Single-Snapshot ESPRIT via Algebraic Diversity

Mitchell A. Thornton
Darwin Deason Institute for Cyber Security
Department of Electrical and Computer Engineering
Southern Methodist University, Dallas, Texas, USA
mitch@smu.edu

Technical Area: E. Array Processing and Multisensor Systems
Topic: E.2 DOA Estimation

***Abstract*—The ESPRIT algorithm estimates directions of arrival (DOAs) from the signal subspace of the sample covariance matrix. At $L = 1$ snapshot, the sample covariance has rank 1, which is insufficient to resolve $K > 1$ sources. This paper applies the algebraic diversity (AD) framework to replace the sample covariance with a group-averaged covariance that has rank M (the array size) from a single snapshot, enabling ESPRIT to operate in the single-snapshot regime. For a 16-element uniform linear array with 3 sources at 6 dB SNR, conventional ESPRIT at $L = 1$ fails 58% of the time due to rank deficiency; AD-ESPRIT succeeds 100% of the time with an RMSE of 3.4 degrees, matching conventional ESPRIT at $L \approx 4$ snapshots. In spatially colored noise ($\rho = 0.9$), conventional ESPRIT at $L = 1$ fails 62% of the time; AD-ESPRIT maintains 100% success with an RMSE of 4.7 degrees. The approach requires no modification to the ESPRIT algorithm itself; only the covariance input changes.**

Extended Summary

Problem and Prior Work

ESPRIT [1] estimates DOAs by exploiting the shift-invariance structure of a uniform linear array (ULA). It requires the sample covariance matrix to have rank at least $K + 1$ to separate the K-dimensional signal subspace from the noise subspace. The standard sample covariance from L snapshots has rank $\min(L, M)$. At $L = 1$, the rank is 1—insufficient for any multi-source scenario—and ESPRIT fails. This snapshot bottleneck limits ESPRIT in radar, sonar, and cognitive radio applications where the environment changes faster than snapshots can be collected.

Several techniques address this limitation but each introduces tradeoffs. Forward-backward averaging doubles the effective snapshots for centrosymmetric arrays but still requires $L > 1$. Spatial smoothing trades aperture for rank, reducing resolution. Toeplitz completion methods require iterative optimization. Sparse reconstruction approaches (ℓ_1-SVD) provide single-snapshot capability but require a discretized angular grid and are computationally expensive. AD-ESPRIT is unique in providing single-snapshot operation with full aperture, no angular grid, no iterative solver, and—as described below—built-in noise characterization.

Approach: AD-ESPRIT

We replace the rank-1 sample covariance with the group-averaged covariance from the algebraic diversity (AD) framework [2, 3]. For the cyclic group $\mathbb{Z}_M$ acting on the snapshot via cyclic shifts, the group-averaged covariance is a circulant (Toeplitz) matrix formed from the circular autocorrelation of a single snapshot. This matrix has rank M regardless of L, providing ESPRIT with a full-rank input. We call this increase from rank 1 to rank M the *rank-lift*. The cyclic group is the natural choice for a ULA because the Vandermonde steering vectors commute with circulant shifts, making the group exactly matched in white noise. The ESPRIT algorithm itself is completely unmodified; only the covariance input changes.

The rank-lift has a second, equally important consequence: it provides a noise subspace from a single snapshot. The conventional rank-1 covariance has $M - 1$ zero eigenvalues that carry no noise information. The AD covariance has M nonzero eigenvalues, so the $M - K$ smallest eigenvalues characterize the noise structure. This enables two capabilities that are impossible with conventional processing at $L = 1$:

Colored noise detection. The *algebraic coloring index* α, defined as the coefficient of variation of the noise eigenvalues, distinguishes white noise ($\alpha \approx 0$, uniform eigenvalues) from colored noise ($\alpha > 0$, spread eigenvalues) without a noise-only calibration interval.

Single-snapshot pre-whitening. When α exceeds a threshold α_0, the noise covariance is estimated directly from the noise subspace as $\hat{\mathbf{R}}_n = \mathbf{U}_n \mathbf{\Lambda}_n \mathbf{U}_n^H$. The snapshot is then pre-whitened ($\tilde{\mathbf{x}} = \hat{\mathbf{R}}_n^{-1/2} \mathbf{x}$), and the AD covariance is recomputed from $\tilde{\mathbf{x}}$. In the whitened domain, the noise eigenvalues become uniform ($\alpha \approx 0$), restoring the matched condition and the full group gain.

The complete AD-ESPRIT procedure operates from a single snapshot in four steps: (0) compute the AD covariance; (1) eigendecompose, estimate K via AIC/MDL, compute α, and if $\alpha > \alpha_0$ then estimate $\hat{\mathbf{R}}_n$, pre-whiten, and re-eigendecompose; (2–3) run standard ESPRIT on the resulting signal subspace. When noise is white, step 1 skips pre-whitening and the procedure reduces to a simple covariance substitution followed by standard ESPRIT.

Results

We simulate a ULA with $M = 16$, half-wavelength spacing, $K = 3$ sources at $\theta = \{-15^\circ, 5^\circ, 25^\circ\}$, and SNR = 6 dB over

500 Monte Carlo trials.

White noise (Fig. 1, top row). Panel (a) shows DOA estimates from 80 trials for three configurations. Conventional ESPRIT at $L = 1$ (red) produces widely scattered estimates, with many trials yielding DOAs far from the truth. AD-ESPRIT at $L = 1$ (green) clusters tightly around the true DOAs (gold dashed), comparable to conventional ESPRIT at $L = 16$ (blue). Panel (b) shows RMSE vs. L: AD-ESPRIT from one snapshot (RMSE $= 3.4°$) matches conventional at $L \approx 4$. Panel (c) shows the success rate: conventional ESPRIT at $L = 1$ succeeds only 42% of the time because the rank-1 covariance cannot support a 3-dimensional signal subspace; AD-ESPRIT succeeds 100% because the rank-lift provides rank $M = 16$.

Colored noise (Fig. 1, bottom row). With spatially correlated AR(1) noise ($\rho = 0.9$), conventional ESPRIT at $L = 1$ degrades further to 38% success because the non-uniform noise eigenvalues corrupt the signal/noise subspace boundary. AD-ESPRIT maintains 100% success with RMSE $= 4.7°$. The slight RMSE increase from $3.4°$ reflects the nonzero coloring index ($\alpha > 0$): the noise covariance is not exactly circulant. Panel (c) compares all three configurations at 6 dB, confirming that AD-ESPRIT from one snapshot outperforms conventional ESPRIT at $L = 1$ by a factor of $3.5\times$ in RMSE while achieving 100% vs. 38% success.

Significance

The rank-lift enables two distinct capabilities from a single snapshot: signal subspace estimation for DOA extraction, and noise subspace characterization for automatic colored-noise detection and pre-whitening. This dual role does not exist in any prior single-snapshot approach. The method is a drop-in replacement applicable to any subspace method (MUSIC, root-MUSIC, unitary ESPRIT) that takes a covariance matrix as input, requiring no modification to the downstream algorithm.

References

[1] R. Roy and T. Kailath, "ESPRIT—Estimation of signal parameters via rotational invariance techniques," *IEEE Trans. ASSP*, vol. 37, no. 7, pp. 984–995, 1989.

[2] M. A. Thornton, "Algebraic diversity: Group-theoretic spectral estimation from single observations," arXiv:2604.03634, 2026.

[3] M. A. Thornton, "Single-snapshot spectral estimation via algebraic diversity," *Proc. 60th Asilomar Conf.*, 2026, submitted.

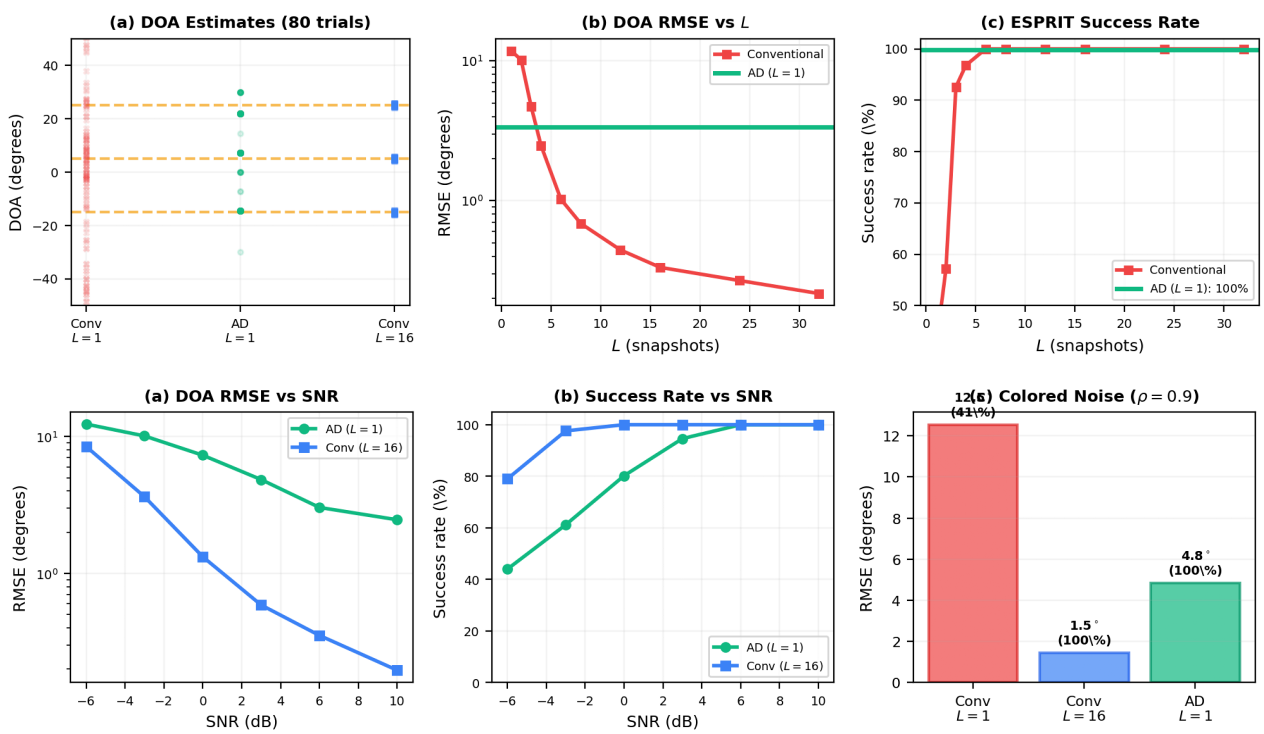

Fig. 1: Top: White noise. (a) DOA estimates from 80 trials. (b) RMSE vs L: AD matches conventional at $L \approx 4$. (c) Success rate: AD 100% vs. conventional 42% at $L = 1$. Bottom: (d) RMSE vs SNR. (e) Success rate vs SNR. (f) Colored noise ($\rho = 0.9$) at 6 dB: AD maintains 100% success.