Single-Snapshot Spectral Estimation 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.6 Robust Methods

***Abstract*—The law of large numbers requires L independent snapshots to reduce the variance of a sample statistic by $1/L$. This paper shows that this requirement is an artifact of ignoring the algebraic structure present in the data. By replacing temporal averaging with a group-averaged estimator that acts on a single snapshot vector $\mathbf{x} \in \mathbb{C}^M$, we achieve a variance reduction factor of $d_{\text{eff}} \le |G| \le M$ from one snapshot. The result is a (G, L) continuum of estimation strategies: one algebraically diverse (AD) snapshot at $d_{\text{eff}} = M$ is statistically equivalent to M conventional snapshots at $d_{\text{eff}} = 1$. Conventional processing is the special case $G = \{e\}$ (the trivial group). However, the group must be matched to the signal; a mismatched group is not merely unhelpful but destructive. We introduce the spectral concentration ψ as a single snapshot diagnostic that distinguishes matched from mismatched groups without prior knowledge of the signal. Simulations with a periodic signal and a chirp signal at $M = 64$ demonstrate both the matched-group advantage and the mismatch diagnostic.**

Extended Summary

Problem and Motivation

Spectral estimation from limited snapshots is a central bottleneck in array processing. The sample covariance matrix $\hat{\mathbf{R}}_L = \frac{1}{L}\sum_\ell \mathbf{x}_\ell \mathbf{x}_\ell^H$ has rank $\min(L, M)$ and variance proportional to $1/L$. Subspace methods such as MUSIC and ESPRIT typically require $L \ge 2M$ snapshots for reliable operation. In snapshot-limited environments (radar, sonar, single-trial BCI), this requirement is often the binding constraint on system performance.

Approach: Group-Averaged Estimation

This paper presents a fundamentally different approach based on the algebraic diversity (AD) framework [1]. Rather than averaging over L independent snapshots, we exploit the algebraic structure within a *single* snapshot. A finite group G with unitary representation $\pi : G \to U(M)$ acts on the snapshot vector $\mathbf{x}$, producing $|G|$ group-rotated copies. The *group-averaged* estimator averages a statistic f over these copies:

$$\hat{\mathbf{\Theta}}_G = \frac{1}{|G|}\sum_{g\in G} \pi_g(\mathbf{x})\,\pi_g(\mathbf{x})^H. \tag{1}$$

When the group is *matched* to the signal (i.e., the signal covariance commutes with the group representation), the variance of each entry of $\hat{\mathbf{\Theta}}_G$ is reduced by a factor of $1/d_{\text{eff}}$, where $d_{\text{eff}} \le |G|$ is the effective group order. For the cyclic group $\mathbb{Z}_M$, this yields a factor of $1/M$—equivalent to averaging M independent snapshots.

Key Results

Trivial Group Embedding (TGE). Conventional temporal averaging is the special case $G = \{e\}$ of a general (G, L) continuum. The combined estimator has variance $\mathrm{Var}/(d_{\text{eff}} \cdot L)$. Conventional processing operates at $d_{\text{eff}} = 1$; AD operates at $d_{\text{eff}} = M$. The law of large numbers is what happens when the trivial group is used.

Rank-lift. The group-averaged estimator produces a rank-M spectral estimate from a single snapshot, whereas the sample covariance has rank 1. This enables single-snapshot subspace detection.

Group gain. The SNR improvement is $\mathcal{G}(G, f) = 10\log_{10}(d_{\text{eff}})$ dB. For $M = 64$, this is 18.1 dB—identical to the array gain of a 64-element ULA beamformer. This identity is not coincidental: beamforming is the group-averaged estimator for $\mathbb{Z}_M$ acting on the spatial steering vector.

Spectral concentration diagnostic (ψ). The group must be matched to the signal. A mismatched group is destructive: it can increase MSE by a factor of 4 compared to no processing. We introduce $\psi(G, \mathbf{x}) = \max_k |\lambda_k|^2 / \sum_k |\lambda_k|^2$ as a single snapshot diagnostic. High ψ indicates a matched group; low ψ indicates mismatch. This is computable from one snapshot with no prior knowledge of the signal.

Simulation Results

We demonstrate both aspects on two signals at $M = 64$, SNR = 6 dB.

Example 1: Periodic signal (matched $\mathbb{Z}_M$). Fig. 1(top) shows that a single AD snapshot matches the waveform MSE of 4 conventional snapshots and produces a full-rank spectral estimate resolving all three tones.

Example 2: Chirp signal (mismatched $\mathbb{Z}_M$, matched $C_\beta \cdot \mathbb{Z}_M$). Fig. 1(bottom) shows that the cyclic group destroys the chirp (MSE = 1.0, worse than raw 0.25), while the conjugated cyclic group recovers it (MSE = 0.28). The ψ diagnostic cleanly separates the two groups ($\psi = 0.04$ vs. 0.39) without knowledge of the true signal.

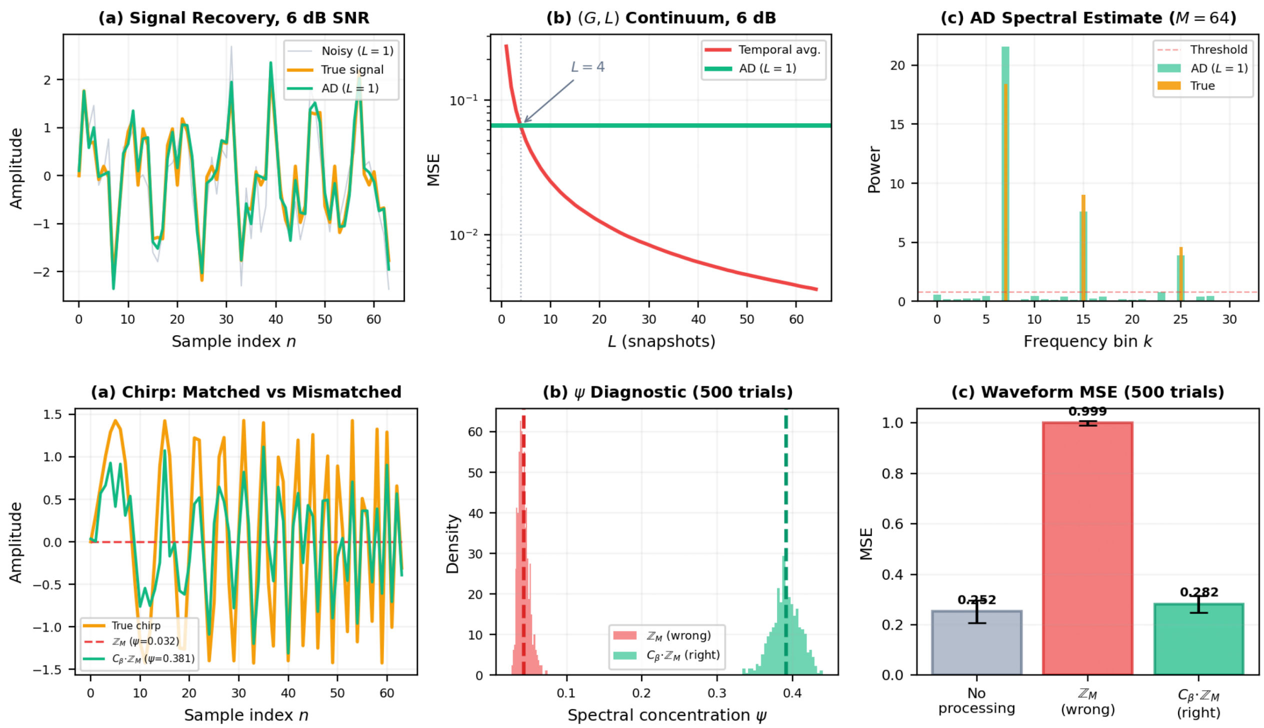

Fig. 1: Top: Matched group ($\mathbb{Z}_M$) for a 3-tone periodic signal at 6 dB. (a) Waveform recovery. (b) (G, L) continuum crossover at $L = 4$. (c) Spectral estimate. Bottom: Mismatched ($\mathbb{Z}_M$) vs. matched ($C_\beta \cdot \mathbb{Z}_M$) for a chirp at 6 dB. (d) Waveform. (e) ψ diagnostic. (f) MSE comparison.

Significance

The paper establishes that the DFT, DCT, beamforming, and temporal averaging are all points on a single (G, L) continuum parameterized by the algebraic group. The law of large numbers is the degenerate $d_{\text{eff}} = 1$ case. The spectral concentration ψ provides a practical, computable diagnostic for group selection. This framework is relevant to any snapshot-limited estimation problem in array processing, radar, sonar, and communications.

References

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