Partially Constrained Adaptive Beamforming for Super-Resolution at Low SNR

Erik Hornberger[1], Shannon D. Blunt[1], Thomas Higgins[2]

[1] *Radar Systems & Remote Sensing Lab (RSL), University of Kansas, Lawrence KS, USA*
[2] *Radar Division, US Naval Research Laboratory, Washington DC, USA*

Abstract—**The reiterative super-resolution (RISR) algorithm was previously developed to enable adaptive beamforming with as few as one time snapshot, is robust to temporally correlated signals, and accounts for array calibration errors. Here a gain-constrained version (denoted GC-RISR) is derived followed by a partially-constrained version (PC-RISR). It is shown that an interesting trait of the latter is spatial super-resolution at SNR values lower than is typical for adaptive beamforming techniques as a trade-off for requiring more iterations to converge. The PC-RISR formulation is controlled by a selectable parameter that serves a role similar to that of an adaptive step-size which balances between convergence speed and accuracy.**

I. Introduction

Direction of Arrival (DOA) estimation is the process of determining the spatial angle from which an incident signal impinges on a sensor array. In most applications, several signals may be present simultaneously, some of which may possess a high degree of temporal correlation (e.g. multipath). Classical DOA estimation techniques such as Capon beamforming, MUSIC, and ESPRIT [1] rely on the determination of a sample covariance matrix (SCM) from a set of spatial snapshots. This SCM estimate may use forward / backward averaging [2] and/or spatial smoothing [3] to address space-time coupling of temporally correlated signals.

The reiterative super-resolution (RISR) algorithm [4,5] was developed from the reiterative minimum mean-square error (RMMSE) framework [6] to provide accurate DOA regardless of signal correlation by employing a structured covariance matrix that avoids calculation of the SCM, thus making it suitable for applications requiring low sample support due to non-stationarity. Further, RISR can incorporate unknown calibration errors based on known tolerances. A modified version was subsequently demonstrated on measured magnetoencephalography (MEG) data to facilitate high resolution functional brain imaging [7].

Paralleling the formulation in [8], here a gain-constrained version of RISR is derived. It observed that the resulting GC-RISR realizes resolution that is somewhat better than standard non-adaptive beamforming, but is inferior to RISR as well as previous SCM-based techniques in this regard. However, when RISR and GC-RISR are combined, a new partially-constrained version (PC-RISR) enables spatial super-resolution at lower SNR than is typical for SCM-based methods (generally at least 10 dB, particularly for low sample support [1]). Simulations to characterize performance show that it is possible to achieve super-resolution at single digit SNR values (inclusive of beamforming gain).

II. Re-Iterative Super-Resolution (RISR)

Consider a uniform linear array with N elements and half wavelength spacing (noting arbitrary arrays are also feasible [5]). For $K < N$ incident signals satisfying the narrowband assumption, and with calibration and mutual coupling effects included to the degree known, the received signal for the ℓth snapshot can be expressed as the $N \times 1$ complex vector

$$\begin{aligned}\mathbf{y}(\ell) &\triangleq [\mathbf{S}\mathbf{x}(\ell)] \odot \mathbf{z} + \mathbf{v}(\ell) \\ &= \mathbf{S}\mathbf{x}(\ell) + \mathbf{v}(\ell) + \mathbf{v}_{\mathrm{z}}(\ell),\end{aligned} \tag{1}$$

where $\mathbf{x}(\ell)$ is an $M \times 1$ vector, the $M \gg N$ elements of which are the complex received amplitudes as a function of spatial angle. The vector $\mathbf{v}(\ell)$ is additive noise of arbitrary distribution, and $\mathbf{S}$ is an $N \times M$ matrix of spatial steering vectors parameterized over the array manifold. The term $\odot$ is the Hadamard product and the $N \times 1$ vector $\mathbf{z}$ accounts for array modelling errors with the nth element represented as

$$z_n = \left[1 + \Delta_{\mathrm{a},n}\right] \mathrm{e}^{\mathrm{j}\Delta_{\varphi,n}}, \tag{2}$$

where $\Delta_{\mathrm{a},n}$ is the error in amplitude and $\Delta_{\varphi,n}$ is the phase error, both with arbitrary distributions (though errors are assumed to be relatively small). The second line in (1) arises from the model in (2) where $\mathbf{v}_{\mathrm{z}}(\ell) = (\mathbf{z} - \mathbf{1}_{N\times 1}) \odot [\mathbf{S}\mathbf{x}(\ell)]$.

It is shown in [5] that minimizing the MMSE cost function

$$J = \mathbb{E}\{\|\mathbf{x}(\ell) - \mathbf{W}^H(\ell)\, \mathbf{y}(\ell)\|^2\} \tag{3}$$

yields the $N \times M$ filter bank $\mathbf{W}(\ell)$, having the mth column

$$\mathbf{w}_m(\ell) = \mathbf{P}_{m,m}(\ell)\left(\mathbf{S}\mathbf{P}(\ell)\mathbf{S}^H + \mathbf{R} + \mathbf{R}_{\mathrm{z}}(\ell)\right)^{-1} \mathbf{s}_m. \tag{4}$$

The $M \times M$ diagonal matrix $\mathbf{P}(\ell)$ contains the incident powers for the M spatial directions for the ℓth snapshot, the $N \times N$ matrix $\mathbf{R}$ is the noise covariance, and the $N \times N$ matrix $\mathbf{R}_{\mathrm{z}}(\ell)$ is the calibration error covariance defined as [5]

$$\mathbf{R}_{\mathrm{z}}(\ell) = \mathbf{I}\sigma_{\mathrm{z}}^2 \odot [\mathbf{S}\mathbf{P}(\ell)\mathbf{S}^H], \tag{5}$$

for $\mathbf{I}$ an $N \times N$ identity matrix and σ_{z}^2 the calibration error variance (which may be known based on design tolerances). Note that (5) is dependent on the incident signals. The terms $\mathbf{s}_m$ and $\mathbf{P}_{m,m}(\ell)$ in (4) specify the mth column of $\mathbf{S}$ and the mth diagonal element of $\mathbf{P}(\ell)$, respectively.

Clearly the values in $\mathbf{P}(\ell)$ are not known and must be estimated. An initial estimate is obtained by applying the standard (non-adaptive) beamformer as [5]

$$\hat{\mathbf{x}}(\ell) = \mathbf{S}^H\mathbf{y}(\ell), \tag{6}$$

This work was supported in part by the Radar Division of the US Naval Research Laboratory and the University of Kansas Self Fellow program.

from which the matrix of power estimates is determined via

$$\hat{\mathbf{P}}(\ell) = [\hat{\mathbf{x}}(\ell)\,\hat{\mathbf{x}}^H(\ell)] \odot \mathbf{I}\,. \quad (7)$$

Using (4), the adaptive filter bank can be computed and subsequently applied to revise the estimates of

$$\hat{\mathbf{x}}(\ell) = \mathbf{W}^H(\ell)\,\mathbf{y}(\ell) \quad (8)$$

that are subsequently used to update (7). This process of applying (4), (8), and (7) is repeated until acceptable convergence is achieved. It is also discussed in [5] how the signal estimates over multiple snapshots can be combined non-coherently within RISR, which serves to average out variations to further enhance super-resolution capability.

III. Incorporating a Gain Constraint

In [8] a gain constraint was introduced for RMMSE-based adaptive pulse compression of radar as a way to improve robustness to mismatch effects. For RISR, which relies on the same RMMSE framework, such mismatch effects (due to imperfect calibration and mutual coupling) translate into small spurious peaks. The calibration error term of (5) addresses this condition to some degree. However, it is worth considering how further robustness may be achieved.

A unity gain constraint can be readily incorporated by modifying the cost function of (3) for the mth column of $\mathbf{W}(\ell)$ as

$$J_m = \mathbb{E}\{\|\mathrm{x}_m(\ell) - \mathbf{w}_m^H(\ell)\,\mathbf{y}(\ell)\|^2\} + \lambda^*(\mathbf{w}_m^H(\ell)\,\mathbf{s}_m - 1) \quad (9)$$

where λ is the Lagrange multiplier. It is straightforward to show that, like [8], the gain-constrained RISR (GC-RISR) takes the form

$$\mathbf{w}_{\mathrm{GC},m}(\ell) = \left(\frac{1}{\mathbf{s}_m^H\,\mathbf{D}(\ell)\,\mathbf{s}_m}\right)\mathbf{D}(\ell)\,\mathbf{s}_m \quad (10)$$

in which

$$\mathbf{D}(\ell) = \left(\mathbf{S}\mathbf{P}(\ell)\mathbf{S}^H + \mathbf{R} + \mathbf{R}_{\mathrm{z}}(\ell)\right)^{-1}. \quad (11)$$

The original RISR formulation from (4) could likewise be expressed using (11) as

$$\mathbf{w}_m(\ell) = \mathbf{P}_{m,m}(\ell)\,\mathbf{D}(\ell)\,\mathbf{s}_m\,. \quad (12)$$

Figure 1 illustrates the distinction between RISR and GC-RISR for a uniform linear array of N = 10 antenna elements (so Rayleigh resolution of 36° in electrical angle) and for the case of two sources (of SNR = 30 dB after array gain) separated by an electrical angle of 20°. It is observed that GC-RISR still separates the two signals (unlike the standard beamformer of (6)), though the rounded peaks so often associated with a gain constraint clearly imply less super-resolution capability than RISR. That said, the flattened response exhibited by GC-RISR further away from the two signals also provides a more physically meaningful noise floor for subsequent detection than is achieved by RISR, which tends to "over-suppress" the spatial directions that are not present. In fact it is this very attribute that prevents identification of low SNR signals that we shall remedy here.

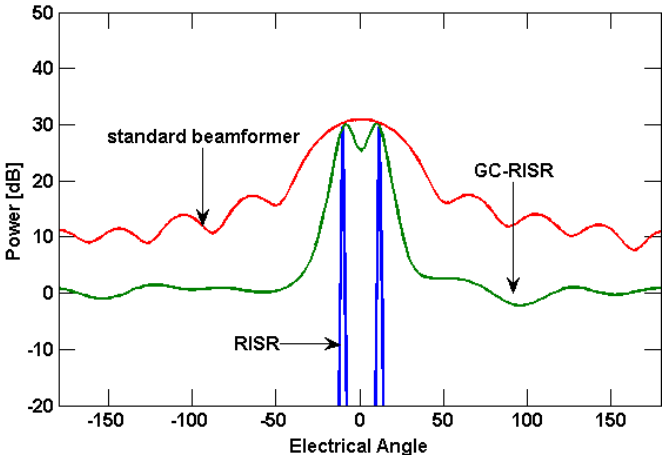

Fig. 1. Comparison of RISR, GC-RISR, and standard beamformer for two sources with SNR of 30dB and located at $\pm 10^{\circ}$.

IV. Partial Gain Constraint

The responses for RISR and GC-RISR as demonstrated in Fig. 1 can be viewed as the two extremes that could be achieved for this manner of DOA estimation. Examination of the filter implementations in (10) and (12) reveal that, for the mth filter, these formulations only differ by an adaptive scale factor. Given the potential for significant dynamic range among received signals, a prospective approach to facilitate a selectable "dial" between these two extremes is

$$\mathbf{w}_{\mathrm{PC},m}(\ell) = \left[\left(\frac{1}{\mathbf{s}_m^H\,\mathbf{D}(\ell)\,\mathbf{s}_m}\right)^{\alpha}\left(\mathbf{P}_{m,m}(\ell)\right)^{1-\alpha}\right]\mathbf{D}(\ell)\,\mathbf{s}_m \quad (13)$$

where the exponent $0 \le \alpha \le 1$ is a weighting factor that can be used to balance between RISR ($\alpha = 0$) and GC-RISR ($\alpha = 1$). We shall refer to this exponentially weighted version as partially-constrained RISR (PC-RISR). The value of α controls the convergence speed and super-resolution capability of PC-RISR. While (13) is a rather simple modification, the prospective benefit it provides is quite pronounced. Table I illustrates the general behavior of PC-RISR as it has been observed thus far.

TABLE I
Behavior of PC-RISR for Different Values of α

α =	0– 0.35	0.35 – 0.45	0.45 – 0.5
Required Iterations	10 ~ 20	20 ~ 50	50 ~ 300
Utility	Suitable for mid-high SNR	Suitable for low-mid SNR	Limit of low SNR performance
Behavior	Low SNR signals are suppressed	Modest SNR signals observed	Low SNR signals observed after long convergence

For values of α below 0.5, PC-RISR converges to point solutions such as that of RISR in Fig. 1. As α increases, the speed of convergence decreases (akin to a step-size parameter) and in turn the ability to separate closely spaced signals is enhanced. Based on simulation trials, for α up to about 0.4, the ability of RISR to operate in the low SNR regime is

significantly enhanced with little risk of spurious peaks. Setting α between 0.4 and 0.5 yields the best probability of separating closely spaced signals, but can produce small spurious peaks above the noise at low sample support. PC-RISR typically converges in 30 iterations, with diminishing returns in signal separability for further iterations as α approaches 0.5.

If α is above 0.5, PC-RISR does not converge to point solutions, but instead resembles GC-RISR from Fig. 1, albeit with narrower peaks and deeper nulls. In these cases, the converged results may exhibit loss in the peak values relative to RISR ($\alpha = 0$) or GC-RISR ($\alpha = 1$) (see Fig. 2).

Figure 2 shows how, for $\alpha = 0.35$, PC-RISR converges iteratively to point solutions while suppressing the response from directions in which no signal is incident. For this same 10 dB two-signal scenario and 30 iterations, Fig. 3 illustrates how different values of α affect the final estimate. Note that full convergence has been obtained for the $\alpha = 0.6$ and $\alpha = 0.75$ cases, which exhibit SNR loss, while the $\alpha = 0.45$ case would require many more iterations to converge to the same point solutions obtained in the $\alpha = 0.35$ case.

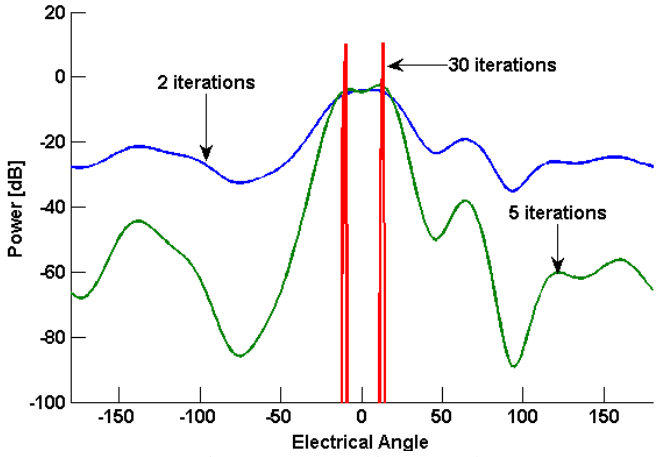

Fig. 2. Illustration of the convergence of PC-RISR for α = 0.35 and two signals of SNR = 10 dB (after array gain) and 20˚ separation.

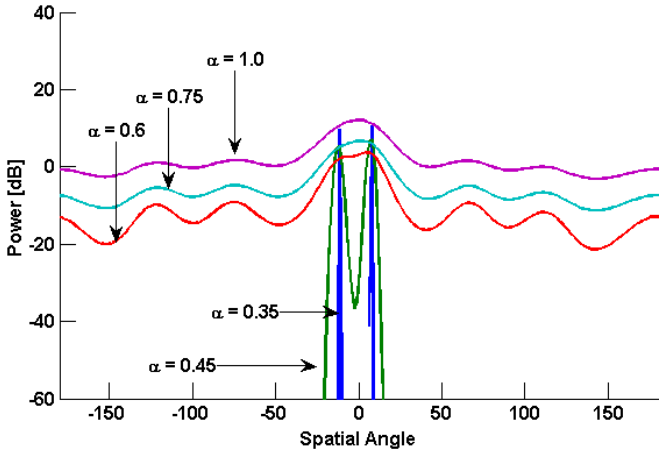

Fig. 3. Example of PC-RISR convergence as a function of α for two signals with SNR = 10 dB and 20˚ separation after 30 iterations. The $\alpha = 0.45$ case would converge to point solutions for higher iterations.

Figure 4 provides an anecdotal result for two signals with SNR = 15 dB and 20° separation for $N = 10$ antenna elements and 10 snapshots. After 30 iterations, PC-RISR ($\alpha = 0.4$) can separate signals that RISR [5] and MUSIC [1] cannot.

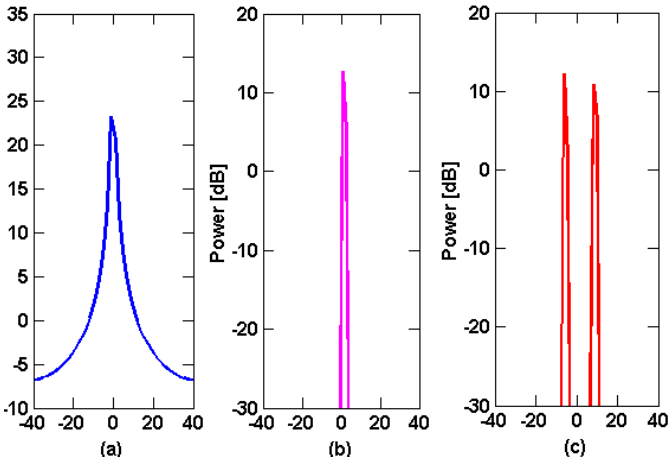

Fig. 4. Anecdotal comparison showing PC-RISR can separate signals when RISR and MUSIC cannot. Two signals at -10° and 10° with SNR = 15dB, N = 10, and 10 snapshots. (a) MUSIC, (b) RISR, (c) PC-RISR with $\alpha = 0.40$

V. Monte Carlo Analysis

Consider $N = 10$ antenna elements in a uniform linear array with half-wavelength spacing (Rayleigh resolution = 36°). All trials are based on $L = 10$ independent snapshots, which, for RISR and its variants, are combined non-coherently as described in [5]. Different values of α are considered and the results compared with root-MUSIC [9] as a benchmark.

Stated SNR values include the array gain of N = 10. Noise is modeled as additive white Gaussian. Incident signals have uniformly distributed random phase and constant amplitude. The array calibration terms $\Delta_{a,n}$ and $\Delta_{\varphi,n}$ from (2) are each uniformly distributed such that the gain and phase error of each antenna element may be as much as 1% of the expected value. The resulting overall calibration error variance is $\sigma_z^2 = 1.5 \times 10^{-3}$.

For each trial, 500 Monte Carlo runs were performed with two signals present. The first arrives from electrical angle $\varphi_1 = 0^{\circ}$ (antenna boresight) and the other is placed at some predetermined offset φ_2. For these results both signals have equal power with random phases and are uncorrelated in time.

A simple criterion is used to decide if signals are separable for the RISR-based methods. Once convergence is halted, all values of the spatial power estimate below the noise floor are set equal to the noise floor. If precisely two peaks remain above the noise floor, and both are within one-half beamwidth $(2\pi/N)$ of their corresponding true DOA values φ_1 and φ_2, then they are deemed correctly separated.

Root-MUSIC was implemented using forward-backward averaging [2] and the Bayesian Information Criteria (BIC) [10] was used to estimate the dimensionality of the signal subspace. For root-MUSIC the signals are deemed separable if they occur within one-half beamwidth of their true DOA and the BIC order estimate equals 2.

Figure 5 depicts the probability of separation of the two signals for an SNR of 5 dB (includes array gain). Here it is observed that the SNR is simply too low for root-MUSIC, yet after 30 iterations PC-RISR ($\alpha = 0.4$) is able to separate the signals at about 0.67 the Rayleigh resolution ($\sim 24^\circ$ separation) for 50% of the trials.

The SNR = 10 dB case is shown in Fig. 6, where root-MUSIC outperforms PC-RISR ($\alpha = 0.2$), which behaves more like RISR [5], while PC-RISR ($\alpha = 0.4$) now separates signals about 0.28 the Rayleigh resolution ($\sim 10^\circ$ separation) in 50% of the trials. In both of these cases, RISR [5] and GC-RISR do not separate the signals due to the low SNR.

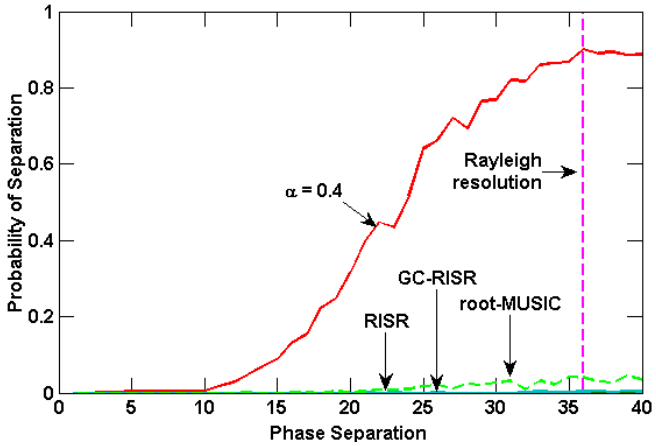

Fig. 5. Probability of separation as a function of signal separation for SNR of 5dB for PC-RISR (α = 0.4) and root-MUSIC.

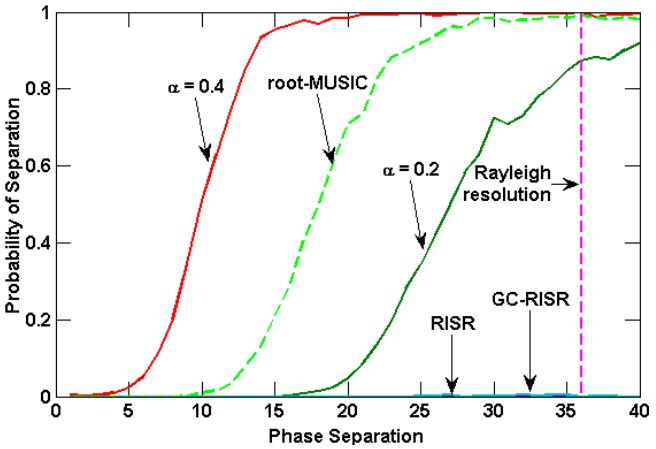

Fig. 6. Probability of separation as a function of signal separation for SNR of 10dB for root-MUSIC, PC-RISR (α = 0.2) and PC-RISR (α = 0.4).

Finally, Fig. 7 illustrates the relationship between α and the number of iterations employed for PC-RISR as a function of SNR (in the low regime) for signals separated by half the Rayleigh resolution. It is observed that $\alpha = 0.2$ performs qualitatively similar to root-MUSIC. For 30 iterations, $\alpha = 0.4$ clearly outperforms $\alpha = 0.49$ for PC-RISR. However, increasing the number of iterations by a factor of 10 to 300 enables the $\alpha = 0.49$ case to separate the two signals at very low SNR values. The 0.49 case does not reach a probability of 1 because a small spurious peak sometimes occurs (which fails the separation metric used here) arising from the low sample support of 10 (relative to *N*). Increasing the sample support (result not shown) alleviates this effect.

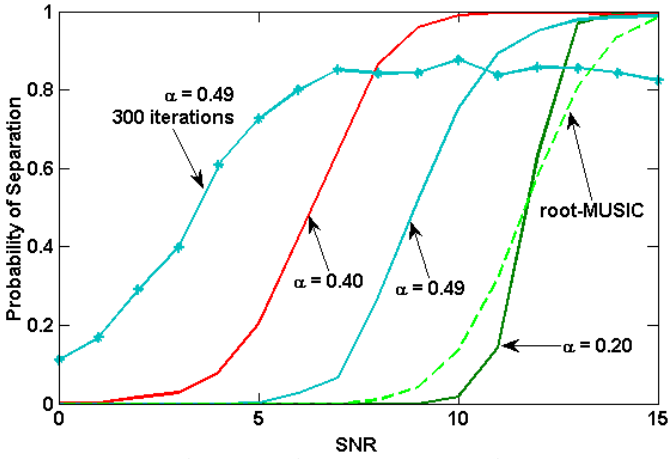

Fig. 7. Probability of separation for PC-RISR vs. SNR for α = 0.2, 0.4, and 0.49 for 30 iterations. Signal separation is half the Rayleigh resolution.

VI. Conclusions

Leveraging the original RISR algorithm and a subsequent gain-constrained version, a partially-constrained formulation denoted as PC-RISR has been demonstrated to enable super-resolved signal separation at low SNR. By virtue of leveraging the RISR structure, this formulation is robust to array calibration errors, temporal correlation of signals, and can be employed with very few data snapshots (even just 1). The trade-off factor α serves as a form of step-size to slow down convergence enough to identify closely spaced signals at low SNR that would otherwise have been suppressed by the original RISR algorithm.

References

[1] H. Van Trees, *Optimum Array Processing*, New York: Wiley, 2002.

[2] J.E. Evans, J.R. Johnson, and D.F. Sun, "Application of advanced signal processing techniques to angle of arrival estimation in ATC navigation and surveillance systems," *MIT Lincoln Lab Tech. Report 582*, June 1982.

[3] T. Shan, M. Wax, and T. Kailath, "On spatial smoothing for direction-of-arrival estimation of coherent signals," *IEEE Trans. Acoustics, Speech, and Signal Processing*, vol. ASSP-33, no. 4, pp. 806-844, Aug. 1985.

[4] S.D. Blunt, T. Chan, and K. Gerlach, "A new framework for direction-of-arrival estimation," *IEEE Sensor Array & Multichannel Signal Processing Workshop*, 21-23 July 2008.

[5] S.D. Blunt, T. Chan, and K. Gerlach, "Robust DOA estimation: the reiterative super resolution (RISR) algorithm," *IEEE Trans. Aerospace & Electronic Systems*, vol. 47, no. 1, pp. 332 – 346, Jan. 2011.

[6] S.D. Blunt and K. Gerlach, "Adaptive pulse compression via MMSE estimation," *IEEE Trans. Aerospace & Electronic Systems*, vol. 42, no. 2, pp. 572-584, July 2006.

[7] M. Popescu, S.D. Blunt, and T. Chan, "Magnetoencephalography source localization using the source affine image reconstruction (SAFFIRE) algorithm," *IEEE Trans. Biomedical Engineering*, vol. 57, no. 7, pp. 1652-1662, July 2010.

[8] T. Higgins, S.D. Blunt, and K. Gerlach, "Gain-constrained adaptive pulse compression via an MVDR framework," *IEEE Radar Conf.*, May 2009.

[9] A.J. Barabel, "Improving the resolution performance of eigenstructure-based direction-finding algorithms," *IEEE Intl. Conf. Acoustics, Speech, and Signal Processing*, pp. 336-339, Apr. 1983.

[10] P. Chen, T.J. Wu, and J. Yang, "A comparative study of model selection criteria for the number of signals," *IET Radar, Sonar & Navigation*, vol. 2, no. 3, pp. 180 – 188, 2008.