

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-Paal¹

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

¹*Jülich Supercomputing Centre, Forschungszentrum Jülich (FZJ), Wilhelm-Johnen Straße, 52428 Jülich (Germany)*

^{a)} max.pargmann@dlr.de

Abstract. Each solar tower power plant is designed for a pre-calculated optimal flux density distribution. Any deviation from this has a direct impact on the output power as well as the durability of the components. An accurate knowledge of the current and predicted flux density is therefore essential. Also, because this is one of the most important input variables for all subsequent power plant processes. But due to individual errors of each heliostat, this theoretical flux density is very difficult to obtain. This includes, that common methods for measuring the flux density are either inaccurate, complicated, or expensive. Although raytracers exist, which can predict a flux density by analytical calculations, such an approach does not reflect reality sufficiently. We present a novel AI based method to predict the flux density map, which is capable to include heliostat specific errors, without having to measure the heliostats surface. Furthermore, we compare the advantages and disadvantages of different network structures for this approach and show first results achieving at best a Peak Signal to Noise Ratio (PSNR) value of up to 27.8 using neural radiance fields (NeRFs).

INTRODUCTION




FIGURE 1. Vanilla NeRF prediction of a single heliostats focal spot under different sun positions (azimuth = 0° - 160° , elevation= 12°) from the test dataset. At most angles the shape as well as the intensity are very close to the targets. Nevertheless, the prediction quality is worse at the edges of the image. This is probably due to the fact that the calibration data can only approximate the hard environmental conditions on the scene of the original NeRF implementation. The overall PSNR is 27.8.

The receiver is the heart of every solar tower power plant. Here the focal spots of all heliostats are concentrated. From the distribution of the focal spots and the total incoming power per area, the radiometric quantity, the flux density distribution is obtained. The surface thus of the receiver can reach 700° Celsius [1] and the target of current research is even higher [2]. 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. Since, common methods for measuring the flux density directly are either inaccurate, or expensive, simulations are used instead.




FIGURE 2. shows the setup of the camera-target method. 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.

For this, one common approach is raytracing. For such a simulation a model of the whole environment is setup inside a virtual volume. This includes sun, heliostat and receiver models (as well as other global parameters, like atmospheric losses etc.). In the standard case each model is represented by an ideal geometric representation (compare Fig. 6).

However, without an a priori knowledge from measurements at the power plants these simulations are too inaccurate for an irradiance prediction in power plant operation. They rather serve as a guide in the design phase and for power plant operators, but do not allow for a precise prediction. For heliostats two sources of error are mainly responsible for the deviation between the real irradiation and the such an idealized solution: the tracking error and surface deformations.

The tracking error can be determined and corrected by a fast and cost-efficient method, the camera target calibration. Deformations of the heliostat that cause a deviation from an ideal flat geometry can be measure with the deflectometry technique, but his is up for today not fully automated and still associated with high costs.

In this work, machine learning (ML) approaches are presented how the measurement of the tracking error can also be used to determine the mirror deformations and thus significantly reduce costs, measurement time and inaccuracy of the simulations.

In the next chapters, we will first discuss state-of-the-art methods, the camera target calibration and the deflectometry. Subsequently, it will be shown how the deflectometry measurement can be dispensed with by means of neural networks and the camera-target calibration. In the following chapter the methods and the data set will be presented. In the last chapters the results are shown and discussed. 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). It arises from torsion, stress deformation, gear ratios etc.. In most commercial power plants, it is counteracted by the camera-target method (or Stone-method, named after its inventor [3]). In this process, for each heliostat individually its focal point is shifted from the receiver of the tower to a white surface (the target) (Fig. 2). There, the position of the focal spot is detected by a camera. An algorithm then measures the centroid of the focal spot and determines the misalignment of the heliostat via the deviation between the expected and the observed focal spot position. The 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 for large fields the data set tends to be too small to

cover all errors. Also, the underlying function template is hard to balance. It is either too simple to cover all errors or too complex and tends to overfitting [4]. In addition, many other calibration concepts are known [5] but because of its widespread 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




FIGURE 2. (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 other main errors are mirror deformations. These are unwanted convex and concave dents and bulges. This can happen due to, incorrect canting or 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 [6] method for this purpose. Thereby a fringe pattern is projected onto the heliostat or the target and the reflection is observed (compare Fig. 2 (a)). The normal vectors of the mirror surface can then be derived from the deviations of the reflected fringe pattern using geometric optics. Although 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 [7], but due to various problems (e.g. morning dew, dust, wrong calibration, camera settings[6]) the measurement is still slow, associated with high maintenance effort and thus high costs. 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

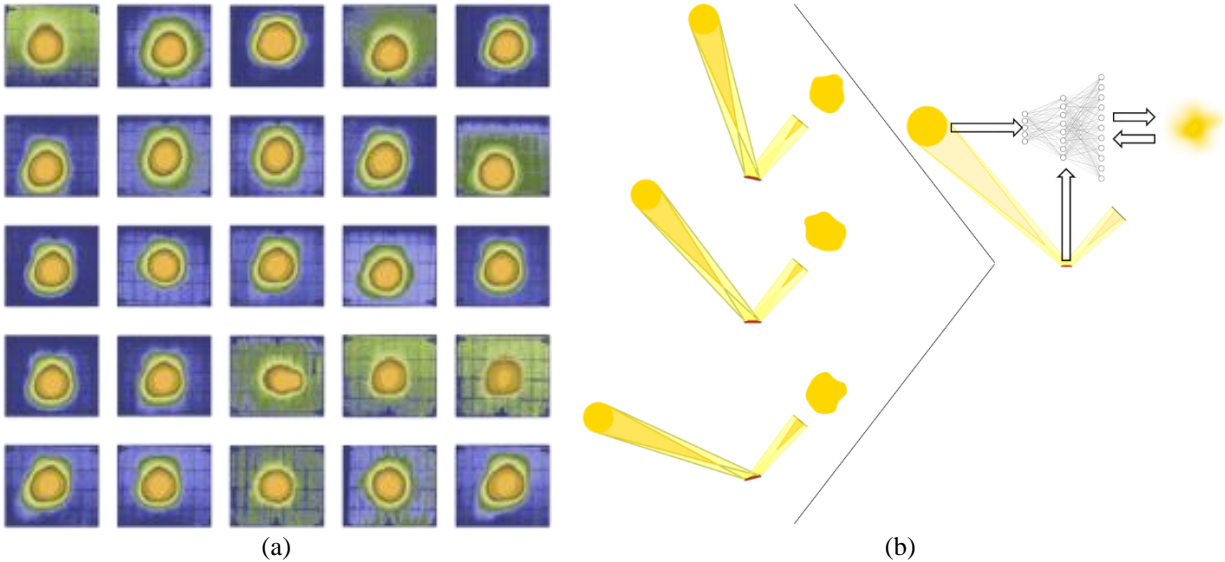


FIGURE 3. (a) Measured focal spots of a single heliostat under different sun angles. (b) Schematic drawing of the calibration and the training process. On the left side the calibration process for one single heliostat (red) is shown. The heliostat is redirecting the incoming sun light to the target. The focal spot has a different shape depending on the suns angle and the heliostats surface. On the right-side the network is trained using the suns positions as an input and the focal spots 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.

If the problem is considered from a physical-mathematical point of view, it quickly becomes apparent that the reflection results from the convolution of solar distribution and the heliostats surface. Even if the fringe pattern is missing, all information about the heliostat must still be present in the image (compare Fig. 2b).

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, as well as to the real expansion of the sun (blur). From all this, results that conventional (geometric) solvers have no more applicability.

By using modern machine learning (ML) algorithms, the aforementioned problems can be cleverly 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 not even a need for direct reconstruction of the heliostats surface. If the network is trained with the calibration images and the corresponding sun positions, the network takes over the process of convolution and the surface information of the heliostat is stored (inaccessible) inside the network. To predict the flux density, they can directly map the position of the sun to the focal spots shape, without any need for reconstruction of the heliostats surface. The more different sun positions are included in the training process, the more information the network has about the heliostat (compare Fig. 3 (b)). The bijectivity can be restored at least partially. Once trained, these networks can replace the classical ray tracing process. For this purpose, the generated images are overlaid including the misalignment derived from the same calibration measurements. From the distribution function generated in this way, the flux density can be determined directly.

In the following, network types are discussed which are suitable for this task. These are, on the one hand general adversarial neural networks (GANs) for computer vision on the other hand neural radiance fields (NeRFs) for new view synthesis.

Generative Adversarial Networks - GANs

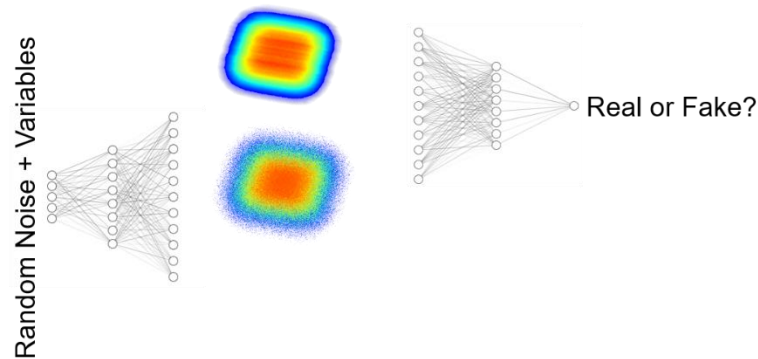


FIGURE 4. Schematic drawing of the two Networks inside a GAN. The left one (the generator) creates an image, either from random noise or conditional variables. The network on the right side (the discriminator) gets real and generated images and tries to reveal the fakes. Together they are optimizing themselves step by step.

GANs are the leading algorithms in the field of image generation. Even more modern structures like adversarial latent autoencoders (ALAE) [8] achieve their incredible results with an underlying GAN structure. A GAN consist of two neural networks.

The first network (Fig. 4 on the left), is called the generator. It generates an artificial image and takes random noise (unconditioned) or additionally variables on which the generated image depends (conditioned) as its input. Each random vector corresponds to a different generated image.

In the case considered here, the generator creates a heliostat specific focal spot image from a given sun position. Since we consider a univocal function (one sun position fits exactly to one focal spot image) the random vector can be omitted in the conditioned case.

The second network (Fig. 4 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 whether the image given to it was real or generated.

Thus, the generator tries to fool the discriminator with false images and the discriminator tries to detect the fakes. The information how well this has worked is returned to the networks after each iteration. 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, this is also where the difficulties begin. Without such a balance, the prediction accuracy drops considerably. Two networks double the number of hyperparameters that have to be tuned. In addition to the usual NN problems like over-, underfit and suitable regularizations, there are mode collapses [9] and in most cases a loss without a meaningful explanation. In addition, controllability, i.e., influencing the generated image via inputs, is still part of active research [8].

For focal spot generation the Deep Convolutional GAN (DCGAN) [10] 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 (Big-GAN, Style-GAN, etc.) 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




FIGURE 5. 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.

NeRFs [11] are new to the game and are the next big player when it comes to novel view generation. Starting with high demands on the scene being observed (static scene and light conditions) photorealistic 3D reconstructions were possible with as few as 100 images. In a very short time, both the conditions on the scene [12] and the images required [13] could be drastically reduced.

In the original release, NeRF consisted of a simple Multi Layer Perceptron (MLP), which has the task of reconstructing a 3D scene from the existing 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, discretized rays are sent through a virtual volume. Each of the points along these rays is then given to the network, which attempts to predict the color and density of the point in space. By summing the points and solving the rendering equation [14], the loss is then determined and the network is updated.

In order to be applied at the solar tower, the calibration process must be reinterpreted. The network's goal is to reconstruct the surface of one single heliostat inside of the virtual volume (compare Figure 5). Instead of taking a photo of the scene from different viewing directions, the focal spot is viewed from different sun positions. Since NeRF assumes a static scene, a coordinate transformation must be applied. Every time the sun moves 1° the heliostat is rotated by $1/2^\circ$, to keep the focal spot on the right place. Considering this rotation directives, an identical image with static heliostat but moving target can be found for each reorientation of the heliostat.

Despite these adjustments, the calibration scene cannot be transferred 1:1 to the vanilla NeRF setting. First there is the constant light assumption. This condition is not met, because after the coordinate transformation, the sun as well as the target is moving. In addition, no direct photo of the heliostat is taken, but its (blurred) focal spot is given to the network. Thus, the 3D reconstruction is not the real heliostat, but a blurred object. This object has no real physical counterpart, but it includes all information to reconstruct the focal spot under different sun angles.

Nevertheless, the setting interpreted in this way can be assumed to be a good approximation. Again, the following results serve as feasibility assessment and can be used as baselines for more suitable algorithms.

DATASETS




FIGURE 6. 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*.

For testing the network structures two different artificial datasets are used. The datasets are generated using the raytracer STRAL [15]. All data sets consist of image files containing the heliostats focal spot with a resolution of 64x64 pixels and the corresponding sun positions. The datasets differ in their solar distribution and the type of heliostat used.

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 interpolation capability of the network is tested with images between these steps and the extrapolation at an azimuth angle of 12° .

The second data set is a bit more realistic. The heliostat consists of 24 facets. The data set itself consists of only 100 images which were randomly drawn from the entire semi sphere above the heliostat, which is why the dataset was named *Semi Sphere*. Unlike the *Sanity Check* dataset, a point sun was used for the simulation here. Although the resulting image is more unrealistic, it is assumed that the sharp edges in the image are more difficult for the networks to create. Thus, the quality of the image can be better determined by sight.

RESULTS FOR GANS

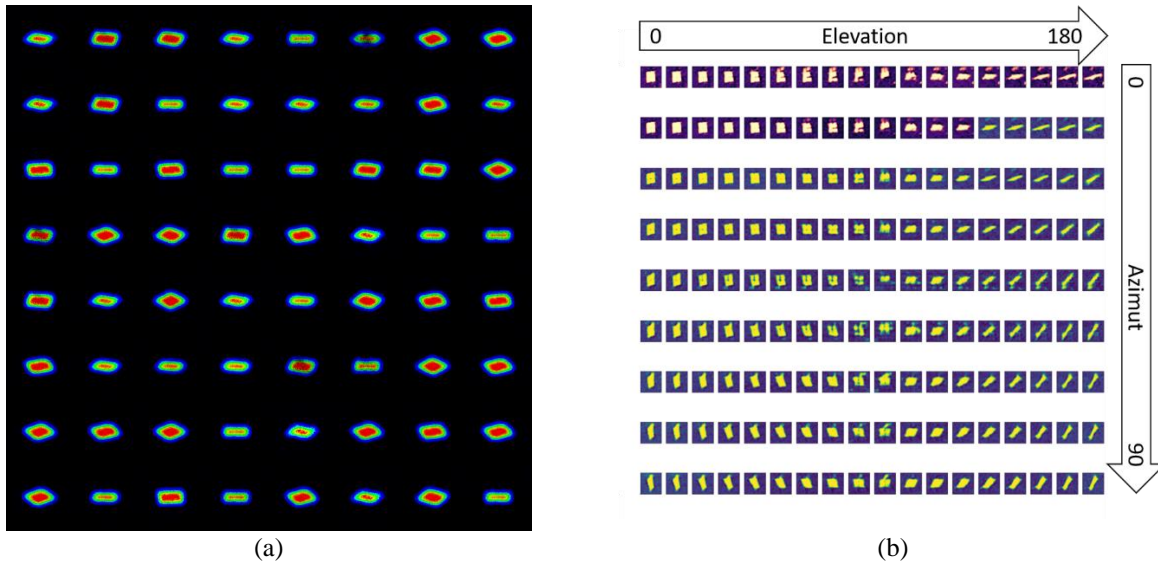


FIGURE 7. (a) shows the results of the unconditioned GAN for the *Sanity Check* data set. The generator only gets random noise as input, so it is not directly possible to control the output. The quality and the variability of all images are very good. (b) shows the results of the conditioned GAN for the *Semi Sphere* data set. The colors indicate whether the created datapoint is intra- (green) or extrapolated (red).

The simplest case to consider for GANs is the one without conditioning variables. The images in Fig. 7 (a) were created with a generator that had only random noise as 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 shall serve as a basic proof of concept, that GANs are capable to produce realistic looking focal spots on small datasets.

Fig. 7 (b) shows the GAN results using the sun position as the input vector, trained on the *Semi Sphere* dataset. The green images are from the interpolation test data set, the red 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 unrealistic. A complete controllability could not be shown for both data sets.

RESULTS FOR NERFS




FIGURE 8. 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. 8 shows the results of the vanilla NeRF network. First of all, it can be noted that the quality of the images is similar for both the Sanity Check dataset and for Semi Sphere. This is also shown by the quantitative results. *Semi Sphere* achieves a Peak Signal to Noise Ratio (PSNR) value of 27.1 on the entire test data set, *Sanity Check* achieves 27.8. As in the original publication, a data set of 100 images is quite sufficient for good convergence. For many angles (for example the red boxes), the network also achieves very close results. Although the edges in the *Semi Sphere* dataset are a bit blurred, the basic shape with the right intensity was met. At the same time, there are plenty of angles where the image quality drops significantly (compare blue boxes). This occurs equally often in both data sets and both times at larger azimuth values (about 90° degrees). Our data set does not fulfill all conditions that must be met for NeRF. Most importantly, this applies to the assumption that the illumination of the scene is constant. In our case, the opposite applies. The different images correspond to different sun position, so the illumination changes severely. So, especially at extreme angles, we expect significant deviation of the prediction of the NeRF model.

CONCLUSION

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	Vanilla 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

The three network types studied had very different results, which is why they have been summarized once again in Table 1. While NeRF did not have the highest quality results, with medium to good image quality and unbeaten controllability, it is still the closest to a final solution. At the same time, this is the method where the next steps are most clearly defined. On the one hand, the illumination of the scene is disregarded so far. While in the original NeRF the direction to the light source always remains identical and the camera angle changes, in the heliostat case both angles are changed. That the relighting or shading can be included in Nerf has already been shown [12]. This should be adapted first.

On the other hand, since it is known what kind of object is being observed, and how it is positioned, 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.

The conditioned GAN performed by far the worst. Neither high quality nor high controllability could be achieved with the *Semi Sphere* data set. Moreover, there are very few examples of GANs with such small data sets in the literature. Training from scratch is most likely not possible. More advanced techniques, such as feature embedding [16] or pretraining, will probably have to be used.

The results of the unconditioned GAN are the most difficult to assess; qualitatively, the results appeared to be of very high quality also on a data set small as 100 datapoints. However, since both the possibility of comparison and quantitative analyses are lacking, no final conclusion can be drawn here.

Considering that newer NeRF structures already provide similar results to the original paper with less than 10 images [13], it is very unlikely data sets for GAN training can be reduced this far.

OUTLOOK

Besides the network types mentioned here, there are other approaches with their own advantages and disadvantages. So far, only networks that were specifically tailored to a single heliostat have been considered. In principle, however, the data of several heliostats can be used to increase the data set. Here it would be less the goal to learn the heliostat surface, but the general transformation function. Furthermore, differentiable rendering stands out here. 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 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. K. W. Stone *Automatic heliostat track alignment method* (Google Patents, 1986)
4. 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
5. C. Stattler et al. *Review of heliostat calibration and tracking control methods* (Solar Energy Vol. 207), p 110-132
6. 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
7. S. Ulmer *Automated high resolution measurement of heliostat slope errors* (Solar Energy, Vol 85, No 4) p 681-687
8. S. Pidhorskyi et al. *Adversarial Latent Autoencoders* (arXiv:2004.04467v1, 2020)
9. D. Bau et al. *Seeing what a gan cannot generate* (Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019) p 4502-4511
10. A. Radford et al. *Unsupervised representation learning with deep convolutional generative adversarial networks* (arXiv preprint arXiv:1511.06434, 2015)
11. 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).
12. 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) [Online]. Available: <https://arxiv.org/abs/2008.02268>.
13. 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>.
14. 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).
15. J. T. Kajiya, *The rendering equation* (Proceedings of the 13th annual conference on Computer graphics and interactive techniques, 1986), pp. 143–150.
16. 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)