Spoil-Constrained Reiterative Estimation (SCoRE) to Address Excessive Super-Resolution

David G. Felton, Christian C. Jones, Patrick M. McCormick, Shannon D. Blunt
Radar Systems Laboratory (RSL), University of Kansas

***Abstract*—The application of structure-based adaptive estimation to random pulse repetition interval (PRI) staggering was recently shown to facilitate practical expansion of unambiguous Doppler, demonstrated via a form of reiterative minimum mean-square error (RMMSE) estimation on open-air measurements. Further analysis revealed that such methods do introduce a trade-space between straddling effects and possible losses due to excessive super-resolution (ESR). Here we explore the reason behind this trade-space and propose an alternative structure-based adaptive estimator that employs multiple eigenvector constraints to mitigate ESR by collectively "spoiling" the beam to the nominal Doppler resolution, with the general framework being directly extensible to a variety of such problems.**

Index Terms*—PRI staggering, Doppler estimation, adaptive processing

I. Introduction

Randomly changing the PRI on a pulse-to-pulse basis is a form of waveform diversity [1] that enables the unambiguous Doppler span to be increased beyond the nominal ±PRF/2 that one otherwise obtains for a uniform PRI, with pulse repetition frequency (PRF) the inverse of the PRI. The cost for this expansion is a Doppler sidelobe floor that at best approaches $1/M$, for M pulses in the coherent processing interval (CPI), and the inability to apply standard windowing methods to suppress sidelobes (over the expanded Doppler span) [2-4].

A variety of processing methods (see references in [4]) have been developed to address this type of problem, which is a form of non-uniform sampling. As a representative of such methods that can operate on high dynamic range radar data, the re-iterative super-resolution (RISR) form of RMMSE was examined in [3,4] for the specific application to Doppler expansion via randomized PRI staggering. While the ensuing experimental results were physically meaningful and could in turn be employed for downstream detection processing, it was observed that a trade-space emerged due to an effect that can be synopsized as excessive super-resolution (ESR). The following provides a deeper explanation of this effect and then proposes an alternative structure-based processing approach denoted as SCoRE that is formulated to compensate for ESR. The most notable impact of SCoRE is reduced variance in the estimation of amplitude and Doppler frequency for movers that would otherwise be completely masked by high-power clutter.

II. Stagger Modeling & Processing Review

In [1], a detailed signal model for random PRI staggering was developed, while [3,4] formulated and experimentally demonstrated receive processing that enables adaptive sidelobe suppression over an expanded Doppler span. The following provides a summary of their salient components.

For M arbitrarily staggered pulses in a CPI, Doppler frequency $f_{\rm D}$ corresponds to the slow-time steering vector

$$\mathbf{v}(f_{\rm D}) = \left[1 \;\; e^{j[2\pi f_{\rm D} T_{\rm acc}(2)]} \;\; \cdots \;\; e^{j[2\pi f_{\rm D} T_{\rm acc}(M)]}\right]^T \tag{2}$$

where $T_{\rm acc}(m)$ for $m = 1, 2, \ldots, M$ is the accumulated slow-time from the beginning of the CPI to the mth PRI, with $T_{\rm acc}(1) = 0$ for the first pulse. At the mth PRI (for $m > 1$), this accumulation is the sum of the preceding PRI intervals, which are individually denoted as T_m and have $T_{\rm avg}$ as their average. In the uniform PRI case, each $T_m = T_{\rm avg}$. While not further considered here, slow-time phase coding [1] could also be included by introducing another distinct phase term on each element of (2).

Subsuming beamforming and pulse compression, and discretizing in both range and Doppler, we can write the M-pulse slow-time data vector for the ℓth range sample as the approximation [1]

$$\mathbf{z}(\ell) \cong \mathbf{V}\mathbf{x}(\ell) + \mathbf{n}(\ell) \tag{1}$$

for $M \times N$ matrix $\mathbf{V}$ the discretized Doppler steering vectors (in the columns) with corresponding scattering in the $N \times 1$ vector $\mathbf{x}(\ell)$, and $\mathbf{n}(\ell)$ an $M \times 1$ vector of noise samples. Denote $K = (N/M) > 1$ due to oversampling in Doppler (i.e., $N > M$). The nominal Doppler span of $\mathbf{V}$ would be ±PRF/2 for uniform PRI, which is the same as average PRF in this case. Staggering permits expansion beyond this limit, with $\pm\beta$ (PRF/2) denoting the expanded Doppler span of interest by factor $\beta > 1$ (see [1]).

Standard Doppler processing (DP) can be performed via

$$\begin{aligned}\hat{\mathbf{x}}_{\rm DP}(\ell) &= \mathbf{W}_{\rm DP}^H \mathbf{z}(\ell) \\ &\cong \mathbf{W}_{\rm DP}^H \mathbf{V}\mathbf{x}(\ell) + \mathbf{W}_{\rm DP}^H \mathbf{n}(\ell)\end{aligned} \tag{3}$$

for $(\bullet)^H$ the complex-conjugate transpose (Hermitian) operation and $\hat{\mathbf{x}}_{\rm DP}(\ell)$ the ensuing Doppler response estimate, which at best will possess sidelobes on the order of $1/M$ if properly designed for the expanded span of $\pm\beta$ (PRF/2) [2]. The $M \times N$ matrix $\mathbf{W}_{\rm DP} = \mathbf{V}$ is also over-sampled ($N > M$) to limit Doppler straddling effects, though doing so means that adjacent columns in $\mathbf{V}$ become correlated, which was shown in [4] to cause

This work was supported by the Office of Naval Research under Contract #N00014-23-C-1053. DISTRIBUTION STATEMENT A. Approved for Public Release.

excessive super-resolution (ESR) effects that translate into variability in the complex amplitude and frequency estimates that could impact downstream detection and tracking.

While [4] also addresses clutter cancellation, here we shall focus solely on adaptive estimation, which was illustrated using the reiterative super-resolution (RISR) algorithm [5] derived from the mean-square error (MSE) cost function

$$J_{\text{MSE}}(\ell, f_{\text{D}}) = E\left[\left|x(\ell; f_{\text{D}}) - \mathbf{w}^H(\ell, f_{\text{D}})\,\mathbf{z}(\ell)\right|^2\right]. \quad (4)$$

The unconstrained (U) RISR algorithmic solution to (4) involves updating the estimates $\hat{x}(\ell; f_{\text{D}})$ across Doppler (collected into vector $\hat{\mathbf{x}}(\ell)$) via repeated application of

$$\mathbf{P}_i(\ell) = [\hat{\mathbf{x}}_{i-1}(\ell)\,\hat{\mathbf{x}}_{i-1}^H(\ell)] \odot \mathbf{I}_{N\times N}, \quad (5)$$

$$\mathbf{W}_{\text{U},i}(\ell) = \left(\mathbf{V}\,\mathbf{P}_i(\ell)\mathbf{V}^H + \mathbf{R}_{\text{nse}}\right)^{-1}\mathbf{V}\,\mathbf{P}_i(\ell), \quad (6)$$

and

$$\hat{\mathbf{x}}_i(\ell) = \mathbf{W}_{\text{U},i}^H(\ell)\,\mathbf{z}(\ell), \quad (7)$$

for iteration index $i = 1, \ldots, I$ and noise covariance matrix $\mathbf{R}_{\text{nse}}$. This process is initialized by performing $\hat{\mathbf{x}}_0(\ell) = \mathbf{W}_{\text{U},0}^H(\ell)\,\mathbf{z}(\ell)$ for $\mathbf{W}_{\text{U},0}(\ell) = \mathbf{W}_{\text{DP}}$. Each column of (6), denoted as $\mathbf{w}_{\text{U},i}(\ell, f_{\text{D}})$, could alternatively be replaced by the unity gain-constrained (GC) form of RISR [6]

$$\mathbf{w}_{\text{GC},i}(\ell, f_{\text{D}}) = \frac{(\mathbf{V}\mathbf{P}_i(\ell)\mathbf{V}^H + \mathbf{R}_{\text{nse}})^{-1}\,\mathbf{v}(f_{\text{D}})}{\mathbf{v}^H(f_{\text{D}})\,(\mathbf{V}\mathbf{P}_i(\ell)\mathbf{V}^H + \mathbf{R}_{\text{nse}})^{-1}\,\mathbf{v}(f_{\text{D}})}, \quad (8)$$

which provides more robustness in practice. In [4] the general equivalence between RISR, a form of expectation-maximization (EM) [7], and the iterative adaptive approach (IAA) [8] was also discussed.

III. Excessive Super-Resolution (ESR)

To understand how/why the unconstrained form in (6) suffers from ESR, and is only partially addressed by (8), consider the simplistic scenario of a single large scatterer in white noise. This scatterer corresponds to Doppler steering vector $\mathbf{v}(f_{\text{D}} = f_{\text{true}})$, which we shall denote as $\mathbf{v}_{\text{true}}$. For now, let this steering vector perfectly align with one that exists in $\mathbf{V}$, which we refer to as $\mathbf{v}_0$ at frequency f_0. Shortly these will differ, hence the need for separate designations.

In $\mathbf{V}$ there are adjacent steering vectors $\mathbf{v}(f_0 + \Delta f = f_{+\Delta})$ and $\mathbf{v}(f_0 - \Delta f = f_{-\Delta})$, where $\pm\Delta f$ depends on the degree of Doppler oversampling ($K > 1$). For convenience, we refer to $\mathbf{v}(f_{+\Delta})$ and $\mathbf{v}(f_{-\Delta})$ as $\mathbf{v}_{+1}$ and $\mathbf{v}_{-1}$, respectively. For any amount of oversampling the vectors $\mathbf{v}_0$ and $\mathbf{v}_{+1}$ (or $\mathbf{v}_{-1}$) possess some non-zero correlation determined by $|\mathbf{v}_0^H\mathbf{v}_{+1}|^2 / (\|\mathbf{v}_0\|^2 \cdot \|\mathbf{v}_{+1}\|^2)$ that approaches unity as K increases (likewise for $\mathbf{v}_{-1}$). For $K > 1$, these adjacent steering vectors coincide with roll-off portions of the Doppler mainlobe centered at f_0.

Thus, the unconstrained RMMSE filter (6) for $\mathbf{v}_0$ can alternatively be written as

$$\mathbf{w}_{\text{U},i}(f_0) \cong \left(\sum_{n=-1}^{+1} p_i(f_n)\mathbf{v}_n\mathbf{v}_n^H + \tilde{\mathbf{V}}\,\tilde{\mathbf{P}}_i\,\tilde{\mathbf{V}}^H + \sigma_{\text{nse}}^2\mathbf{I}\right)^{-1}\mathbf{v}_0\; p_i(f_0), \quad (9)$$

where range index ℓ has been omitted for clarity. The tilde on $\tilde{\mathbf{V}}$ and $\tilde{\mathbf{P}}_i$ indicates removal of the $n = -1$, 0, and +1 terms that have been extracted (into the summation) for illustration purposes. After initialization using standard Doppler processing by applying $\mathbf{W}_{\text{DP}} = \mathbf{V}$, the initial power estimates $p_{i=1}(f_0)$, $p_{i=1}(f_{+1})$, and $p_{i=1}(f_{-1})$ via (5) would all be large given a large scatterer having Doppler f_0. The inverse matrix in (9) therefore suppresses signals corresponding to $\mathbf{v}_0$, $\mathbf{v}_{+1}$, and $\mathbf{v}_{-1}$. However, the $p_{i=1}(f_0)$ scaling outside the inverse matrix in (9) serves to compensate the $n = 0$ term (since $\mathbf{v}_0 = \mathbf{v}_{\text{true}}$), thus avoiding self-cancellation during the subsequent application of $\mathbf{w}_{\text{U},i=1}(f_0)$ for the updated estimation of scattering per (7).

Forming the adjacent filters $\mathbf{w}_{\text{U},i=1}(f_{+1})$ and $\mathbf{w}_{\text{U},i=1}(f_{-1})$ only entails changing the last two terms in (9) to $\mathbf{v}_{\pm1}$ and $p_{i=1}(f_{\pm1})$ since the inverse matrix is the same for all Dopplers in $\mathbf{V}$. Thus, the entire estimated Doppler span experiences a null centered at f_0. The initial $p_{i=1}(f_{+1})$ and $p_{i=1}(f_{-1})$ adjacent terms, while large, would nonetheless be smaller than $p_{i=1}(f_0)$ since $\mathbf{v}_{+1}$ and $\mathbf{v}_{-1}$ are marginally misaligned with $\mathbf{v}_{\text{true}}$ (= $\mathbf{v}_0$ currently), meaning their respective filters experience less compensation. Successive application of (5)-(7) over iteration index i therefore yields successively less compensation for $p_i(f_{+1})$ and $p_i(f_{-1})$ that would otherwise offset the effects of the nearby null centered at f_0, ultimately leading to suppression of these adjacent terms. The ensuing super-resolution effect serves as an implicit manifestation of sparsity because the result becomes a collection of point responses at the given oversampling granularity.

Now generalize to the physically meaningful scenario where the true Doppler value does <u>not</u> lie on the quantized grid specified by $\mathbf{V}$ (since Doppler exists on a continuum), which is essentially always true for every scatterer in reality. This "straddling" (or cusping) condition is well-known in radar (see Chap. 14 of [9]), with more recent work in compressive sensing also referring to it as the "off grid" problem. To be precise, the *lack of straddling* is a rare occurrence, meaning straddling by some unknown amount should be treated as the natural state of scattering. We can quantify a "straddling loss" as the amount of deviation from unity for correlation $|\mathbf{v}_0^H\mathbf{v}_{\text{true}}|^2 / (\|\mathbf{v}_0\|^2 \cdot \|\mathbf{v}_{\text{true}}\|^2)$ when $\mathbf{v}_{\text{true}} \neq \mathbf{v}_0$ but $\mathbf{v}_0$ is still the closest steering vector. Increasing the degree of oversampling, which reduces the separation (in frequency) between adjacent steering vectors, therefore reduces this loss.

Revisiting (9), this straddling misalignment leads to an expected amount of loss for $p_i(f_0)$, with the super-resolution result progressing in the same manner as before. However, estimation of the rest of the Doppler span now experiences a null at $\mathbf{v}_0$ instead of $\mathbf{v}_{\text{true}}$. Thus, misalignment of the null with the large scatterer at f_{true} ultimately leads to a mismatch effect that degrades sensitivity to other smaller scatterers. While this effect can be remediated (somewhat) by increasing the degree of oversampling, doing so is not a panacea. The natural Doppler-spread of clutter, which comprises a continuum of large scattering, still leads to (generally incorrect) point-like sparsity for unconstrained adaptive filtering because each filter is independently compensated, leading to a "local competition" among highly similar steering vectors and their respective

power estimates. This overly sparse manifestation of ESR induces an attendant loss of sensitivity when dynamic range is high, as illustrated in Fig. 1 (red trace), where many movers are missed, spurious peaks emerge, and the noise floor is absent.

Now consider the GC version in (8) via the same scenario of a single large scatterer in white noise. In this case the scaling term at the end of (9) now becomes the denominator

$$\hat{p}_i(f_0) \to \frac{1}{\mathbf{v}_0^H\left(\sum_{n=-1}^{+1} p_i(f_n)\mathbf{v}_n\mathbf{v}_n^H + \tilde{\mathbf{V}}\,\tilde{\mathbf{P}}_i\tilde{\mathbf{V}}^H + \sigma_{\text{nse}}^2\mathbf{I}\right)^{-1}\mathbf{v}_0} \qquad (10)$$

in which the entire Doppler span estimate is included. Under the idealistic condition that $\mathbf{v}_0 = \mathbf{v}_{\text{true}}$, the same point-like response occurs as it did for the unconstrained case, which would be the correct result in the case of perfect alignment. In the general straddling scenario, however, the inclusion of all estimated scattering in the normalization term of (10) instead tends to lead toward a "local weighted average" effect in which a combination of $\hat{p}_{i=0}(f_0)$, $\hat{p}_{i=0}(f_{+1})$, and/or $\hat{p}_{i=0}(f_{-1})$ frames a slightly less super-resolved mainlobe estimate. Consequently, the GC version more naturally addresses Doppler-spread scattering (e.g. clutter) by avoiding point-like sparsity, thus mostly mitigating the previous sensitivity degradation (see black trace in Fig. 1) because steering vectors $\mathbf{v}_0$ and $\mathbf{v}_{+1}$ (or $\mathbf{v}_{-1}$) around $\mathbf{v}_{\text{true}}$ collectively provide a slightly broader null, which addresses null misalignment that degraded sensitivity for estimation of the rest of the Doppler span (for unconstrained).

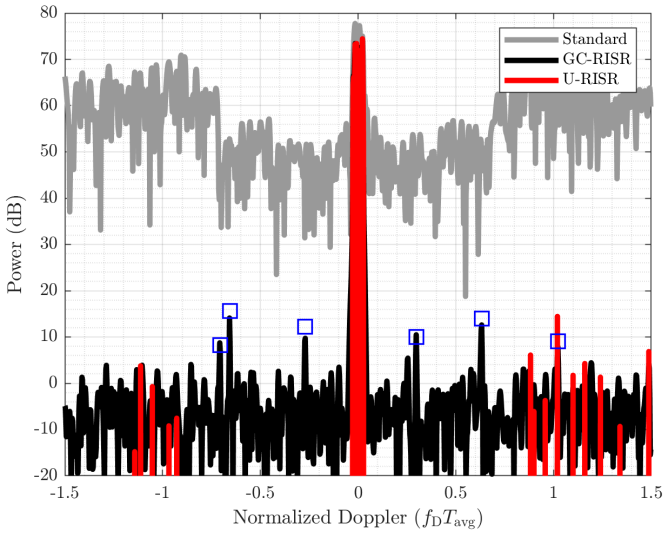

Fig. 1. Illustration of high-power clutter and movers using standard Doppler processing, U-RISR, and GC-RISR, with the gain constraint preserving a noise floor and sensitivity to lower power scattering

With that said, for GC adaptive filtering of high dynamic range scenarios, some ESR effects do remain. Where the unconstrained filters suppressed adjacent components through successively less compensation, the preservation of some Doppler width for GC filters means that adjacent components are collectively nulling/compensating. At modest oversampling (e.g. $1 < K < 2$) the effect of ESR is minor, but straddling is worse (possibly a few dB). At higher oversampling, which defrays much of the straddling loss, the frequency span of the "local weighted average" likewise narrows since $\mathbf{v}_0$ and $\mathbf{v}_{+1}$ (or $\mathbf{v}_{-1}$) around $\mathbf{v}_{\text{true}}$ are closer together. In the presence of different random instantiations of high-power clutter, the ensuing super-resolved response can evidence some frequency variability (yet still within the nominal mainlobe) as well as a few dB in amplitude variability. The reason is because estimation of the clutter Doppler span again becomes overly sparse; less severe than the unconstrained case but still performance-limiting (see [4]). It therefore stands to reason that mitigating ESR should provide less variability for amplitude and frequency estimation in these high dynamic range scenarios.

IV. Spoil-Constrained Reiterative Estimation (SCoRE)

To address both straddling and ESR effects, we reformulate structure-based adaptive filtering to preserve the nominal Doppler mainlobe (i.e. prevent super-resolution) by applying additional linear equality constraints. To simplify the extension to multiple constraints, replace the MSE cost function of (4) with the minimum power distortionless response (MPDR) [10]

$$J_{\text{MPDR}}(\ell, f_{\text{D}}) = E\left[\left|\mathbf{w}^H(\ell, f_{\text{D}})\,\mathbf{z}(\ell)\right|^2\right] \qquad (11)$$

subject to the constraints

$$\mathbf{C}^H(f_{\text{D}})\,\mathbf{w}(\ell, f_{\text{D}}) = \mathbf{g} \qquad (12)$$

in which $M \times L$ matrix $\mathbf{C}(f_{\text{D}})$ contains L linear equality constraint vectors (in the columns) and $L \times 1$ vector $\mathbf{g}$ are the L gain values. In the case that $L = 1$ and $\mathbf{C}(f_{\text{D}})$ is the single vector $\mathbf{v}(f_{\text{D}})$, this problem yields the familiar GC form in (8) [6].

For $L > 1$, we extend to multiple Lagrange multipliers [10], thereby yielding the multiple-gain-constrained (MGC) filter formulation

$$\mathbf{w}_{\text{MGC},i}(\ell, f_{\text{D}}) = \mathbf{R}_{\text{z}}^{-1}\,\mathbf{C}(f_{\text{D}})\left(\mathbf{C}^H(f_{\text{D}})\,\mathbf{R}_{\text{z}}^{-1}\,\mathbf{C}(f_{\text{D}})\right)^{-1}\mathbf{g}\,. \qquad (13)$$

In keeping with the reiterative update framework of (5)-(7), the structured covariance matrix in (13) is

$$\mathbf{R}_{\text{z}} = \left(\mathbf{V}\,\mathbf{P}_i(\ell)\mathbf{V}^H + \mathbf{R}_{\text{nse}}\right), \qquad (14)$$

the same as appears in (6) and (8). The "Doppler frequency steered" constraint matrix in (12) and (13) is obtained via

$$\mathbf{C}(f_{\text{D}}) = \overline{\mathbf{C}} \odot [\mathbf{v}(f_{\text{D}})\mathbf{1}_{L\times 1}^T], \qquad (15)$$

with the nominal constraint vectors in $\overline{\mathbf{C}}$ defined relative to zero Doppler.

Given that we wish to preserve the Doppler mainlobe to prevent ESR, the direct approach to parameterize $\overline{\mathbf{C}}$ would be to collect L equally-spaced steering vectors spanning $\pm\Delta_{\text{ML}}$ in Doppler (including zero). Here, Δ_{ML} denotes the mainlobe extent being subsumed, such as first-null or 3-dB. Of course, the potential problem with this approach is that these L steering vectors spanning some portion of the Doppler mainlobe are highly correlated, which could translate to ill-conditioning effects when used in (13).

An alternative approach, and the one that we shall employ, is to first construct the $M \times M$ structured covariance matrix

$$\mathbf{A} = \int_{-\Delta_{\mathrm{ML}}}^{+\Delta_{\mathrm{ML}}} \mathbf{v}(f_D)\,\mathbf{v}^H(f_D)\,df_D = \mathbf{Q}\boldsymbol{\Lambda}\mathbf{Q}^H, \tag{16}$$

that involves integration over a desired portion of the zero-centered Doppler mainlobe for steering vectors specific to the given stagger sequence. The subsequent eigen-decomposition in (16) realizes a set of eigenvalues in $\boldsymbol{\Lambda}$ assumed to be in descending order. Thus, for λ_1 to λ_L deemed to be the dominant eigenvalues, the corresponding eigenvectors in $\mathbf{Q}$ are used to form the constraint matrix as

$$\overline{\mathbf{C}} = \left[\mathbf{q}_1 \;\cdots\; \mathbf{q}_L\right]. \tag{17}$$

This approach is reminiscent of "eigenvalue beamforming" that uses sample covariance matrix (SCM) estimation [11].

For a given constraint matrix $\overline{\mathbf{C}}$, the associated gain values in $\mathbf{g}$ can be readily obtained by noting that we wish to preserve the mainlobe structure. Consequently, we can pose (12) as

$$\begin{aligned} \mathbf{g} &= \mathbf{C}^H(f_{\mathrm{D}})\,\mathbf{w}(\ell, f_{\mathrm{D}}) \\ &\doteq \mathbf{C}^H(f_{\mathrm{D}})\,\mathbf{v}(f_{\mathrm{D}}) \\ &= \overline{\mathbf{C}}^H\mathbf{v}(0) \end{aligned} \tag{18}$$

where the last line used (15) to remove frequency dependence.

V. Additional Implementation Aspects

It was observed in [4] that while adaptive estimation using GC-RISR preserves a semblance of the noise floor, the level is roughly β times lower than reality for this PRI-staggered/ Doppler-expanded application. The reason is due to the fixed noise energy in $|\mathbf{n}(\ell)|^2$ being estimated over a $\beta\times$ greater span via $\mathbf{P}_i(\ell)$ in (5), with the noise effectively being double-counted in (14). This effect also carries over to the use of multiple mainlobe constraints per Sect. IV.

It was likewise observed in [4] that the analytical evaluation of signal-to-interference-plus-noise (SINR) ratio involves a loss that is revealed when the scaling term in (10) gets normalized out. While this scaling term is an intrinsic part of the estimation process via (8), it translates into an undesired noise gain at the final stage, which can become significant for scenarios with high dynamic range. In short, we have another aspect that can lead to misestimation of the noise, with the combination of these aspects potentially degrading downstream detection.

These effects can be remediated through a couple of simple modifications. First, consider the fact that the only components of $N \times N$ diagonal matrix $\mathbf{P}_i(\ell)$ that need to be included are those having high enough power to produce sidelobes. Consequently, the double-counting of noise in (14) can be addressed by replacing (5) in each iteration with "sidelobe significance pruning" (S2P) via the two-step procedure

$$\tilde{\mathbf{P}}_i(\ell) = [\hat{\mathbf{x}}_{i-1}(\ell)\,\hat{\mathbf{x}}_{i-1}^H(\ell)] \odot \mathbf{I}_{N\times N} \tag{19}$$

and

$$\mathbf{P}_{i,nn}(\ell) = \begin{cases} \tilde{\mathbf{P}}_{i,nn}(\ell), & \text{if } \tilde{\mathbf{P}}_{i,nn}(\ell) > (\sigma_{\mathrm{n}}^2/\mathrm{PSL}_{\mathrm{abs}}) \\ 0, & \text{otherwise} \end{cases} \tag{20}$$

for $n = 1, 2, \cdots, N$ and with $\mathrm{PSL}_{\mathrm{abs}}$ (<1) the Doppler peak sidelobe level on an absolute scale (i.e. not dB) for the given stagger sequence over the Doppler span of interest.

The other modification involves employing the lesson learned from the analytical SINR evaluation in [4], where the iterative estimation process realizes an intrinsic scaling that also scales the noise. Thus, after the last iteration, "terminal noise normalization" (TN2) is performed via

$$\overline{\mathbf{w}}_i(f_{\mathrm{D}}) = \frac{\mathbf{w}_i(f_{\mathrm{D}})}{\left\|\mathbf{w}_i(f_{\mathrm{D}})\right\|} \tag{21}$$

so noise is passed with unity gain when (21) is applied to (1).

Collecting the various pieces together, the SCoRE algorithm involves selecting the constraint terms $\mathbf{C}$ and $\mathbf{g}$ via (16)-(18) and is then initialized by performing $\hat{\mathbf{x}}_0(\ell) = \mathbf{W}_{\mathrm{DP}}^H(\ell)\,\mathbf{z}(\ell)$. For each subsequent iteration $i = 1, \ldots, I-1$ we apply (19) and (20) to obtain the S2P version of $\mathbf{P}_i(\ell)$, determine $\mathbf{R}_{\mathrm{z}}$ via (14), and then the MGC filterbank $\mathbf{W}_{\mathrm{MGC},i}(\ell)$ by evaluating (13) with (15) for each discretized Doppler. This filterbank is then applied like in (7) to realize the new estimate $\hat{\mathbf{x}}_i(\ell)$. When iteration $i = I$, the TN2 in (21) is performed prior to the filterbank being applied to $\mathbf{z}(\ell)$ like (7), thereby realizing a final result that is suitable for a subsequent detection stage.

VI. Simulation Results

To illustrate the benefits of limiting ESR during adaptive estimation, 1000 Monte Carlo trials were performed using a clutter-to-noise ratio of 80 dB, with independent random instantiations of clutter and noise for each trial (both complex Gaussian). To permit suitable averaging across trials, a single random instantiation of PRI staggering was generated (the same as depicted in Fig. 1) and reused for all 1000 trials. The noise power was set to unity.

We examine the mean power and standard deviation on a per-Doppler-bin basis across all trials, comparing the impact of the degree of oversampling by using a modest (K = 1.25) and high (K = 4) value. The latter addresses the straddling loss incurred by the former, but in turn encounters ESR effects that SCoRE has been developed to address.

We also assess Doppler frequency estimation accuracy for each mover in the simulated scene. For simplicity, we ignore the (clearly necessary) detection stage and search for the maximum value over a null-to-null mainlobe width around the true Doppler of each mover (for each processing method and for each trial). Once this local peak is determined, the values on either side comprise a 3-tuple for which fine Doppler estimation is performed via the method described in Sect. 17.5.5 of [9], specifically using equation (17.42). The resulting estimate is then compared to the true Doppler to obtain a root-MSE that is normalized by the mainlobe peak-to-null width, thereby realizing a "percent of mainlobe" RMSE.

First, for K = 1.25 and 4, respectively, Figs. 2 and 3 show zoomed-in Doppler estimation responses for two proximate movers when applying GC-RISR and SCoRE with L = 4 constraints for (16) constructed using the 3-dB mainlobe width. What becomes evident in these plots is that, especially at higher oversampling K, SCoRE provides much more consistent mover

responses by avoiding the variability that comes with ESR.

Again for $K = 1.25$ and 4, respectively, Figs. 4 and 5 depict the mean power estimates and standard deviations on a per-Doppler-bin basis across the 1000 trials zoomed-in on these same two proximate movers, using GC-RISR and SCoRE with different numbers of eigenvector constraints. In so doing we find that there is a sweet spot in the setting of L, with too low a value not providing sufficient dynamic range operation and too high a value having the effect of introducing sidelobes. We see that L between 3 and 5 provides good results here, with 4 arguably the best in terms of consistently high power and low standard deviation. Also, Fig. 5 shows that a better power estimate is achieved for the 15 dB SNR mover when higher oversampling is used, though GC-RISR suffers from greater variability in that case. In contrast, SCoRE ($L = 4$ especially) reduces this variability while also improving on the power estimate.

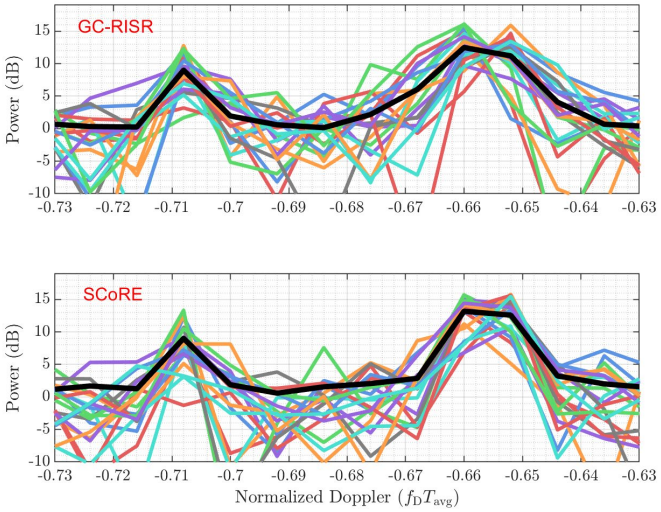

Fig. 2. Adaptive Doppler estimation responses for 20 independent trials (random clutter/noise) for (top) GC-RISR and (bottom) SCoRE with $L = 4$ zoomed to two proximate movers, for $K = 1.25$ (modest) oversampling

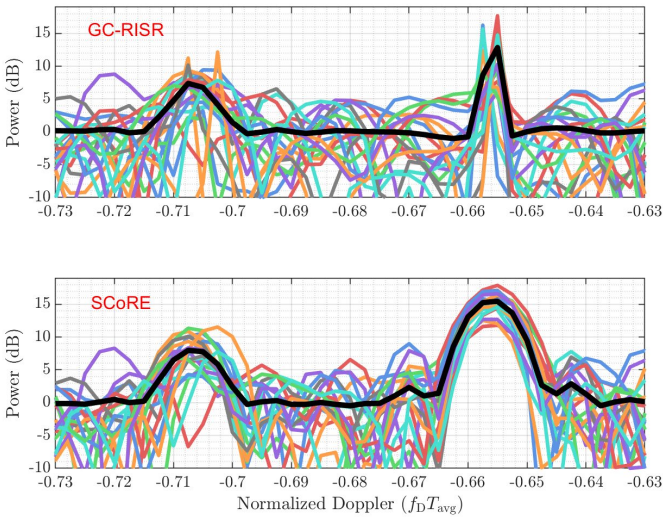

Fig. 3. Adaptive Doppler estimation responses for 20 independent trials (random clutter/noise) for (top) GC-RISR and (bottom) SCoRE with $L = 4$ zoomed to two proximate movers, for $K = 4$ (high) oversampling

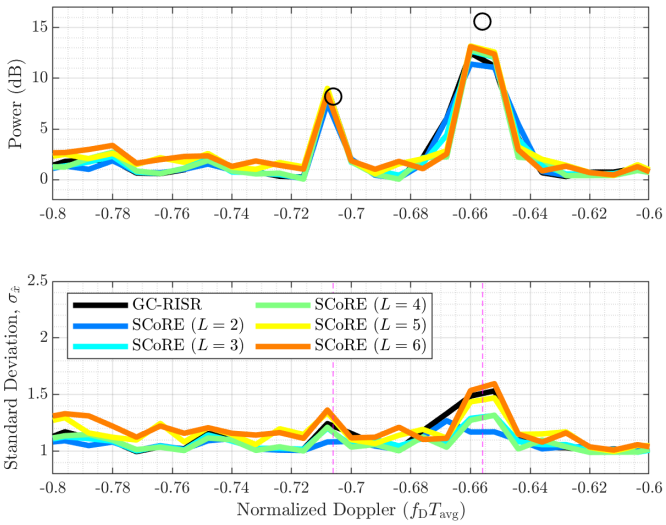

Fig. 4. Adaptive estimation comparison for two proximate movers for (top) mean power and (bottom) standard deviation over 1000 independent trials (random clutter/noise) for $K = 1.25$ (modest) oversampling

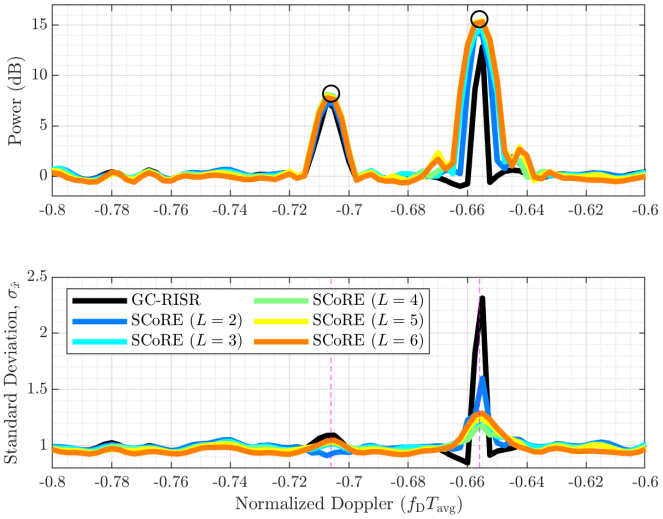

Fig. 5. Adaptive estimation comparison for two proximate movers for (top) mean power and (bottom) standard deviation over 1000 independent trials (random clutter/noise) for $K = 4$ (high) oversampling

Tables 1 and 2 list the Doppler "percent of mainlobe" RMSE values for $K = 1.25$ and 4, respectively, for each of the 6 movers in the scene (per Fig. 1), averaged over the 1000 trials. As a note, Figs. 2-5 depict movers #1 and #2. We see that movers #1 and #6 have SNR values (after all gain) of 8.2 and 9.1 dB, which would likely not be detectable on a consistent basis. Because this assessment ignores detection, it is not surprising that such high frequency error occurs. Movers #2-5 have SNR values from 10 to 15.6 dB, which would be increasingly detectable, though the presence of high clutter would prevent doing so without some means of adaptive estimation or clutter cancellation. What is particularly noteworthy is the significant reduction in error when oversampling is increased (from Table 1 to Table 2), especially for SCoRE. Put another way, by preventing ESR and the variability it entails, the nominal resolution spoiling enforcement of SCoRE leads to improved Doppler estimation with higher oversampling for high dynamic range scenarios.

Table 1. Doppler estimation error via (17.42) in [9], averaged over 1000 independent trials (random clutter/noise) and normalized by first-null width to yield Doppler mainlobe percentage for $K = 1.25$

	SNR	True Doppler	GC-RISR	SCoRE (L=3)	SCoRE (L=4)	SCoRE (L=5)
Target 1	8.2	-0.706	27.6%	29.4%	31.9%	33.8%
Target 2	15.6	-0.656	18.1%	18.4%	17.2%	22.3%
Target 3	12.2	-0.271	26.2%	24.1%	24.5%	24.8%
Target 4	10	0.2982	21.9%	23.7%	24.2%	26.4%
Target 5	14	0.6319	20.6%	20.2%	23.2%	25.9%
Target 6	9.1	1.0234	39.0%	43.0%	44.4%	46.2%

Table 2. Doppler estimation error via (17.42) in [9], averaged over 1000 independent trials (random clutter/noise) and normalized by first-null width to yield Doppler mainlobe percentage for $K = 4$

	SNR	True Doppler	GC-RISR	SCoRE (L=3)	SCoRE (L=4)	SCoRE (L=5)
Target 1	8.2	-0.706	25.6%	25.3%	26.7%	26.9%
Target 2	15.6	-0.656	16.0%	13.1%	10.8%	10.3%
Target 3	12.2	-0.271	18.6%	14.8%	13.2%	13.9%
Target 4	10	0.2982	18.7%	18.9%	17.5%	18.0%
Target 5	14	0.6319	14.3%	13.1%	10.7%	10.5%
Target 6	9.1	1.0234	32.9%	34.9%	37.5%	40.2%

VII. Experimental Results

Finally, in the interest of showing that SCoRE does indeed work on real data, Figs. 6 and 7 depict standard Doppler processing and SCoRE (using $L = 4$), respectively, applied to randomly staggered PRI data illuminating a traffic intersection in Lawrence, KS. The CPI is 400 ms and comprises 200 pulses modulated with a 100 MHz LFM waveform having a 6 μs pulse-width. Random PRIs span 1 ms to 2.9 ms, with a 2 ms mean. The dashed purple lines indicate where ±PRF/2 would be for the uniform case. Further details of these open-air measurements can be found in [4].

Even for this relatively low dynamic range scenario (due to lower power transmit) we see prominent sidelobe "streaks" occur for standard processing (Fig. 6) that are completely suppressed by SCoRE to reveal the actual movers (Fig. 7). Consequently, SCoRE provides enhanced visibility for the expansion of Doppler using PRI staggering; the point being that the improved estimation consistency demonstrated in simulation is expected to translate to hardware operation.

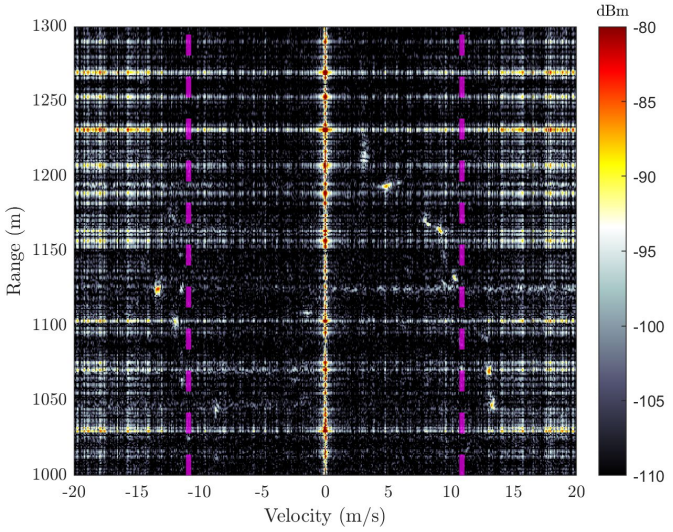

Fig. 6. Standard Doppler processing of open-air experimental data collection involving random PRI staggering to expand unambiguous Doppler

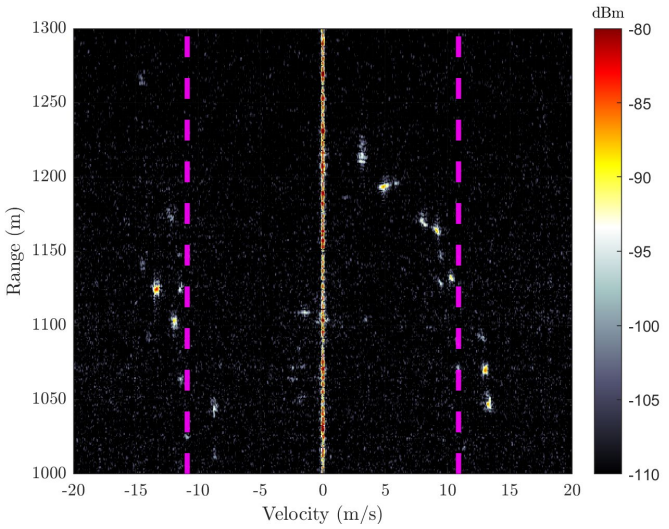

Fig. 7. Adaptive processing via SCoRE of open-air experimental data collection involving random PRI staggering to expand unambiguous Doppler

VIII. Conclusions

The recent evaluation of GC-RISR on open-air data involving random PRI staggering revealed a trade-space between straddling loss and ESR effects, the latter translating into undesired variability in terms of estimated amplitude and fine Doppler estimation. It has been shown that reformulating structure-based adaptive estimation to accommodate multiple eigenvector constraints encompassing the mainlobe yields the ability to reduce this variability, thereby providing more consistent downstream detection and tracking performance.

References

[1] S.D. Blunt, L.A. Harnett, B. Ravenscroft, R.J. Chang, C.T. Allen, P.M. McCormick, "Implications of diversified doppler for random PRI radar," *IEEE Trans. AES*, vol. 59, no. 4, pp. 3811-3834, Aug. 2023.

[2] R.J. Chang, C.C. Jones, J.W. Owen, S.D. Blunt, "Gradient-based optimization of pseudo-random PRI staggering," *IEEE Trans. Radar Systems*, vol. 1, pp. 249-263, July 2023.

[3] L.A. Harnett, B. Ravenscroft, S.D. Blunt, C.T. Allen, "Experimental evaluation of adaptive doppler estimation for PRI-staggered radar," *IEEE Radar Conf.*, New York City, NY, USA, Mar. 2022.

[4] L.A. Harnett, B. Ravenscroft, D.G. Felton, P.M. McCormick, J.W. Owen, S.D. Blunt, "Adaptive doppler estimation for random PRI staggering with clutter cancellation," *IEEE Trans. Radar Systems*, vol. 4, 2026.

[5] S.D. Blunt, T. Chan, K. Gerlach, "Robust DOA estimation: the reiterative superresolution (RISR) algorithm," *IEEE Trans. AES*, vol. 47, no. 1, pp. 332-346, Jan. 2011.

[6] E. Hornberger, S.D. Blunt, T. Higgins, "Partially constrained adaptive beamforming for super-resolution at low SNR," *IEEE Intl. CAMSAP Workshop*, Cancun, Mexico, Dec. 2015.

[7] M.A.T. Figueiredo, "Adaptive sparseness for supervised learning," *IEEE Trans. Pattern Analysis & Machine Intelligence*, vol. 25, no. 9, pp. 1150-1159, Sept. 2003.

[8] P. Stoica, J. Li, H. He, "Spectral analysis of nonuniformly sampled data: a new approach versus the periodogram," *IEEE Trans. Signal Processing*, vol. 57, no. 3, pp. 848-58, Mar. 2009.

[9] M.A. Richards, J.A. Scheer, W.A. Holm, eds., *Principles of Modern Radar: Basic Principles*, 1st ed., SciTech Publishing, 2010.

[10] H.L. Van Trees, *Optimum Array Processing*, Wiley-Interscience, 2002, Chap. 6.

[11] A. Pezeshki, B.D. Van Veen, L.L. Scharf, H. Cox, M.L. Nordenvaad, "Eigenvalue beamforming using a multirank MVDR beamformer and subspace selection," *IEEE Trans. Signal Processing*, vol. 56, no. 5, pp. 1954-1967, May 2008,