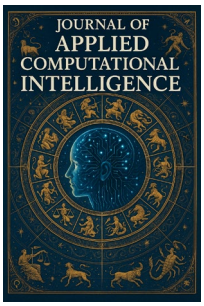
UNIVERSITY OF WEST ATTICA

eISSN: 3136-6937

Journal of Applied Computational Intelligence

Vol. 4, No. 2, April, 2026



Noise Covariance Adaptive OAMP-Net for MIMO-OFDM Detection under Colored Interference

Hsin-hao Kuai[1], Mei-yu Teng[1], Mei-hao Ou[2,*]

[1] *College of Electrical Engineering and Computer Science, National Taiwan University of Science and Technology, Taipei, 10607, China*

[2] *School of Information Science and Technology, Tamkang University, New Taipei, 25137, China*

Received 29 December 2025; Revised 19 April 2026; Accepted 27 April 2026; Published 30 April 2026

Abstract

Spatially colored interference alters the geometry of the residual signal in a MIMO-OFDM receiver. A detector calibrated only with the average noise power may assign excessive confidence to interference-dominated directions. A noise-covariance adaptive orthogonal approximate message-passing network with pilot-residual estimation, shrinkage regularization, shared whitening and confidence-gated feedback is presented in this paper. Trace-normalized linear estimators and finite-alphabet posterior denoisers are used to keep an explicit dependence on the channel, and bounded-layer parameters regulate the step-size, damping and variance correction. The numerical dataset is a constructed demonstration of the reporting protocol across interference correlation, pilot availability, covariance drift and channel mismatch; independent detector measurements and completed training outcomes remain to be established. The reported bit error rates at a signal-to-noise ratio of 18 dB for covariance adaptation, frozen whitening and a white noise network are 0.0062, 0.0128 and 0.0230, respectively. Increasing pilot support from 16 to 128 snapshots reduces normalized covariance error from 0.32 to 0.09. Confidence-threshold analysis identifies a trade-off between residual contamination and accepted sample support, and the lowest reported error rate is at an intermediate threshold. The architecture provides a structured interface between covariance estimation and unfolded inference. The evaluation framework links detection error to conditioning, confidence calibration, transient recovery and computational cost, enabling analysis of receiver behavior under changes in the spatial disturbance metric.

Keywords: Noise Covariance Adaptation, Colored Interference, Orthogonal Approximate Message Passing, MIMO-OFDM Detection, Model-Driven Learning

1. Introduction

Recent MIMO-OFDM detection research exploits waveform structure within accelerated expectation-maximization and deep-unfolding receivers [1]. Spatial interference correlation estimation can exploit angle-of-arrival information to recover the directional structure of the disturbance [2]. Graph-based unfolding models represent channel and interference relationships explicitly in multiuser MIMO networks [3]. Joint channel and interference-plus-noise covariance estimation has been developed for co-channel interference rejection in OFDM systems [4]. Frequency-selective interference also motivates examining whether a covariance estimated over one resource group remains representative of neighboring observations. A practical detector must accommodate this structure with limited pilots, uncertain channel estimates and a finite computational budget. Limited training observations constrain the accuracy of individual covariance entries, while interference can change before a long averaging window becomes statistically beneficial.

* Corresponding author: meihao.o@sit.tku.edu.tw

http://dx.doi.org/10.68406/jaci.2026.vol4iss2no6:74-89

➢ This is an open access article under the CC BY-NC-SA 4.0 license (https://creativecommons.org/licenses/by/4.0/).

© Authors retain all copyrights.

These conditions create a receiver design problem where estimation reliability, detection accuracy and adaptation speed are coupled. The engineering objective is to determine how much disturbance structure can be inferred from current observations and how strongly that information should influence symbol detection.

Iterative model-based detectors provide a useful foundation because their internal states retain an explicit relationship to the channel and transmitted constellation. Orthogonal approximate message passing has been analyzed for spatially coupled linear models with right-orthogonally invariant measurements [5]. He et al. [6] developed a generalized projected-gradient method and unfolded it into a learned detector to improve massive MIMO detection performance and convergence efficiency. Model-driven and data-driven MIMO detection methods have been developed for hardware impairments using noisy training labels [7]. Scalar-noise calibration reduces the directional uncertainty to a single power estimation. If the two disturbance covariances have the same trace but different eigenvectors or eigenvalue spreads, a scalar variance gives both the same description despite their different effects on spatial separation. A learned step size can reduce mismatch within the training distribution, while an explicit spatial metric identifies unreliable receive directions. Adding a raw sample covariance may increase estimation errors in the case of a small snapshot count. Therefore, a reliable adaptation mechanism needs to model colored interference and reduce the uncertainty in poorly observed covariance directions.

Several strands of prior research address components of this problem. Shrinkage estimation has also been formulated for distributions on the manifold of symmetric positive-definite matrices [8]. Residual-based noise covariance estimation in high-dimensional linear models requires correction for bias associated with estimated regression coefficients [9]. Partial-learning MIMO detectors combine learned symbol decisions, interference cancellation and conventional detection within an iterative soft-information framework [10]. Control of the feedback dynamics is required for their integration. The covariance estimator changes the metric used by every layer, while the denoiser modifies the residual supplied to the estimator; the resulting feedback loop may reinforce erroneous decisions. Another problem is that covariance normalization conceals an actual rise in disturbance power and thus provides the detector with an accurate covariance shape but an incorrect power scale. Research into robust unfolded inference focuses on the requirement to investigate these discrepancies outside of the nominal training conditions. Evaluation protocols of adaptive receivers also differentiate between stationary accuracy and behavior immediately after an environmental change. The research problem concerns the integration of regularized covariance estimation, model-consistent whitening and guarded feedback, with explicit treatment of the assumptions underlying finite-dimensional unfolded inference.

The receiver architecture developed in this study integrates pilot-based covariance estimation and unfolded symbol inference into a common processing system to adapt the observation metric to changes in colored interference and maintain the reliability of decision feedback. Both the received vector and its corresponding channel matrix are whitened in the same way. Then, the unfolded layers use a trace-normalized linear update and a constellation-aware posterior denoiser with bounded learned parameters. The covariance state is treated as a receiver input with its own support, conditioning and update constraints. This separation allows the impact of directional disturbance information to be tracked in the inference layers. The analysis distinguishes the effects of pilot support, feedback selection and covariance reuse to enable the evaluation of architectural choices under specific operating conditions. Section 2 relates the design to existing detection and covariance estimation methods and identifies the assumptions requiring validation. Section 3 introduces the observation model, covariance construction and whitening interface. Section 4 presents the unfolded inference architecture, reliability controls and proposed training procedure. Section 5 presents the evaluation protocol and analyzes detection, covariance estimation, adaptation and computational cost. Section 6 presents the conclusions, limitations and directions for future research and independent verification. The contribution lies in coordinating covariance estimation, shared whitening and guarded feedback within a compact unfolded receiver. These components are established individually; their coupling defines the architectural contribution examined here.

2. Related Work

2.1. Model Based and Unfolded Spatial Detection

Predictable computation is provided by linear spatial equalizers; however, the uncertainty of finite-alphabet symbols has been reduced to that of a roughly continuous estimation problem. Performance may deteriorate when the channel is poorly conditioned or interference aligns with a useful spatial mode. Successive cancellation introduces discrete decisions in the processing chain; however, the order of these decisions and the reliability of early estimates are also design parameters. Comparative MIMO detection research evaluates learned, approximate and linear detectors in terms of detection performance and computational complexity [11]. These comparisons motivate including a covariance-aware linear reference alongside the white-noise baseline in subsequent implementations. The evaluation separates the contributions of the observation metric and learned inference. This distinction is particularly relevant when multiple subcarriers share a covariance state and the factorization cost is amortized over a resource group.

Orthogonal approximate message passing employs the relationship between a linear estimator and a nonlinear estimator to control the spread of estimation errors. Asymptotic arguments rely on matrix ensembles, estimator properties and dimensional limits whose applicability to a practical multicarrier channel requires examination. Capacity-optimality analysis of OAMP explicitly considers sensing-matrix structure and signaling assumptions [12]. Unfolded implementation can maintain the general iteration structure and relax some scalar parameters; the state evolution under a colored, estimated, and time-varying disturbance requires separate analysis. Specifically, whitening an imperfectly estimated covariance alters the channel's singular values and residual statistics. Although a finite-alphabet posterior obtained with a scalar effective variance can be used, the probabilities need to be calibrated. Orthogonal message passing provides the architectural foundation, while the finite-dimensional behavior of the damped learned layers requires separate validation.

Unfolded spatial detection networks are situated between fixed iterative solvers and large unconstrained neural receivers in the current work. Layer-dependent steps, variance corrections and damping coefficients can be learned while maintaining explicit channel dependence. Model-driven unfolded channel estimation has been developed to address interference and computational cost in MIMO backscatter reception [13]. Therefore, these results motivate a compact learned controller whose behavior depends on the covariance conditions encountered during evaluation. A network trained on a single interference eigenbasis may exploit statistical regularities that change when the interferer moves. Therefore, the evaluation should modify both the eigenvectors and correlation magnitudes, and it should present a frozen-parameter test prior to any environment-specific fine-tuning. Matching label access and training conditions across comparators isolates the contribution of adaptation. Beamforming, backscatter estimation and covariance recovery provide related methodological evidence; their results establish background principles rather than direct validation of the present detector.

2.2. Covariance Estimation and Interference Representation

An unrestricted sample covariance may show detailed spatial patterns; however, the reliability of this covariance is affected by the quantity and independence of the snapshots employed. Estimation becomes particularly difficult when the receive dimension approaches the effective snapshot support. Shrinkage covariance estimation from correlated sub-Gaussian samples has been studied through non-asymptotic estimation-error bounds [14]. A diagonal target maintains antenna-specific power and suppresses uncertain off-diagonal terms; an isotropic target reduces both correlation and power imbalance. The target should reflect the intended operating environment. A receiver having significantly uneven front-end noise levels may lose valuable information due to excessive isotropic shrinkage. Retaining all diagonal fluctuations can preserve estimation noise when pilot support is very short. These alternatives should be evaluated with the same observation budget.

Low-rank interference models take advantage of the fact that a small number of dominant interferers may be in only a few receive directions. Low-rank positive-semidefinite covariance matrices can be recovered from structured rank-one measurements with rank-dependent measurement requirements [15]. This assumption becomes less suitable if diffuse interference, filtering artifacts, or several low-power transmitters add significant power. Low-rank models may underestimate the disturbance energy outside the chosen subspace and thus be overly confident in the remaining receive directions. Full-covariance shrinkage is a more conservative option for the first study but has a higher factorization cost. This trade-off requires reporting covariance accuracy, numerical conditioning and the results of an incorrect interference structure selection.

Residual-based adaptation expands the observation set by employing the received data after desired signal reconstruction. This approach reduces reliance on dedicated pilot resources. The primary risk is circularity; an erroneous detector may produce a residual whose covariance exhibits the behavior of desired signal cancellation errors instead of external interference. Beyond decision selection, label-free wireless receiver adaptation has been formulated using physics-informed constraints [16]. Confidence calibration compares posterior scores with observed symbol correctness before those scores control feedback. Correlated channels and mismatched effective variances may result in an overconfident posterior. Therefore, to safeguard the feedback path, alongside a probability threshold, one also needs a minimum accepted sample count, an update rate limit and a fallback state for insufficient current evidence.

2.3. Adaptation and Evaluation Gaps

Online adaptation differentiates between environment estimation and detector training. Updates of covariance employ unlabeled data; parameter retraining uses transmitted labels. Uzlaner et al. [17] proposed modular drift detection to identify when deep receivers require adaptation and which receiver components should be retrained. The proposed architecture fixes the learned layer parameters at evaluation and modifies only the covariance state. This choice makes the proposed source of adaptation traceable. Nonlinear hardware distortion and channel model errors may require additional state variables. If the receiver is exposed to disturbance statistics beyond its second-order description, extra training or a different likelihood model may be necessary.

Evaluation methodology presents a further research gap. Reporting the average error rate over a long test sequence may hide a short period of substantial degradation following an interference change. Beam-tracking comparisons examine model-based adaptive filters and learning-based approaches using time-indexed estimation error and outage-related performance measures [18]. Useful tests for covariance-adaptive detection include short pilot windows, sudden eigenbasis rotation, erroneous channel estimates, changing interference power and acceptance gates that allow for a small proportion of incorrect decisions. Resource allocation, information access and stopping conditions of the comparison must be held constant. The design addresses these gaps with a shared-whitening interface and an intentionally limited feedback loop.

3. Observation Model and Covariance Interface

3.1. Frequency Domain Observations and Disturbance Assumptions

Consider an uncoded multicarrier receiver with several transmit streams and at least as many receive antennas. The cyclic prefix exceeds the modeled channel memory, and synchronization accuracy supports a per-subcarrier observation model. Residual carrier offset, intercarrier leakage, and amplifier nonlinearities remain outside the initial likelihood formulation. The following observation equation describes the received signal on an active subcarrier, with the disturbance combining thermal noise and external interference:

$$\boldsymbol{y}_k = \boldsymbol{H}_k \boldsymbol{x}_k + \boldsymbol{n}_k, \mathbb{E}\left[\boldsymbol{n}_k \boldsymbol{n}_k^H\right] = \boldsymbol{R}_k. \tag{1}$$

Here, $\mathbf{y}_k$ is the received vector, $\mathbf{H}_k$ is the channel matrix, and $\mathbf{x}_k$ contains the transmitted symbols. The disturbance is assumed to have zero mean, and its covariance $\mathbf{R}_k$ is Hermitian positive definite after inclusion of the thermal component. Symbols are drawn independently from a normalized square quadrature amplitude modulation alphabet. The spatial disturbance model accommodates unequal interference power across receive directions. Material coupling between subcarriers requires an enlarged observation vector or a corresponding extension of the likelihood. Neighboring subcarriers share a covariance state when their disturbance statistics remain sufficiently similar. This grouping reduces estimation and factorization overhead, while introducing a potential bias when interference varies appreciably within the group.

Channel knowledge and covariance knowledge are treated as distinct receiver inputs. A separate pilot estimator supplies the channel estimate, and all practical comparators receive the same channel information. The reference disturbance covariance is reserved for the oracle configuration. Covariance shape and total power are stored separately to preserve changes in disturbance strength during normalization. Each state record contains the regularized matrix, its trace scale, a timestamp, the supporting snapshot count, and a numerical condition indicator. These quantities support covariance reuse decisions using information available at the receiver. The nominal sixteen-receive, eight-transmit configuration therefore uses a sixteen-dimensional covariance and a sixteen-by-eight channel matrix. A resource group shares covariance information while retaining its own per-tone channel observations.

3.2. Pilot Residual Estimation and Regularization

Pilot residuals are obtained by subtracting the channel-based reconstruction of known pilot symbols from their received observations. Their outer products are averaged across the selected covariance group and update window. Channel estimation errors contribute to these residuals, giving the resulting matrix an effective disturbance interpretation. This interpretation supports metric construction when the residual statistics remain representative of data reception. Differences between pilot and data power, or changes in the channel estimator between these phases, require corresponding treatment in the receiver model. Sample estimation and regularization are combined in the covariance equation below: Pilot reconstruction errors contribute to the effective disturbance estimate. Pilot-to-data transfer requires compatible power normalization and channel estimation conditions; changing either condition calls for renewed covariance validation.

$$\boldsymbol{S} = \frac{1}{M}\sum_{m=1}^{M} \boldsymbol{e}_m \boldsymbol{e}_m^H, \widehat{\boldsymbol{R}} = (1-\lambda)\boldsymbol{S} + \lambda \frac{tr(\boldsymbol{S})}{N_r}\boldsymbol{I} + \varepsilon \boldsymbol{I}. \tag{2}$$

The residual vector $\mathbf{e}_m$ provides one covariance snapshot, M denotes the snapshot count, and N_r is the receive dimension. The shrinkage coefficient λ lies between zero and one, while the loading parameter ε remains strictly positive. An isotropic target provides a well-conditioned initial design; a diagonal target offers an alternative when antenna-specific power differences carry useful information. Loading scales with average residual power and includes an absolute numerical floor for weak observations. Observable diagnostics, including snapshot support and eigenvalue spread, determine the shrinkage strength or serve as inputs to a bounded learned controller. Each candidate matrix is symmetrized, checked for finite entries, and factorized after its smallest eigenvalue satisfies the configured floor.

Adjacent tones and successive pilot residuals can be dependent, making the effective information content smaller than the nominal snapshot count. The estimation protocol therefore records the grouping policy together with the raw support. Limited independent support increases shrinkage and strengthens retention of the previous valid state. Sudden power changes trigger a controlled assessment of adaptation requirements. Numerical validity and desired-signal contamination are assessed separately through residual energy and posterior reliability. Resource grouping trades statistical support against frequency selectivity. The grouping width and covariance age should be reported together, since a larger snapshot count may include dependent observations or stale interference statistics.

Figure 1 connects the pilot path to the covariance state and identifies the admission point for data feedback. Disabling feedback preserves the forward processing path, allowing the pilot-only ablation to retain the same detector and initialization. Rejected updates preserve the previous valid state and generate a diagnostic record. State replacement occurs at a frame boundary, maintaining a consistent observation metric throughout the unfolded processing of each received vector. This schedule introduces a defined delay between an environmental change and application of the revised covariance.

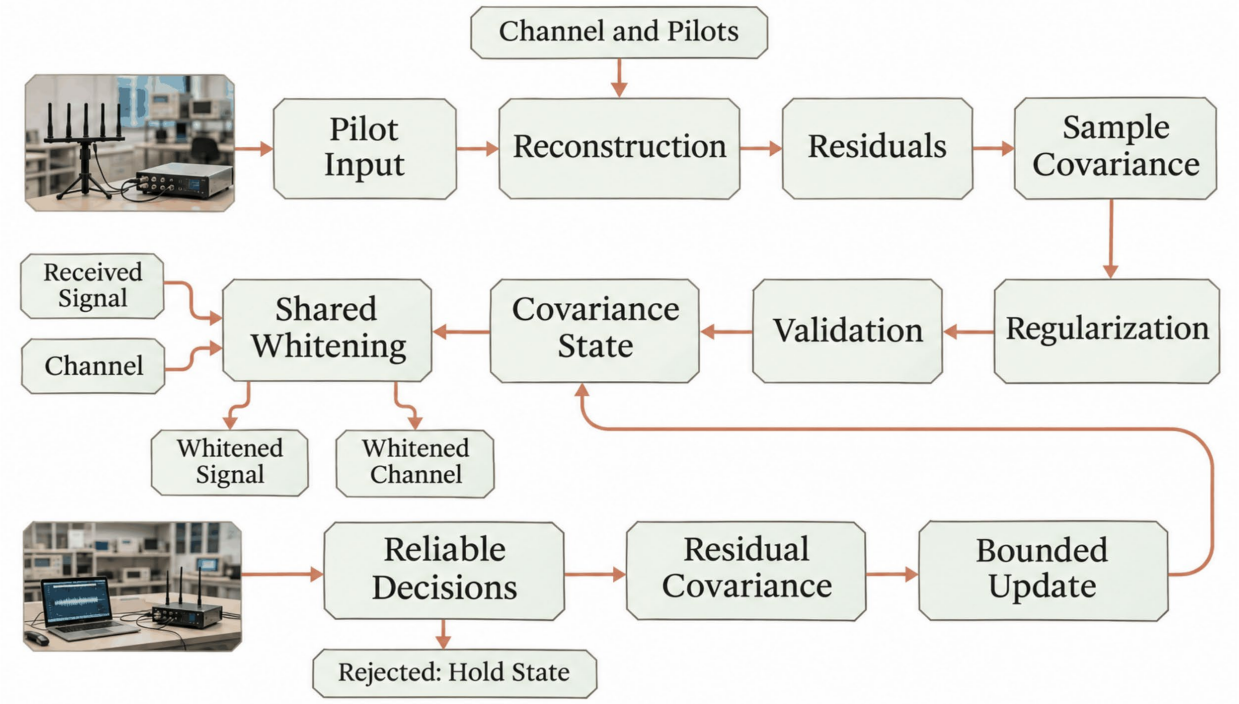


Figure 1. Covariance acquisition and metric construction workflow

3.3. Shared Whitening and Consistency Checks

Whitening applies the inverse Cholesky factor of the regularized covariance to the received vector and its associated channel matrix. Triangular solves provide the required transformation efficiently. Reusing a factor across an admissible resource group reduces overhead while preserving a common metric for the corresponding observations. The shared-whitening equation makes this coupling explicit:

$$\widehat{\boldsymbol{R}} = \boldsymbol{L}\boldsymbol{L}^{H}, \tilde{\boldsymbol{y}} = \boldsymbol{L}^{-1}\boldsymbol{y}, \widetilde{\boldsymbol{H}} = \boldsymbol{L}^{-1}\boldsymbol{H}. \tag{3}$$

The lower triangular matrix **L** defines both transformations. Applying an identical factor preserves the relationship between the received samples and the channel used by the detector. Exact covariance knowledge would produce identity disturbance covariance in the transformed space. Estimation error leaves residual correlation and unequal directional variance, so the unfolded detector retains a bounded effective variance correction. This correction addresses remaining scale mismatch, while the covariance state supplies the directional information. Sensitivity to whitening orientation is examined through changes in the interference eigenbasis. Covariance estimation error leaves a transformed disturbance with residual spatial structure. The bounded scalar variance correction addresses average scale mismatch; directional mismatch remains a separate covariance estimation issue.

Condition checks protect the triangular solve and quantify the regularization required by uncertain covariance directions. An excessively large condition number triggers a more strongly shrunk candidate before the receiver commits the metric. The additional loading is recorded because it changes the spatial weighting applied during detection. Fallback frequency provides a further diagnostic of estimator suitability. Frequent fallback can indicate insufficient pilot support, excessive variation within a resource group, or interference dynamics faster than the selected update interval. Recording

these events together with covariance age and accepted support helps distinguish sampling limitations from delayed tracking. Loading, condition-number limits and reset triggers are implementation parameters that require explicit recording. Their selection is a specification requirement, and the present analysis does not establish a universally optimal numerical setting.

4. Covariance Adaptive Unfolded Detector

4.1. Linear Estimation and Finite Alphabet Denoising

The detector starts from a zero-symbol estimate and an initial uncertainty consistent with the normalized constellation. Each layer forms a residual in the whitened receive space and applies a regularized linear estimator. Trace normalization aligns the average estimator gain with the transmit dimension. A positive regularization parameter reflects the current uncertainty, while a matrix factorization or an iterative solve implements the linear operation under a controlled numerical tolerance. These operations are expressed through the linear-estimation equation:

$$\boldsymbol{B}_t = \widetilde{\boldsymbol{H}}^H\left(\widetilde{\boldsymbol{H}}\widetilde{\boldsymbol{H}}^H + \alpha_t \boldsymbol{I}\right)^{-1}, \boldsymbol{W}_t = \frac{N_t \boldsymbol{B}_t}{tr\left(\boldsymbol{B}_t \widetilde{\boldsymbol{H}}\right)}. \tag{4}$$

Here, t denotes the layer index, N_t is the transmit dimension, and α_t controls regularization. The trace denominator is checked against a positive numerical floor. An observation with negligible channel energy returns a low-confidence prior. Stable linear solves implement the inverse operator, with the factorization dimension determined by the selected algebraic form. Computational accounting distinguishes these operations from covariance estimation and state management. Layer parameters may be shared to reduce storage, although the initial architecture permits distinct bounded parameters at each depth. The regularization and trace floor belong to numerical stabilization, while the learned step controls the subsequent residual correction. Recording these quantities separately distinguishes solver behavior from learned inference.

The linear output supplies constellation posterior probabilities through a scalar effective variance. This representation preserves a transparent finite-alphabet denoiser while allowing uncertainty correction within each layer. Damping controls the influence of the updated posterior means when residual statistics or covariance estimates are imperfect. The resulting symbol-update equation combines the learned linear step with the damped denoising operation:

$$\begin{aligned} \boldsymbol{r}_t &= \boldsymbol{x}_t + \gamma_t \boldsymbol{W}_t\left(\widetilde{\boldsymbol{y}} - \widetilde{\boldsymbol{H}}\boldsymbol{x}_t\right), \\ \boldsymbol{x}_{t+1} &= (1-\beta_t)\boldsymbol{x}_t + \beta_t D(\boldsymbol{r}_t; \tau_t). \end{aligned} \tag{5}$$

The step multiplier γ_t adjusts the linear correction, β_t controls damping, and τ_t is the effective denoising variance. The operator D returns the constellation posterior mean. During probability evaluation, subtracting the minimum squared distance before exponentiation improves numerical stability. Posterior variances support uncertainty propagation, while posterior maxima contribute to feedback confidence. The effective variance depends on current symbol uncertainty, the linear error propagation matrix, and transformed disturbance power, followed by a bounded learned correction and a positive floor. Gaussianity and independence of layer errors remain approximations whose adequacy depends on channel correlation and covariance accuracy. The variance calculation and probability calibration should be assessed separately: a low symbol error rate does not by itself establish reliable confidence scores. The posterior probabilities used by the gate require validation under the same covariance mismatch conditions as detection.

Figure 2 separates inference within a frame from the slower covariance update loop. Every layer processing the current vector uses the same covariance factor, and learned parameters remain fixed during evaluation. The initial benchmark uses a fixed depth to preserve a common iteration budget across covariance configurations. An adaptive stopping rule can subsequently be assessed against this reference under a matched error criterion. Posterior probabilities also provide soft outputs for extensions involving coded reception and decoder feedback.

4.2. Confidence Gated Covariance Feedback

Following detection, hard symbol reconstructions generate candidate residuals for observations satisfying the confidence requirement. A vector-level gate accounts for the influence of an unreliable stream on all receive dimensions of its reconstruction residual. Accepted support and residual energy provide additional controls before an update is admitted. Large, coherent energy increases trigger stronger loading and a bounded response. The frame-boundary update follows the gated covariance equation:

$$\widehat{\boldsymbol{R}}_{next} = \begin{cases} (1-\eta)\widehat{\boldsymbol{R}} + \eta\widehat{\boldsymbol{R}}_{data}, & c \geq c_0 \; and \; M_{acc} \geq M_{min}, \\ \widehat{\boldsymbol{R}}, & otherwise. \end{cases} \tag{6}$$

The confidence score c is the minimum stream posterior maximum for an observation vector. The threshold c_0 controls acceptance, and M_{acc} records the accepted support against the requirement M_{min}. Before mixing, the data residual covariance $\widehat{\mathbf{R}}_{data}$ undergoes shrinkage and loading. The bounded gain η determines its contribution to the next state. A limit on relative matrix change further constrains abrupt metric rotation, while prolonged rejection activates a pilot reset path. All control inputs originate from received observations, pilot information, detector probabilities, and stored receiver state. A gate rejection retains the current covariance for the next frame. Acceptance requires both confidence and sample-support conditions; the minimum stream confidence prevents one unreliable stream from being hidden by high confidence on other streams.

Selection bias remains relevant even when accepted decisions are usually correct. Confidence filtering favors easier observations, whose disturbance realizations may have lower energy than those affecting the complete data stream. Their residual covariance can consequently underestimate overall disturbance power. The architecture therefore combines pilot and accepted data information while recording support and trace-scale changes separately. A strict threshold improves accepted-decision purity but reduces the information available for tracking. A more permissive threshold increases support while exposing the estimator to cancellation errors. Varying the confidence threshold and update gain independently separates these effects and helps identify operating conditions requiring stronger pilot retention. The threshold sweep therefore concerns an end-to-end selection effect, not decision purity alone. A reduction in accepted support can increase covariance age and tracking delay even as the accepted decisions become more reliable.

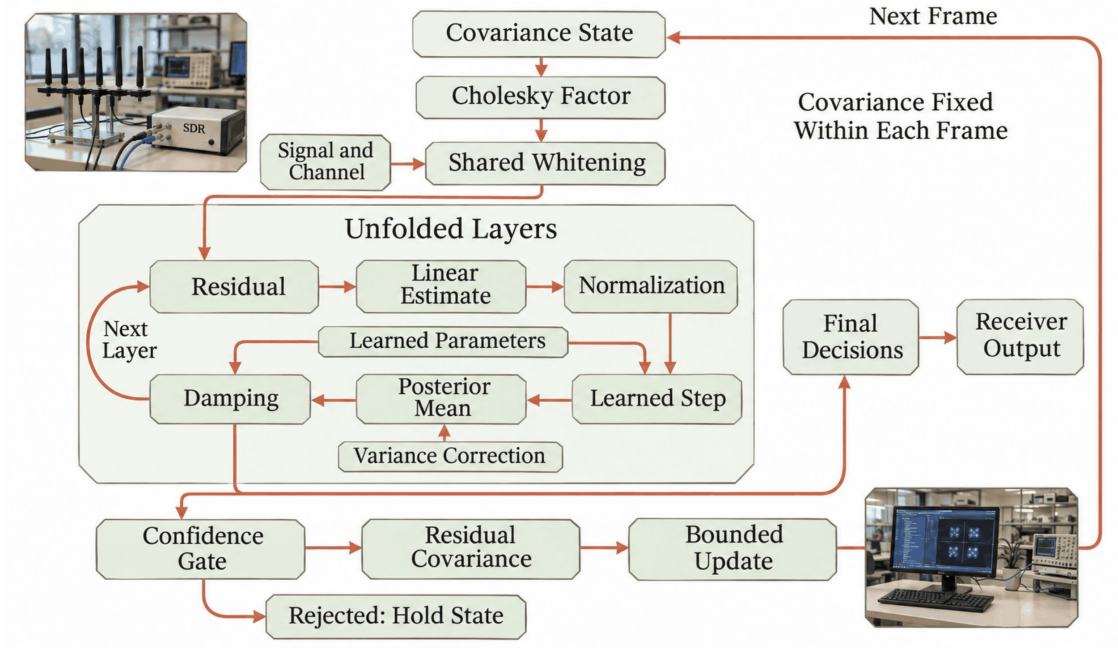


Figure 2. Covariance adaptive unfolded inference workflow

4.3. Proposed Training and Implementation Procedure

Training is organized into two stages. The first uses reference covariance matrices to calibrate the unfolded symbol estimator across channels and disturbance powers. The second introduces pilot-derived covariance estimates and varies snapshot support, shrinkage, channel estimation quality, and interference orientation. A layer-weighted symbol reconstruction objective is combined with a probability calibration term. Validation data determine their relative weights, which are recorded before testing. This procedure links parameter selection to both symbol estimation and the confidence values subsequently used by the feedback gate. These stages specify a proposed training procedure. Optimizer settings, batch sizes, stopping criteria and fitted parameter values are required implementation records; completed training runs are not established by the numerical reporting dataset.

The learned quantities comprise a step multiplier, damping coefficient, and variance correction at each layer, together with a compact bounded covariance controller. Positive parameterizations preserve valid regularization and variance values; while damping and update gains remain within their admissible intervals. Positive definiteness is maintained through the covariance construction and validation procedure. Model selection uses independent validation seeds and includes stationary conditions alongside controlled shifts. Channel realizations define the evaluation split, keeping neighboring tones from the same channel within a common partition. Validation and test partitions should be separated by independent channel and interference realizations. Sharing neighboring tones from one realization across partitions would overstate generalization through correlated observations.

Implementation records include random seeds, constellation normalization, channel generation, covariance families, pilot placement, parameter bounds, solver tolerances, and stopping criteria. Failure to converge within the solver budget is retained as a failed detection event or a documented fallback. Practical comparators receive identical channel estimates and pilot observations. The oracle configuration uses the reference covariance to quantify the gap associated with covariance estimation, while frozen whitening retains its initial estimated state throughout the sequence. Online covariance adaptation updates the receiver state with fixed learned parameters. Maintaining this distinction allows changes in detection behavior to be attributed to the covariance update mechanism under a common information budget.

5. Experimental Protocol and Numerical Analysis

5.1. Benchmark Design and Detection Behavior

The numerical analysis examines the connection between the disturbance metric and symbol detection through complementary operating conditions. The nominal configuration establishes a common spatial dimension, modulation order, pilot budget, and network depth. Separate sweeps then vary the signal to noise ratio, interference correlation, covariance support, confidence threshold, and update policy. Error rates are interpreted alongside covariance accuracy and computational cost so that a change in detection behavior can be associated with a specific receiver mechanism. Stationary tests characterize the quality of the observation metric, while transient tests examine the consequences of covariance drift and delayed updates. Comparators share the same channel inputs and pilot resources, with true covariance information reserved for the oracle reference. Each performance coordinate is retained together with its corresponding configuration to support consistent comparisons across the figures and tables.

The proposed nominal configuration uses eight transmit streams, sixteen receive antennas, 256 active subcarriers, normalized 16-point quadrature amplitude modulation, and eight unfolded layers. A resource group contains sixteen neighboring tones. The pilot support sweep covers 16, 32, 64, 96, and 128 residual snapshots. Nominal receive correlation is 0.7, while its robustness sweep spans 0.1 through 0.9. A complete detector study should test more than one covariance family because a single Toeplitz model can favor methods that implicitly match its structure. Covariance-reconstruction beamforming methods explicitly account for interference-plus-noise structure when suppressing jamming in wireless communication systems [19]. The comparison includes a white noise network, a frozen whitening network, and the proposed adaptive variant.

Table 1 separates fixed nominal settings from the sweeps that probe estimator support and adaptation. A future run should use a minimum number of independent channel blocks in addition to the error-based stopping rule, because many decoded bits within one correlated block do not supply the same information as independently generated blocks. Blockwise bootstrap procedures have been developed for inference with locally stationary time series and time-varying coefficient models [20]. For an executed study, a block bootstrap across independent channel and interference realizations would be preferable to treating every bit as independent when the disturbance is shared across a resource group. The listed stopping criterion defines a prospective sampling rule. Error counts and independent block counts are required to construct uncertainty intervals; the reported rate coordinates alone do not determine their sampling precision.

Table 1. Proposed experimental configuration and reporting obligations

Item	Specified value	Interpretation
Spatial dimensions	8 transmit and 16 receive	Nominal uncoded scenario
Waveform	256 active tones and 16 tone groups	Spatial covariance shared within each group
Constellation and depth	16-point QAM and 8 layers	Fixed depth comparison
Signal to noise ratio	6 to 24 dB in 3 dB steps	Seven operating points
Pilot support	16, 32, 64, 96, 128	Five covariance support settings
Confidence threshold	0.80, 0.85, 0.90, 0.95, 0.98	Five gate settings
Stopping criterion	At least 200 frame errors or 100000 frames	Error based stopping rule

Figure 3(a) reports adaptive error rates of 0.145, 0.086, 0.041, 0.017, 0.0062, 0.0021, and 0.00072 across the seven signal-to-noise ratios. The white noise network has error rates of 0.190, 0.137, 0.086, 0.048, 0.0230, 0.0110, and 0.0055 at the same points. The widening separation at higher ratios illustrates the signature that an interference metric mismatch could produce, and provides a pattern for subsequent empirical verification. Figure 3(b) includes five configurations with rates of 0.0230, 0.0128, 0.0084, 0.0062, and 0.0048 at 18 dB, corresponding to white noise, frozen whitening, pilot only adaptation, gated adaptation, and oracle whitening. The adaptive value is 73.0 percent below the white noise value and 51.6 percent below the frozen value within this comparison.

Sensitivity to interference structure is shown in Figure 3(c), where the five correlation settings of 0.1, 0.3, 0.5, 0.7, and 0.9 have adaptive error rates of 0.0035, 0.0040, 0.0048, 0.0062, and 0.0108. The corresponding normalized covariance errors are 0.08, 0.09, 0.12, 0.16, and 0.25. This scatter view separates the correlation coordinate from the estimator diagnostic and makes the deterioration near the most correlated condition visible. A covariance error averaged over all matrix entries compresses directional information; the positive association should therefore be tested with directional diagnostics in an actual implementation. Figure 3(d) supplies five frame error distributions, each containing twenty values. Their spread grows with correlation so that the upper tail can be inspected alongside the central tendency. Wireless outage-prediction research examines how confidence calibration and classification thresholds affect reliability-related decisions [21]. Estimation of outage probability requires independently sampled trials in addition to these distributions.

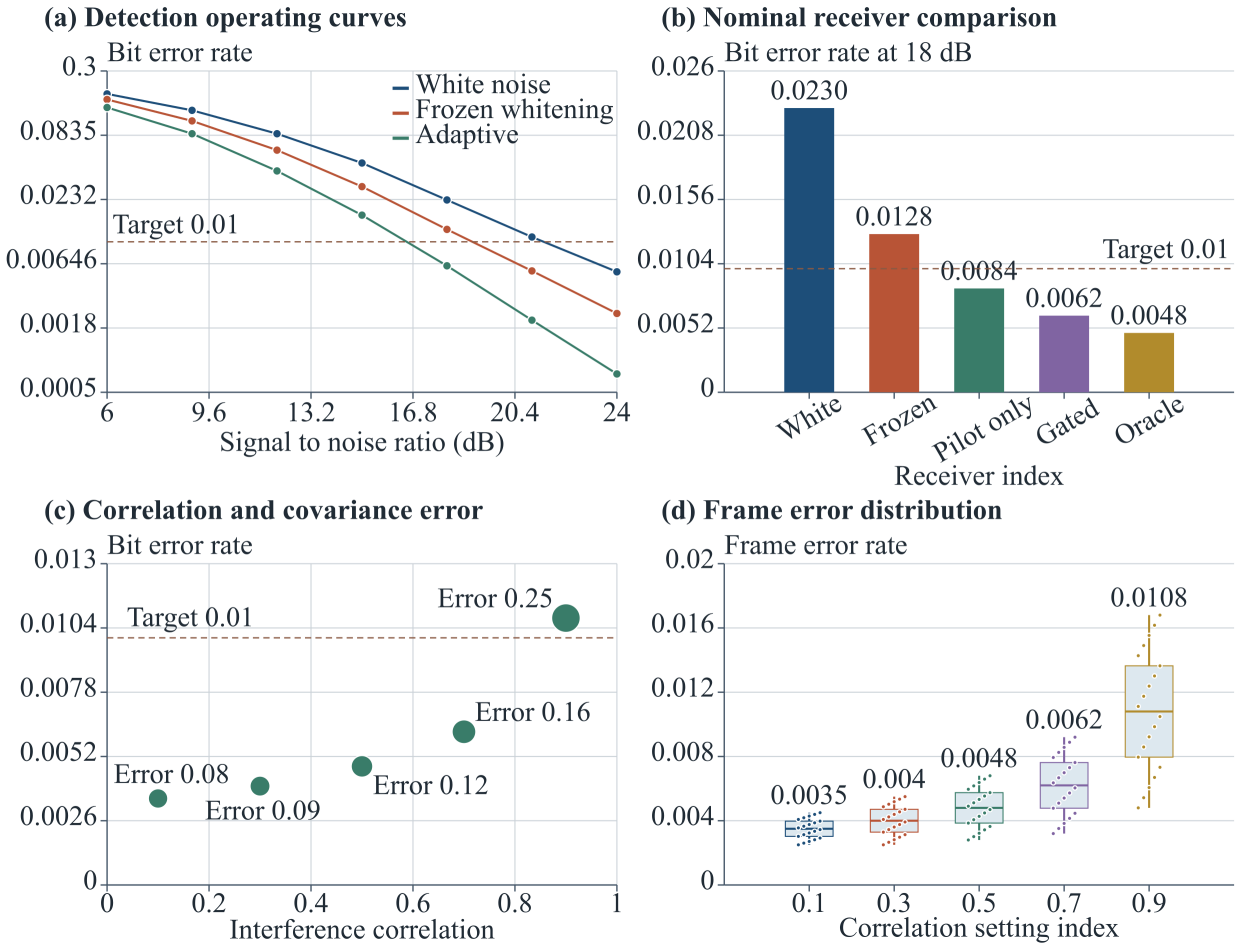


Figure 3. Detection operating characteristics. (a) Bit error rate against signal to noise ratio; (b) bit error rate at 18 dB for five receiver configurations; (c) interference correlation versus bit error rate and normalized covariance error; (d) distribution of frame level error rates under five correlation settings

The operating point sweep should be expanded in an executed study to separate interference power from spatial correlation. Holding total disturbance power constant while changing correlation examines geometry; increasing total power while holding the covariance shape fixed examines scale. Changing both at once can make a scalar noise estimator appear responsible for an error that is actually caused by a directional mismatch. A balanced design would also compare normalized covariance error with transformed residual correlation, because likelihood calibration also depends on the directions of the remaining matrix error. Channel estimation error should be generated with a recorded mechanism rather than injected as an unspecified scalar penalty. These controls determine whether a measured advantage can reasonably be attributed to the covariance interface.

5.2. Covariance Support and Adaptation Mechanisms

Pilot scarcity is a central stress condition because an adaptive receiver cannot rely indefinitely on an accurately measured covariance. Sample support affects the eigenvectors and eigenvalues differently, and strong sample fluctuations can create an apparently dominant interference direction that is not present in the generating process. Large-dimensional precision-matrix eigenvalue estimation studies explicitly relate spectral estimation to the dimension-to-sample-size regime [22]. The support sweep therefore reports normalized covariance error, an eigenvalue profile, a

shrinkage response surface, and a distribution of update gains. Their purpose is to expose whether the estimator is becoming more informative or merely more aggressive. The covariance error is evaluated against the reference matrix, while practical detection uses the estimated covariance.

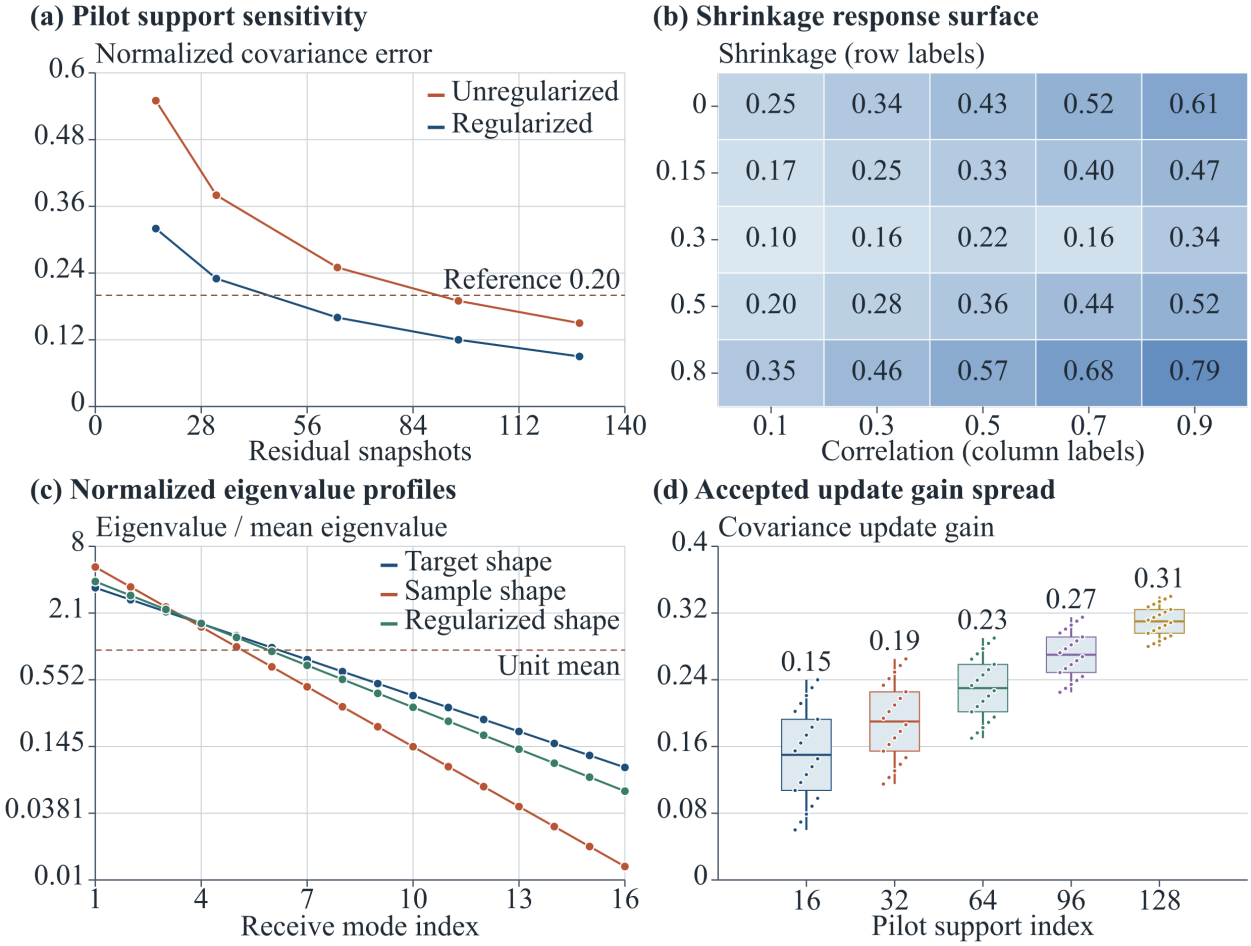


Figure 4. Covariance estimation diagnostics. (a) Covariance error against pilot support; (b) shrinkage and correlation response heatmap; (c) receive disturbance eigenvalue profiles; (d) distributions of accepted covariance update gains

Figure 4(a) reports normalized covariance errors of 0.32, 0.23, 0.16, 0.12, and 0.09 for the five pilot support values for the regularized estimator. The corresponding unregularized values are 0.55, 0.38, 0.25, 0.19, and 0.15. The difference is largest at the shortest support, reflecting the sensitivity of short support estimates to regularization. Figure 4(b) maps error values across five shrinkage levels and five correlation settings, with its lowest nominal correlation 0.7 value of 0.16 at shrinkage 0.3. The response rises at both weak and excessive shrinkage in that column. Weak shrinkage retains sample fluctuations, whereas excessive shrinkage suppresses actual directional structure. The heatmap resolves the interaction between covariance structure and regularization strength.

Figure 4(c) compares three sixteen mode eigenvalue profiles: a reference covariance, an unregularized sample estimate, and a regularized estimate. All profiles are normalized to unit average eigenvalue so that the comparison concerns shape rather than total power. The unregularized profile contains a larger leading mode and smaller trailing modes than the target profile. This pattern represents the whitening risk associated with an exaggerated eigenvalue spread: small estimated eigenvalues cause stronger amplification of the corresponding transformed directions. Anti-jamming beamforming has been designed using covariance reconstruction under limited prior knowledge [23]. Eigenvector alignment complements the spectral view by indicating the orientation of the whitening metric. Equal covariance traces isolate spectral shape in this panel. Interpretation also requires eigenvector alignment because two matrices with similar eigenvalues can impose different spatial detection metrics.

The accepted update gain distributions in Figure 4(d) use twenty values for each of the five pilot support levels. Their medians are approximately 0.15, 0.19, 0.23, 0.27, and 0.31, while the associated spreads narrow as support grows. The distributions describe a controller that permits stronger updates when its covariance information becomes more reliable. Under fast drift, a short support window may require a stronger update despite its greater sampling variance; the policy therefore requires evaluation against the disturbance change rate. Confidence-estimation research distinguishes probability calibration from the ability to separate correct predictions from failures [24]. That check includes accepted

but incorrect vectors together with the overall fraction of accepted samples. Errors concentrated in one spatial mode can compromise an otherwise favorable confidence distribution.

Transient testing begins from a stationary covariance and rotates its dominant subspace after a specified frame boundary. The pilot sequence, desired channel, and interference power are held fixed around that boundary in the first test so that the measured response can be associated with the eigenbasis change. A second test changes power without rotating the eigenbasis, and a third changes both. The proposed protocol uses independent shift locations across sequences to avoid rewarding a controller that implicitly learns a periodic reset. Actual adaptation time should be computed separately for every trial, with a recorded threshold and a rule preventing a transient threshold crossing from being counted as sustained recovery.

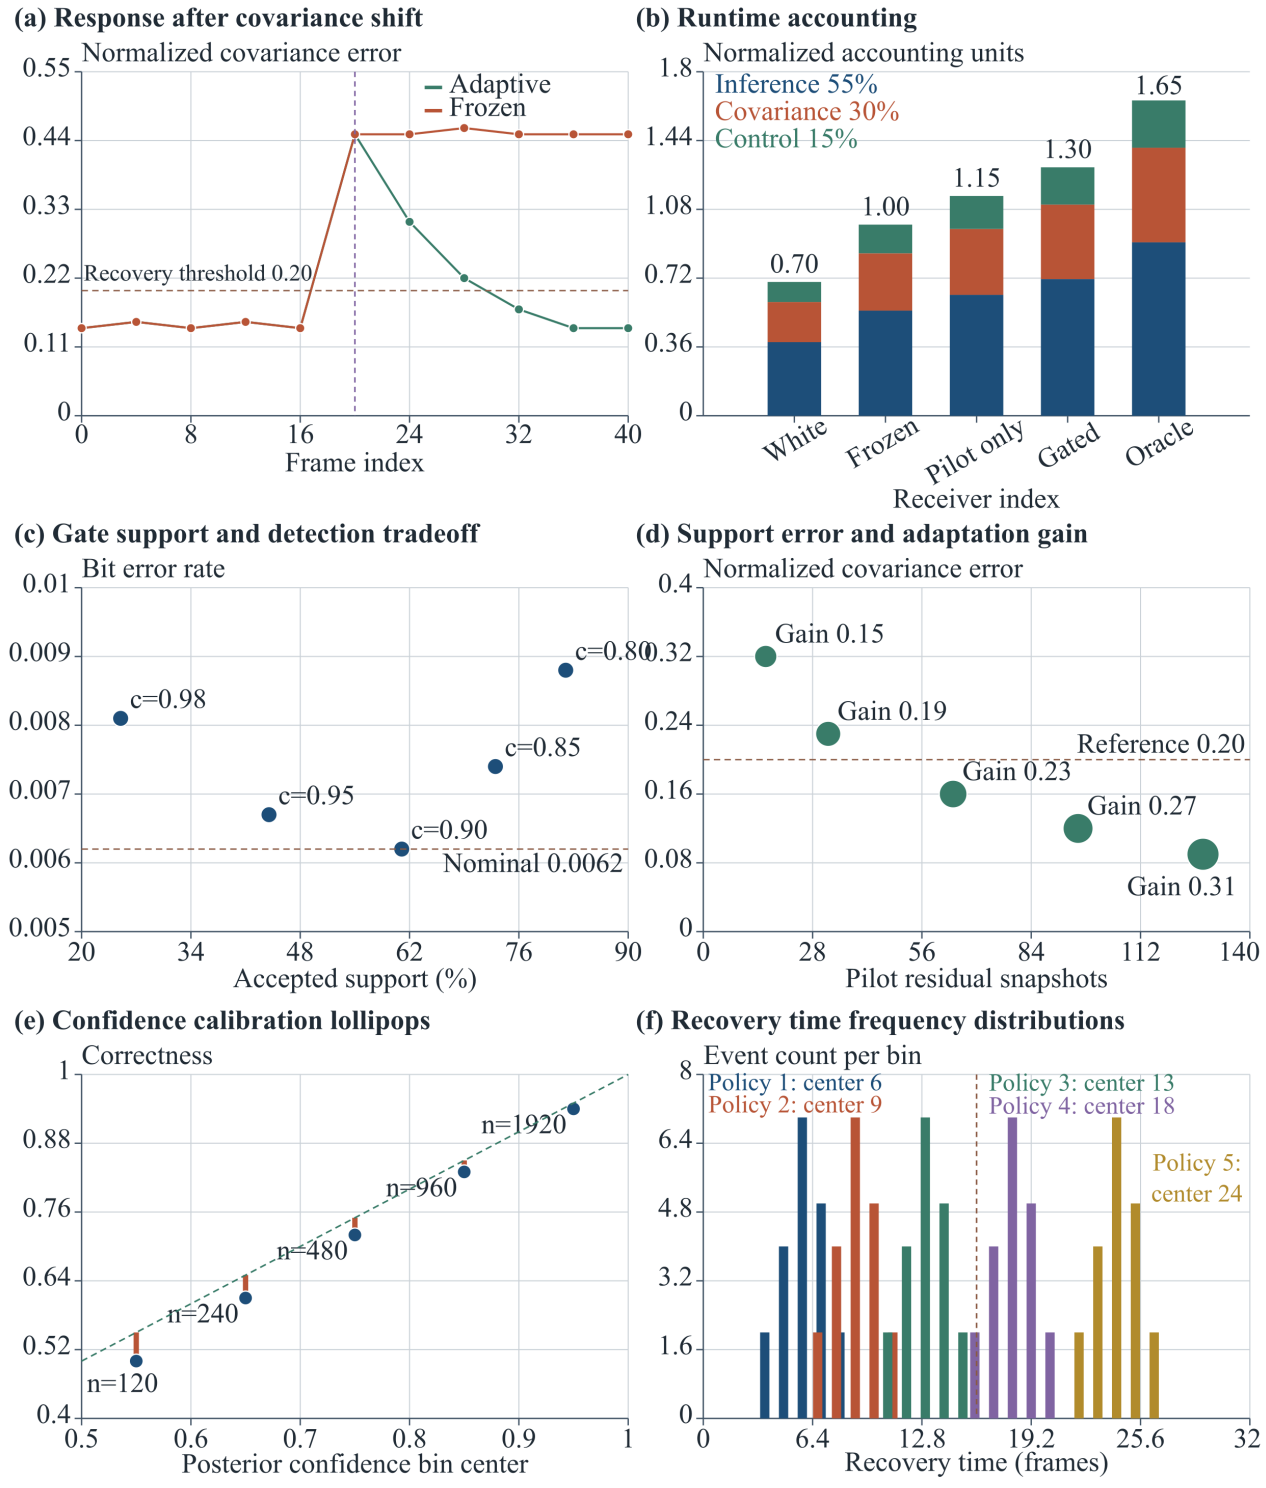


Figure 5. Adaptation and implementation diagnostics. (a) Post shift covariance recovery; (b) normalized runtime components; (c) confidence threshold tradeoff between accepted support and error rate; (d) pilot support and covariance error with adaptation gain encoded by marker size; (e) calibration reliability across confidence bins; (f) distribution of recovery times for five update policies

Figure 5(a) places an interference change at frame 20 and reports adaptive covariance error values 0.45, 0.31, 0.22, 0.17, and 0.14 at frames 20, 24, 28, 32, and 36. A frozen state remains near 0.45 over the same interval. The recovery threshold is 0.20; the displayed samples bracket its first post-change crossing between frames 28 and 32. The plotted observation times are discrete; claiming a more precise recovery time would require additional samples or a declared interpolation rule. Tracking analysis of adaptive filters evaluates behavior under nonstationary conditions alongside steady-state estimation properties [25]. The benchmark therefore retains pre change variation and late post change behavior to expose tracking speed and residual variability. A controller that recovers rapidly but remains noisy could impose a persistent cost on the detector.

The normalized runtime breakdown in Figure 5(b) reports totals of 0.70, 1.00, 1.15, 1.30, and 1.65 for white noise, frozen whitening, pilot only adaptation, gated adaptation, and oracle configurations, respectively. The oracle total includes reference matrix handling in a separate implementation; oracle cost decreases when estimation is bypassed. Figure 5(c) reports accepted supports of 82, 73, 61, 44, and 25 percent across thresholds from 0.80 to 0.98, together with error rates of 0.0088, 0.0074, 0.0062, 0.0067, and 0.0081. The interior minimum reflects the competing effects of confidence filtering and reduced update support. Increasing the threshold reduces contaminated updates while decreasing the volume of fresh information. These normalized quantities are accounting coordinates rather than measured processor latency. Oracle reference handling is included in the displayed implementation budget, so its position should be interpreted within that accounting convention.

Figure 5(d) combines the five pilot supports with their covariance errors and controller gains, allowing sampling quality and update strength to be examined jointly. Marker sizes represent gains of 0.15, 0.19, 0.23, 0.27, and 0.31; numerical axis ticks provide a direct quantitative interpretation alongside marker area. Figure 5(e) uses confidence bins centered at 0.55, 0.65, 0.75, 0.85, and 0.95 and reports correctness values of 0.50, 0.61, 0.72, 0.83, and 0.94 for the adaptive receiver. Message-passing calibration studies relate model confidence to aggregation depth and evaluate methods for correcting both overconfidence and underconfidence [26].

Figure 5(f) compares twenty recovery times for each of five update policies with nominal central values of 6, 9, 13, 18, and 24 frames. A vertical threshold at sixteen frames marks a prospective tracking requirement. The frequency distribution view reveals how recovery events lie relative to that requirement. Independent covariance change events provide the appropriate resampling unit for uncertainty estimation. Recovery distributions should also distinguish failed recovery from slow recovery; right-censored cases require explicit representation in the plot. Recovery events exceeding the observation window require explicit censoring indicators alongside the distribution of successful recovery times. Recovery is defined as a post-change covariance error at or below the stated threshold, evaluated at the displayed observation times. A failed crossing remains a censored recovery event, and temporal interpolation requires an explicit rule.

5.3. Ablation Interpretation and Robustness Boundaries

The ablation design separates the observation metric, the learned symbol estimator, and the feedback path. A white noise network establishes sensitivity to directional mismatch, a frozen whitening network isolates the effect of having a colored metric without online tracking, and a pilot only adaptive network tests whether decision feedback provides information beyond the pilot stream. An oracle covariance reference indicates how much improvement may remain available from better covariance knowledge. MIMO box-decoding research evaluates the trade-off between detection performance and implementation efficiency, including suitability for parallel processing [27]. Each ablation should therefore retain the same constellation, channel inputs, test blocks, and learned depth. When one variant uses fewer operations, the comparison should show both error versus operating condition and error versus cost, with operation budgets reported explicitly. The white-noise baseline isolates metric mismatch, frozen whitening isolates the absence of tracking, and pilot-only adaptation isolates the data-feedback contribution. These contrasts organize causal interpretation within the specified comparison budget.

Table 2 gives a compact cross check of the nominal values used in the detection and cost plots. The covariance error entries describe the metric state, while the bit error rates describe symbol detection; each diagnostic represents a distinct aspect of receiver performance. The pilot only to gated change is from 0.0084 to 0.0062, an arithmetic reduction of 26.2 percent accompanied by a cost increase from 1.15 to 1.30 units. The suitability of this trade-off depends on the receiver error target and computational budget. Assessment of real time feasibility requires timing measurements, energy measurements, and a target scheduling budget. The oracle row also has zero covariance error by definition, while thermal noise and finite constellation ambiguity still produce detection error.

Table 2. Ablation values at the nominal 18 dB operating point

Configuration	Bit error rate	Covariance error	Normalized cost
White noise network	0.0230	0.42	0.70

Frozen whitening network	0.0128	0.28	1.00
Pilot only adaptation	0.0084	0.20	1.15
Gated covariance adaptation	0.0062	0.16	1.30
Oracle whitening reference	0.0048	0.00	1.65

Figure 6 places the five nominal configurations at cost and error coordinates of (0.70, 0.0230), (1.00, 0.0128), (1.15, 0.0084), (1.30, 0.0062), and (1.65, 0.0048). A horizontal error threshold of 0.01 and a vertical cost threshold of 1.40 define a prospective acceptance region. Pilot only and gated adaptation lie inside that region in the dataset, while the oracle exceeds its cost allowance. The geometric placement illustrates a design decision under the specified error and cost allowances. Low-complexity soft-output MIMO detection research evaluates near-optimum detection performance together with implementation-oriented complexity [28]. An implementation assessment should report factorization frequency, per-tone linear-solve cost and batch scheduling effects before drawing deployment conclusions. Robustness tests should include deliberately unfavorable conditions alongside variations close to the training generator. One test uses a rank deficient interference component superimposed on thermal noise. Another uses a full rank covariance with a flatter spectrum but a rapidly rotating eigenbasis. A third introduces a channel estimation bias that persists across the pilot and data phases. The confidence gate may fail differently in each case: it may lack enough accepted vectors, accept overconfident wrong vectors, or preserve a stale but well-conditioned covariance. Recording only average bit error would obscure these distinctions. Useful logs include rejection counts, reset counts, accepted residual energy, covariance trace, condition number, and the fraction of accepted decisions later found incorrect by the evaluation procedure. The error allowance and cost allowance are design thresholds for the displayed comparison. Their intersection defines the admissible region; a different application budget changes that region without changing the underlying plotted coordinates.

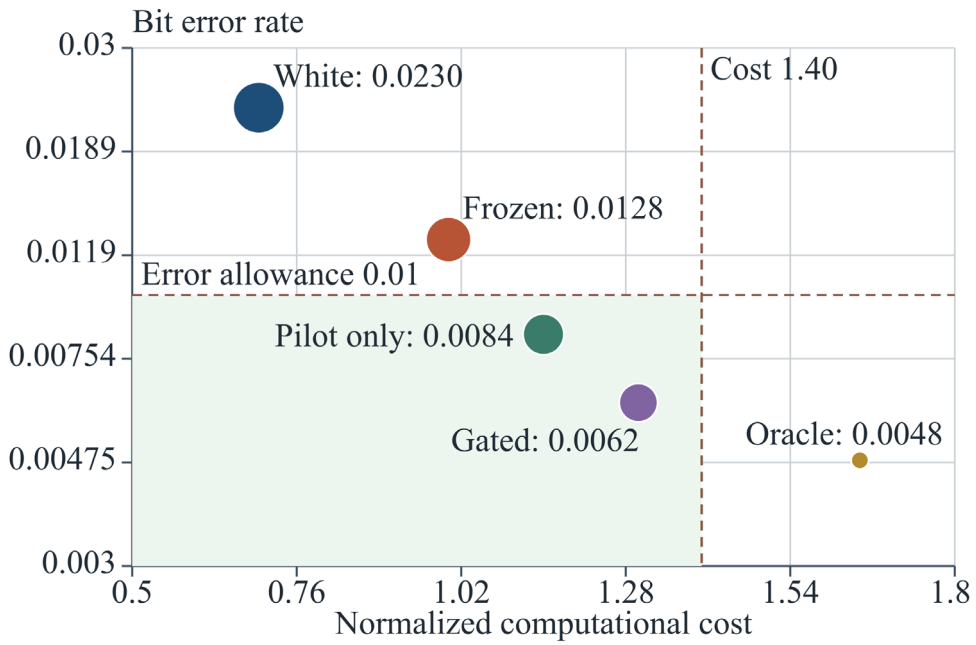


Figure 6. Detection and cost frontier with one panel

Training distribution shift requires particular care because a covariance controller can appear robust while the denoiser has learned a narrow constellation or channel prior. Evaluation should freeze all learned parameters and change the receive correlation family, interference eigenbasis, modulation order, and channel delay profile one at a time before combining them. A modulation change alters the denoiser itself, so receiver information changes even when the parameter count remains constant. Any retraining or recalibration must be labeled and applied comparably across learned baselines. The initial architecture also assumes proper complex disturbance; strongly improper interference would require an augmented representation. Frequency coupling, impulsive interference, and nonlinear hardware effects each require additional likelihood modeling beyond second order spatial whitening. A persistent channel estimation bias presents a different adaptation failure from a rotating interference subspace. The reconstruction residual then contains a structured component correlated with the desired transmission, and the covariance controller may interpret that component as external disturbance. Stronger whitening can attenuate receive directions that also carry useful signal energy, producing an apparently well-conditioned metric with degraded spatial separation. A targeted test should therefore hold external interference constant while increasing channel bias through five independently specified levels. The evaluator would compare residual power, accepted decision purity, covariance eigenvector alignment, and detection

error on the same blocks. If covariance error improves while detection deteriorates, the effective disturbance approximation requires examination at the pilot and data interface. The covariance state can then be evaluated against a pilot reset policy with identical channel inputs and update opportunities.

A second control experiment should decouple feedback purity from feedback volume. At a fixed confidence threshold, the evaluator can subsample accepted vectors to five common support levels while retaining the same pilot residuals. This control measures the contribution of sample availability without changing the detector confidence distribution. A complementary experiment keeps the accepted count fixed and varies the fraction of incorrect reconstructions admitted by the evaluator. That intervention belongs to the offline evaluation protocol. Its value is diagnostic: it measures how cancellation errors distort the estimated disturbance subspace and identifies the contamination level at which pilot only adaptation becomes preferable. Recording the angle between the contaminated covariance eigenvectors and the desired channel columns would help explain why equal covariance errors can produce different symbol error rates. The comparison should also retain the covariance trace, since an accepted batch biased toward easy vectors may underestimate disturbance power even when its estimated eigenvectors remain accurate. The prospective cost limits require a separate sensitivity test because factor reuse changes with both resource grouping and interference dynamics. A group containing more tones amortizes covariance factorization over more symbol estimates, but it also spans a wider frequency interval over which the interference covariance may vary. An implementation study should examine at least five group widths while preserving the total number of processed tones. For each width, the receiver would record factorization calls, triangular solves, state transfers, and the frequency of covariance resets. The relevant operating condition is the point at which savings from reuse are offset by metric mismatch within the group. The cost value of 1.30 for gated adaptation describes one accounting configuration; changing the grouping policy changes that coordinate and can move the receiver across the prospective cost limit. Therefore, the actual comparison needs to consider both the per-group overhead and the per-tone inference cost under the same batch processing strategy.

Interference changes during an individual frame also challenge the selected update schedule. A covariance held fixed across the unfolded layers provides a consistent metric for one observation, while a frame boundary update introduces delay for symbols received immediately after a change. The evaluator should place covariance shifts at the beginning, middle, and end of a frame and inspect error rates by symbol position. This test separates the effect of a delayed state update from failure of the denoiser under a valid metric. Shorter frames can reduce the delay but increase covariance processing frequency; longer frames increase the number of data residuals available for an update while exposing more symbols to a stale state. For each frame length, accepted sample support and shift location should be reported together with the distribution of recovery events. An update policy that satisfies the average cost allowance may still exceed a short scheduling deadline when a pilot reset and a covariance factorization occur in the same frame.

6. Conclusion

This study developed a noise covariance adaptive OAMP-Net architecture for MIMO-OFDM detection under colored interference, with the numerical comparison describing the balance between detection accuracy and adaptation reliability under the specified covariance and feedback conditions.

1. Detection performance. At a signal to noise ratio of 18 dB, the covariance adaptive configuration has a reported bit error rate of 0.0062, compared with 0.0128 for frozen whitening and 0.0230 for the white noise network. These values correspond to reductions of 51.6% and 73.0%, respectively. The comparison highlights the contribution of updating the spatial disturbance metric as interference conditions change.

2. Covariance estimation and feedback behavior. Increasing pilot support from 16 to 128 snapshots reduced normalized covariance error from 0.32 to 0.09. Feedback performance exhibited an intermediate optimum: increasing the confidence threshold from 0.90 to 0.98 reduced accepted sample support from 61% to 25%, while the bit error rate increased from 0.0062 to 0.0081. This behavior reflects the competing effects of reconstruction reliability and information availability for covariance tracking.

3. Applicability and limitations. The architecture integrates shared whitening, bounded layer parameters, and guarded covariance updates within an interpretable receiver structure. Its effectiveness depends on representative residual observations, sufficient accepted support, and covariance stability within each processing group. Channel estimation bias, rapid interference changes, and departures from the spatial second order disturbance model require additional treatment. Additional limitations concern confidence calibration, the choice of covariance grouping interval and the absence of independently measured execution cost. The numerical comparison supports interpretation of the architecture under the stated conditions, with broader validation remaining necessary.

Future work should refine confidence calibration, adapt covariance grouping to interference dynamics, and assess implementation cost under practical receiver scheduling constraints.

7. Declarations

7.1 Author Contributions

Mei-hao Ou contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Mei-yu Teng contributes to validation, analysis, investigation, data collection, draft preparation. Hsin-hao Kuai contributes to conceptualization, validation, analysis. All authors have read and agreed with the manuscript before its submission and publication.

7.2 Funding and Acknowledgments

This research received no specific financial support from any funding agency.

7.3 Institutional Review Board Statement

Not applicable.

7.4 Informed Consent Statement

Not applicable.

7.5 Date Availability

Not applicable.

7.6 Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

7.7 Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors utilized Grammarly for linguistic polishing, logical organization and draft revision. The generative AI tool was not involved in research design, data collection, data analysis, or the equation of core conclusions. All outputs generated by AI were reviewed, revised and verified manually by the authors. The authors take full responsibility for all content of this manuscript.

8. References

[1] Shao, M., Ma, W.-K., Liu, J., & Huang, Z. (2024). Accelerated and deep expectation maximization for one-bit MIMO-OFDM detection. IEEE Transactions on Signal Processing, 72, 1094–1113. https://doi.org/10.1109/tsp.2024.3359083

[2] Tomasin, S., Hasler, R., Tulino, A. M., & Sánchez-Fernández, M. (2024). Estimation of interference correlation in mmWave cellular systems. IEEE Transactions on Wireless Communications, 23(3), 1771–1782. https://doi.org/10.1109/twc.2023.3291917

[3] Chowdhury, A., Verma, G., Swami, A., & Segarra, S. (2024). Deep graph unfolding for beamforming in MU-MIMO interference networks. IEEE Transactions on Wireless Communications, 23(5), 4889–4903. https://doi.org/10.1109/twc.2023.3323207

[4] Kuang, Y., Xu, J., & Liu, L. (2025). Joint maximum likelihood channel and covariance estimation for co-channel interference rejection in OFDM systems. IEEE Transactions on Vehicular Technology, 74(10), 15718–15727. https://doi.org/10.1109/tvt.2025.3567300

[5] Takeuchi, K. (2024). Orthogonal approximate message-passing for spatially coupled linear models. IEEE Transactions on Information Theory, 70(1), 594–631. https://doi.org/10.1109/tit.2023.3311408

[6] He, L., Wang, Z., Yang, S., Liu, T., & Huang, Y. (2024). Generalizing projected gradient descent for deep-learning-aided massive MIMO detection. IEEE Transactions on Wireless Communications, 23(3), 1827–1839. https://doi.org/10.1109/twc.2023.3292124

[7] Kwon, J., Jeon, S., Jeon, Y.-S., & Poor, H. V. (2024). MIMO detection under hardware impairments: Learning with noisy labels. IEEE Transactions on Wireless Communications, 23(6), 6015–6029. https://doi.org/10.1109/twc.2023.3329521

[8] Yang, C.-H., Doss, H., & Vemuri, B. C. (2024). An empirical Bayes approach to shrinkage estimation on the manifold of symmetric positive-definite matrices. Journal of the American Statistical Association, 119(545), 259–272. https://doi.org/10.1080/01621459.2022.2110877

[9] Tan, K., Romon, G., & Bellec, P. C. (2024). Noise covariance estimation in multi-task high-dimensional linear models. Bernoulli, 30(3), 1695–1722. https://doi.org/10.3150/23-bej1644

[10] Babulghum, A., Xu, C., Ng, S. X., & El-Hajjar, M. (2024). Partial learning-based iterative detection of MIMO systems. IEEE Open Journal of Vehicular Technology, 5, 1600–1610. https://doi.org/10.1109/ojvt.2024.3482008

[11] Daha, M. Y., Khurshid, K., Ashraf, M. I., & Hadi, M. U. (2024). Optimizing signal detection in MIMO systems: AI vs approximate and linear detectors. IEEE Communications Letters, 28(10), 2387–2391. https://doi.org/10.1109/lcomm.2024.3451655

[12] Liu, L., Liang, S., & Ping, L. (2024). On capacity optimality of OAMP: Beyond IID sensing matrices and Gaussian signaling. IEEE Transactions on Communications, 72(5), 2519–2535. https://doi.org/10.1109/tcomm.2024.3354201

[13] Zhou, Y., Li, X., Zhang, X., Hui, X., & Chen, Y. (2024). Model-driven channel estimation for MIMO monostatic backscatter system with deep unfolding. IEEE Open Journal of the Communications Society, 5, 6697–6712. https://doi.org/10.1109/ojcoms.2024.3479234

[14] Dong, J., Cui, W., & Liu, Y. (2025). Covariance matrix estimation from correlated sub-Gaussian samples via the shrinkage estimator. IEEE Signal Processing Letters, 32, 841–845. https://doi.org/10.1109/lsp.2025.3541427

[15] Wang, P., Fang, J., Ma, B., Wang, B., & Leus, G. (2025). Low-rank covariance matrix recovery from rank-one measurements: An analytical solution. IEEE Signal Processing Letters, 32, 2674–2678. https://doi.org/10.1109/lsp.2025.3584252

[16] Cheng, L., Liu, X., Kong, L., Xiong, J., Zhao, H., & Wei, J. (2026). Label-free environment adaptation for wireless receivers via physics-informed constraints. IEEE Transactions on Wireless Communications, 25, 23250–23264. https://doi.org/10.1109/twc.2026.3730873

[17] Uzlaner, N., Raviv, T., Shlezinger, N., & Todros, K. (2025). Asynchronous online adaptation via modular drift detection for deep receivers. IEEE Transactions on Wireless Communications, 24(5), 4454–4468. https://doi.org/10.1109/twc.2025.3533959

[18] Shallal Abbas Anooz, R., Pourrostam, J., & Al-Ibadi, M. (2024). Adaptive filters versus machine learning based beam tracking techniques for millimeter-wave wireless communications systems. IEEE Access, 12, 165878–165888. https://doi.org/10.1109/access.2024.3491870

[19] Chen, S., Qi, X., Peng, M., & Yan, S. (2025). Anti-main-lobe jamming beamforming with covariance matrix reconstruction for wireless communication systems under jamming attacks. IEEE Transactions on Vehicular Technology, 74(4), 6168–6179. https://doi.org/10.1109/tvt.2024.3515972

[20] Lin, Y., Song, M., & van der Sluis, B. (2025). Bootstrap inference for linear time-varying coefficient models in locally stationary time series. Journal of Computational and Graphical Statistics, 34(2), 654–667. https://doi.org/10.1080/10618600.2024.2403705

[21] Raina, R., Simmons, N., Simmons, D. E., Yacoub, M. D., & Duong, T. Q. (2026). To trust or not to trust: On calibration in ML-based resource allocation for wireless networks. IEEE Transactions on Network Science and Engineering, 13, 5961–5977. https://doi.org/10.1109/tnse.2025.3636073

[22] Zhou, J., Xie, J., & Chen, J. (2026). Direct estimation of eigenvalues of large dimensional precision matrix. IEEE Transactions on Signal Processing, 74, 183–199. https://doi.org/10.1109/tsp.2025.3647215

[23] Pan, Z., Zhang, B., Wang, H., Ma, W., & Guo, D. (2024). Beamforming design for anti-jamming IoT communication based on covariance matrix reconstruction with little prior knowledge. IEEE Transactions on Consumer Electronics, 70(4), 7494–7504. https://doi.org/10.1109/tce.2024.3420450

[24] Zhu, F., Zhang, X.-Y., Cheng, Z., & Liu, C.-L. (2024). Revisiting confidence estimation: Towards reliable failure prediction. IEEE Transactions on Pattern Analysis and Machine Intelligence, 46(5), 3370–3387. https://doi.org/10.1109/tpami.2023.3342285

[25] Khalili, A., Rastegarnia, A., Farzamnia, A., Sanei, S., & Alghamdi, T. A. H. (2024). Tracking analysis of maximum Versoria criterion based adaptive filter. IEEE Access, 12, 30747–30753. https://doi.org/10.1109/access.2024.3370471

[26] Wang, M., Yang, H., Huang, J., & Cheng, Q. (2024). Moderate message passing improves calibration: A universal way to mitigate confidence bias in graph neural networks. Proceedings of the AAAI Conference on Artificial Intelligence, 38(19), 21681–21689. https://doi.org/10.1609/aaai.v38i19.30167

[27] Qureshi, S. F., Damjancevic, S., Matus, E., van der Wolf, P., Utyansky, D., & Fettweis, G. (2025). Low-complexity MIMO detection in 5G IoT: The box decoding approach. Journal of Signal Processing Systems, 97(2-4), 173–184. https://doi.org/10.1007/s11265-025-01963-w

[28] Hu, J., Song, S., & Wang, Z. (2024). A low-complexity soft-output massive MIMO detector with near-optimum performance. IEEE Transactions on Circuits and Systems I: Regular Papers, 71(12), 5445–5456. https://doi.org/10.1109/tcsi.2024.3435361