A Neural Network Target Detector with Partial CA-CFAR Supervised Training

Jabran Akhtar and Karl Erik Olsen
Norwegian Defence Research Establishment (FFI)
Box 25, 2027 Kjeller, Norway
Email: jabran.akhtar@ffi.no, karl-erik.olsen@ffi.no

***Abstract*—Radar target detection is a fundamental process in order to differentiate objects of interest and separate them from background noise. A desired objective in detection operations is to satisfy two very contradictory requirements: offer a high probability of detection with a low false alarm rate. In this paper, we propose the utilization of an artificial neural network for the target detection procedure. The artificial neural network is proposed trained under a conventional constant false alarm rate (CFAR) detector to return positive outcomes conditionally when a true target also exists at a cell under test; otherwise the network aims to identify false positive detections and not to return positive results. Through detailed simulations it is shown that such a scheme can easily be used to train a neural network for fluctuating Swerling targets with a very comparable detection performance as of CFAR but with a significantly lower false alarm rate.**

Keywords: Constant false alarm rate (CFAR), detection, Swerling targets, neural network, machine learning

I. Introduction

The ability to discover targets at a great distance is one of the key features of radar systems. Disclosing potential objects in a reliable manner with a high probability of detection (P_D) and a low false alarm rate (P_{FA}), however, still remains a compelling task. A well-studied approach for target detection in demanding noisy environments is to employ so-called constant false alarm rate (CFAR) detectors [1], [2]. The received radar signals are then analyzed on a cell to cell basis. Each cell is compared against a select few neighbouring cells and if the value in the test cell exceeds its neighbours by a given margin then a detection is declared. A large number of such detection techniques have been proposed where the more prominent one includes cell averaging (CA) CFAR. CA-CFAR is also known to be optimal for detecting Swerling 1 targets [2] under certain conditions [3].

The last decade has also seen a significant increase in available computational power and the use of machine learning algorithms, particularly those based on deep neural networks (NN). These networks can be trained to be exceptionally good at classification of various features nested in signals and images [4], [5]. The use of machine learning methods has also been discussed in radar settings to perform target detection [6], [7], [8], [9], [10] using assorted strategies. Other works include [11] where a trained neural network is employed to switch between select CFAR techniques while [12] uses neural network in a detection strategy to adjust for clutter banks and treatment of multiple targets. Among more recent papers [13] proposes a peak shape recognition process thorough neural networks in frequency domain for continuous wave radars. Despite these advances, the available literature on neural networks and their use in radar detection applications still remains relative sparse and lacks high level maturity. For example, a common aspect found in many of the above cited papers is that the neural networks considered are of very moderate size; the networks proposed in [9], [10], [13] all only contain one single hidden layer. The progress in computational power, noticeably linked to graphical processing units, today permits the training and utilization of extremely large and deep neural networks on massive amounts of training data; the benefits of which have so far not been studied in details in relation to radar and target detection applications. This also implies that newer, simpler but practical brute force neural network training strategies are required who may easily be adapted for legacy systems and already available and processed data.

In this article, we therefore reexamine the problem of performing CFAR type of detection for fluctuating targets through the use of trained neural networks. The objective is to completely replace the traditional CFAR processor by a neural network to improve the overall detection process. Many previous works such as [7], [9], [10] attempt to emend the target detection performance beyond the one given by traditional CFAR methods. Although improving the probability of detection remains a very important factor, this work instead concentrates on a different approach where we attempt to retain the detection performance, P_D, as already given by CFAR detection strategies. The emphasis is therefore on lowering the false alarm rate which can just as well be beneficial to many practical radar setups. Simulations in various settings with fluctuating target models and noisy backgrounds are carried out to demonstrate the improvement in performance obtained through trained neural networks and evaluated against conventional CFAR methods.

II. Radar and Signal model

We consider a standard pulsed radar system where a waveform $p(t)$ is emitted at regular intervals. After emission of each pulse the incoming echoes are sampled with a given rate and a pulse compression is performed through a matched-filtering operation resulting in,

$$\hat{r}(t) = p^*(-t) * \sum_n \sigma_n p(t - \Delta_n) + w(t), \qquad (1)$$

where $t = 1, 2, ..., R$. t is a discrete parameter corresponding to different time delays and consequently range cells while R indirectly points to the the maximum radar range. In the incoming target echoes, $\sigma_n \in \mathbb{C}$ are the reflectivity levels and Δ_n is the delay associated with reflector n. $w(t)$ is white Gaussian noise and $*$ specifies convolution. The targets are assumed to be slowly fluctuating and follow a Swerling 1 distribution where the values of σ_n vary randomly from dwell to dwell but with a given mean signal-to-noise ratio (SNR). The radar may also emit multiple pulses in a coherent processing interval (CPI) and the received echoes may be integrated together coherently or incoherently to form $\hat{r}(t)$. In any case, to disclose potential targets within the received signal a detection procedure needs to be applied, typically using a CFAR detector on $\hat{r}(t)$.

A. CFAR detector

A CFAR detector processes inputs coming from the square law range samples of $r(t) = |\hat{r}(t)|^2 \ \forall \ t$ and each cell is analyzed on a cell by cell basis. A sliding window of size $2N + 2G + 1$ is chosen and moved across all possible range cells of $r(t)$, excluding potential edges. The $2N + 2G + 1$ samples specified by the window are extracted in $x(u),\ u = 1, 2, ..., 2N + 2G + 1$ and the cell in the middle of the window, $x(N + G + 1)$ cell under test (CUT), is compared against a scaled average composed of $2N$ averaging cells. G number of guard cells immediately to the right and left of CUT are ignored. The average, γ, in cell averaging is computed as mean of all $2N$ cells, $\gamma = \frac{1}{2N}(\sum_{k=1}^{N} x(k) + \sum_{k=N+G+2}^{2N} x(k))$. A detection is declared if

$$x(t)_{|t=CUT} > \gamma\, K, \tag{2}$$

where K (dB) is a defined threshold selected to satisfy requirements regarding probability of detection and false alarm rate. CFAR techniques can with advantageous be applied in noise only cases as there is then no need to separately estimate the noise and interference level which can change as a function of frequency, temperature and / or range.

B. A neural network detector

As it has been shown in the literature [6], [8], [9], [10] a target detection process for radar data may be implemented via a neural network and different strategies have been put forward on how to train shallow networks. An uncomplicated training strategy can consist of utilizing simulated or collected radar data where one a priori is aware of target locations. The network can be educated to only return a positive detection when a target is at a cell under test. The problem with this approach is that it is difficult to translate into some type of indirect classification scheme which a neural network can readily be trained to recognize. For example, if a fluctuating target ends up with a large negative SNR then such a target is highly unlikely to get identified in a general situation regardless how many such scenarios are included during training.

What we next develop is therefore an alternative training strategy where a traditional CFAR detector is employed to supervise and teach a network to differentiate out true targets and only then return a positive outcome. The training process for returning a positive outcome is hence partly based on a traditional and well-understood concept. In addition to that, the neural network must learn to handle situations where the CFAR can erroneously return positive outcomes.

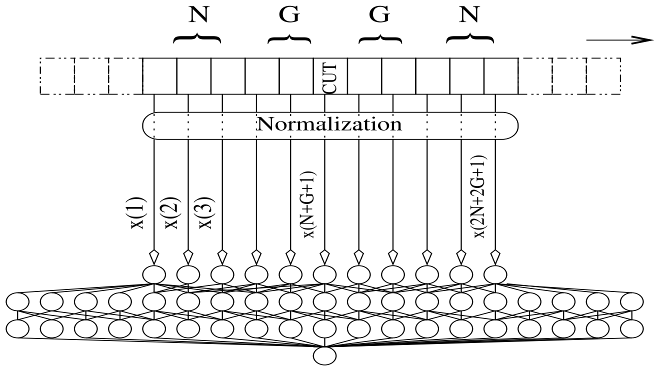

Fig. 1: Neural network detector structure

For the artificial neural network we consider a setup where the sliding CFAR window is moved across the signal, as previously; but the $2N + 2G + 1$ inputs are normalized and fed directly into a fully-connected feed-forwarding network. The normalization is carried out via standard min max normalization to map the samples in entire CFAR window to the range between 0 to 1, $\hat{x}(t) = \frac{x(t) - \min(x(t))}{\max(x(t)) - \min(x(t))}$. The number of deep layers and nodes may be varied, though generally it is understood that more flexibility is achieved with greater number of layers. Taking note of recent work on neural network [14], where it is proposed to decrease the depth and rather increase the width of networks, we suggest to use slightly wider type of networks. Having roughly at least the same number of nodes in each layer as the number of CFAR samples, or up to twice that of; and minimum two or three hidden layers seems to present largely satisfactory outcomes with good generalization. The output layer contains only a single node as a binary detection estimate is desired. The process is depicted in figure 1 where it is assumed that a wider type of network is employed.

Proper training of the neural network is an important aspect and for this we assume that L number of independent radar signal realizations $r(t)_1, ..., r_L(t)$ have been gathered wherein the number of targets and their bin positions are known. The whole signal may be utilized for training, or alternatively to reduce the computational burden representative regions within each signal may be selected where the sliding CFAR window is executed. The regions selected should contain the target at an established known position, the areas within the vicinity of the targets and samples only containing noise. Furthermore, consideration needs to be given to the noise floor level. If the noise plane is fixed at a specific value across all data then the trained neural network may not adapt well to changing noise environments. With this in mind, the average noise floor should therefore vary over a desired interval in training data to educate properly over a larger set of statistics.

The main parameter otherwise assumed locked during training is the threshold value K. This determines the P_D for the standard CA-CFAR algorithm and is the objective the network aims for. Different values of K thus lead to distinctively trained network.

The desired output from each window block of training data is 0 (no detection) or 1 for detection. Formally we propose the training conditions of output from the final node as following:

- 1: a true target is known to be present at cell under test *and* CA-CFAR returns a positive detection.
- 0: otherwise

The positive training corresponds with a traditional CFAR method for target detection, however, just as important is the fact that the network now trains to return a zero output in the absence of targets regardless the CFAR outcome. At best, we can expect the trained neural network to offer the same P_D performance as of the original CA-CFAR method but with a lower false alarm rate. The training process is visualized in figure 2. With increasingly threshold values the number of false detections arising from the CFAR process decreases, thus the network can be expected to converge to the standard CA-CFAR scheme when K is sufficiently large. The technique is therefore most suitable for low thresholding values where standard CFAR can return considerable number of false detections.

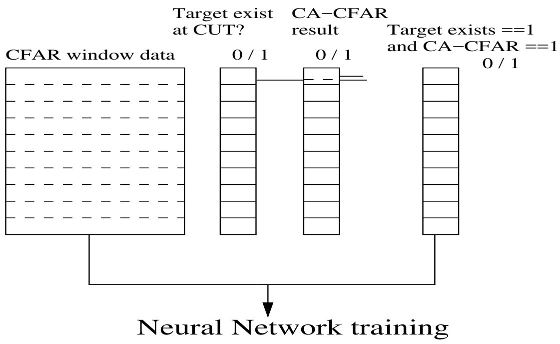

Fig. 2: Schematic description of the training process

This final performance of the trained neural network is determined by running a large untrained data set thorough it. The next section demonstrates these concepts on concrete examples with detailed simulations.

III. Simulations and Results

To compare and analyze the performance of neural networks trained with the proposed technique we model a type of pulsed radar system. A single pulse is presumed emit which is reflected back from two independent targets placed at simulated range bins 75 and 275. The total number of bins is 400. The pulse itself is of unit norm and does not introduce any additional compression gain and any beneficial antenna gains are also not incorporated. The targets reflectivity is assumed to follow standard Swerling 1 model where the mean is varied randomly during training to mirror different power levels following the exponential distribution. This is indirectly assumed to take into account impact of propagation and other environmental effects. For clarity, we point out that due to the single pulse nature of this scenario the scatterer can also be regarded as a Swerling 2 target. A total of $L = 10000$ such random signals are generated where the target average SNR varied from a minimum of -100 dB to 65 dB uniformly. It was assumed that any sidelobes are below the noise floor level. The noise floor was not fixed rather ranged between -70 dB to -35 dB also following a uniform distribution from pulse

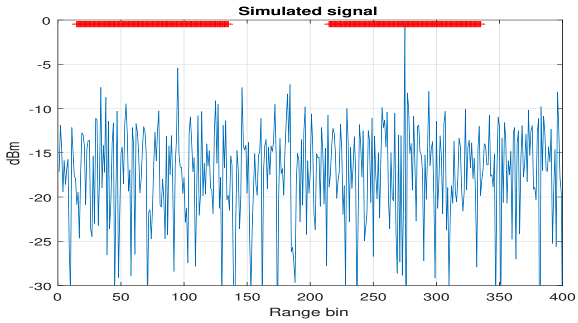

Fig. 3: Example of simulated signal with training region

to pulse. Figure 3 shows an example plot of how the two targets stand out in an high SNR situation. The regions marked by red are where CUT are taken and the CFAR tests are executed and extracted for the training process. This gives a ratio of more than 2:300 between cells containing target and no targets. Greater ratios can be more favorable for training, particularly for higher threshold values of K but care must be taken as not to overtrain the network. The CFAR parameters are set as $G = 4$ guards cells and $N = 5$ averaging cells on each side or alternatively as $G = 2$ guards cells and $N = 7$ averaging cells. The data is put to use to train fully-connected feed-forwarding networks with 2 hidden layers and 32 nodes in each layer using the hyperbolic tangent as the activation function. Training on the selected regions with 10000 simulated pulses gives rise to 20000 targets within a total of 2 million CFAR window testing samples. The training was carried out through the scale conjugate gradient algorithm until a convergence was achieved with the gradient below $1 \cdot 10^{-8}$ or performance error under $1 \cdot 10^{-8}$.

After the initial training process a total of ten million new pulses were generated; using the aforementioned methodology, each again containing two targets with random Swerling 1 distribution but now with a predetermined mean power value. The noise floor also varied, as described previously, randomly between -70 dB to -35 dB on a pulse to pulse basis. This can be seen as modeling a radar system where the noise varies over time. The CFAR sliding window was repeatedly run through all signals in full and the performance computed with traditional CA-CFAR technique and the one obtained through the trained neural network. Ten million pulses and complete CFAR evaluation provides with a total of about 1.9 billion CFAR tests containing a total of 20 million targets, which should give a reasonable good estimate of the final performance. This process was repeated with varying average target power to obtain probability of detection P_D and the false alarm rate P_{FA} curves with respect to mean target SNR. P_D was computed as the number of correctly detected targets relative to the total number of simulated targets while P_{FA} as the number of incorrectly detected targets proportionate to correctly detected targets. One benefit of a neural network is that the end result may not necessarily be either 0 or 1, instead it may be seen as a measure of probability of detection. For the described simulations, detection thresholds of an output from the final node of being greater than $d > 0.5$ and $d > 0.8$ were considered, i.e. the probability of target detection to have greater probability than no target to the point of credible object exposure.

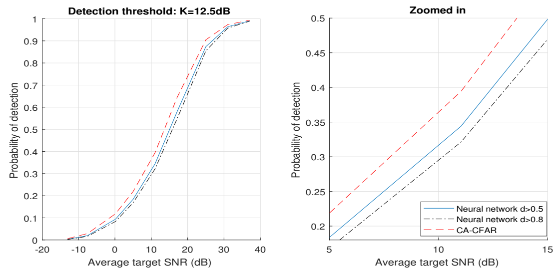

Fig. 4: P_D, trained NN (2×32) and CA-CFAR, $K = 12.5$ dB

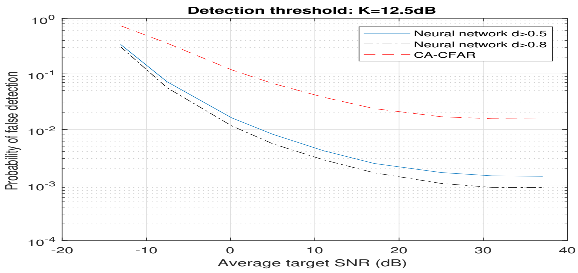

Fig. 5: P_{FA}, trained NN (2×32) and CA-CFAR, $K = 12.5$dB

The training process was repeated for varying thresholding levels, selected as $K = 12.5$ and $K = 14$. At large enough values of K, the number of false detections with CA-CFAR becomes virtually non-existing to the point that a neural network using the given training setup does not bring forth any advantage.

Figure 4 shows probability of detection curves generated through the described simulation conditions under the relative low threshold of $K = 12.5$ dB and $G = 4, N = 5$. The left side shows the full plot, while the right side represents a smaller region between 5 and 15 dB scaled up. The x-axis gives the SNR in dB which is the average of target power and the average noise level. The P_D performance level for the trained neural network detector follows the same general form as that for CA-CFAR but with a slight loss corresponding to 1-2 dB. The loss is about 3 dB if the detection parameter is set to $d > 0.8$. The false alarm rate for the same set is displayed in figure 5. In contrast to CA-CFAR, the neural network here shows a significant decrease in the number of false disclosures. Both CA-CFAR and the neural network reach a lower floor at high SNR values, however, the neural network starts off much better at low SNR figures and retains the advantage throughout.

Figures 6 and 7 depict simulations results with the threshold now set at $K = 14$ dB and $G = 2, N = 7$. As with the previous case, there is a marginal loss with regard to CA-CFAR in the probability of detection. This is nevertheless greatly compensated with a large decrease in the false alarm rate. A constant advantage of 3-6 dB can be observed for any given value of P_{FA} depending upon the threshold d.

The indicated results and the underlying training is based on the fact that the mean noise level is not constant, rather varies between a set interval. To evaluate the performance under a determined noise floor we utilized the previously trained networks (trained under a range of noise levels) and repeated the simulations with a fixed mean noise floor of -45 dB. The results of this can be seen in figures 8 and 9 for $K = 12.5, G = 4, N = 5$ and figures 10 and 11 for $K = 14, G = 2, N = 7$. The outcomes are very similar to the ones observed in the previous plots where a slight loss in Pd can be observed though the neural network manages to return far fewer number of false detections. By training over a narrower noise floor range the results could probably be improved, though this aspect is not pursued further here.

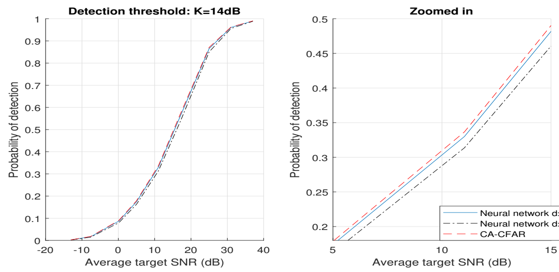

Fig. 6: P_D, trained NN (2×32) and CA-CFAR, $K = 14$ dB

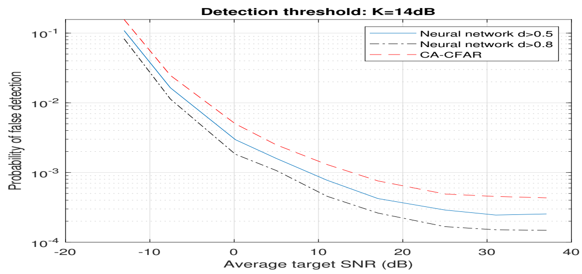

Fig. 7: P_{FA}, trained NN (2×32) and CA-CFAR, $K = 14$ dB

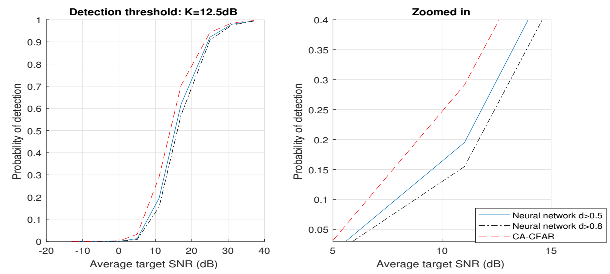

Fig. 8: P_D, trained NN (2×32) and CA-CFAR, $K = 12.5$ dB, fixed noise floor

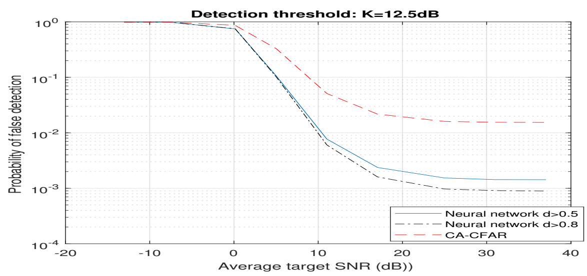

Fig. 9: P_{FA}, trained NN (2×32) and CA-CFAR, $K = 12.5$dB, fixed noise floor

A. Multiple cell targets

We point out that the strategy put forth may easily be modified to account for other type of CFAR detectors and

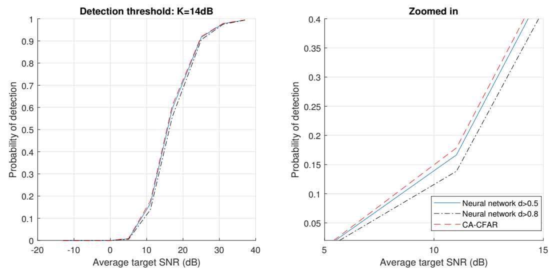

Fig. 10: P_D, trained NN (2×32) and CA-CFAR, $K = 14$ dB, fixed noise floor

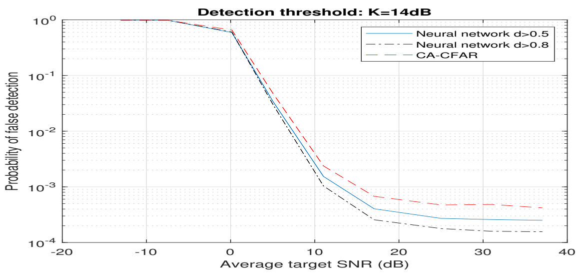

Fig. 11: P_{FA}, trained NN (2×32) and CA-CFAR, $K = 14$ dB, fixed noise floor

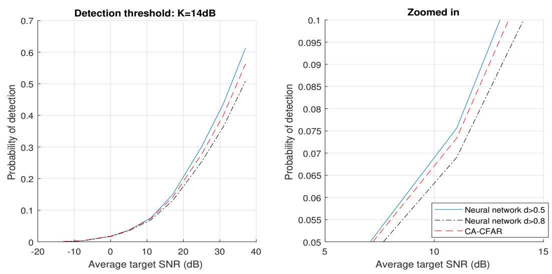

Fig. 12: P_D, target with sidelobes, trained NN (3×48) and CA-CFAR, $K = 14$ dB

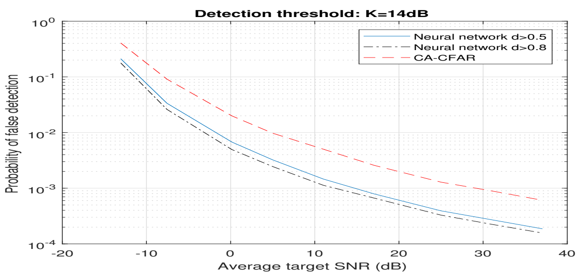

Fig. 13: P_{FA}, target with sidelobes, trained NN (3×48) and CA-CFAR, $K = 14$ dB

by incorporating other form of signal sources in radar data; and if necessary by increasing the width and deepness of the network to improve convergence. The previous cases without the inclusion of target sidelobes are important to manifest the performance in a theoretical best case situation when the objects are highly localized in range. In this section we accordingly expand the previous target model to also include sidelobes at both target ends with levels of -20 and -26 dB below peak. For the training process, the sidelobes are now also considered to be equitable to target positions and the training criteria for output from the final node therefore adjusted to:

- 1: a true target or its sidelobes are known to be present at cell under test *and* CA-CFAR returns a positive detection.
- 0: otherwise

Beside these modifications, the same data generating and training process is applied on a slightly bigger network, now with 3 hidden layers and 48 nodes in each layer. Smaller type of networks do not seem to generalize that well to more complicated models. The final outcomes for $K = 14$ dB with $G = 4$ and $N = 5$ can be seen in figures 12 and 13. With the presence of sidelobes, the broader and deeper neural network most likely has a complementary advantage of being able to identify the patterns generated by sidelobes and the leverage of this can be seen in the probability of detection curves. The overall performance is very close to CA-CFAR but can actually be improved upon slightly by using the node threshold of $d > 0.5$, particularly at high SNRs. Notice that that P_D does not go all the way to 1 as all target sidelobe cells are now calculated as potentially legitimate targets. The P_{FA} is, as previously for target models without sidelobes, significantly lower than standard CA-CFAR further validating the usefulness of trained neural networks.

For an even further extended type of target model, four uncorrelated scatterers were simulated localized at range bins 74,75,76 and 275 each with an independent Swerling 1 fading and sidelobe levels of -20 and -26dB each. This can be seen as modeling a bigger target with several independent reflectors, or multiple closely spaced targets at the first cluster. Figures 14 and 15 display the curves obtained under the detection threshold value of $K = 12.5$ dB with CA-CFAR and a neural network trained specifically under this scenario ($G = 4, N = 5$). Under this low threshold value CA-CFAR, as expectantly, has a large false alarm rate. The trained neural network manages to lower this greatly, though with some limited degradation in the P_D especially at high SNRs. This thus offers a very functional alternative to CA-CFAR with a different kind of trade-off.

The subject of how an artificial neural network trains up to recognize different patterns is beyond the scope of this article. Some insights into this can nevertheless improve the understanding of the underlying detection process. A possible approach for this is to examine some example plots on what differentiates a false CA-CFAR positive detection while a trained neural network returns a correct negative response. Figure 16 provides averaged plots of samples taken from 500 normalized CFAR windows all resulting in an identical outcome. The samples are taken from the last simulation run of 4 modeled targets and a detection threshold of $K = 12.5$ dB. The top plot shows the standard case when both CA-CFAR and the neural network return a true positive outcome. In this case the target is in the middle and exceeds the scaled mean value given by the averaging cells with a gradual intensity decline due to the sidelobes. This is very much what one would expect. The lower plot demonstrates the case when CA-CFAR returns a false positive outcome while the neural network succeeds in correctly not offering a detection ($d < 0.8$). The problem with CA-CFAR seems to be that it absolutely ignores all data in the

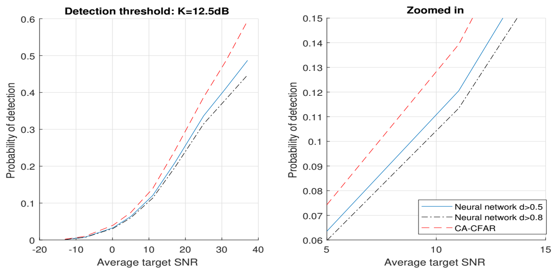

Fig. 14: P_D: Multiple target trained NN (3×48) and CA-CFAR, $K = 12.5$ dB

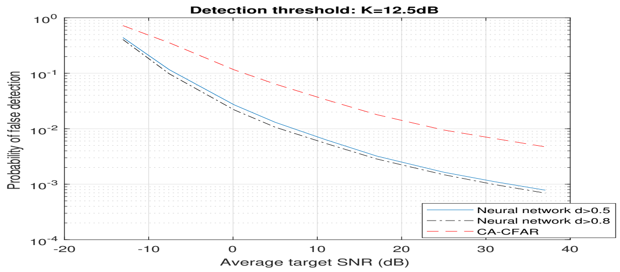

Fig. 15: P_{FA}: Multiple target trained NN (3×48) and CA-CFAR, $K = 12.5$ dB

guard cells. If the guard cells contain high amplitude data then that information is not taken into account by CA-CFAR. The neural network utilizes the whole window for assessment and the presence of very large values in guard cells not showing a consistent drop from the peak is taken into consideration. The target and noise characteristics based on teaching data can then imply that the cell under test is less likely to be a target.

Overall, the general results strongly point to the fact that as a substitute for a classical CFAR detector one may easily train an appropriately sized artificial neural network with sufficient data to perform the task of target detection.

IV. Conclusion

This paper considered the application of artificial neural network for target detection with respect to fluctuating targets. For the training of the neural network a new type of, simple but practicable, scheme was prospected based partly on the conventional CA-CFAR detector. The network employed standard CFAR methods for learning when to return positive outcomes combined with when a target was also known to be at the evaluated position. In all other cases, the network was educated not to return positive outcomes. It was shown that such a trained neural network can offer comparable, or slightly lower, performance as of traditional CA-CFAR for target detection but with a noticeably lower false alarm rate. Detailed simulations were carried out to demonstrate this in a statistically conclusive manner where the major benefits were found to be under low thresholding values. The forthright implementation of the proposed training strategy makes for easy adaptation.

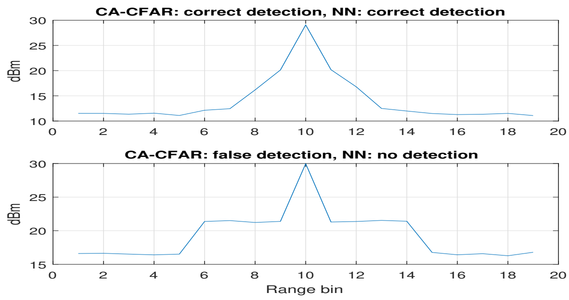

Fig. 16: Averaged plots over normalized CFAR window data

References

[1] P. P. Gandhi and S. Kassam, "Analysis of CFAR processors in non-homogeneous background," *IEEE Aerospace and Electronic Systems Magazine*, vol. 24, no. 4, pp. 427–445, July 1988.

[2] W. L. Melvin and J. A. S. (Eds.), *Principles of Modern Radar*. SciTech Publishing, 2013.

[3] P. P. Gandhi and S. A. Kassam, "Optimality of the cell averaging CFAR detector," *IEEE Trans. Inf. Theory*, vol. 40, no. 4, pp. 1226–1228, Nov. 1994.

[4] R. Reed and R. J. Marks, *"Neural Smithing, Supervised Learning in Feedforward Artificial Neural Networks"*. MIT Press, 1999.

[5] C. M. Bishop, *"Pattern Recognition and Machine Learning"*. Springer, 2006.

[6] F. Amoozegar and M. Sundareshan, "A robust neural network scheme for constant false alarm rate processing for target detection in clutter environment," in *Proc. American Control Conference*, 1994.

[7] P. P. Gandhi and V. Ramamurti, "Neural networks for signal detection in non-gaussian noise," *IEEE Trans. Signal Processing*, vol. 45, no. 11, pp. 2846–2851, Nov. 1997.

[8] N. Galvez, J. Pasciaroni, O. Agamennoni, and J. Cousseau, "Radar signal detector implemented with artificial neural networks," in *Proc. XIX Congreso Argentino de Control Automatico*, 2004.

[9] K. Cheikh and F. Soltani, "Application of neural networks to radar signal detection in K-distributed clutter," *IET Radar, Sonar & Navigation*, vol. 153, no. 5, pp. 460–466, Oct. 2006.

[10] D. Mata-Moya, N. del Rey-Maestre, V. Peláez-Sánchez, M.-P. Jarabo-Amores, and J. M. de Nicolás, "MLP-CFAR for improving coherent radar detectors robustness in variable scenarios," *Elsevier Expert Systems with Applications*, vol. 42, no. 11, pp. 4878–4891, July 2015.

[11] B. Rohman, D. Kurniawan, and M. Miftahushudur, "Switching CA/OS CFAR using neural network for radar target detection in non-homogeneous environment," in *Proc. International Electronics Symposium*, 2015.

[12] N. Galvez, J. Cousseau, J. Pasciaroni, and O. Agamennoni, "Efficient non homogeneous CFAR processing," in *Proc. Latin American Applied Research*, 2011.

[13] S. Wunsch, J. Fink, and F. K. Jondral, "Improved detection by peak shape recognition using artificial neural networks," in *Proc. VTC Fall*, 2015.

[14] S. Zagoruyko and N. Komodakis, "Wide residual networks," in *Proceedings of the British Machine Vision Conference (BMVC)*, 2016.