Machine Learning Based Fourier Phase Retrieval for Planar Near-Field Antenna Measurements

M. Dirix[1]
[1]Antenna Systems Solutions
C/ Isabel Torres, 9.-1º
Santander, 39011, Cantabria, Spain,
mdirix@asysol.com

S.F. Gregson[2, 3]
[2]Next Phase Measurements
11521 Monarch St, Garden Grove, CA, USA
[3]Queen Mary University of London, London, UK,
stuart.gregson@qmul.ac.uk

***Abstract*—The success and efficiency of many classical iterative plane-to-plane based phase retrieval algorithms is to a large extent dependent upon the fidelity of the initializing, *i.e.* guiding, phase estimation [1], [2]. This is especially so when using these techniques to recover the phase of active electronically scanned array antennas such as those employed within beam-steering mm-wave Massive MIMO antenna systems intended for 5G New Radio applications where the performance of the algorithm, and its ability to not become trapped within one of the (possibly many) local minima, is particularly dependent upon the quality of the initializing guess where access to a phase reference is not always convenient, or even possible. Many traditional phase recovery iterative Fourier methods employ simulation or passive measurement supported phase initialization [1], however this information is not always readily available, or in the measurement may require a destructive, invasive, examination of the device under test (DUT). In this work we address this issue by presenting a proof of concept which employs a machine learning based neural network [3] to estimate the initializing phase function based on the assessment of the measured amplitude only near-field pattern. Here, we show that there is sufficient information contained within the difference between the two near-field amplitude only scans to be able to determine the antenna beam steering characteristics. A simplified beam steering case with electronic scanning in one, or more, scanning axes is demonstrated and verifies the power of the novel method, as well as illustrating its inherent resilience to noise within the amplitude only measurements, and verification of the robustness of the approach thereby extending the range of measurement applications for which this class of iterative Fourier algorithms may be successfully deployed [4].**

***Index Terms*—Phase-less, Phase Retrieval, Near-Field, Machine Learning, Artificial Intelligence, Neural Network.**

I. Introduction

The phase-retrieval problem arises in many applications of electromagnetics in which the wave phase is apparently lost, or is impractical to measure, with only intensity data available. Near-field to far-field transforms require phase information to be able to accurately predict the equivalent far-field pattern from data acquired in the near vicinity of the radiator. The motivation for adopting amplitude only, *i.e.* phase-less, measurements has typically been predicated upon the difficulty of accurately acquiring phase information at high frequencies, either through a lack of a phase reference or through an inability to measure the phase with the requisite accuracy. This awkward situation has been further compounded by the increasing need to acquire measurements of integrated antenna systems intended for communication system applications where the ever closer integration of active electronics and antenna components means dedicated RF test ports are becoming increasingly scarce. Additionally, these systems tend to operate using complex waveforms meaning the CW signals, as required by standard VNAs and receivers are also unavailable. The adoption of array technology, such as that employed within modern massive multiple input multiple output (MIMO) antennas means that there is also a strong desire to perform back-projection imaging which provides non-destructive, non-invasive, antenna aperture diagnostics which also requires phase information to yield useful images of the element excitations.

Phase-less measurements can be broadly divided into three categories: multiple magnitude measurement techniques, typically requiring four separate intensity measurements, indirect holographic techniques, which require separate range illuminator(s) and far higher spatial sampling rates, and lastly multiple scan based techniques [4]. Of these, multiple scan based post-processing approaches have largely emerged as the prevailing strategy for tackling these types of measurement problems which is most probably a result of the absence of the need for additional specialist measurement equipment, such as RF circuits *etc*., they can utilise the same sampling rate, and they leverage straight forward but fairly intensive numerical post-processing. Thus, this non-linear inverse problem is generally approached using an iterative alternating projection based processing scheme which utilises the efficient fast Fourier transform. Here, the complex wave amplitude is derived from magnitude only measurements that are taken over a scanning surface where some specific change has been introduced. This difference between the respective intensity patterns drives the convergence of the algorithm. This is clearly a demanding task as there are an infinite number of complex functions with the same intensity distributions with only the band-limited, spatially limited nature of the fields to reduce the complexity and size of the search domain. Although the recovery algorithm can take many forms, most commonly the AUT-to-probe separation is changed between scans and the analytical relationship that exists between the respective probe signals is used to determine the missing phase information. This is important as these sorts of non-convex problems generally converge fairly slowly, are sensitive to the initialisation of the boundary values, are liable to producing false solutions (*i.e.* local optima) and can be limited by

measurement truncation effects. When using a typical iterative Fourier phase retrieval algorithm, the critical steps can be summarised as illustrated in Figure 1, [5].

1. Measure the amplitude of the field over plane 1.
2. Measure the amplitude of the field over plane 2.
3. Guess the initial phase over plane 1.
4. Calculate the propagated AUT aperture fields to plane 2 from plane 1.
5. Replace the amplitude estimation at plane 2 with the measured amplitude at plane 2.
6. Calculate the back-propagated fields to plane 1 from plane 2.
7. Replace the amplitude estimation at plane 1 with the measured amplitude at plane 1.
8. Repeat steps 4 to 7 until reconstructed amplitude on plane 1 (or plane 2) has converged to within a prescribed tolerance.
9. Transform the reconstructed complex fields to the far-field using a standard algorithm.

Figure 1: Plane-to-plane phase retrieval algorithm.

These routines are generally sensitive to the degree of difference between the two intensity only near-field measurements, and the quality of the initialising phase estimate. So, as the ability of the measurement engineer to choose the form of the intensity measurements is comparatively limited, it is then very important that the initial guess of the phase function does not introduce any non-physical effects that will either slow the convergence of the algorithm, or result in the algorithm becoming stagnated in some undesirable, non-optimum solution. Obtaining a reliable initialising phase estimate is important since differing phase functions with the same amplitude data will inevitably result in encountering different algorithm performance. One such example of this behaviour is that of the electronically scanned array antenna when the main-beam is not aligned with the positive *z*-axis of the antenna measurement system, *i.e.* the normal to the planar acquisition surface. Here, it can be easily shown that the absence of a priori knowledge of the scan direction will often result in the incorrect recovery of the near-field phase and an inability to reliably predict the corresponding far-field antenna pattern. This is particularly true of the orthogonal, cross-polarised component where there is a further question relating to the relative phase relationship *between* the two tangential orthogonal electric field components. The lack of a grantee of the existence of, or ability to determine, a unique solution is of particular concern since, as noted above the measurement of active Massive MIMO antennas falls into exactly this class of measurement problems. This is clear motivation for the development of robust strategies for the postulation of appropriate initialising phase functions to which the remainder of this paper is devoted.

II. Machine Learning

During the past few years there has been a very significant growth in interest in Artificial Intelligence (AI) and Machine Learning (ML), and more specifically in the use of artificial neural networks (ANN). The technique of ANN employs algorithms to parse data, learn from that data, and then make informed decisions based on what has been learnt. ANNs utilise algorithms, specifically matrices, in layers, to create artificial neural networks that can learn and make "intelligent" decisions. ANN based techniques have been shown to be highly effective in the fields of image classification, speech processing, *etc*. There are numerous examples where ANN aided research has been found to be capable of detecting minute details with images which were not discernible by the human eye, and which have been utilised, for example, in the medical field [5]. These have also been shown to be of utility in tackling complex electromagnetic tasks such as the design and optimal parameter extraction for antennas, beam-forming algorithms for adaptive antenna arrays, and data interpretation for radar *etc*. Largely however, in the field of electromagnetics the application has been focused on the acceleration of the simulation core, allowing optimisations by replacing the full-wave simulation with a previously initialized neural network [6]. However at the time of writing, the authors are not aware of any contributions to open literature devoted to its application in the field of antenna metrology. However, direction of arrival (DoA) estimation is a problem that has been successfully tackled using ML approaches with several papers being devoted to this area of research one of which successfully used a deep neural network (DNN) to sense signals arriving within a 120 degree sector from a linear array antenna [7]. Encouraged by this, and by way of establishing a proof of concept, the problem of deriving an estimation of the initialising phase distribution within an iterative, Fourier phase recovery algorithm for use with planar near-field antenna measurements was chosen as it is closely related to the DoA problem. An active, electronic beam-steering, mm-wave, Massive MIMO, antennas for 5G New Radio applications operating at FR2 was selected as being the test antenna which, as noted above, is a well know application where industry is searching for solutions for phase-less antenna measurements. ML was used to provide as accurate an initialising estimation for the phase function as was practicable with this being purely predicated upon the intensity only measurements. The main advantage of this approach is expected to be the reduced computational time and increased stability when compared to conventional numerical algorithms, *e.g.* least squares function fitting *etc*., and the absence of the need for detailed antenna information. That is to say, we are using the ML to train algorithms to learn from and act on data without being explicitly programmed for that specific task.

III. Convolutional Neural Network

A complete discussion of the AI algorithm and its development is clearly beyond the scope of this paper, however we shall now provide an overview of the technique highlighting some of the more important aspects. For the proof of concept a convolutional neural network (CNN) based on deep learning techniques has been implemented. Considering the nature of an amplitude only planar measurement, the input of the CNN is modelled around a typical image input processing layer structure, *i.e.* each measurement point corresponds to a single pixel within the input image. Thus, the so called "receptive field" size is determined by the number of measurement points available in the amplitude data. The receptive field [8] is then further condensed by a sequence of convolution layers comprising a multilayer perception, which is one of the more straight forward approaches.

The output of the CNN can be configured, among others, as a regression or a classification type output. The advantage of

the regression output type is that it allows a higher resolution result than the training input data provided. However, the complexity, especially in the case of multiple output parameters rises exponentially. Therefore, in the present case the network was chosen to have an output where the dimensionality of the data was categorical. This means the algorithm chooses between a discrete set of previously defined categorical elements, *e.g.* as in the classical example of choosing between a bird, dog or cat, or recognition of hand written alphanumeric characters.

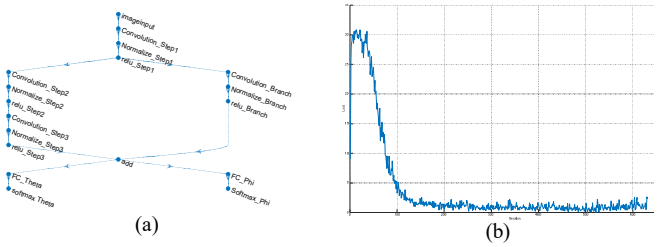

(a) (b)

Figure 2: ANN Layer graph showing hidden layers and skip branch (a) and ANN loss over training over iterations.

The concept of a neural network was inspired by the brain only here the neurons are numbers, containing values between zero and one, in a vector [10]. The value zero corresponds to inactivity whilst one denotes activity. With this first, *i.e.* input, layer of our network comprising the M by N elements within the near-field amplitude only measurement. The last layer will contain K by L neurons representing the discretized set of possible scan angles. The activation of these neurons, which again comprises a number between zero and one, represents the likelihood, *i.e.* probability, that a given near-field image corresponds to a particular scan angle. There are a number of layers in-between, called the hidden layers, and each of these will contain a number of neurons. The number of hidden layers within the network and the number of neurons is largely arbitrary, and can be varied to maximise the effectiveness of the overall artificial network. In our case however, each successive hidden layer contains fewer neurons than the last thereby condensing the large number of inputs to the smaller constringed number of possible outputs. In this implementation we used three hidden layers and these can be seen presented in Figure 2(a), shown in the left hand branch.

The training process is used to configure, *i.e.* optimise, the way in which one layer influences the next. This process assigns weights to the connections between the various neurons in one layer and the next. These weights will be positive or negative to denote the strength of the attraction or repulsion. The summation of these weights multiplied by the activation of the respective neurons provides a measure of how strongly the neurons in the first layer are connected to a neuron in the next layer. However, we require this weighted sum to fall within the range 0 to 1, and so we use a function to convert the various weights into this range. Initially the Sigmoid function was used for this task which is a function that exhibits a characteristic "S"-shaped or sigmoid curve. More recently however, the rectified linear unit, ReLU, function has been adopted for this normalisation task as it has been found to yield networks that are far more easily trained, and that is especially true for the case of deep neural networks such as those used here. Additionally, and following the biological analogy, we would like for the neuron to be activated only if the weighted sum is larger than some defined bias value before the neuron becomes meaningfully active. We therefore add this bias before passing the weighted sum through the ReLU function. This task is performed for every neuron in this first hidden layer, each of which will have a different bias. This is then repeated for every layer within the artificial neural network. This provides a great deal of different weights and biases that can be tuned during the learning process. Mathematically, this can be expressed as a weight matrix multiplied by the column vector of measurements which is added to the column vector of the biases. The ReLU function is then applied to every element within the resulting column vector to produce the transition between the neurons in one layer and next. This form of linear algebra can be seen to be of the form of convolution, and hence the name, *i.e.* convolutional neural network (CNN) configuration which can be numerically implemented very efficiently. The final step in the analysis is to compute the probability that the identified categories of discrete azimuth and elevation scan angle is the true answer, and is denoted by the Softmax layer in Figure 2. The body of the network is then divided in a main branch containing two basic steps and a skip branch containing only one basic step. The division in branches and adding skip connections improves the convergence of the model by decreasing the vanishing gradient issue and thus improving the stability of the back propagation [8].

The training process comprises determining the right set of weights and biases that allows us to classify a given near-field intensity measurement to a given discrete azimuth and elevation scan angle. We then hope that the set of weights and biases that we have determined from the known training data generalises to other data that the ANN has not been exposed to. Typically we would randomly assign vales to the weights and biases. We then pass the data through the ANN and compare the output to the known correct output and measure the error. This cost function can be the sum of the squares of the difference between the output vector and the true answer vector. We can then compute the average cost across *all* of the training examples to measure the success of the ANN and then search for the minimum by adjusting the weights and biases to optimise the effectiveness of the network, as we are not permitted to adjust the activations, *i.e.* the values contained within the neurons in the various layers in the ANN. This optimisation process is called back propagation and uses efficient stochastic gradient descent. We then repeat this for each of the training examples to train the ANN. This difference is tracked and indicated as the Loss parameter as shown in Figure 2 (b). A loss of 0 meaning the output vector and the true answer are equal. However, a more efficient strategy is to break the data up into subsets and to optimise over each of these sets. This process is termed an epoch and we may work through many such epochs during the overall training process. One aspect of this is that for the ANN to be effective; we need a *lot* of training data and even then, the tightly constrained nature of this training setup means the ANN can confidently predict the wrong answer!

IV. Traning and Testing the CNN

The proof of concept presented here is based on a simulation of the planar near-field measurement of a beam

scanning antenna at 28 GHz. The antenna diameter was set to 0.254 m (10 inches), and the beam scanning was limited to 0.5º steps up to 17.5º in the polar θ direction, and 10º steps in the azimuthal ϕ direction where we have used conventional polar spherical coordinate system with the positive z-axis aligned with the normal to the array antenna's aperture plane. This results in a dataset of 1,296 different scan angles. The training was then performed to a maximum of 32 epochs. Figure 3 shows a random selection of the amplitude distribution plots. From inspection of these, it is clearly a very difficult task for a human, by eye, to determine which measurement corresponds to which scan angle. Here, the antenna measurement was simulated with the AUT to probe distance of 3 wavelengths (λ) and at 36λ having a length and width of approximately 0.762 m (30") at $\frac{1}{6}\lambda$ sample spacing. Here, the nearer acquisition plane is denoted as plane 1 and the further away case as plane 2.

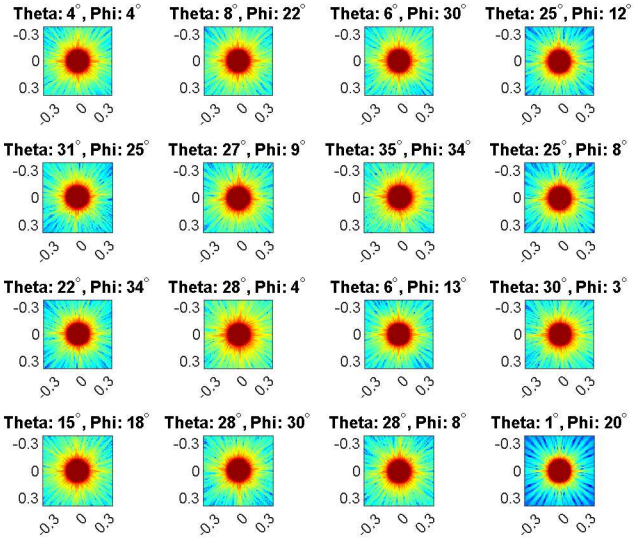

Figure 3: Random selection of amplitude distribution for various electronic scan angles.

For the purposes of efficiency, the training of the ANN was performed utilising the workstation GPU. The GPU memory requirement of the network training scales exponentially with the size of the input array, *i.e.* image pixel count, and thus it makes sense to limit the number of points. This reduction can be achieved by limiting both the length and width as well as increasing the step size of the simulated acquisition. It was found that for this particular proof of concept, error free beam scanning angle detection was possible, in the noiseless case, when sampling using half wavelength data-point spacing and that this fitted within the available GPU memory. It should be noted here that although the data presented to the CNN must have the same sample spacing and geometric locations of the measured points, however the real measured data can be under sampled electronically to match the CNN required input grid, as a higher sampling rate can be beneficial for the further processing steps.

The trained CNN was tested with various beam scanning angles and was found to be flawless in the noiseless case. While adding increasing amounts of white noise it was found that classification errors started to be encountered at a noise level of 63 dB below peak, with an overall error of 1% and which grew to an error of 18% at a noise level of 53 dB below peak. Figure 4 shows the resulting confusion matrix for each θ-angle (with flattened φ-angle dimension). This is a commonly used method of assessing how well a machine learning ANN performed in terms of classifying a given dataset [9]. A confusion matrix presents a tabular layout of the different possible outcomes of the prediction and results of a classification problem, and attempts to help visualise the outcomes. Thus, it plots a table of all the predicted and actual values of the classifier.

Here, it can be seen that for some beam scanning angles in the θ-direction the algorithm starts to provide an incorrect answer in the presence of increased noise levels. However, when further investigated it was found that for many cases the wrong answer is limited to jumping to a *neighbouring* class. Thus, providing there are sufficient angles within the detection layer, the difference between the true and estimated scan angle will be small resulting in an initialisation phase that will be very close to the desired phase function. This leads to the question how good a phase-guess is required for the phase-recovery algorithm to still converge and to run efficiently. This question couples conversely with the question as to how dense the classification grid is required to be to result in a sufficiently low error in the corresponding far-field amplitude distribution.

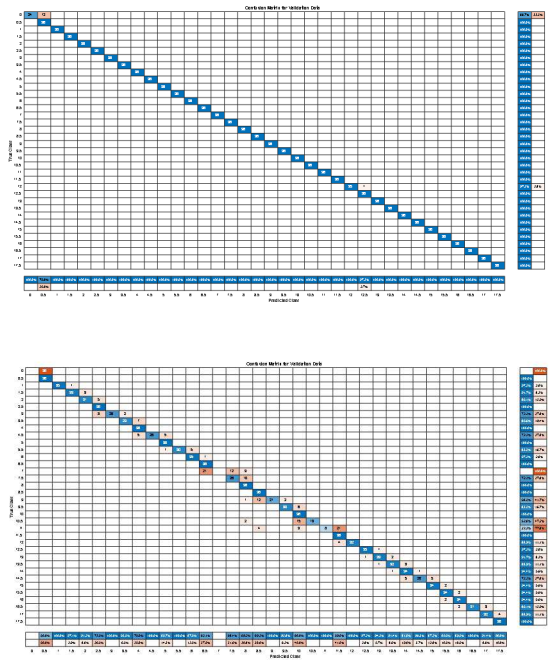

(a)

(b)

Figure 4: Confusion matrix with white noise at (a top) 63 and (b bottom) 53 dB below peak.

V. Validation of Applicability to the Fourier Method

Even though, arguably, in the case of a correct guess of the beam scanning angle the far-field amplitude reconstruction achieves the highest accuracy. That is to say, if the phase distribution is merely incorporated with the measured amplitude and a single plane is used for the near-field to far-field transformation, as was discussed above in section IV, it is ill-advised in practice, due to the nature of the classification the

angle prediction has only limited resolution and further in the presence of noise can result in misclassifications. Therefore, we are still required to employ the iterative plane-to-plate phase retrieval algorithm after the inclusion of the estimated phase distribution in the initialisation step.

In the test case the antenna beam scanning angle was set to 7º in θ and 0º in ϕ. This configuration allows for easy extraction and comparison of the antenna main beam pointing by merely displaying θ-cuts. The plane-to-plane algorithm was configured such that it stops when either a predefined convergence has been reached, or a maximum of 10,000 iterations have been performed. The convergence was calculated as the change rate of the phase distribution when expressed in dB. It was found that if this rate has reached a level of -125 dB, there are no further remarkable changes in the far-field amplitude pattern. It is noted here that the mentioned convergence is not an indication for a successful recovery, it merely indicates that no further improvements were gained and is an artefact of the nature of this inverse problem. As noted above, the phase retrieval algorithm was not the primary purpose of this paper, and the open literature contains examples of other convergence criteria and measures of adjacency that have been explored [5].

In all of the following results the far-field amplitude and phase distribution is calculated both from the near-field data of plane 1. For the reference calculation the ideal amplitude and phase distribution is used and for the recovered calculation the amplitude and phase distribution after the application of the algorithm is used (where the algorithm dictates the amplitude to be the same as the measured data, here the ideal case). The difference between the reference and the recovered far-field amplitude patterns are represented by the equivalent multipath level (EMPL) [5].

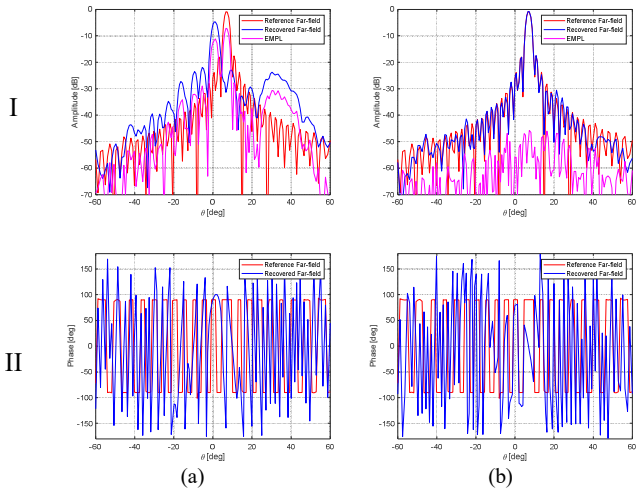

Figure 5: Recovered Amplitude (I) and Phase (II) Pattern, without (a) and with (b) initial beam scanning angle detection.

In Figure 5, the recovered amplitude and phase of an amplitude only measurement is shown on the left hand side together with a correctly guessed 7º beam scanning angle on the right. It can be seen that without any initial phase guess the plane-to-plane algorithm fails completely, but with the initial phase guess the EMPL is at a very low level. While reviewing the far-field phase distribution, Figure 5 (II), it is noted that also with an initialized phase distribution the plane-to-plane algorithm cannot be used to recover the far-field phase in a completely reliable way and thus any further processing of the far-field polarisation vector will most likely provide erroneous results.

The recovered near-field phase distribution compared to the reference phase distribution, *cf.* Figure 6, shows a quite interesting result. It can be seen here that without any initial beam scanning angle detection, the phase guess completely fails, and for the case including initialisation of the phase works only well within the region of space that the antenna physically occupies. What is also interesting is that within the region the antenna physically occupies the recovery seemingly has higher accuracy than what would be initially guessed while reviewing the far-field phase Figure 5 (II).

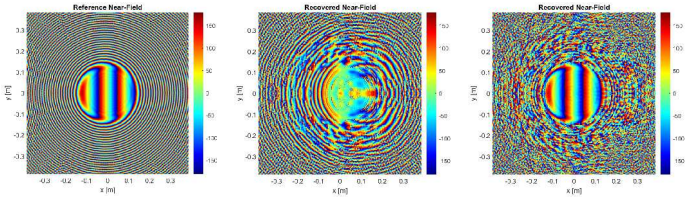

Figure 6: Near-Field Phase Distribution at Plane 1, Reference (a), Recovered without (b) and with (c) initial beam scanning angle detection

In Figure 7 the far-field amplitude patterns are shown again however in this case the impact of an erroneous classification of the beam scanning angle is examined. On the left hand side the result is shown for the same 7º beam scanning angle being classified as 6.5º and on the right hand side it being classified as 7.5º. Here, it can be clearly seen that even in presence of an erroneous classification, the accuracy of the resulting far-field amplitude pattern still far exceeds the result without initial guess thereby further confirming the utility and robustness of the ANN based processing.

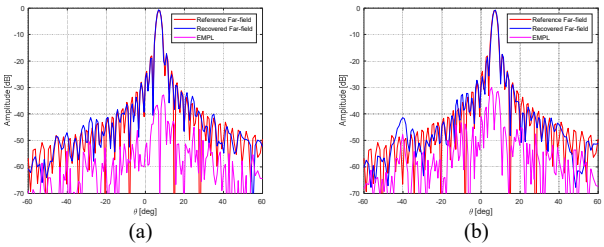

Figure 7: Recovered Amplitude Pattern with erroneous beam scanning detection at 6.5º (a) and 7.5º (b) angle

The calculated far-field, plane wave spectrum over the UV-plane, for both the measured (ideal) and the recovered case with initial phase guess, show some further interesting details, Figure 8. Here, it can be seen that for the amplitude (I) in the recovered case the main beam and its location are correctly derived, while further away from the main beam interference starts to increase. More interestingly though, while looking at the phase distribution, the measured case shows the expected parabolic phase taper in concentric rings around the pattern peak, while the recovered pattern seems to show both the expected concentric rings together with a dominant interference pattern which appears to originate from a copy of the same

concentric ring pattern at an offset angle in the same direction of the beam scanning angle. One might consider this as originating from an aliasing effect which needs further research, but is something that is not aligned with the primary thrust of the AI research focus of this paper.

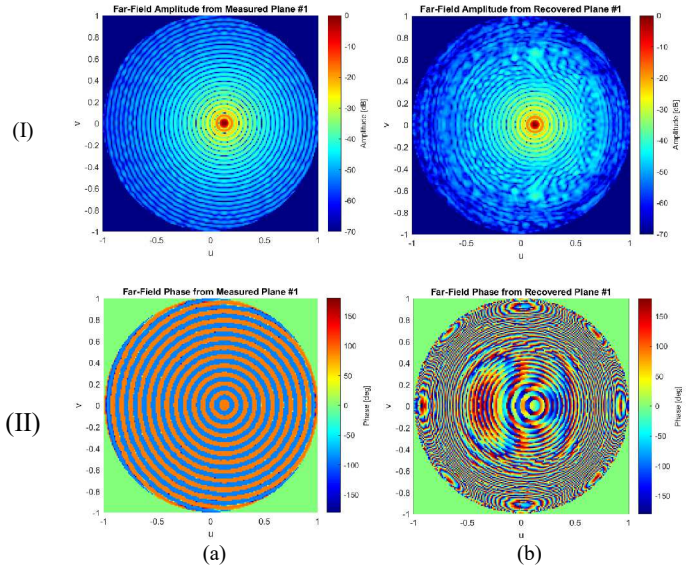

Figure 8: Far-Field amplitude (I) and phase (II) over UV-plane transformed from measured (a) and recovered (b) plane 2 Near-Field data

VI. Summary and Outlook

In this work a ML based approach is suggested to provide the initialisation of the phase distribution of a phase-less near-field antenna measurement. To show the applicability as a proof of concept a beam scanning application was used to demonstrate the ability to classify the scanning angle. The detection was further tested under noisy conditions. It was further shown that without any initialisation the plane-to-plane algorithm fails at recovering the far-field amplitude distribution. However, using an AI generated initial phase estimation the standard iterative plane-to-plane algorithm is able to successfully recover the far-field antenna pattern, even if the recovered beam scanning angle has minor deviations from the true scan angle.

While inherently a regression based output instead of the currently implemented classification method increases the complexity of the algorithm, this will eventually lead to a higher accuracy of the detected beam scanning angle, as in reality minute differences in production lead to unexpected scanning angles which are an inherent reason for the application of antenna production testing. While the expansion of the output to a regression type has merit in terms of beam scanning angle detection, the classification output can continue to promise to be more valuable in the case when antenna detection is extended to the detection of different antenna types and configurations. The inverse DL algorithms which are able to generate images can then be considered to be interesting to generate the initializing phase distribution directly. As this represents the first contribution to the open literature in this field, a great deal of work is still required. This will include examining the use of the ANN with different antenna types, configurations and refining and optimisation the ANN learning procedure.

References

[1] A. Anderson and S. Sali, "New Possibilities for Phaseless Microwave Diagnostics," in *Inst. Elect. Eng.*, 1985.

[2] S. Razavi en Y. Rahmat-Samii, „A New Look at Phaseless Planar Near-Field Measurements: Limitations, Simulations, Measurements, and a Hybrid Solution," *IEEE Antennas and Propagation Magazine,* vol. Vol. 49, nr. No. 2, April 2007.

[3] M. Diesenroth, A. Aldo Faisal and C. Ong, Mathematics for Machine Learning, Cambridge University Pr..

[4] J. Fienup, "Phase Retrieval Algorithms: A Comparison," *Applied Optics,* pp. 2758-2769, 15 August 1982.

[5] S. F. Gregson, J. McCormick and C. G. Parini, Principles of Plana Near-Field Antenna Measurements, 2nd ed., IET Press, 2023.

[6] D.-M. Koh, N. Papanikolaou, U. Bick, R. illing, C. E. Kahn Jr., J. Kalpathi-Cramer, C. Matos, L. Martí-Bonmatí, A. Miles, S. Ki Mun, S. Napel, A. Rockall, E. Sala, N. Strickland en F. Prior, „Artificial intelligence and machine learning in cancer imaging," *Nature,* 27 October 2022.

[7] O. Khatib, S. Ren and J. a. P. W. J. Malof, "Informed Deep Learning in Metamaterials," in *International Applied Computational Electromagnetics Society Symposium (ACES)*, Monterey/Seaside, CA, USA, 2023.

[8] A. Gotsis, K. Siakavara and J. N. Sahalos, "On the direction of arrival (DoA) estimation for a switched-beam antenna system using neural networks," *IEEE Transactions on Antennas and Propagation,* vol. 57, pp. 1399-1411, 2009.

[9] N. Adaloglou, "Intuitive Explanation of Skip Connections in Deep Learning," [Online]. Available: https://theaisummer.com[/receptive-field/;/skip-connections/]. [Accessed 01 07 2023].

[10] M. Li and M. Salucci, Applications of Deep Learning in Electromagnetics, IET, 2022.

[11] A. Gad, "Evaluating Deep Learning Models: The Confusion Matrix, Accuracy, Precision, and Recall," KDnuggets, 19 02 2021. [Online]. Available: https://www.kdnuggets.com/2021/02/evaluating-deep-learning-models-confusion-matrix-accuracy-precision-recall.html. [Accessed 01 05 2023].