„A Data-Driven Methodology for Precision Flux Density Predictions for Heliostat Fields“

„Eine datengetriebene Methodik zur präzisen Flussdichtenvorhersage für Heliostatenfelder“

Von der Fakultät für Maschinenwesen der Rheinisch-Westfälischen Technischen Hochschule Aachen zur Erlangung des akademischen Grades eines Doktors der Ingenieurwissenschaften genehmigte Dissertation

vorgelegt von

Mathias Angelo Kuhl

Berichter:
Univ.-Prof. Dr.-Ing. Robert Pitz-Paal
Univ.-Prof. Marek Behr, Ph. D.

Tag der mündlichen Prüfung: 17.12.2025

Abstract

The imperative transition to renewable energy sources necessitates advancements in technologies such as Concentrated Solar Thermal (CST) systems. A significant challenge in CST is the accurate characterization of heliostat beams, traditionally reliant on detailed surface deformation measurements—a process that is both costly and time-intensive. This dissertation introduces an innovative data-driven methodology that bypasses the need for such extensive physical modeling.

Central to this approach is a generative neural network framework trained directly on focal spot data. By abstracting heliostat properties and learning the variations in focal spots under different conditions—such as sun angles and aimpoints—the model predicts flux distributions with high accuracy. This unified predictive pipeline aggregates data from the entire heliostat field, capturing shared patterns while adapting to individual heliostat variations.

Validation using real-world data from the Solar Tower Jülich demonstrated the model's robustness. The data-driven pipeline achieved a prediction error of 11% for individual focal spots, surpassing current state-of-the-art methods that rely on extensive measurements, in both accuracy and scalability. When considering factors such as tracking errors, Direct Normal Irradiance (DNI), and reflectivity uncertainties, the total flux prediction error, aggregated from all heliostats, was approximately 5%.

System-level analysis revealed that the enhanced accuracy of focal spot predictions translates into significant efficiency gains. The data-driven pipeline demonstrated efficiency improvements of over 4–16%, depending on receiver complexity.

In summary, this research presents a scalable, cost-effective, and highly accurate framework for heliostat flux prediction. By eliminating the dependency on extensive physical measurements, it offers a practical alternative for CST systems, paving the way for more efficient and sustainable solar energy technologies.

Zusammenfassung

Der notwendige Übergang zu erneuerbaren Energien erfordert technologische Fortschritte in Bereichen wie konzentrierter Solarthermie (CST). Eine zentrale Herausforderung in der CST-Technologie besteht in der präzisen Charakterisierung der Strahlenbündel von Heliostaten. Traditionell basiert diese auf detaillierten Messungen der Oberflächenverformung – ein Verfahren, das sowohl kostenintensiv als auch zeitaufwendig ist. Diese Dissertation stellt einen innovativen, datengetriebenen Ansatz vor, der eine aufwendige physikalische Modellierung überflüssig macht.

Im Mittelpunkt dieses Ansatzes steht ein generatives neuronales Netzwerk, das direkt auf Brennfleckdaten trainiert wird. Durch die Abstraktion der Heliostateigenschaften und das Erlernen der Variationen der Brennfleckverteilungen unter verschiedenen Bedingungen – etwa Sonnenwinkel und Zielpunktlage – prognostiziert das Modell die Flussverteilungen mit hoher Genauigkeit. Die einheitliche Vorhersagepipeline integriert Daten aus dem gesamten Heliostatfeld, erkennt übergreifende Muster und passt sich gleichzeitig an individuelle Heliostatabweichungen an.

Die Validierung mit realen Daten aus dem Solarturm Jülich bestätigt die Robustheit des Modells. Die datengetriebene Pipeline erreicht für einzelne Brennpunkte einen Vorhersagefehler von lediglich 11% und übertrifft damit etablierte, auf umfangreichen Messungen basierende Methoden sowohl hinsichtlich Robustheit als auch Skalierbarkeit. Unter Berücksichtigung relevanter Einflussfaktoren – darunter Nachführfehler, direkte Normalstrahlung (DNI) und Reflexionsunsicherheiten – beträgt der aggregierte Fehler der Gesamtflussvorhersage über alle Heliostaten hinweg etwa 5%.

Die Systemanalyse zeigt, dass die verbesserte Genauigkeit der Flussdichteprognosen signifikante Effizienzsteigerungen ermöglicht. Je nach Komplexität des Strahlungsempfängers führt die datengetriebene Pipeline zu Effizienzgewinnen von 4–16%.

Zusammenfassend präsentiert diese Arbeit ein skalierbares, kosteneffizientes und hochpräzises Verfahren zur Vorhersage von Flussdichteverteilungen für Heliostatfelder. Durch den Verzicht auf umfangreiche physikalische Messungen bietet der Ansatz eine praxistaugliche Alternative für CST-Systeme und trägt zur Entwicklung effizienterer und nachhaltigerer solarthermischer Technologien bei.

Statement on AI Use

ChatGPT was used during the preparation of this dissertation to support the linguistic revision and refinement of text passages across all chapters. Its use was limited to wording, readability, and stylistic improvement. All AI-assisted text was carefully reviewed, revised where necessary, and approved by the author. Full responsibility for the scientific content, argumentation, results, and conclusions remains solely with the author.

Contents

List of Abbreviations

Table 1.: List of Abbreviations

Abbreviation	Description
AFD	Allowable Flux Dsitribution
APCA	Aim Point Control Algorithm
CNN	Convolutional Neural Network
CSP	Concentrated Solar Power
CST	Central Solar Tower
DA	Dataset Abbreviation
DNI	Direct Normal Irradiance
DLR	Deutsches Zentrum für Luft- und Raumfahrt (eng.: German AeroSpace Center)
EOD	Edge of Distribution
IR	Infrared
LCOE	Levelized Cost of Energy
ML	Machine Learning
MPC	Model Predictive Control
MSE	Mean Squared Error
NN	Neural Network
OVR	One-vs-Rest
PDF	Probability Density Function
PV	Photovoltaic
SGD	Stochastic Gradient Descent
STJ	Description placeholder
STP	Standard Test Procedure
VAE	Variational Autoencoder

List of Mathematical Symbols

Table 2.: List of Mathematical Symbols

Symbol	Description	Unit
A	Area	m^2
B	Background or empty target image	-
Cov	Covariance	-
E	Error	%
G	Generator Neural Network	-
H_i	Heliostat with number i	-
$\tilde{H}_i$	Aggregated properties of heliostat i	-
I	Target image	-
L	Loss for neural network training and evaluation	-
R	Spatial target reflectivity	-
T_i	Calibration target with number i	-
U	UNet image-to-image network	-
X	Flux density distribution	$\frac{W}{m^2}$
α	Aim or reflectance angle	radians
ρ	Constant reflectivity value	-
σ	Uncertainty standard deviation	-
ϵ	Emissivity or error parameter	-
$\dot{Q}$	Energy flow	W
$\dot{q}$	Energy flow per area	$\frac{W}{m^2}$
$\mathcal{D}$	Dataset	-
$\tilde{H}$	Aggregated heliostat properties	-
$\vec{A}d$	3D vector of aim direction	-
$\vec{A}p$	3D vector of aim point	-
$\vec{S}$	3D vector of sun position	-
$\vec{n}$	3D heliostat surface normal vector	-
c	Scaling factor	-
d	Heliostat distance to tower	m
n	Number of elements	-
s	Projection depth	m

Chapter 1.

Introduction

1.1. Motivation

The accelerating impacts of global warming and the finite nature of fossil fuels underscore the urgent need for renewable energy sources [1]. Renewable energy technologies are indispensable for mitigating climate change and ensuring a sustainable, secure energy supply for future generations [2].

Among renewable energy technologies, wind and photovoltaic (PV) power have advanced rapidly due to their cost-effectiveness and ability to efficiently convert renewable energy into electricity [3]. However, their reliance on electrical power presents significant challenges in applications where direct utilization is not feasible. Balancing supply and demand necessitates extensive energy storage solutions, which are both technically complex and economically demanding [4]. Additionally, decarbonizing sectors such as heavy transportation, aviation, and industrial processes remains a significant challenge due to the inefficiencies of converting electricity into energy carriers or high-temperature process heat [5].

In this context, Concentrated Solar Technology (CST) provides a fundamentally different approach compared to PV and wind power. Instead of directly generating electricity, CST systems use mirrors to concentrate sunlight onto a central receiver, producing high-temperature heat. This thermal energy serves as a versatile resource, facilitating not only direct electricity generation but also a wide range of additional applications.

CST's primary application has been electricity generation, where it excels due to its inherent ability to store and dispatch energy efficiently. By leveraging thermal energy storage, CST can deliver power on demand, reducing dependence on battery storage, which remains costly and difficult to scale [6]. Additionally, CST systems can complement other renewable energy sources by functioning as Carnot batteries—storing surplus electricity from PV or wind as thermal energy and converting it back into electricity when required [7, 8]. This integration enhances grid stability and ensures a reliable supply of renewable energy during periods of low solar or wind availability.

The practical benefits of CST's storage capabilities are evident in recent hybrid solar parks in China, where CST systems are integrated with PV installations to optimize power output [9, 10]. While PV systems provide cost-effective, continuous electricity generation, CST complements them by addressing peak demand and enabling dispatchable electricity [11]. As renewable energy adoption accelerates globally, CST's role within the energy mix is expected to expand, driven by its unique ability to provide storage and peak shaving capabilities [12].

Beyond electricity generation, CST is an essential solution for the heat sector, which accounts for nearly half of global energy consumption yet lags in decarbonization efforts [13]. Achieving high-temperature heat for industrial processes poses significant challenges for electrification due to the inefficiencies and limitations of electric heaters. CST provides a viable alternative, supplying temperatures exceeding 800°C for applications like chemical processing and cement production, offering a pathway to decarbonize industrial sectors that are otherwise difficult to electrify.

CST also plays a critical role in solar fuel production [14]. High-energy-density sectors such as aviation and heavy transportation heavily rely on liquid fuels. While PV can contribute to solar fuel production through electrolysis, CST's high-temperature heat enables chemical reactions directly, offering a potentially more efficient pathway. This capability supports thermochemical cycles for producing hydrogen, ammonia, and other synthetic fuels, bypassing the inefficiencies of converting electricity into chemical energy [15].

— By addressing the critical challenges of energy storage, peak shaving, high-temperature process heat, and solar fuel production, CST emerges as a versatile and indispensable technology in the renewable energy landscape. Its ability to complement PV and wind by providing dispatchable electricity, enabling industrial decarbonization, and driving efficient solar fuel production underscores its pivotal role in achieving a sustainable and reliable energy future.

1.2. Challenges in CST Adoption

Building upon CST's versatility and indispensable role in the renewable energy landscape, it becomes imperative to address the challenges that hinder its widespread adoption. While CST demonstrates unique advantages over other renewable technologies, its economic viability is constrained by significant technical and operational hurdles. These challenges, particularly those associated with the heliostat field and flux distribution control, demand innovative solutions to unlock CST's full potential. The following section outlines the broader challenges of CST systems before narrowing the focus to the specific problem addressed in this work: accurate flux density prediction and its critical role in CST efficiency and safety.

Economic Challenges in CST Systems and the Role of Intelligent Algorithms

Despite its significant potential, the adoption of CST is significantly hindered by its relatively high Levelized Cost of Energy (LCOE). One of the primary cost drivers in CST tower plants is the heliostat field, accounting for approximately 50% of the total capital investment [16]. While inherently expensive, precisely controllable heliostats are essential for maximizing system efficiency. Tracking inaccuracies not only diminish thermal efficiency but also lead to increased receiver wear and higher maintenance demands, ultimately impacting the overall performance and longevity of the power plant [17].

Addressing these economic challenges requires innovative approaches that enhance the efficiency and reliability of CST systems. Intelligent algorithms present a promising solution by improving predictions and control mechanisms within CST plants. These advanced algorithms optimize heliostat operation, enhancing thermal efficiency and reducing material wear. They extend the service life of critical components and lower

maintenance costs, creating a more sustainable operational framework [18, 19]. Furthermore, accurate prediction and forecasting capabilities enable CST plants to adapt to dynamic solar conditions, improving operational stability and facilitating the efficient management of thermal energy storage and power output [20–23].

The potential of intelligent algorithms extends beyond operational improvements. By compensating for the reduced accuracy of lower-cost heliostats with sophisticated control strategies, these algorithms offer substantial cost savings during the design and construction phases. This approach enables CST systems to achieve high efficiencies while utilizing less expensive, less precise components—an essential step toward making CST a competitive and scalable solution in the renewable energy market.

In summary, intelligent algorithms are poised to play a pivotal role in overcoming the economic barriers to CST adoption. By enhancing performance, reducing costs, and improving scalability, they contribute to making CST a more competitive and viable technology within the evolving energy market, strengthening its role in the transition toward a sustainable and resilient energy infrastructure.

Challenges in Flux Density Prediction for Solar Tower Plants

The operational efficiency and safety of Solar Tower Plants (STP) critically depend on the accurate prediction of the flux density distribution on the receiver surface. This distribution, a measure of the spatially distributed energy flux, directly impacts both the thermal efficiency and the longevity of the receiver. To maintain safety, durability, and optimal performance, CST receivers require an optimal flux density distribution that adheres to strict thermal thresholds. Achieving this distribution is further complicated by the variable nature of solar radiation and cloud passages, which necessitate precise control over the flux density to optimize energy absorption and avoid damage.

A major challenge in controlling flux distribution is the accurate prediction of focal spots produced by individual heliostats. These focal spots, the concentrated light reflected by each heliostat, combine to form the overall flux density distribution on the receiver. The position of these spots can be adjusted via heliostat aiming strategies, but the effectiveness of any aimpoint distribution algorithm depends on precise knowledge of each heliostat's focal spot characteristics.

Existing methods for predicting these focal spot characteristics present several limitations. Simplified models, with idealized heliostat assumptions, often fail to accurately represent the behavior of the heliostat, resulting in poor flux distribution predictions. Alternatively, advanced measurement techniques like deflectometry or photogrammetry can significantly improve prediction accuracy but are time-intensive, complex, and expensive, making them impractical for large-scale commercial CST plants.

To address these limitations, current research focuses on developing beam characterization methods that use operational data, particularly calibration data, to characterize heliostats. Heliostat calibration, a routine process in many CST plants, provides images of isolated focal spots, offering an opportunity to bypass the need for additional measurements. However, these calibration-based methods face their own challenges. The flux concentration between the heliostat mirrors and the calibration target leads to information loss, making it difficult to reconstruct the heliostat mirror properties with the same level of detail as more direct measurement techniques. Moreover, methods that attempt to derive complex surface deformations from calibration data often struggle with

robustness and stability, particularly over large distances.

— The challenge lies in developing a methodology that can accurately predict flux density while addressing the drawbacks of current approaches. The method must leverage operational data to provide precise predictions without additional measurements. It must also be robust, scalable, and capable of generalizing across a wide range of heliostat conditions and operational scenarios.

1.3. Scientific Objective

Addressing the outlined challenges requires a paradigm shift in how heliostat characterization and flux density predictions are approached. Traditional methodologies relying on extensive measurements and intricate modeling have proven to be costly, time-intensive, and difficult to scale. Therefore, the critical question becomes: can a purely data-driven approach, leveraging routinely collected calibration data, achieve the predictive accuracy necessary for practical deployment in operational CST plants?

This research develops and validates a novel methodology leveraging generative neural networks to predict flux distributions. By eliminating the dependency on detailed physical models and additional measurement campaigns, this approach focuses on scalability and cost-effectiveness while maintaining high predictive accuracy.

A key innovation of this methodology lies in its ability to overcome the limited information contained in focal spot data by integrating heliostat characterization and flux prediction within a unified, end-to-end framework. Instead of separately addressing these processes, the proposed framework enables the network to learn abstract representations of heliostat behavior while simultaneously predicting the flux density distribution. This integration allows the model to compensate for the information loss inherent in focal spot-based characterization, leveraging shared patterns and relationships across the entire heliostat field in one unified field model to enhance predictive performance.

By transforming the heliostat flux prediction process into a data-driven paradigm, this work aims to demonstrate that machine learning techniques can not only match but potentially surpass the performance of traditional methods in terms of accuracy, robustness, and scalability. The following sections outline the methodology, key steps, and scientific contributions of this research.

Methodology

As illustrated in Figure 1.1, the proposed methodology leverages calibration data obtained from routine heliostat calibration procedures using dedicated calibration targets. These images, which capture isolated focal spots, form the foundation for training a neural network capable of accurately predicting flux distributions. The methodology is structured into the following key steps:

1. **Data Acquisition, Prefiltering, and Preprocessing:** Calibration images depicting the concentrated solar flux emitted by individual heliostats are collected as part of the routine calibration process. These raw images undergo a prefiltering followed by preprocessing to extract relevant features to serve as optimal inputs for neural network training.

2. **Neural Network Training:** The neural network is designed to learn abstract representations of heliostat characteristics, capturing essential information required

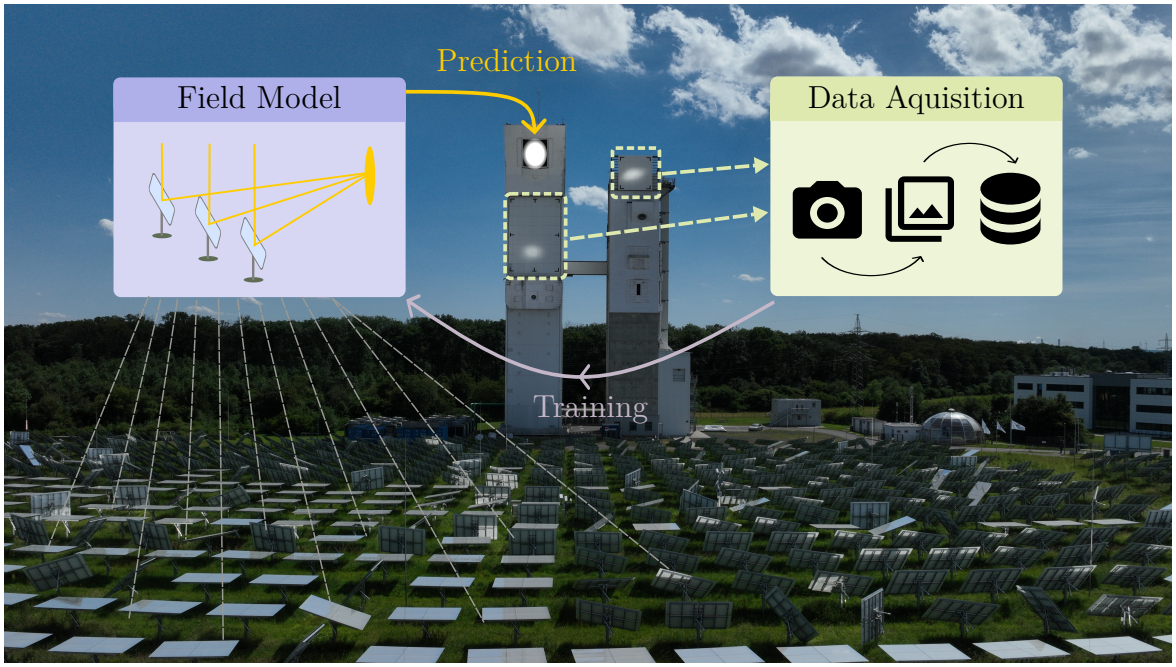

Figure 1.1.: Schematic of the proposed methodology. Calibration target images are used to train a heliostat field model, enabling accurate prediction of flux distribution on the receiver surface across varying operating conditions.

for flux prediction. Generative models are employed to address the information loss inherent in the transition from heliostat mirrors to the calibration target. This training process allows the network to encapsulate the unique and shared characteristics of individual heliostats while generalizing across diverse conditions.

3. **Prediction and Generalization:** Once trained, the neural network predicts the flux distribution for the entire heliostat field under varying operational conditions. The robustness and generalizability of the model are evaluated using data from the Solar Tower Jülich. This step ensures the model's reliability across a wide range of heliostats and input scenarios, including previously unseen conditions.

This methodology significantly reduces reliance on physical modeling, which is often error-prone and difficult to reconcile with real-world data. By focusing on operational data such as calibration images, the approach offers a practical and scalable solution for predicting flux distributions in large heliostat fields. It bridges the gap between traditional, measurement-intensive methods and modern, data-driven techniques, paving the way for more efficient and cost-effective CST operations.

Scientific Contributions

This research advances the field of CST by proposing a accurate and cost-effective solution for heliostat characterization and flux prediction. The key scientific contributions include:

- Development of a purely data-driven methodology for predicting heliostat flux distributions, eliminating the dependency on detailed physical models and additional measurements, thereby reducing complexity and cost.

- Demonstration of the ability of generative neural networks to mitigate information loss inherent in focal spot-based characterization approaches. The proposed methodology achieves high accuracy by focusing on abstract representations of heliostat behavior, rather than relying on direct physical modeling.
- Introduction of a unified, field-wide generative model that integrates data from the entire heliostat field. This approach enables synergies through shared knowledge of common errors and characteristics across heliostats, significantly enhancing scalability and robustness. For the first time, this generalization capability allows a single model to combine and leverage data from the entire heliostat field, yielding efficiencies not achievable with individual heliostat models.
- Validation of the proposed methodology using real-world data from the Solar Tower Jülich. This demonstrates the method's practical applicability and ensures that it meets the operational demands of CST plants.

By emphasizing data-driven approaches and minimizing reliance on physical modeling, this research not only offers a robust and practical solution to heliostat characterization but also establishes a foundation for broader applications of machine learning in CST and related fields. The proposed methods highlight the transformative potential of data-driven techniques in replacing traditional, measurement-intensive approaches, particularly in scenarios where physical models struggle to reconcile with real-world complexities. This contribution extends to other domains of solar energy and engineering, paving the way for scalable, efficient, and accurate predictive frameworks.

1.4. Thesis Structure

This thesis is organized to provide a logical progression from the identification of challenges in CST to the development, evaluation, and impact of the proposed methodology. The structure is outlined as follows:

- **Chapter 1: Introduction** - Presents the motivation, problem statement, and scientific objectives. It highlights the challenges in heliostat characterization and introduces the proposed data-driven approach.
- **Chapter 2: Machine Learning Fundamentals** - Explores the foundational principles of machine learning relevant to this research, including dataset preparation, loss functions, and generalization. It also introduces key architectural concepts.
- **Chapter 3: State of the Art** - Reviews current technologies and methodologies in CST, with a focus on heliostat calibration and flux prediction. It critically examines the limitations of existing approaches.
- **Chapter 4: Methodology** - Details the proposed data-driven framework, including the transformation of raw images into flux distributions, generative model training, and techniques for enforcing specific predictive behaviors.
- **Chapter 5: Dataset Characteristics, Preparation, and Test Scenarios** - Provides an in-depth description of the datasets used for model training and validation, encompassing both calibration data and artificially generated datasets. It also covers the design of test scenarios.

- **Chapter 6: Evaluation** - Evaluates the proposed methodology, focusing on prediction accuracy, generalization capabilities, and robustness under diverse conditions.
- **Chapter 7: System-Level Impact Analysis** - Evaluates the impact of the proposed methodology on the performance of CST plants, emphasizing improvements in flux prediction accuracy for the aggregated flux density of all heliostats and its implications for overall plant efficiency and operational optimization.
- **Chapter 8: Conclusion and Outlook** - Summarizes the key findings, compares the results to the state of the art, and discusses limitations and potential improvements. Directions for future research are also proposed.

This structure provides a clear and coherent narrative, guiding the reader through the identification of challenges, the development and validation of the proposed methodology, and its broader implications for CST systems.

Chapter 2.

Machine Learning Fundamentals

This work aims to overcome existing challenges in heliostat characterization and flux prediction by employing advanced machine learning (ML) methodologies. By leveraging ML-driven approaches, the proposed framework seeks to enhance predictive accuracy, robustness, and scalability, ensuring adaptability to diverse operational conditions in real-world CST systems.

ML has revolutionized computational science, providing powerful tools for extracting patterns and learning complex relationships from data. Among these tools, neural networks (NNs) have proven especially effective due to their capacity to approximate intricate functions and capture non-linear dependencies [24]. This chapter introduces foundational concepts in supervised learning, standard methodologies for training and evaluation, and key neural network architectures. These concepts form the basis for the data-driven approach developed in this thesis, addressing the unique challenges of heliostat flux prediction. Additionally, the chapter explores the critical concept of generalization, emphasizing its importance in ensuring the scalability and robustness of ML models for heliostat flux prediction tasks.

2.1. Learning Complex Functions with Neural Networks: Supervised Learning

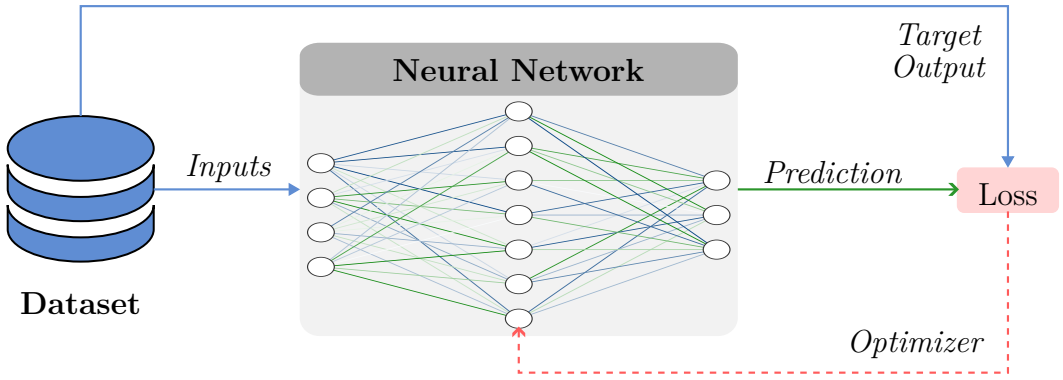

Figure 2.1.: Illustration of the supervised learning procedure, highlighting the interplay between input and output data, predictions, loss function, and weight update via the optimizer.

At the heart of neural networks is their ability to learn complex, high-dimensional functions directly from data. A neural network consists of interconnected layers of artificial neurons, parameterized by weights and biases. These parameters are adjusted during training to minimize the error between the network's predictions and the ground truth, which refers to the known correct output, in the case of supervised learning.

The learning process involves defining a loss function to quantify this error and employing optimization algorithms to iteratively adjust the parameters of the neural network. Through this process, the network approximates an underlying function that maps inputs to outputs, enabling predictions on unseen data. For instance, in a flux prediction context, the input might include sun positions and heliostat properties, while the output represents the flux distribution as an image.

2.2. Supervised Learning: Standard Workflow

Supervised learning is a cornerstone of machine learning (ML), where models are trained on labeled datasets to establish mappings between inputs and outputs [24]. Unlike other paradigms, supervised learning is particularly suited for applications requiring specific behavior and controlled outputs, making it the field closest to traditional engineering approaches for prediction tasks. By leveraging clearly defined input-output pairs, supervised learning enables precise and predictable model behavior, a critical factor in applications such as flux density prediction where accuracy and reliability are paramount. The workflow, illustrated in Figure 2.1, involves several interconnected components, each critical to the development of robust and accurate predictive models.

Datasets, Preprocessing, and Filtering

The foundation of supervised learning lies in the quality of the dataset, which consists of input-output pairs such as numerical vectors, images, or sequences. Effective model training requires well-preprocessed and curated data. Preprocessing and filtering steps often vary depending on the nature and source of the data but generally include:

- **Source-Specific Adjustments:** Tailored modifications such as cropping, resizing, or filtering signals to standardize inputs for subsequent processing.
- **Normalization:** Scaling features to consistent ranges to enhance numerical stability and training efficiency.
- **Outlier Detection and Removal:** Identifying and excluding anomalous data points that could skew the model's learning process.
- **Data Augmentation:** Introducing variations through transformations like rotations, flips, and intensity changes to expand dataset diversity and improve model robustness, particularly for image data.

Once prepared, the dataset is typically split into subsets for training, validation, and testing. The training set is used to optimize model parameters, the validation set monitors performance and aids in hyperparameter tuning, and the test set evaluates the model's generalization capability to unseen data. Properly structured datasets not only improve learning efficiency but also ensure reliable and unbiased evaluations.

Model Creation and Components

A neural network's architecture defines its structure and ability to learn patterns. At the most basic level, a model comprises:

- **Layers:** Such as convolutional layers for feature extraction and fully connected layers for global reasoning [24].
- **Activation Functions:** Introduce non-linearity to capture complex patterns [25].
- **Connectivity:** Determines data flow, e.g., in feedforward networks or attention-based models [26].

At a more abstract level, advanced models are constructed by combining components such as encoders, decoders, and attention layers to enable specific data flows. Selecting a suitable model architecture is often an iterative process, informed by domain knowledge and experimental results.

Loss Function and Optimization

The loss function quantifies the difference between a model's predictions and the ground truth, guiding the optimization process. Commonly used loss functions include:

- **Mean Squared Error (MSE):** Suitable for regression tasks, measuring the squared difference between predictions and actual values.
- **Cross-Entropy Loss:** Used in classification problems to compare predicted probabilities with true class labels.
- **Pixel-Wise Loss:** Applied in image tasks to measure differences on a per-pixel basis between predicted and target images.

Optimization involves iterative updates to the model's parameters to minimize the loss function. Techniques such as stochastic gradient descent (SGD) and Adam are commonly employed, with the choice often influenced by the problem's complexity, dataset characteristics, and computational resources.

Training and Evaluation

Training involves passing the dataset through the network multiple times, with each complete pass refining the model's parameters. Validation metrics, such as accuracy or loss, are monitored during training to detect overfitting or underfitting, enabling adjustments to hyperparameters like learning rates or regularization factors.

Final evaluation is conducted on the test set to measure the model's performance on unseen data. Metrics are chosen based on the task; for image-based tasks, pixel-wise accuracy or structural similarity measures might be used. A thorough evaluation ensures the model's robustness and practical applicability.

2.3. Neural Network Architectures

Neural network architectures define how data flows through a model, enabling it to perform specific tasks. The choice of architecture determines the network's capacity to process information and solve complex problems. In this section, we introduce three key architectural paradigms relevant to image data: image-to-image networks, conditional image generation, and encoder-decoder architectures. These paradigms employ distinct strategies to process and transform data, forming the foundation for the methods developed in this work.

Image-to-Image Networks

Image-to-image networks are designed to transform input images into corresponding output images, maintaining spatial alignment between the two. Among their applications, **segmentation** is particularly relevant for this work. Segmentation involves converting input images into pixel-wise masks, labeling regions of interest within the data [27]. This process is crucial for preprocessing and structuring data for downstream tasks, ensuring that meaningful features are extracted accurately.

A well-known example of an image-to-image network is the **UNet** architecture [27], which is particularly effective for segmentation tasks. UNet employs skip connections that link corresponding encoder and decoder layers, preserving spatial details while incorporating higher-level abstractions. This design enables precise transformations, making it especially suitable for tasks requiring pixel-level accuracy.

Image-to-image networks, and segmentation tasks in particular, form an essential foundation for processing and analyzing image data. Their ability to extract and structure relevant information prepares datasets for advanced machine learning applications, as will be discussed in subsequent chapters.

Conditional Image Generation

Conditional image generation focuses on creating realistic images guided by specific input conditions. This approach transforms low-dimensional input vectors (e.g., abstract representations or encoded conditions) into high-dimensional image outputs, ensuring alignment with the provided conditions. Prominent models in this paradigm include **StyleGAN** [28], **Latent Diffusion** models [29], and conditional generative adversarial networks (**cGANs**) [30].

For instance, **StyleGAN** generates images by progressively refining features across multiple layers, guided by a latent vector representing desired characteristics like style or texture. Conditional image generation is especially useful for tasks requiring precise control over output properties, such as creating images that adhere to specific conditions or constraints. This paradigm enables the synthesis of high-quality images tailored to specific input scenarios.

Encoder-Decoder Architectures

Encoder-decoder architectures are versatile tools for tasks that involve dimensionality reduction and transformation across data domains. The encoder compresses the input data into a latent representation, capturing essential features in a compact form. The decoder

reconstructs the output from this representation, enabling tasks like data generation, transformation, or restoration.

Variational Autoencoders (VAEs): VAEs are a probabilistic type of encoder-decoder architecture designed for dimensionality reduction and data generation [31]. The encoder maps input data into a latent space, while the decoder reconstructs the data from a sampled point in this space. VAEs are particularly well-suited for applications requiring smooth interpolation and controlled synthesis of new data by leveraging probabilistic latent representations. These characteristics make VAEs applicable to tasks involving abstract feature encoding and reconstruction.

Transformers: Transformers extend the encoder-decoder paradigm with attention mechanisms, enabling the modeling of complex relationships between data elements [26]. Unlike VAEs, which focus on dimensionality reduction, Transformers excel at capturing contextual dependencies. In image-based tasks, Transformers process sequences of image patches or encoded feature vectors, enabling fine-grained analysis and generation. Their modular design supports diverse tasks that require a combination of contextual understanding and flexible transformations, such as complex scene generation or data-to-image mappings.

— A thorough understanding of these neural network paradigms is critical for developing tailored architectures that address specific challenges, such as those encountered in CST applications. Each paradigm—whether focused on feature extraction, conditional generation, or dimensionality reduction—provides unique strengths that can be leveraged to design efficient and accurate solutions. By combining components such as encoders, decoders, and attention mechanisms, these architectures lay the groundwork for creating end-to-end frameworks.

In the context of CST, where accurate modeling and prediction are essential, leveraging these state-of-the-art architectures enables the development of tailored solutions for heliostat characterization and flux prediction. A deep understanding of their underlying principles allows for the adaptation and fine-tuning of neural network designs to meet the specific demands of CST applications.

2.4. Generalization in Machine Learning

Generalization is the cornerstone of effective machine learning models. It refers to a model's ability to perform well on unseen data by capturing underlying patterns in the training data rather than memorizing the examples. A well-generalized model should be robust to variations in input conditions and capable of adapting to scenarios not explicitly represented during training.

Generalization is particularly crucial in real-world applications where models must handle unseen conditions, noise, and variability in the data. For instance, a model trained to predict heliostat flux distributions should perform well across a wide range of sun positions and heliostat configurations, even when specific combinations were not part of the training dataset.

Factors Influencing Generalization

Achieving generalization involves careful consideration of several factors:

- **Data Diversity:** Ensuring the training dataset captures sufficient variability in

input conditions is fundamental to generalization. For heliostat flux prediction, this includes incorporating diverse sun positions, heliostat orientations, and aim points. Diverse data allows the model to infer patterns and relationships rather than relying on specific examples.

- **Regularization Techniques:** Regularization mitigates overfitting by introducing constraints or penalties, employing methods such as dropout, weight decay, and data augmentation to enhance the model's ability to generalize [32].
- **Model Design:** Designing the model architecture plays a critical role in achieving generalization by balancing flexibility to learn complex patterns with domain-specific constraints that guide and optimize the learning process.

Measuring Generalization

Evaluating generalization requires assessing the model's performance on unseen data. Standard practices include:

- **Validation and Test Sets:** Separate datasets ensure the model's performance is evaluated on data not used during training.
- **Cross-Validation:** Partitioning the dataset into multiple subsets to train and test iteratively, reducing dependence on a single train-test split.
- **Edge-of-Distribution (EOD) Analysis:** Evaluating performance on rare or extreme input conditions to ensure robustness in edge cases.

Generalization in Heliostat Flux Prediction

In the context of heliostat flux prediction, generalization refers to a model's ability to perform accurately under diverse and unseen operational conditions. Specifically, a well-generalized model should:

- **Adapt to Varying Sun Positions:** Accurately predict flux distributions for sun positions that differ from those represented in the training data.
- **Handle Heliostat-Specific Variations:** Account for differences in heliostat properties, such as mirror quality, reflectivity, and canting, without requiring extensive data for each heliostat.

Generalization in this context ensures that a single predictive model can be applied across the entire heliostat field and operate effectively under changing environmental and system conditions. This capability allows the model to make full use of the diverse data available from heliostat field operations, laying a strong foundation for robust and scalable applications in flux prediction.

—

This chapter has introduced foundational concepts of machine learning, focusing on supervised learning, neural network architectures, and the critical principle of generalization. These concepts form the backbone of predictive modeling, enabling the development of robust and accurate solutions tailored to complex tasks. A detailed understanding of supervised learning workflows, along with the strengths and applications of various neural network paradigms, provides the necessary groundwork for designing and implementing models capable of addressing real-world challenges.

Chapter 3.

State of the Art

A thorough understanding of the state of the art in solar tower technology and heliostat flux prediction is essential for identifying existing challenges and driving advancements in the field. This chapter provides a structured overview of established methodologies and technologies, with a particular focus on the heliostat field—a critical component of solar tower systems.

Section 3.1 introduces the fundamental principles and components of a solar tower plant, outlining its operational framework and the unique challenges associated with heliostat fields. This serves as a foundation for Section 3.2, which delves into the plant's operational levels, emphasizing the interplay of subsystems required to ensure safe and efficient performance.

The focus then shifts to heliostat modeling for flux prediction. Section 3.3 examines heliostat calibration procedures, which provide the operational data essential to this research. Building on this, Section 3.4 presents a comprehensive review of existing methods for heliostat flux prediction, structured around three core aspects:

- **Focal Spot Computation Methods**, which simulate the spatial distribution of concentrated flux by a heliostat.
- **Heliostat Characterization Techniques**: These methods involve deriving specific properties of heliostatstrough additional measurements to enhance the accuracy of flux predictions
- **Efficiency Modeling**, which simulates optical losses to enable the precise scaling of absolute flux values, ensuring realistic flux predictions.

By systematically analyzing these methodologies, the chapter identifies their limitations and outlines key research gaps. These insights underscore the need for a purely data-driven approach to heliostat flux prediction, providing the foundation for the novel methodologies developed in subsequent chapters.

3.1. The Solar Tower Plant

Solar Tower Power Plants (STPs) represent a pivotal advancement in Concentrated Solar Power (CSP) technologies, offering higher concentration ratios and facilitating higher operating temperatures compared to other CSP technologies such as Fresnel or parabolic trough systems. Using a field of mirrors to concentrate sunlight onto a central receiver atop a tower, as shown in Figure 3.1, this configuration enables efficient concentration of solar energy onto a receiver. The receiver, positioned atop a central tower, absorbs this concentrated solar irradiation for various purposes, including electricity generation

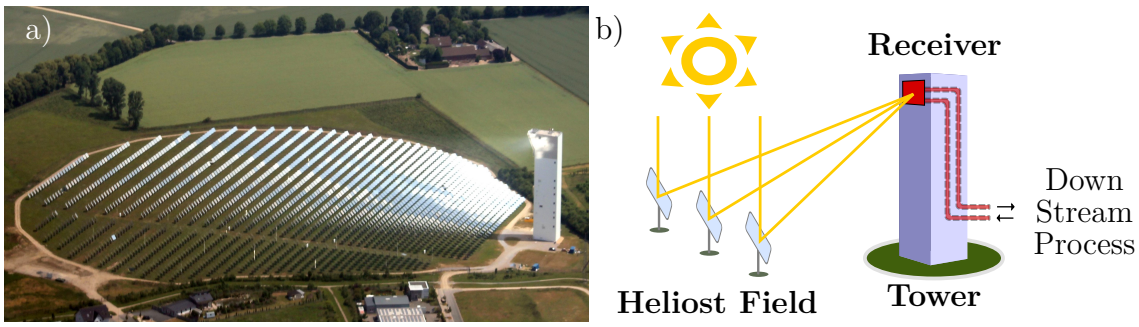

Figure 3.1.: a) The Solar Tower Jülich (STJ), a research facility. b) Schematic illustration of the operating principle of a solar tower plant.

through thermal power blocks, fuel production via chemical reactors, or other downstream processes utilizing the generated high-temperature heat. These advantages make STPs a prominent focus of research and development.

The Solar Tower Jülich (STJ)[33], depicted in Figure 3.1a), is a dedicated research facility operated by the German Aerospace Center (DLR). Built in 2008 in Jülich, Germany, the STJ serves as a testbed for advancing solar tower technology and provides the foundation for the experimental and methodological investigations in this work. The facility features two towers: the left tower houses a cylindrical open volumetric receiver operated with hot air, along with calibration targets for heliostat alignment and experimental installations, while the right tower includes additional research levels and a calibration target at its summit. Given the scope of this work, the focus remains on the heliostat field, excluding downstream processes such as heat utilization.

Heliostat Field

The heliostat field is a critical component of any Solar Tower Plant (STP), comprising hundreds or even thousands of heliostats, depending on the plant's size and design. Each heliostat consists of a support structure holding multiple mirror facets, controlled by actuators that allow precise adjustment along two axes. These adjustments enable accurate solar tracking throughout the day, ensuring that sunlight is consistently focused on the receiver for optimal energy concentration.

To achieve high concentration ratios, the precise alignment of heliostat facets is essential. This alignment is achieved through two primary adjustments, which significantly influence the concentrating properties [34]:

- **Canting:** The alignment of individual mirror facets to converge reflected light at a common focal point.
- **Focusing:** The adjustment of facet curvature to match the focal distance between the heliostat and the receiver.

At the Solar Tower Jülich (STJ), the heliostat field consists of over 2,000 heliostats arranged in a specific layout, as depicted in Figure 3.1a). Each heliostat is equipped with four flat mirror facets, collectively providing a total reflective area of 7.68 m^2. The facets are designed to be flat, except for mirror errors, with each facet canted to align with the receiver.

The Optical Properties of Solar Concentrators

The efficiency of heliostat fields is heavily influenced by optical errors, which directly impact the performance of solar concentrators. These errors, outlined below, stem from manufacturing imperfections, environmental factors, and operational wear [35, 36]:

- **Shape Errors:** Deviations in the mirror surface from its optimal geometric form, resulting from manufacturing defects, thermal expansion, wind loads, or gravitational effects. These inaccuracies misdirect reflected rays, reducing the overall concentration efficiency.
- **Canting Errors:** Misalignments of individual mirror facets, occurring during installation or due to structural deformations over time. Precise canting is critical for ensuring that sunlight is accurately directed toward the receiver.
- **Specularity Errors:** Surface roughness or degradation of reflective materials scatters light, reducing the flux concentration on the receiver. These errors are often exacerbated by factors such as dust, scratches, or material aging.
- **Tracking Errors:** Inaccuracies in the heliostat's tracking system, caused by mechanical faults, software issues, or environmental influences such as wind. These errors hinder the heliostat's ability to consistently maintain focus on the receiver.

While high-precision optical components could mitigate these errors, their cost makes them impractical for large-scale commercial CST plants. Consequently, a balance between cost and performance must be struck, with these imperfections carefully accounted for in the design and operation of solar concentrators [35].

3.2. Operational Levels of a Solar Tower Plant

After sunlight is concentrated by individual heliostat mirrors, further concentration occurs at the system level through the superposition of focal spots from multiple heliostats onto the receiver. The resulting total flux density distribution is a critical factor influencing the operational efficiency, durability, and safety of a CSP plant. The primary challenge lies in achieving the optimal flux distribution despite heliostat inaccuracies and dynamic environmental conditions. To address this challenge, CSP plants rely on a control loop that continuously monitors and adjusts the flux distribution, as illustrated in Figure 3.2. This operational control loop comprises two primary components: the measurement side and the prediction side.

Control Loop for Flux Density Management

On the **Measurement Side**, the condition of the receiver is continuously monitored. The incident flux from the heliostat field can be measured using flux gauges located in the radiation shield near the receiver [37], or through more recent techniques that analyze the reflection of flux from the receiver surface [38]. Additionally, an infrared (IR) camera is often employed to monitor the receiver's surface temperature, detecting hotspots that could lead to damage [39].

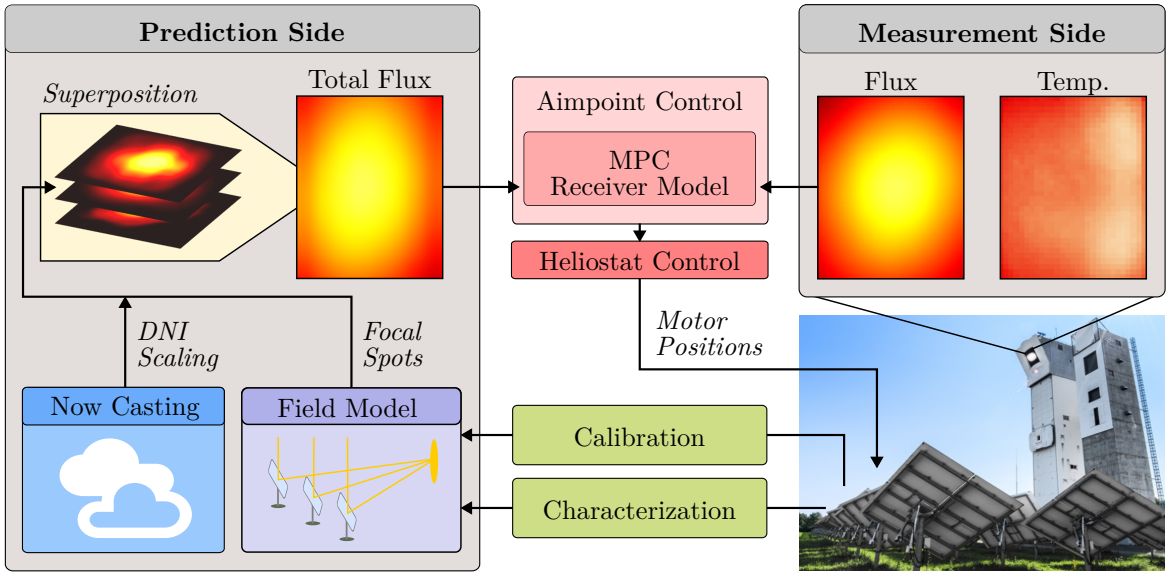

Figure 3.2.: Process levels of a CSP plant. The image illustrates how measurements and predictions are integrated into a control loop to dynamically adjust heliostat aiming to optimize the flux density and overall plant efficiency.

On the **Prediction Side**, a heliostat field model is used to simulate the heliostat focal spots, which are then superimposed to predict the total flux distribution on the receiver. Nowcasting systems track cloud shadows to predict spatially resolved direct normal irradiation (DNI) on the ground, enabling adjustments to the intensity of the heliostat focal spots for more accurate flux predictions [40, 41].

Both measurement and prediction outputs feed into an **Aimpoint Control Algorithm (APCA)**, which serves as the core decision-making component. The primary objectives of the APCA are to maximize receiver efficiency while avoiding localized flux and temperature peaks on the receiver surface [20, 21, 42]. The APCA achieves this by dynamically adjusting the orientation of the heliostats, thereby controlling the spatial distribution of the total incident flux.

The APCA typically incorporates a thermal model of the receiver to simulate the impact of different flux distributions on thermal performance. Using model predictive control (MPC) [18, 19], the APCA determines optimal aimpoints for all heliostats, ensuring the receiver operates within its thermal constraints while maintaining high efficiency. These calculated aimpoints are then transmitted to the heliostat control system, which adjusts the motor positions of each heliostat to achieve the desired orientation.

Importance of Accurate Flux Prediction Models

While measurements provide real-time feedback, flux prediction models are essential for achieving precise control and maximizing efficiency. Their importance is underscored by two key factors:

- **Limitations of Measurement Systems:** Flux measurement systems are inherently prone to significant uncertainties, typically ranging from 10% to 40% [37],

making them insufficiently reliable for precise aimpoint control. Additionally, direct flux measurements are often impractical for certain receiver configurations, such as cavity receivers or reactors, where geometric constraints limit the applicability of camera-based measurement techniques.

- **Optimizing APCA Performance:** Accurate predictions of heliostat focal spot shapes significantly improve the APCA's ability to achieve optimal flux distributions. Studies show that incorporating accurate focal spot data can enhance APCA efficiency by up to 20% [43].

To address these challenges, a robust digital twin of the heliostat field is essential. Such a model must not only predict focal spots with high accuracy but also account for temporal variations in heliostat properties caused by environmental factors.

Heliostat Models for Flux Prediction

Heliostat field models typically consist of individual models for each heliostat, with each model comprising two primary components:

- **Alignment Model:** This component calculates the position of the focal spot center by translating heliostat motor positions into precise mirror orientations. To ensure alignment accuracy, regular calibration procedures, such as the camera-target method (detailed in Section 3.3), are performed. These calibrations account for changes in heliostat geometry caused by environmental factors such as weather or mechanical wear, ensuring the model reflects current heliostat conditions accurately.

- **Flux Model:** This component predicts the spatial distribution of flux around the focal spot center, incorporating factors such as heliostat mirror properties, sun positions, and heliostat orientations. Unlike alignment models, flux models often rely on additional, non-automated measurements—such as deflectometry or photogrammetry—to account for mirror errors. While these advanced techniques provide higher accuracy, they are complex, costly, and challenging to implement on a large scale, as further discussed in Section 3.4.

Both components are integral to accurate flux prediction and will be elaborated upon in the subsequent sections, including their methodologies and roles in operational CSP plants.

3.3. Heliostat Calibration Procedure

The heliostat calibration procedure is a cornerstone of ensuring precise heliostat alignment, which is critical for the efficient operation of solar power plants. Among the available calibration methods, the widely adopted *Camera-Target Method* is renowned for its accuracy and reliability and is also the most widely used approach in commercial applications [44].

In this approach, each heliostat is sequentially redirected from its operational aimpoint on the receiver to a nearby white calibration target. A strategically positioned camera captures images of the focal spots generated by the heliostat under varying sun positions throughout the year. These images are used to determine the center of intensity of the

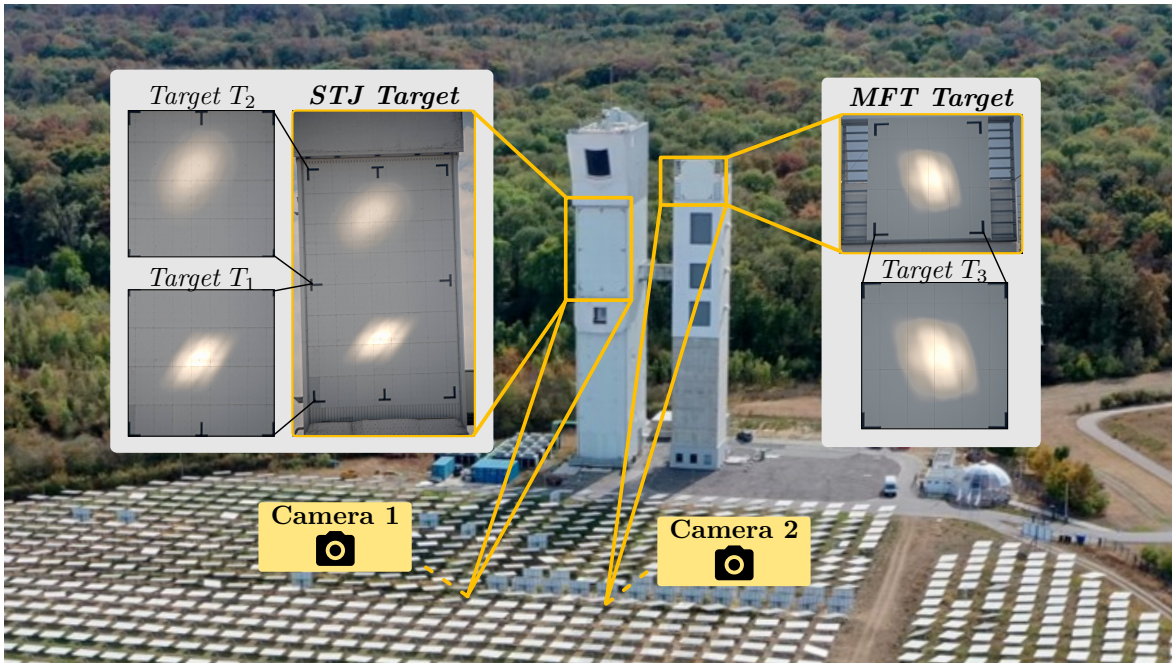

Figure 3.3.: Aerial view of the STJ showing the two towers of the plant. The locations of the calibration cameras are indicated. Insets show sample images from both cameras and the processed cropped images for each calibration target.

reflected sunlight, or focal point, enabling the measurement of the heliostat's pointing deviation—the discrepancy between the aimpoint set by the control system and the actual measured point.

The calibration process typically involves fitting a rigid-body regression model to the collected calibration data points. This widely established method adjusts the heliostat model's geometric parameters, correcting misalignments and ensuring precise solar tracking. Recently, alternative approaches, such as neural network-based calibration models, have emerged as promising research areas. Regardless of the specific methodology, frequent calibration is essential to account for ongoing changes in heliostat geometry caused by environmental factors or mechanical wear. By collecting data across various sun positions throughout the year, the calibration procedure maintains accurate heliostat tracking and adapts to dynamic variations in heliostat performance.

As illustrated in Figure 3.3, the STJ is equipped with three calibration targets, enabling the simultaneous calibration of multiple heliostats at varying aimpoints. The data collected during this calibration process serves as the foundation for the data-driven methodology employed in this work.

3.4. Heliostat Flux Prediction Methods

Accurate flux prediction is a critical aspect of heliostat modeling, complementing alignment prediction by determining the spatial distribution and intensity of reflected sunlight on the receiver. As illustrated in Figure 3.4, the flux prediction process can be divided into three interconnected components: *Focal Spot Computation Methods*, *Heliostat Characterization Techniques*, and *Heliostat Efficiency Modeling*.

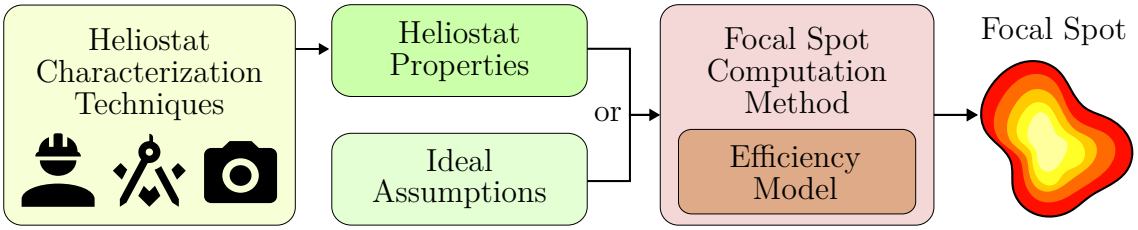

Figure 3.4.: Flux prediction process: Breakdown into its individual components.

The **first component**, *Focal Spot Computation Methods* (Section 3.4.1), encompasses methods to compute the shape and intensity distribution of focal spots, or flux distributions. These methods range from simplified geometric models to advanced approaches that account for optical imperfections. Those methods may rely on idealized assumptions or incorporate additional measurements to enhance accuracy.

The **second component**, *Heliostat Characterization Techniques* (Section 3.4.2), focuses on obtaining accurate data about real-world heliostat properties, such as shape and canting errors. These techniques, including deflectometry and photogrammetry, play a pivotal role in bridging the gap between theoretical models and actual operational conditions.

The **third component**, *Heliostat Efficiency Modeling* (Section 3.4.3), addresses the scaling of predicted flux to realistic values by accounting for energy losses caused by shading, blocking, and atmospheric extinction. These models ensure that predictions reflect the true performance of the heliostat field under varying environmental conditions.

Together, these components form a cohesive framework for generating precise and reliable heliostat flux predictions. The following sections delve into each of these components, examining the current methodologies, their limitations, and potential improvements.

3.4.1. Focal Spot Computation Methods

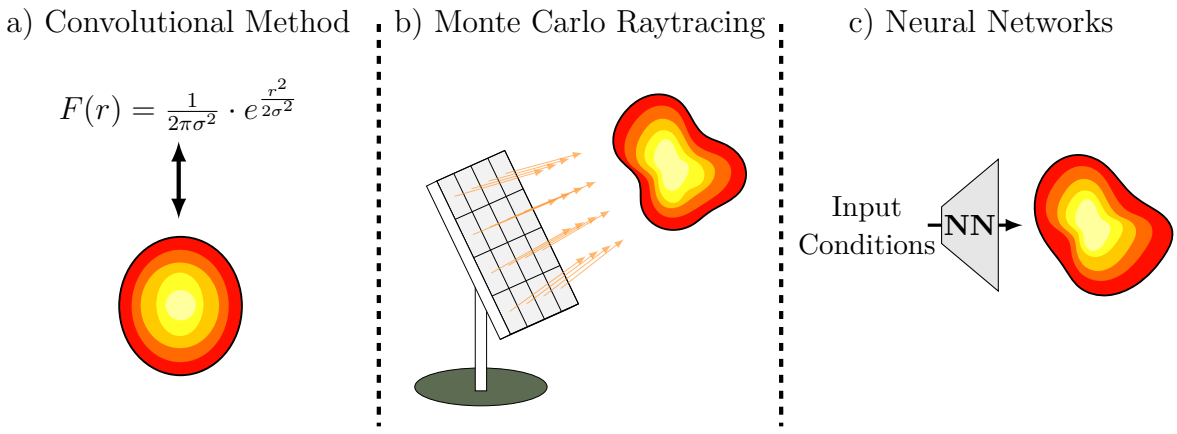

Figure 3.5.: Heliostat Focal Spot Computation Methods: a) Convolutional Methods, b) Monte Carlo Raytracer, c) Generative Neural Networks

Focal Spot Computation Methods are essential tools for simulating the optical behavior of heliostats and predicting variations in flux distributions based on a range of input

conditions, such as sun position and heliostat orientation. These methods are pivotal in optimizing the performance of solar power tower systems by providing accurate and reliable flux predictions.

The primary objective of these methods is to model how heliostats reflect and focus sunlight onto a target receiver. This involves accounting for factors that influence flux distribution, including the geometry of the heliostat, the size, position, and shape of mirror facets, as well as the canting and alignment of these facets. Predictions may be generated under idealized assumptions or enhanced through additional measurements that capture mirror shape and canting errors.

Under idealized assumptions, methods typically consider perfect geometric configurations, assuming that mirror facets are perfectly shaped and precisely positioned. However, real-world conditions often involve deviations caused by manufacturing tolerances, installation inaccuracies, and operational factors. To address these challenges, advanced methods incorporate empirical data and measurements, which account for these imperfections and result in more accurate flux predictions.

The Focal Spot Computation Methods discussed in this section include *Convolutional Methods*, *Monte Carlo Raytracing*, and *Neural Networks*. Each approach provides a unique methodology for predicting flux, addressing the complexities and subtleties of heliostat performance. Figure 3.5 summarizes these methods through illustrative examples, which are elaborated upon in the following sections.

Convolutional Methods

Convolutional methods for heliostat simulation utilize simplified mathematical models to approximate the radiative flux distribution on receivers in solar thermal power plants. These methods employ convolution integrals to model the distribution of reflected sunlight, enabling efficient and scalable simulations of heliostat fields.

The modeling process begins by representing the distribution of incident sunlight, often using an idealized Gaussian distribution or more precise sun shapes that account for the non-uniform intensity of sunlight. Subsequently, various error sources, such as tracking errors, astigmatism, and surface curvature-induced beam quality degradation, are incorporated. These errors are modeled as a broadening of the focal spot shapes.

A notable example of convolutional methods is the HFLCAL model [45], which simplifies the reflected image of each heliostat into a circular normal distribution. This approach aggregates multiple optical errors, including mirror quality and tracking inaccuracies, into a single characteristic parameter, σ. The model assumes well-canted concentrating facets with spherical curvature, and the flux distribution is effectively described by a single circular normal distribution. This simplification is supported by the Central Limit Theorem, which states that the aggregation of multiple arbitrarily distributed quantities converges to a normal distribution under certain conditions.

The convolution process maps the reflected light from each heliostat onto an image plane before projecting it onto the receiver. This intermediate step captures geometric distortions and variations in incident angles with greater accuracy. For instance, the iCauchy model [46] improves upon traditional convolutional methods by accurately projecting the effective reflecting surface onto the image plane and applying a quasi-Cauchy kernel for convolution. This refinement results in more precise flux distribution simulations, particularly for varying incident angles.

Convolutional methods can operate under idealized assumptions or incorporate empir-

ical measurements to enhance accuracy. Tracking errors and mirror quality, for example, can be directly measured and integrated into the model, thereby improving the precision of the simulations.

The primary advantages of convolutional methods lie in their computational efficiency and scalability, making them valuable tools for designing, optimizing, and monitoring the performance of large-scale heliostat fields. These methods also facilitate real-time simulation adjustments, which are crucial for the dynamic operation of solar power plants. However, they may not capture intricate details of flux distributions, particularly in scenarios involving significant geometric distortions of heliostat mirrors. Simplified assumptions, such as idealized distributions and isotropic scattering, may overlook complex interactions, and the accuracy of these simulations heavily depends on the precise estimation of model parameters [47, 48].

Monte Carlo Ray Tracing

Monte Carlo ray tracing is another powerful simulation method used to model heliostat focal spots in solar thermal power plants. Unlike convolutional methods, which rely on mathematical simplifications, Monte Carlo ray tracing adopts a probabilistic approach to simulate the interaction of light rays with reflective surfaces, offering a detailed and highly flexible analysis of optical systems [49, 50].

This method begins by defining a large number of rays, which are then traced through the optical system, interacting with surfaces through processes such as reflection, refraction, and scattering. At each interaction, the ray's path is modified based on the surface properties it encounters. By replicating real photon interactions, Monte Carlo ray tracing achieves a high level of accuracy, making it particularly suitable for modeling complex optical systems where precise flux distributions are required.

The light source in Monte Carlo ray tracing can be modeled using various sun shapes, such as Gaussian or pillbox distributions, or even user-defined intensity profiles. This flexibility enables the accurate representation of solar irradiation under diverse atmospheric conditions and angles of incidence. The rays interact with heliostat surfaces modeled with detailed properties, including slope errors, surface roughness, and tracking inaccuracies. These imperfections are incorporated as probabilistic deviations, broadening the reflected beam and impacting the resulting flux distribution on the receiver.

Rays are typically generated on each heliostat surface. According to the sun's position and the normal vector of the reflecting surface element, the direction of reflection is determined. For this direction, many rays from the chosen sun distribution are sampled. Each ray is tracked from the heliostat surface to the aperture plane or receiver surface, allowing for detailed shading and blocking analysis between the heliostats. The consideration of single rays also allows for the analysis of complex receiver surface geometries, as the exact impact point of each ray with the receiver surface can be calculated. Additionally, blocking of rays within the receiver or surrounding support structures or multiple reflections within the receiver can be tracked, enabling the analysis of complex receivers like cavity receivers or reactors. Ray tracers can work either with ideal assumptions about the heliostats or incorporate measurements of surface deformations into the ray generation process. This adaptability enhances the accuracy and relevance of the simulations.

The primary advantages of Monte Carlo ray tracing lie in its ability to handle complex, non-ideal optical systems and its high accuracy in predicting flux distributions. It accommodates detailed surface properties, shading and blocking effects, and multiple re-

flections, making it invaluable for in-depth analysis of heliostat fields and receiver designs. However, this method is computationally intensive, with simulation time increasing significantly as system complexity and the number of rays grow. This high computational demand makes Monte Carlo ray tracing less suitable for quick optimizations or real-time simulations, particularly in systems with large heliostat fields [47, 48].

Generative Neural Networks

Generative neural networks present a novel approach for predicting heliostat focal spots, leveraging input conditions such as sun positions to estimate flux distributions. These networks are trained on images of focal spots, enabling them to learn intricate patterns and relationships without requiring detailed physical modeling of the heliostat, as showcased by Pargmann et al. [51].

In contrast to convolutional methods and Monte Carlo ray tracing, neural networks function as a data-driven black box model. They do not rely on explicit physical principles but instead learn from examples of focal spots under varying conditions. Pargmann et al. trained neural networks on synthetic data generated using a ray tracer. While the dataset used represented an ideal heliostat and was limited in size, the results revealed the significant potential of this approach for flux prediction.

The black box nature of neural networks presents both advantages and disadvantages. A major benefit is the elimination of the need for detailed physical models of the heliostats. This can significantly simplify the simulation process and reduce computational costs. However, the main drawback is the dependency on large amounts of high-quality data. Without extensive and diverse training datasets, the neural network may struggle to generalize to new conditions or configurations, affecting its accuracy and reliability.

3.4.2. Heliostat Characterization Techniques

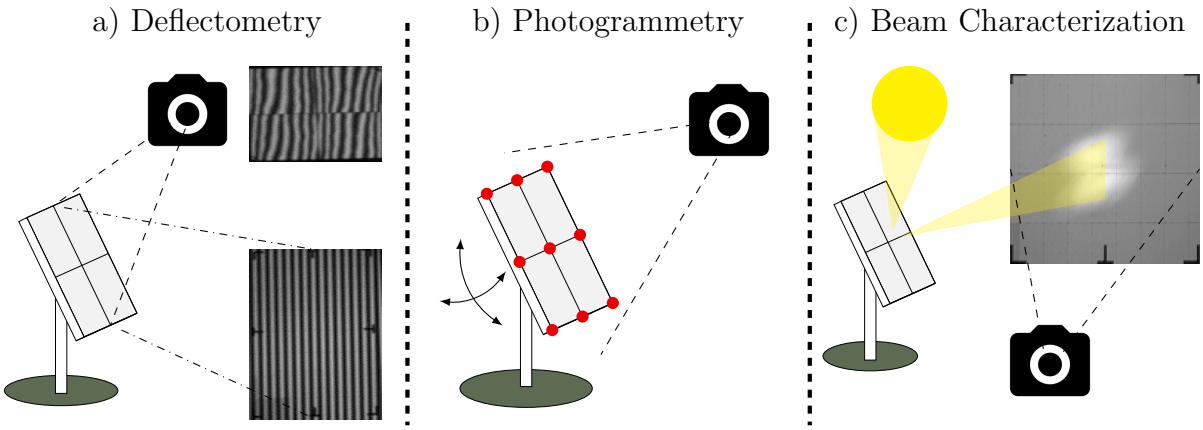

Figure 3.6.: Different techniques to characterize a heliostat mirror surface: a) Deflectometry - Measure deformations of a stripe pattern; b) Photogrammetry - Track fixed points while moving the heliostat; c) Beam Characterization - Measure deformations in the reflected focal spot.

Flux prediction methods as described before, while valuable, often rely on idealized assumptions about heliostat properties such as perfect geometry, mirror facet size, positions, and shapes. These assumptions can limit the accuracy and reliability of flux

predictions, particularly in real-world applications where various imperfections and operational conditions come into play. To enhance the precision of flux predictions and ensure optimal performance of solar thermal power plants, it is crucial to incorporate additional measurements that characterize the actual properties of heliostats and track changes in these properties over time.

Heliostat characterization involves detailed examination and measurement of components relevant to the flux prediction process. These measurements provide critical insights into practical deviations and errors, such as surface deformations and other imperfections. By understanding and quantifying these factors, flux prediction models can be refined to better reflect real-world conditions, leading to more accurate and reliable flux predictions and optimizations.

In this section, we will explore key techniques used for heliostat characterization, including *Deflectometry*, *Photogrammetry*, and *Beam Characterization*. These methods provide insights into mirror surface conditions, geometric accuracy, and focal spot quality, respectively. Each technique is illustrated in Figure 3.6. The following subsections will delve into each of these techniques, explaining their methodologies, applications, and potential to enhance flux prediction.

Deflectometry

Deflectometry is a precise technique for measuring slope errors and deformations in heliostat mirrors. This method involves projecting a stripe pattern onto a reflective surface and analyzing the distortions in the reflected pattern to reconstruct the mirror's surface profile (see Figure 3.6a). By capturing and processing these distortions, deflectometry provides detailed surface information, which is utilized for accurate flux prediction [52]. Precise surface profiles are primarily employed in ray tracing computations, while the mean slope error derived from these measurements can also be used in the convolutional approach.

The measurement procedure begins with projecting a regular stripe pattern onto the heliostat mirror, followed by capturing the reflection with a camera. Analyzing the deviations in the reflected pattern reveals surface irregularities and slope errors. This process involves calculating surface normals from the pattern distortions, enabling a comprehensive reconstruction of the mirror's surface. Accurate results necessitate precise calibration and are mathematically demanding [53].

Deflectometry is distinguished by its high accuracy, offering measurement resolutions of approximately 1 million points per heliostat with uncertainties below 0.2 milliradians. This granularity allows for the detection of minute surface irregularities, facilitating the creation of detailed heliostat models. These surface profiles can be integrated into Monte Carlo ray tracing simulations, enhancing flux predictions by accounting for real-world heliostat imperfections [52].

Despite its advantages, field application of deflectometry presents significant challenges. The process is complex, time-consuming, and requires meticulous setup and calibration, limiting its efficiency in large heliostat fields. Moreover, the method's sensitivity to environmental factors such as dew, dust, and wind can impair measurement accuracy and reliability in field conditions. These challenges hinder the widespread adoption of deflectometry in operational settings [53].

In summary, while deflectometry provides high-resolution surface characterization for advanced flux prediction, its practical application in the field is constrained by envi-

ronmental sensitivities and the complexity of the measurement process, restricting its scalability [35, 52, 53].

Photogrammetry

Photogrammetry is a non-contact, highly accurate technique for characterizing the surface shape of heliostat mirrors by capturing three-dimensional surface data. Unlike deflectometry, photogrammetry relies on simple photographs taken from multiple angles or during heliostat movement to reconstruct the mirror's three-dimensional geometry.

The photogrammetry process involves capturing a series of images from different perspectives around the heliostat. These images are processed using specialized software to detect and match points across the images, reconstructing the three-dimensional coordinates of the heliostat surface. The result is a detailed point cloud representing the surface geometry, which can be further processed to generate a high-resolution surface model [54].

One of photogrammetry's key advantages is its ability to provide highly detailed and accurate surface measurements. Studies have shown that photogrammetry can achieve a standard deviation of less than 0.2 mm when measuring heliostat surfaces, making it effective for detecting surface deformations and misalignments that impact optical performance. [54, 55]. Similar to deflectometry, these surface profiles are primarily utilized in ray tracing computations, while the mean slope error obtained from such measurements can also be applied in the convolutional approach.

However, field application of photogrammetry presents challenges. The process is time-consuming, requiring careful setup and multiple image captures from different angles. Traditional methods often use retroreflective targets to improve accuracy and resolution, but these must be manually placed on the heliostat surface, adding several hours to the preparation and measurement process, which limits scalability [55]. Additionally, environmental conditions such as lighting and weather can affect image quality and measurement accuracy [55].

Recent advancements have sought to streamline photogrammetry and reduce its complexity. For example, techniques have been developed to automatically detect heliostat facet corners in digital images, eliminating the need for retroreflective targets. This automation significantly reduces measurement time while maintaining acceptable accuracy levels for practical applications. Edge detection methods, for instance, allow rapid and automated photogrammetric analysis by identifying facet vertices [55].

Despite these advancements, photogrammetry still faces limitations in robustness and efficiency, especially under harsh field conditions. The method's dependence on high-quality images makes it vulnerable to disturbances such as shadows, reflections, or adverse weather, which can compromise measurement accuracy. While automated methods improve speed, they sacrifice accuracy compared to traditional approaches using retroreflective targets, as they only determine facet edge positions. Consequently, photogrammetry remains impractical for commercial applications, whether using retroreflective targets or not.

Beam Characterization Methods

Beam Characterization encompasses a wide range of techniques that utilize the reflected focal spot from heliostats to gain insights into their properties. These methods are instrumental in assessing the optical performance of heliostats and improving flux prediction

methods, especially in large solar thermal power plants. The focal spot, the concentrated radiation of the sun formed on a calibration target or other surface, contains valuable information about the heliostat's mirror shape, alignment, and surface quality. Beam Characterization techniques range from measuring basic focal spot shapes to optimizing heliostat facet canting and reconstructing detailed surface profiles.

At the simplest level, Beam Characterization can determine the basic shape of the focal spot, providing a quick assessment of the heliostat's general optical performance. Methods like those introduced by Kiera [56] use convolution models to approximate the shape of the focal spot by superimposing error distributions from an individual heliostats and adjusting the parameters accordingly to the measured focal spot. This approach, while straightforward and computationally efficient, often suffers from oversimplifications that limit its precision. It can predict the general shape of the focal spot but does not account for detailed surface deformations or specific heliostat errors.

Moving to a more detailed level, Beam Characterization can be used to optimize the canting parameters of heliostat facets. Techniques such as those developed by Sánchez-González et al. [57] use optimization algorithms to fit simulated flux distributions to those captured on a target, accounting for different facets of the heliostat. This method provides a higher level of detail and accuracy compared to simple shape measurements, as it directly determines correct alignment errors of facets.

The most advanced Beam Characterization techniques involve reconstructing the detailed mirror surface profile of heliostats, like methods by Pargmann et al. [22] and Martinez et al. [58].

For instance, the work of Pargmann et al. demonstrated a method where a limited number of focal spot measurements from the standard calibration target used for heliostat calibration are used to reconstruct the surface deformations of heliostats. This approach involves optimizing the heliostat surface through differentiable ray tracing, which allows for precise simulations of heliostat behavior under various conditions. Although this method provides highly accurate surface profiles and enhances flux predictions significantly, it is computationally intensive, making it less practical for extensive heliostat fields. Moreover, the methodology faces limitations for heliostats located farther from the calibration target due to information loss occurring due to concentration between the mirror surface reflection and the focal spot's impact on the target. The investigation was conducted solely on simulated data; hence, the robustness of the approach may decrease when applied to real calibration data. [22]

Martinez et al. employed a similar technique optimizing small surface element orientations to a measured focal spot, but in comparison to Pargmann et al. use a movable target which can be placed directly in front of the heliostat. The shorter distance between heliostat and target reduces the concentration factor and therefore the information loss, enabling surface reconstruction for all heliostats independent of their distance to the receiver. However, this method does not work with operational data obtained from heliostat calibration, necessitating an additional movable target, which increases the complexity and cost of this method significantly [58].

In conclusion, Beam Characterization offers a comprehensive approach to heliostat characterization, from basic shape assessments to detailed surface reconstructions at low costs. The primary benefit of Beam Characterization is its potential to provide detailed and accurate assessments of heliostat performance based on operational data. By lever-

aging focal spot images, these methods can enhance the precision of flux predictions at a low cost. However, the complexity and computational demands of advanced methods remain significant challenges, particularly for large-scale implementations.

3.4.3. Heliostat Efficiency Modeling

The methods discussed in previous sections primarily address the geometric simulation of focal spot shapes, focusing on the spatial distribution of flux. However, accurate flux prediction also requires modeling the various optical losses along the path of sunlight to calculate flux values in Watts per square meter. This section delves into the Efficiency Models, which encompass these calculations and loss mechanisms.

Sunlight undergoes several losses from its source to the receiver. Atmospheric extinction and cloud shadowing constitute the initial losses occurring within the atmosphere and up to the heliostat. Since the Direct Normal Irradiance (DNI) is usually measured on the ground near the heliostat field or determined through nowcasting methods [40, 41], these atmospheric losses are generally excluded from heliostat modeling.

Starting from the DNI, the maximum possible incident energy flux on the receiver from one heliostat is given by the mirror surface area A_{hel}:

$$\dot{Q}_{max} = DNI \cdot A_{hel}$$

For convolutional methods, the mirror is typically approximated as a continuous plane. In some instances, the mirror is divided based on its facet geometry, and the energy flux is calculated for each segment accordingly. In contrast, Monte Carlo ray tracers divide the mirror surface into many theoretical surface elements. The energy is then calculated for each element, allowing for the scaling of rays sampled from each element to accurately simulate the flux distribution.

In real-world scenarios, this maximum energy flow is reduced by several losses, represented by the efficiencies η_i, as shown in the following equation:

$$\dot{Q}_{real} = DNI \cdot A_{hel} \cdot \eta_{cos} \cdot \eta_{\rho} \cdot \eta_{SB} \cdot \eta_{\epsilon} \cdot \eta_{S}$$

where:

- η_{cos} is the cosine efficiency,
- η_{ρ} is the reflection efficiency,
- η_{SB} is the shading and blocking efficiency,
- η_{ϵ} is the atmospheric extinction efficiency,
- η_{S} is the spillage efficiency.

These efficiencies are described in detail in the following paragraphs.

Reflection and Cosine Efficiency

Closely related to the heliostat itself are the reflection efficiency η_ρ and the cosine efficiency η_{cos}. The reflection efficiency describes losses caused by the mirror reflectivity $\rho < 1$, leading to less flux being reflected than incoming. The reflection efficiency can be directly described by the heliostat mirror surface reflectivity:

$$\eta_\rho = \rho$$

Cosine efficiency η_{cos}, on the other hand, describes losses related to a smaller effective mirror area due to the angled incident direction of the sunlight on the mirror. As the sun moves, the heliostats track its position to maintain focus on a fixed point, the receiver. This tracking causes varying incident angles of sunlight on the mirror, affecting the cosine efficiency, which is expressed as:

$$\eta_{cos} = \cos(\theta)$$

where θ is the angle of incidence of sunlight on the mirror surface. [59]

Shading and Blocking Losses

Shading and blocking are significant sources of efficiency losses in heliostat fields. Shading occurs when a heliostat is obstructed by another heliostat or the tower before the reflection, while blocking happens when a heliostat obstructs the reflected rays from a neighboring heliostat after reflection. Using Monte Carlo ray tracers, shading and blocking can be determined for each individual ray through collision tests with surrounding heliostats and objects. Although this method is highly accurate, it is also computationally expensive. In convolutional methods, shading and blocking losses are typically calculated using analytical approaches. First, potential blocking heliostats for the heliostat under analysis are identified. Then, the projection of the outer contour of the heliostat is used to determine the proportion of radiation that gets blocked based on the direction of the reflection onto the other heliostat. While convolutional methods can account for shading and blocking losses in the total transferred energy calculation, they cannot account for changes in focal spot distribution due to blocked radiation as accurately as Monte Carlo Raytracers can. [60–62]

Atmospheric Extinction

Atmospheric extinction efficiency η_ϵ accounts for the attenuation of solar radiation between the heliostat and the receiver, resulting from absorption and scattering by air molecules, aerosols, and other particles. This efficiency decreases as the distance between the heliostat and the receiver increases. It is generally expressed as:

$$\eta_\epsilon = e^{-k \cdot d}$$

where k is the atmospheric attenuation coefficient (dependent on atmospheric conditions such as humidity and aerosol concentration), and d is the distance between the heliostat and the receiver. The exponential nature of the relationship emphasizes the significant impact of longer distances on the transmitted energy, making η_ϵ a critical factor for heliostat fields with large tower-to-heliostat distances.

Spillage

Spillage losses occur when radiation reflected from the heliostat misses the receiver. For Monte Carlo ray tracers, this can be determined directly and with high accuracy by tracking individual rays, enabling precise identification of both the rays that reach the receiver and those that miss. In convolutional methods, the calculated focal spot is projected onto the aperture plane of the receiver to determine the proportion of radiation that hits or misses the receiver. However, this method is less precise than the ray tracing approach.

— In conclusion, while heliostat efficiency models are essential for accurately predicting the flux on the receiver by accounting for various optical losses—such as reflection, cosine losses, shading, blocking, and spillage—these aspects are not the focus of the current research, as they are already well-established and highly accurate, leaving limited scope for further improvement.

3.5. Conclusion and Research Gaps

The state of the art in solar tower technology demonstrates remarkable progress in leveraging advanced algorithms to improve the efficiency and cost-effectiveness of Concentrated Solar Power (CSP) plants. A key operational value driving this progress is flux density prediction, which plays a pivotal role in optimizing efficiency and ensuring the longevity of critical components such as the receiver.

In the domain of heliostat flux prediction, two critical trade-offs are evident:

- **Flux Prediction: Accuracy vs. Scalability:** Existing methods are either highly accurate but computationally expensive and slow, or they are computationally efficient but lack precision. This limitation makes it difficult to balance the need for detailed flux predictions with the operational demands of large-scale solar tower plants.
- **Heliostat Characterization: Cost vs. Applicability** Current characterization techniques are either accurate but time-consuming and expensive, or they are less costly but insufficiently precise. This trade-off restricts the feasibility of achieving high accuracies when applying these methods to commercial-scale power plants.

Addressing these trade-offs is crucial to fully realize the potential of solar tower technology. Future research should prioritize heliostat characterization and flux prediction methods that balance high accuracy, scalability, and cost-efficiency. Beam characterization, leveraging operational data, stands out as a promising approach, particularly when combined with advanced computational techniques.

In conclusion, while the field has made significant strides, integrative solutions are essential to overcome the limitations identified in flux prediction and heliostat characterization. The next chapter presents a novel methodology that addresses these challenges by utilizing beam characterization and neural network-based models, offering a pathway toward scalable, accurate, and cost-efficient solutions for Concentrated Solar Tower systems.

Chapter 4.

Methodology

In the State of the Art Chapter, critical gaps in existing flux prediction and heliostat characterization methods were identified. Current flux prediction techniques face a trade-off between accuracy and scalability, while heliostat characterization methods often rely on costly and time-intensive measurements that are impractical for large heliostat fields. Addressing these limitations is crucial for developing efficient, scalable, and accurate heliostat flux prediction methods, which are fundamental to the optimization of Concentrated Solar Power (CSP) plants.

Among heliostat characterization techniques, Beam Characterization was identified as particularly promising due to its operational efficiency and cost-effectiveness. Using calibration images, this method eliminates the need for additional measurements, reducing costs and making it scalable to larger heliostat fields. However, existing methods based on focal spot measurements are limited by their reliance on oversimplified models or complex surface reconstructions, which often suffer from robustness and stability issues due to information loss occurring during the concentration between the heliostat and the calibration target. These challenges become more pronounced for heliostats positioned farther from the calibration target.

To address these issues, we propose a data-driven approach that replaces traditional heliostat characterization and flux prediction methods with a unified methodology. This novel approach directly utilizes focal spot data to predict flux distributions without requiring detailed physical models of heliostats. The method introduces abstract heliostat representations learned through neural networks, enabling a more flexible and efficient characterization of heliostat properties. By focusing on the variables of interest—focal spots—this approach ensures accurate flux predictions while maintaining scalability and robustness.

According to the classification in Section 3.4, this approach combines Beam Characterization for heliostat characterization with Neural Networks for focal spot computation. Rather than treating these two processes separately, a single neural network is utilized to seamlessly integrate both tasks, providing a streamlined and efficient solution for predicting heliostat flux density based on straightforward focal spot measurements.

The proposed methodology is illustrated in Figure 4.1, which outlines the key steps from raw image acquisition to the final prediction of flux distribution on the receiver. The chapter is organized as follows:

- **U - Translation of Raw Images to Flux Density Distributions (Section 4.1)**: This section introduces an Image-to-Image neural network that preprocesses raw calibration target images and translates camera pixel intensities into flux density distributions, enabling accurate and reliable data representation.

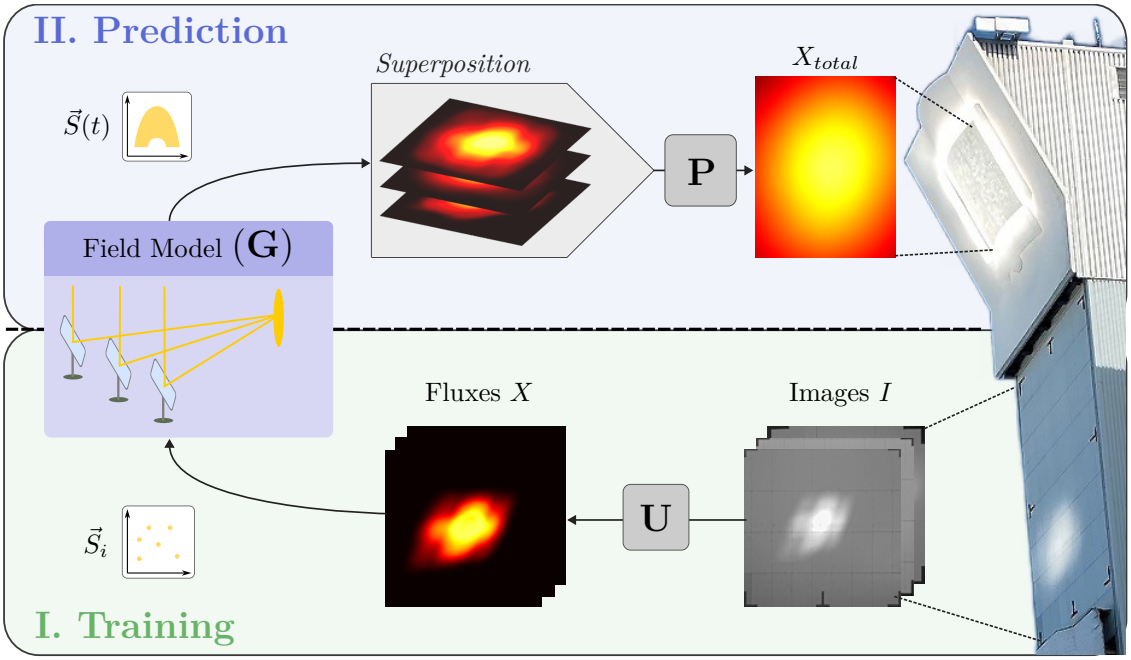

Figure 4.1.: Illustration of the overall methodology for data-driven flux prediction using target images, highlighting each individual step in the process, from raw image acquisition to flux prediction on the receiver. (Adapted from [63])

- **G - Generalizing Flux Prediction Models (Section 4.2)**: This part discusses the encoding of latent heliostat representations and their use in a generative neural network to predict focal spot variations under diverse conditions.
- **P - Projection, Intensity Scaling and Superposition on the Receiver (Section 4.3)**: This section explains the process of scaling predicted focal spots to absolute flux values and projecting them onto the receiver surface to compute the total flux distribution.

By addressing the challenges of scalability, accuracy, and cost-effectiveness, this methodology provides a comprehensive framework for advancing heliostat characterization and flux prediction in CST plants. The structured presentation ensures that each component is transparent, reproducible, and rigorously evaluated, laying the groundwork for practical implementation.

4.1. Translation of Raw Images to Flux Density Distributions

Starting from raw calibration images, the initial steps involve cropping and rectifying these images using marker detection, resulting in the cropped images shown in Figure 4.2. To date, these raw target images have been primarily utilized to determine the focal spot center, which suffices for optimizing the alignment model. These calculations are typically performed using simple algorithms, such as thresholding and contour detection. However, these methods are incapable of deriving accurate flux density distributions, which are essential for a data-driven flux prediction methodology. Consequently, there

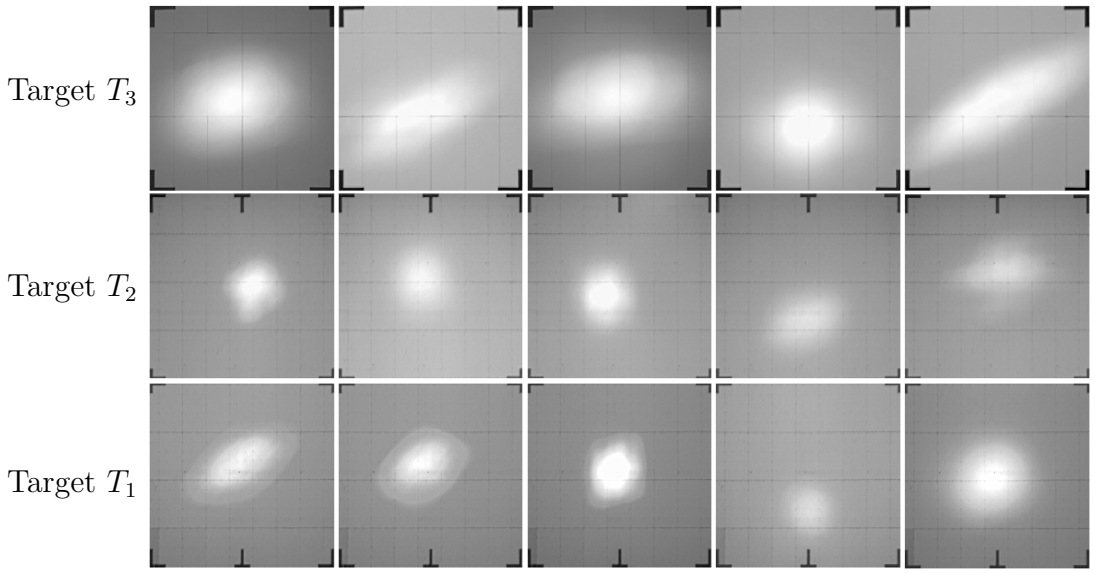

Figure 4.2.: Sample target images from the Solar Tower Jülich (STJ) obtained through the standard calibration process (described in Section 3.3). Random samples from all three calibration targets are displayed.

is no existing method for extracting high-resolution flux distributions from calibration images, creating a critical gap in the current workflow.

This section addresses this gap by introducing a novel methodology to translate raw target images into flux density distributions. The proposed approach leverages an image-to-image neural network (NN), specifically a UNet3+ architecture [64], to bypass the limitations of conventional image processing algorithms. These traditional methods struggle to account for the complex characteristics of calibration images, such as variability in background illumination, non-uniform reflectivity, and diverse focal spot intensities. In contrast, the NN approach directly extracts the desired flux density distributions by learning from realistic artificial training data.

Part of this section has already been published in Kuhl et al. [65], where this method was first introduced in combination with an efficiency model to derive accurate flux distributions. This work represents the first application of neural networks to calibration images, addressing limitations in existing image processing algorithms.

Challenges in Flux Density Derivation

Extracting flux density distributions from the raw images presents significant challenges. The camera intensity in these images is influenced by both the focal spot and background irradiation reflected off the calibration target. Variations in the intensity and position of the focal spots, combined with changing background irradiation, complicate the use of traditional image processing algorithms. These algorithms often fail to reliably differentiate between background illumination and heliostat-reflected radiation, hindering automated flux density extraction and seamless integration into the calibration workflow.

Common approaches, such as subtracting background irradiation using reference images of the empty target, provide only partial solutions. Temporal fluctuations in irradiation and the extended calibration time required for capturing multiple images further

limit their effectiveness.

Problem Definition of Target Image Segmentation

The key challenge addressed by the UNet is the separation of observed target images, I, into two distinct components: the flux distribution, X, and the background illumination, B. This separation is critical for accurately determining the flux density distribution reflected by a heliostat.

An observed image I is composed of a superposition of two principal components: the ambient background illumination I_A and the incident heliostat flux X. Both components are modulated by the spatial reflectivity R of the target, leading to the relationship:

$$I = R \cdot (X + I_A) = R \cdot X + B.$$

Here, $B = R \cdot I_A$ denotes the reflected background illumination, while X corresponds to the relative flux distribution contributed by the heliostat. The relative flux intensities are normalized and maintain proportionality to their absolute counterparts. Consequently, the absolute flux distribution X_a can be expressed as:

$$X = c \cdot X_a,$$

where c is a scaling factor that maps the relative flux distribution X to absolute flux values.

The translation of pixel intensities into absolute flux values in $\mathrm{W\,m^{-2}}$ is hindered by the absence of detailed camera settings, such as shutter speed and ISO. However, for the purposes of this study, the relative flux distributions X are sufficient. The generative model employed later in the methodology (see Section 4.2) operates exclusively on relative flux distributions. Absolute scaling is subsequently applied in a later step by the efficiency model described in Section 4.3.

Image-to-Image Neural Network for Flux Derivation

A UNet is employed to map images I to pixel-wise outputs for X (relative flux distribution) and B (background), as illustrated in Figure 4.3(right). The UNet3+ architecture [64], recognized for its remarkable capability in segmentation tasks, is used for this purpose:

$$X, B = \mathbf{U}(I).$$

While the primary objective is to extract X, including B as an auxiliary output enhances the network's performance. Predicting B supports robust background subtraction, a critical step in flux calculation, and improves the network's ability to differentiate between background illumination and heliostat flux. As supported by Liebel et al. [66], carefully chosen auxiliary tasks can significantly benefit the main task, leading to more accurate predictions of X.

Generating Artificial Data for Training

A key challenge in training the UNet is the lack of labeled ground truth data for the observed images I. To overcome this, artificial target images are generated by simulating each component of the target images, X, B, and R, as follows:

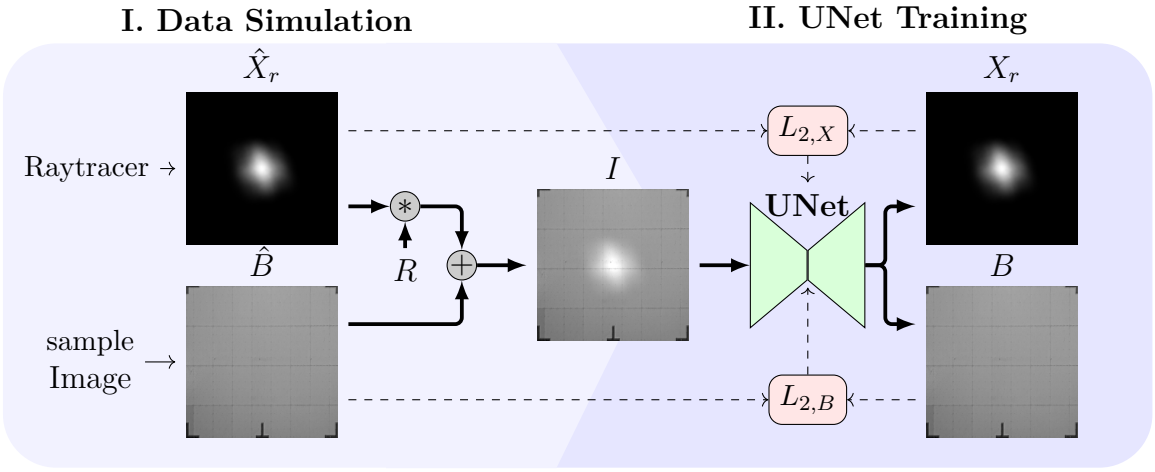

Figure 4.3.: UNet Method. A ray tracer generates focal spots $\hat{X}$, which are multiplied by the target's spatial reflectivity R and combined with a randomly sampled empty target image $\hat{B}$. The UNet is trained to decompose the simulated image I into X and B, with pixel-wise L2 losses $L_{2,i}$ calculated against the ground truth for training. [65]

- **Background Illumination (B):** Background images B are obtained from empty target captures, representing reflected ambient light. Over 1000 such images were collected throughout the year, capturing a wide range of irradiation conditions. These images are randomly sampled during training to simulate diverse background scenarios.
- **Spatial Reflectivity (R):** The spatial reflectivity R is calculated from averaged background images under the assumption of Lambertian reflectivity. It is computed as $R = \overline{B}/\max(\overline{B})$, where $\overline{B}$ is the pixel-wise mean of all background images, normalized by its maximum value.
- **Simulated Flux Distributions (X):** Relative flux distributions X are generated using a ray tracer, incorporating surface measurements from deflectometry. These distributions are scaled to align with the intensity ranges observed in real calibration images, ensuring consistency between simulated and real-world data.

The artificial data generation process is illustrated in Figure 4.3(left), with the resulting simulated images shown in Figure 4.4, providing a comprehensive dataset for training the UNet.

Training and Loss Function

The UNet is trained on the artificial dataset using a pixel-wise L2 loss function:

$$\mathcal{L} = L_{2,X} + L_{2,B},$$

$$L_{2,X} = \frac{1}{N}\sum_{i=1}^{N}(X_i - \hat{X}_i)^2, \quad L_{2,B} = \frac{1}{N}\sum_{i=1}^{N}(B_i - \hat{B}_i)^2,$$

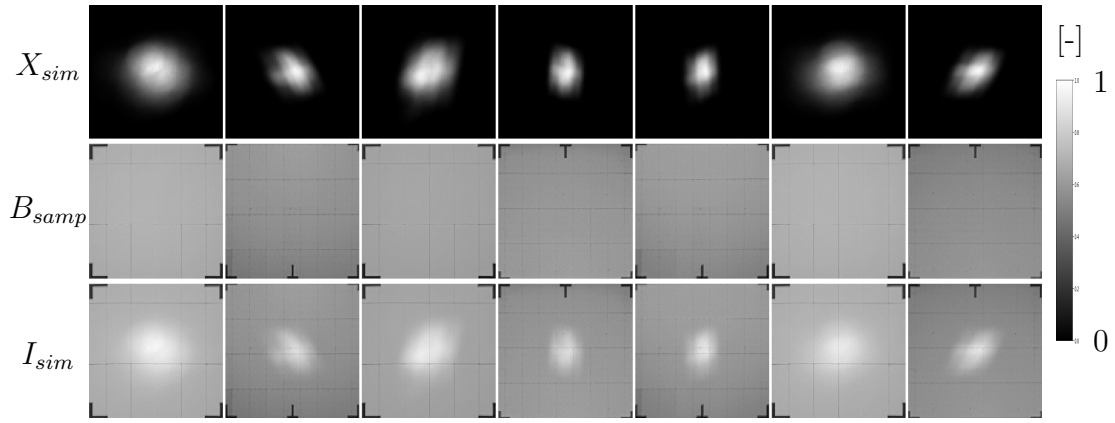

Figure 4.4.: Samples of artificially generated target images I_{sim}, created from simulated focal spots X_{sim} and sampled background images B_{samp}.

where $L_{2,X}$ and $L_{2,B}$ compute the mean squared error between the predicted and artificial ground truth values for X and B, respectively. The whole process is illustrated in Figure 4.3.

— By leveraging a UNet trained on simulated data, this methodology bypasses the limitations of conventional image processing techniques. The innovation lies in the creative generation of realistic artificial data with known properties, using available data sources. This enables robust and automated extraction of flux density distributions, which forms the basis for the data-driven heliostat characterization and flux prediction detailed in subsequent sections.

4.2. Generalizing Heliostat Flux Prediction Approach

The next step in the proposed methodology (Figure 4.1) builds on the flux measurements derived in Section 4.1, shifting the focus from single heliostat characterization to generalizing flux predictions across the heliostat field. This section introduces a data-driven approach that combines heliostat characterization and flux prediction into a unified framework, leveraging neural networks to directly model focal spot changes under varying input conditions, such as sun position and aim point

Following the definitions provided in Section 3.4, the proposed approach is separated into two key components: the Heliostat Characterization Technique and the Flux Prediction Method. The Heliostat Characterization Technique derives abstract heliostat properties, denoted as $\tilde{H}$, from sample flux measurements. In machine learning terms, this step corresponds to encoding the image data into a latent heliostat representation. Subsequently, the Flux Prediction Method uses these properties, along with input conditions, to predict focal spot changes through a generative neural network. In machine learning terminology, this corresponds to generating images based on the learned or encoded properties and specified conditions. Mathematically, this is expressed as:

$$\tilde{H}_i = \mathbf{E}(X_j \in \mathcal{X}_i) \tag{4.1}$$

$$X = \mathbf{G}(\tilde{H}_i, \vec{S}, \vec{Ap}) \tag{4.2}$$

Here, the encoder E transforms discrete sample focal spots X_j into latent heliostat representations $\tilde{H}$, while the generator G uses these representations, along with sun position $\vec{S}$ and aim point $\vec{Ap}$, to produce continuous focal spot predictions X for any given input condition.

The key innovations of this approach are as follows:

- **Combined Framework for Heliostat Characterization and Flux Prediction:** This approach integrates heliostat characterization and flux prediction into a single, unified framework, eliminating the need for intermediate steps and reliance on detailed physical parameters, as required by other methods. The framework leverages latent heliostat representations learned directly from focal spot measurements, enabling end-to-end optimization from measurement to flux prediction. This streamlined process ensures high prediction accuracy while maintaining a direct focus on the target value—flux distributions.
- **Generalization Across the Entire Heliostat Field:** The framework introduces a unique ability to generalize across the entire heliostat field by training the model on data from all heliostats—something no prior heliostat characterization technique could account for. By leveraging data from the full field, the network encodes and shares information about common defects, errors, and features across heliostats. This shared knowledge fosters a unified representation of the heliostat field, enhancing scalability and directly improving the accuracy of flux predictions for individual heliostats.

In the following sections, the encoding process of heliostat properties, generative model architectures, and their adaptations for this task are described in detail, highlighting the unique approach of this methodology.

4.2.1. Identifying Encoding Techniques for Heliostat Properties

The central challenge in this methodology lies in defining an abstract yet generalizable representation of heliostat properties that can be effectively integrated into a machine learning framework. The success of the data-driven framework hinges on the network's ability to distill key heliostat characteristics from focal spot measurements. These encoded representations must be both compact and expressive, enabling accurate flux prediction across the entire heliostat field. Two encoding approaches have been developed to address these requirements, as detailed in the following paragraphs.

Encoding through Backpropagation through the Generator

In this approach, depicted in Figure 4.5a, the heliostat properties $\tilde{H}$ are represented as learnable latent vectors, one for each heliostat. These latent vectors are initialized randomly and iteratively refined during training via backpropagation. The generator G predicts the focal spot X based on the current values of $\tilde{H}$, sun position $\vec{S}$, and aim point $\vec{Ap}$. The loss between the predicted focal spot and its ground truth is then backpropagated through the generator to update both the generator weights and the specific latent vector $\tilde{H}_i$ corresponding to the heliostat under consideration.

This iterative learning process allows the network to infer heliostat-specific properties directly from focal spot samples without requiring predefined physical parameters. By

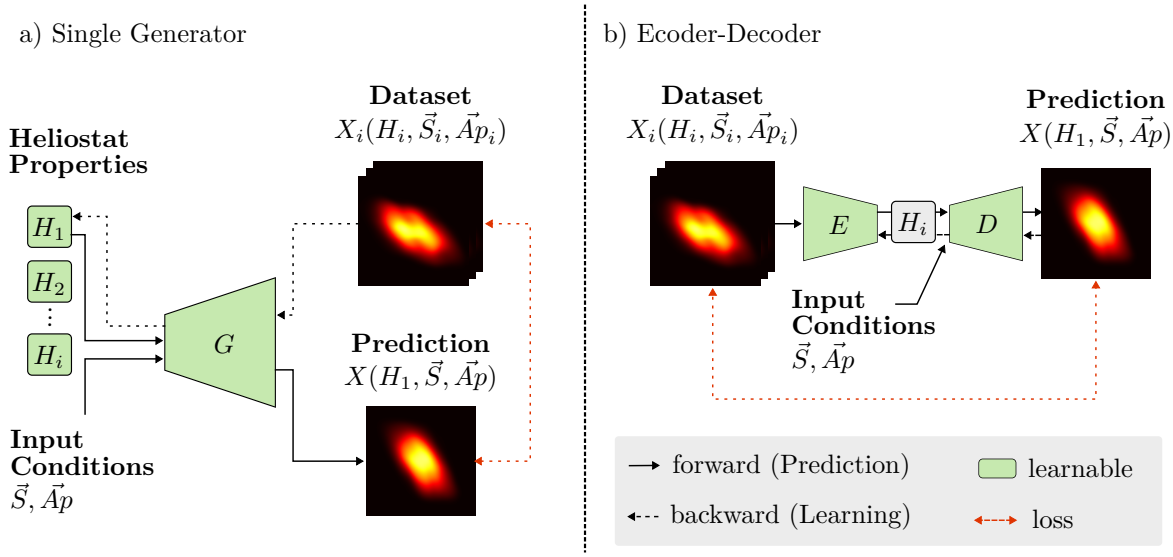

Figure 4.5.: Illustration of the two encoding techniques for heliostat properties $\tilde{H}$. (a) Single generator approach, where heliostat parameters are learned iteratively via back-propagation during training. (b) Encoder-decoder structure, where heliostat properties are explicitly output by an encoder.

training the generator on data from many heliostats with diverse conditions, the network is forced to generalize the flux generation logic, while the latent vectors capture heliostat-specific details.

This approach offers a straightforward implementation and is particularly effective when the relationship between focal spot measurements and heliostat properties is complex. However, adding new calibration points or heliostats requires fine-tuning or re-training the associated latent vectors, which can limit scalability during field operation. The StyleGAN architecture chosen for this approach will be discussed in detail in Section 4.2.2.

Encoding with an Explicit Encoder/Decoder Structure

This method, illustrated in Figure 4.5b, adopts a more direct approach by incorporating a distinct encoder network to derive the heliostat properties $\tilde{H}$ from focal spot samples. The encoder processes input calibration points associated with a heliostat, transforming them into abstract representations of its properties. These encoded properties are then provided to the generator, along with the input conditions, to predict the corresponding focal spot. The parameters of the encoder and decoder are updated after a single forward pass - the prediction of a focal spot based on other focal spots and the subsequent loss calculation on the predicted X. Consequently, the heliostat properties are not explicitly learnable but are directly produced as the output of the encoder.

The key advantage of this explicit encoder-decoder structure is its flexibility during training. For instance, it enables the network to explicitly enforce specific behaviors, such as learning from a small set of calibration points or improving spatial extrapolation, as demonstrated in later sections. This direct encoding also makes it easier to update heliostat properties when new calibration points are available, as only a single forward pass through the encoder is needed.

However, the separation of encoding and generation increases the complexity of the

model, and the network must learn the mapping from focal spots to heliostat properties within a single forward pass, which can be challenging for highly complex relationships. This approach uses a modified transformer architecture, which is explicitly explained in Section 4.2.3.

Differences Between Encoding Methodologies

While both approaches rely on abstract, latent representations of heliostat properties, their mechanisms and implications differ significantly:

- **Learnability and Flexibility:** The backpropagation-based approach allows for iterative learning of heliostat properties, making it well-suited for scenarios where the encoding process is complex. In contrast, the encoder-decoder approach requires the encoding to be learned in a single pass but provides greater flexibility to enforce specific training behaviors and adapt to new data.

- **Scalability and Operational Efficiency:** The encoder-decoder method is more efficient for field operations, as new calibration points can be processed without retraining the model. In contrast, the backpropagation-based approach requires retraining or fine-tuning of the latent vectors for new calibration data, which may limit its scalability.

In summary, the backpropagation-based method (StyleGAN) demonstrates exceptional performance in tackling complex encoding tasks due to its iterative learning process of heliostat properties. Conversely, the encoder-decoder approach (Transformer) provides enhanced versatility and efficiency, particularly in scenarios demanding frequent updates or specialized training tasks. While both methods exhibit distinct advantages, the optimal approach for the STJ heliostat field and this specific application will be identified in the evaluation chapter. This ensures that the selected method aligns with the practical requirements of the heliostat flux prediction framework.

4.2.2. Single Generator Method: StyleGAN Architecture

The Single Generator Method is built upon a StyleGAN-inspired architecture to predict focal spots X based on input conditions $\tilde{H}_i$, $\vec{S}$, and $\vec{Ap}$, as formulated in Equation 4.1. A key aspect of this approach was identifying a model with a generator guided by vectorial inputs and tailoring it to meet the specific requirements of heliostat flux prediction.

The StyleGAN architecture, originally proposed by Karras et al. [28], was selected as a foundation for this task due to its ability to synthesize high-resolution images from multidimensional input vectors. Its hierarchical structure effectively balances large-scale and fine-grained details, making it particularly suitable for capturing the intricate features of focal spots. However, significant modifications were necessary to adapt the architecture for this unique application. Part of this section has already been published in Kuhl et al. [63], where this method was first introduced and the concept of purely data-driven flux prediction was applied to real measured data for the first time.

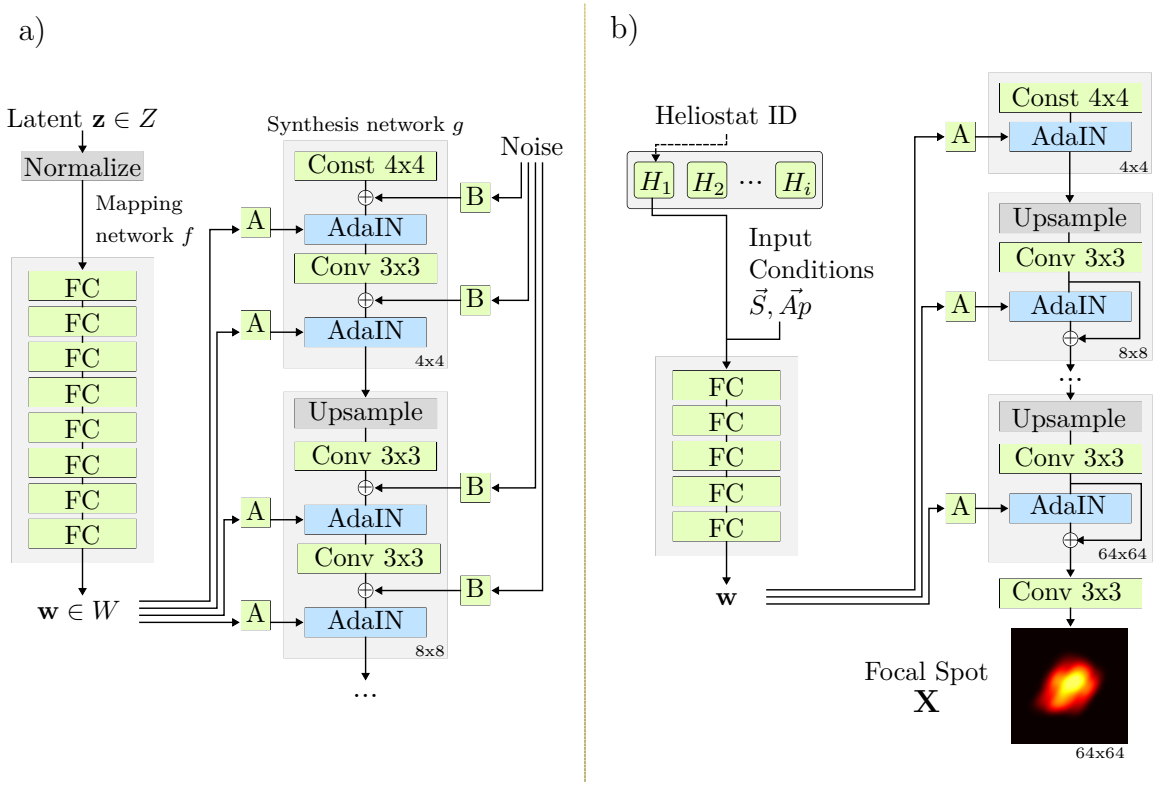

Figure 4.6.: Single Generator Method. Left: Original StyleGAN architecture [28]. Right: Modified architecture adapted for focal spot prediction.

Adapting StyleGAN for Focal Spot Prediction

The original StyleGAN architecture (Figure 4.6, left) was designed for generating diverse image datasets with inherent randomness, incorporating noise layers and a random input vector. For focal spot prediction, where a heliostat's properties deterministically map to its focal spot under specific conditions, these components were replaced or removed to ensure a one-to-one mapping (Figure 4.6, right). Key adaptations included:

- **Targeted Adaptation for Heliostat Parameters:** Unique learnable latent vectors $\tilde{H}_i$ (as described before) were introduced for each heliostat. These vectors abstractly encode heliostat-specific properties while leveraging shared knowledge through the generator. This innovative use of abstract parameters ensures efficient representation and enhances a generalizing generator.

- **Replacement of Random Input Vectors:** Random input vectors were substituted with deterministic inputs: heliostat properties ($\tilde{H}_i$), solar position ($\vec{S}$), and aim point ($\vec{Ap}$). This preserves the deterministic relationship between inputs and outputs, which is critical for accurate focal spot prediction.

- **Elimination of Noise Layers:** Noise layers, which introduce stochasticity for diversity in typical StyleGAN tasks, were removed to ensure consistency and deterministic output for each unique input configuration.

- **Elimination of the Discriminator:** The discriminator, typically used in adversarial training, was omitted as perceptual quality was not the focus. This simplifi-

cation allowed training to prioritize pixel-wise accuracy, directly aligning with the flux prediction application's objectives.

Adapted Training Procedure

To ensure robust predictions across a diverse heliostat field, a generalizing training approach was adopted. Each focal spot sample in the dataset is tagged with a heliostat ID (H_i), solar position ($\vec{S}$), and aim point ($\vec{Ap}$), along with its corresponding focal spot (X). The training process involves:

1. Selecting the learnable vector $\tilde{H}_i$ corresponding to the heliostat ID (H_i).

2. Combining $\tilde{H}_i$ with the solar position ($\vec{S}$) and aim point ($\vec{Ap}$) to form the generator input.

3. Using the generator G to predict the focal spot (X) for the given input conditions.

The predicted focal spot (X) is evaluated against the measured focal spot ($\hat{X}$) using a pixel-wise loss function. This loss drives both the training of the generator G and the updates to the learnable heliostat vectors $\tilde{H}_i$. While heliostat vectors vectors are updated solely by the sampels of the specific heliostat, the generator is updated by all measurements, which leads to the desired generalizing generator.

— The adaptations to the StyleGAN architecture demonstrate how a general-purpose generative model can be tailored to meet the specific challenges of heliostat flux prediction. By introducing learnable heliostat-specific parameters and replacing stochastic elements with deterministic inputs, the architecture achieves robust, scalable, and accurate predictions, contributing to the broader goal of a unified data-driven framework for heliostat characterization and flux prediction.

4.2.3. Encoder-Decoder Method: Transformer Architecture

In addition to the Single Generator Approach, a second strategy for generalizing heliostat flux predictions was explored: the Encoder-Decoder Method, which leverages a modified Transformer architecture. This approach tackles the challenge of encoding heliostat properties while adapting a versatile, general-purpose architecture to the specific requirements of image-based flux prediction.

The Transformer architecture, originally proposed by Vaswani et al. [26], was selected due to its strengths in handling variable-length input sequences and its attention mechanisms, which effectively weight and contextualize input data. These capabilities are particularly suitable for this task, as heliostats can have varying amounts of calibration points, and the weighting of individual calibration points can significantly enhance prediction accuracy. However, applying the Transformer to heliostat flux prediction required substantial adaptations, as the nature of the task—mapping between focal spots for varying conditions—differs significantly from its original use in language processing.

Part of this section has already been published in Kuhl et al. [67], where the modified transformer architecture for flux prediction was introduced.

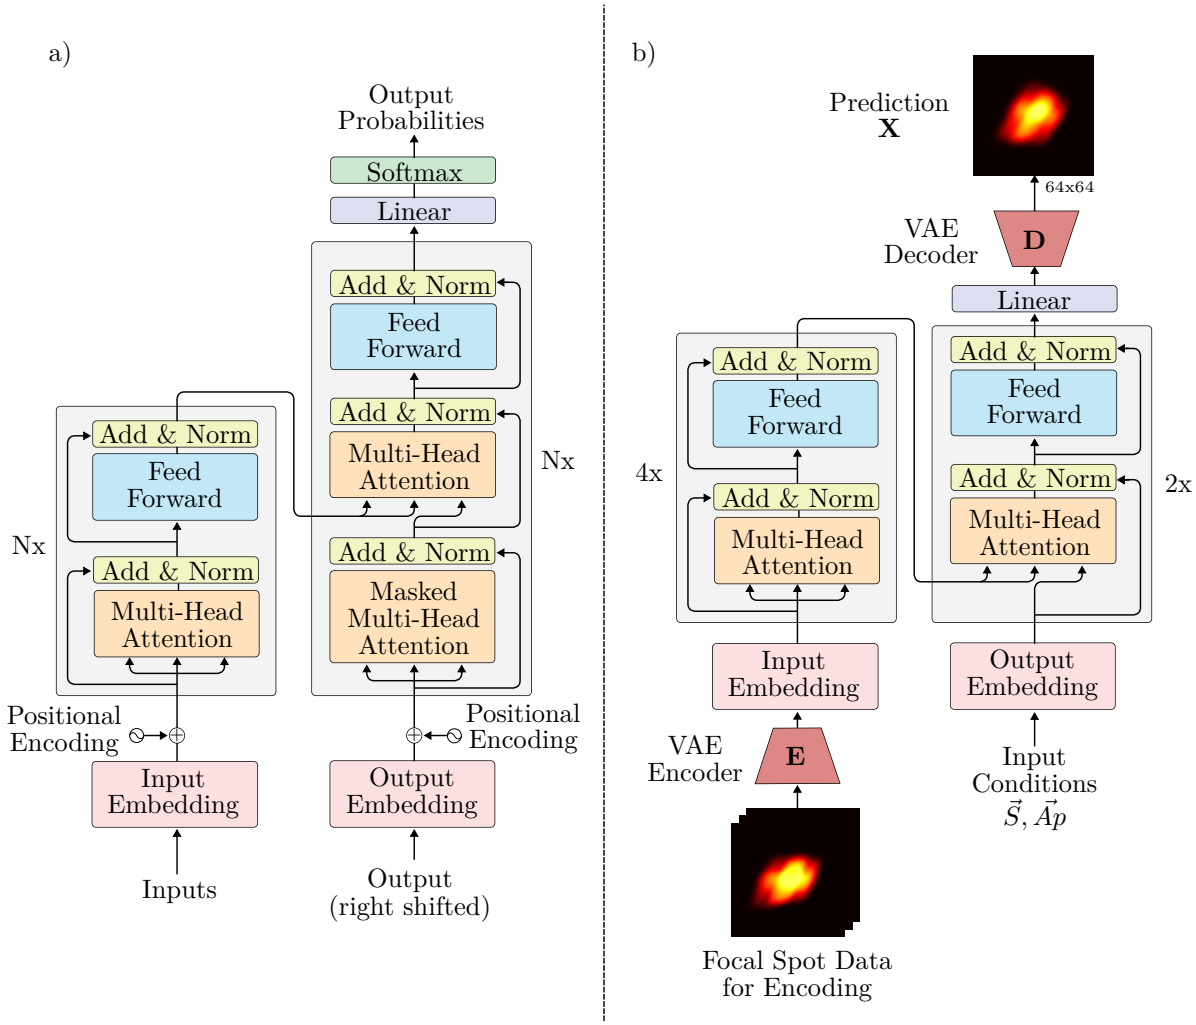

Figure 4.7.: Encoder-Decoder Method. Left: Original Transformer architecture [26]. Right: Modified Transformer architecture combined with a VAE for this approach.

Modifications to the Transformer Architecture

To adapt the Transformer for heliostat flux prediction, several key modifications were made, illustrated in Figure 4.7:

- **Integration with a Variational Autoencoder (VAE):** A VAE was incorporated to preprocess focal spot images into vector representations compatible with the Transformer. The VAE encoder compresses input images into a compact latent space, which is fed into the Transformer encoder. The VAE decoder transfers Transformer's predictions back into image space. While the VAE is pretrained on the compression task, it remains trainable during Transformer training to ensure optimal convergence and task-specific refinement.
- **Elimination of Positional Encoding:** Unlike language tasks, where positional encoding is essential to sequence understanding, the focal spot data is uniquely defined by input conditions such as solar position and aim point. Positional encoding was replaced by directly concatenating these input conditions with the latent vectors, allowing the network to guide the weighting and processing of calibration points effectively.

- **Refinement of Attention Mechanisms:** The Transformer's self-attention mechanism was retained in the encoder to enable contextual weighting of calibration points based on their relevance to the prediction conditions. In contrast, self-attention was removed from the decoder since the prediction points (e.g., focal spots under varying conditions) are independent of one another. Instead, the cross-attention mechanism in the decoder directly processes the input conditions, guiding how the calibration points are weighted and used for predictions.

Training Procedure and Tailored Strategies

The training procedure for the Transformer-based architecture differs significantly from the Single Generator Method. Instead of processing randomly selected samples across all heliostats, the encoder processes all calibration focal spots from a single heliostat in a single forward pass, allowing it to learn an effective weighting between calibration points. The training process involves:

1. Dividing all focal spot images of a heliostat into two groups: one for encoding and one for forming the prediction loss.
2. Using the encoder to generate a latent representation of the heliostat's properties from the encoding group.
3. Feeding this encoded representation and input conditions into the decoder to predict flux distributions for the prediction group.
4. Applying the loss to the predicted group, updating both the encoder-decoder and the VAE components.

This procedure enables tailored training strategies that address specific challenges and optimize the network's behavior. For example:

- **Learning from Limited Calibration Points:** To improve the model's ability to extract heliostat properties from sparse data, the encoder can be trained on a single focal spot and tasked with predicting all others for the same heliostat. This enforces robust property extraction from minimal calibration samples.
- **Extrapolation Across Input Conditions:** The network can be trained to extrapolate predictions across diverse input conditions. For instance, it can be tasked with predicting focal spots in the evening based on calibration data from the morning, enhancing its ability to generalize spatially and temporally.

The Encoder-Decoder Method highlights the potential of adapting general-purpose architectures, such as the Transformer, to address the specific challenges of heliostat flux prediction. By integrating a Variational Autoencoder (VAE) for image compression, enhancing attention mechanisms, and tailoring the training procedure, this method provides a robust alternative to the Single Generator Approach.

— Both approaches represent complementary strategies for modeling heliostat fields and predicting flux distributions, addressing different aspects of scalability, flexibility, and precision. Their comparative evaluation, presented in the Evaluation Chapter, will provide insights into the strengths and applicability of each approach, guiding the selection of the most effective method for this domain.

4.2.4. Generalizing Across Spatial Input Conditions: Aim Point Extrapolation

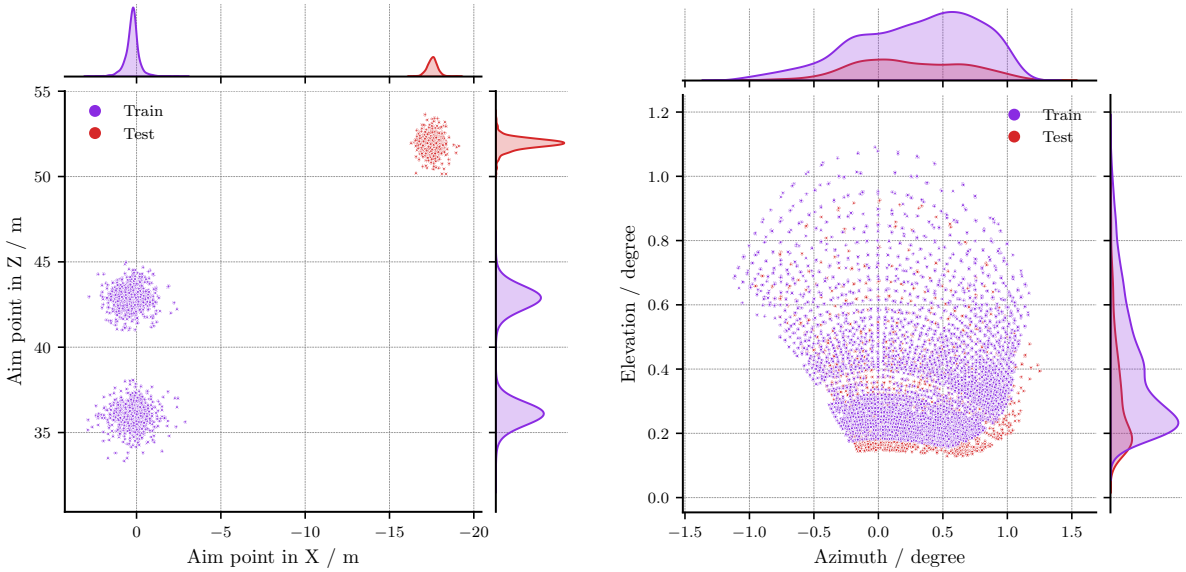

Figure 4.8.: Distribution of input conditions between training and prediction cases, represented as aim point $\vec{Ap}$ (left) and aim direction $\vec{Ad}$ (right). [67]

Building on the overall methodology outlined in Figure 4.1, we have explored the translation of raw calibration images into flux densities and developed generalizing models to predict flux distributions. A critical extension of this methodology involves enabling the network to spatially extrapolate aim points between training and prediction conditions. This capability is crucial because the calibration target data used for training differs significantly from the operational receiver configuration. Specifically, the aim points and corresponding incident angles vary between the calibration target and the receiver, necessitating a model that can generalize across these differing conditions.

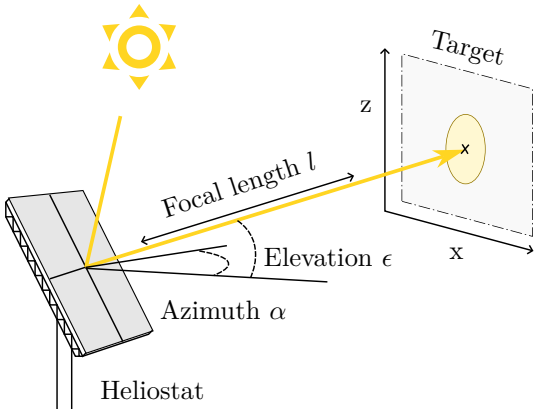

Figure 4.9.: Reparametrization of aim point coordinates ($\vec{Ap}$) into aim direction, defined by azimuth α and elevation ϵ angles.

The unique setup of the solar tower in Jülich provides a valuable opportunity to validate spatial aim point extrapolation. With two towers and multiple calibration targets, it is possible to train the network on data from one calibration target and test its extrapolation capabilities on another. This setup ensures precise validation while maintaining the high accuracy of flux measurements achievable with Lambertian calibration targets, compared to the receiver surface.

To address the challenge of spatial extrapolation, two complementary approaches have been developed and are discussed in the following sections. The first approach focuses on reparametrizing input conditions to better align training and prediction distributions, making it generally applicable to any dataset. The second approach, which is optional,

leverages artificial data to actively enhance extrapolation capabilities when such data is available. Parts of this section have already been published in Kuhl et al. [67], where these approaches were introduced to enable a data-driven framework for aimpoint generalization.

Reparametrization Approach

When describing the focal spot impact point as an aim point ($\vec{Ap}$) in three-dimensional coordinates, it becomes evident that the training calibration target (purple, Figure 4.9) does not overlap with the input condition range of the prediction calibration target (red) or the receiver. This disparity introduces unseen input conditions during prediction, increasing the risk of model inaccuracies.

To address this, we reformulate the aim point as an aim direction ($\vec{Ad}$), defined by the azimuth and elevation angles of the connecting vector between the heliostat and the focal spot impact point, as well as the vector's length. This reparametrization is illustrated in Figure 4.9. By expressing aim conditions in this general format, significant overlap is achieved between the training and prediction input condition ranges (Figure 4.8).

This reformulation provides several advantages:

- It reduces the likelihood of encountering entirely unseen input conditions during prediction, mitigating undesired behavior.
- The general formulation complements the generalizing capabilities of the prediction network by leveraging natural variation in aim directions present in training data.
- It enables the network to learn from variations within calibration targets, particularly in power plants where distinct areas on calibration targets are commonly used for parallel calibration of different heliostats.

Consequently, the network is equipped to learn how changes in aim direction influence focal spots, even for input conditions extending beyond the calibration target's surface.

Training with Artificial Data

An additional approach to enhancing spatial extrapolation involves incorporating artificial data into network training, provided such data is available. This artificial data, generated through ray tracing enhanced with deflectometric measurements, enables the creation of highly realistic focal spot samples. These simulated datasets can be tailored to include a wide range of input conditions, covering the entire aim point range across both calibration targets and the receiver.

Artificial datasets ensure comprehensive coverage of input conditions, particularly those that extend beyond the calibration target. This extended coverage enables the network to generalize effectively to unseen scenarios. Moreover, the network's ability to transfer knowledge learned from artificial data to real-world applications is particularly advantageous. For instance, deflectometric measurements from a small subset of heliostats can be used to train the network, which can then apply this learned behavior to predict flux distributions for the entire heliostat field.

In addition to enhancing generalization, artificial data introduces controlled variability into the training process, reducing the risk of overfitting to specific calibration targets. This results in a more robust network capable of handling diverse and challenging input

conditions. However, the feasibility of this approach relies on the availability of deflectometric measurements, which are critical for generating the realistic surface profiles required to simulate accurate focal spots.

— The Reparametrization Approach and the use of artificial data represent two complementary strategies for addressing the challenge of spatial extrapolation. Reparametrization, applied universally across all models, ensures a better representation of input conditions by aligning training and prediction ranges, thereby improving the consistency and reliability of network predictions. On the other hand, the incorporation of artificial data serves as an optional enhancement, providing an opportunity to actively improve the network's extrapolation capabilities when deflectometric measurements are available. Both approaches are evaluated to assess their individual and combined effectiveness in addressing the diverse requirements of heliostat flux prediction.

4.3. Intensity Scaling and Superposition on the Receiver

As illustrated in Figure 4.1, the preceding sections focused on predicting heliostat focal spot distributions, transitioning from raw calibration target images to flux distributions in the aperture plane of the receiver or reactor. However, predicting the total flux density distribution on the receiver surface requires additional steps: (1) projecting the flux from the aperture plane onto the receiver surface and (2) superimposing individual heliostat focal spots with intensity scaling. These steps are elaborated in the following subsections.

Parts of this section have already been published in previous works. The efficiency model was first introduced in Kuhl et al. [65], while the projection approach was originally presented in [67].

4.3.1. Efficiency Model

Up to this point, the predicted focal spot distributions have been relative flux distributions, denoted as X. This approach was chosen to enable the generative model to focus exclusively on predicting the shape of the flux distribution. For training, all focal spots were normalized to a mean pixel value of 0.1, ensuring uniformity and isolating the shape as the primary prediction target. To transition from these relative flux values X to absolute flux values X_a in $\mathrm{W\,m^{-2}}$, a scaling factor c is introduced, as defined by:

$$c \cdot X = X_a \tag{4.3}$$

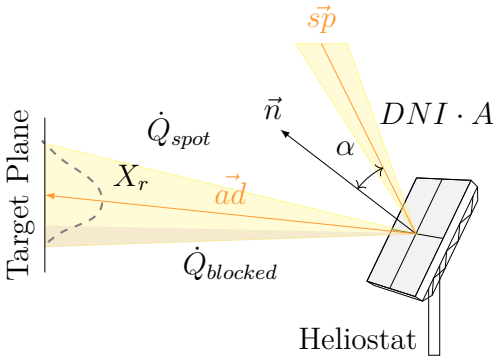

Figure 4.10.: Heliostat model used for intensity scaling. The relative flux distribution X is scaled by the total focal spot energy $\dot{Q}_{spot}$. [65]

The scaling factor c is determined based on the total energy flux of the focal spot, ensuring accurate translation to absolute values. For this purpose, a methodology common in state-of-the-art heliostat field efficiency models [59, 68] is adapted. This methodology scales the relative intensity distribution by calculating the total incident energy $\dot{Q}_{spot}$ reflected by a heliostat.

Figure 4.10 depicts the heliostat model used in this approach. The total incident energy flux $\dot{Q}_{spot}$ of the focal spot is calculated using the Direct Normal Irradiance (DNI), the effective mirror area A_{eff}, and the mirror reflectivity r:

$$\dot{Q}_{spot} = r \cdot DNI \cdot A_{eff} \tag{4.4}$$

The effective mirror area A_{eff} accounts for the angle of incidence α of sunlight on the heliostat surface, derived from the sun's position vector $\vec{sp}$ and the heliostat's aim direction vector $\vec{ad}$. Assuming the heliostat mirror normal $\vec{n}$ aligns with the reflection, α is calculated as half the angle between the incident and reflected directions:

$$\alpha = \frac{1}{2} \cdot \angle(\vec{sp}, \vec{ad}) \tag{4.5}$$

$$A_{eff} = \cos(\alpha) \cdot A_H \tag{4.6}$$

Here, A_H is the total mirror area projected in the direction of the incident light. Canting effects are negligible for the calculation of the effective mirror area at the Solar Tower Jülich (STJ), where heliostats are minimally canted. DNI values are continuously measured at concentrated solar power (CSP) plants, ensuring their availability for this methodology.

To compute the total flux $\dot{Q}_{image}$ represented by the normalized flux distribution X, pixel intensities are treated as flux per unit area $\dot{q}$ in $\mathrm{W\,m^{-2}}$. The relationship between $\dot{Q}_{image}$, $\dot{Q}_{spot}$, and the scaling factor c is expressed as:

$$\dot{Q}_{image} = A_F \cdot \dot{q}_{image,avg} \tag{4.7}$$

$$c = \frac{\dot{Q}_{spot}}{\dot{Q}_{image}} = \frac{r \cdot \cos(\alpha) \cdot DNI \cdot A_H}{A_T \cdot \dot{q}_{image,avg}} \tag{4.8}$$

Here, A_F represents the area defined by the focal spot image dimensions, and $\dot{q}_{image,avg}$ is the average pixel value of X.

4.3.2. Shading and Blocking

The contribution of $\dot{Q}_{\text{blocked}}$ to the efficiency model is neglected in this work. While most efficiency models include blocking analysis [60–62], it is not relevant for this study's focal spot analysis. This is because the analysis focuses on normalized focal spot distributions, avoiding external scaling factors. Typically, $\dot{Q}_{\text{blocked}}$ becomes significant in the context of field layout optimization, where the objective is to minimize blocked energy or maximize receiver power by optimizing aimpoint distributions. As these aspects are beyond the scope of this work, $\dot{Q}_{\text{blocked}}$ is not included in the efficiency model.

However, while shading and blocking are omitted from the efficiency model, they are implicitly considered in the generative model (G) that predicts the relative flux distributions. Since the generative model is trained on real measured data of focal spots, the resulting flux distributions inherently account for the effects of shading and blocking. The neural network learns these effects implicitly, as they are embedded in the focal spot shapes corresponding to different sun positions and aimpoints. This approach ensures that the impact of shading and blocking is effectively integrated into the predictions, even without explicit modeling in the efficiency framework.

4.3.3. Projection Technique and Superposition of Focal Spots

The generative model, trained exclusively on data from flat calibration targets, is inherently limited to predicting focal spot distributions on flat surfaces, such as the receiver's aperture plane. While these predictions are valuable for operational insights, they are insufficient for accurately predicting flux distributions on more complex receiver geometries. Consequently, an efficient method for projecting flux distributions from flat surfaces onto arbitrary complex surfaces was required.

To meet this requirement, a projection technique based on ray tracing is developed and implemented. Unlike traditional ray tracing, which samples rays from the heliostat surface, this approach samples rays from the predicted flux distribution within the aperture plane.

For generating rays suitable for ray tracing, three fundamental properties must be determined: the ray origin, the ray direction, and the energy transported by the ray. In this approach, one ray is sampled per pixel of the focal spot in the aperture plane, with the origin of each ray corresponding to the center of the respective pixel. The energy associated with each ray is calculated based on the flux value of the corresponding pixel and its surface area. The direction of the rays can be defined using one of three approaches:

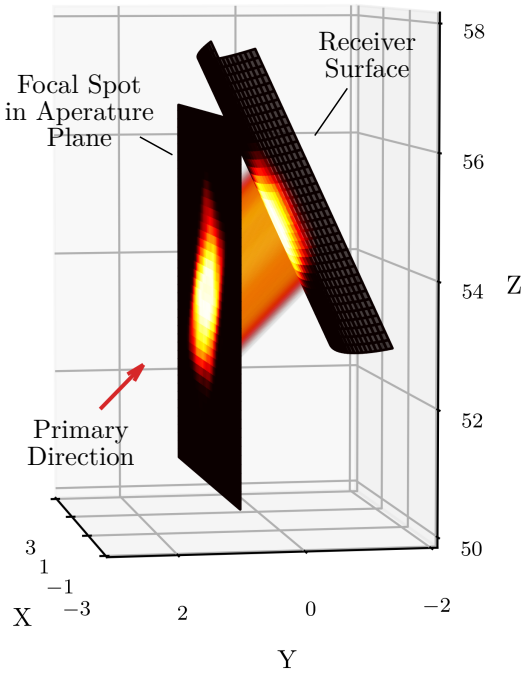

Figure 4.11.: Projection method: The focal spot in the aperture plane is converted into rays, which are traced to the receiver surface at the Solar Tower Jülich (STJ). [67]

1. **Parallel:** Utilize the primary direction, defined as the vector connecting the heliostat center to the aim point, as the direction for all rays.

2. **Angled:** Compute individual connecting vectors between the heliostat center and each pixel within the aperture plane, based on their spatial positions.

3. **CNN-based:** Employ a Convolutional Neural Network (CNN) to predict the direction of rays for each pixel.

Given that the ray sampling originates within the aperture plane, the primary direction is inherently determined by the connecting vector between the heliostat and the aim point. Consequently, deviations in ray directions are expected to be minor and primarily influenced by the heliostat's canting and focusing characteristics. The **Parallel** and **Angled** methods represent two boundary conditions: the former assumes all rays are parallel, while the latter assumes maximum divergence. In practice, the actual heliostat behavior will lie between these edge cases. However, deriving a more precise model for ray directions without detailed knowledge of the orientation of individual mirror facets is not possible.

To address this challenge, a third approach is introduced, leveraging a neural network to predict ray directions directly within the aperture plane. The proposed neural network is compact, comprising only four sequential convolutional layers. Its input features include the intensity distribution of the focal spot, the primary direction vector, and the focal distance between the heliostat and the aim point. The network outputs direction vectors for each pixel of the input image, enabling spatially resolved predictions of directional variations. This allows the network to identify regions within the focal spot where rays converge or diverge. For example, it can adaptively infer how ray directions should vary in scenarios where the focal spot is compact and the heliostat is located at a significant distance from the aperture plane.

The neural network is trained using artificial data generated from deflectometry measurements of a limited number of heliostats. These measurements are utilized to simulate flux distributions across several parallel planes positioned behind the aperture, providing a diverse dataset for learning ray direction patterns. The selection of the method for defining ray directions depends on the canting and focusing properties of the heliostat. A detailed evaluation of the optimal strategy is presented in the evaluation chapter.

Once the rays are defined, standard ray tracing operations, including intersection tests, are performed to calculate the energy impact on the receiver surface. This process is depicted in Figure 4.11, where the focal spot image is projected onto the complex receiver surface of the Open Volumetric Receiver at the Solar Tower Jülich (STJ). Initially, rays are sampled for all heliostats, followed by the calculation of all ray impact points on the receiver. This results in the aggregated total flux density distribution across the receiver surface, combining contributions from all heliostats.

Although this technique utilizes standard ray tracing principles, it differs significantly from traditional ray tracing methods. Traditional approaches start at the mirror surface and simulate the entire heliostat beam, requiring extensive sampling to capture the solar distribution and surface curvature effects. In contrast, this new projection technique begins with a pre-predicted flux distribution in the aperture plane, effectively treating it as a cross-section of the heliostat beam. This approach significantly reduces computational complexity, as the relatively short distance (0–3 meters for most receivers) between the aperture plane and the receiver surface allows for a simplified model of ray directions. As a result, our method requires less than 0.1-1% (dependent on the focal spot resolution) of the rays used in conventional ray tracing, underlining its computational efficiency.

For the assignment of rays to receiver bins, we employ the differentiable binning method proposed by Pargmann et al. [22], which maintains accuracy while minimizing the number of rays required. This method also facilitates differentiable aim point optimization, as discussed in Chapter 7. For scenarios demanding higher accuracy, such as those involving complex reactor geometries with shading or blocking effects, additional rays can be sampled by upsampling the focal spot image, thereby achieving the required level of precision.

4.3.4. Total Error Estimation Framework

The overall data-driven pipeline for flux prediction (compare Figure 4.1) comprises multiple interconnected components, each contributing its own error to the final output. To comprehensively evaluate the pipeline's performance, it is essential to deploy an approach that considers not only the individual errors of these components but also their interac-

tions and how errors propagate through the system. This ensures a holistic understanding of the pipeline's accuracy and reliability.

Error Propagation Model

The total error E_{total} of the pipeline is estimated by modeling error propagation across its stages and assessing their combined impact on the final output. The total error E_{total} is expressed as:

$$E_{\text{total}} = \sqrt{E_U^2 + E_G^2 + E_P^2 + 2 \cdot \text{Cov}(E_U, E_G) + 2 \cdot \text{Cov}(E_G, E_P)}, \tag{4.9}$$

where:

- E_U: Error from the UNet, responsible for translating raw images to flux densities.
- E_G: Error from the generative model, predicting focal spots on the aperture plane.
- E_P: Error from the projection stage, mapping fluxes to the receiver surface.
- $\text{Cov}(E_U, E_G)$: Covariance between E_U and E_G, reflecting interactions between the UNet and Transformer errors.
- $\text{Cov}(E_G, E_P)$: Covariance between E_G and E_P, capturing how Transformer errors influence projection errors.

Key Assumptions

The error propagation model is based on the following assumptions:

1. Errors at each stage follow Gaussian distributions, supported by central limit theorem considerations [69].
2. First-order covariances ($\text{Cov}(E_U, E_G)$, $\text{Cov}(E_G, E_P)$) dominate error interactions, while higher-order dependencies are negligible.
3. Errors propagate sequentially, with no direct influence from non-adjacent stages (e.g., E_U does not directly affect E_P).

Covariance Estimation

Covariance terms are estimated using Monte Carlo simulations. Realistic perturbations are introduced at intermediate stages, and their effects on subsequent stages are quantified. The covariances are calculated as:

$$\text{Cov}(E_U, E_G) = \text{mean}\left((\text{Perturbation Magnitude in } X) \cdot (\text{Error in } X(s))\right), \tag{4.10}$$

$$\text{Cov}(E_G, E_P) = \text{mean}\left((\text{Perturbation Magnitude in } X(s)) \cdot (\text{Error in } X_p)\right). \tag{4.11}$$

The perturbation methodology used to evaluate these covariances is detailed in Appendix A.3.

— By integrating individual error contributions and their interdependencies, the total error estimation framework offers a robust assessment of the pipeline's overall performance, spanning from raw image processing to flux prediction on the receiver. This methodology facilitates a comprehensive evaluation of the system's strengths and limitations, thereby informing and guiding future enhancements in each pipeline component.

4.4. Overall Workflow for Operation

After explaining the individual steps outlined in Figure 4.1, this section summarizes how these components can be employed during the operation of a solar tower power plant. By this stage, we assume sufficient data has been acquired and both the UNet for target image segmentation and the generative model for focal spot prediction have been successfully trained. The overall operational workflow is as follows:

1. **Translation of Target Images to Flux Distributions:** Raw target images obtained from the calibration process are processed by the UNet to derive relative flux distributions X (4.1).

2. **Encoding Heliostat Properties:** All available focal spot measurements for each heliostat are utilized to encode its latent heliostat properties $\tilde{H}$ using one of the encoding approaches discussed earlier (4.2.1).

3. **Setting Aim Points:** Aim points, which serve as inputs for the generative model, for the current time step are determined either through a pre-defined aim point distribution metric or an aim point control algorithm.

4. **Focal Spot Prediction:** For each heliostat, the encoded properties $\tilde{H}$, the current sun position, and the assigned aim point are input into the generative model to predict the relative focal spot distribution at the given conditions (4.2).

5. **Intensity Scaling and Ray Tracing Preparation:** The predicted relative flux distributions for each heliostat are scaled into absolute flux distributions X_a using the Intensity Model (4.3.1). These absolute distributions are then transferred into the ray tracing domain for further processing (4.3.3).

6. **Flux Assignment to Receiver Surface:** Using ray tracing techniques, the rays derived from each focal spot are projected onto the receiver surface. The intersection between rays and the receiver mesh is calculated to assign flux values to the mesh elements. This step provides the total flux density distribution on the receiver surface (4.3.3).

7. **(Optional) Aim Point Optimization:** Based on the total superposed flux density distribution, optimization criteria can be applied to refine the aim point distribution across heliostats, ensuring operational goals such as flux uniformity or thermal stress minimization are met.

It is important to note that the steps described in points 1 and 2—processing calibration images with the UNet and encoding heliostat properties—need to be repeated only when new calibration images are acquired or during specific events such as maintenance or recalibration. In contrast, steps 3 through 7 are executed within the power plant's operational control loop, reflecting the real-time dynamics of solar tower operation.

— This workflow highlights how the methodology integrates all components into a cohesive operational framework, bridging data-driven focal spot prediction with practical considerations for solar tower control and optimization.

4.5. Computational Framework and Implementation

The development and evaluation of all methods in this work were carried out using Python 3.10, ensuring compatibility with modern libraries and tools for machine learning and numerical analysis. The primary framework for implementing neural networks was PyTorch, which provides extensive support for constructing, training, and evaluating deep learning models. One of PyTorch's key advantages is its `autograd` module, which enables automatic differentiation, streamlining the optimization process for neural networks and ensuring efficient gradient computation.

The Adam optimizer was consistently employed across all neural network training and evaluation processes. This choice was motivated by Adam's adaptive learning rate properties, which allow for efficient and stable optimization across a wide range of model architectures and data distributions.

Chapter 5.

Dataset Characteristics, Preparation, and Test Scenarios

A robust and meticulously curated dataset serves as the foundation for the success of any machine learning methodology. In this chapter, the two primary sources of data employed for the training and evaluation of the proposed approach are presented. The characteristics and inherent properties of these data sources are analyzed in detail. Furthermore, the preprocessing, filtering, and organization steps are described, highlighting how the raw data were systematically refined into distinct datasets specifically tailored to address various test cases and scenarios within this work.

The first two sections focus on the main data sources: (1) measured focal spot data from the Solar Tower Jülich (STJ), obtained during the heliostat calibration process (Section 5.1), and (2) simulated data generated through a combination of deflectometry measurements and ray tracing (Section 5.2). Building on these data sources, specific test cases are defined, each associated with a Dataset Abbreviation (DA). These abbreviations concisely describe the data used for training and evaluation. The details of this abbreviation convention are presented in Section 5.3.

All data, including calibration images, deflectometry measurements, kinematic settings, and weather information, from the Solar Tower Jülich (STJ) power plant has been published in the PAINT database [70]. This database provides operational data of concentrating solar power plants in accordance with FAIR data principles, ensuring findability, accessibility, interoperability, and reusability. Furthermore, all data relevant to this work is included, supporting transparency, reproducibility, and alignment with open science guidelines.

5.1. Focal Spot Data from Heliostat Calibration Procedure

The primary dataset used in this work comprises focal spot data measured at the STJ during heliostat calibration procedures (as described in Section 3.3). During calibration, raw images of focal spots are collected from dedicated calibration targets. For our purposes, this raw data is processed further to extract focal spot information, as detailed in Section 4.1. This processed data is subsequently used to train and evaluate the predictive models.

Since the calibration data was not originally acquired for training flux prediction models, it is crucial to evaluate its characteristics and apply pre-filtering methods to create balanced and representative datasets. These datasets are tailored for specific test cases, ensuring robust model training and evaluation.

5.1.1. Preparation of Raw Images

The raw images are rectified and cropped to the target edges using marker detection. Following this, the cropped raw images are translated into flux density distributions using the UNet framework described in Section 4.1. From the resulting flux distribution, the center of the focal spot is determined using a weighted center of intensity calculation. The image is then further cropped around the calculated center to dimensions corresponding to 6 m × 6 m. This ensures that the focal spot data is independent of tracking errors, as all focal spot images are consistently centered around their respective centers of intensity.

5.1.2. Data Characteristics

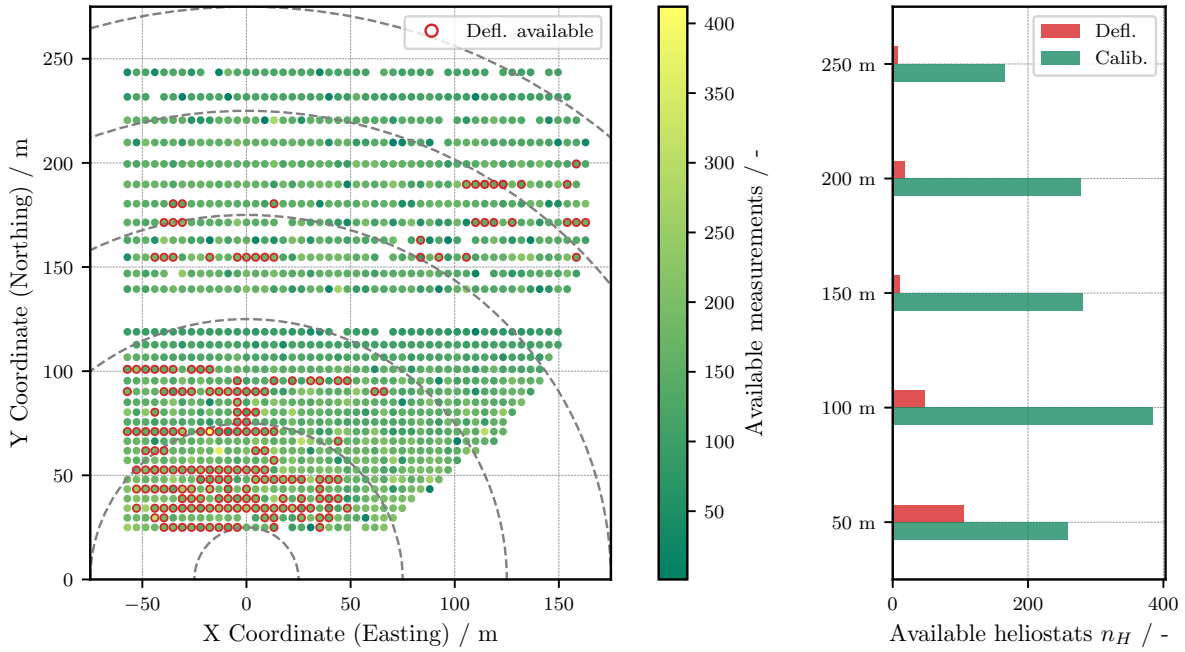

Figure 5.1.: Left: Overview of the heliostat field showing data point distribution and available deflectometry measurements per heliostat. Right: Heliostat field divided into groups based on distance from the tower, with available heliostats per group.

The dataset consists of over 200,000 data points collected between 2023 and 2024. Each data point corresponds to a focal spot image X, recorded under specific input conditions: aim point $\vec{Ap}$, sun position $\vec{S}$, and heliostat ID H_i, which indicates the heliostat from which the focal spot originates. The data was gathered during routine heliostat calibration procedures conducted at the STJ. To ensure the dataset's suitability for training and evaluation, its key characteristics, including the total number of data points, their distribution across heliostats, and the variation in input conditions is analyzed.

Figure 5.1 provides an overview of the heliostat field, highlighting the data availability for each heliostat. The left panel shows the field layout, where each dot represents a heliostat, colored according to the number of available data points. The dataset includes approximately 1,350 heliostats, each with at least 10 data points. However, the distribution is uneven: some heliostats have as many as 600 data points, while others have fewer than 20. This imbalance must be addressed to prevent biased training and evaluation.

To facilitate evaluation, the heliostats are grouped based on their distance from the receiver. Previous studies have shown that this distance significantly influences flux prediction and target image-based heliostat characterization [22, 58]. The heliostat field is divided into 50-meter intervals, from 50 meters to 300 meters. As shown in the right panel of Figure 5.1, each distance group contains a sufficient number of heliostats for robust analysis.

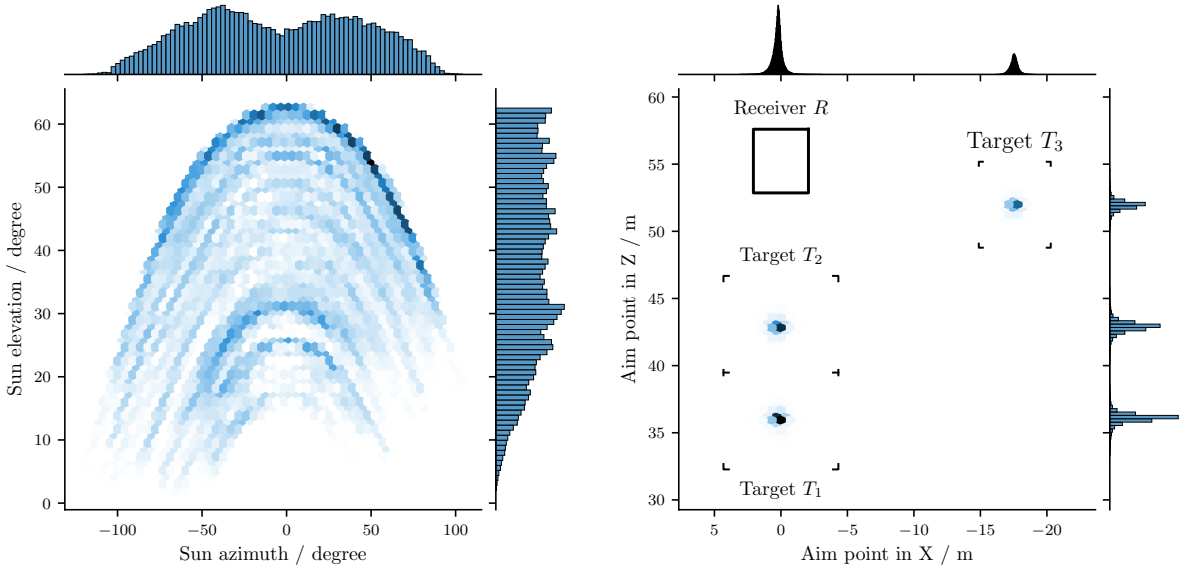

Figure 5.2.: Distribution of input conditions: Sun position $\vec{S}$ (left) and aim point $\vec{Ap}$ (right) across the dataset. The right plot also indicates the positions of the calibration targets and the receiver.

Figure 5.2 illustrates the distribution of input conditions in the dataset. The left panel shows the distribution of sun positions, with coverage spanning the full range of relevant conditions over the year. The right panel displays the aim point distribution, which clusters around the three calibration targets used during the calibration procedure. Notably, although heliostats are intended to aim at the centers of these targets, aim point inaccuracies introduce variability around the target centers. This natural variability is valuable for evaluating the model's robustness to input variation.

While the dataset comprehensively covers all relevant input conditions, it is important to account for biases, such as the increased number of samples near the solstices due to measurement campaigns. Such imbalances must be mitigated during dataset preparation to ensure unbiased training and evaluation.

5.1.3. Data Prefiltering and Dataset Creation

To create balanced and representative datasets, pre-filtering procedures are implemented to ensure an equal distribution of data across heliostats and input conditions. The following sections outline the methods used to achieve this balance.

Splitting Data per Heliostat and Calibration Target

To ensure equal representation, the number of training and test samples per heliostat is standardized. This approach prevents the overrepresentation of heliostats with a larger

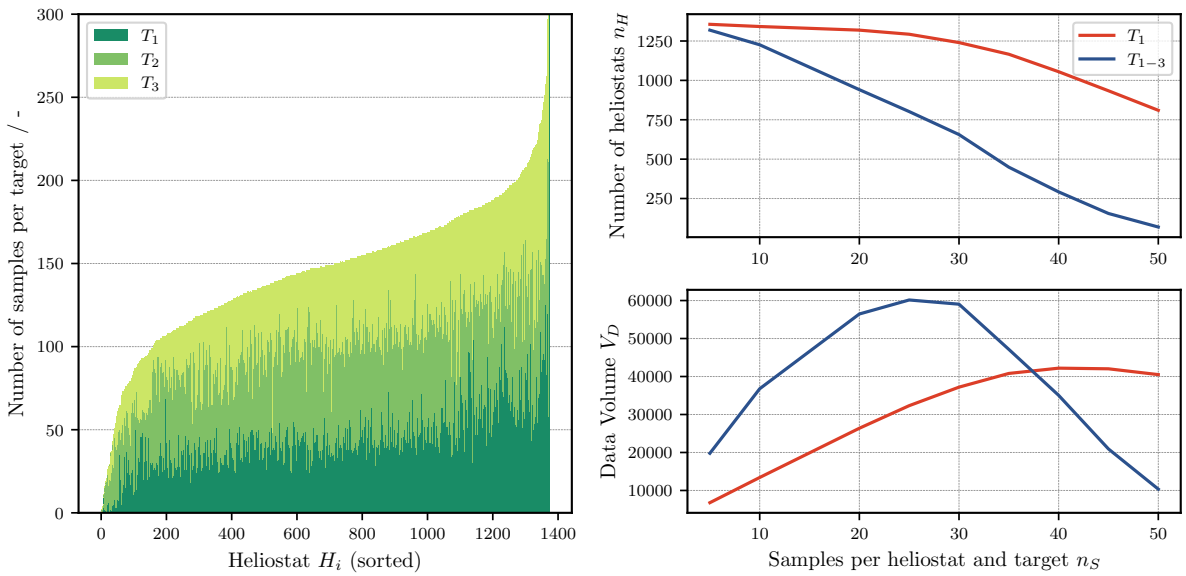

Figure 5.3.: Left: Number of measurement samples per target for each heliostat, sorted by the total number of measurements. Upper right: Number of heliostats n_H meeting the condition of having more than n_S samples. Lower right: Resulting data volume V_D for a given threshold n_S.

number of samples, which could otherwise introduce biases into the training process. Although weighting heliostat samples based on their availability is a potential alternative, it risks amplifying existing biases. Therefore, a uniform sample size per heliostat is maintained to promote fairness and consistency in the dataset.

For cases involving multiple calibration targets, equal representation across targets is ensured. Figure 5.3 (left) illustrates the number of samples per heliostat for each target, highlighting an uneven distribution that underscores the necessity of this pre-filtering step.

To determine the optimal threshold n_S (the number of samples per heliostat per target), a balance must be struck between two key objectives: maximizing the dataset size and ensuring sufficient samples per heliostat for accurate predictions. The choice of n_S directly influences both the number of eligible heliostats n_H and the resulting data volume V_D. As shown in Figure 5.3 (right), increasing n_S results in a more stringent criterion, leading to fewer heliostats meeting the required number of samples. While a higher n_S ensures more robust and representative data for each heliostat, it reduces the number of eligible heliostats and the overall dataset size. Conversely, a lower n_S allows for the inclusion of more heliostats but may compromise the statistical reliability of the data for individual heliostats.

For the case T_1, an optimal value of $n_s = 40$ was selected, with $n_{\text{test}} = 15$ of the 40 samples allocated for validation and the remaining 25 samples used for training. For the case T_{1-3}, a threshold of $n_s = 25$ was chosen, with $n_{\text{test,target}} = 5$ samples per target, resulting in a total of 15 test samples across all targets. An investigation revealed that further splitting the validation set into separate test and validation subsets did not yield any difference, as the loss values were identical for both subsets.

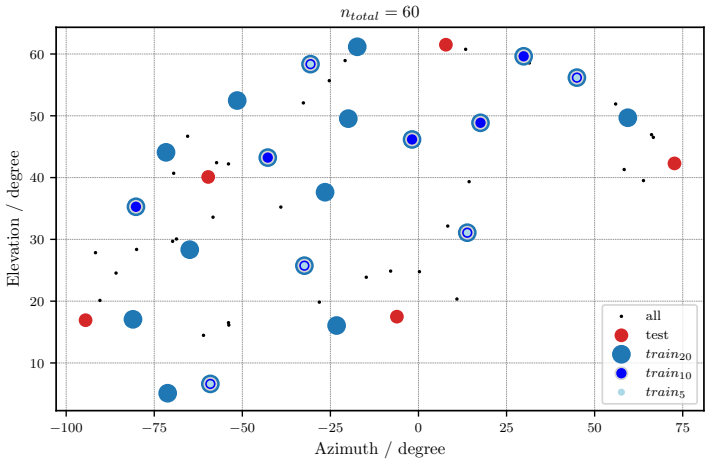

Figure 5.4.: Sun positions for one heliostat shown in azimuth and elevation. Maximum Distance Sampling was applied to generate a test dataset and three training datasets: $D_5 \subseteq D_{10} \subseteq D_{20}$.

Maximum Distance Sampling for Sun Positions

To ensure balanced sun position coverage, a Maximum Distance Sampling algorithm is developed. The algorithm begins by randomly selecting an initial data point and then iteratively selects subsequent points that maximize the angular distance in azimuth and elevation from the points already chosen. This approach generates a well-distributed subset of data points across the sun position range, as illustrated in Figure 5.4. Furthermore, this ordered selection process ensures that smaller subsets (e.g., D_5) are nested within larger subsets (e.g., D_{10}), maintaining consistency across different subset sizes.

This strategy ensures broad coverage of sun positions during evaluation and enables robust analysis of the impact of sample size on prediction accuracy.

— The calibration dataset serves as the primary data source for the training and evaluation of the generative models. For most evaluation cases and test scenarios, this dataset is used exclusively.

5.2. Artificial Dataset through Raytracing

In addition to measured calibration data, simulated focal spot data is utilized, generated through a combination of deflectometry measurements and ray tracing. Deflectometry provides highly accurate heliostat surface profiles (as described in Section 3.4.2), which are used as input for ray tracing to simulate realistic focal spots under various conditions.

Evaluation and Comparison

Although deflectometric measurements are not strictly required for the methodology presented in this work, they can serve as a valuable tool for evaluation and validation. Deflectometry combined with ray tracing is considered one of the most accurate methods for flux prediction[71]. Comparing the predictive accuracy of our proposed approach against these highly accurate simulations allows for an objective assessment of its performance and robustness. This evaluation also highlights the potential of our methodology as a viable alternative that requires less detailed heliostat surface information.

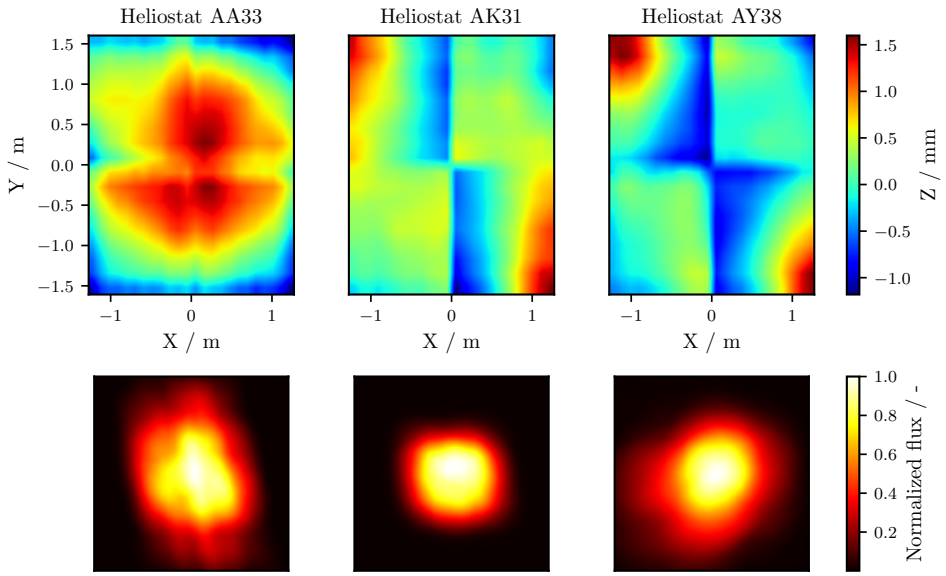

Figure 5.5.: Samples of deflectometry-measured surfaces (top row) and their corresponding simulated focal spots (bottom row) for three heliostats at solar noon.

Artificial Datasets for Training

Deflectometry and ray tracing can also be leveraged to create artificial datasets for training. By simulating focal spots under arbitrary input conditions, such as diverse sun positions and aim points, artificial datasets can augment real-world data. This is particularly useful for training models to extrapolate behavior for unseen input conditions.

For instance, as discussed in Section 4.2.4, artificial datasets can be used to enhance spatial aim point extrapolation capabilities. By simulating a wide range of aim points, including those outside the calibration target range, the network can learn to generalize beyond the observed data. This is especially beneficial in scenarios where direct measurements on a receiver surface are infeasible.

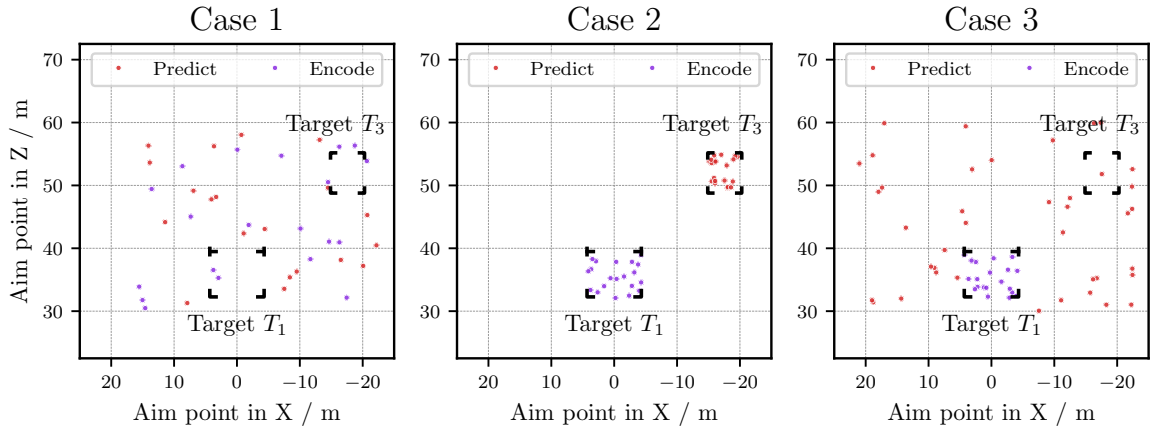

Figure 5.6.: Aim point distributions for different artificial dataset cases. Purple: Aim points used for encoding. Red: Aim points used for prediction loss calculation.

Three distinct artificial datasets are defined to evaluate the influence of aim point distributions on extrapolation performance (illustrated in Figure 5.6):

- **Case 1:** Simulated focal spots cover the entire aim point range, including calibration targets, the receiver, and intermediate regions.
- **Case 2:** Data from Target T_1 is used for encoding, while predictions are enforced on Target T_3 to explicitly evaluate spatial extrapolation capabilities.
- **Case 3:** Calibration target data is used for encoding, while randomly sampled aim points across the broader range are used for prediction loss evaluation. This scenario balances aim point generalization with extrapolation.

These cases are tested in Section 6.3.2 to identify the best approach for enhancing spatial extrapolation accuracy.

— The artificial datasets are utilized for the Spatial Aim Point Extrapolation Enhancement (Section 4.2.4) and for training the neural network (*NN*) used in the Projection approach (Section 4.3).

5.3. Dataset Abbreviation Convention

Given the wide variety of test scenarios and datasets discussed, a clear and concise convention for dataset annotation is essential. In Chapter 6, these scenarios are evaluated using the datasets described here. To streamline discussions and maintain clarity a Dataset Abbreviation Convention that explicitly defines the training and evaluation sets for each case is introduced.

$$\text{Dataset:} \quad \mathcal{D}\left[\underbrace{nT_{x_i}(+A_c)}_{\text{Train}} \rightarrow \underbrace{nT_{x_i}}_{\text{Eval}}\right] \tag{5.1}$$

A dataset $\mathcal{D}$ is defined by its training and evaluation sets, separated by an arrow. The training set specifies the data used to train the model, while the evaluation set defines the data points used to assess its performance. The abbreviation includes:

- T_{x_i}: Specifies the calibration target(s) from which the data originates (e.g., T_1 for Target 1).
- n: Indicates the number of samples per heliostat from each target.
- A_c: Denotes the inclusion of artificial data. The subscript c refers to the aim point case (1, 2, or 3) used for artificial data generation, as described in Section 5.2.

For example:

$$\mathcal{D}\left[10T_{1,2} + A_2 \rightarrow 5T_3\right] \tag{5.2}$$

This notation indicates that the model is trained on 10 samples per heliostat from Targets 1 and 2, augmented with artificial data (aim point case 2), and evaluated on five focal spots per heliostat from Target 3. This convention provides a standardized method to describe the training and evaluation configurations used throughout the experimental analysis.

—

In this chapter, the characteristics and preparation of the datasets used in this work were described. The database includes real-world calibration data from the STJ and artificial datasets generated through deflectometry-enhanced ray tracing. Methods to ensure balanced representation of heliostats, sun positions, and aim points in the datasets, along with a standardized notation for describing test scenarios were introduced. Together, these datasets form the foundation for the training and evaluation of the proposed methodology, supporting rigorous testing under diverse conditions.

Chapter 6.

Evaluation and Analysis of Data-Driven Flux Prediction

This chapter evaluates the performance and capabilities of the developed models, focusing on three key aspects: **model accuracy** and **robustness and generalization**. The evaluation is designed to systematically assess the models' ability to predict and adapt across various stages of the overall workflow and under diverse conditions.

The first focus of this chapter is **model accuracy** (Section 6.3), which examines how well the proposed methods predict the flux distribution of focal spots at different stages, including segmentation, prediction, and projection. Following the workflow outlined in Figure 4.1, each stage is investigated and evaluated separately to provide a deeper understanding of the accuracy and performance at every step. This detailed breakdown allows for the identification of strengths and limitations specific to each component, paving the way for targeted improvements. Finally, a comprehensive total error estimation is performed by modeling error propagation across the pipeline stages to evaluate the overall accuracy and reliability of the proposed approach.

The second focus is **robustness and generalization** (Section 6.4), exploring the models' ability to maintain performance across varying heliostat positions and conditions, like sun positions or aimpoints. Particular attention is given to how well the models generalize to unseen scenarios, such as underrepresented sun positions or previously unseen heliostats. These evaluations aim to establish whether the proposed methods can handle real-world variability and deliver reliable results in operational contexts.

By addressing these two research questions, this chapter provides a comprehensive understanding of the models' strengths, limitations, and potential for future enhancements. The evaluation methodology, results, and insights presented here form the foundation for validating the applicability of the proposed approaches in practical solar field operations.

Certain results presented in this chapter have previously been published in peer-reviewed works. The findings related to the UNet-based flux derivation method, discussed in Section 6.2.1, were initially introduced in Kuhl et al. [65]. The evaluation of the generative modeling approach based on StyleGAN was presented in Kuhl et al. [63], while the results pertaining to the Transformer-based methodology, including spatial aimpoint extrapolation and receiver projection accuracy, were reported in Kuhl et al. [67]. Although these publications provide foundational insights into the respective methodologies, the present chapter extends the scope of analysis by offering a more comprehensive evaluation. This includes additional performance assessments, comparative studies, and an in-depth discussion of the broader implications of the results, thereby contributing to a more nuanced understanding of data-driven flux prediction framework.

6.1. Evaluation Metrics and Loss Function

For visual comparison, we present the focal spot intensities using a "hot" colormap. All focal spots are normalized such that their summed pixel intensity is identical, representing focal spots with equivalent energy. Consequently, colorbars are omitted from the plots, as the normalization ensures consistent intensity scaling across all focal spots.

To facilitate a quantitative comparison between focal spots, we introduce a pixel-wise loss function, L, that measures the discrepancy between the predicted focal spots and the ground truth. This loss function is defined as the summed pixel deviation between predicted and ground-truth flux values, normalized by the summed intensity:

$$L = \frac{\sum_i |X_i - \hat{X}_i|}{\sum_i \hat{X}_i} \tag{6.1}$$

Here, X denotes the predicted flux at pixel i, and $\hat{X}$ represents the corresponding ground truth flux value. Both focal spots are normalized before comparison.

This pixel-wise loss function was selected because it is more closely aligned with the specific requirements of flux prediction in this study. Unlike isoline-based comparison methods commonly used in similar works, which evaluate the contour shapes or boundaries of focal spots, the pixel-wise intensity comparison directly measures spatially resolved flux errors. By evaluating the intensity values at each pixel, this loss function provides a more precise and problem-relevant metric.

Additionally, this loss function was explicitly designed to be ideal for focal spot comparison. Its formulation ensures that proportional discrepancies in flux are captured across all regions of the focal spot, regardless of spatial resolution or image scale. Because the numerator and denominator scale equally with resolution, the loss remains invariant to changes in image resolution. Furthermore, regions with negligible flux have a minimal impact on the metric, as all deviations are normalized by the total flux intensity. This ensures consistent comparison of focal spots across heliostats positioned at varying distances in the field, regardless of differences in spot size or intensity distribution.

An essential aspect of this evaluation lies in its exclusive reliance on relative comparisons. By normalizing the flux distribution through the summed intensity, the loss calculation effectively disregards the absolute intensity of the focal spot. This approach is particularly appropriate for the objectives of this study, as the primary focus is on assessing the spatial distribution of flux rather than its absolute magnitude. By concentrating on relative differences, the loss function robustly captures the accuracy of spatially resolved predictions, while remaining unaffected by external variations in absolute flux levels, such as those arising from fluctuations in direct normal irradiation (DNI) or changes in reflectivity.

This carefully designed pixel-wise loss function, therefore, provides a robust and meaningful metric for evaluating the accuracy of predicted focal spots in the context of flux modeling and prediction.

6.2. Flux Derivation and Refinement: Establishing a Basis for Predictive Models

Accurate and reliable data preparation and feature extraction are essential for the effectiveness of predictive models. This section evaluates two novel methodologies: **UNet-based flux calculation** and **automated outlier detection**. These methods ensure the precision and integrity of the focal spot data, establishing a robust foundation for both model training and evaluation.

6.2.1. UNet Image Segmentation: Precise Flux Derivation from Raw Images

The proposed UNet segmentation method is an innovative approach for extracting flux distributions from target images while accounting for background radiation and reflectivity. Precise flux distributions are critical not only for providing high-quality input data to predictive models but also for enabling accurate evaluation of model predictions.

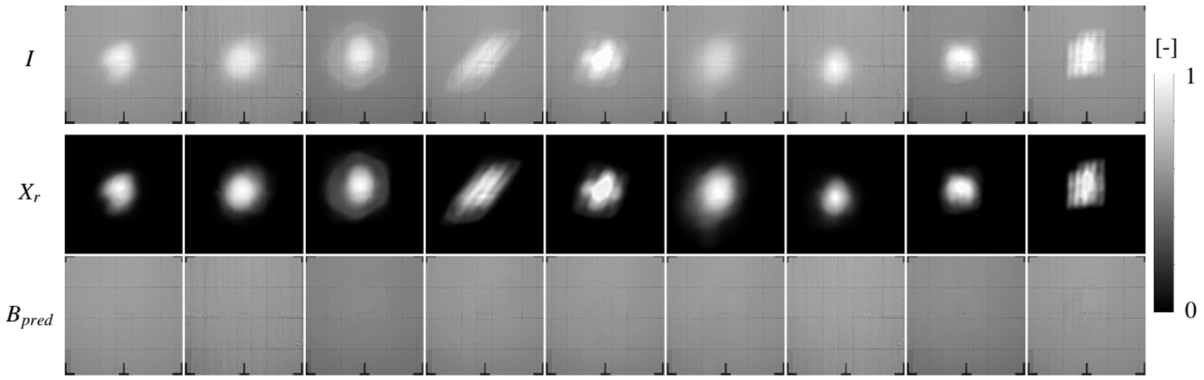

Figure 6.1.: Application of UNet Model to nine Sample Images I: The rows labeled X_r and B display the UNet's predictions for relative flux and background irradiation, respectively. [65]

Figure 6.1 illustrates the segmentation results generated by the UNet model, which was trained on simulated data and applied to real target images from the Solar Tower Jülich (STJ). The first row presents sample target images I under various conditions. The subsequent rows display the calculated relative flux distributions X and the predicted background radiation B. The flux distributions X are smooth and capture fine details of the focal spots, such as regions with low illumination or slight brightness variations. Similarly, the predicted background B accurately reflects varying conditions of background intensity.

The accuracy of detecting X from raw images is evaluated using the pixel-wise loss metric L described in Section 6.1. When applied to a test set excluded from training, the UNet achieved a relative **mean pixel deviation of 1.5%**, highlighting its precision in detecting flux distributions in superposed images. This high level of accuracy ensures that the segmented focal spots are well-suited for both predictive model training and serving as a reliable benchmark for evaluating model predictions.

6.2.2. Automated Outlier Detection for Ensuring Data Reliability

Given the scale of the dataset—comprising over 200,000 images—manually identifying and filtering erroneous samples is infeasible. An automated outlier detection method was therefore developed during training to ensure data integrity. This step eliminates problematic samples, enabling more robust and reliable evaluations.

Several anomalies were observed in the dataset, attributable to common error sources during data acquisition:

- **Multiple Focal Spots:** Focal spots from malfunctioning heliostats can appear on the calibration target while another heliostat is being calibrated, leading to overlapping or unrelated focal spots.
- **Shadowing and Blocking:** Shadows from clouds or obstructions caused by other heliostats alter the focal spot shapes unpredictably.
- **Incorrect Focal Spot Assignment:** A focal spot from a non-calibrated heliostat may cross the calibration target but be incorrectly attributed to the heliostat under calibration.

To identify and filter such anomalies, an initial predictive model (either Transformer or StyleGAN) is trained on the full unfiltered dataset. This model predicts the most likely focal spot shape for each heliostat for all calibration points. Samples with prediction errors exceeding a defined threshold—set as 2 times the standard deviation above the median—are flagged as outliers.

Figure 6.2 illustrates the outlier detection process for a sample heliostat. Most samples cluster within the median error range, while a few—four in this case—exhibit significantly higher deviations. Upon closer inspection, these outliers correspond to the described error sources.

This automated detection ensures that only reliable focal spot samples are retained for further evaluation. By integrating this process, we maintain the consistency and accuracy of the dataset without the need for impractical manual sorting.

— These foundational steps—precise flux calculations via the UNet segmentation and robust anomaly filtering through automated outlier detection—form the backbone of the predictive model evaluation framework. Together, they ensure that the models are trained and evaluated using high-quality, reliable data, enabling a rigorous assessment of their performance.

6.3. Generative Model Evaluation: Prediction Accuracy

This section serves as a central focus of the evaluation, aiming to determine whether the proposed data-driven methods can match or surpass state-of-the-art approaches in focal spot prediction accuracy. The accuracy evaluation framework is organized as follows: first, all focal spot prediction methods—both state-of-the-art approaches and the proposed data-driven methods—are systematically compared based on their prediction accuracy. Next, the performance of individual components within the data-driven pipeline is analyzed in detail. Finally, the cumulative error of the entire data-driven pipeline is estimated, providing a comprehensive assessment of its overall effectiveness.

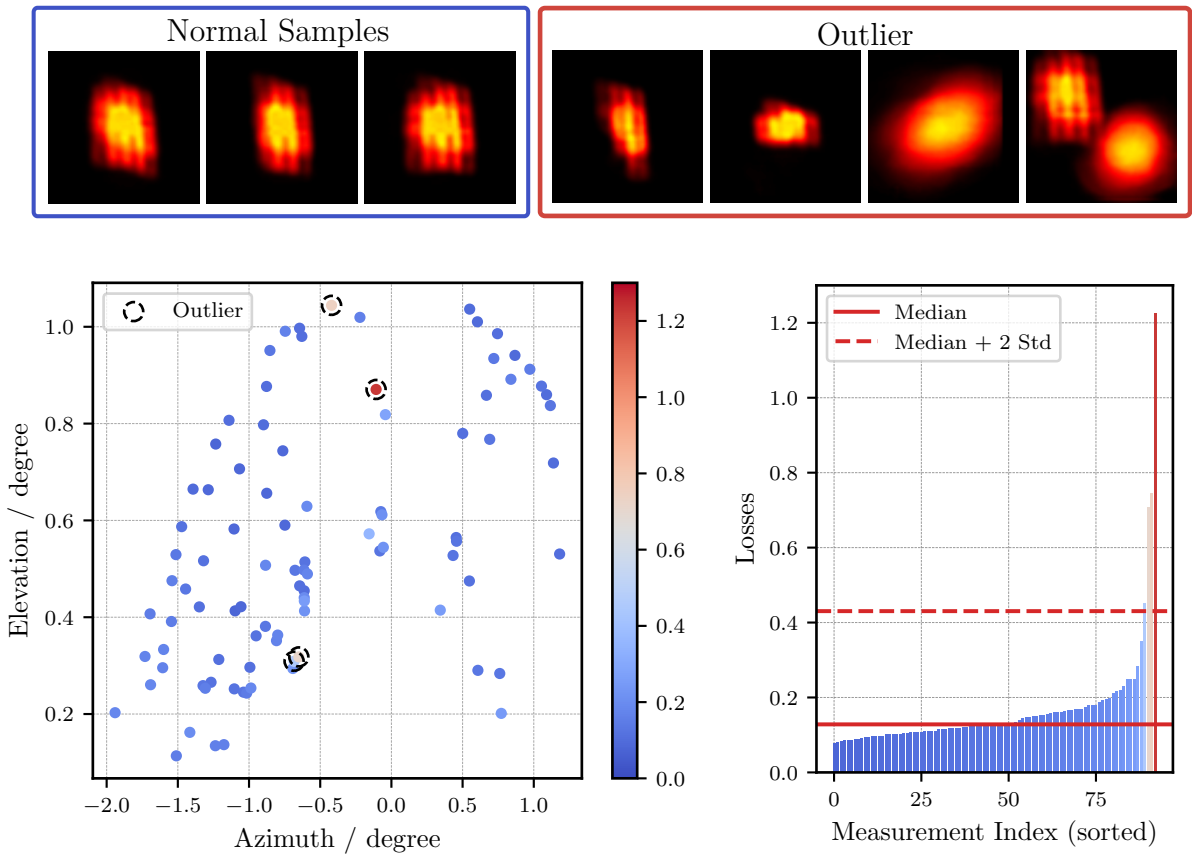

Figure 6.2.: Normal Samples: Examples of typical focal spot samples with errors close to the median. Outlier: Focal spots that exceed the defined error threshold, with identified error sources (from left to right): Unusual Shading and Blocking (2x), Incorrect Focal Spot Assignment, and Multiple Focal Spots. Left Plot: Scatter plot displaying sun positions for all calibration measurements, with error magnitudes represented by color; outliers are highlighted with circles. Right Plot: Bar plot showing sorted errors for all measurements, providing an overview of the error distribution.

6.3.1. Model Comparison

To address the lack of comparative analyses of existing heliostat flux prediction methods on a standardized dataset, a benchmark evaluation was conducted to compare the proposed data-driven approaches with current state-of-the-art techniques. Notably, such a direct comparison has not been performed previously for the existing state-of-the-art methods. This evaluation encompasses the following approaches:

- **Raytracing (Ideal):** A raytracer employing ideal heliostat assumptions, including perfect geometry with optimal canting, while neglecting mirror errors.
- **Convolutional Method:** A convolutional approach where a single parameter, σ, is fitted for each heliostat based on measured target images to model the Gaussian error distribution.
- **Raytracing (Deflectometry):** A raytracer that incorporates measured surface properties obtained through deflectometry-based heliostat characterization.

- **Data-Driven Approaches:** Two novel data-driven approaches developed in this work: the *Single Generator* (StyleGAN-based) and the *Encoder-Decoder Approach* (Transformer-based).

A detailed description of the implementation and setup of the first three state-of-the-art methods is provided in Appendix A.1.

The evaluation commences with a qualitative visual comparison of focal spot shapes across the various methods to identify discernible differences in their predicted distributions. This is subsequently complemented by a quantitative analysis utilizing a pixel-wise flux loss metric, offering an objective and robust assessment of prediction accuracy.

Focal Spot Comparison

Figure 6.3 provides a visual comparison of focal spot predictions from various flux prediction methods against ground truth measurements obtained from the calibration target. The UNet is used to translate the raw calibration target images into flux distributions, which serve as the ground truth. Predictions are evaluated for five heliostats positioned at different distances from the calibration target.

Observations from Figure 6.3:

- **Raytracing (Ideal):** For the closest heliostat (H_1), the assumption of flat mirror facets results in overly uniform focal spot shape, lacking the observed deformations in the ground truth. This discrepancy diminishes for heliostats at greater distances (H_4 and H_5), as the sun's opening angle reduces the prominence of local mirror deviations. However, deviations in focal spot orientation and size persist across all distances.

- **Convolutional Method:** For nearer heliostats (H_1 to H_3), deviations in focal spot shape are observed due to the inherent Gaussian shape assumption, which fails to capture more complex geometries. For heliostats farther from the tower (H_4 and H_5), the predictions align more closely with the measured shapes. The focal spot size is predicted accurately for all heliostats.

- **Raytracing (Deflectometry):** This method predicts the overall shape and orientation of the focal spots with high accuracy. However, discrepancies in size are noted for some heliostats, such as H_2 and H_3. These differences may arise from unaccounted environmental factors (e.g., atmospheric scattering) or heliostat-specific conditions (e.g., soiling or microscopic surface imperfections).

- **Data-Driven Approaches (StyleGAN and Transformer):** Both data-driven approaches provide accurate predictions of focal spot size and shape and demonstrate robustness to variations in environmental conditions. However, fine details, such as sharp edges in the focal spot for H_1, are less pronounced compared to deflectometry. These limitations may stem from measurement constraints, such as overexposure during data acquisition, which make those features less prominent in the measurements.

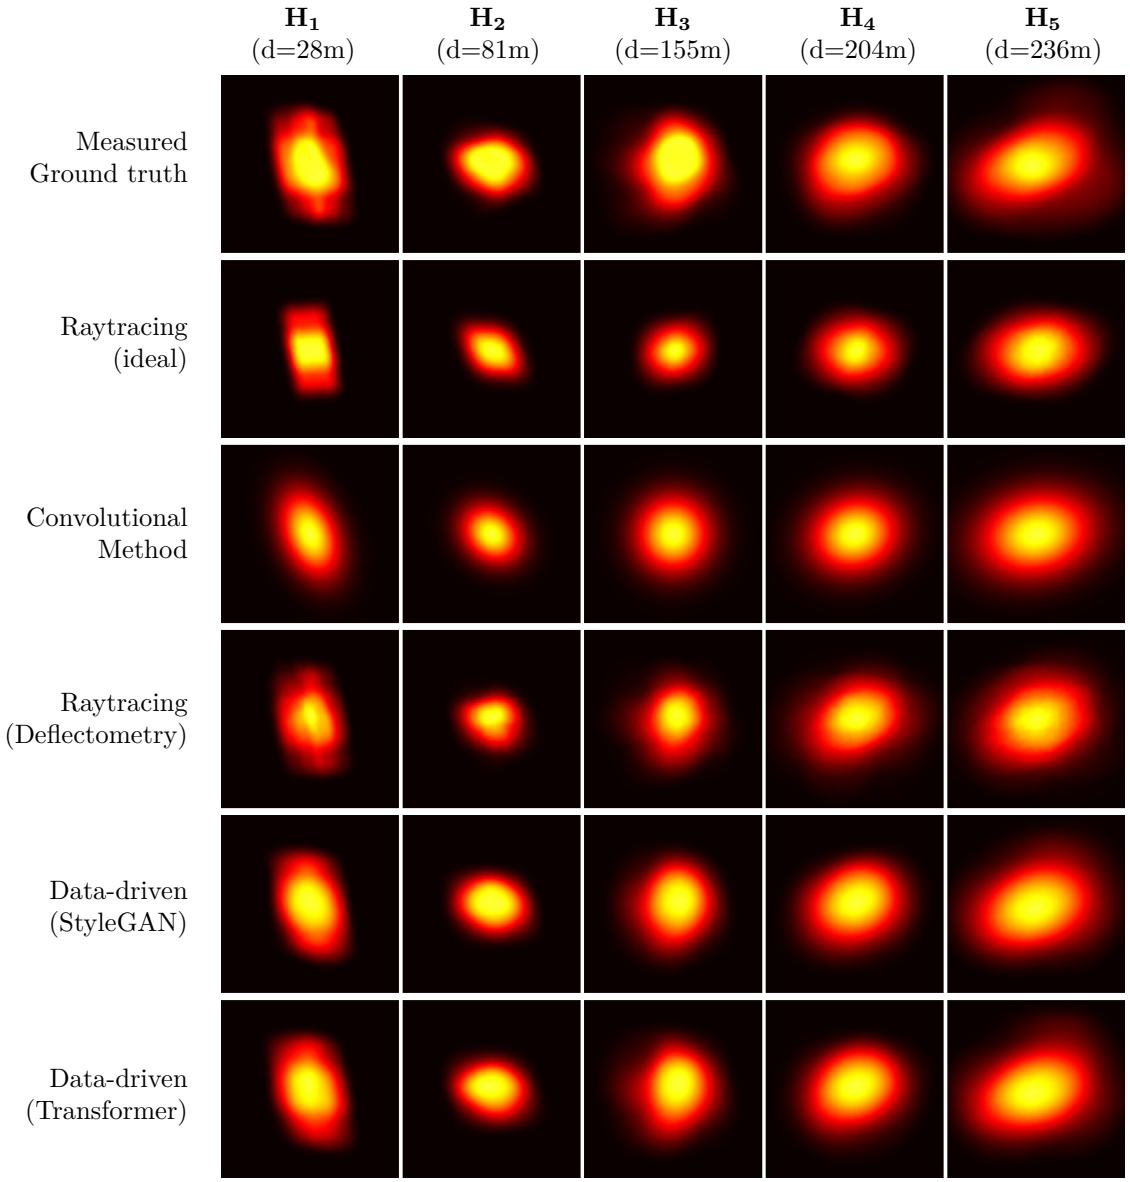

Figure 6.3.: Comparison of focal spot predictions from five different methods (rows) with measured ground truth from the calibration target (first row). Predictions are shown for five heliostats (columns) at varying distances d.

— In Conclusion the data-driven approaches (StyleGAN and Transformer) and the deflectometry-based raytracing method demonstrate the highest accuracy in focal spot prediction, albeit with different strengths. Data-driven methods excel at predicting focal spot size, as they are directly fitted to focal spot measurements, making them inherently robust to environmental variations. Deflectometry, on the other hand, achieves superior accuracy in capturing fine details and sharp edges, owing to its reliance on precisely measured surface profiles. In contrast, the ideal raytracing and convolutional methods are less effective, particularly in representing complex focal spot shapes.

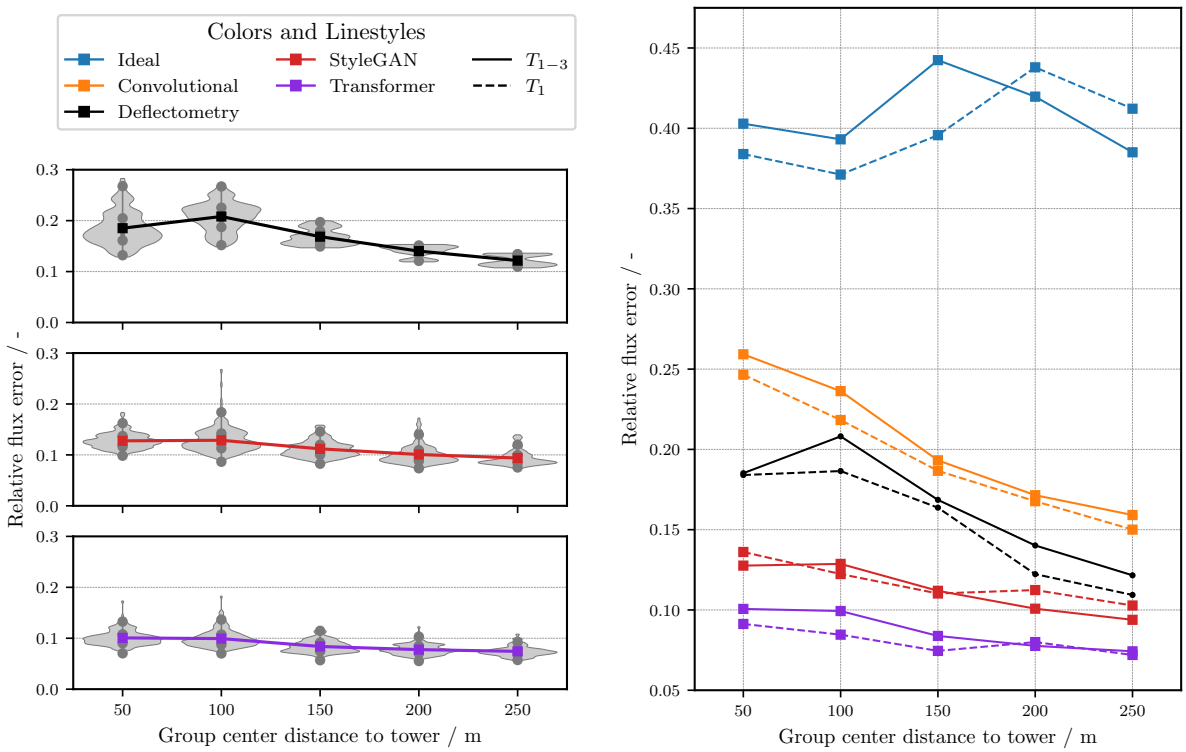

Figure 6.4.: Left: Loss curves for deflectometry and data-driven methods, grouped by distance, showing error distributions for each distance group. Right: Comparative loss curves for all prediction methods, evaluated on two datasets: one using data from a single calibration target and the other incorporating data from all three calibration targets. Datasets: $T_1 = \mathcal{D}[25T_1 \rightarrow 15T_1]$ and $T_{1-3} = \mathcal{D}[20T_{1-3} \rightarrow 5T_{1-3}]$.

Quantitative Model Accuracy

To quantitatively compare the methods, the pixelwise flux loss metric L defined in Section 6.1 is used. Losses are evaluated on calibration targets across all heliostats, grouped by their distance from the tower. The results are presented in Figure 6.4:

- **Error Distributions:** The error distributions of a heliostat group shown on the left side indicate that, for different methods and distance groups, errors are evenly distributed around the mean error. This consistency allows the accuracy of each method on a given dataset to be effectively summarized as a single error line in the plot, as illustrated in the right plot.

- **General Trends:** For most methods (excluding the ideal model), prediction accuracy improves with increasing distance from the tower. This trend can be attributed to the sun's opening angle, which causes focal spots at greater distances to more closely approximate the sun's shape. As a result, the reduced irregularities in shape make these focal spots easier to predict.

- **Ideal Model:** The ideal raytracing model shows no consistent performance trend across distance groups and generally performs poorly in flux prediction accuracy at all distances.

- **Convolutional Method:** Less accurate for nearby heliostats but shows significant improvement for distant ones, as its Gaussian assumptions become more accurate with increasing heliostat distance.
- **Deflectometry-Based Raytracing:** Among state-of-the-art methods, deflectometry-based raytracing demonstrates the highest accuracy overall. It performs particularly well for distant heliostats but exhibits increased errors for closer ones. This suggests that, while deflectometry accurately captures fine focal spot details and deviations for close heliostats (Figure 6.3), slight shifts and misalignments can still impact accuracy.
- **StyleGAN vs. Transformer:** Both data-driven methods outperform deflectometry and other state-of-the-art methods across all distance groups. Of the two, the Transformer demonstrates superior accuracy with consistently lower errors.
- **Aimpoint Variations:** Both data-driven approaches effectively predict the influence of aimpoint changes on focal spots, as demonstrated by the consistent error levels across datasets T_1 and T_{1-3}. This consistency indicates that the approaches can accurately account for aimpoint variations without compromising prediction accuracy.

As demonstrated in Appendix A.2, the error curves have been shown to be stable across variations in sun positions and heliostat subsets. This stability ensures that the loss curves reliably represent robust field-level accuracies, further validating the consistency and reliability of the evaluation methodology.

Among the state-of-the-art methods, deflectometry-based raytracing demonstrates the highest accuracy, representing the benchmark for traditional approaches. Meanwhile, the data-driven Transformer surpasses all methods, including deflectometry, in predictive accuracy. Consequently, further investigations focus on deflectometry as the state-of-the-art reference and the Transformer as the data-driven approach.

— Although the Transformer model is selected for further evaluations due to its superior performance, it is important to note that the StyleGAN-based approach also achieves accuracies surpassing state-of-the-art methods. While the Transformer demonstrates higher accuracy in the context of this study, there may be scenarios involving different heliostat types or operational conditions where the iterative learning approach of StyleGAN could provide advantages and potentially outperform the Transformer.

Accuracy on Different Calibration Targets

Given that slight deviations are observed between the accuracy results on different datasets (T_1 and T_{1-3}) without a clear trend, a detailed evaluation of accuracy on individual calibration targets is warranted. Consequently, the performance of the transformer and deflectometry methods is further evaluated across each calibration target individually. The loss curves for each target are shown in Figure 6.5.

- **Deflectometry:**
  - In the 100 m range, higher losses are observed for T_2 and T_3. This is because most heliostats in this range are canted differently from the rest of the field.

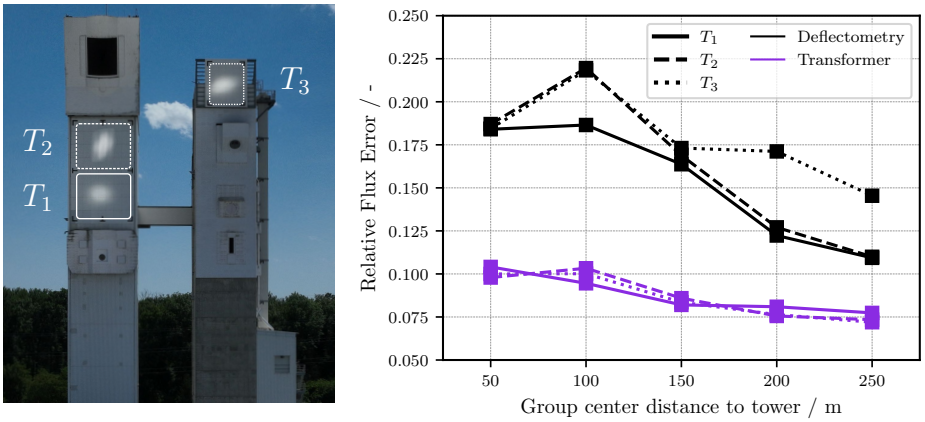

Figure 6.5.: Loss curves for deflectometry and the transformer network across three calibration targets: T_1, T_2, and T_3. For Dataset $\mathcal{D}[20T_{1-3} \rightarrow 5T_{1-3}]$, prediction loss evaluated for each target T_i individually.

Their facets are aligned with a research level located below target T_1 instead of the receiver, leading to slightly different prediction behavior.

– For distant heliostats (200–250 m), higher losses are observed for T_3 because the smaller size of this target truncates the measured focal spots, introducing additional error. It is important to note that this increased error arises from measurement errors rather than prediction method itself.

- **Transformer:**
  - Demonstrates consistent accuracy across all targets, unaffected by variations in canting for heliostats in the 100 m range.
  - Is unaffected by truncated focal spots for target T_3 in the 200–250 m range.

On the one hand, the data-driven approach demonstrates robustness to heliostat-specific parameters, such as variations in canting. Since it is directly trained on measured focal spots, it operates independently of facet canting and other external influences.

However, this also presents potential drawbacks. In the 200–250 m range, where the focal spots measured on T_3 are known to be less accurate due to truncation, the data-driven model minimizes prediction loss across all targets uniformly. This can result in predictions that attempt to balance the discrepancies between truncated focal spots (from T_3) and non-truncated focal spots (from T_1 and T_2), potentially producing intermediate predictions that do not perfectly reflect either case.

— This highlights a critical consideration for data-driven models: the quality and consistency of input and evaluation data must be rigorously assessed to prevent unintended behaviors and ensure accurate predictions.

6.3.2. Spatial Aimpoint Extrapolation Capabilities

While previous investigations accounted for aimpoint variations, they did not explicitly address the ability to extrapolate focal spot predictions to aimpoints beyond the training range. The capability to generalize predictions to unseen aimpoints is a fundamental requirement for accurate receiver-level flux predictions derived from calibration target images.

Direct measurements of focal spots on receiver surfaces are often infeasible due to low reflectivity and irregular surface properties. Therefore, to evaluate spatial extrapolation capabilities, this analysis utilizes measurements from dedicated calibration targets that are excluded from the training process (either T_2 or T_3).

The evaluation proceeds in two stages. First, the optimal artificial dataset, a critical factor in enhancing extrapolation capabilities, is identified. Subsequently, two distinct scenarios involving varying combinations of calibration targets for training and evaluation are examined in detail.

Selection of the Best Artificial Dataset

Artificial datasets play a important role in improving the extrapolation capabilities of predictive models by including data from diverse and unmeasured aimpoint ranges. As discussed in Chapter 5, multiple strategies are available for generating artificial datasets, particularly regarding the distribution of aimpoints for encoding and prediction in the Transformer network.

To identify the most effective dataset, three separate Transformer networks were trained using real data from target T_1 combined with distinct artificial datasets. The accuracy of these networks was then evaluated by predicting focal spots on target T_3 and comparing them to corresponding real measurements, using training data exclusively from T_1. Figure 6.6 illustrates the comparative accuracy of predictions on T_3 of the networks trained with the three artificial dataset cases.

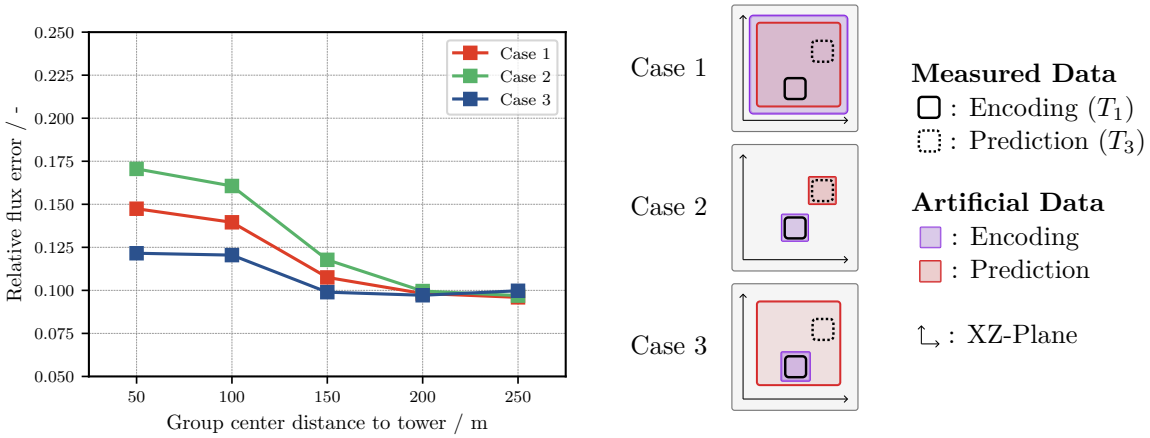

Figure 6.6.: Comparison of loss curves for the Transformer network trained with three artificial dataset cases (Case 1–3, described in Chapter 5). The plot illustrates the capability of each artificial dataset to enhance the Transformer's aimpoint extrapolation performance, enabling accurate focal spot predictions on T_3 based on images from T_1. The artificial datasets differ in the aimpoint ranges used for encoding and prediction during the Transformer's training. Datasets: $\mathcal{D}[20T_1 + A_{Case} \rightarrow 5T_3]$.

- **General Observations:** Differences in performance among the cases are primarily observed for close heliostats. Beyond the 200 m range, all cases yield comparable accuracy, as the extrapolation task becomes less complex at greater distances due to the less significant influence of the aimpoint change on the focal spot shape.

- **Case 2 - Aimpoints Restricted to Specific Targets for Encoding and Prediction:** This approach yields the lowest accuracy, as the overly restrictive aimpoint range impedes the network's ability to generalize effectively to unseen conditions.
- **Case 1 - Uniformly Distributed Aimpoints:** This case achieves moderate accuracy improvements but lacks a specific focus on aimpoint extrapolation. The limited variation between the encoding and prediction aimpoints diminishes its effectiveness for the task-specific goal of aimpoint extrapolation.
- **Case 3 - Aimpoints Restricted to Target T_1 for Encoding with a Broader Range for Prediction:** This case demonstrates the highest accuracy, effectively balancing task-specific focus with enhanced generalization capabilities across varied aimpoints.

Based on this analysis, Case 3 is identified as the most suitable artificial dataset configuration, offering superior extrapolation accuracy. This configuration is selected for subsequent evaluations to enhance the predictive performance of the Transformer network.

Evaluation Framework for Test Scenarios

To analyze the extrapolation characteristics of the predictive model, two distinct scenarios were developed. Each scenario evaluates the performance of different cases, categorized and visually represented in the plots by distinct **line styles**:

- **Baseline Model Accuracy (solid violet line):** Evaluates the model's capability to predict focal spots on the same target used during training, without involving aimpoint changes, excluding artificial datasets.
- **Extrapolation Case (dashed/dotted violet line):** Assesses the model's ability to extrapolate predictions to Target $T_{2,3}$ using only measured data, excluding artificial datasets.
- **Enhanced Extrapolation with Artificial Data (dashed/dotted blue line):** Analyzes the effect of augmenting the training dataset with artificial data to improve extrapolation accuracy.
- **Deflectometry (solid black line):** Serves as a baseline for comparison with the current state-of-the-art methods.

The color of the **markers** indicates the target on which the predictions are evaluated: purple for T_1, orange for T_2, and red for T_3.

Scenario 1: Impact of Extrapolation Distance

In this scenario, the impact of extrapolation distance on prediction error is analyzed. Measurements from Target T_1 are used to predict focal spots on Targets T_2 and T_3, with and without the inclusion of artificial datasets. Target T_3 represents a more challenging extrapolation task due to its greater spatial separation from T_1. Figure 6.7 presents the loss curves for this analysis.

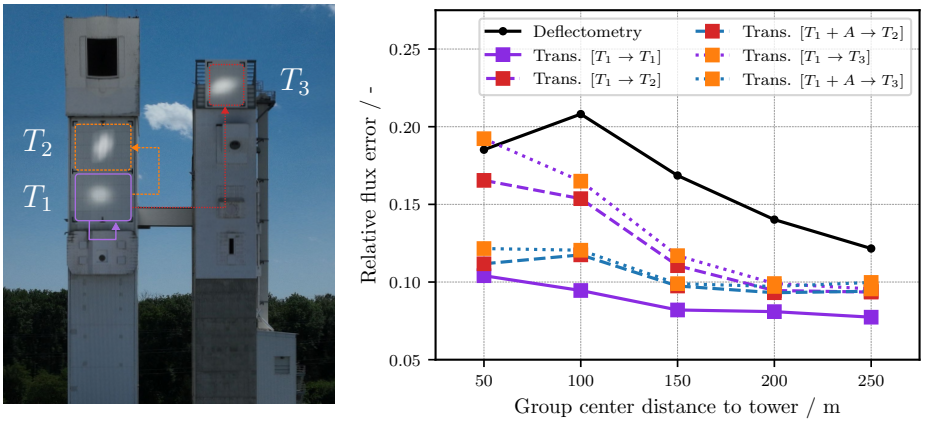

Figure 6.7.: Extrapolation accuracy for predictions on Targets T_2 and T_3 based on training with T_1. Artificial datasets (blue lines) are included to enhance predictions. Datasets: $\mathcal{D}[20T_i + (A_3) \rightarrow 5T_j]$.

- **Increased Extrapolation Distance—Higher Errors:** As expected, the prediction error increases with extrapolation distance due to the growing discrepancy in geometric and angular conditions between training and evaluation targets.
- **Closer Heliostats:** Extrapolation losses are higher for closer heliostats due to the pronounced effects of aimpoint changes on heliostat orientation, the incidence angle of radiation on the target, and the distance between the heliostat and the target plane. These factors introduce significant variations in focal spot characteristics, increasing the complexity of accurate extrapolation.
- **Distant Heliostats:** Extrapolation losses decrease for distant heliostats as the impact of aimpoint changes on orientation and incidence angle diminishes. Additionally, focal spots exhibit simpler, Gaussian-like distributions at greater distances, reducing the complexity of the extrapolation task.
- **Effect of Artificial Data:** Including artificial datasets substantially reduces extrapolation losses for closer heliostats, often achieving accuracy levels comparable to the baseline performance on T_1. However, the benefit diminishes for heliostats beyond 200 m, where the model already demonstrates strong predictive accuracy without artificial data augmentation.

This analysis highlights the utility of artificial datasets in improving extrapolation accuracy for heliostats at closer ranges, where variations in orientation and incidence angle present the greatest challenges.

Scenario 2: Influence of Aim Point Diversity in Measured Data

This scenario investigates the impact of aim point diversity in measured focal spot data on the model's extrapolation capabilities. Two cases are compared: one where the training dataset includes measurements from a single target (T_2), and another where data from two targets ($T_{1,2}$) is used. Notably, T_2 was selected as the single-target case to ensure that the minimum distance to T_3 remains consistent across both scenarios. The results, depicted in Figure 6.8, are summarized below:

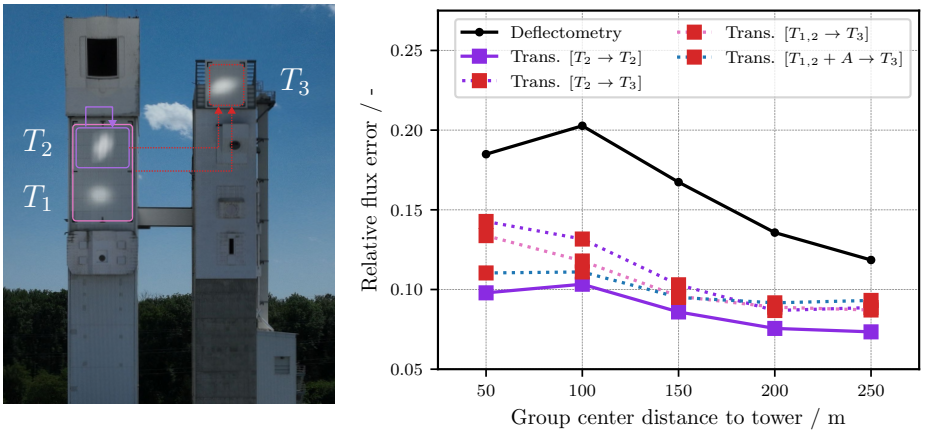

Figure 6.8.: Extrapolation accuracy for predictions on T_3, comparing training with data from T_2 alone versus a combination of T_1 and T_2. Datasets: $\mathcal{D}[20T_i + A_3 \to 5T_j]$. [67]

- **Enhanced Extrapolation with Increased Aim Point Diversity:** Models trained on a combination of T_1 and T_2 demonstrate superior extrapolation accuracy on T_3 for heliostats up to 200 m compared to models trained solely on T_2. The broader range of aim points in the training data equips the model with greater flexibility to generalize to unseen configurations.
- **Artificial Data Augments Performance:** Incorporating artificial datasets into the training process reduces prediction losses further for heliostats within 100 m of the tower. This improvement is most pronounced in closer heliostats, where geometric complexity and variations in incidence angle pose greater challenges.

These findings underscore the critical role of aim point diversity in measured data for enhancing extrapolation capabilities.

Conclusion: Enhancing Extrapolation Capabilities

This section presents the key findings from the evaluation of spatial aimpoint extrapolation capabilities:

- **Optimal Artificial Dataset Configuration:** Case 3, which restricts aimpoints for encoding to regions with measurable real data while extending the range for prediction, provides the highest accuracy. This approach effectively balances task-specific focus with broader generalization.
- **Impact of Artificial Datasets:** Artificial datasets significantly improve extrapolation accuracy for close-range heliostats by addressing variations in geometry and incident angles. However, their impact diminishes for heliostats beyond 200 m, where the baseline accuracy is already high.
- **Aimpoint Diversity in Measured Data:** Incorporating data from multiple calibration targets (T_1 and T_2) enhances extrapolation performance compared to single-target datasets (T_2 alone). The variety of aimpoints in the training data equips the model with better generalization capabilities. In practical applications, calibration areas in power plants should be fully utilized to introduce sufficient aimpoint diversity.

These findings demonstrate that combining artificial datasets with diverse measured data significantly enhances the model's extrapolation capabilities, ensuring robust and accurate predictions in operational scenarios where focal spots measured on calibration targets are used to predict flux distributions on the receiver. It is also important to note that artificial data is not essential, as models trained solely on calibration data also achieve high accuracies and outperform state-of-the-art methods.

6.3.3. Receiver Projection Accuracy Evaluation

In the overall pipeline, the steps involved in translating raw images to flux densities and predicting focal spots under varying conditions, including aim point extrapolation using a transformer network, have been addressed. These steps enable the prediction of focal spots on the flat aperture plane in front of a receiver.

This section investigates how the prediction accuracy varies when projecting focal spots from flat aperture planes onto complex receiver structures. The different projection methods, as outlined in Section 4.3, are evaluated in this context. The evaluation of the projection technique is divided into two primary parts:

1. Analysis of the error behavior as a function of projection depth, followed by the identification of the most effective projection method.

2. Application of the projection technique to compute flux distributions on two example receiver configurations—a open volumetric receiver and a cavity receiver—and direct comparison of the results on the receiver surface.

All analyses are based on simulated data derived from ray tracing combined with deflectometry measurements. While the measured surfaces may exhibit inaccuracies due to the limitations of deflectometric measurement, the ray tracing simulations itself derived from these surfaces remain consistent within the model. This consistency ensures that flux evaluations on various planes or surfaces are consistent, providing a robust assessment of the projection methods.

Projection Depth Analysis

To investigate how projection errors vary with depth, focal spots for individual heliostats are simulated at multiple depths. For each heliostat, a series of parallel planes is defined, beginning at the aperture plane ($s = 0$) and extending in 1-meter increments along the aim axis connecting the heliostat to the aperture. Ray tracing is used to simulate the focal spots at these planes. The projection method is then applied to the focal spot at $s = 0$, and the results are compared to the simulated focal spots at various depths $X(s)$. Pixelwise loss is computed for each depth, resulting in the loss curves shown in Figure 6.9a, with each curve representing a single heliostat.

The loss increases nearly linearly with projection depth for all heliostats. Therefore, this relationship can be succinctly captured by a single gradient, expressed as $\frac{\%}{\mathrm{m}}$, which quantifies the percentage loss per meter of projection depth.

—

In Chapter 4, three distinct approaches for calculating ray directions during projection were introduced: 'angled', 'parallel', and 'NN' (neural network). Figure 6.9b presents scatter plots comparing the loss gradients for each heliostat across these projection methods:

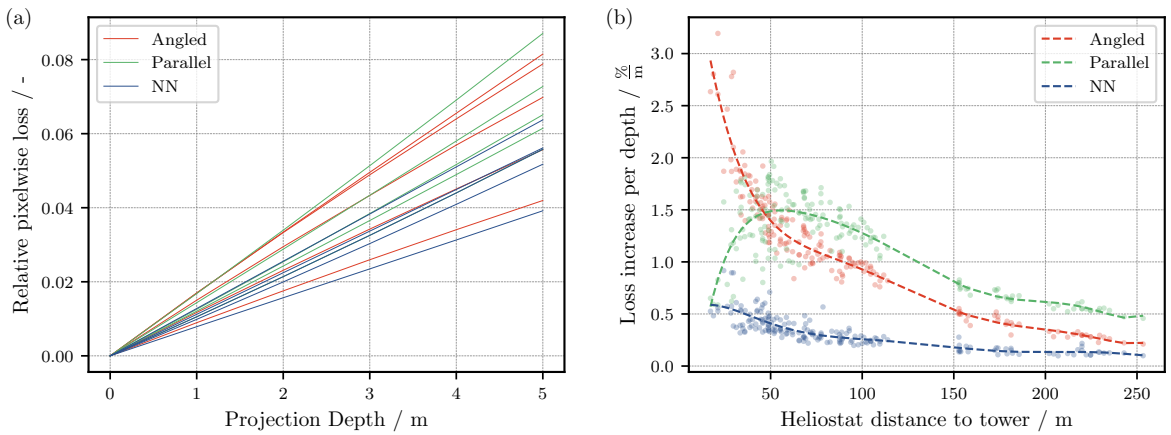

Figure 6.9.: (a) Illustration of projection depth analysis, where each line represents a single heliostat, for a few sample heliostats. (b) Scatter plot of simulated heliostats showing the gradients of the depth loss curve (from a) for each heliostat, along with trendlines for the different projection methods. (Adapted from [67])

- The parallel projection method yields higher accuracy for close heliostats, where the focal spot beam remains compact due to the short distance.
- The angled projection method performs better for distant heliostats, where the focal spot begins to spread due to the angular divergence of sunlight.
- The neural network (NN) method consistently outperforms both the parallel and angled methods, demonstrating superior accuracy across the full range of heliostat distances.

Based on these findings, the NN-based projection approach is selected for further investigation.

Evaluation on Receiver Surfaces

Building on the projection depth analysis conducted on flat planes, the projection technique is further evaluated on two distinct receiver configurations:

1. An open volumetric receiver (OVR), as installed at the Solar Tower Jülich.
2. A cavity receiver, modeled as a cylindrical opening, representing a typical configuration for solar reactors employed in solar fuel production.

Figure 6.10 illustrates the geometries of the two receiver configurations. Using deflectometry measurements, heliostats are simulated to evaluate flux distributions both on the aperture plane and on the receiver surfaces. The projection method is applied to the focal spots on the aperture plane, and the resulting flux distributions on the receiver surfaces are compared to the simulated ground truth:

- For the open volumetric receiver (OVR), the mean deviation between the projected and simulated flux distributions is **2.59%**.
- For the cavity receiver, the mean deviation is **2.36%** (2.82% with Spillage).

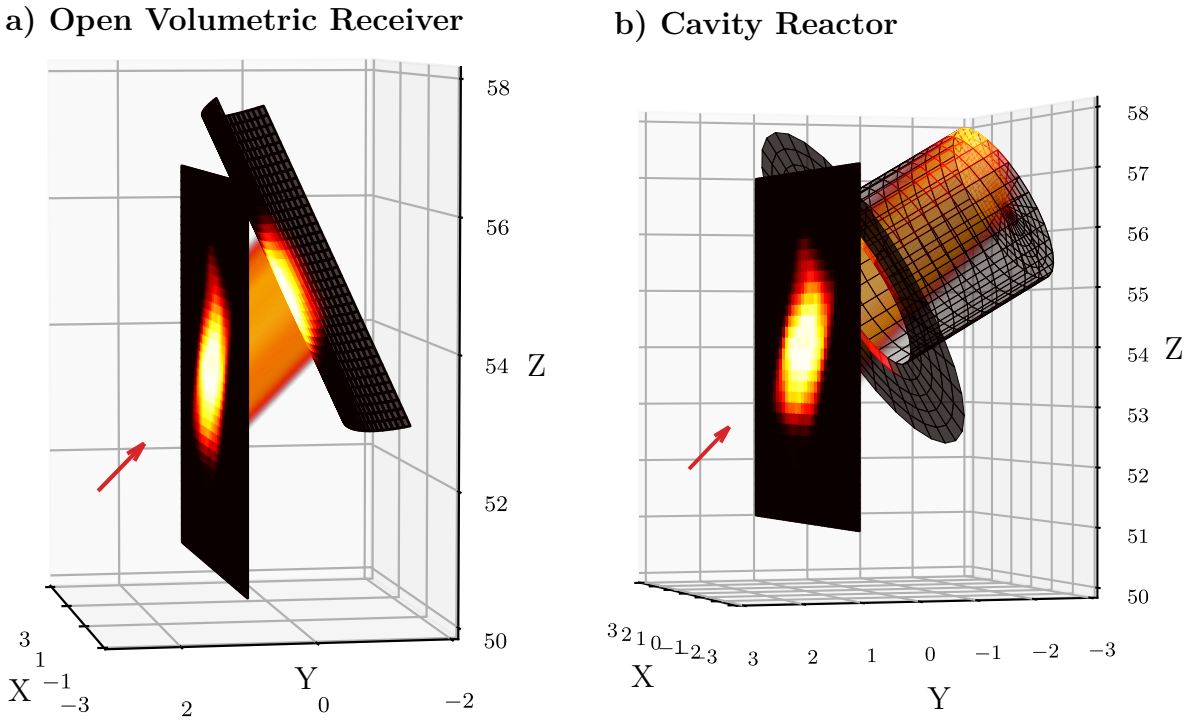

Figure 6.10.: Receiver configurations used for evaluation with focal spot in aperature projected along the primary direction (red arrow). a): Open volumetric receiver. b): Cylindrical cavity receiver. (Adapted from [67])

- These deviations are comparable to the per-meter loss gradients observed in the projection depth analysis; however, the angled geometry of the surface can introduce a small additional error.

The slightly lower projection errors observed for the cavity receiver compared to the open volumetric receiver (OVR), despite the greater projection distance, can be attributed to the coarser resolution of the circular mesh used for the cavity receiver.

— The results demonstrate that the novel projection method allows for accurate projection of focal spot predictions from the aperature plane onto various receiver surfaces, with minimal loss in accuracy, making it applicable to a wide range of receiver configurations.

6.3.4. Total Error Analysis

The evaluation thus far has focused on analyzing the individual errors associated with each stage of the pipeline. These stages include translating raw images to flux densities via the UNet ($\mathbf{U}(I)$), predicting focal spots on the aperture plane using the Transformer ($\mathbf{G}(X_i, \vec{S}, \vec{Ap})$), and projecting fluxes onto the receiver surface ($\mathbf{P}(X(\vec{S}, \vec{Ap}))$). While these analyses provide valuable insights into the accuracy of each component, they do not consider how errors propagate through the pipeline or interact with one another.

This section aims to estimate the total error E_{total} of the pipeline by modeling error propagation across stages and assessing their combined impact on the final output, as described in Section 4.3.4. Such an analysis is crucial for understanding the overall robustness and reliability of the system.

Error Assumptions for Total Error Analysis

The total error analysis integrates the following assumptions for the individual error components and their covariances, based on measured or estimated values from the evaluation:

- E_U: The error from the UNet ($\mathbf{U}(I)$) is assumed to be constant at 1.5%, reflecting its evaluated performance in translating raw images to flux density distributions.
- E_G: The Transformer's focal spot prediction error (E_G) is estimated from the loss curves (Figure 6.8) obtained for the dataset $\mathcal{D}[T_{1,2}+A_3 \rightarrow T_3]$, focusing on aimpoint extrapolation scenarios.
- E_P: The error introduced during the projection stage ($\mathbf{P}(X(\vec{S}, \vec{Ap}))$) is calculated per distance group for the open volumetric receiver (OVR) geometry.
- $\mathrm{Cov}(E_U, E_G)$: The covariance between the UNet and Transformer errors is estimated through a Monte Carlo approach, where random perturbations are introduced in the UNet outputs, and their impact on Transformer predictions is quantified.
- $\mathrm{Cov}(E_G, E_P)$: The covariance between the Transformer and Projection stage errors is similarly evaluated using Monte Carlo simulations. Perturbations in focal spot predictions are propagated through the projection step, and the resulting errors are analyzed.

These error values and covariances are used in the total error calculation to provide a holistic understanding of the pipeline's accuracy and its robustness across multiple stages.

Results

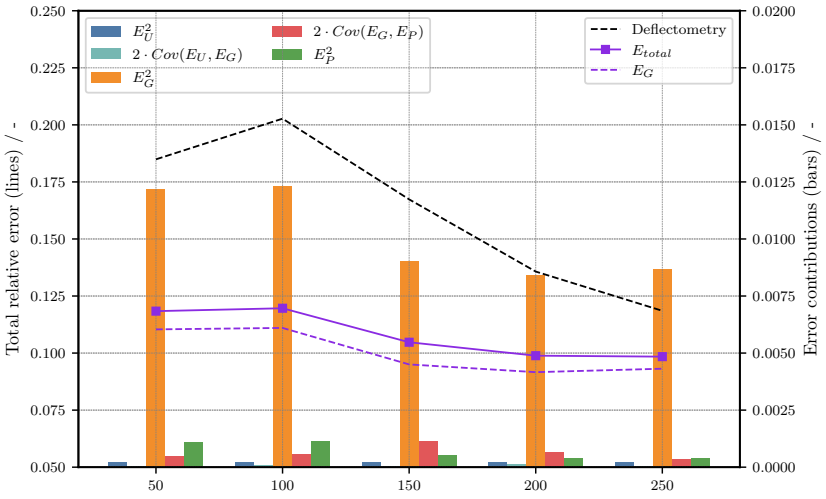

Figure 6.11.: Total error estimation for the complete pipeline, combining aimpoint extrapolation and receiver projection. Bars represent individual error contributions; lines indicate total error. Results are compared to the baseline deflectometry method. [67]

Figure 6.11 summarizes the total error analysis, highlighting the following key findings:

- **Dominant Error Source:** The largest contributor to E_{total} is the Transformer (E_G), reflecting the complexity of focal spot prediction. This suggests future efforts should prioritize enhancing the Transformer's accuracy.
- **Projection Error:** Errors from the projection stage (E_P) and their covariance with Transformer errors ($\text{Cov}(E_G, E_P)$) are relatively small, contributing minimally to the total error.
- **Error Independence:** The negligible covariance term $\text{Cov}(E_U, E_G)$ demonstrates the Transformer's robustness in mitigating upstream errors, confirming its ability to process multiple focal spots effectively.
- **Validation of Assumptions:** The analysis validates the assumption that first-order dependencies dominate error propagation. Higher-order interactions were found to be insignificant, justifying their exclusion from the model.
- **Comparison with Deflectometry:** The total error of the proposed pipeline remains significantly lower than that of deflectometry-based methods, even under challenging scenarios combining aimpoint extrapolation and receiver projection.

The total error analysis confirms the robustness and accuracy of the proposed data-driven pipeline, which outperforms deflectometry in terms of total error. The modular design effectively mitigates error propagation, ensuring high precision in flux density predictions across all stages of the pipeline.

6.4. Robustness and Generalization Evaluation

Building on the prior analysis, which established that the data-driven methodology demonstrates superior accuracy and surpasses state-of-the-art approaches, this section shifts focus to the evaluation of robustness and generalization. These factors are critical for assessing the practical applicability and scalability of the methodology to real-world scenarios.

This section systematically examines the robustness and generalization capabilities of the prediction model across a range of scenarios. Given that the transformer-based model exhibited the highest accuracy in previous evaluations, it is chosen as the primary subject of this analysis. The evaluation encompasses key aspects such as performance under edge-of-distribution (EOD) conditions, the impact of dataset size on model effectiveness and the ability of the model to generalize to unseen scenarios.

6.4.1. Edge-of-Distribution (EOD) Accuracy

EOD accuracy evaluates the model's ability to handle extreme or rare input conditions, such as unusual sun positions. These conditions are underrepresented in the training data due to the uneven distribution of sun angles during calibration. Figure 6.12 shows the distribution of calibration points across azimuth and elevation angles and their corresponding mean losses losses for this segment:

- Rare conditions, such as high/low azimuth and low elevation angles, exhibit higher losses due to their underrepresentation during training.

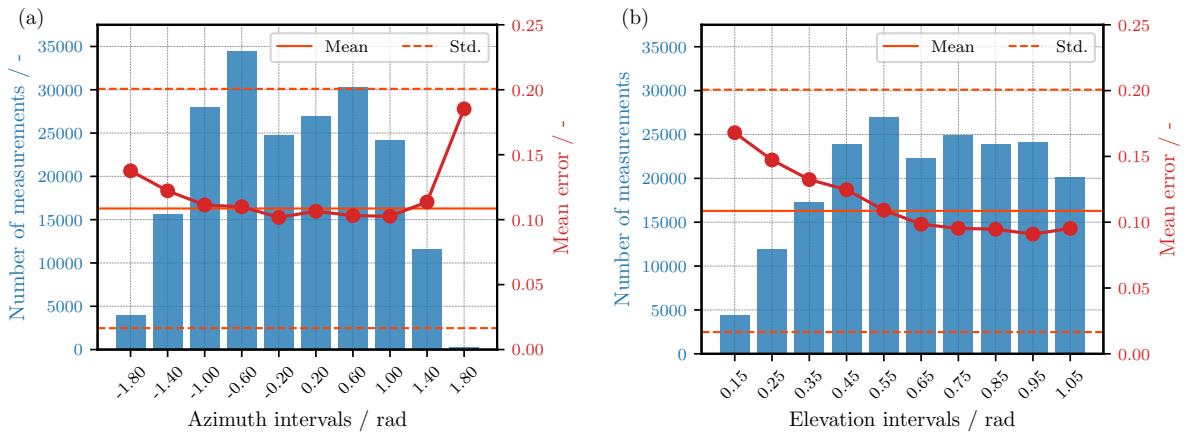

Figure 6.12.: Edge-of-Distribution (EOD) accuracy. (a) Losses by azimuth angle. (b) Losses by elevation angle. Red lines indicate loss values; blue bars show the distribution of calibration points.

- But, loss values for EOD cases remain within the standard deviation of all measurements, indicating controlled degradation of accuracy.

Although reweighting samples based on their frequency could mitigate the effects of an unbalanced data distribution, it is deemed unnecessary, as these edge cases correspond to less operationally relevant periods (e.g., early morning or late evening). The analysis indicates that while the loss increases for EOD cases, the impact is not significant and remains within acceptable limits for practical applications.

6.4.2. Impact of Dataset Size on Model Performance

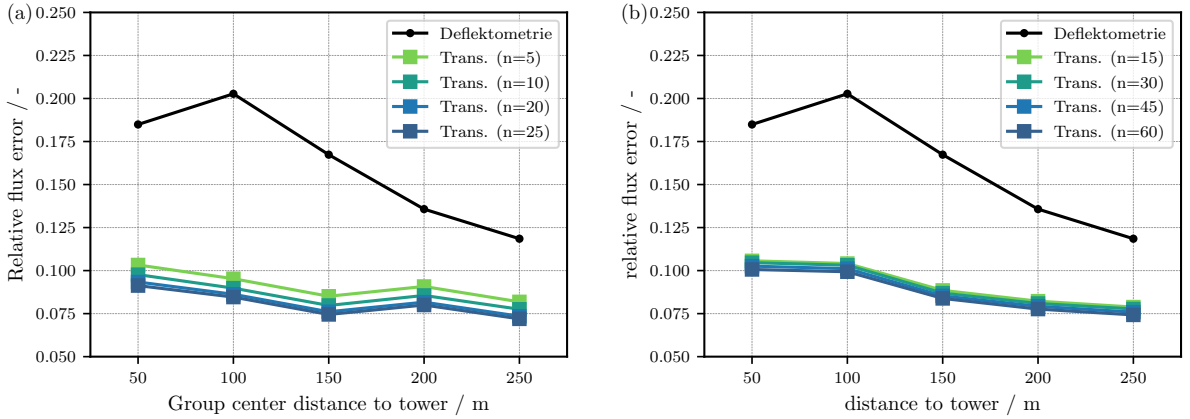

Figure 6.13.: Dataset size dependency, without artifical data. (a): $\mathcal{D}[nT_1 \rightarrow 5T_1]$ - Loss curves for a single target. (b): $\mathcal{D}[nT_{1-3} \rightarrow 5T_{1-3}]$ - Loss curves for all three targets.

Thus far, all evaluations have been conducted using a minimum dataset of 25 focal spot images per heliostat. This section investigates whether the approach remains applicable with fewer images, a critical consideration for ensuring the scalability of the methodology to large heliostat fields, where the number of available data points per heliostat is often limited.

To evaluate the model's performance with limited training data, we train models with varying numbers of focal spot measurements per heliostat. The sampling algorithm (explained in Section 5.1.3) ensures that smaller datasets are strict subsets of the larger datasets, maintaining consistency and comparability. Evaluation is performed using a consistent test set for all cases to ensure reliable comparisons. Figure 6.13 illustrates the results for two scenarios: training on a single target (left) and training across all three targets (right):

- Reducing the dataset size increases prediction losses, but the effect is minimal for $n \geq 5$ measurements per heliostat.
- Loss increases uniformly across distance groups, with no significant dependency on heliostat distance.

The Transformer model demonstrates remarkable robustness, maintaining high accuracy even with minimal training samples per heliostat. While a slight decrease in accuracy is observed as the number of available images decreases, the performance remains significantly superior to deflectometry-based methods, even with as few as 5 images per heliostat. Notably, these 5 images are readily available through standard heliostat calibration procedures, ensuring that the data-driven flux prediction approach can operate effectively without requiring additional measurements.

6.4.3. Generalizing Capabilities

The ability to generalize is a fundamental requirement for developing a robust flux prediction model for large-scale heliostat fields. In this context, **generalization** refers to the model's capacity to exchange information between heliostats, leveraging the combined dataset to improve the prediction accuracy for all heliostats, including those with limited training data or previously unseen conditions. By integrating data from the entire field into a single model, the Transformer architecture captures global patterns and dependencies that enhance its predictive performance.

Observed Generalization in Aim Point Extrapolation The generalizing capabilities of the model have already been demonstrated in the context of aim point extrapolation. As detailed in Section 6.3.2, the inclusion of artificial datasets, designed to introduce unseen conditions, improved the model's accuracy for real measured data. This improvement was possible because the model leveraged the shared patterns in both artificial and measured data, effectively transferring knowledge from simulated conditions to real-world heliostats. This result underscores the model's ability to generalize by combining diverse data sources.

Experiment: Generalization to Unseen Heliostats To further evaluate the model's generalizing capabilities, an experiment was conducted designed to test its performance on previously unseen heliostats. In this extreme scenario, the model is provided with a single focal spot image for a given heliostat and is tasked with predicting focal spots for other sun positions for the same heliostat. Figure 6.14 illustrates the results for this experiment:

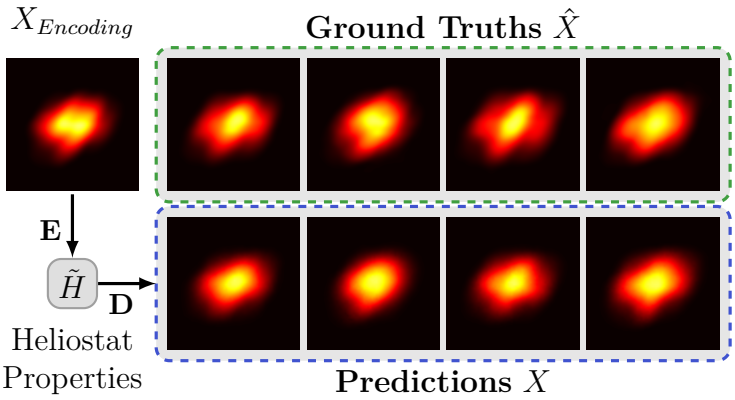

Figure 6.14.: Experiment: Deriving heliostat properties from a single focal spot and testing prediction accuracy under varying conditions. [63]

- Even with only one focal spot image for deriving heliostat properties, the model successfully predicts focal spots for other sun positions. This demonstrates that the model has learned to generalize heliostat-specific properties from minimal data.
- The model's ability to predict for unseen heliostats is a direct result of its global training approach, which integrates data from the entire heliostat field. By leveraging shared patterns and dependencies, the Transformer can infer missing information for new heliostats.
- This experiment highlights the model's scalability and adaptability, ensuring that it remains effective even in scenarios where calibration data is sparse or incomplete.

The Transformer model's ability to generalize across the heliostat field enhances prediction accuracy by leveraging shared information, enabling robust extrapolation to unseen conditions with limited data. This scalability reduces dependency on extensive datasets and underscores its practicality for large-scale heliostat field deployment.

6.5. Summary and Conclusion of the Evaluation of Data-Driven Flux Prediction

This chapter systematically demonstrated the effectiveness and innovative nature of the proposed data-driven methodology for heliostat flux prediction, addressing the research gaps and objectives outlined in the introduction. The key conclusions are summarized as follows:

1. **Model Accuracy and Total Error Analysis:**
   - The Transformer-based data-driven pipeline achieved superior accuracy in predicting spatially resolved flux distributions, outperforming deflectometry—the most accurate state-of-the-art method.
   - Comprehensive evaluation of the pipeline components (UNet, Transformer, and Projection) and total error propagation analysis demonstrated the accuracy of the proposed methodology, even in complex scenarios involving aimpoint extrapolation and receiver projection, with significantly lower errors than traditional methods.

2. **Robustness, Generalization, and Scalability:**
    - The Transformer model demonstrated resilience under diverse operating conditions, including edge-of-distribution (EOD) scenarios with rare sun positions, exhibiting only minor accuracy degradation.
    - The model consistently maintained strong performance with minimal training data, requiring as few as five focal spot images per heliostat to deliver accurate predictions. This scalability highlights its practical applicability to large heliostat fields.
    - The methodology's generalization capabilities enable accurate predictions for previously unseen heliostats and sun positions by leveraging shared patterns across the heliostat field. This reduces the need for extensive individual calibration and enhances predictive accuracy through field-wide synergy.

The evaluation demonstrated the superior accuracy, robustness, and scalability of the proposed data-driven methodology for heliostat flux prediction. The Transformer-based pipeline outperformed state-of-the-art methods, maintained resilience under diverse conditions, and proved its practicality for large-scale heliostat fields through effective generalization and minimal data requirements.

Chapter 7.

System-Level Impact Analysis

Building on the demonstrated accuracy of the data-driven pipeline in predicting individual heliostat focal spots, this chapter extends the evaluation to analyze its implications at the system level for Concentrated Solar Thermal (CST) plant operations. The focus shifts from single heliostat evaluations to the aggregated impact of focal spot prediction accuracy on the superposed total flux density across the receiver, accounting for the aiming inaccuracies of heliostats.

While accurate flux distribution predictions for individual heliostats are critical for assessing model performance, the operational efficiency of a CST plant ultimately depends on the total flux distribution generated by the entire heliostat field. This chapter investigates how prediction accuracies for individual focal spots propagate to system-level metrics, such as flux uniformity and peak flux constraints, which are key factors in receiver performance and durability.

Furthermore, the chapter examines the total absolute flux error by incorporating additional sources of uncertainty, such as Direct Normal Irradiance (DNI) and heliostat reflectivity. This comprehensive error analysis highlights the precision of the Transformer-based method relative to traditional approaches. Finally, the broader system impact of enhanced flux prediction accuracy is quantified by evaluating its influence on the receiver utilization and its direct contribution to improved receiver power output and overall CST plant efficiency.

7.1. Evaluation Framework for Total Flux Analysis

To evaluate the total flux prediction accuracy in a controlled and reproducible manner, a hypothetical receiver geometry was positioned directly behind calibration target T_3. The ground truth flux distribution was constructed by projecting measured focal spots from target T_3 onto this receiver geometry using the projection method described in Section 6.3.3, which has been shown to introduce only minimal error. This approach provides several key advantages:

- **Defined Ground Truth:** Measured focal spots ensure precise and well-defined flux distributions for comparison.
- **Controlled Parameters:** Tracking errors, aimpoint deviations, and other parameters can be systematically varied, allowing for detailed sensitivity analyses.
- **Avoidance of Measurement Uncertainties:** Total flux measurements on real receivers often exhibit inaccuracies of 10–40% and are currently unavailable at the Solar Tower Jülich (STJ). This controlled scenario eliminates these limitations while providing insights into the accuracy of the prediction methodology.

A detailed description of this total flux simulation approach is provided in Appendix B.1. This setup ensures that the total flux evaluation is primarily influenced by focal spot prediction errors. Additionally, the DNI and heliostat reflectivity are assumed to be error-free in the first part of this analysis, as they affect all prediction methods equally. This allows the evaluation to focus solely on the accuracy of the prediction methods. A comprehensive assessment of the absolute error, incorporating DNI and reflectivity uncertainties, is presented in Section 7.4.

Prediction Methods Compared

Four focal spot prediction methods were evaluated to analyze total flux accuracy and system performance, including one optimal baseline model and three comparative approaches applicable for different receiver geometries:

- **Optimal Baseline:** This method uses measured focal spots directly, representing accurate ground truths for the focal spots. While it ensures error-free focal spots, superposed flux deviations arise solely from tracking inaccuracies in individual heliostats. In scenarios with perfect tracking (0 mrad), this baseline exhibits no error. It serves as a benchmark to evaluate how tracking errors alone influence the prediction performance when focal spots are error-free.
- **Transformer-Based Method:** The proposed data-driven approach leverages a Transformer model trained on focal spots from targets T_1 and T_2 combined with CNN-based projection technique.
- **Convolutional Method:** This simpler benchmark employs a Gaussian-based model, fitting a single parameter σ per heliostat to measured focal spot distributions. The Convolutional Method is primarily used as a benchmark for straightforward receiver geometries, such as the Open Volumetric Receiver (OVR).
- **Ideal Heliostat (Raytracing):** This raytracing-based benchmark assumes perfect heliostat geometry and neglects mirror errors and tracking inaccuracies. Raytracing is specifically applied to more complex receiver designs, such as cavity receivers, where the Convolutional Method is not applicable. The raytracing approach provides a comparative baseline under idealized assumptions for such geometries.

It is important to note that deflectometry-based raytracing was excluded from this evaluation due to insufficient characterization data for a system-wide analysis. Detailed descriptions of the Convolutional Method and Ideal Heliostat Method are provided in Appendix A.1, along with their specific applications and assumptions.

Incorporating Tracking Error

Tracking accuracy plays a critical role in total flux prediction accuracy, as aiming deviations can shift focal spots unpredictable on the receiver surface. To assess this impact, tracking errors were simulated using tracking error distributions representative for the heliostat field at the STJ. Three cases were analyzed:

- **Current Tracking Accuracy:** Deviations with a mean tracking error of **2 mrad**, representing the current field accuracy at the STJ.

- **Improved Tracking Error:** Deviations with a mean tracking error of **1 mrad**, achieveable by recent enhancements in tracking models [72].
- **No Tracking Error:** Ideal alignment (**0 mrad**), optimal case when there were no tracking deviations.

Details on the tracking error simulation methodology are provided in Appendix B.2. By incorporating tracking inaccuracies, the evaluation provides a comprehensive understanding of combined errors from focal spot prediction and alignment errors.

Aimpoint Optimization

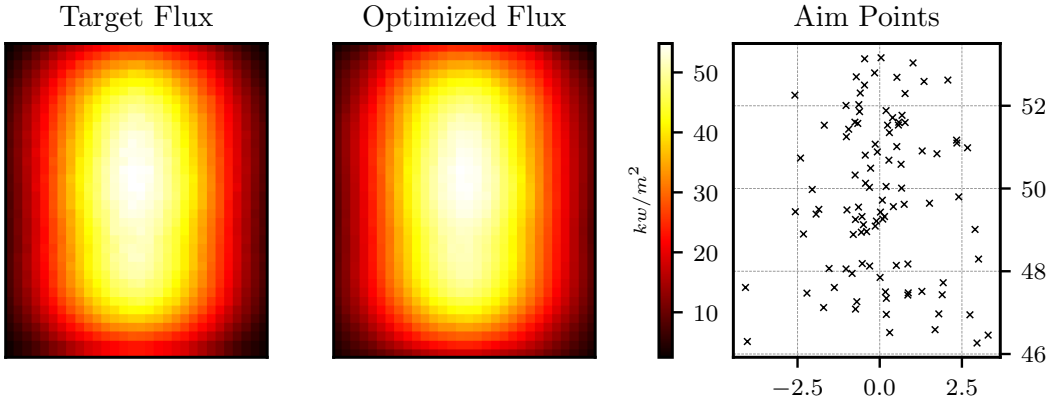

Figure 7.1.: The Target Flux represents the defined flux distribution for optimization. The Optimized Flux is the resulting flux distribution achieved by optimizing aimpoints using the optimal baseline focal spots. The Aim Points plot shows the optimized aimpoint positions for all heliostats ($n = 100$) in the xz-plane (in meters).

To evaluate the total flux density distribution, aim points for the individual heliostats must be defined. To achieve this, a target flux distribution is specified, and the aim points are optimized to align the predicted flux with the desired target flux. A differentiable aim point optimizer was developed and implemented for this purpose (details regarding the optimizer are provided in Appendix B.3). Figure 7.1 illustrates the defined target flux, the optimized flux distribution (based on the optimal baseline focal spots), and the corresponding aim points. The results demonstrate that the optimizer achieves a high degree of agreement between the target flux and the predicted flux with optimized aim points.

Evaluation Metrics

The aim points used in the analysis are subject to uncertainty, as the actual aim points may differ due to tracking inaccuracies. To address this, a Monte Carlo approach is employed, where 1,000 variations are simulated for each case by introducing random aim point errors drawn from the tracking error distribution. The expected ground truth flux distribution is computed as the mean of all these variations.

To compare a prediction method with the ground truth, two evaluation metrics are defined:

- **Mean Deviation:** This metric calculates the pixelwise mean deviation between the expected flux distributions of the prediction method and the ground truth. It provides an overall measure of the prediction method's accuracy relative to the expected ground truth.
- **Max Deviation:** This metric determines the maximum pixelwise deviation by comparing all variations of the ground truth flux to the expected value of the prediction method. It identifies whether extreme deviations can occur for certain aim point variations, offering insights into the robustness of the prediction method under tracking uncertainties.

This approach ensures a comprehensive evaluation of both the accuracy and robustness of the prediction method when subjected to realistic tracking inaccuracies.

7.2. Flux Prediction Accuracy

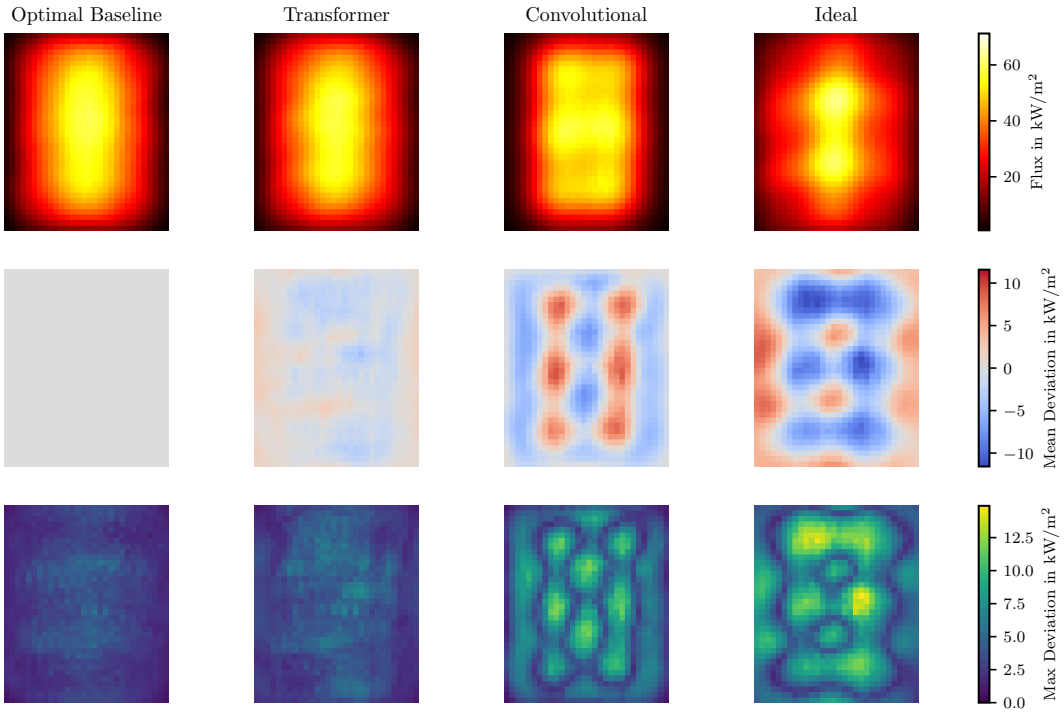

Figure 7.2.: First Scenario (Case 1). Total flux predictions for different methods (columns) with $n = 100$ heliostats and a tracking error of 1 mrad. The first row shows the flux density distributions, the second row presents the mean deviation metric, and the third row displays the max deviation metric. Results for $n = 1000$ can be found in Appendix C.

The first scenario considers an aimpoint distribution optimized for the ground truth flux (Optimal Baseline, as illustraded in Figure 7.1), ensuring ideal alignment with the desired target flux on the receiver surface. The same aimpoints were then used to compute the total flux for each prediction method. Figure 7.2 presents the resulting total flux distributions and their deviations:

- **Optimal Baseline:** The mean deviation is zero, as the expected flux values for both the prediction and the ground truth are identical. However, even with error-free focal spots, slight deviations occur in some variations due to tracking inaccuracies, as observed in the *Max Deviation* plot.

- **Transformer Method:** The data-driven pipeline achieved small deviations, with slight overestimations in the central receiver area and underestimations at the edges. This indicates a high level of accuracy in focal spot prediction, with minor systematic errors.

- **Convolutional Method:** Significant deviations were observed, characterized by overestimated flux in the central regions and underestimations at the periphery. This pattern reflects the Gaussian assumption's tendency to overpredict peak intensities, leading to inaccuracies in total flux predictions.

- **Ideal Heliostat:** Deviations were also substantial but exhibited an inverted error pattern compared to the Convolutional Method. The focal spots were predicted to be too small, resulting in underestimated central flux and overestimated peripheral flux.

- **General Observation:** Across all cases, there is a strong agreement between the flux values in the *Mean Deviation* and *Max Deviation* metrics. The deviation patterns for both metrics overlap significantly, with the *Max Deviation* approximately 25% higher than the *Mean Deviation*.

For the Convolutional Method, and particularly for the Ideal Heliostat Method, systematic errors in focal spot predictions (e.g., inaccuracies in size or shape) compound, resulting in large deviations in the total flux. These findings highlight the importance of accurate focal spot modeling to minimize system-level errors.

7.2.1. Influence of Heliostat Field Size

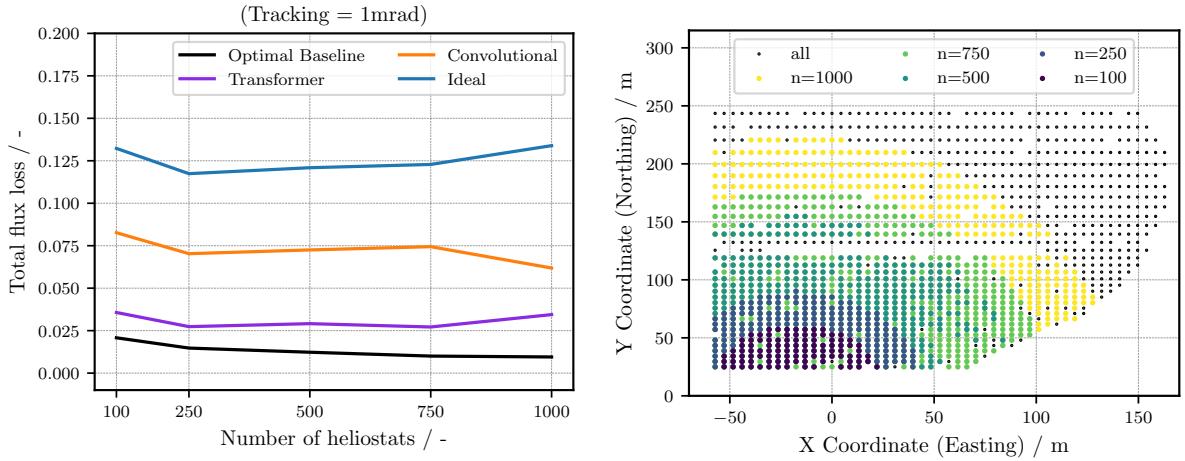

Figure 7.3.: Left: Loss curves for each prediction method across varying numbers of heliostats. Right: Positions of the heliostats selected for each case.

To analyze the influence of field size, or the number of heliostats, the field size was varied from 100 to 1,000 heliostats to evaluate its effect on total flux prediction accuracy. Figure 7.3 presents the results as loss curves:

- **Optimal Baseline:** For the optimal baseline, increasing n resulted in lower errors in total flux predictions. With optimal focal spot predictions, this is attributed to the diminishing impact of tracking inaccuracies as the overlap of focal spots increases with more heliostats.
- **Transformer Method:** The Transformer method demonstrated consistently low total flux errors across all field sizes. Minor increases were observed at the smallest field size (100 heliostats) and the largest field size (1,000 heliostats). Notably, this method achieved the highest accuracy among all prediction approaches. The slight increase in error at $n = 1,000$ is due to the inclusion of more heliostats from the far eastern part of the field, which represent edge cases for the data-driven approach.
- **Convolutional Method:** The Convolutional method exhibited decreasing total flux errors with increasing field size. This trend correlates with improved focal spot prediction accuracy for distant heliostats and the statistical benefits of the convolutional assumptions, where the superposition of focal spots aligns with the Central Limit Theorem.
- **Ideal Method:** The ideal method assumptions performed the worst across all field sizes, with significant errors observed irrespective of the field size.

The analysis shows that total flux prediction accuracy remains mostly consistent across field sizes, with the Transformer method maintaining the highest performance.

7.2.2. Impact of Tracking Error

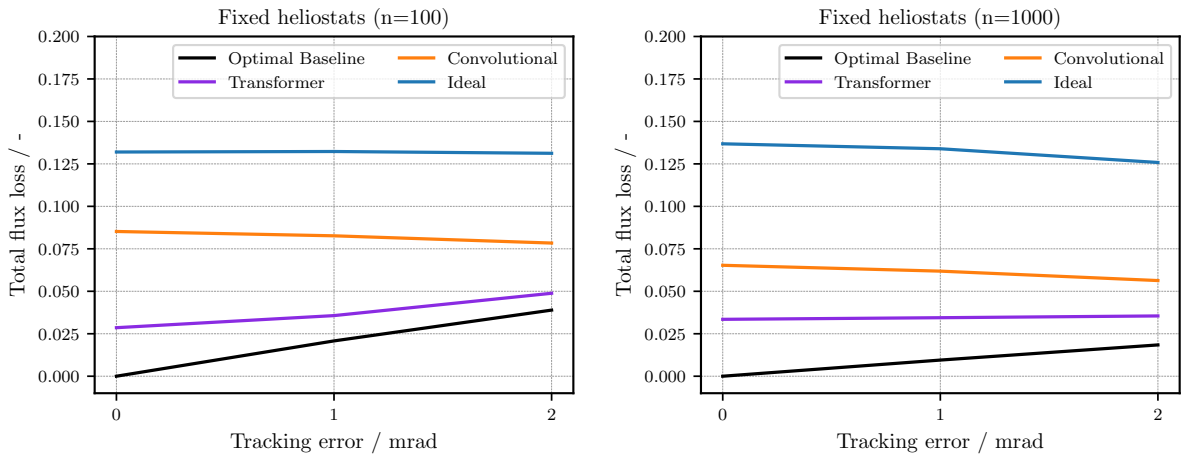

Figure 7.4.: Impact of tracking errors on total flux prediction accuracy. Left: Results for $n = 100$ heliostats. Right: Results for $n = 1,000$ heliostats.

The impact of tracking errors was analyzed for two different field sizes (100 and 1,000 heliostats), with tracking errors ranging from 0 to 2 mrad, as described earlier. Figure 7.4 illustrates the relationship between tracking error magnitude and total flux prediction accuracy:

- **Optimal Baseline:** For perfectly predicted focal spots, the total flux error decreases linearly with decreasing tracking error, reaching zero in the case of perfect tracking.
- **Transformer Method:** The Transformer method consistently outperformed all others, achieving total flux deviations close to the optimal baseline in all scenarios.
- **Convolutional and Ideal Methods:** For the Convolutional and Ideal methods, higher tracking accuracy sometimes resulted in increased prediction errors. This is due to the dominance of systematic focal spot prediction errors when tracking inaccuracies are minimized. Notably, the Convolutional method performed better in scenarios with less accurate heliostat field aiming, as higher tracking errors caused the expected focal spots in the ground truth to widen and become more Gaussian-like, aligning well with the Convolutional method's assumptions.
- **Field Size Sensitivity:** Smaller heliostat fields (100 heliostats) exhibited greater sensitivity to tracking errors, while larger fields (1,000 heliostats) mitigated the impact of tracking inaccuracies through averaging effects.

It is important to note that this analysis focuses solely on the influence of tracking errors on flux prediction accuracy at the receiver. Other advantages of improved tracking accuracy, such as spillage reduction, are not considered here.

7.3. Impact on Aimpoint Optimization

In the first scenario (7.2), the ground truth focal spots were used to optimize the aimpoints to match the target flux, and the same aimpoints were applied across all methods for comparison. In this second scenario, the focal spots of each prediction method are used to individually optimize its aimpoint configuration. This optimized configuration is then applied to compute the underlying ground truth flux using the exact focal spots. Figure 7.6 illustrates the resulting total flux distributions for this second scenario. The results differ from the first scenario due to the distinct error propagation patterns resulting from the new aimpoint distributions. For both the Optimal Baseline and the Transformer Method, the error patterns remain similar. For the Ideal Heliostat Method, however, errors remain significantly high. In contrast, the Convolutional Method exhibits reduced errors compared to the first scenario. The quantitative analysis, along with the spread between Case 1 and Case 2, is presented in Figure 7.5. These results highlight several key observations:

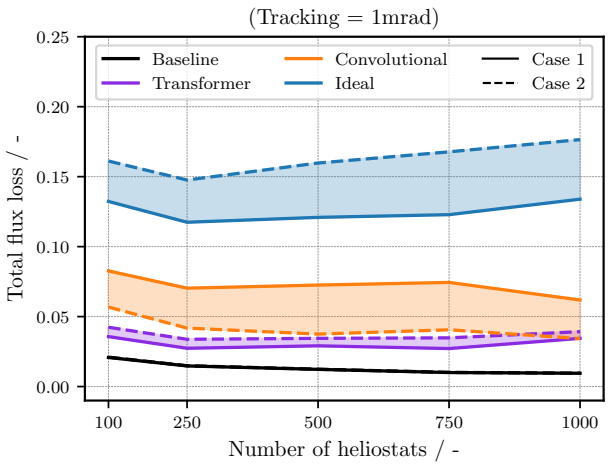

Figure 7.5.: Quantitative analysis of total flux errors for the second scenario, including error spreads between Scenario 1 and Scenario 2.

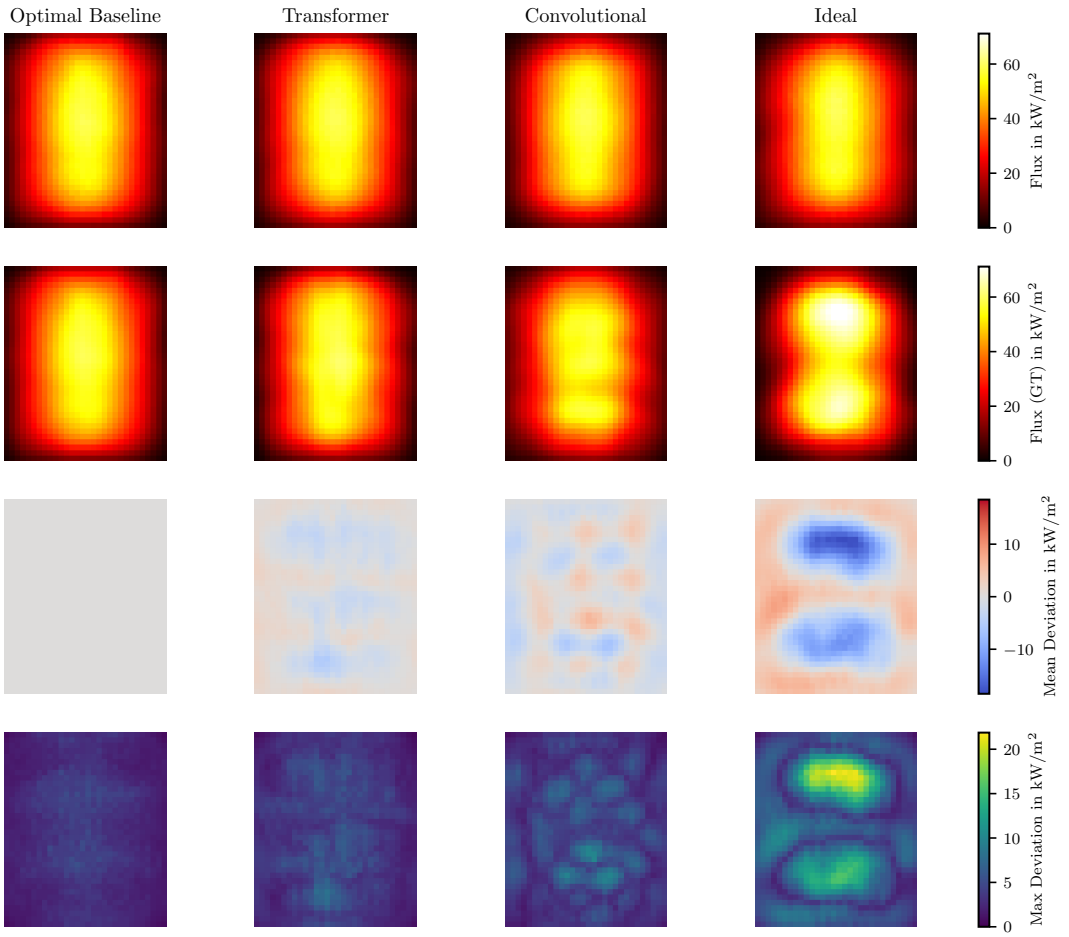

Figure 7.6.: Second scenario (Case 2) with $n = 100$ heliostats and a tracking error of 1 mrad. First row: Optimized expected total flux predictions for each method. Second row: Ground truth flux generated using the optimized aimpoints from optimized flux (first row) simulated with ground truth focal spots. Third and fourth rows: Deviation metrics. Results for $n = 1000$ can be found in Appendix C.

- **Dependency on Aimpoint Distribution:** The results demonstrate a clear dependency on the aimpoint distribution, indicating that the configuration of the aimpoints directly affects the accuracy of the total flux predictions for each method.
- **Performance of the Transformer Method:** Among all methods and across all field sizes, the Transformer Method consistently exhibits the smallest error spread, confirming its robustness and superior performance compared to other approaches.

A comprehensive analysis involving multiple aimpoint configurations, sun positions, and cloud cover scenarios would provide deeper insights. However, such an investigation is beyond the scope of this work. As a result, the presented findings should be interpreted as indicative trends and expected ranges of total flux errors for each method.

7.4. Estimation of Total Absolute Error

To estimate the absolute total flux error of the flux density distribution, two additional sources of uncertainty must be considered: the Direct Normal Irradiance (DNI) of the

sun and the reflectivity of the heliostats. Both quantities can reasonably be assumed to behave independently. While reflectivity could potentially interact with the total flux prediction—since individual heliostat reflectivity variations could influence the relative flux distribution—this interaction is neglected for this initial estimation. Consequently, Gaussian error propagation is applied to calculate the total error as:

$$E_{\text{tot}} = \sqrt{E_T^2 + E_{\text{DNI}}^2 + E_\rho^2}, \tag{7.1}$$

where $E_T = 4.5\%$ represents the mean error associated with the variations in total flux prediction using the Transformer-based method (Figure 7.5), $E_{\text{DNI}} = 1\%$ corresponds to the uncertainty in the DNI, and $E_\rho = 2\%$ accounts for the uncertainty in heliostat reflectivity. Substituting these values, the total error is calculated as:

$$E_{\text{tot}} = \sqrt{(4.5\%)^2 + (1\%)^2 + (2\%)^2} \approx 5\%. \tag{7.2}$$

This result quantifies the combined total absolute error in the flux density distribution, incorporating uncertainties from the flux prediction method, DNI, and heliostat reflectivity. The estimated error demonstrates the precision of the proposed flux prediction methodology, even when accounting for multiple sources of uncertainty.

7.5. Impact on Power Plant Efficiency

The system-level impact of the proposed data-driven flux prediction method can be evaluated by analyzing its influence on the operational efficiency of the receiver. From a material perspective, the receiver's performance is constrained by the Allowable Flux Density (AFD), which represents the maximum flux the components can absorb without sustaining damage. An aim point distribution must account for uncertainties in flux prediction to ensure that the AFD limit is respected, even under extreme conditions. More accurate flux predictions enable a more efficient utilization of the receiver's maximum allowable flux, directly influencing the power output and overall efficiency of the plant. This section examines how prediction accuracy contributes to improved receiver performance and overall system efficiency.

Receiver Utilization of Flux Prediction Methods

The receiver utilization (η_{AFD}) quantifies how effectively a flux prediction method can utilize the AFD while accounting for uncertainties and avoiding exceedance of the maximum allowable flux. Incorporating the maximum flux deviation ($\Delta\Phi_{\max}$) for each method, η_{AFD} is calculated as:

$$\eta_{\text{AFD}} = \frac{\text{AFD} - \Delta\Phi_{\max}}{\text{AFD}} = 1 - E_{\text{tot,max}} = 1 - \sqrt{E_{\max}^2 + E_{\text{DNI}}^2 + E_\rho^2}, \tag{7.3}$$

where the AFD represents the maximum flux density that the receiver can tolerate without risking material degradation or operational inefficiency. The maximum flux deviation, $\Delta\Phi_{\max}$, corresponds to the largest local error introduced by the prediction method, thus constraining the usable fraction of the AFD. In this analysis, $E_{\max}$ is calculated for each

method by determining the 99% quantile of maximum deviations across all aim point variations for a tracking accuracy of $\sigma = 1\,\text{mrad}$, averaged over all heliostat field sizes. The maximum error of the flux prediction method, $E_{\max}$, is then combined with uncertainties in DNI and reflectivity to yield the total error $E_{\text{tot,max}}$.

Comparison of Prediction Methods

Using the maximum deviations of all cases presented in Figure 7.5, the receiver utilization for each flux prediction method is calculated. The following Table summarizes the results:

Method	E_{max} (%)	**Receiver Utilization ($\eta_{\mathbf{AFD}}$) (%)**
Transformer Method	5.64	94.36
Convolutional Method	9.30	90.70
Ideal Heliostat	18.83	81.17

Table 7.1.: Comparison of receiver utilization for different flux prediction methods based on their maximum deviations.

The Transformer-based approach achieves significantly higher receiver utilization compared to state-of-the-art methods, such as the Convolutional Model and the Ideal Heliostat Method. This improvement underscores the accuracy of the data-driven pipeline, which allows high receiver utilization.

Impact on Receiver Power and Plant Efficiency

When all other boundary conditions remain constant, the receiver utilization directly correlates with the receiver power output. The efficiency gain ($\Delta\eta$) provided by the Transformer method compared to other approaches is calculated as the relative improvement in η_{AFD}:

$$\Delta\eta_{\text{Transformer vs. Convolutional}} = \frac{\eta_{\text{AFD, Transformer}}}{\eta_{\text{AFD, Convolutional}}} - 1 \approx \frac{94.36}{90.70} - 1 \approx 4.04\%. \tag{7.4}$$

$$\Delta\eta_{\text{Transformer vs. Ideal}} = \frac{\eta_{\text{AFD, Transformer}}}{\eta_{\text{AFD, Ideal}}} - 1 \approx \frac{94.36}{81.17} - 1 \approx 16.25\%, \tag{7.5}$$

This gain represents the potential increase in power output enabled by more accurate flux predictions. For simpler receiver geometries, such as the OVR, where the convolutional method is applicable, the new data-driven pipeline achieves an efficiency gain of 4.04%. For more complex receiver designs, such as cavity reactors where the convolutional method cannot be applied and the ideal raytracing serves as the benchmark, the data-driven approach demonstrates an even higher efficiency gain of 16.25%.

Although the efficiency gain or the increased absorbable power is typically considered during the design phase of the power plant, achieving higher gains during peak hours would also require a larger heliostat field capable of generating this additional power. However, these cross-influences are neglected in this analysis.

— The system-level impact analysis highlights the substantial benefits of the Transformer-based flux prediction method, which achieves significantly higher receiver utilization compared to state-of-the-art methods. This improvement directly translates into increased receiver power and overall plant efficiency.

7.6. Summary and Conclusion of the System-Level Analysis

The system-level analysis highlights the importance of focal spot prediction accuracy in achieving reliable total flux predictions. Key conclusions are as follows:

- **Transformer-Based Pipeline Performance:** he higher accuracy in focal spot predictions directly translates to improved total flux predictions. The Transformer-based pipeline achieves superior accuracy in total flux prediction, consistently outperforming state-of-the-art methods under varying tracking errors and field sizes.
- **Absolute Flux Error Potential:** The estimated absolute flux error of approximately 5% demonstrates the high precision of the flux prediction method. Compared to direct flux measurement approaches, which exhibit uncertainties ranging from 10% to 40% [37], this prediction methodology offers a highly accurate and viable alternative. This is particularly advantageous in scenarios, such as reactors, where direct flux measurements are not feasible.
- **Receiver Utilization and Efficiency Gains:** The Transformer-based method achieves significantly higher receiver utilization (94.36%) compared to the Convolutional Model (90.70%) and the Ideal Heliostat Method (81.17%). This corresponds to efficiency gains of 4.04% and 16.25%, respectively, highlighting the practical benefits of accurate flux prediction for increasing receiver power output and overall plant efficiency.

The findings highlight the potential of the proposed data-driven methods for practical applications in Concentrated Solar Thermal (CST) systems, offering a reliable and accurate approach to flux prediction and optimization. The comparison of the two cases revealed that the aimpoint configuration significantly influences prediction accuracy. Future research should explore the impact of a wider range of aimpoint configurations, sun positions, and cloud passages to comprehensively quantify the variability in total flux error.

Chapter 8.

Conclusion and Outlook

At the outset of this thesis, the central question was whether purely data-driven methodologies could overcome the inherent information loss associated with traditional heliostat beam characterization techniques. Specifically, the objective was to develop a unified framework, trained directly on focal spot data, capable of achieving flux prediction accuracies comparable to, or exceeding, state-of-the-art methods such as deflectometry-based raytracing, while eliminating the need for costly and time-intensive measurements.

To address this challenge, the concept of heliostat modeling was approached in a completely new way. Traditional approaches rely on detailed surface deformation measurements to characterize each heliostat—deformations that are challenging to reconstruct accurately from focal spot data. In contrast, this work introduces an abstraction of heliostat properties, focusing on learning how focal spots—essentially cross-sections of the heliostat beam—change under varying conditions. By predicting focal spot properties and their variations directly, the pipeline bypasses the need for detailed physical surface representations of heliostats.

To realize this abstraction, a generalizing heliostat field model was developed. This model aggregates all data from the entire heliostat field into a single predictive framework, leveraging shared patterns and relationships across heliostats. The approach not only captures collective field behavior but also maintains flexibility to adapt to individual heliostat properties. By learning directly from focal spot data, the data-driven pipeline provides a robust and scalable solution to flux prediction. Tailoring advanced machine learning techniques to these specific requirements, the proposed pipeline eliminates the dependency on costly and complex measurements, offering a practical and highly accurate alternative for CST systems.

This innovative approach enables robust and scalable heliostat characterization and flux prediction, independent of heliostat distance or specific properties. By relying exclusively on standard calibration target images, the method is both cost-efficient and scalable to large heliostat fields. The data-driven methodology not only matched but surpassed the accuracy of state-of-the-art techniques such as deflectometry-based raytracing, which rely on extensive and complex measurements. For individual focal spots, the data-driven pipeline achieved a prediction error of 11% when evaluated across the entire heliostat field.

The system-level analysis revealed that the superior accuracy of focal spot predictions also translates into higher prediction accuracies for the total flux density aggregated across all heliostats. When accounting for tracking errors, Direct Normal Irradiance (DNI), and reflectivity uncertainties, the data-driven pipeline achieved a total flux prediction error of approximately 5%. This represents a remarkable level of accuracy, outperforming current total flux measurement methods based on cameras, which exhibit uncertainties ranging from 10–40%. Such precision underscores the potential of the data-driven pipeline as a

superior alternative to direct flux measurements, particularly for applications involving cavity receivers or reactors, where accurate in-cavity flux measurements are inherently challenging.

When translating total flux prediction accuracy into overall system efficiency, the data-driven pipeline demonstrated efficiency gains of over 4% for simple receiver geometries (compared to the convolutional method) and approximately 16% for complex receiver designs (compared to ideal raytracing). These gains highlight the transformative potential of this approach, paving the way for more efficient and scalable CST plant operations.

Summary of Research Goals and Contributions

This research addressed key limitations of traditional heliostat characterization and flux prediction methods, such as reliance on expensive measurements and complex physical modeling, by introducing the following major contributions:

- **Transformer-Based Predictive Pipeline:** Developed a high-accuracy predictive pipeline leveraging calibration data from standard operation, capable of generating precise flux distributions on receiver-level that surpass the accuracy of state-of-the-art methods.
- **Unified Data-Driven Framework:** Established an integrated framework for heliostat characterization and flux prediction, mitigating the information loss inherent in conventional beam characterization approaches.
- **Generalizing Field Model:** Introduced the first method capable of utilizing data from the entire heliostat field to enhance prediction accuracy and scalability, fostering synergies across individual heliostat datasets.
- **Validation and Applicability:** Demonstrated the robustness and practical applicability of the proposed methodology using real-world data from the Solar Tower Jülich (STJ), showcasing its adaptability to diverse operational conditions.

These contributions highlight the transformative potential of the proposed data-driven methodology, positioning it as a scalable, cost-effective, and highly accurate framework for heliostat flux prediction. This work represents a significant advancement in CST plant optimization and establishes a foundation for further innovation in solar energy technologies.

Overcoming Adoption Barriers and Ensuring Accessibility

While the proposed machine learning approaches represent a significant advancement in CST flux prediction, they introduce methodologies that are still emerging within the broader CST research community. To support adoption and ensure reproducibility, code related to the methodology has been made publicly available in a ready-to-use format [73, 74].

An essential component of this approach is the dataset itself, alongside the careful preprocessing, prefiltering, and feature extraction steps that ensure high-quality input for the machine learning models. To support the community, the dataset, along with detailed preprocessing methods and the full data pipeline, will also be published [70].

This ensures that researchers can replicate and build upon the work presented in this thesis. By addressing these challenges, the work not only advances CST research but also lowers the barriers to entry for applying machine learning techniques in CST systems. The open-source availability of the code and dataset offers significant potential to accelerate the development of advanced methods, fostering innovation and making CST systems more competitive.

8.1. Future Research Directions for Focal Spot Prediction

The total error analysis presented in Chapter 6 identified the generative model as the primary contributor to overall prediction errors. To further improve the accuracy of the transformer-based generative model, additional data sources can be integrated into its training. The unique data-driven architecture of the model allows for the inclusion of diverse data types, facilitating hybrid approaches. For instance, combining traditional heliostat characterization techniques, such as deflectometry or differentiable ray tracing [22], with the generalizing heliostat field model could harness the strengths of both methodologies, potentially leading to even more accurate predictions.

However, as demonstrated in Chapter 7, the performance of the transformer model is already approaching the optimal baseline—where no focal spot errors are present—when predictions are aggregated into the total flux for the entire heliostat field. Therefore, future research should shift focus to leveraging these accurate focal spot predictions for automated aimpoint control. By integrating accurate flux predictions with a receiver model and aimpoint optimization algorithms, further investigations could explore how precise total flux predictions and control strategies enhance power plant efficiency at the receiver level, to develop taylored control strategies.

8.2. Broader Applications of Developed Methods

The methods developed and evaluated in this work primarily focused on data-driven flux prediction. However, their utility extends beyond this specific application and can be adapted for other domains within and beyond Concentrated Solar Thermal (CST) systems.

Preprocessing and Prefiltering Pipeline The importance of preprocessing, prefiltering, and feature extraction for data-driven flux prediction was highlighted in Section 6.2. These techniques are also highly beneficial for other workflows that rely on calibration target images, such as heliostat alignment models. Accurate focal spot extraction and the exclusion of outliers can significantly improve the quality of alignment model training.

At the STJ, the UNet framework has already been integrated into the calibration workflow, enabling accurate focal spot segmentation even under challenging conditions such as strong background illumination or low-intensity focal spots. This represents a significant improvement over the previous algorithms and ensures more robust calibration.

Camera-Based Flux Measurements The UNet framework demonstrated its capability to accurately derive flux distributions from raw calibration targets. This methodology

could be extended to other applications, such as camera-based receiver-level flux measurements. Neural networks could map raw images directly to flux intensity distributions, overcoming the challenges and time-intensive processes of optical receiver calibration.

Additionally, the total flux predictions from the data-driven pipeline—achieving a high accuracy of approximately 5% as shown in Chapter 7—could be used to support or recalibrate camera-based measurement systems. Although predictions inherently have errors, these errors are expected to be distributed around the correct values, particularly for parameters like DNI and reflectivity. Over time, the integration of these predictions as calibration data could improve the accuracy of camera-based systems, which have currently accuracies in the range of 10-40%.

Heliostat Alignment Models The transformer network demonstrated strong potential for developing generalizing heliostat field models for flux prediction. This same approach could be adapted to heliostat alignment calibration, enabling the sharing of error patterns across multiple heliostats and reducing the data requirements for individual calibration.

Incorporating the predicted focal spot shapes from the flux prediction pipeline into the alignment process could also enhance the accuracy of centroid determinations, especially in cases with shading or blocking. Additionally, unifying the alignment and flux prediction models into a single framework could provide mutual benefits, as these processes are currently treated independently but have significant interdependencies.

Applied AI in Engineering This work highlights the transformative potential of machine learning in engineering applications. Tailored neural networks, such as the UNet framework for target image segmentation, the transformer and StyleGAN networks for flux prediction, and the projection network for determining ray directions in the aperture plane, consistently outperformed their physically modeled counterparts. These results underscore the value of machine learning approaches in scenarios where data availability is high, and physical models struggle to align with real-world complexities.

— The methods and insights developed in this thesis not only advance CST research but also demonstrate the broader applicability of machine learning in solving complex engineering problems. By fostering synergies between data-driven and traditional approaches, future research can build on these foundations to drive further innovations in CST systems and beyond.

8.3. Concluding Remarks

This thesis demonstrates that data-driven methodologies can achieve unprecedented levels of accuracy and scalability in predicting heliostat flux distributions, effectively addressing critical challenges in the optimization of Concentrated Solar Thermal (CST) systems. By integrating machine learning techniques with engineering domain knowledge and innovative approaches, this work bridges the gap between traditional physical modeling and modern data-driven paradigms, offering a transformative framework for enhancing CST plant operations.

The evaluation chapter establishes a solid foundation for the practical adoption of these methods, showcasing their superior performance in accuracy, robustness, and scalability when compared to state-of-the-art techniques. Furthermore, the system-level analysis

highlights the broader potential of these approaches, including their integration with aimpoint control to enhance receiver-level efficiency.

This research underscores the value of combining advanced data science methodologies with engineering expertise to address complex challenges in renewable energy systems. The open-source availability of the code and dataset fosters collaboration and innovation, enabling the broader CST research community to build upon this foundation and further refine these methodologies.

Appendix A.

Supplementary Methods for Focal Spot Evaluation

This chapter provides supplementary methods that support the focal spot evaluation in Chapter 6. These include the heliostat models used for benchmarking, an analysis of loss curve stability of the data-driven approach under varying conditions, and perturbation operations designed to simulate realistic errors in focal spot measurements used for total error estimation.

A.1. Heliostat Models for Comparison

This section details the heliostat models used to benchmark the proposed approach. Each model represents a different level of approximation, ranging from an idealized heliostat to a high-fidelity representation incorporating surface errors. These models provide a comparative framework to evaluate heliostat performance against current state-of-the-art methods.

Ideal Heliostat

The ideal heliostat model assumes a perfectly manufactured and aligned heliostat, free from surface imperfections or misalignments. This model is commonly used to simulate heliostat fields when measurement data on surface errors is unavailable and is currently employed at the STJ for flux prediction.

Flux distributions are computed using a raytracer with the following inputs:

- **Geometric Parameters:** The heliostat's dimensions, orientation, and position within the solar field.
- **Ideal Canting Configuration:** Canting is applied based on the heliostat's distance to the receiver, as set during construction.
- **Sun Shape:** The expected sun shape for the STJ location is used.

The ideal heliostat model is the simplest, requiring no additional measurements beyond standard field parameters.

Convolutional Heliostat

The convolutional heliostat model approximates the reflected beam as a Gaussian distribution projected onto the target plane. This method simplifies modeling while preserving key characteristics of the focal spot shape.

The Gaussian beam profile is defined by the total standard deviation σ_{total}, which combines intrinsic heliostat errors with distance-dependent scaling:

$$\sigma_{\text{total}} = \sigma_0 + k \cdot d,$$

where σ_0 is the intrinsic error of the heliostat, k is a distance-dependent scaling factor, and d is the distance from the heliostat to the aimpoint.

The Gaussian distribution undergoes rotation and stretching to account for the heliostat's position and aimpoint. The stretch factor s adjusts for the angle θ between the heliostat's aim direction and the normal to the target plane:

$$s = \frac{1}{\cos(\theta) + \epsilon},$$

where ϵ is a small constant to ensure numerical stability.

The Gaussian is rotated in the XZ-plane by an angle θ_{xz}, calculated as:

$$\theta_{xz} = \tan^{-1}\left(-\frac{x_{\text{aim}}}{z_{\text{aim}}}\right),$$

where x_{aim} and z_{aim} are the aim direction components.

The flux distribution is then given by:

$$f(x', z'') = \frac{1}{2\pi\sigma_{\text{total}}^2} \exp\left(-\frac{x'^2 + z''^2}{2\sigma_{\text{total}}^2}\right),$$

where the transformed coordinates x' and z'' are computed as:

$$\begin{aligned} x' &= \cos(\theta_{xz}) \cdot x - \sin(\theta_{xz}) \cdot z, \\ z' &= \sin(\theta_{xz}) \cdot x + \cos(\theta_{xz}) \cdot z, \\ z'' &= z' \cdot s. \end{aligned}$$

Although convolutional models are typically applied with global estimates for σ_0 and k based on limited photogrammetry or deflectometry data, in this case, σ_0 was fitted individually for each heliostat, while k was fitted for the entire field to match measured focal spots. This adjustment improves comparability with the data-driven approach.

Deflectometry Heliostat

The deflectometry-based heliostat model integrates measured surface error maps for high-fidelity simulation of heliostat performance. These error maps, obtained using deflectometry techniques, capture detailed surface deviations, including slope errors and curvature inconsistencies.

The flux distribution is computed using a raytracer with the following inputs:

- **Geometric Parameters:** The heliostat's dimensions, orientation, and canting configuration.
- **Deflectometry Error Map:** High-resolution surface error maps used to adjust reflected ray directions.
- **Sun Shape:** The expected sun shape for the STJ location is used.

The deflectometry heliostat model is the most accurate representation of heliostat performance. However, the measurement process is time-intensive, and currently, only around 180 deflectometry measurements are available.

A.2. Loss Curve Stability

To ensure that the loss curves of the data-driven pipeline accurately represent field-level prediction accuracy, their stability is assessed under two conditions:

(a) Variability in sampled sun positions for a fixed set of heliostats.

(b) Variability in the selection of heliostats from the field.

Stability Across Sun Positions

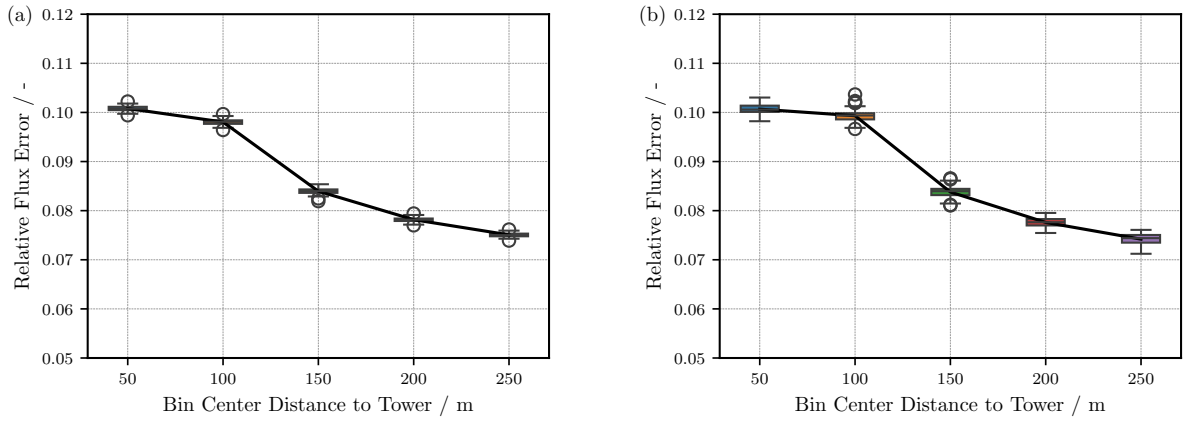

Figure A.1.: Stability of loss curves under different conditions: (a) Variability in sampled sun positions for a fixed set of heliostats. (b) Variability in heliostat subsets.

For a fixed set of heliostats, we assess the impact of randomly sampled sun positions on the loss curves. Using the distance-sampling method described in Chapter 5, five sun positions are randomly selected for each heliostat during evaluation. To test the stability of this approach, the trained model is evaluated 1,000 times with different random sun positions, and the resulting loss curves are compared. Figure A.1(a) shows the mean loss curves and their standard deviations:

- The average standard deviation of loss values across distance groups is **0.05%**, indicating negligible dependency on specific sun positions.
- This demonstrates that evaluating with five randomly sampled sun positions per heliostat is sufficient to capture field-level accuracy.

Stability Across Heliostat Subsets

We also assess whether the loss curves are representative of the entire heliostat field, including heliostats excluded during training or evaluation. From a model trained on 850 heliostats, subsets of 425 heliostats (50%) are randomly selected, and loss curves are calculated while maintaining the distance distribution. This process is repeated 1,000 times, and the resulting loss curves are analyzed. Figure A.1(b) presents the mean values and standard deviations for each distance group:

- The average standard deviation across distance groups is **0.1%**, indicating that the loss curves are representative of the entire heliostat field.

- This confirms that the selected heliostat subsets provide a reliable reflection of the model's performance across the entire field.

A.3. Focal Spot Perturbation Operations for Error Analysis

The perturbation operations described in this section are specifically designed to simulate realistic error conditions observed in focal spot measurements. These perturbations serve as the basis for estimating how individual errors propagate through the pipeline, ultimately contributing to the total system error (E_{total}) as described in Section 6.3.5. By calibrating the parameters of the perturbation operations to match the observed error ranges of actual focal spot predictions, this methodology ensures that the simulated errors accurately reflect real-world conditions.

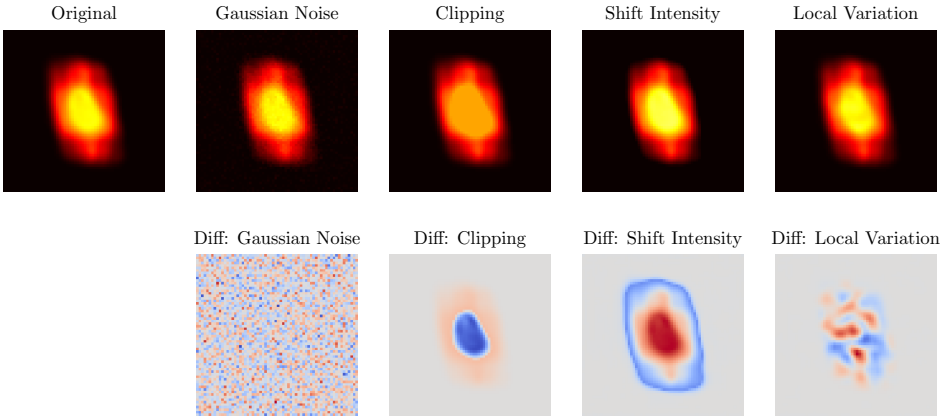

Figure A.2.: Examples of focal spot perturbations. Each column corresponds to a different perturbation operation: Gaussian Noise, Clipping, Shift Intensity, and Local Variation. The first row shows the perturbed image, and the second row displays the difference image (relative to the unperturbed focal spot). Parameters for the perturbations are amplified to highlight their effects.

Perturbation Operations The perturbation operations applied to focal spots include:

- **Gaussian Noise:** Simulates electronic sensor noise by adding random Gaussian fluctuations to the intensity values, reflecting the stochastic noise inherent to imaging systems. This results in a noisy appearance, as shown in the difference images.
- **Clipping:** Models overexposure effects by capping the intensity to a fraction of its peak value, simulating saturation in the sensor's dynamic range. This operation results in a truncated intensity range, visible in the perturbed and difference images.
- **Shift Intensity:** Introduces a random intensity offset and clips values below zero, modeling background subtraction errors of the UNet framework and preprocessing imperfections. This reduces the overall intensity and alters the background, as illustrated in the visualized perturbations.

- **Local Variation:** Applies Gaussian-blurred noise to simulate spatially varying intensity deviations, mimicking errors introduced during UNet flux detection or generative model flux prediction. This results in a distorted and uneven focal spot, as illustrated in the difference images.

Figure A.2 illustrates examples of these perturbations applied to a focal spot image. The first row shows the perturbed focal spot, while the second row highlights the difference between the perturbed and unperturbed images. Parameters for the perturbations are deliberately amplified to visually emphasize their effects, including the spatial distribution and magnitude of the introduced changes.

Appendix B.

Total Flux Simulation and Optimization Framework

This chapter outlines the methodologies used to optimize aimpoints and simulate total flux distributions for the analysis of heliostat field performance in Chapter 7. These approaches form the foundation for evaluating how focal spot prediction errors propagate into receiver-level flux distributions, enabling a robust assessment of total flux accuracy and system-level impacts.

B.1. Total Flux Simulation

To assess how focal spot prediction errors propagate into the total superposed flux density of a heliostat field at the receiver level, a comparative framework was developed. This framework utilizes direct focal spot measurements from target T_3 as the ground truth and compares them to focal spots predicted by various methods, also evaluated on target T_3. Rather than analyzing the total flux on the target area itself, all focal spots are projected onto a hypothetical receiver geometry to enable receiver-level analysis. This hypothetical receiver replicates the geometry of the open volumetric receiver at the STJ but is positioned directly behind the calibration target. The setup accounts for the receiver's complex geometry, including its curvature, which influences the projection and spatial distribution of focal spots. The process follows these steps:

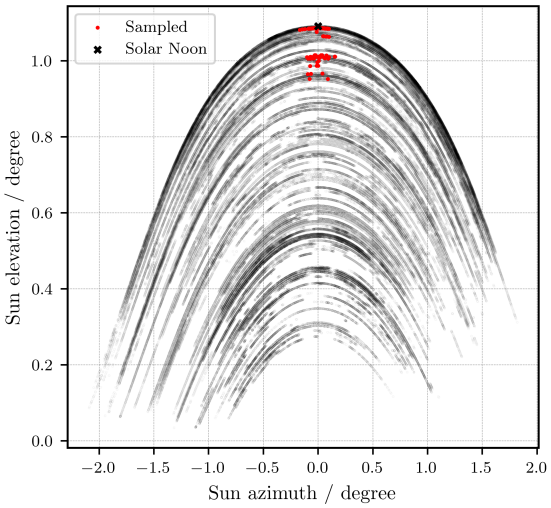

Figure B.1.: All calibration measurements, highlighting the defined sun position (solar noon on June 21st) and the closest calibration points selected for 100 heliostats.

1. Define focal spots in the XZ-aperture plane (the plane corresponding to target T_3) in front of the hypothetical receiver. Focal spots are obtained either through direct measurements or prediction methods.

2. Translate the focal spot to its designated aimpoint in the XZ-plane.

3. Scale the focal spot using the efficiency model described in Section 4.3, assuming a DNI of 1000 $\mathrm{W/m^2}$ and a reflectivity of 0.9.

4. Project the focal spot from the aperture plane onto the receiver geometry, assigning the resulting flux to discrete receiver bins.

For ground truth data, direct measurements from the calibration target are employed. In simulating the heliostat field, the aim is to evaluate all heliostat focal spots under identical sun positions or times. Since calibration data are not available for all heliostats at the exact same sun position, a specific sun position is defined—in this case, solar noon on June 21st. For each heliostat, the calibration measurement with the closest angular distance to this sun position is selected. If no sufficiently close measurement exists, the heliostat is excluded from the analysis. An example of this sampling process is illustrated in Figure B.1. This framework relies on the following assumptions:

- Focal spots do not change significantly within the aimpoint range (approximately $\pm 2\,\mathrm{m}$).
- Focal spots remain stable across the sampled sun positions.
- The projection approach (Section 4.3) introduces negligible additional error.

These assumptions are anticipated to introduce only minor errors into the analysis, comparable to typical flux prediction errors. Consequently, this framework provides a robust and reliable basis for evaluating total flux prediction accuracy, while leveraging known focal spot properties of individual heliostats to facilitate an in-depth analysis.

B.2. Tracking Error Simulation

While tracking errors could be neglected during single focal spot analysis—since all focal spots were centered around their center of intensity—for total flux analysis, they must be considered as they directly affect the impact point of each focal spot. Figure B.2(a) shows the distribution of tracking errors for the heliostat field at the STJ, obtained from the calibration procedure. This distribution represents the density of deviations between the aimpoint set by the heliostat controller and the measured impact point. The deviation is quantified as the angular difference between the connecting vector from the heliostat to the aimpoint and the vector to the measured impact point.

The Probability Density Function (PDF) of the Rayleigh distribution is defined as:

$$f(x;\sigma) = \frac{x}{\sigma^2} e^{-\frac{x^2}{2\sigma^2}}, \quad x \geq 0,$$

where σ represents the scale parameter, corresponding to the tracking accuracy of the heliostat field. Three different values of σ (0, 1, and 2) are considered to represent varying levels of field accuracy. For stochastic evaluation, the aimpoint of each heliostat is modified with an error randomly drawn from the Rayleigh distribution, resulting in a distribution of aimpoints that includes tracking errors. For each scenario, 1,000 variations with different randomly drawn errors are generated, and the corresponding total flux is calculated. The expected flux is then determined as the mean value across all variations. To evaluate extreme cases, the ground truth variations are compared to the

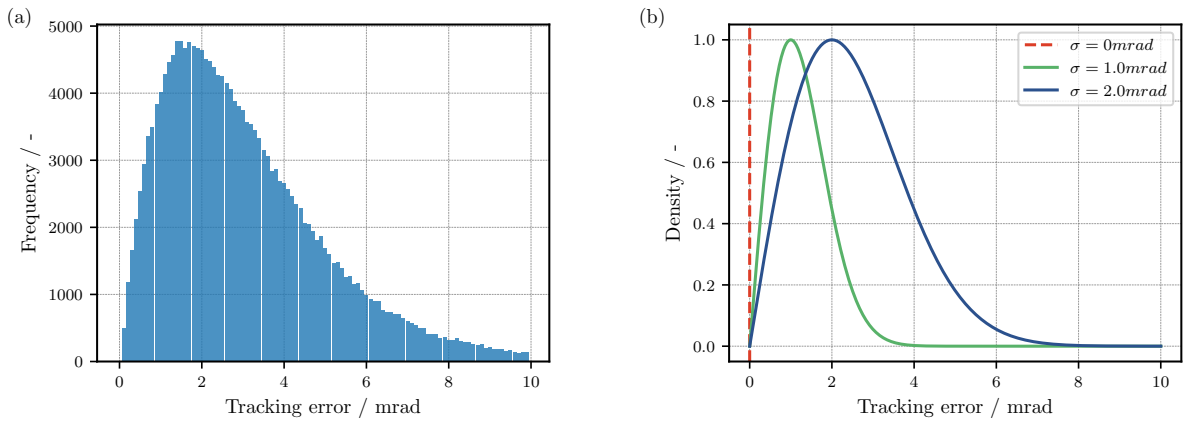

Figure B.2.: (a) Distribution of tracking errors obtained from calibration data at the STJ. (b) Simulated Rayleigh distributions for different tracking accuracies ($\sigma = 0$, $\sigma = 1$, $\sigma = 2$).

mean predicted flux, capturing the potential impact of specific aimpoint configurations on total flux.

For aimpoint optimization, the focal spots are adjusted to account for tracking errors by generating 1,000 samples of each focal spot, each with random tracking errors. The expected focal spot is calculated as the mean of all these variations, and this averaged focal spot is subsequently used for aimpoint optimization.

B.3. Aim Point Optimization Algorithm

The aim point optimization algorithm adjusts the aim points of heliostats to minimize the deviation between the resulting flux distribution on the receiver and a predefined target flux. By leveraging the differentiability of the formulation, the algorithm employs gradient-based optimizers such as Adam or RAdam, significantly improving computational efficiency and convergence speed compared to heuristic methods like genetic algorithms.

Objective

The optimization minimizes the total loss:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{flux}} + \lambda_{\text{AP}} \mathcal{L}_{\text{AP}},$$

where:

- $\mathcal{L}_{\text{flux}}$: Quantifies the deviation between the current flux distribution, F_{current}, and the target flux, F_{target}:

$$\mathcal{L}_{\text{flux}} = \frac{1}{A} \sum_{b \in \text{bins}} \left(\frac{F_b}{\alpha} - F_{\text{target},b} \right)^2,$$

  where F_b is the flux in bin b, α is a normalization factor, and A is the total flux area.

- $\mathcal{L}_{\text{AP}}$: An optional penalty term encouraging aim points ($\vec{Ap}$) to remain close to their initial estimates:

$$\mathcal{L}_{\text{AP}} = \frac{1}{N} \sum_{i=1}^{N} \|\vec{Ap}_i - \vec{Ap}_{\text{init}}\|^2,$$

  where N is the number of heliostats, and $\vec{Ap}_{\text{init}}$ are the initial aim points.

In this implementation, λ_{AP} is set to zero, focusing entirely on minimizing $\mathcal{L}_{\text{flux}}$.

Key Steps in the Optimization Process

1. Initialization

- Initial aim points ($\vec{Ap}_{\text{init}} = [x, z]$) are sampled uniformly within physically valid bounds defined by minimum ($\vec{Ap}_{\text{min}}$) and maximum ($\vec{Ap}_{\text{max}}$) permissible values:

$$\vec{Ap}_{\text{init}} = \vec{Ap}_{\text{min}} + r \cdot (\vec{Ap}_{\text{max}} - \vec{Ap}_{\text{min}}), \quad r \sim \text{Uniform}(0, 1).$$

- The target flux distribution, F_{target}, is loaded from a predefined file.

2. Flux Calculation The current flux distribution (F_{current}) is computed using:

- **Heliostat Focal Spots:** Adjusted focal spots scaled by the heliostats' power output (Q_{tot}).
- **Impact Points:** The intersection of reflected rays with the receiver plane, calculated using the current aim points.

The flux distribution is discretized into bins over the receiver surface:

$$F_{\text{current}} = \text{Projection}(F_{\text{spots}}, \text{positions}, \vec{Ap}),$$

where `calcBins` aggregates contributions from all heliostats.

3. Gradient-Based Optimization Using the differentiability of the flux computation, the RAdam optimizer directly updates the aim points. Key steps include:

- Gradients of $\mathcal{L}_{\text{total}}$ with respect to $\vec{Ap}$ are computed via backpropagation.
- Aim points are updated iteratively:

$$\vec{Ap} \leftarrow \vec{Ap} - \eta \cdot \nabla \mathcal{L}_{\text{total}},$$

  where η is the learning rate.

- A learning rate scheduler dynamically reduces η to improve convergence.

4. Stability Enhancements To ensure numerical stability:

- Gradients are clipped to prevent exploding values:

$$\|\nabla \mathcal{L}_{\text{total}}\|_\infty \leq \text{max\_norm}.$$

- Aim points are clamped within permissible bounds:

$$\vec{Ap}_{x,z} = \text{clip}(\vec{Ap}_{x,z}, \vec{Ap}_{\text{min}}, \vec{Ap}_{\text{max}}).$$

5. Iterative Updates During each iteration:

- The aim points ($\vec{Ap}$) are reconstructed into full 3D coordinates:

$$\vec{Ap} = [x, \text{fixed } y, z],$$

  where y corresponds to the receiver aperture plane.

- The flux loss ($\mathcal{L}_{\text{flux}}$) is recalculated, and aim points are updated to minimize the loss.

Appendix C.

Additional Data

This chapter provides supplementary visualizations of the total flux predictions discussed in the system-level analysis. These figures compare flux distributions and deviation metrics for two scenarios: Case 1 (uniform aimpoints) and Case 2 (optimized aimpoints for each prediction method) for $n = 1000$ heliostats.

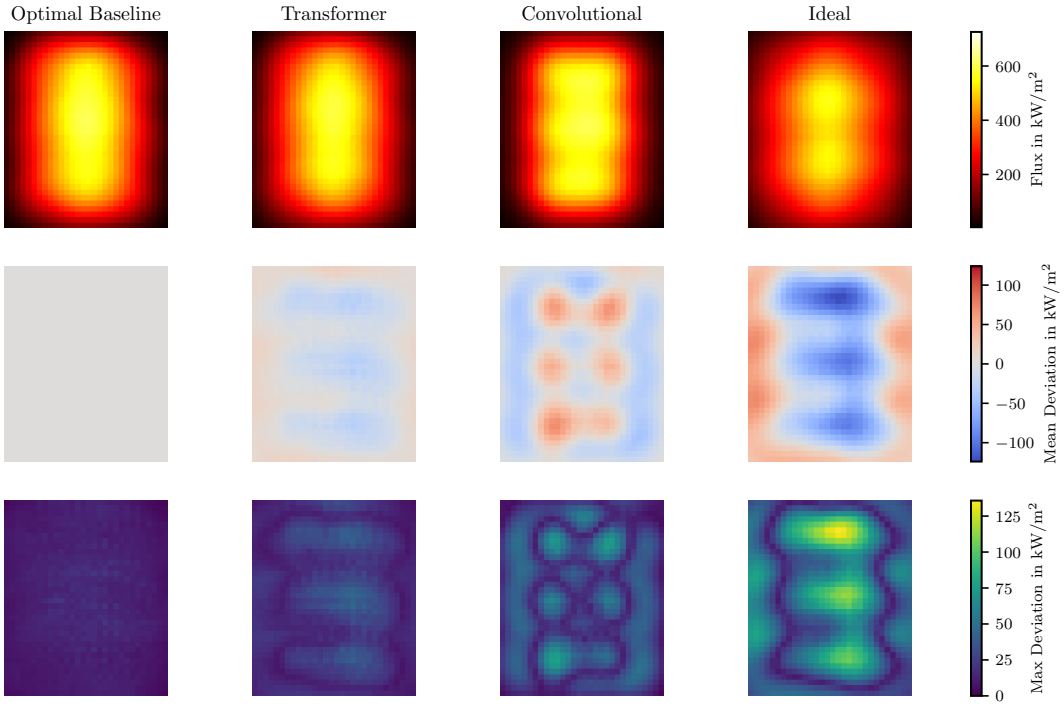

Figure C.1.: First Scenario (Case 1). Total flux predictions for different methods (columns) with $n = 1000$ heliostats and a tracking error of 1 mrad. The first row shows the flux density distributions, the second row presents the mean deviation metric, and the third row displays the max deviation metric.

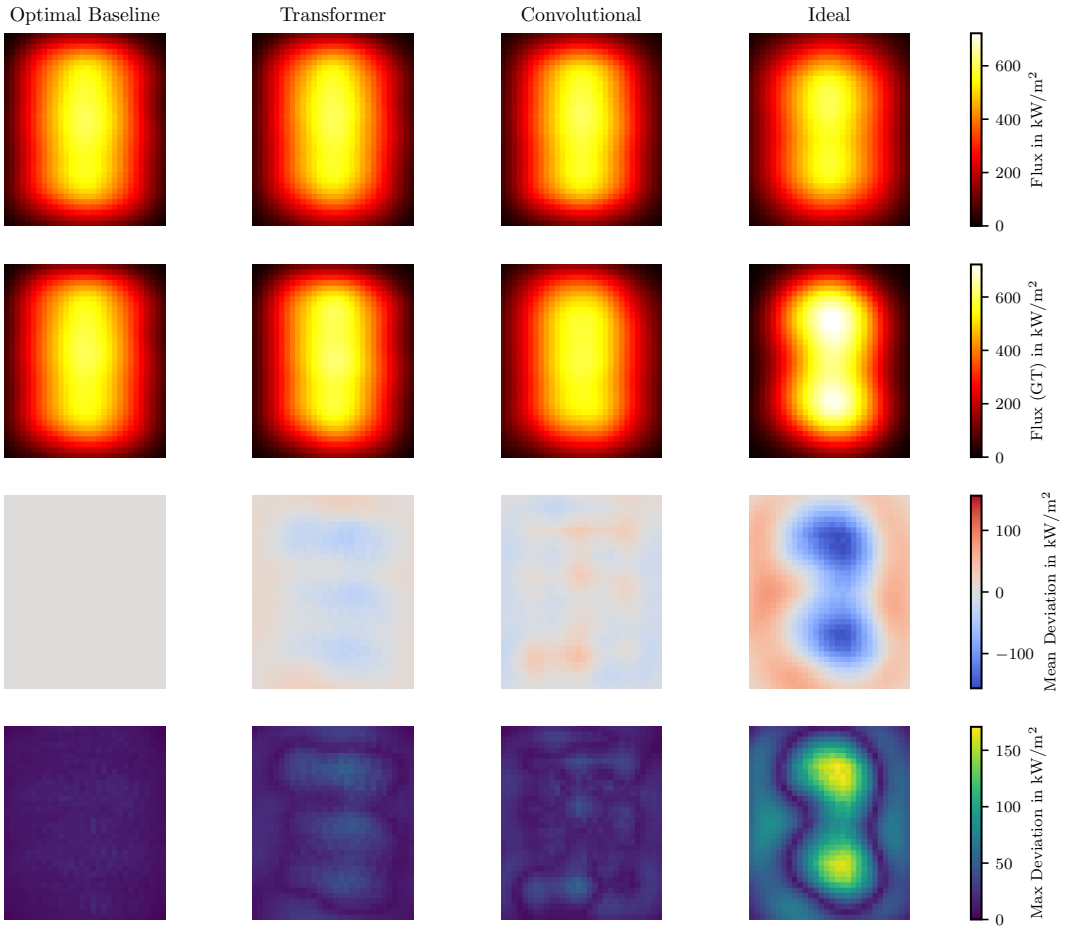

Figure C.2.: Second Scenario (Case 2). Total flux predictions for $n = 1000$ heliostats with a tracking error of 1 mrad. The first row shows optimized expected total flux predictions for each method. The second row displays the ground truth flux calculated using the optimized aimpoints from the first row. The third and fourth rows present the mean deviation and max deviation metrics, respectively.

Bibliography

[1] IPCC. "Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change". In: (2021). URL: https://www.ipcc.ch/report/ar6/wg1/.

[2] International Energy Agency (IEA). *World Energy Outlook 2021*. Paris: OECD/IEA, 2021. URL: https://www.iea.org/reports/world-energy-outlook-2021.

[3] International Renewable Energy Agency (IRENA). "Renewable Power Generation Costs in 2021". In: (2021). URL: https://www.irena.org/publications/2022/Jul/Renewable-Power-Generation-Costs-in-2021.

[4] Peter D Lund et al. "Review of energy system flexibility measures to enable high levels of variable renewable electricity". In: *Renewable and sustainable energy reviews* 45 (2015), pp. 785–807.

[5] International Energy Agency (IEA). *The Future of Hydrogen: Seizing Today's Opportunities*. Report prepared by the IEA for the G20, Japan. Paris: OECD/IEA, 2019. URL: https://www.iea.org/reports/the-future-of-hydrogen.

[6] Clifford K. Ho and Brian D. Iverson. "Review of high-temperature central receiver designs for concentrating solar power". In: *Renewable and Sustainable Energy Reviews* 29 (2014), pp. 835–846. DOI: 10.1016/j.rser.2013.08.099.

[7] Huili Zhang et al. "Thermal energy storage: Recent developments and practical aspects". In: *Progress in Energy and Combustion Science* 53 (2016), pp. 1–40.

[8] Olivier Dumont et al. "Carnot battery technology: A state-of-the-art review". In: *Journal of Energy Storage* 32 (2020), p. 101756.

[9] "Luneng 900MW PV+100MW CSP Hybrid Project Started Construction". In: *China Solar Thermal Alliance* (2022). Accessed: 2024-12-20. URL: http://en.cnste.org/html/csp/2022/0804/1280.html.

[10] Vincent Shaw. "SDIC Gansu commissions 750 MW CSP-PV plant in China". In: *pv magazine International* (2024). Accessed: 2024-12-20. URL: https://www.pv-magazine.com/2024/12/18/sdic-gansu-commissions-750-mw-csp-pv-plant-in-china/.

[11] International Energy Agency (IEA). "World Energy Outlook 2022". In: (2022). URL: https://www.iea.org/reports/world-energy-outlook-2022.

[12] International Energy Agency. "Solar Energy: Tracking Progress". In: *IEA Tracking Clean Energy Progress* (2022). Accessed: 2023-11-20. URL: https://www.iea.org/reports/solar-energy.

[13] International Energy Agency. *Global Energy Review: Heat 2021*. https://www.iea.org/reports/global-energy-review-2021. 2021.

[14] Aldo Steinfeld and Anton Meier. "Solar fuels and materials". In: *Encyclopedia of energy*. Vol. 5. Elsevier Academic Press, 2004, pp. 623–637.

[15] Deepak Yadav and Rangan Banerjee. “A review of solar thermochemical processes”. In: *Renewable and Sustainable Energy Reviews* 54 (2016), pp. 497–532.

[16] Gregory J Kolb et al. “Heliostat cost reduction study”. In: *Sandia National Laboratories, Albuquerque, NM, Report No. SAND2007-3293* 103 (2007).

[17] Yan Luo et al. “Analysis of thermal stress, fatigue life and allowable flux density for the molten salt receiver in solar power tower plants”. In: *International Journal of Low-Carbon Technologies* 17 (2022), pp. 1385–1398.

[18] Adolfo Juan Sánchez del Pozo Fernández. “Model predictive control and optimization of solar thermal energy plants”. In: (2019).

[19] Rudolf Popp et al. “A comparison between model predictive and PID-based control of a molten salt solar tower receiver”. In: *AIP Conference Proceedings*. Vol. 2815. 1. AIP Publishing. 2023.

[20] Yi’an Wang, Zhe Wu, and Dong Ni. “Real-Time Optimization of Heliostat Field Aiming Strategy via an Improved Swarm Intelligence Algorithm”. In: *Applied Sciences* 14.1 (2024), p. 416.

[21] Alberto Sánchez-González, María Reyes Rodríguez-Sánchez, and Domingo Santana. “Aiming factor to flatten the flux distribution on cylindrical receivers”. In: *Energy* 153 (2018), pp. 113–125.

[22] Max Pargmann et al. “Automatic heliostat learning for in situ concentrating solar power plant metrology with differentiable ray tracing”. In: *Nature Communications* 15.1 (2024), p. 6997.

[23] Navid Mohammadzadeh et al. “A Stochastic-MILP dispatch optimization model for Concentrated Solar Thermal under uncertainty”. In: *Sustainable Energy, Grids and Networks* (2024), p. 101458.

[24] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. “Deep learning”. In: *nature* 521.7553 (2015), pp. 436–444.

[25] Ian Goodfellow. *Deep learning*. 2016.

[26] Ashish Vaswani et al. “Attention is all you need”. In: *Advances in neural information processing systems* 30 (2017).

[27] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. “U-Net: Convolutional networks for biomedical image segmentation”. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention* (2015), pp. 234–241.

[28] Tero Karras, Samuli Laine, and Timo Aila. “A Style-Based Generator Architecture for Generative Adversarial Networks”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 4401–4410.

[29] Robin Rombach et al. “High-resolution image synthesis with latent diffusion models”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 10684–10695.

[30] Mehdi Mirza and Simon Osindero. “Conditional generative adversarial nets”. In: *arXiv preprint arXiv:1411.1784* (2014).

[31] Diederik P Kingma and Max Welling. “Auto-encoding variational Bayes”. In: *arXiv preprint arXiv:1312.6114* (2013).

[32] Nitish Srivastava et al. “Dropout: A simple way to prevent neural networks from overfitting”. In: *Journal of Machine Learning Research* 15.1 (2014), pp. 1929–1958.

[33] Gerrit Koll et al. “The Solar Tower Jülich - A Research and Demonstration Plant for Central Receiver Systems”. In: *Proceedings*. Sept. 2009. URL: https://elib.dlr.de/60306/.

[34] R. Buck and E. Teufel. “Comparison and Optimization of Heliostat Canting Methods”. In: *Journal of Solar Energy Engineering* 131.1 (Jan. 2009), p. 011001. DOI: 10.1115/1.3027500.

[35] Camilo A Arancibia-Bulnes et al. “A survey of methods for the evaluation of reflective solar concentrator optics”. In: *Renewable and Sustainable Energy Reviews* 69 (2017), pp. 673–684.

[36] Massaab El Ydrissi, Hicham Ghennioui, Abdi Farid, et al. “A review of optical errors and available applications of deflectometry technique in solar thermal power applications”. In: *Renewable and Sustainable Energy Reviews* 116 (2019), p. 109438.

[37] Marc Röger et al. “Techniques to measure solar flux density distribution on large-scale receivers”. In: *Journal of Solar Energy Engineering* 136.3 (2014), p. 031013.

[38] Matthias Offergeld et al. “Flux density measurement for industrial-scale solar power towers using the reflection off the absorber”. In: *AIP Conference Proceedings*. Vol. 2126. 1. AIP Publishing. 2019.

[39] Sanjoy Chatterjee, Ravikant Ravikant, and V Narayanan. “Estimation of temperature profile of a receiver in solar central tower”. In: *AIP Conference Proceedings*. Vol. 2815. 1. AIP Publishing. 2023.

[40] Bijan Nouri et al. “Nowcasting of DNI maps for the solar field based on voxel carving and individual 3D cloud objects from all sky images”. In: *AIP Conference Proceedings*. Vol. 2033. 1. AIP Publishing. 2018.

[41] Bijan Nouri et al. “Probabilistic solar nowcasting based on all-sky imagers”. In: *Solar Energy* 253 (2023), pp. 285–307.

[42] Laurin Oberkirsch et al. “GPU-based aim point optimization for solar tower power plants”. In: *Solar Energy* 220 (2021), pp. 1089–1098.

[43] Ruidi Zhu et al. “Heliostat field aiming strategy optimization with post-installation calibration”. In: *Applied Thermal Engineering* 202 (2022), p. 117720.

[44] Johannes Christoph Sattler et al. “Review of heliostat calibration and tracking control methods”. In: *Solar Energy* 207 (2020), pp. 110–132. ISSN: 0038-092X. DOI: https://doi.org/10.1016/j.solener.2020.06.030. URL: https://www.sciencedirect.com/science/article/pii/S0038092X20306447.

[45] Peter Schwarzbözl, Robert Pitz-Paal, and Mark Schmitz. “Visual HFLCAL-A software tool for layout and optimisation of heliostat fields”. In: *Proceedings*. 2009.

[46] Caitou He et al. “Analytical radiative flux model via convolution integral and image plane mapping”. In: *Energy* 222 (2021), p. 119937.

[47] Pierre Garcia, Alain Ferriere, and Jean-Jacques Bezian. “Codes for solar flux calculation dedicated to central receiver system applications: A comparative review”. In: *Solar Energy* 82.3 (2008), pp. 189–197.

[48] Nicolas C Cruz et al. “Review of software for optical analyzing and optimizing heliostat fields”. In: *Renewable and Sustainable Energy Reviews* 72 (2017), pp. 1001–1018.

[49] Tim Wendelin. “SolTRACE: a new optical modeling tool for concentrating solar optics”. In: *International solar energy conference*. Vol. 36762. 2003, pp. 253–260.

[50] Nils Ahlbrink et al. “STRAL: Fast ray tracing software with tool coupling capabilities for high-precision simulations of solar thermal power plants”. In: *Proceedings of the SolarPACES 2012 conference*. 2012.

[51] Max Pargmann et al. “Towards a neural network based flux density prediction–Using generative models to enhance CSP raytracing”. In: *AIP Conference Proceedings*. Vol. 2815. 1. AIP Publishing. 2023.

[52] Steffen Ulmer et al. “Automated high resolution measurement of heliostat slope errors”. In: *Solar Energy* 85.4 (2011), pp. 681–687.

[53] Paul M Scott and Greg Burgess. “Measurement of mirror panels using coloured pattern deflectometry”. In: *Proceedings of SolarPACES* (2010).

[54] AM Bonanos et al. “Heliostat surface shape characterization for accurate flux prediction”. In: *Renewable energy* 142 (2019), pp. 30–40.

[55] Marc Röger, Christoph Prahl, and Steffen Ulmer. “Fast determination of heliostat shape and orientation by edge detection and photogrammetry”. In: *Proc. 14th CSP SolarPACES Symposium 2008*. SolarPACES. 2008.

[56] M Kiera and W Schiel. “Measurement and analysis of heliostat images”. In: (1989).

[57] Alberto Sánchez-González et al. “Determination of heliostat canting errors via deterministic optimization”. In: *Solar Energy* 150 (2017), pp. 136–146.

[58] Alejandro Martínez-Hernández et al. “Advanced surface reconstruction method for solar reflective concentrators by flux mapping”. In: *Solar Energy* 266 (2023), p. 112162.

[59] Bruce L Kistler. *A user’s manual for DELSOL3: a computer code for calculating the optical performance and optimal system design for solar thermal central receiver plants*. Tech. rep. Sandia National Lab.(SNL-CA), Livermore, CA (United States), 1986.

[60] Mustafa E Elayeb, Rajab A Haman, and Fuad MF Siala. “Calculation of the blocking factor in heliostat fields”. In: *Energy procedia* 57 (2014), pp. 291–300.

[61] Guillermo Ortega and Antonio Rovira. “A fast and accurate methodology for the calculation of the shading and blocking efficiency in central receiver systems”. In: *Renewable energy* 154 (2020), pp. 58–70.

[62] Manish Raj and Jishnu Bhattacharya. “An accurate and cheaper method of estimating shading and blocking losses in a heliostat field through efficient filtering, removal of double counting and parallel plane assumption”. In: *Solar Energy* 243 (2022), pp. 469–482.

[63] Mathias Kuhl et al. “Flux density distribution forecasting in concentrated solar tower plants: A data-driven approach”. In: *Solar Energy* 282 (2024), p. 112894. ISSN: 0038-092X. DOI: https://doi.org/10.1016/j.solener.2024.112894.

[64] Huimin Huang et al. "Unet 3+: A full-scale connected unet for medical image segmentation". In: *ICASSP 2020-2020 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE. 2020, pp. 1055–1059.

[65] Mathias Kuhl et al. "In-situ UNet-based heliostat beam characterization method for precise flux calculation using the camera-target method". In: *Solar Energy* 279 (2024), p. 112811. ISSN: 0038-092X. DOI: https://doi.org/10.1016/j.solener.2024.112811.

[66] Lukas Liebel and Marco Körner. "Auxiliary tasks in multi-task learning". In: *arXiv preprint arXiv:1805.06334* (2018).

[67] Mathias Kuhl et al. "Accurate and Scalable Receiver-Level Flux Prediction: A Fully Data-Driven Solution". In: *Solar Energy* (2025). Preprint submitted.

[68] Francisco J Collado. "Quick evaluation of the annual heliostat field efficiency". In: *Solar Energy* 82.4 (2008), pp. 379–384.

[69] Athanasios Papoulis. *Probability, Random Variables, and Stochastic Processes*. 3rd. McGraw-Hill, 1991. ISBN: 978-0070484771.

[70] Kaleb Phipps, Mathias Kuhl, Marie Weiel, et al. *PAINT Database*. 2025. URL: https://paint-database.org/.

[71] Boris Belhomme et al. "A new fast ray tracing tool for high-precision simulation of heliostat fields". In: (2009).

[72] Max Pargmann et al. "Enhancing heliostat calibration on low data by fusing robotic rigid body kinematics with neural networks". In: *Solar Energy* 264 (2023), p. 111962.

[73] Mathias Kuhl. *UTIS-HeliostatBeamCharacterization: UNet-Based Target Image Segmentation for Camera Target Method in Solar Tower Plants*. 2024. URL: https://github.com/DLR-SF/UTIS-HeliostatBeamCharacterization.

[74] Mathias Kuhl. *HelioBeamAI-Artificial Intelligence for Heliostat Beam Flux Prediction*. 2025. URL: https://github.com/DLR-SF/HelioBeamAI.