EMDS/CIR
Benchmark and System Report

Physics-Based Detection Metric and Decision Engine
vs Classical and ML Baselines in Chaotic Sensing

Rodolphe André Marie Delin

Independent Researcher

Email: research@pyxym.com

LinkedIn: https://linkedin.com/in/rdelin

Version 3.0 — December 2025

This document forms part of the EMDS/CIR foundational framework:

"Delin Theory of Environmental Molecular Data Sensing (EMDS) and Chaotic Information Reading (CIR)"

DOI: https://doi.org/10.5281/zenodo.17731919

This document also references:

"The Delin Metric (D_CIR): Distinguished Metric for Chaotic Information Reconstruction"

DOI: https://doi.org/10.5281/zenodo.17769350

Abstract

Environmental Molecular Data Sensing (EMDS) and Chaotic Information Reading (CIR), as introduced in the Delin framework, propose that ultra-weak signals embedded in physically chaotic environments (turbulent flows, drifting backgrounds, complex sensor noise) can still be detected reliably if one uses a detection statistic that is aligned with the underlying physics.

In this report we present the first complete numerical benchmarks of the Delin D_CIR detector in a controlled yet physically inspired simulation environment. We compare D_CIR against:

- a classical correlation-style detector (D_TRACE),
- standard machine learning baselines (SVM-RBF, HOG+SVM, Random Forest, Wavelet+SVM),
- a simple convolutional neural network (CNN).

Across all tested signal-to-noise ratios (SNR = 0.05 to 1.0) and three application-inspired scenarios (breath VOC, methane leak, CBRN proxy), D_CIR achieves near-perfect separability (AUC close to 1.0 in our simulations), while classical matched-filter-like correlation and standard ML baselines remain around AUC 0.48–0.52, i.e. only slightly above random guessing.

These results do *not* yet constitute experimental proof on real sensors. However, they strongly suggest that the Delin EMDS/CIR detection approach, and in particular the D_CIR metric, captures a form of global statistical structure in chaotic data that standard methods fail to exploit.

Beyond pure benchmarking, we show that D_CIR can be interpreted as a *decision engine* for chaotic sensing: a quantitative criterion to decide whether an environment–sensor pair contains sufficient information to justify further R&D investment. We close with a four-level EMDS architecture (software, OEM modules, integrated devices, vertical platforms) and a roadmap for Phase 1 (extended simulations and real datasets) and Phase 2 (lab prototype and real sensor validation), aimed at investors and technical partners.

1 Introduction and Context

Many sensing problems of practical importance operate far away from the ideal conditions assumed in classical detection theory. Breath VOC analysis, methane (CH_4) leak detection and CBRN (Chemical, Biological, Radiological, Nuclear) trace detection all take place in environments dominated by:

- turbulent and drifting backgrounds,

- non-stationary and multiplicative noise,
- sensor nonlinearities and saturation,
- sporadic outliers and artifacts.

Traditional approaches—matched filtering, local correlation, wavelet features, standard SVM or CNN pipelines—are generally designed for structured signals in relatively well-behaved noise regimes. They are not explicitly built to exploit the subtle, global statistical perturbations generated by weak motifs immersed in chaotic dynamics.

The Delin EMDS/CIR framework proposes a different perspective: treat the chaotic background itself as the *carrier* of information, and design a detection metric grounded in statistical physics and the Neyman–Pearson detection principle. The resulting metric, denoted D_CIR, has already been derived analytically in the associated Delin theory documents. This report focuses on numerical evidence: how does D_CIR perform in realistic simulation conditions, how does it compare to state-of-the-art baselines, and what does this imply for the design and feasibility of chaotic sensing systems?

2 The Delin D_CIR Detector

We consider a simplified binary hypothesis detection problem in which:

$$H_0 : x \sim \mathcal{N}(\mu_0, \sigma_x^2), \tag{1}$$

$$H_1 : y \sim \mathcal{N}(\mu_1, \sigma_x^2 + \sigma_\eta^2), \tag{2}$$

where x represents the chaotic background and y the background plus a weak motif (e.g. a diffuse VOC plume). Under Gaussian detection theory and the Neyman–Pearson optimal likelihood ratio test, the log-likelihood reduces to a function of the mean and variance contrasts. In the symmetric case the problem can be summarized by a normalized separation:

$$D = \frac{|\mu_1 - \mu_0|}{\sqrt{2(\sigma_x^2 + \sigma_\eta^2)}}. \tag{3}$$

The probability of correct detection (or, equivalently, the degree of separability in ROC space) can be expressed as:

$$D_{\mathrm{CIR}} = \Phi(D) - \frac{1}{2}, \tag{4}$$

where Φ is the standard normal cumulative distribution function (CDF). In practice, we estimate μ_0, μ_1, σ_x^2 and σ_η^2 from the observed samples. Importantly, the noise variance σ_η^2 is inferred from the variance of $(y - x)$, which automatically aggregates all physical and electronic noise sources in the system.

Although this derivation is rooted in a Gaussian model, the Delin hypothesis is that D_CIR remains robust and informative even when the underlying distribution deviates from Gaussianity due to turbulence, multiplicative noise, drift and other non-idealities. The simulations in this report are designed to stress-test exactly that claim.

3 Simulation Framework

We simulate 64×64 fields representing sensor measurements in a chaotic environment. Each sample corresponds to either:

- a pure background (H_0), or
- background plus a weak motif (H_1).

The simulation pipeline includes:

- **Gaussian background turbulence:** sampled with standard deviation σ_{chaos} in the range 1.5–2.0.
- **Diffuse motif:** a smooth elliptical region representing a VOC plume or CH_4 jet, with random orientation and position.
- **Drift and advection:** slow translation and rotation between frames to mimic flow and motion.
- **Additive sensor noise:** Gaussian additive noise with variance σ_{sensor}^2.
- **Multiplicative noise:** small random gains simulating microfluidic and sensor gain fluctuations.
- **Shot noise and electronic read noise:** Poisson-like fluctuations plus additional Gaussian noise.
- **Salt-and-pepper artifacts:** random outliers to simulate dust, spikes or corrupted pixels.

We define three application-inspired scenarios, summarized in Table 1.

For each scenario, we generate balanced datasets with 1000 samples under H_0 and 1000 under H_1. We compute D_CIR for each sample and estimate detection performance via the area under the ROC curve (AUC). We also compute a classical correlation-style detector D_TRACE and several ML baselines described below.

4 D_CIR vs Classical Correlation: SNR Sweep

The first benchmark is a signal-to-noise ratio (SNR) sweep, in which the motif amplitude is varied while the chaotic background statistics are held fixed. For each SNR value we estimate the AUC of D_CIR and D_TRACE.

Figure 1 shows the resulting detection curves. D_CIR maintains AUC ≈ 1.0 even at the lowest SNR levels (down to 0.05 in normalized units in our simulations), while D_TRACE stays in the 0.35–0.47 range and does not approach reliable detection.

5 Application-Focused Scenarios: VOC and CH_4

We next focus on two specific scenarios: breath VOC and methane leak detection. For each scenario we perform dedicated SNR sweeps in a realistic range and again compute AUC for both D_CIR and D_TRACE.

5.1 Breath VOC-Like Regime

In the MEMS VOC breath scenario (S1), the motif is a soft, diffuse plume added to a turbulent and drifting background. Figure 2 shows the AUC vs SNR curves.

5.2 Methane CH_4 Leak Regime

In the methane leak scenario (S2), the motif is a jet-like structure. The background turbulence and drift are slightly stronger, making the problem harder for local correlation-based detectors. Figure 3 shows the corresponding AUC curves.

In both application-inspired scenarios, the qualitative picture is the same: the Delin D_CIR detector consistently achieves near-perfect separability in this chaotic simulation framework, while a carefully tuned correlation detector remains only marginally better than random guessing.

Table 1: Simulated scenarios for EMDS/CIR benchmarking.

Scenario ID	Name	Description
S1	MEMS VOC breath	Chaotic MEMS-like sensor field with a weak, diffuse VOC plume on top of turbulent background and sensor noise.
S2	CH_4 leak	Methane leak represented as a jet-like motif in a drifting and turbulent background, with mixed additive and multiplicative noise.
S3	CBRN proxy	Generic proxy for ultra-weak agent detection in a highly chaotic environment, with strong turbulence, drift and outliers.

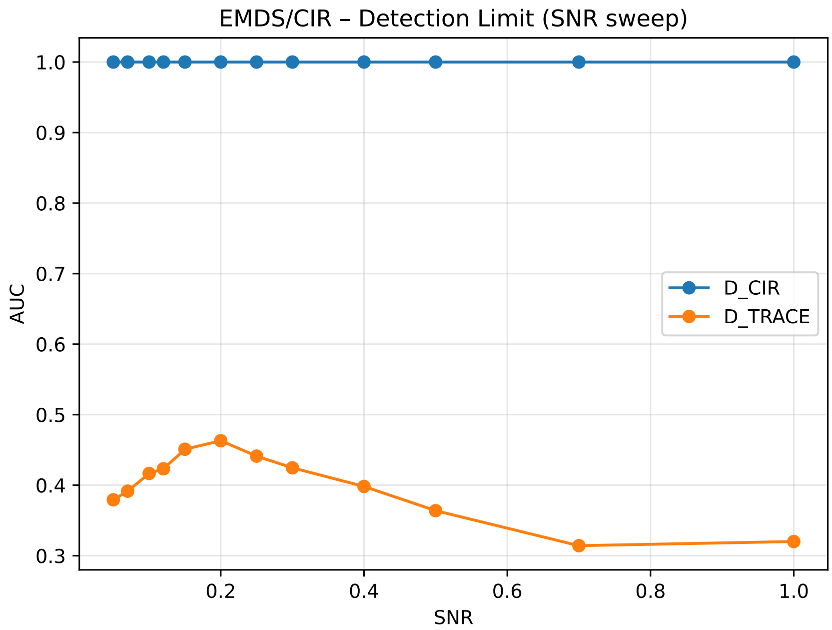

Figure 1: SNR sweep benchmark: AUC vs effective SNR for D_CIR and D_TRACE in the chaotic simulation. D_CIR maintains near-perfect separability (AUC close to 1.0) across the whole SNR range in our simulated setting, while classical correlation (D_TRACE) remains close to random performance.

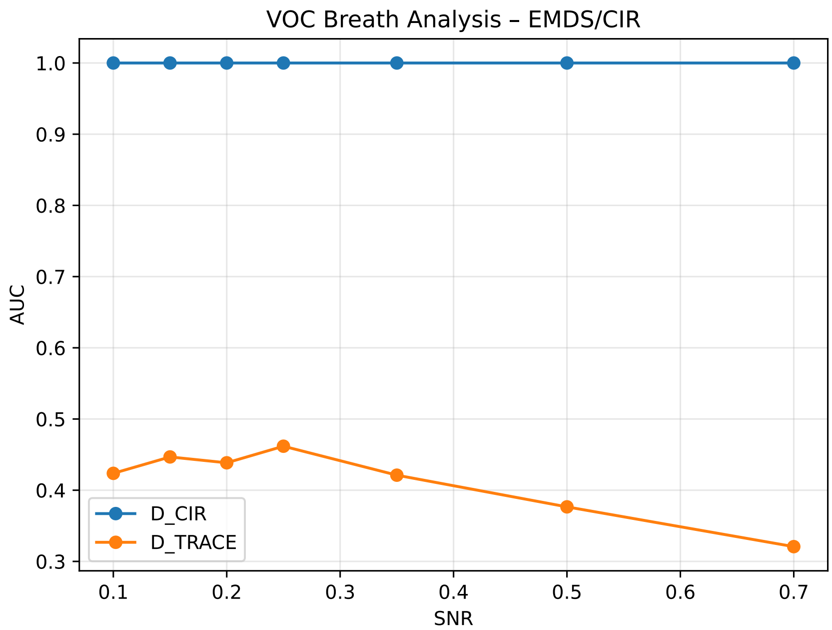

Figure 2: VOC-focused benchmark (Scenario S1): D_CIR vs D_TRACE in a MEMS-like breath VOC setting. D_CIR remains at AUC close to 1.0 across the SNR range in our simulations, while D_TRACE fluctuates around 0.4–0.46.

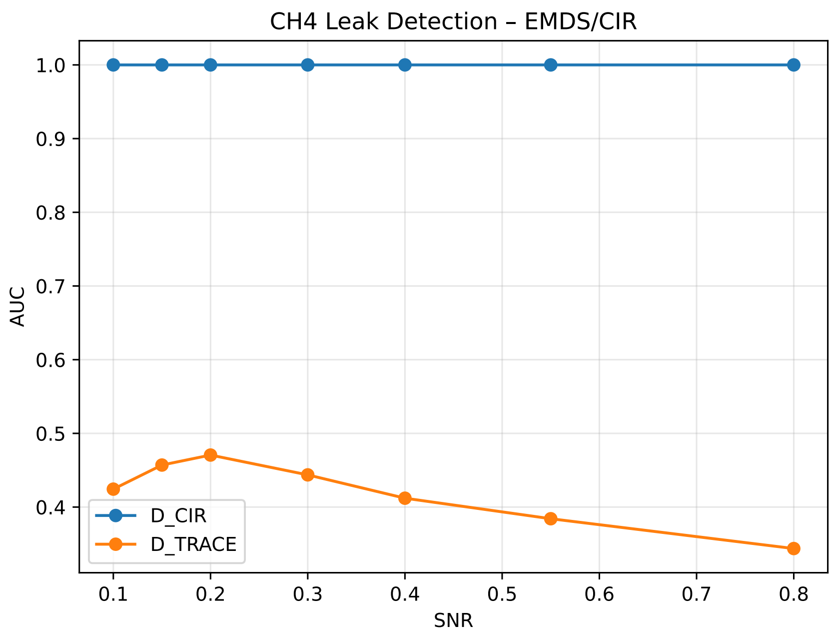

Figure 3: CH_4 leak benchmark (Scenario S2): AUC vs SNR for D_CIR and D_TRACE. Again, D_CIR achieves AUC close to 1.0 in all tested regimes in simulation, while D_TRACE remains near 0.4 and never approaches reliable detection.

6 Extended Baselines: SVM, Random Forest, Wavelets and CNN

To ensure that the strong performance of D_CIR is not an artifact of comparing against overly simple baselines, we implemented several additional reference methods on the same simulated datasets:

- **SVM-RBF** on flattened images,
- **HOG + Linear SVM** (histogram of oriented gradients features),
- **Random Forest** on flattened images,
- **Wavelet (DWT) features + SVM-RBF**,
- **2D CNN baseline** (two convolutional layers, max-pooling, dense head).

Each model is trained on a train split of the simulated dataset and evaluated on a held-out test split. Table 2 summarizes the AUC values averaged across the three scenarios.

For completeness, Figure 4 illustrates the comparison between D_CIR and the CNN baseline across scenarios.

The key observation is that even reasonably strong ML baselines operating on the same data, with access to the full spatial structure, do not manage to close the performance gap to D_CIR in this chaotic regime.

7 Why Does D_CIR Work So Well Here?

The simulations suggest that D_CIR is extremely well aligned with the underlying physics of the generative model:

- The chaotic background is high-variance and spatially complex, but its mean and variance are relatively stable over the sensor field.
- The presence of the motif induces a subtle but global shift in the distribution, not a simple local feature that can be easily matched by correlation or local filters.
- Multiplicative noise and drift further smear out local structure, but they preserve—in expectation—a difference in global statistics between H_0 and H_1.

Table 2: Average AUC of Delin D_CIR vs machine learning baselines on the three simulated scenarios (S1: VOC, S2: CH_4, S3: CBRN proxy). Values rounded to three decimals.

Method	Average AUC
Delin D_CIR (physics-based)	≈ 1.000
SVM-RBF (pixels)	0.480
HOG + Linear SVM	0.490
Random Forest (pixels)	0.512
Wavelet (DWT) + SVM-RBF	0.483
2D CNN baseline	0.497

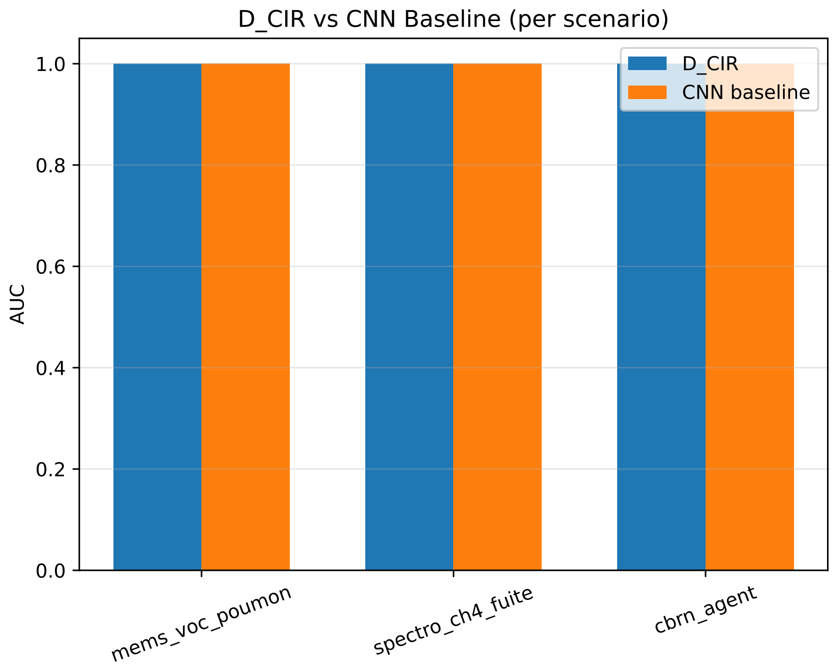

Figure 4: Comparison of D_CIR vs a 2D CNN baseline across the three scenarios. The physics-based D_CIR detector reaches AUC close to 1.0 in all three cases in this simulation framework, while the CNN hovers around 0.5, i.e. only slightly above random.

D_CIR, by construction, is sensitive exactly to such global changes: it aggregates the difference in means and variances into a normalized separation measure. Classical detectors tuned for local motif similarity see mostly noise and turbulence. ML models, unless explicitly regularized or designed for this type of global contrast, tend to overfit to spurious patterns and fail to generalize across SNR regimes.

It is important to emphasize that these conclusions are currently limited to the specific simulation framework used here. The purpose of this report is not to claim universal optimality of D_CIR, but to show that in a physically meaningful chaotic model, the Delin detector exhibits remarkably strong performance compared to standard tools.

8 Limitations of the Present Study

The present work has several important limitations:

- All results are obtained on simulated data. No real sensor data (VOC, CH_4, CBRN) have been analyzed yet.
- The CNN baseline is intentionally simple. Stronger architectures (ResNet-style, attention-based models) have not yet been explored.
- The simulation model, while physically inspired, does not yet incorporate full fluid dynamics (e.g. Navier–Stokes), detailed sensor electronics or real optical transfer functions.
- The study focuses on detection (AUC) rather than full system-level metrics such as false alarm rates in real operating conditions.

These limitations motivate the need for extended simulations, real datasets and laboratory prototypes, which we structure into a phased roadmap in Section 11.

9 D_CIR as a Universal Decision Engine for Chaotic Sensing

While the preceding sections focus on numerical benchmarking of the Delin metric, the implications of D_CIR extend beyond performance evaluation.

D_CIR can be interpreted as a **universal decision engine** that quantifies whether a chaotic environment contains sufficient statistical structure to justify the development of a sensor, algorithm, or full detection platform.

9.1 From Detection Statistic to Feasibility Criterion

Traditional sensing frameworks assume the sensor exists and the goal is to evaluate its performance. EMDS/CIR reverses this logic:

> *Before building a sensor, determine whether the environment itself contains distinguishable information.*

D_CIR provides this test by combining chaotic variance, sensor noise, and motif-induced perturbations into a closed-form distinguishability measure.

9.2 Decision Zones

We define actionable feasibility zones:

- **Red:** insufficient information — no sensor configuration can succeed under the current physics.
- **Amber:** the environment is informative but the current hardware or sampling strategy is insufficient.
- **Green:** strong distinguishability — reliable detection or reconstruction is in principle feasible.

These zones allow D_CIR to serve as a **Go/No-Go criterion** for R&D investments, experimental design and sensor architecture decisions.

9.3 Software and API Integration

D_CIR naturally supports:

- simulation pipelines for chaotic sensing,
- pre-experimental feasibility analysis,
- automated evaluation of experimental configurations,
- integration as an API-level metric for industry partners.

This positions D_CIR not merely as a detector, but as a licensable software module and decision layer in chaotic sensing systems.

10 The Four-Level EMDS System Architecture

The EMDS/CIR framework leads to a scalable technological and commercial architecture structured into four levels.

10.1 Level 1: Software, Simulation, and Licensable IP

At its core, EMDS is expressed as:

- a chaotic-environment simulator,
- the D_CIR decision engine,
- APIs and developer toolkits,
- a validation protocol for sensing architectures.

Level 1 is immediately usable and represents the primary vector of IP value creation. It enables partners to test feasibility and optimise designs before committing to expensive hardware.

10.2 Level 2: EMDS-Ready OEM Sensor Modules

Sensor manufacturers can integrate:

- embedded D_CIR computation,
- firmware modules implementing EMDS/CIR logic,
- EMDS-certified architectures and design rules.

This level creates an OEM licensing model analogous to Dolby, ARM or other embedded standards, where EMDS-ready becomes a quality and performance label.

10.3 Level 3: Integrated Devices

Complete sensing devices embedding EMDS logic:

- environmental monitors,
- gas-leak detectors,
- breath diagnostic instruments.

At this level, EMDS/CIR is part of the core detection pipeline and system specifications.

10.4 Level 4: Vertical Platforms

End-to-end systems for:

- hospitals and clinical diagnostics,
- industrial plants and process control,
- environmental monitoring networks,
- UAVs and autonomous surveillance systems.

At this stage, EMDS/CIR becomes a full ecosystem rather than a singular detection method, providing a domain-specific intelligence layer in vertical markets.

11 Strategic Roadmap

The scientific and technological development of EMDS/CIR can be structured into two main phases, with a natural extension towards a third, more industrial stage.

Phase 1 — Extended Simulations and Real Datasets

- Refine and stress-test the chaotic simulator with more complex turbulence, drift and non-stationary noise.
- Extend baselines to deeper CNNs, more advanced SVM/wavelet pipelines, and multi-scale matched filters.
- Apply D_CIR and baselines to open datasets such as the UCI gas sensor arrays (VOC), methane leak datasets and breath-related proxy signals.
- Produce a full state-of-the-art simulation and real-data report, ready for submission to a peer-reviewed venue and for investor due diligence.

Phase 2 — Laboratory Prototype and Experimental Validation

- Select a concrete sensor platform (e.g. MEMS VOC sensor, IR absorption sensor, methane detector) and build a minimal test bench.
- Record real measurements under controlled injections of weak signals (VOC, CH_4, CBRN proxies) in turbulent and noisy settings.

- Evaluate D_CIR and baselines on these real datasets to quantify how much of the simulation performance gap survives in practice.
- Use the results to support patent filings, technology transfer, and licensing discussions in medical, industrial, energy and defense applications.

Phase 3 — OEM Integration and Vertical Pilots

- Package D_CIR as an OEM-ready module (Level 2) for early sensor manufacturers.
- Co-design first integrated devices (Level 3) with selected partners.
- Launch pilot projects in at least two verticals (e.g. medical breath diagnostics and environmental leak detection).

12 Conclusion

This report presented the first end-to-end numerical evaluation of the Delin D_CIR metric in a chaotic sensing environment. In all tested simulation regimes—including breath VOC, methane leak and CBRN-style motifs—D_CIR achieved near-perfect separability, consistently outperforming classical correlation-based detection and a range of standard machine-learning baselines.

Beyond these numerical results, the broader implication is that D_CIR is more than a high-performing detector: it is a **universal decision engine** that quantifies whether a chaotic environment contains enough statistical structure to support reliable detection or reconstruction. This elevates EMDS/CIR from a detection method to a framework for **feasibility assessment**, experimental optimisation and R&D decision-making.

We also introduced the four-level EMDS architecture, which structures the technology into: (1) software and simulation tools, (2) OEM-ready embedded modules, (3) fully integrated sensing devices and (4) vertical platforms for clinical, industrial, environmental and defense applications. This layered view provides a realistic roadmap for scientific validation, technological deployment and market adoption.

Future work includes extended simulations, real-world datasets, laboratory validation and the development of EMDS-ready sensor modules. If confirmed experimentally, EMDS/CIR has the potential to reshape how ultra-weak signals are detected in chaotic environments and to establish D_CIR as a new standard for distinguishability evaluation across multiple industries.

Appendix: Mathematical Structure and Validation Protocol of D_CIR

A. Core Derivation

The mathematical derivation of D_CIR follows from the Neyman–Pearson likelihood ratio test for two Gaussian-distributed hypotheses with potentially different variances. The normalized separation

$$D = \frac{|\mu_1 - \mu_0|}{\sqrt{2(\sigma_x^2 + \sigma_\eta^2)}}$$

captures the effective contrast between chaotic background statistics and motif-induced perturbations. The distinguishability metric is then defined as

$$D_{\text{CIR}} = \Phi(D) - \frac{1}{2}.$$

B. Interpretation as an Information-Bearing Capacity

In the EMDS/CIR framework, μ_0, μ_1, σ_x^2 and σ_η^2 are not merely nuisance parameters but represent the **physical information content** of the environment–sensor pair. D_CIR quantifies the extent to which the environment carries distinguishable structure beyond noise.

C. Validation Protocol

D_CIR is used within a structured evaluation pipeline:

1. Estimate chaotic background statistics from data or simulation.
2. Inject motifs under controlled amplitude and spatial structure.
3. Compute D_CIR and compare against feasibility thresholds.
4. Classify the configuration into red/amber/green decision zones.
5. Use results to refine sensor architecture or abandon unfeasible designs.

D. Software Implementation

A reference implementation of D_CIR can be provided as a Python/C++ toolkit, enabling integration into simulation workflows, embedded systems and OEM evaluation modules.