

RESEARCH ARTICLE | OCTOBER 06 2023

Towards a neural network based flux density prediction – Using generative models to enhance CSP raytracing FREE

Max Pargmann ; Daniel Maldonado Quinto; Stefan Kesselheim; Jan Ebert; Robert Pitz-Paal



AIP Conf. Proc. 2815, 030015 (2023)

<https://doi.org/10.1063/5.0148765>



22 April 2024 17:59:22



Lock-in Amplifier



Boxcar Averager



Zurich Instruments

Boost Your Optics and Photonics Measurements

[Find out more](#)

Towards a Neural Network Based Flux Density Prediction – Using Generative Models to Enhance CSP Raytracing

Max Pargmann^{1, a)}, Daniel Maldonado Quinto¹, Stefan Kesselheim², Jan Ebert²,
Robert Pitz-Paál¹

¹*Institute of Solar Research, German Aerospace Center (DLR), Linder Hoehe, 51147 Cologne (Germany)*

²*Jülich Supercomputing Centre, Forschungszentrum Jülich (FZJ), Wilhelm-Johnen Straße, 52428 Jülich (Germany)*

^{a)} Corresponding author: max.pargmann@dlr.de

Abstract. Each solar tower power plant is designed for a pre-calculated optimal flux density distribution on the receiver. Any deviation from it has a direct impact on the component durability, energy efficiency and all downstream processes. An accurate knowledge of the current and predicted flux density is therefore essential. However, the existing measurements to obtain the flux density are either inaccurate, complicated or expensive. Moreover, to predict upcoming flux density maps, simulations are indispensable. Also, because there is still no cost-efficient way to measure heliostat surfaces on an industrial scale, these simulations neglect the individual surface errors of each heliostat and are therefore too inaccurate for a reliable prediction. We present a novel method based on artificial intelligence (AI) to include real heliostat surface errors into the flux density simulation, only by measuring the heliostats focal spot. The method needs only data generated by the already well-established camera-target heliostat calibration. Using the AI method as a supplement to this calibration, within this measurement, it is possible to simultaneously detect both main heliostat errors, the surface deformations and the misalignment. In this work, different neural network (NN) architectures are trained with artificially generated data and studied for their applicability to the described methodology. The network types used are conditioned and unconditioned generative adversarial networks (GANs) as well as neural radiance fields (NeRFs). The latter archives at best a Peak Signal to Noise Ratio (PSNR) of up to 27.8. Afterwards the network results are compared qualitatively in terms of image quality, controllability and dataset size.

INTRODUCTION




FIGURE 1. NeRF prediction for a test dataset (azimuth = 0° - 160° , elevation= 12°) of a single heliostat's focal spot. At most angles the shape as well as the intensity are very close to the targets. Since the training data set only contains elevations up to 10° , all images shown are slightly extrapolated. For the interpolation the results are marginally better. The overall test PSNR is 27.8.

The receiver is the heart of every solar tower power plant. Here the focal spots of all heliostats are concentrated. The surface of the receiver can reach 700°C [1] and the target of current research is even higher [2]. From the distribution of the focal spots and the total incoming power per area, the radiometric quantity, the flux density distribution is obtained. Each solar tower power plant is designed for an optimal (theoretical) flux density distribution. The closer the incoming radiation is to the designed flux density distribution, the higher is the power plant performance. In contrast, a deviation can lead not only to lower performance but also to component damage due to temperature peaks. An accurate knowledge of the current and upcoming flux density is therefore essential for power

plant operators. Since, common methods for measuring the flux density directly are either inaccurate, or expensive[3][4][5], simulations are used instead.

For this, the most common approach is raytracing. Such simulations set up the whole environment including receiver, sun and heliostats (as well as other global parameters, like atmospheric losses etc.). In the standard case each component, especially the heliostats are represented by an ideal geometric representation. However, measured components can be included in the simulation. Then rays emitted by the sun, represented by arrays are generated and reflected until they hit the receiver. In this way an intensity distribution is created.

Even though these simulations are close to reality, they deviate from it [7][8] depending on the errors of each individual heliostat in the field. Especially, without an a priori knowledge from measurements these simulations are not able to predict temperature peaks or misalignment. This forces power plant operators to operate conservatively at higher safety standards, which effects the overall efficiency of the power plant.

For heliostats two error sources are mainly responsible for the deviation between the real irradiation and simulation: the tracking error [9] and surface deformations[10]. While the tracking error can be determined and corrected by a fast and cost-efficient method, the camera target calibration, surface deformation measurements, like deflectometry or photogrammetry, are up for today not fully automated and still associated with high costs.

In this work a machine learning (ML) approach is presented, which is able to deduce the individual surface errors and to map them to a heliostat specific flux density map. For the method only information and images are needed, which are generated by the camera target calibration. Therefore, this approach is able to eliminate the need of complex surface measurements and to detect the two main error sources of heliostats in one already established measurement. It also reduces (personnel) costs and measurement time, while raising simulation accuracies.

Furthermore, a neural network structure suitable for the task is sought. Therefore, conditioned as well as unconditioned generative adversarial networks (GANs) and neural radiance fields (NeRFs) are tested. The best results were archived by NeRF reaching a peak signal to noise ratio (PSNR) of 27.8 shown in Fig. 1.

In the next chapters, we will first discuss the state-of-the-art methods, the camera target calibration and the deflectometry. Subsequently, the ML approach and the NN architectures as well as the artificial dataset for training and evaluation will be presented. In the last chapters the results are shown, discussed and the network types are compared to each other qualitatively. Finally, an outlook on future work is given.

HELIOSTAT CALIBRATION USING THE CAMERA-TARGET METHOD

The sun tracking error describes the difference between the designated position $\vec{n}_{designated}$ and the actual position $\vec{n}_{is}$ of the heliostat (compare Fig. 2).




FIGURE 2. shows the setup of the camera-target method [6]. The focal spot of one heliostat is moved from the receiver to the target. There, a camera detects the discrepancy between the *designated* and the *is* position. With this information a linear regression solver is predicting the error parameters. Beside the misalignment, the heliostat can also have mirror deformations, which are not considered by this method.

It arises from misplacement, torsion, stress deformation, gear ratios, wind loads etc. In most commercial power plants, it is counteracted by the camera-target method (or Stone-method, named after its inventor [6]). In this process, each heliostat's focal spot is shifted individually from the receiver of the tower to a white surface (the target) (see Fig. 2). There, the position of the focal spot is detected by a camera. An algorithm measures the centroid of the focal spot and determines the deviation between the expected and the observed position. Afterwards a function template is fitted using regression.

The Stone-method is fast, cheap and with a sufficiently large data set, it is characterized by a high accuracy. The method also has its weaknesses, for example the underlying function template has to bridge the gap between high accuracy (many free variables) and data efficiency (few free variables). Especially for large fields the data set tends to be too small to cover all errors and the function template tends to overfit [11]. In addition, many other calibration concepts are known [9] but because of its widespread use and accuracy the Stone-method serves as a baseline model for all other calibration algorithms and is up for today state-of-the-art.

SURFACE ERROR DETECTION USING STRIPE DEFLECTOMETRY

The other main heliostat error are mirror deformations. These are unwanted convex and concave dents and bulges. This can happen due to incorrect canting, mechanical stress during build up or aging. The defects are difficult to locate due to the reflective surface. Deflectometry is the most commonly used and recommended [12] method for this purpose. Thereby a fringe pattern is projected onto the heliostat or the target and the reflection is observed (compare Fig. 3 (a)).




FIGURE 3. (a) schematic drawing of the deflectometric measurement. A stripe pattern is projected on the heliostat or a white surface e.g. using a video projector and the reflection is observed. (b) shows the same setup but using the sun distribution as the projected pattern. The resulting image is blurred and has less information than the stripe pattern, but it is still a convolution of the sun's shape and the heliostat's surface including all necessary information.

The normal vectors of the mirror's surface can then be derived from the deviations of the reflected fringe pattern using geometric optics. Although knowledge about the surface of the mirrors is relevant information for power plant operators and this method has been known for quite some time (first publications of automated deflectometry in the context of solar power plants date back to 2011 [13]), the measurement is still slow, associated with high maintenance effort and thus high costs. This is due to various problems e.g. morning dew, dust, wrong calibration, camera settings[12]. In addition, the measurements have to be repeated regularly, since the deformation changes over the years. A fully automated use at power plants is still pending.

NEURAL NETWORKS FOR FOCAL SPOT DEFLECTOMETRY

Considered from a physical-mathematical point of view, the image of the reflection of the sun at the heliostat results from a convolution. Even if the fringe pattern is missing, all information about the heliostat must still be present in the image (compare Fig. 3b).

Without the stripe pattern, there is a much worse local resolution of the image. Furthermore, information is lost due to the non-bijective mapping of the 3D heliostat surface onto the 2D target surface and the sun blur. Without the stripe pattern conventional (geometric) solvers have no more applicability.

By using modern machine learning (ML) algorithms, the aforementioned problems can be circumvented. For example, the ability of neural networks to serve as a universal function approximator allows them to map arbitrary input variables to any outputs. If they are used for flux density prediction, there is no need for a direct reconstruction of the heliostat's surface. While training, using the sun position as the input and the corresponding calibration images as the output, the network is able to take over the convolution and the surface information is stored (inaccessible) inside the network. The more different sun positions are included in the training process, the more information the network has about the heliostat (compare Fig. 4).




FIGURE 4. (a) Measured focal spots of a single heliostat under different sun angles. (b) Schematic drawing of the calibration and the training process. On the left side the calibration process for one single heliostat (red) is shown. The heliostat is redirecting the incoming sunlight to the target. The focal spot has a different shape depending on the sun's angle and the heliostat's surface. On the right side the network is trained using the sun's positions as an input and the focal spot's shape (as an image) as an output. The difference between predicted and calibration image is then passed back to the network. On this way the necessary heliostat information for the convolution is stored inside the network.

On this way, the bijectivity can be restored at least partially. Once trained, one network can replace the classical ray tracing process for one heliostat.

To generate the overall flux density map, the raytracer results are overlaid with the ones from the neural network. The focal spots of each measured heliostat are generated by the neural network, including the surface errors and misalignment. All non-measured heliostats are treated classically using raytracing. Depending on the chosen network architecture this can also be a calculation time acceleration.

Simple deconvolutional (transposed convolutional) networks for image generation, which map few (tabular) inputs to completely new images, have proven to be very inefficient in the past. Which is why in these cases more complex structures are used. In the following the two network types (GANs and NeRFs) are presented, which could achieve incredible results under comparable tasks and are therefore potentially applicable for the presented methodology.

Generative Adversarial Networks - GANs

GANs are the leading algorithms in the field of image generation. Even more advanced structures like adversarial latent autoencoders (ALAE) [14] achieve their incredible results with an underlying GAN structure. In general, a GAN consists of two neural networks.




FIGURE 5. Schematic drawing of the two networks inside a GAN. The left one (the generator) creates an image, either from random noise or additional conditional variables. The network on the right side (the discriminator) gets real and generated images and tries to reveal the fakes. Together they are optimizing themselves step by step.

The first network (Fig. 5 on the left) is called the generator. It generates an artificial image and takes random noise (unconditioned GAN) or additionally variables (conditioned GAN or cGAN) as its input. The noise provides a greater variability of the generated images.

The second network (Fig. 5 on the right) is called the discriminator. It takes images (as well as the control variables in the conditioned case) as an input and has one single output. The network tries to distinguish real and generated fake images.

Thus, the generator tries to fool the discriminator with false images and the discriminator tries to detect the fakes. The networks are updated using the generator's as well as the discriminator's loss. In an ideal training run, a Nash equilibrium occurs between the two networks, and improvement of one network leads directly to improvement of the other in the next step.

However, without such a balance, the prediction accuracy drops considerably. Two networks double the number of hyperparameters that have to be tuned. In addition to usual neural network (NN) problems like over-, underfit and suitable regularizations, there are mode collapses [15] and in most cases a loss without a meaningful explanation. In addition, the ability to influence the generated images depending on input variables (controllability), is still part of active research [14].

For the application at the solar tower the random vector of the generator is not needed, because the convolution of heliostat surface and incoming radiation corresponds to a univocal function. For each sun position, there is only one heliostat specific focal spot. Random variables given to network e.g. in the form of atmospheric noise are neglected in the following discussion.

As the architecture a deep convolutional GAN (DCGAN) [16] is used. The discriminator consists of convolution layers 64-(32-16-8-4-2)-1 with LeakyReLU as its activation function and batch norm layers in between. It takes 64x64 RGB images as an input. The generator is symmetric using transposed convolutional layers and batch norm as well.

The GAN is trained, both in the conditioned as well as in the unconditioned variant. In principle, more advanced structures (like Big-GAN, Style-GAN, ALAE) can also be used for this purpose. However, especially with the limited data set of the calibration, smaller, simpler network types are suitable for a proof-of-concept. They can also be used as a baseline model for subsequent attempts.

Neural Radiance Fields - NeRFs

NeRFs [17] are new to the game and are the next big player when it comes to novel view generation. With high demands on the scene (static scene and light conditions) photorealistic 3D reconstructions were possible with as few as 100 images. In a very short time, the requirements could be drastically reduced. For example the conditions on the scene [18][19] or the dataset [20].

NeRF consists of a simple multi-layer perceptron (MLP), which has the task of reconstructing a 3D scene from 2D images and their position in space. For simplicity, here the difference between dense and fine network is neglected. However, instead of training the network with the pixels directly, raycasting is used. Starting from the viewer's point (where the image was taken), discretized rays are sent through each pixel into a volume. Each of the points along these rays are then given to the network. This attempts to predict the color and density of the point in space. By summing the points and solving the rendering equation [21], the loss is then determined and the network is updated. After training a voxel like representation of the (arbitrary) object inside the volume is obtained.

The disadvantage of this method lies clearly in the speed, since for each newly generated image, the entire volume must be rasterized. However, NeRF networks are also making considerable progress in this area[22].

In order to be applied at the solar tower, the calibration process must be reformulated to match the NeRF conditions. Thus the sun position will be treated as the viewer's point and the heliostat's focal spot as the image of the object itself (compare Figure 6).




FIGURE 6. To use NeRF for focal spot prediction the sun positions are taken as the viewing angle and the images are interpreted as lying on the surface of a virtual sphere. Except for constant illumination, all original NeRF environmental conditions can be met.

Therefore, to keep the scene static a coordinate transformation must be applied. In reality, every time the sun moves 1° the heliostat is rotated by $1/2^\circ$, to keep the focal spot on the right place. Considering these rotation directives, an identical image can be observed using a static heliostat and a moving target.

By means of this transformation, the condition of a static scene can be realized. This does not apply to static light conditions. After the transformation, the sun as well as the target are moving.

In addition, the ray paths of NeRF are not adapted to beam broadening (sun blur). Thus, NeRF in its original implementation [17] is not able to create a real 3D construction of the heliostat, but only a likewise blurred virtual replica of the heliostat. Nevertheless, this replica includes all information to reconstruct the focal spot under different sun angles.

While not all assumptions are met, the setting interpreted in this way can still be assumed as a good approximation. The following results serve only as a feasibility assessment and can be used as baselines for more suitable algorithms.

DATASETS

For testing the network structures, they are trained with two different artificial datasets. The datasets are generated using the raytracer STRAL [23]. All data sets consist of image files containing the heliostat's focal spot with a resolution of 64x64 pixels and the corresponding sun positions. The datasets differ in their solar distribution and the heliostat type (compare Fig. 7).




FIGURE 7. On the left side is a schematic drawing of the simulation setting. On the right side are the 2 different heliostat prototypes used for raytracing and their corresponding focal spot as well as key features of the created datasets. The left heliostat is used for the *Sanity Check* dataset, the right one for *Semi Sphere*.

The first data set is called *Sanity Check* and uses a simple square heliostat with one facet. The sun has a fixed azimuth angle of 10° and an elevation that is scanned in 200 steps over 180° degrees.

The second data set is called *Semi Sphere*. The heliostat consists of 24 facets. The amount of data is reduced to 100 images, which were randomly drawn from the entire semi sphere above the heliostat. Unlike the *Sanity Check* dataset, a point sun was used for the simulation here. This results in sharper focal spots especially at their edges. Also, the individual facets become recognizable in the target image. Point sun images can be treated as an approximation, which holds only for heliostats very close to the target. Nevertheless, the edges of the resulting focal spot can be used as a qualitative metric for the generated images. It is assumed that the blurring of the image due to sun blur is a simpler problem, since less details have to be reconstructed.

RESULTS FOR GANS




FIGURE 8. (a) shows the results of the unconditioned GAN for the *Sanity Check* data set. The generator only gets random noise as input, so it is not directly possible to control the output. The quality and the variability of all images are very good. (b) shows the results of the conditioned GAN for the *Semi Sphere* data set. The colors indicate whether the created datapoint is intra- (red) or extrapolated (green).

The simplest case to consider for GANs is without conditioning variables. The images in Fig. 8 (a) were created with a generator that had only random noise as its input. The *Sanity Check* data set was chosen as the target. Because there is no direct controllability, it is difficult to make a quantitative comparison. Qualitatively, the images achieve a very high quality and variety. They are visually comparable with the training data. Apparently, a mode collapse did not occur. In principle, influencing the input vector (the latent space) could also affect image generation and a sun based control could be possible. For now, the unconditioned training should serve as a basic proof of concept, that GANs are capable to produce realistic looking focal spots on small datasets.

Fig. 8 (b) shows the GAN results using the sun position as the input vector, trained on the *Semi Sphere* dataset. The red images are from the interpolation test data set, the green images are from the extrapolation set. Because of the small amount of data (100 images) and the higher complexity of the problem, the quality of the results suffers considerably. Even if basically a change depending on the elevation and azimuth can be observed, the images behave very unrealistically. A controllability could not be shown for both data sets.

RESULTS FOR NERFS




FIGURE 9. shows the NeRF results for both datasets. There is only a small quality loss between both data sets, which indicates that the amount of data is sufficient. But there is a lack of quality at some angles most likely to the approximations made.

Fig. 9 shows the results of the NeRF network. The images correspond to an elevation of 12° and an azimuth between 0° to 160° . It should be noted that in the *Sanity Check* dataset only contains elevations up to a maximum of 10° . Thus, all images shown in the upper row are (slightly) extrapolated. Qualitatively the images look very similar for both datasets. This is also shown by the quantitative results. *Semi Sphere* achieves a PSNR value of 27.1 on the entire test data set, *Sanity Check* achieves 27.8. As in the original publication by Mildenhall [17], a data set of 100 images is quite sufficient for good convergence.

However, the quality depends on the angle. For many angles (for example the red boxes), the network achieves very good results. Although the edges in the *Semi Sphere* dataset are a bit blurred, the principle shape and the correct intensities were met. But there are also angles where the image quality drops significantly (compare blue boxes). This occurs equally often in both data sets and both times at nearly the same azimuth values (about 90° degrees). Since the two data sets contain very different sun angles, but the image generation breaks down at similar angles, it is very likely that the approximations made do not hold in this range.

CONCLUSION

Independent of the chosen method, the environmental conditions (shape of the heliostat, point or real sun) do not seem to have a great influence on the image quality. Since the AI algorithms do not rely on reconstructing the surface of the heliostat, but directly learn its focal spot for different sun positions, it is very likely that the complexity of the heliostat or the environment plays only a minor role. The case considered in the data set *Sanity Check* with a blurred image is therefore most likely easier to reconstruct, since only low-frequency details have to be reproduced. Besides that, the results of the shown methods are very different.

For GANs, the number of images is an absolutely relevant factor. With only 100 images, almost no usable results could be obtained. Moreover, there are very few examples of GANs with such small data sets in the literature. Training from scratch using only 100 samples is most likely not possible or needs a much better regularization. More advanced techniques, such as feature embedding [24] or pretraining, will probably have to be used.

200 images, on the other hand, could provide acceptable accuracy, at least for the unconditioned GANs. NeRF, in contrast, was able to deliver results with both datasets. Thus, it is very likely that follow-up networks such as DietNerf [20] can also reduce the NeRF data requirements for the solar tower application.

However, the quality of the NeRF results is still far from reality. Blurred edges and quality breakdowns under certain angles strongly suggest that the approximations are insufficient. Here, the unconditioned GAN produced the most realistic looking images, with clear sharp edges. However, the quality breaks down directly using the cGAN, which performed by far the worst. Neither high quality nor high controllability could be achieved.

Also, if the unconditioned GAN has the best-looking results, there is no direct way to control its outputs.

While NeRF did not have the highest quality results, it outperforms the GAN Networks in its controllability. NeRF shows in Fig. 9 that it is capable of both inter- and (slight) extrapolation. Thus, despite the lower quality, NeRF is the network with the highest applicability for heliostat reconstruction

None of used networks for focal spot reconstruction were close to reality. Furthermore, the results vary greatly in terms of controllability, data set size and image quality which is why a quantitative analysis is neglected. Nevertheless, in order to give an overview of the results, they have been summarized qualitatively in Table 1.

TABLE 1. Summary of the different findings of all network types analyzed. The controllability of the unconditional GAN model was given as None, since it was not investigated to what extent the results can be influenced by the latent space. In principle, however, this is possible.

Quality Criteria	NeRF	Unconditioned GAN	Conditioned GAN
Image Quality “Sanity Check”	Medium-High	High-Very High	Medium
Image Quality “Semi Sphere”	Medium-High	Medium	Low
Controllability “Sanity Check”	Very High	None*	Medium
Controllability “Semi Sphere”	Very High	None*	Low

OUTLOOK

For GANs, conditioned or not, training from the sketch with less than 100 data is most likely not possible. Nevertheless, there are methods that could enable GANs for this task.

An adequate pre-training with similar artificial data could significantly reduce the required data. However, it remains very unlikely for the unconditioned GAN to radically reduce the amount of data required while achieving a high degree of controllability via the latent space.

A similar situation applies to the cGAN. Here, either the data set would have to be increased first, an appropriate regularization or a pretraining would have to be found to achieve an acceptable quality at all. In order to further reduce the data set, completely new approaches are needed.

A better chance is to train the GAN not with one but with the data of the whole heliostat field. A training with data from many heliostats would increase the complexity of the problem, but also the dataset significantly. Especially with large datasets it is easier to use the strengths of NNs.

NeRF did not have the best results, but the combination of less needed data, the possibility of extrapolation and medium image quality are already convincing. In addition, many approximations and reformulations were made for the proof-of-concept where necessary for the considered NeRF network, but are dispensable in follow-ups, e.g. Dynamic NeRF [19], which allows movements within the observed scene. Same counts for relighting and shading [18].

In addition, NeRF uses a physical backbone which is tailored to the direct observation of objects. Instead of raycasting, a raytracing algorithm should be used, which also takes ray widening into account. Furthermore, the observed scene in the heliostat case is very well defined. While NeRF tries to reconstruct an arbitrary object, the heliostat shape, as well as its position and alignment is well known. So, the entire opacity prediction network section inside NeRF can be neglected in the future. This will significantly reduce both training and rendering time. Especially with regard to the potential ray tracing, such a speedup is needed.

But there are also promising options beyond the observed networks. Since the NeRF approach will also move in the direction of ray tracing, differentiable rendering stands out here in particular. With this it is possible to infer from 2 dimensional images to 3 dimensional structures. i.e. it should be possible to derive the heliostat surface directly with very few calibration images.

REFERENCES

1. Z. Wang *Design of Solar Thermal Power Plants* (Academic Press, 2019), pp 1-46
2. P. Doron , *A high temperature receiver for a solarized micro-gas-turbine* (AIP Conference Proceedings 2303, 030012 (2020)), <https://doi.org/10.1063/5.0028527>
3. Röger, M., Herrmann, P., Ulmer, S., Ebert, M., Prahl, C., Göhring, F., *Techniques to measure solar flux density distribution on large-scale receivers* (Journal of Solar Energy Engineering, 2014, Vol. 136, p. 031013)
4. Ulmer, S., *Messung der Strahlungsflussdichte-Verteilung von punktkonzentrierenden solarthermischen Kraftwerken* (PhD Thesis, VDI Verlag, Düsseldorf, 2004)
5. Göhring, F., Bender, O., Röger, M., Nettelau, J., Schwarzbözl, P., *Flux density measurement on open volumetric receivers* (Proceedings of SolarPACES 2011, Granada, Spain, Sept. 20–23)
6. K. W. Stone *Automatic heliostat track alignment method* (Google Patents, 1986)
7. Offergeld, Matthias und Röger, Marc und Stadler, Hannes und Gorzalka, Philip und Hoffschmidt, Bernhard (2019) *Flux density measurement for industrial-scale solar power towers using the reflection off the absorber.* (AIP Conference Proceedings, 2126 (110002), Seiten 1-8. SolarPACES Conference, 8-12. Okt. 2018, Casablanca, Morocco. doi: 10.1063/1.5117617).
8. Belhomme, Boris und Pitz-Paal, Robert und Schwarzbözl, Peter und Ulmer, Steffen (2009) *A New Fast Ray Tracing Tool for High-Precision Simulation of Heliostat Fields.* Journal of Solar Energy Engineering, 131 (3). American Society of Mechanical Engineers (ASME). ISSN 0199-6231.
9. C. Stattler et al. *Review of heliostat calibration and tracking control methods* (Solar Energy Vol. 207), p 110-132

10. Xiao, Jun, et al. *A review of available methods for surface shape measurement of solar concentrator in solar thermal power applications*. (*Renewable and Sustainable Energy Reviews* 16.5 (2012): 2539-2544.)
11. M. Pargmann et al. *High accuracy data-driven heliostat calibration and state prediction with pretrained deep neural networks* (*Solar Energy*, Vol 218, 2021), p 48.56
12. T. März et al. *Validation of Two Optical Measurement Methods for the Qualification of the Shape Accuracy of Mirror Panels for Concentrating Solar Systems* (*Journal of Solar Energy Engineering*, Vol 133, No 3), 2011
13. S. Ulmer *Automated high resolution measurement of heliostat slope errors* (*Solar Energy*, Vol 85, No 4) p 681-687
14. S. Pidhorskyi et al. *Adversarial Latent Autoencoders* (*arXiv:2004.04467v1*, 2020)
15. D. Bau et al. *Seeing what a gan cannot generate* (Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019) p 4502-4511
16. A. Radford et al. *Unsupervised representation learning with deep convolutional generative adversarial networks* (arXiv preprint [arXiv:1511.06434](https://arxiv.org/abs/1511.06434), 2015)
17. B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, *NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis* (CoRR abs/2003.08934, 2020).
18. R. Martin-Brualla, N. Radwan, M. S. M. Sajjadi, J. T. Barron, A. Dosovitskiy, and D. Duckworth, *NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections* (CoRR, vol. abs/2008.02268, 2020)
19. Pumarola, Albert, et al. *D-nerf: Neural radiance fields for dynamic scenes*. (Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021)
20. A. Jain, M. Tancik, and P. Abbeel, *Putting NeRF on a Diet: Semantically Consistent Few-Shot View Synthesis* (CoRR, vol. abs/2104.00677, 2021), [Online]. Available: <https://arxiv.org/abs/2104.00677>.
21. J. T. Kajiya, *The rendering equation* (Proceedings of the 13th annual conference on Computer graphics and interactive techniques, 1986), pp. 143–150.
22. Garbin, Stephan J., et al. *Fastnerf: High-fidelity neural rendering at 200fps*. (arXiv preprint [arXiv:2103.10380](https://arxiv.org/abs/2103.10380), 2021). [Online]. Available: <https://arxiv.org/abs/2008.02268>.
23. N. Ahlbrink, B. Belhomme, R. Flesch, D. Maldonado Quinto, A. Rong, and P. Schwarzbözl, *STRAL: Fast ray tracing software with tool coupling capabilities for high-precision simulations of solar thermal power plants* (Proceedings of the SolarPACES 2012 conference, 2012).
24. M. Tancik et al., *Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains* (Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020)