






The Bird in the Swarm: How to Detect and Resolve Closely-Spaced Targets in RADAR

Benedict Rui Yang Lee¹ , Shriya Peruri² , and Jing Tong Teo² 

¹ NUS High School of Math and Science, 20 Clementi Avenue 1, Clementi 129957, Singapore
benedictleeruiyang@gmail.com

² Raffles Institution (Junior College), 1 Raffles Institution Lane, Bishan 575954, Singapore
peruri.shriya@gmail.com, jingtongteo@gmail.com

Abstract. The ability of radars is best harnessed for national defence by detecting and localising incursive objects in its airspace. However, signals of closely-spaced targets such as swarms of drones often smear into a single, indistinguishable signal, leading to the incorrect deduction that only a single target is present using the conventional Cell Averaging-Constant False Alarm Rate (CA-CFAR). Hence, this project aims to create a detection algorithm to increase the accuracy of detecting closely-spaced targets. We propose an improved algorithm combining Ordered Statistic-Constant False Alarm Rate (OS-CFAR), Mean-Shift clustering and CLEAN algorithms. Algorithm testing was done through simulations. These results were analysed using the Receiver Operating Characteristic (ROC) graph which showed a large improvement in efficacy, achieving a perfect ROC for the improved algorithm. This is also observed upon validating the improved algorithm on real data collected through experiments mimicking the behaviour of two closely-spaced targets where the target with higher SNR completely obscures the other target. There was a substantial decrease in the probability of false alarm for the improved algorithm of 21.1% on average, reaching 32.1% for an experiment documenting the greater sidelobe interference. Thus, the improved algorithm proves to be a promising alternative.

Keywords: RADAR · Closely-spaced targets · Modified CLEAN algorithm · OS-CFAR · Mean shift clustering

1 Introduction

Radars utilise the reflections of radio waves—long-wavelength electromagnetic radiation—off of target bodies to determine their presence, as well as parameters such as range, velocity, and angle. The all-weather, day-night, and long-range surveillance capabilities provided by radars make them essential for time-critical applications like airspace surveillance.

One problem in radar signal processing is resolving closely-spaced targets. Radar signal processors do not have prior knowledge of the number of targets present in the surveillance region and their characteristics. This coupled with the fundamental limits in parameter resolution due to the radar's wavelength results in neighbouring targets interfering with each other. Hence, neighbouring radar targets often smear into a single, indistinguishable signal, causing less targets to be detected. This problem is exacerbated when a target has a significantly higher signal-to-noise ratio (SNR) than another—the weaker target is obscured under strong sidelobes, making it difficult to accurately detect it.

In real-world scenarios, this problem arises when trying to detect closely-spaced formations of planes or swarms of drones with incursive and adversarial intent. The accuracy of individually localising such targets is crucial for an adequate response. This holds extreme importance for national defence during conflicts—during the Gulf War, concentrated formations of coalition aircraft overwhelmed the extensive air defences of Iraq. While this could be attributed to many factors, a key factor was the inability to accurately detect close targets, limiting their ability to strike back. Hence, this project aims to design an algorithm tailored to the automatic detection and localisation of closely-spaced targets to achieve significant improvement in efficacy over a conventional algorithm. This will be done using a modified CLEAN algorithm to iteratively remove strong signals. Through this method, the detectability of weaker neighbouring targets is hypothesised to increase, addressing the aforementioned challenge.

2 Materials and Methods

In radar, targets are detected and localised on a range-velocity (RV) map (Fig. 1). To form a RV map, the reflected target echoes are processed against the emitted radar signal to compare their shifts in time and Doppler, which correspond to range and velocity respectively. The colour magnitude (as labelled by the colour bar in Fig. 1) represents the SNR in each RV cell. The higher the SNR of a cell, the higher the probability of a target being present. In Fig. 1, a strong target at 16 m range and 0.5 m/s velocity can be observed, as well as strong returns along the zero-velocity line from non-target, stationary objects like buildings.

To develop algorithms to detect swarm targets, datasets with closely-spaced targets were obtained through simulation and experiment. Targets are defined to be closely spaced in the RV map if they are separated by one cell in either range, velocity or both axes.

For initial algorithm development, closely-spaced targets were simulated in a clean, ideal environment. This was done using a pulse-Doppler radar simulator, which used chirp signals to simulate the radar-emitted signal and the target echoes against a background of additive white Gaussian noise. By varying the range, velocity and SNR of six targets, three distinct groups of closely-spaced targets of three, two and one targets respectively, were simulated.

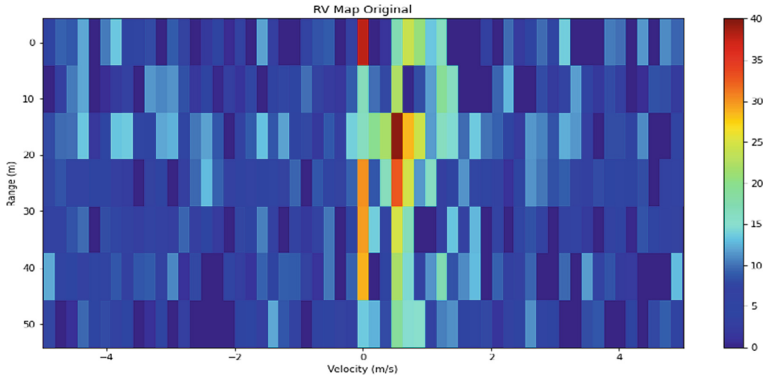


Fig. 1. RV map

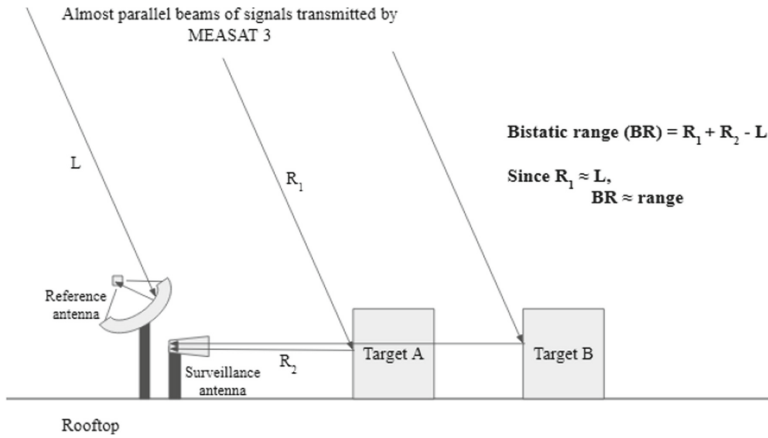


Fig. 2. Bistatic radar set-up used for the experiment

For algorithm validation, a trial experiment was conducted to obtain real radar data. A bistatic radar was used for this experiment (Fig. 2), which differs from typical monostatic radar where one antenna acts as both the transmitter and receiver. The two radio frequency channels consisted of an elliptical dish antenna¹ that acts as a reference antenna to receive radio signals from the geostationary MEASAT-3 satellite, and a horn antenna² that acts as a surveillance antenna to receive target reflections over the region of surveillance. The signals were collected at 11896 at 50 MHz sampling rate and digitised using a software-defined radio.

¹ Images of the experimental equipment can be found in the Appendix.

² See Footnote 1.

Two targets were used in the experiment: a big metal plate on a trolley (Target A¹), and a small metal plate on a rack (Target B¹). In Experiment 1, to obtain data varying in range and velocity, both targets started at close range and were pushed away and towards the surveillance antenna at different speeds for 10 repetitions. In Experiment 2, to obtain data varying in range only, both targets started 8 m away from each other and were pushed away and towards the surveillance antenna at similar speeds for 10 repetitions. To vary SNR, the metal plate was removed from Target A¹ and both experiments were repeated as Experiments 3 and 4. These experiments were designed to mimic the movement of two closely-spaced targets and receive data where the targets are close enough in either range or velocity for their sidelobes to interfere.

Due to hardware limitations, only one surveillance channel was available. Thus, the angle parameter was not processed due to the inability to perform beamforming. Moreover, a realistic interpretation of closely-spaced targets in the real world would be targets that are closely spaced in the range, velocity and angle domain. Therefore, it is fair to assume that the targets are closely-spaced in all domains and that the angle domain is not a key distinguishing factor.

3 Algorithm Development

The conventional detection algorithm, CA-CFAR, calculates the noise level by finding the average power of cells around the cell under test (CUT) in an “×” shape (Fig. 3). This excludes the cells in a “+” shape (Fig. 4) around the CUT which contains the sidelobes of a potential target, hence preventing interference with the average noise level [1]. Target detection occurs when the CUT is sufficiently above the noise level by a threshold value.




Fig. 3. “×” shape sampling around CUT

However, for closely-spaced targets, sidelobes of a target interfere with the cells around a CUT of another target. This raises the calculated threshold level above the CUT, causing the real target to be undetected. In addition, while an “×” shape for threshold calculation is desirable for normal targets, it results in sidelobes being detected as separate targets, giving a high false alarm rate. This interferes with the determination of which detected targets are actual closely-spaced targets. Figure 5 illustrates the conventional algorithm against the improved algorithm proposed.




Fig. 4. “+” shape sampling around CUT




Fig. 5. Overview of radar signal processing and integrated detection algorithms

3.1 Change 1: OS-CFAR and “+” Shape Sampling

To overcome the aforementioned limitations, “ \times ” shape sampling was changed to “+” shape sampling. The prior estimates noise more accurately while the latter estimates target sidelobes more accurately. The “+” shape (Fig. 4) was determined to be more appropriate for this algorithm as it allows the threshold of the cells containing sidelobes to be disproportionately high, making only the highest SNR cell at the target centre to be detected as a target. To further remove low noise outliers, OS-CFAR was adopted. The threshold for detection was calculated by ordering the cells in the “+” shape by power magnitude and taking only the 50–100th percentile of values for averaging. Hence, fewer sidelobes are detected as targets, improving the false alarm rate significantly. It has also been proven that for clutter or multi-target situations, OS-CFAR outperforms other types of CFAR algorithms [2].

3.2 Change 2: Mean-Shift Clustering

To remedy the issue of multiple detections for individual targets caused by sidelobes, the Mean-Shift clustering method [3] was adopted. Unlike other centroid-based methods which require the number of clusters to be pre-specified, Mean-Shift only requires a bandwidth to determine the size of clusters formed. This agrees with the fundamental notion of this project; there is a lack of ground truth to determine the number of clusters to be formed. After some simulations, an optimum bandwidth was calculated as the detections of target sidelobes were generally within a certain number of cells from the actual target location in the RV map.

However, OS-CFAR and Mean-Shift do not enhance the accuracy of closely-spaced target detection as the target with lower SNR may be undetected if it is obscured by the sidelobes of the target with higher SNR. Even if the target with lower SNR is successfully detected, Mean-Shift will cluster it together with the target with higher SNR due to their close proximity. To overcome this, an adapted CLEAN algorithm was used.

3.3 Change 3: CLEAN Algorithm

The CLEAN algorithm was initially developed for radio astronomy to overcome the problem of bright objects having responses that spread out and obscure dimmer objects [4]. Although several papers have applied the CLEAN algorithm to radar signal processing, it has only been shown to be effective in simulations [5–8]. This project attempts to prove the effectiveness of the CLEAN algorithm in detecting closely-spaced targets on real-life data too.

The CLEAN algorithm works by first identifying the target with the highest SNR in each cluster found after Mean-Shift clustering. Secondly, it simulates the signal of those targets using the range, velocity and SNR data extracted. The algorithm then subtracts the simulated signal from the original signal received by the radar to completely remove these targets and their sidelobes that may have potentially obscured closely-spaced targets with significantly lower SNR. This new signal obtained will undergo another round of signal processing to display the otherwise hidden targets previously obscured by the sidelobes of the high SNR target (Figs. 6, 7). This algorithm runs iteratively until all the hidden targets are uncovered and no targets are detected. Lastly, all removed targets are added back to an RV map to reveal all detections.




Fig. 6. Initial detection (only highest SNR target)




Fig. 7. Detection after iterations of CLEAN

4 Results

The performance of the improved algorithm was tested against the conventional CA-CFAR algorithm on simulated radar data. Receiver Operating Characteristic (ROC) curves were used to conduct analysis of the effectiveness of both algorithms. The ROC curve plots the probability of detection (PD) and probability of false alarm (PFA) for each algorithm at different threshold values. 45 average values of PD and PFA were obtained with 45 different thresholds over multiple runs. The simulated targets had an SNR ranging from 10 to 22 dB (Fig. 8).

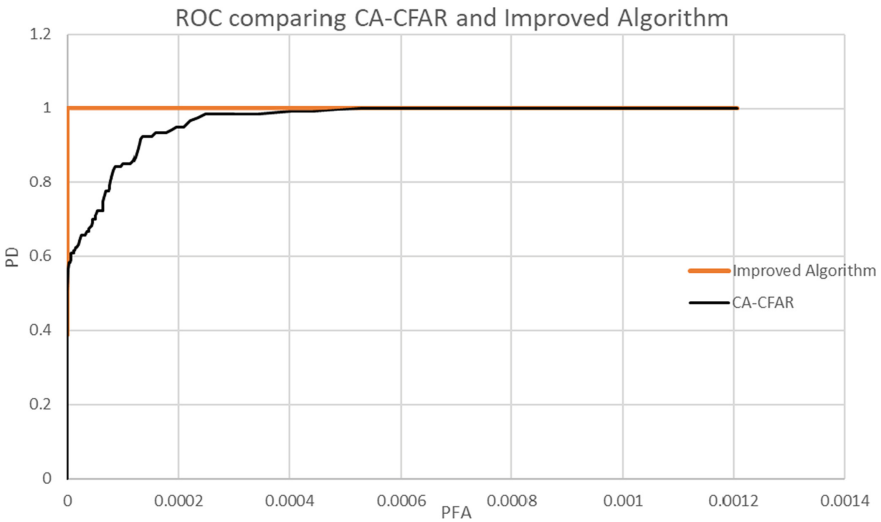


Fig. 8. ROC comparison curve for simulation of 6 targets

The ROC curve for the simulation of the 6 targets shows a clear improvement in efficacy for the improved algorithm developed. The improved algorithm is more effective as it has higher or equal PD for every PFA as compared to CA-CFAR.

The performance of each algorithm was also compared using real radar data collected from the experiment. For this data, a ROC was not generated as the experiment cannot be controlled to get sufficient points for plotting ROC at an SNR meaningfully. Instead, the PD was fixed at 1 by using a base threshold that ensures both targets are detected across all scans and the subsequent average PFA value for each experiment was compared for both algorithms (Fig. 9).

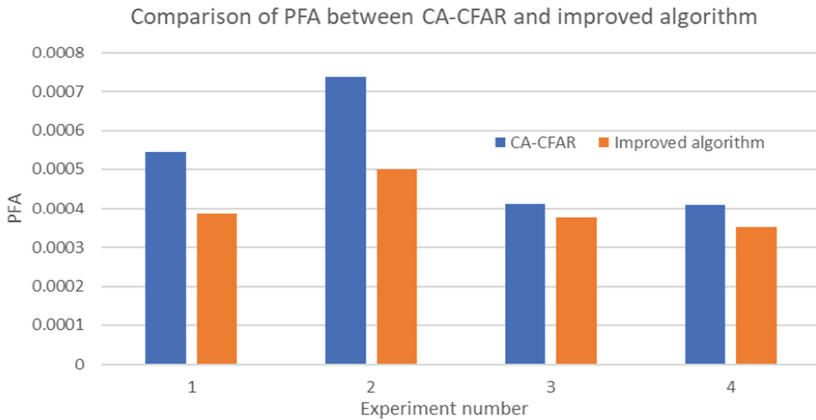


Fig. 9. Bar graph of PFA comparing performance of algorithms

Experiment 3 had a 9.1% decrease of PFA from CA-CFAR to the improved algorithm as compared to 29.1% in Experiment 1. Meanwhile, Experiment 4 had a 14.2% decrease in PFA from CA-CFAR to the improved algorithm as compared to 32.1% in Experiment 2.

5 Discussion and Conclusion

In simulated data, the improved algorithm provided the highest possible theoretical efficiency, achieving a perfect ROC curve. This is because the signal of the target with the highest SNR could be re-simulated well as the parameters of the target were extracted accurately. This allowed for a full subtraction of the target's signal from the received signal. However, the perfect effectiveness of CLEAN on simulated data cannot be replicated on real data due to 2 reasons: the data file captures target movement over a period of 0.5 s and the velocity resolution is too fine at 0.2 m/s per cell. This results in smearing across several velocity cells which causes inaccurate extraction of the target parameters. Hence, the re-simulation of the target signal is inaccurate and consequently, the target signal cannot be completely subtracted.

Despite these limitations, the improved algorithm proved effective over CA-CFAR for closely-spaced targets evidenced by the fact that the average improvement in PFA was 21.1%. Furthermore, in Experiment 1 and 2, there is greater sidelobe interference due to the presence of a target with a significantly higher SNR. Therefore, the effectiveness of the improved algorithm over CA-CFAR was better observed: a staggering improvement in the PFA of the improved algorithm in Experiment 1 and 2 compared to 3 and 4. In conclusion, the improved algorithm is more effective in detecting closely-spaced targets in both a simulated and real environment.

Another limitation is that this algorithm requires more time to compile returns due to having to simulate targets as part of the CLEAN algorithm and OS-CFAR requires higher computational power. This results in a slightly worse than “live” refresh rate of the system.

Our algorithm is methodologically curated for optimum performance in closely spaced targets. This could work in concert with a conventional radar signal processing system to run in parallel. This would enable an overall encompassing system of both extreme, precise fidelity of our curated and optimised signal processing code and the high refresh rate of the conventional system. This would provide a radar operator with significant improvement in effective data such that an optimum response to a convoluted threat can be efficiently mounted.

Acknowledgements. We would like to thank our project mentor, Mr Alastair Wee, for his guidance throughout this project. We would also like to thank the DSO SR department for their help in providing resources and conducting our experiment. Lastly, we like to express our appreciation to the DSO HR department for their help with financing our project and providing us with the logistics we require.

Appendix

See Figs. 10, 11, 12, 13 and 14.



Fig. 10. Surveillance antenna (horn antenna)



Fig. 11. Reference antenna (dish antenna)



Fig. 12. Trolley (Target A)




Fig. 13. Trolley (Target A with metal plate)



Fig. 14. Rack (Target B)

References

1. Shrivathsa, V.S.: Cell averaging: constant false alarm rate detection in radar. *Int. Res. J. Eng. Technol.* **7**, 2433–2438 (2018)
2. Hatem, G.M., Abdul Sadah, J.W., Saeed, T.R.: Comparative study of various CFAR algorithms for non-homogenous environments. *IOP Conf. Ser. Mater. Sci. Eng.* **433**, 012080 (2018)
3. Zhu, R., Wang, C., Jin, S., Zhang, J., Lu, M., Li, H.: Radar signal pre-sorting based on mean-shift algorithm. In: *Proceedings of the 2022 7th International Conference on Intelligent Computing and Signal Processing (ICSP) (2022)*. <https://doi.org/10.1109/icsp54964.2022.9778459>
4. Cornwell, T.J.: Multiscale CLEAN deconvolution of radio synthesis images. *IEEE J. Select. Top. Sig. Process.* **126**, 793–801 (2008)
5. Foreman, T.L.: Reinterpreting the CLEAN algorithm as an optimum detector. In: *Proceedings of the 2006 IEEE Conference on Radar*, pp. 769–775 (2006)
6. Foreman, T.L.: Application of the CLEAN detector to low signal to noise ratio targets. In: *Proceedings of the 2010 IEEE Radar Conference*, pp. 150–155 (2010)
7. Haliloglu, O., Yilmaz, A.O.: Successive target cancelation in pulse compression radars. In: *Proceedings of the 2007 IEEE Radar Conference*, pp. 885–890 (2007)
8. Kulpa, K.: The CLEAN type algorithms for radar signal processing. In: *Proceedings of the 2008 Microwaves, Radar and Remote Sensing Symposium Proceedings*, pp. 152–157 (2008)