# Sensitivity Optimisation of Real-Time Adaptive Range-Doppler Processing in FM Passive Radar

Joshua L. Sendall
CSIR
Pretoria, South Africa

Francois D. Maasdorp
CSIR
Pretoria, South Africa

Craig A. Tong
CSIR
Pretoria, South Africa

***Abstract*—This paper analyses the performance of adaptive processing algorithms, applied in the range-Doppler domain, for a frequency modulation (FM) based passive radar system, in terms of the processing efficiency, sensitivity and ultimately, detection performance. The algorithms are intended to be applied in a permanently installed passive-radar network and to this end, real-time signal processing considerations are also discussed. The processing methods are evaluated on a data set measured using the passive-radar network and the performance and shortcomings of each method are analysed. It is found that a hybrid Matching Pursuit (MP) method consistently produces the best results, while maintaining reasonable processing and memory requirements for the real-time processing chain. Using Orthogonal Matching Pursuit (OMP) resulted in highest overall probability of detection (Pd), 31% higher (SNR increased by 11.8 dB) than a conventional matched filter.**

***Index Terms*—passive radar, PCL, adaptive processing, range-Doppler processing, FM**

## I. Introduction

Interest in passive radar has been revived in recent years as a result of the congested radio spectrum [1], [2] and the potentially reduced cost of passive radar systems when compared to their active counterparts. However, passive radar is still a developing technology due to a lack of qualification and performance data. To this end, the Council for Scientific and Industrial Research (CSIR) is in the process of the developing a permanent installation frequency modulation (FM)-based passive radar test bed [3] to characterise, over a period of several years, the performance and suitability of passive radar for air traffic control.

The system is comprised of a network of passive radar receivers, located around O.R. Tambo International Airport in the metropolis area of the Gauteng Province in South Africa. The urban environment in which the system is located brings with it the challenges of a dense radio-spectrum environment with signals of interest ranging from below -120 dBm for smaller aircraft at long ranges to signals above -20 dBm for high powered transmitters near receivers in the network.

Passive radars, exploiting broadcast illuminators of opportunity (such at FM radio and digital television), experience direct path interference (DPI), which is the reception of the transmitted signal in the surveillance channel [4]. DPI is typically more than $90\,\mathrm{dB}$ above target returns [4]. Thus, particularly in FM-based passive radar, DPI is removed with an additional processing step known as DPI cancellation [5].

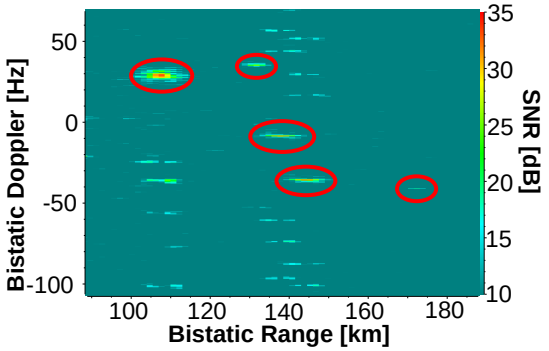


Fig. 1: False targets created by sidelobes shown in the amplitude range Doppler (ARD) space. True targets are circled in red. All peaks of false targets are at least 15 dB above the noise floor and will be declared as detections by a typical constant false-alarm rate (CFAR) detection scheme.

Large aircraft that are detected near to receivers create strong returns with sidelobe interference, which in turn masks other weaker targets [5], [6]. This ultimately decreases the sensitivity of the receiver. For extremely strong targets, sidelobes result in peaks large enough to be falsely declared as targets by detection algorithms. An example of this phenomena is shown in Fig. 1. These sidelobes resemble targets and create multiple false targets per real target in the detection stage. In a typical processing chain, the false targets are then propagated to the tracking filter and can cause the tracking filter to treat them as a swarm of targets. This is found to be more common as the sensitivity of the radar increases and the coherent processing gain increases. It is particularly aggravated when the coherent processing interval (CPI) is lengthened.

In order to deal with DPI, large targets and stationary clutter, the extensive cancellation algorithm (ECA) has been proposed [5] (an Orthogonal Matching Pursuit (OMP) like scheme), as well as various compressive sensing techniques [7]. Both techniques have high complexity and will not achieve real-time operation on current processors. As such, they may not be directly suitable for the requirements of real-time systems.

Adaptive range-Doppler processing is another contender to mitigate the detrimental effects of DPI, clutter and strong target returns. This paper analyses several adaptive range-Doppler processing schemes focussing on their implementation in real-time systems and their effectiveness on measured data. The

algorithms are analysed using FM passive radar data of aerial targets. The algorithms investigated are conjugate-gradient least squares (CGLS), Matching Pursuit (MP), OMP and a hybridised version of OMP termed Hybridised Matching Pursuit.

The paper proceeds as follows: Section II introduces the theory of adaptive range-Doppler map construction; Section III addresses the real-time implementation of the algorithms under investigation; Section IV reviews the performance of the algorithms on measured data and finally the conclusions are presented in Section V.

## II. Background to Adaptive Range-Doppler Map Construction Techniques

This section presents the mathematical background to each of the adaptive range-Doppler generation methods addressed in this paper.

### A. Matching Pursuit

MP [8] has been applied to FM passive radar to mitigate some of the detection constraints [6], and attempts to estimate the range-Doppler map as a sparse system. The algorithm is comparable to the CLEAN algorithm [9]. MP is a sub-optimal approximator but it can be desireable due to it's lower computational complexity and intuitive nature compared to other sparse-system estimators.

MP operates by defining a dictionary of possible vectors (or atoms), $\mathbf{A}$. In this case these atoms are time-frequency shifted vectors of the reference signal, which are used to create the ARD, or at least a section of the ARD. At each iteration ($k$) the dictionary is calculated and from it one or more atoms are selected to be representative of the sparse signal space $\mathbf{\Phi}_k$, where $\mathbf{\Phi}_k = [\phi_0\ \phi_1\ \cdots\ \phi_k]$. The dictionary is calculated using

$$\mathbf{D}_k = \mathbf{A}^H \mathbf{x}_{s(k)} \tag{1}$$

where $\mathbf{x}_{s(k)}$ is the surveillance signal after the $k$th atom has been added to $\mathbf{\Phi}_k$, and $\mathbf{D}_k$ are the atom weights for iteration $k$. The residual signal is then updated by,

$$\mathbf{x}_{s(k+1)} = \mathbf{x}_{s(k)} - w_k \phi_k \tag{2}$$

$w_k$ is defined as

$$w_k = \frac{\mathbf{x}_{s(k)} \phi_k^H}{|\phi_k|^2} \tag{3}$$

and $\phi_k$ is the time and frequency shifted version of the reference signal corresponding to the $k$th atom in $\mathbf{\Phi}_k$.

Calculating the dictionary correlation, (1), is the most computationally demanding aspect of the MP approach, and as such it is common place to select more than one atom in each iteration.The selection process can select the $P$ largest contributors [8] from the dictionary. Alternative schemes can be used to select the atoms such as CFAR [6]. It should be noted, however, that in cases where returns exhibit strong sidelobes in the dictionary space, a degradation in estimation performance can be expected by selecting false contributors [6]. The contribution of each atom is then iteratively removed starting from the largest atom in descending order.

### B. Orthogonal Matching Pursuit

Like the MP algorithm, OMP [10] removes each signal in the range-Doppler map sequentially. However, at each iteration the contribution of the additional atom ($a_k$) is solved in conjunction with all of the selected atoms. Thus, at each iteration the weights are updated by

$$\mathbf{x}_{s(k+1)} = \mathbf{x}_s - \mathbf{\Phi}_k^H \mathbf{w}_k \tag{4}$$

where $\mathbf{w}_k$ is a vector of the calculated weights for iteration $k$. The weights are calculated by solving

$$\mathbf{\Phi}_k \mathbf{w}_k = \mathbf{x}_s. \tag{5}$$

This sparse representation does not scale well when there are a high number of contributing atoms due to the complexity of (5) which must be solved at each iteration. Solving (5) has a complexity of $\mathcal{O}(k^2 M)$ when solved using QR factorisation [11], where $M$ is the length of $\phi_k$.

The first ECA [5] is essentially this algorithm expressed in the terms of passive radar. However due to the computational complexity of implementing it, an alternative method is presented, ECA-batches (ECA-B).

This alternative is able to reduce the computational and memory requirements of the algorithm by splitting $\mathbf{A}$ and $\mathbf{x}_{s(k)}$ into a number of temporal batches which can be processed in parallel. The batching has the added effect of widening the cancellation notch in the Doppler dimension [12], but can also cause false targets to appear on the ARD surface when a target's bistatic range rate is low [12].

### C. Conjugate gradient least-squares

As an alternative to the $l^1$-norm solvers, this paper proposes a method for real-time range-Doppler map construction using a least squares (LS) filter. A generic LS solver (such as QR factorisation) is currently impractical to implement, in terms of both computational load and memory constraints. As such, the method here leverages the CGLS solver. The CGLS filter is ideal for this case as it can provide a solution to badly conditioned and sparse matrices [13] (which $\mathbf{A}$ tends to be), can be initialised with an estimated solution, and has a significantly reduced memory footprint when compared to direct LS solvers [14].

The Conjugate gradient method has been explored for the purpose of direct-path and clutter cancellation [15]. However, in this implementation it is used not just to estimate the clutter and direct path, but the entire range-Doppler surface. The CGLS algorithm used in this paper is based on [16].

Theoretically CGLS is well suited to range-Doppler map construction for a number of reasons. However, of particular importance in this application, is the preservation of the sparsity of the solution. This is achieved by leveraging regularisation.

Within the CGLS algorithm, regularisation is performed implicitly [17]. In this case the number of iterations in the

algorithm controls the regularisation, which is increased as the number of iterations are decreased. This is due to the algorithm's property of tending to estimate the components with the highest Eigen values first [17]. Therefore, contrary to intuition, a better estimation of the signal space is not necessarily achieved by increasing the number of iterations.

## III. Real-time implementation strategies and stability

### A. Interior point optimisation

It is worth noting that all the algorithms discussed in this paper use the same $\mathbf{A}$ matrix, and that the largest computational load in many of the algorithms is contained in the two matrix-vector multiplications using $\mathbf{A}$. Namely

$$\mathbf{A}\mathbf{x} \tag{6}$$

and

$$\mathbf{A}^H\mathbf{y}, \tag{7}$$

where $\mathbf{A}$ is an $M \times N$ matrix with $N = N_D N_R$. $N_D$ is the number of Doppler bins considered, and $N_R$ is the number range bins. In cross-correlation the ARD is calculated using (7).

We restrict the definition of $\mathbf{A}$ such that the Doppler bins considered lie on the interval $f_m = mf_s/M$ where $m \in [-(M-1)/2, M/2]$, where $f_s$ is the sampling frequency, and the range bins $r_n = cn/f_s$ where $n \in \mathbb{Z}$ and $c$ is the speed of light. Essentially restricting the range and Doppler bins to their discretized sample points in both time and frequency. By doing this these two multiplications can be expressed as convolutions, and implemented using fast Fourier transforms (FFTs). Furthermore, both (6) and (7) can be calculated using convolution, which relaxes the memory requirements of the system as $\mathbf{A}$ is not explicitly stored. This is common practice in the calculation of the range-Doppler map when using cross-correlation.

Using conventional matrix multiplication (6) and (7) have a complexity of $\mathcal{O}(MN)$. By contrast, using frequency domain correlation implemented using an FFTs, (6) and (7) have a complexity of $\mathcal{O}(M \log_2(M) N_R)$. Assuming a typical FM radio based passive radar $M = 200 \times 10^3$, $N_D = 1600$, and $N_R = 300$, the convolution method requires approximately $90$ times fewer computations. With $M$ increased to $800 \times 10^3$ the difference decreases to approximately $82$ times.

### B. Matching Pursuit

Matching Pursuit has successfully been implemented in passive radar [6], along with strategies to implement it in real time. Key among these is selecting multiple targets in each iteration, as each iteration requires the calculation of (7). In this approach, the selected atoms' contributions are then removed using (2) in descending order of their magnitude.

While this strategy works well in general, it was found that the removal of the contribution from strong atoms can have a significant impact on other atoms. In some instances this occurred to the point where the resulting ARD had a higher noise floor compared to simple cross-correlation and the detection performance of the system was reduced. Therefore a stopping criteria was introduced to prevent degradation of the solution,

$$\left\|\mathbf{x}_{s(i+1)}\right\|_2^2 \leq \left\|\mathbf{x}_{s(i)}\right\|_2^2 \tag{8}$$

### C. Orthogonal Matching Pursuit

Orthogonal Matching Pursuit results in a more stable implementation than MP. However, it was found that even once they have been included in the cancellation mask, atoms can retain largest magnitude in $\mathbf{D}_k$. This indicates the ill-conditioned nature of the measured data, and the limits caused by numerical accuracy. Nonetheless, this can be avoided by a simple check for uniqueness of the selected atoms in $\mathbf{D}_k$ before searching for new components.

A further limitation in OMP is the increased computational load resulting from the need to solve (5) and compute (4) at each iteration. The computational burden can be improved by exploiting the factorisation of the previously defined $\mathbf{\Phi}_{k-1}$ [18]. This can be achieved defining

$$\mathbf{C}_k = \mathbf{\Phi}_k^H \mathbf{\Phi}_k \tag{9}$$

and

$$\mathbf{d}_k = \mathbf{\Phi}_k^H \mathbf{x}_s \tag{10}$$

each OMP iteration can be turned into an update of $\mathbf{C}$ and its Cholesky decomposition [18]. Therefore, updating the weights and the residual in the $k$th iteration has a complexity of $\mathcal{O}(Mk + k^2)$ [18]. This is as opposed to the full solution of (5) and (4) which has a complexity of $\mathcal{O}(Mk^2)$. One drawback with this implementation is the reduction in numerical stability, which results from the Cholesky decomposition and the lack of column pivoting.

The use of the Cholesky decomposition, although computationally efficient, is not ideal as it is only able to safely decompose positive-definite matrices. Due to its Hermitian structure, $\mathbf{C}$ is only guaranteed to be semi-positive definite. This could be solved by resorting to another matrix decomposition, such as LDLT or LU. However, in these algorithms, numerical stability is achieved by performing column pivoting [11]. The above mentioned update strategy impedes this as each iteration only introduces a single row to $\mathbf{C}$. Therefore, if one of the columns is poorly conditioned, it is not pivoted.

A level of tolerance can be re-introduced, through two means. Firstly, by selecting more than one atom for each iteration of the solution, column pivoting can be applied to the new subset. Secondly, numerically unstable columns are pivoted to be lower in the decomposition. Therefore, recomputing the final few columns of the previous iteration can add additional numerical stability. However, this does require the re-computation of those rows.

On measured data sets, poorly conditioned $\mathbf{C}_k$ matrices were encountered, and more frequently as the number of atoms selected at each iteration increased. The OMP algorithm was, therefore implemented to recalculate a single row, and apply column pivoting to each subset. $\mathbf{C}_k$ was decomposed using

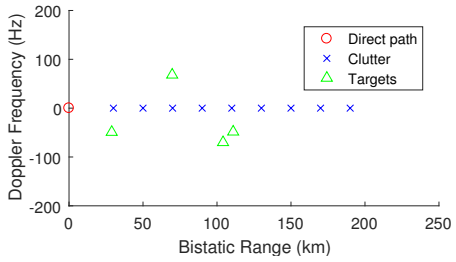


Fig. 2: Idealised Range-Doppler map for a passive bistatic radar.

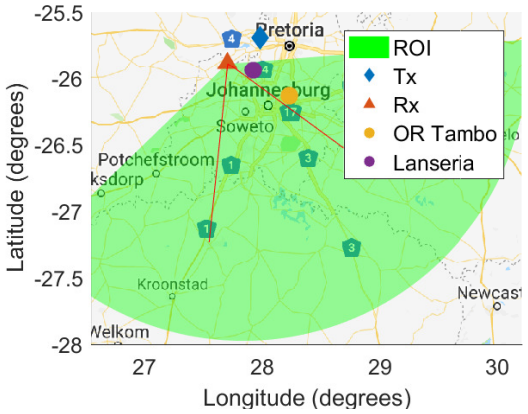


Fig. 3: Configuration of the experiment.(Map sourced from AfriGIS (Pty) Ltd, Google)

LU factorisation.

The update in each iteration for exceedingly sparse signals, i.e. when $K << M$ is dominated by the update of $\mathbf{C}_k$ rather than the factorisation of thereof. This is usually the case in FM passive radar. In order to help alleviate the computational burden, $\mathbf{C}_k$ can be updated as a batch of components rather than a single one. This is achieved by computing the lowest $\Delta k$ rows of $\mathbf{C}_k$ through

$$\mathbf{\Phi}_k^H \mathbf{\Phi}_{\Delta k}, \tag{11}$$

where $\mathbf{\Phi}_{\Delta k}$ consists of the columns of $\mathbf{A}$ added to $\mathbf{\Phi}_{k-1}$ to form $\mathbf{\Phi}_k$. The rest of $\mathbf{C}_k$ can then be filled in using $\mathbf{C}_k$'s Hermitian symmetry. This method does not decrease the complexity of the update, but uses a matrix-matrix multiplication, instead of multiple matrix-vector multiplications. On processors such as graphics processing units (GPUs), matrix-matrix operations are significantly faster than matrix-vector operations, due to more efficient data accessing.

### D. Hybrid Matching Pursuit

The computational requirement of OMP increases as the number of iterations cubed. In order to relieve this, a number of known elements (i.e. the direct path and clutter) can be estimated and removed using an adaptive filter. These components can be efficiently estimated using a CGLS filter [15] or using a Toeplitz solver. These components can then be removed from $\mathbf{x}_s$, and OMP applied to the residual signal. This initial step emulates an MP update, but with a vector of weights instead of a single weight. This method is denoted as Hybrid Matching Pursuit (HMP) for the duration of this paper.

The energy received at a passive radar node is typically composed of 3 main components shown in Fig. 2 the direct path, static multi-path (reflections from large static objects), and reflections from moving platforms. The direct path and static multi-path contain the vast majority of the signal's energy. As illustrated in in Fig. 2, for a static configuration the direct path and clutter reside on the zero-Doppler line. Due to this, the clutter and direct path regions are selected for the initial weight calculation.This method not only allows the initial number of iterations to be efficiently estimated, but it also reduces the number of iterations required in the OMP step. The price paid for this is a sub-optimal solution, as orthogonality of all the components is lost. Another practical constraint that can be alleviated using HMP is that a smaller $\mathbf{\Phi}$ matrix can be stored, which is the main contributor to the memory requirements for OMP.

### E. Conjugate-gradient least squares

Two key properties of CGLS can be exploited to reduce its computational requirements.

Firstly, the CGLS can take advantage of the optimisation discussed in Section III-A. The reduction in computational complexity, however, is still not sufficient to support real-time operation. The major limitation is that the solution requires many iterations to converge to the point where targets are revealed.

Secondly, the CGLS algorithm can be initialised with an estimate of the range-Doppler map. By intelligently selecting the starting weights, the number of iterations necessary to reveal targets can be significantly reduced. This is done in a similar manner as the previous section. The weights on the zero-Doppler range line are estimated. Then, by using those calculated weights to initialise the algorithm, a range-Doppler map is generated with comparatively few CGLS iterations.

It should be noted though, that the range-Doppler map is a sparse signal space in typical FM passive radar configurations. Therefore, regularisation is required to achieve a suitable solution.

## IV. Performance on measured/recorded data

The various methods discussed in Section III were applied to a set of data recorded from the long term-passive radar test facility at the Council for Scientific and Industrial Research (CSIR). The data is a $40\,\mathrm{h}$ recording from a single passive radar node with the geometry shown in Fig. 3. The figure also shows the region of interest (ROI) for the experiment, along with the transmitter site, receiver site, the $60^\circ$ beam-width, and two local airports (O.R. Tambo International Airport, and Lanseria International Airport).

The experiment was conducted using a radio station broadcasting at $94.2\,\mathrm{MHz}$ and had a CPI of $4\,\mathrm{s}$. Two surveillance

channels were used and processed and then non-coherently integrated. The signals were passed through a fixed-tuned radio-frequency (RF) front-end [19], and sampled using an Ettus X310 software-defined radio (SDR). Detection was performed using a censored cell-averaging CFAR in the Doppler dimension, with 7 gaurd cells and 5 reference cells on each side of the cell under test. The baseline between the transmitter and receiver was $35.4\,\mathrm{km}$. The ROI was the area bounded by a bistatic range of $450\,\mathrm{km}$ and an azimuth $80°$ to either side of the surveillance antenna's boresight. For target ground truth data a Kinetic Avionics SBS-3 (a commercial ADS-B receiver) was used to record the position and velocity of aircraft in the vicinity.

### A. Comparison of Algorithms

All of the algorithms are configured to process 320 range bins (although only 276 of these are within the ROI), and 1601 Doppler bins. The configuration of each algorithm reflects the limit at which it could be run in real-time, or for which optimal detection performance was achieved.

For reference, cross-correlation (denoted by XCorr) is included in the set of algorithms which are compared. The cancellation filter necessary for cross-correlation was implemented using a batched Toeplitz least-squares algorithm [20], which removes energy along the zero-Doppler line. The batch length was $200 \times 10^3$ samples, and the cancellation mask extended from range bin -4 to 255. The negative range shift included in the mask was found to provide a significant improvement in cancellation performance, with little additional processing load.

The implementation of HMP used the same Toeplitz filter as cross-correlation coupled to a OMP estimator for 350 iterations. The implementation of OMP used a batch-size of 32 and 400 iterations. The implementation of MP used 200 iterations in batches of 32. The CGLS algorithm was initialised with the weights from the cancellation filter, estimated the range-Doppler space from range bin -4 to 250, and used 13 iterations.

### B. Probability of detection performance

The probability of detection (Pd) was evaluated by interpolating the position and velocity of each of the reference tracks (from ADS-B) to the start of each radar CPI. These were then converted to bistatic range and velocity. From these reference points any detections with a corresponding time and a range-velocity gate of $2\,\mathrm{km}$ and $10\,\mathrm{m\,s^{-1}}$ were considered correct detections. False alarms increase the reported Pd, but with a probability of false alarm calculated to be $1 \times 10^{-5}$ (theoretically), the difference is negligible when compared to the Pd achieved. The mean execution time and overall Pd for each of the algorithms is shown in Table I.

From Table I it can be seen that cross correlation performs sub-optimally when compared to the OMP estimators. OMP achieved the highest Pd, reducing the the number of missed detections by $51\,\%$ when compared to cross correlation. HMP performed the next best, requiring less processing time and a comparable Pd when compared to OMP. Therefore, HMP is an option for systems with tighter processing limitations. MP required less processing time, but resulted in an even lower Pd. CGLS performed marginally well, but required more processing than any of the other algorithms.

TABLE I: Detection performance of adaptive range-Doppler processing methods.

| | Mean execution time | Pd | SNR |
|---|---|---|---|
| XCorr | $179.5\,\mathrm{ms}$ | 0.612 | $20.7\,\mathrm{dB}$ |
| MP | $828.48\,\mathrm{ms}$ | 0.652 | $21.2\,\mathrm{dB}$ |
| HMP | $2102.459\,\mathrm{ms}$ | 0.780 | $29.3\,\mathrm{dB}$ |
| OMP | $2929.742\,\mathrm{ms}$ | 0.801 | $32.5\,\mathrm{dB}$ |
| CGLS | $3667.497\,\mathrm{ms}$ | 0.700 | $24.8\,\mathrm{dB}$ |

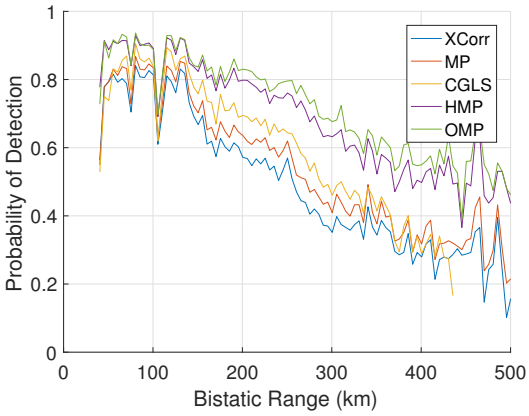


Fig. 4: Probability of detection versus bistatic range.

Without the batch update of $\mathbf{C}_k$ discussed in Section III-C, the OMP implementation ran with a mean processing interval of $12.34\,\mathrm{s}$. This demonstrates that for rectangular and highly sparse problems, the update of $\mathbf{C}_k$ can be the processing bottle neck in an OMP operation.

The Pd as a function of bistatic range is shown in Fig. 4 and exhibits a trend that the probability of detection decreases as the range increases. The relative Pd between the OMP and HMP is negligible at ranges $<160\,\mathrm{km}$, but at longer ranges begins to increase to approximately $9.7\,\%$ at $390\,\mathrm{km}$. The difference between OMP and cross-correlation increases from approximately $37\,\%$ to approximately $83\,\%$ after $260\,\mathrm{km}$. This demonstrates how the long range (and thus weaker) targets are masked by stronger returns, and illustrates the successful recovery of these targets using a sparse estimator. Fig. 4 exhibits a decrease in the Pd at approximately $110\,\mathrm{km}$. This corresponds to the largest airport in the region (O.R. Tambo). Here the majority of the aircraft are either taking-off or landing, and because of their lower altitude, the radar experiences difficulty in detecting the aircraft.

### C. Signal-to-noise ratio analysis

The signal-to-noise ratio (SNR) over bistatic range ($R_b$) is shown in Fig. 5 for the associated detections. The SNR

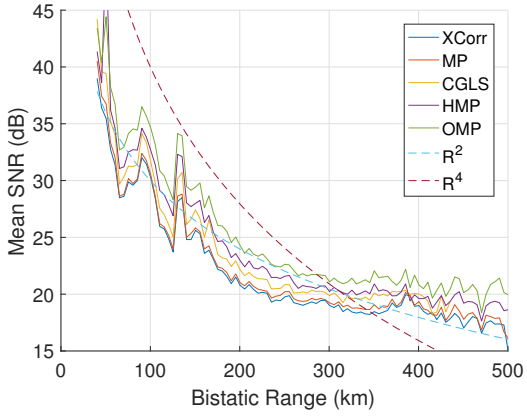


Fig. 5: SNR versus bistatic range.

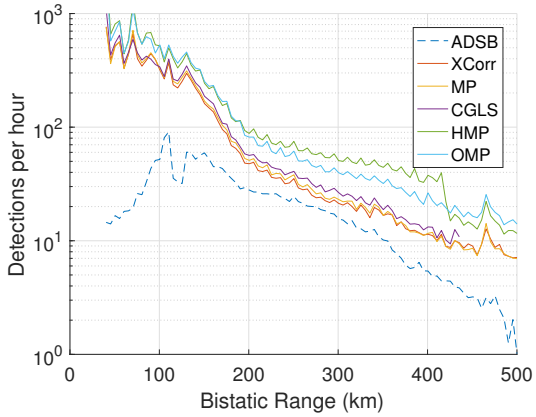


Fig. 6: Detection rate over bistatic range.

was calculated by comparing the target return to the median magnitude of the ARD space for that CPI.

Bistatic range is defined as $R_{Tx} + R_{Rx}$ where $R_{Tx}$ is the range from the transmitter and $R_{Rx}$ is the range to the receiver. The SNR in Fig. 5 decreases approximately with an $R_b^2$ relationship. This relationship is contrary to the SNR predicted by the free-space model, where SNR $\propto R_{Tx}^2 R_{Rx}^2$. While this is unexpected, due to the varying geometries and radar cross sections (RCSs) between targets, the relationship between SNR and bistatic range is not well constrained in this experiment.

The detection rate of the ADS-B receiver and the radar are shown in Fig. 6. Here it can be seen that the detection rate for the radar is much higher at near ranges, but as the range increases, radar's detection rate drops and matches the slope of the ADS-B receiver's detection rate. Interestingly, the radar displays a higher detection rate than the ADS-B at all ranges.

## V. Conclusion

Adaptive ARD generation methods were explored and key implementation constraints were highlighted, allowing for real-time implementation. The adaptive ARD methods were applied to measured data, and it was shown that a number of them were effective in overcoming the short-comings of conventional cross-correlation. It was found that OMP was the most suitable algorithm for the task, achieving the highest Pd, an improvement of $31\,\%$, and an increase in SNR of $11.8\,\mathrm{dB}$. A hybrid matching pursuit algorithm is presented, and is shown to be a suitable candidate for systems with tighter computation constraints. Hybrid matching pursuit displayed a $2.6\,\%$ lower Pd, but required $28.2\,\%$ less processing.

## References


[1] H. Griffiths, "Where has all the spectrum gone?" in *2013 International Conference on Radar*, Sept 2013, pp. 1–5.

[2] D. for Culture Media and Sport, "Enabling UK growth - releasing public spectrum," Tech. Rep., March 2014.

[3] F. D. V. Maasdorp, C. A. Tong, A. Lysko, M. R. Inggs, and D. W. O'Hagan, "The design and development of a FM band passive radar test-bed for long term qualification testing," in *2017 IEEE Radar Conference (RadarConf)*, May 2017, pp. 1515–1520.

[4] P. Howland, D. Maksimiuk, and G. Reitsma, "FM radio based bistatic radar," *IEE Radar, Sonar and Navigation*, vol. 152, pp. 107–115, 2005.

[5] F. Colone, R. Cardinali, P. Lombardo, O. Crognale, A. Cosmi, A. Lauri, and T. Bucciarelli, "Space-time constant modulus algorithm for multipath removal on the reference signal exploited by passive bistatic radar," *IET Radar, Sonar and Navigation*, vol. 3, no. 3, pp. 253–264, June 2009.

[6] K. S. Kulpa and Z. Czekala, "Masking effect and its removal in PCL radar," *IEE Proceedings - Radar, Sonar and Navigation*, vol. 152, no. 3, pp. 174–178, June 2005.

[7] C. R. Berger, B. Demissie, J. Heckenbach, P. Willett, and S. Zhou, "Signal processing for passive radar using OFDM waveforms," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 1, pp. 226–238, Feb 2010.

[8] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, Dec 1993.

[9] J. A. Högbom, "Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines," , vol. 15, p. 417, Jun. 1974.

[10] T. T. Cai and L. Wang, "Orthogonal matching pursuit for sparse signal recovery with noise," *IEEE Transactions on Information Theory*, vol. 57, no. 7, pp. 4680–4688, July 2011.

[11] L. Trefethen and D. Bau, *Numerical Linear Algebra*. Society for Industrial and Applied Mathematics, 1997.

[12] F. Colone, C. Palmarini, T. Martelli, and E. Tilli, "Sliding extensive cancellation algorithm for disturbance removal in passive radar," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 3, pp. 1309–1326, June 2016.

[13] Z. Zlatev and H. Nielsen, "Solving large and sparse linear least-squares problems by conjugate gradient algorithms," *Computers & Mathematics with Applications*, vol. 15, no. 3, pp. 185–202, 1988. [Online]. Available: http://www.sciencedirect.com/science/article/pii/0898122188901708

[14] M. Inggs, C. Tong, R. Nadjiasngar, G. Lange, A. Mishra, and F. Maasdorp, "Planning and design phases of a commensal radar system in the FM broadcast band," *IEEE Aerospace and Electronic Systems Magazine*, vol. 29, no. 7, pp. 50–63, July 2014.

[15] J. Palmer and S. Searle, "Evaluation of adaptive filter algorithms for clutter cancellation in passive bistatic radar," in *IEEE Radar Conference (RADAR)*, May 2012, pp. 493–498.

[16] C. Tong, "A scalable real-time processing chain for radar exploiting illuminators of opportunity," Ph.D. dissertation, University of Cape Town, 2014.

[17] P. C. Hansen, "Regularization tools version 4.0 for Matlab 7.3," *Numerical Algorithms*, vol. 46, no. 2, pp. 189–194, Oct 2007. [Online]. Available: https://doi.org/10.1007/s11075-007-9136-9

[18] B. L. Sturm and M. G. Christensen, "Comparison of orthogonal matching pursuit implementations," in *2012 Proceedings of the 20th European Signal Processing Conference (EUSIPCO)*, Aug 2012, pp. 220–224.

[19] M. Inggs, R. Geschke, J. Coetser, and D. O'Hagan, "High sensitivity fixed tuned direct conversion receiver for fm band commensal radar," in *2015 IEEE Radar Conference*, Oct 2015, pp. 174–179.

[20] J. Sendall and W. du Plessis, "Optimised wiener filtering in overdetermined systems," in *IEEE 2018 International Conference on Radar*, Aug 2018.