Frugal Physics-Informed AI for Adaptive Radar Detection in Challenging Electromagnetic Environments

Tom Portafaix[1,2], Jean-Marc Le Caillec[2], Michel Narozny[1], Jean-François Degurse[3]

[1] Thales LAS, France [2] IMT Atlantique, France [3] Thales DMS, France

Identification of Use Cases

Operational context and problem. Airborne radar platforms increasingly operate in challenging electromagnetic environments, where small, slow or low-RCS targets must still be detected while the receiver is exposed to spatially coherent interference and to platform-induced ground or sea clutter spread over Doppler. The reference is *adaptive* processing: Capon/MVDR beamforming against spatially coherent interference and STAP against Doppler-correlated clutter. Both whiten the snapshot with the inverse of the interference-plus-noise covariance $\mathbf{R}$, through an adaptive weight $\mathbf{w} \propto \mathbf{R}^{-1}\mathbf{s}$ ($\mathbf{s}$: steering vector).

That covariance is unknown, and estimating it imposes two costs. The first is *sample support*: the Reed–Mallett–Brennan rule requires $K \gtrsim 2N$ homogeneous secondary cells to stay within 3 dB of the optimum, which heterogeneous scenes rarely provide. The second is *computation*: inverting the covariance costs $\mathcal{O}(N^3)$ per look direction, with $N = AL$ growing with the antenna channels A and Doppler taps L. On constrained embedded platforms the adaptive chain is therefore degraded or dropped, and the sensor falls back on a non-adaptive sum channel whose detection collapses in scenes where the optimal processor still detects.

Capability, users and benefits. The capability to be improved is to retain near-optimal interference and clutter rejection within the memory, energy and latency budget of an embedded processor. Target users are airborne radar operators and surveillance programmes deploying adaptive processing on constrained hardware, where detection has to be preserved in degraded and congested electromagnetic conditions.

Expected benefits are a preserved detection range under interference, a constant inference cost with no matrix inversion nor secondary data, and physically interpretable internal quantities, a prerequisite for qualification. The module is a drop-in replacement for the adaptive stage of an existing chain, retrainable from a simulator.

Proposed Technical Approach

Technical lock. The lock is hence maintaining adaptive detection when onboard resources make covariance estimation and inversion impractical. Our first attempt was to remove that chain altogether, by training black-box U-Net detectors on range-Doppler (RD) images, including variants with the antenna channels concatenated and a self-attention module across them. All detect well in thermal noise and *collapse* as soon as spatially coherent interference is present: given only pixels and a mask, the network behaves as a generic segmenter and never reconstructs the whitening operators, and such a detector offers nothing to inspect. The design below (Figure 1) answers both: it is inversion-free at inference and interpretable by construction.

Novelty. Rather than learning only a detector output, we explicitly supervise the adaptive quantities $\widehat{\mathbf{R}}_\phi$ and $\widehat{\mathbf{w}}_\phi$, yielding an inversion-free and physically inspectable representation.

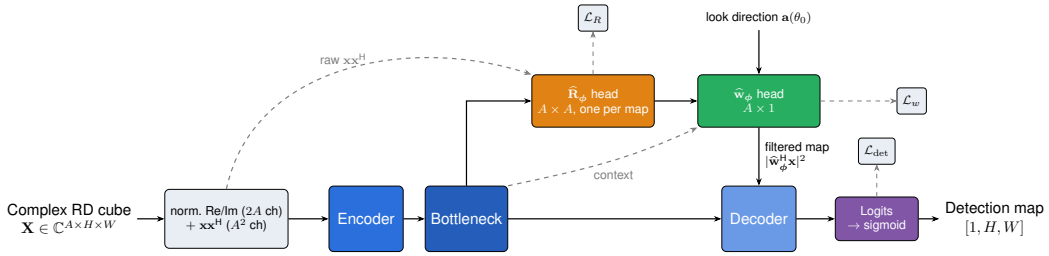

Figure 1: Physics-informed detector: the covariance and weight bottlenecks follow the classical chain, are supervised against analytically computed oracles, and condition the decoder through the filtered map. No matrix inversion is performed.

Proposed design: unroll the adaptive chain and supervise it. Instead of hiding the physics, we expose it. The network is fed the complex RD cube $\mathbf{X} \in \mathbb{C}^{A\times H\times W}$ acquired on $A{=}4$ antenna channels and, at every cell, is given the per-pixel outer product $\mathbf{x}\mathbf{x}^{\mathsf{H}}$ (16 real maps) alongside the normalised real and imaginary parts (8 maps): the exact second-order statistic on which Capon operates, rather than leaving the network to synthesise it. A convolutional encoder then feeds two *supervised, physically interpretable bottlenecks* that follow the classical chain. The covariance head pools the bottleneck features and the raw $\mathbf{x}\mathbf{x}^{\mathsf{H}}$ maps over the whole map into a single $\widehat{\mathbf{R}}_{\boldsymbol{\phi}} \in \mathbb{C}^{A\times A}$, written $\mathbf{L}_{\boldsymbol{\phi}}\mathbf{L}_{\boldsymbol{\phi}}^{\mathsf{H}}$ with $\mathbf{L}_{\boldsymbol{\phi}}$ lower triangular, hence Hermitian positive semi-definite by construction. The weight head maps $\widehat{\mathbf{R}}_{\boldsymbol{\phi}}$, the bottleneck context and the look direction $\mathbf{a}(\theta_0)$ to $\widehat{\mathbf{w}}_{\boldsymbol{\phi}} \in \mathbb{C}^{A}$, under the distortionless constraint $\mathbf{a}(\theta_0)^{\mathsf{H}}\widehat{\mathbf{w}}_{\boldsymbol{\phi}} = 1$. The decoder receives the filtered map $|\widehat{\mathbf{w}}_{\boldsymbol{\phi}}^{\mathsf{H}}\mathbf{x}|^2$, conditioning detection on the learned whitening.

The simulator provides the oracle covariance $\mathbf{R}$ and optimal weight $\mathbf{w}$ for direct physical supervision:

$$\mathcal{L} = \mathcal{L}_{\text{det}} + \lambda_R \left\| \frac{\widehat{\mathbf{R}}_{\boldsymbol{\phi}}}{\|\widehat{\mathbf{R}}_{\boldsymbol{\phi}}\|_F} - \frac{\mathbf{R}}{\|\mathbf{R}\|_F} \right\|_F^2 + \lambda_w \left\| \widehat{\mathbf{w}}_{\boldsymbol{\phi}} - \mathbf{w} \right\|_2^2 .$$

Covariance normalisation focuses learning on interference *geometry* rather than power. At inference, a single forward pass needs no covariance formation, inversion or secondary cells. With about $0.5\,$M parameters, the network is also one to two orders of magnitude smaller than standard vision backbones. It remains *openable*: $\widehat{\mathbf{R}}_{\boldsymbol{\phi}}$ can be fed to an analytical Capon or STAP whitener to quantify the information captured by the learned representation, supporting trusted-AI assessment.

Validation. A hierarchical benchmark built on a four-antenna array isolates each correlation: (1) thermal noise only, (2) spatially coherent interference, (3) Doppler-correlated clutter requiring a space-time formulation. Regimes 1–2 are reported here; regime 3 is under way. Six detectors are compared on identical scenes: Capon oracle (exact covariance), empirical Capon (sample-matrix estimate), sum channel, our end-to-end network, Capon on the learned $\widehat{\mathbf{R}}_{\boldsymbol{\phi}}$, and $\widehat{\mathbf{w}}_{\boldsymbol{\phi}}$ alone; thresholds are CFAR-calibrated per scene, with 200 Monte-Carlo draws per point.

In regime 1 (Figure 2a) the sum channel, optimal in white noise, matches single-pulse Swerling-0 theory ($P_d{=}0.95$ at $13.7\,$dB SNR), which validates the Monte-Carlo evaluation pipeline. The network and Capon on $\widehat{\mathbf{R}}_{\boldsymbol{\phi}}$ coincide with this optimum, and the network learns $\widehat{\mathbf{R}}_{\boldsymbol{\phi}} \approx \mathbf{I}$ to within 5×10^{-5}, which validates the architecture.

Regime 2 (Figure 2b) is the decisive one: the sum channel never detects, while the end-to-end network stays within 2.5–$3.5\,$dB of empirical Capon without any matrix inversion. Capon on the *learned* covariance and the learned weight alone still detect and join the full network at high P_d: the whitening information is carried by the interpretable bottlenecks. The remaining gap to the oracle comes from $\widehat{\mathbf{R}}_{\boldsymbol{\phi}}$, which captures the interference direction but spreads its power over several eigenmodes, limiting the depth of the null.

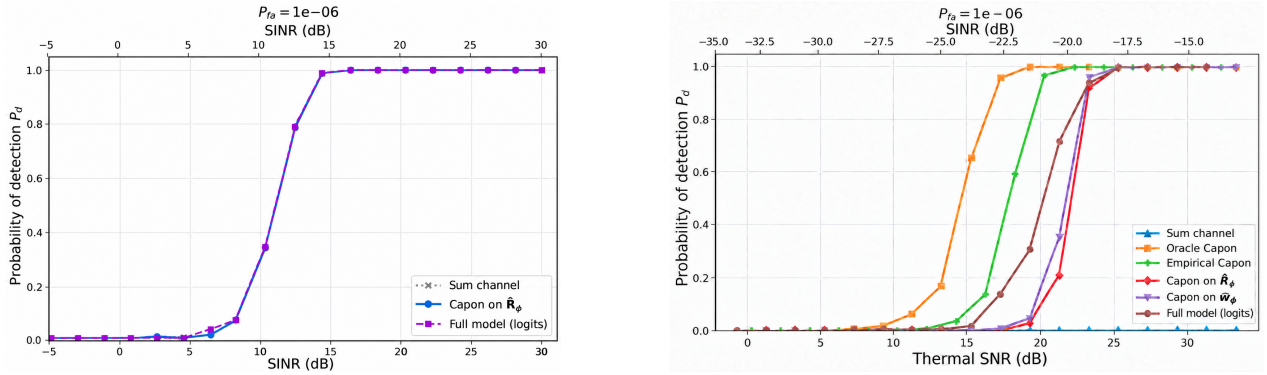

(a) Regime 1: thermal noise only **(b)** Regime 2: spatially coherent interference

Figure 2: Detection probability versus SNR at $P_{\text{fa}} = 10^{-6}$ (in regime 1, Capon oracle = sum channel).

Maturity level. Proof of concept validated in simulation (regimes 1–2; regime 3 in progress), not yet exercised on recorded data or target hardware.

Deployment perspectives. Embedded latency and memory benchmarking against sample-matrix Capon/STAP under limited data; robustness to close or multiple interference sources, weaker targets and recorded data; stronger physical priors on $\widehat{\mathbf{R}}_{\boldsymbol{\phi}}$.