the heliostat's desired aim point as an optimizable variable, enabling its precise determination and improve the accuracy of the calibration. This avenue of investigation warrants further exploration in future studies.

During power plant construction, surface quality measurements of heliostats are typically conducted sparsely and randomly across the field due to the significant time and cost associated with conventional methods such as deflectometry or photogrammetry. Our method overcomes these limitations by enabling the extraction of detailed surface information for every heliostat in the field, facilitating comprehensive monitoring of the field's optical performance. While occasional inaccuracies may occur, this approach allows for a vastly greater number of heliostats to be measured accurately compared to traditional methods, potentially significantly enhancing the efficiency and coverage of surface quality assessments.

To successfully apply the trained model to real-world data, the limited availability of surface data for training poses significant challenges. Relying solely on real data for training is impractical due to the requirement for surface measurements of a large number of heliostats in the corresponding power plant. Consequently, the preferred approach involves transferring the simulated model presented here to real data using deep learning techniques like domain randomization [35–38]. Key latent parameters, such as varying sun shapes [41,47], soiling [48, 49], mirror surface roughness [50], heliostat geometry models, and the disparities between simulated flux density and target images due to non-ideal Lambertian properties, background radiation reflection, varying sun shapes or aberration effects can be randomized during simulation. This can ensure the robustness of the simulated model against differences between simulation and real-world data.

An alternative strategy for transferring the model to real-world data involves incorporating differentiable raytracing formulations or utilizing a deep learning ray tracer within the workflow. This approach would enable close-loop training, using target images exclusively after pretraining on simulated data.

6. Conclusion

We have presented iDLR, an innovative and cost-effective methodology for predicting heliostat surfaces from target images, leveraging deep learning to solve the inverse direction of the raytracing process. A necessity for iDLR is the compact heliostat surface parameterization using NURBS. This approach enables heliostat surfaces to be represented with just 256 parameters, achieving a 99.97% reduction in the number of parameters and a 99.91% reduction in memory requirements, while maintaining minimal loss in accuracy. This significantly improves computational efficiency compared to existing methods. Our findings suggest that NURBS parameterization has the potential to become the new SOTA for heliostat surface representation. Our simulation results demonstrate that iDLR effectively predicts heliostat surfaces with accuracy comparable to deflectometry for the majority of heliostats, with a median MAE of 0.14 mm. When integrating the iDLR surface predictions into a ray-tracing environment to compute flux densities, our method achieves an accuracy of 92%, surpassing the performance of the ideal heliostat assumption by 25%. The method relies solely on software and requires no additional hardware or sensors for inference. It utilizes input data already available from routine power plant

operations, enabling nearly zero-cost inference to monitor the optical performance of the whole heliostat field. Furthermore, integrating the predicted surfaces into raytracing aim-point optimization strategies and the calibration procedure has the potential to enhance the flux density distribution on the receiver, thereby optimizing the overall efficiency of power plant operations.

Declaration of generative AI in scientific writing

During the preparation of this work the author(s) used large language models (LLM) (GPT3.5/4 of openAI and the LLMs by deepL) in order to improve the language. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

CRediT authorship contribution statement

Jan Lewen: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Max Pargmann:** Supervision. **Mehdi Cherti:** Supervision. **Jenia Jitsev:** Supervision. **Robert Pitz-Paal:** Supervision. **Daniel Maldonado Quinto:** Supervision, Project administration, Funding acquisition.

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.

Acknowledgments

We gratefully acknowledge the use of data from the Solar Tower Jülich, a research power plant operated by the German Aerospace Center (DLR). Additionally, we extend our thanks for the funding provided by the Helmholtz Association (HGF) for the GANCASTR project (Grant Number ZT-I-PF-5-069), under which this research was conducted. Furthermore, we thank the Helmholtz Association's Initiative and Networking Fund for funding computational resources on the HAICORE@FZJ partition.

References

- [1] B. Belhomme, R. Pitz-Paal, P. Schwarzbözl, Optimization of heliostat aim point selection for central receiver systems based on the ant colony optimization metaheuristic, *J. Sol. Energy Eng.* 136 (1) (2013) <http://dx.doi.org/10.1115/1.4024738>.
- [2] L. Oberkirsch, D.A.V. Zanger, D. Maldonado Quinto, P. Schwarzbözl, B. Hoffschmidt, Static optimal control: Real-time optimization within closed-loop aim point control for solar power towers, *Sol. Energy* 255 (2023) 327–338, <http://dx.doi.org/10.1016/j.solener.2023.03.051>, URL: <https://www.sciencedirect.com/science/article/pii/S0038092X23002116>.
- [3] S. Wu, D. Ni, Real-time heliostat field aiming strategy generation for varying cloud shadowing using deep learning, in: AIP Conference Proceedings, AIP Publishing, 2023, <http://dx.doi.org/10.1063/5.0149189>.
- [4] R. Zhu, D. Ni, T. Yang, J. Yang, J. Chen, G. Xiao, Heliostat field aiming strategy optimization with post-installation calibration, *Appl. Therm. Eng.* 202 (2022) 117720, <http://dx.doi.org/10.1016/j.applthermaleng.2021.117720>.
- [5] L.L. Vant-Hull, M.E. Izygon, C.L. Pitman, *Real-Time Computation and Control of Solar Flux Density on a Central Receiver (Solar Two)(Protection Against Excess Flux Density)*, Technical Report, American Solar Energy Society, Boulder, CO (United States), 1996.
- [6] L.L. Vant-Hull, M.E. Izygon, C.L. Pitman, *Real-Time Computation and Control of Solar Flux Density on a Central Receiver (Solar Two)(Preheat)*, Technical Report, American Society of Mechanical Engineers, New York, NY (United States), 1996.
- [7] L. Vant-Hull, *Concentrator optics*, in: *Solar Power Plants: Fundamentals, Technology, Systems, Economics*, Springer, 1991, pp. 84–133.
- [8] K. Lovegrove, W. Stein, *Concentrating Solar Power Technology: Principles, Developments and Applications*, Elsevier, 2012.

- [9] C.A. Arancibia-Bulnes, M.I. Peña-Cruz, A. Mutuberría, R. Díaz-Urbe, M. Sánchez-González, A survey of methods for the evaluation of reflective solar concentrator optics, *Renew. Sustain. Energy Rev.* 69 (2017) 673–684, <http://dx.doi.org/10.1016/j.rser.2016.11.048>, URL: <https://www.sciencedirect.com/science/article/pii/S136403211630747X>.
- [10] B. Belhomme, R. Pitz-Paal, P. Schwarzbözl, S. Ulmer, A new fast raytracing tool for high-precision simulation of heliostat fields, *J. Sol. Energy Eng.* 131 (3) (2009) <http://dx.doi.org/10.1115/1.3139139>.
- [11] S. Ulmer, T. März, C. Prah, W. Reinalter, B. Belhomme, Automated high resolution measurement of heliostat slope errors, *Sol. Energy* 85 (4) (2011) 681–687.
- [12] M. Pargmann, J. Ebert, D.M. Quinto, R. Pitz-Paal, S. Kesselheim, In-situ solar tower power plant optimization by differentiable ray tracing, 2023, <http://dx.doi.org/10.21203/rs.3.rs-2554998/v1>, Research Square Platform LLC.
- [13] M. El Ydrissi, H. Ghennioui, A. Farid, et al., A review of optical errors and available applications of deflectometry technique in solar thermal power applications, *Renew. Sustain. Energy Rev.* 116 (2019) 109438.
- [14] A. Bonanos, M. Faka, D. Abate, S. Hermon, M. Blanco, Heliostat surface shape characterization for accurate flux prediction, *Renew. Energy* 142 (2019) <http://dx.doi.org/10.1016/j.renene.2019.04.051>.
- [15] R. Monterreal, R. Enrique, J. Fernández-Reche, An improved methodology for heliostat testing and evaluation at the plataforma solar de almería, in: *AIP Conference Proceedings*, vol. 1850, (no. 1) AIP Publishing, 2017.
- [16] M.R. Shortis, G.H.G. Johnston, Photogrammetry: An available surface characterization tool for solar concentrators, Part I: Measurements of surfaces, *J. Sol. Energy Eng.* 118 (3) (1996) 146–150, <http://dx.doi.org/10.1115/1.2870886>.
- [17] K. Pottler, E. Lu^opfert, G.H.G. Johnston, M.R. Shortis, Photogrammetry: A powerful tool for geometric analysis of solar concentrators and their components, *J. Sol. Energy Eng.* 127 (1) (2005) 94–101.
- [18] K. Pottler, E. Lu^opfert, G.H. Johnston, M.R. Shortis, Photogrammetry: A powerful tool for geometric analysis of solar concentrators and their components, *J. Sol. Energy Eng.* 127 (1) (2005) 94–101.
- [19] M. Röger, C. Prah, S. Ulmer, Heliostat shape and orientation by edge detection, 2010.
- [20] A. Bonanos, M. Faka, D. Abate, S. Hermon, M. Blanco, Heliostat surface shape characterization for accurate flux prediction, *Renew. Energy* 142 (2019) 30–40.
- [21] A. Martínez-Hernández, R. Conceição, C.-A. Asselineau, M. Romero, J. González-Aguilar, Advanced surface reconstruction method for solar reflective concentrators by flux mapping, *Sol. Energy* 266 (2023) 112162, <http://dx.doi.org/10.1016/j.solener.2023.112162>, URL: <https://www.sciencedirect.com/science/article/pii/S0038092X2300796X>.
- [22] M. Pargmann, M. Leibauer, V. Nettelroth, D. Maldonado Quinto, R. Pitz-Paal, Enhancing heliostat calibration on low data by fusing robotic rigid body kinematics with neural networks, *Sol. Energy* 264 (2023) 111962, <http://dx.doi.org/10.1016/j.solener.2023.111962>, URL: <https://www.sciencedirect.com/science/article/pii/S0038092X23005960>.
- [23] T. Karras, S. Laine, T. Aila, A style-based generator architecture for generative adversarial networks, 2019, [arXiv:1812.04948](https://arxiv.org/abs/1812.04948).
- [24] T.B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D.M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language models are few-shot learners, 2020, [arXiv:2005.14165](https://arxiv.org/abs/2005.14165).
- [25] S. Bubeck, V. Chandrasekaran, R. Eldan, J. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y.T. Lee, Y. Li, S. Lundberg, H. Nori, H. Palangi, M.T. Ribeiro, Y. Zhang, Sparks of artificial general intelligence: Early experiments with GPT-4, 2023, [arXiv:2303.12712](https://arxiv.org/abs/2303.12712).
- [26] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, High-resolution image synthesis with latent diffusion models, 2022, [arXiv:2112.10752](https://arxiv.org/abs/2112.10752).
- [27] H. Li, J. Schwab, S. Antholzer, M. Haltmeier, NETT: solving inverse problems with deep neural networks, *Inverse Problems* 36 (6) (2020) 065005, <http://dx.doi.org/10.1088/1361-6420/ab6d57>.
- [28] J. Adler, O. Öktem, Solving ill-posed inverse problems using iterative deep neural networks, *Inverse Problems* 33 (12) (2017) 124007, <http://dx.doi.org/10.1088/1361-6420/aa9581>.
- [29] X. Wang, A. Kumar, C. Shelton, B.M. Wong, Harnessing deep neural networks to solve inverse problems in quantum dynamics: machine-learned predictions of time-dependent optimal control fields, *Phys. Chem. Chem. Phys.* : PCCP (2020) <http://dx.doi.org/10.26434/chemrxiv.12905633>.
- [30] T. Bister, M. Erdmann, U. Köthe, J. Schulte, Inference of cosmic-ray source properties by conditional invertible neural networks, *Eur. Phys. J. C* 82 (2) (2022) <http://dx.doi.org/10.1140/epjc/s10052-022-10138-x>.
- [31] S. Mousavi, G. Beroza, Deep-learning seismology, *Science* 377 (2022) <http://dx.doi.org/10.1126/science.abm4470>.
- [32] D. Colombo, E. Turkoglu, W. Li, E. Sandoval-Curiel, D. Rovetta, Physics-driven deep-learning inversion with application to transient electromagnetics, *Geophysics* (2021) 1–65, <http://dx.doi.org/10.1190/GEO2020-0760.1>.
- [33] S. Arridge, P. Maass, O. Öktem, C.-B. Schönlieb, Solving inverse problems using data-driven models, *Acta Numer.* 28 (2019) 1–174, <http://dx.doi.org/10.1017/S0962492919000059>.
- [34] K.W. Stone, Automatic heliostat track alignment method, 1986, Google Patents, US Patent 4, 564, 275.
- [35] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, P. Abbeel, Domain randomization for transferring deep neural networks from simulation to the real world, 2017, <http://dx.doi.org/10.48550/ARXIV.1703.06907>, URL: <https://arxiv.org/abs/1703.06907>.
- [36] X.B. Peng, M. Andrychowicz, W. Zaremba, P. Abbeel, Sim-to-real transfer of robotic control with dynamics randomization, in: 2018 IEEE International Conference on Robotics and Automation, ICRA, IEEE, 2018, <http://dx.doi.org/10.1109/icra.2018.8460528>.
- [37] F. Sadeghi, S. Levine, CAD2RL: Real single-image flight without a single real image, 2016, <http://dx.doi.org/10.48550/ARXIV.1611.04201>, arXiv, URL: <https://arxiv.org/abs/1611.04201>.
- [38] OpenAI, M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, J. Schneider, S. Sidor, J. Tobin, P. Welinder, L. Weng, W. Zaremba, Learning dexterous in-hand manipulation, 2018, <http://dx.doi.org/10.48550/ARXIV.1808.00177>, arXiv, URL: <https://arxiv.org/abs/1808.00177>.
- [39] M. Pargmann, E. Jan, S. Kesselheim, M. Götz, D. Maldonado Quinto, Differentiable raytracing for solar tower power plants, 2024, <http://dx.doi.org/10.5281/ZENODO.11047454>, Zenodo, URL: <https://zenodo.org/doi/10.5281/zenodo.11047454>.
- [40] A. Deva Prasad, A. Balu, H. Shah, S. Sarkar, C. Hegde, A. Krishnamurthy, NURBS-diff: A differentiable programming module for NURBS, *Computer-Aided Des.* 146 (2022) 103199, <http://dx.doi.org/10.1016/j.cad.2022.103199>.
- [41] D. Buie, A. Monger, C. Dey, Sunshape distributions for terrestrial solar simulations, *Sol. Energy* 74 (2) (2003) 113–122, [http://dx.doi.org/10.1016/s0038-092x\(03\)00125-7](http://dx.doi.org/10.1016/s0038-092x(03)00125-7).
- [42] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, T. Aila, Analyzing and improving the image quality of StyleGAN, 2020, [arXiv:1912.04958](https://arxiv.org/abs/1912.04958).
- [43] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, 2015, [arXiv:1512.03385](https://arxiv.org/abs/1512.03385).
- [44] E. Richardson, Y. Alaluf, O. Patashnik, Y. Nitzan, Y. Azar, S. Shapiro, D. Cohen-Or, Encoding in style: a StyleGAN encoder for image-to-image translation, 2021, [arXiv:2008.00951](https://arxiv.org/abs/2008.00951).
- [45] B. Belhomme, R. Pitz-Paal, P. Schwarzbözl, Optimization of heliostat aim point selection for central receiver systems based on the ant colony optimization metaheuristic, *J. Sol. Energy Eng.* 136 (1) (2013) <http://dx.doi.org/10.1115/1.4024738>.
- [46] J.C. Sattler, M. Röger, P. Schwarzbözl, R. Buck, A. Macke, C. Raeder, J. Götsche, Review of heliostat calibration and tracking control methods, *Sol. Energy* 207 (2020) 110–132.
- [47] A. Neumann, A. Witzke, The influence of sunshape on the DLR solar furnace beam, *Sol. Energy* 66 (6) (1999) 447–457, [http://dx.doi.org/10.1016/s0038-092x\(99\)00048-1](http://dx.doi.org/10.1016/s0038-092x(99)00048-1).
- [48] R. Berg, Heliostat Dust Buildup and Cleaning Studies, Office of Scientific and Technical Information (OSTI), 1978, <http://dx.doi.org/10.2172/6867834>.
- [49] E. Roth, R. Pettit, Effect of Soiling on Solar Mirrors and Techniques Used to Maintain High Reflectivity, Office of Scientific and Technical Information (OSTI), 1980, <http://dx.doi.org/10.2172/5249717>.
- [50] A.M. Bonanos, Error analysis for concentrated solar collectors, *J. Renew. Sustain. Energy* 4 (6) (2012) <http://dx.doi.org/10.1063/1.4768546>.