

Subspace Overlap Predicts Residual Sensing Feasibility in OFDM ISAC Systems

Ramakrishna Pasupuleti

workisfun415@gmail.com

Independent Researcher <https://orcid.org/0009-0008-8418-1430>

Research Article

Keywords: ISAC, OFDM, MMSE, SVD, subspace overlap, residual sensing, 6G

Posted Date: July 9th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10293579/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Abstract

We introduce the subspace overlap metric as a quantitative predictor of whether MMSE equalization residuals in OFDM integrated sensing and communication systems can be exploited for sensing. Across twelve channel models including five 3GPP clustered delay line profiles, residual artifacts are universally low-rank but separability from target structure depends on subspace geometry. When overlap is below 0.05, rank-1 PCA achieves positive differential; above 0.10, separation fails. The proposed K–R approach outperforms all baselines including matched filtering with perfect channel knowledge, validated over 500 Monte Carlo trials at $p < 0.001$. A blind adaptive rank selection algorithm is provided.

I. INTRODUCTION

INTEGRATED sensing and communication (ISAC) is a foundational 6G capability, formally adopted at 3GPP RAN #108 [1]. Current receivers use parallel sensing chains requiring dedicated resources [2], [3]. We ask: can the MMSE equalization residual—conventionally discarded—serve as a sensing resource?

The LISAC framework [4] exploits residuals for communication improvement but not sensing. Data-aided receivers [5] process the full signal. We show the residual contains target structure, but its exploitability depends on artifact-target subspace geometry—quantified by a new metric η_1 . We further provide a blind algorithm for adaptive rank selection that requires no oracle knowledge of the target subspace.

Contributions: (C1) universal low-rank residual artifacts across 12 channels; (C2) the subspace overlap metric η_1 as a statistically significant predictor ($p = 0.043$) whose explanatory power increases to $R^2 = 0.55$ in a multi-variable model; (C3) a blind eigengap-based adaptive rank algorithm extending feasibility to 11/12 channels; (C4) K–R PCA-1 outperforms all baselines including matched filtering—explained by the MMSE equalizer’s ~ 20 dB interference cancellation. Preprint: [6].

II. SYSTEM MODEL AND K–R RECEIVER

Consider OFDM at $f_c = 28$ GHz, $N_{sc} = 256$, $N_{sym} = 64$, $\Delta f = 120$ kHz (NR FR2 [7]). Received: $Y[k, l] = (H_{comm} + H_{target})X[k, l] + W$, where H_{target} is 20–40 dB below H_{comm} . Pilots every 4th subcarrier; LS + interpolation + LMMSE smoothing [8]. The sensing residual $r_s = (Y - \hat{H}\hat{X}_{hard})\hat{X}_{hard}^*/(|\hat{X}_{hard}|^2 + \epsilon)$ decomposes as $r_s = r_{artifact} + r_{target} + r_{noise}$.

K–R Receiver: K-Step = standard MMSE (unmodified); R-Step = SVD, adaptive PCA-k removal (Algorithm 1), 2D-FFT $\rightarrow$ range-Doppler map. Twelve channels tested: 7 parametric + 5 CDL profiles (CDL-A through E).

III. SUBSPACE OVERLAP THEORY

A. Universal Low-Rank Structure

SVD of the artifact-only residual reveals $\sigma_1/\sigma_2 > 1.5$ for all 12 channels (Table I), with the first singular value capturing 28–45% of total energy. The physical origin is frequency-domain interpolation acting as a low-pass filter on the estimation error, producing smooth, correlated artifacts concentrated in the first singular vector regardless of channel structure [8], [13].

TABLE I: SUBSPACE ANALYSIS ACROSS 12 CHANNELS

Ch	σ_1/σ_2	η_1	k^*	Δ^*	CI	DS	Np	
UMi	4.47	.009	1	+1.9	± 3	52	3	✓
S-NLOS	2.68	.127	1	−0.7	± 4	85	3	✗
W-NLOS	4.61	.030	0	+1.0	± 3	42	3	✓
Many-p	4.19	.025	1	+1.1	± 3	68	5	✓
Hi-Dop	1.66	.031	4	+0.7	± 3	50	3	✓
LOS	4.64	.036	0	+0.5	± 3	0	1	✓
Rician	4.62	.039	0	+1.1	± 3	36	3	✓
CDL-A	4.08	.017	1	+2.7	± 4	160	6	✓
CDL-B	4.13	.011	1	+2.5	± 4	110	4	✓
CDL-C	3.36	.025	1	+4.5	± 5	310	5	✓
CDL-D	4.58	.025	0	+0.3	± 2	28	3	✓
CDL-E	4.62	.042	1	+0.3	± 2	18	3	✓

B. Subspace Overlap Metric

We define $\eta_k = \|\mathbf{U}_{A,k} \mathbf{U}_T\|_F^2 / \min(k, r_T)$, where $\mathbf{U}_{A,k}$ contains the first k artifact singular vectors and $\mathbf{U}_T$ the target singular vectors. When $\eta \approx 0$ subspaces are orthogonal; when $\eta \rightarrow 1$, entangled. Single-variable regression across 12 channels (fixed SNR = 20 dB, gain = −10 dB) yields $\Delta^* \approx 2.25 - 26.65 \cdot \eta_1$ ($R^2 = 0.35$, $p = 0.043$). Expanding to 180 data points across 5 SNRs and 3 target gains, a multi-variable model incorporating η_1 , delay spread, target gain, number of paths, Doppler spread, Rician K-factor, and SNR achieves $R^2 = 0.55$ (adjusted $R^2 = 0.53$). The top three predictors by individual R^2 are: delay spread (0.19), target gain (0.14), and number of paths (0.12). Thus, η_1 captures the geometric separability component, while propagation structure and operating conditions explain the remainder. Fig. 1 shows the predictive relationship.

The key insight: σ_1/σ_2 does not predict performance—CDL-C has the lowest ratio (3.36) yet the highest differential (+4.5 dB). Feasibility depends on subspace geometry (η_1), not low-rankness.

C. Blind Adaptive Rank Selection

A practical implementation cannot assume knowledge of the target subspace $\mathcal{U}_T$. We propose Algorithm 1, which selects k^* using only the eigengap structure of the observed residual—no oracle required:

Algorithm 1: Blind Adaptive Rank Selection

Input: Observed residual R_s

```

1:  $[U, \Sigma, V] \leftarrow \text{SVD}(R_s)$ 
2: for  $k = 0, 1, \dots, K_{\max}$  do
3:    $\text{gap}(k) \leftarrow \sigma_k / \sigma_{k+1}$ 
4:  $k^* \leftarrow \arg\max \text{gap}(k)$ 
5:  $R_{\text{clean}} \leftarrow \sum_{i > k^*} \sigma_i u_i v_i^H$ 
6:  $\text{RDM} \leftarrow |\text{FFT2}(R_{\text{clean}})|^2$ 

```

Output: Range-Doppler map, k^*

The eigengap identifies where the artifact energy drops off sharply, enabling blind rank selection. This rescues three channels: High Doppler ($k^*=4$, +1.48 dB), Weak NLOS ($k^*=0$, +1.21 dB), LOS ($k^*=0$, +0.63 dB). With Algorithm 1, 11/12 channels achieve positive differential.

IV. RESULTS

A. Baseline Comparison

Under UMi, SNR = 20 dB, target −10 dB (Table II), K–R PCA-1 achieves +2.11 dB. This superiority over matched filtering with perfect CSI (+0.35 dB) is explained by MMSE interference cancellation: the equalizer removes approximately 20 dB of communication signal power, reducing the target-to-interference ratio in the residual from −20 dB (in Y) to approximately 0 dB. The matched filter operates on Y where the target is buried 20 dB below the communication channel; PCA-1 operates on the residual where most of that interference is already cancelled. PCA-1 does not outperform the theoretically optimal detector—it outperforms the matched filter because MMSE pre-cancels interference.

TABLE II: BASELINE COMPARISON (UMI, SNR=20 DB, $N \geq 200$)

Method	Diff	CI	p	MMSE	
Conv. 2D-FFT	-0.47	± 0.01	<.001	No	⊗
Matched (CSI)	+0.35	± 0.09	<.001	No	⊗
Res. PCA-0	-3.39	± 0.15	<.001	Yes	⊗
PCA-1 (ours)	+2.11	± 0.25	<.001	Yes	✓
Res. PCA-5	-1.10	± 0.33	<.001	Yes	⊗

Over 500 trials, all target gains -5 to -30 dB achieve $p < 0.001$. Peak: $+2.12 \pm 0.16$ dB at -10 dB. MIMO scaling: SISO $+1.62 \rightarrow 4 \times 4 +2.53$ dB (CI narrows 60%). Single-frame CFAR remains unachievable ($H1/H0 \approx 1.0$); multi-frame accumulation bridges the gap to $P_d = 0.98$ at $P_{fa} = 0.01$ with 50 frames (27 ms), detailed in [6]. Fig. 2 visualizes the baseline comparison.

B. Complexity

Benchmarked on a modern CPU: power iteration (rank-1 SVD) costs 0.08 ms versus 1.2 ms for full SVD—93% cheaper. The complete R-step (power iteration + FFT) adds 0.23 ms, representing 27% overhead on the 0.85 ms communication baseline (channel estimation + MMSE). Total K–R receiver: 1.08 ms/frame. Algorithm 1 requires one full SVD per channel-change event, not per frame.

V. CONCLUSION

We introduced η_1 as the first subspace-geometry-based predictor of residual sensing feasibility in OFDM ISAC. Low-rank artifacts are universal across 12 channels, but separability depends on η_1 : below 0.05 enables $+1$ to $+4.5$ dB differential; above 0.10 fails. A multi-variable model incorporating propagation features achieves $R^2 = 0.55$ across 180 operating points. The blind Algorithm 1 selects optimal rank without oracle knowledge, extending feasibility to 11/12 channels. K–R PCA-1 outperforms all baselines because MMSE pre-cancels ~ 20 dB of communication interference. These findings establish the subspace conditions for residual-based sensing without dedicated hardware.

References

1. 3GPP RAN #108 (2025) New SID: Study on ISAC for NR, RP-251861, Jun
2. Liu F et al (2020) Joint radar and communication design. IEEE Trans Commun 68(6):3834–3862
3. Koivunen V et al (2024) Multicarrier ISAC. IEEE Signal Process Mag 41(5):17–30
4. Bian C et al (2025) LISAC, arXiv:2410.10711
5. Temiz M et al (2025) Data-aided ISAC, arXiv:2505.01280
6. Pasupuleti R (2026) K–R ISAC analysis, Zenodo, DOI: 10.5281/zenodo.21201976
7. 3GPP TR (2022) 38.901, v17.0.0

8. Edfors O et al (1998) OFDM channel estimation by SVD. IEEE Trans Commun 46(7):931–939
9. Xiong Y et al (2023) Fundamental tradeoff of ISAC. IEEE Trans Inf Theory 69(9):5723–5751
10. Sturm C, Wiesbeck W (2011) OFDM radar, Proc. IEEE, vol. 99, no. 7, pp. 1236–1259
11. Zhang JA et al (2021) Signal processing for joint comm and radar. IEEE JSTSP 15(6):1295–1315
12. Lu S et al (2024) ISAC: Ten challenges, IEEE IoT-J. 11(11):19094–19120
13. Ozdemir MK, Arslan H (2007) OFDM channel estimation. IEEE COMST 9(2):18–48
14. Liu F et al (2022) ISAC for 6G, IEEE JSAC, vol. 40, no. 6, pp. 1728–1767
15. González-Prelcic N et al (2024) ISAC revolution, Proc. IEEE, vol. 112, no. 7, pp. 676–723
16. Bomfin R, Chafii M, ZP-OFDM ISAC, (2024) IEEE TCOM 72(10):6304–6319
17. Wei Z et al (2023) ISAC techniques. IEEE IoT-J 10(13):11621–11666
18. Mura S et al (2025) OFDM ISAC waveform, IEEE TWC, vol. 24, no. 6, pp. 5241–5257
19. Yuan W et al (2024) OTFS to DD-ISAC. IEEE WCM 31(6):152–160
20. 3GPP TR (2024) 22.837, Release 19

Figures




Figure 1

η_1 predicts residual sensing feasibility ($R^2 = 0.35$ at fixed conditions, $R^2 = 0.55$ in multi-variable model with 180 points). Circles: parametric channels. Squares: 3GPP CDL. Green: $\eta_1 < 0.05$ (feasible). Red: $\eta_1 > 0.10$ (fails).




Figure 2

Baseline comparison: K-R PCA-1 (+2.11 dB) outperforms all methods. The MMSE equalizer cancels ~20 dB of communication interference, giving the residual far better target visibility than the raw signal. Error bars: 95% CI.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [WCLResultsRegression.json](#)
- [WCLSupplementaryREADME.md](#)