engineering proceedings

MDPI

Proceeding Paper

The Energy-Dispersion Index: A Unified Metric for Automated Variational Mode Decomposition in Bearing Fault Diagnosis †

Ikram Bagri *, Achraf Touil and Ahmed Mousrij

Electrical and Mechanical Engineering Department, Faculty of Science and Technology, Hassan 1st University, Settat 26000, Morocco; ac.touil@uhp.ac.ma (A.T.); ahmed.mousrij@uhp.ac.ma (A.M.)

* Correspondence: i.bagri@uhp.ac.ma

† Presented at the 8th edition of the International Conference on Advanced Technologies for Humanity (ICATH 2026), Rabat, Morocco, 15–17 July 2026.

Abstract

Variational mode decomposition (VMD) is a mathematically formulated signal processing method that targets the identification of multi-scale temporal information. It has gained significant interest in vibration-based bearing fault diagnosis, though it suffers from a critical bottleneck: parameter tuning. To this day, there exists no consensus method for automating the selection of the main VMD parameters: the mode count (K), the penalty factor (α) and the reconstruction tolerance (τ). The issue stems from the focus of conventional evaluation metrics on isolated aspects of decomposition, without taking into account the underlying physics of the decomposition. This paper bridges this gap by introducing the Energy-Dispersion Index (EDI). The latter is a composite metric that evaluates information concentration in decomposed modes on the basis of the joint energy–entropy duality. Following its initial formulation, the metric was subjected to a rigorous experimental plan for empirical validation: a full-factorial experiment with 8640 decompositions, three datasets, three noise families, and four SNR levels, comparing EDI to the literature metrics. The statistically analyzed results provided significant proof of the metric's outperformance: high correlation with ground-truth fidelity metrics ($r = 0.681$–0.763), a mean fidelity of 0.847, and a 53.2% improvement in orthogonality over Kurtosis along with significant noise reduction (11–25%, the lowest gain being over Envelope Entropy) compared to literature metrics. Furthermore, EDI demonstrated intrinsic coupling of the K and α parameters; it adjusted its behavior based on noise levels without manual recalibration of α. As such, the metric demonstrated its practicality and robustness for fully automated signal processing frameworks.

Keywords: variational mode decomposition; automated parameter selection; bearing fault diagnosis; physics-informed signal processing; Industry 4.0

Check for updates

Academic Editors: Brahim El Bhiri, Ayman S. Mosallam and Younes Karfa Bekali

Published: 25 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

In manufacturing and industrial contexts, bearing failures account for a disproportionate share of unplanned mechanical downtime. The economic argument for predictive maintenance is well established [1] —detecting a fault before it propagates is almost always cheaper than responding after the fact.

Vibration measurement is the primary tool for bearing condition monitoring. A defective bearing produces repeated impact events at characteristic frequencies determined by the bearing geometry and shaft speed, and these are in principle recoverable from the raw acceleration signal [2]. In practice, this is difficult: the fault-related energy competes with

 https://doi.org/10.3390/engproc2026168027

contributions from adjacent machinery, power-line harmonics, and load-induced spectral variations [3], and the fault transient is typically the weakest component in the recording.

Variational Mode Decomposition (VMD) addresses this by reformulating component extraction as a constrained optimization problem [4]. Rather than the recursive sifting used by empirical mode decomposition (EMD)—which tends to generate spurious modes under noise [5]—VMD searches simultaneously for K band-limited components whose sum matches the input, solving:

$$\min_{\{u_k\},\{\omega_k\}} \left\{ \sum_{k=1}^{K} \left\| \partial_t \left[\left(\delta(t) + \frac{j}{\pi t} \right) * u_k(t) \right] e^{-j\omega_k t} \right\|_2^2 \right\} \quad (1)$$

subject to $\sum_{k=1}^{K} u_k = f(t)$. Here u_k are the modal components, ω_k their center frequencies, and K the mode count. VMD is less sensitive to measurement noise than recursive methods and does not depend on the empirical sifting stopping criteria of EMD.

The limitation is that VMD cannot run without three numerical inputs, and wrong values for any of them produce degraded decompositions (Table 1).

Table 1. VMD hyperparameters and their practical significance.

Parameter	Symbol	Function	Impact of Misspecification
Mode count	K	Number of components to extract	Too few: modes fuse together, losing the fault. Too many: noise gets its own mode, masquerading as a signal.
Penalty coefficient	α	Bandwidth constraint on each mode	Too low: spectral content bleeds across modes. Too high: the fault transient gets squeezed out along with the noise.
Reconstruction tolerance	τ	Strictness of the sum constraint	Too tight: overfitting to noise. Too loose: transient features get blurred out.

For a long time, parameter selection was performed manually: a practitioner would run several decompositions with different parameter combinations, inspect the resulting mode shapes visually, and retain the configuration that appeared most physically consistent [6]. This approach does not scale—it requires domain expertise, is time-consuming, and has no place in automated monitoring systems operating on edge hardware.

The automated selection methods proposed to date address the problem only partially, and a closer examination reveals that each leaves a distinct gap that no single existing metric closes.

Frequency-tracking heuristics [7] presuppose prior knowledge of the fault frequency band, which is rarely available in practice; they also assume that the fault signature remains confined to a stable frequency interval, an assumption that breaks down under variable load and speed conditions. Correlation-based maximization [8] selects modes that correlate well with the raw signal, but this criterion does not discriminate between fault-carrying modes and noise modes that happen to share the same spectral envelope—a noise-dominated mode can score highly on correlation while contributing nothing diagnostically useful. Envelope entropy minimization captures temporal concentration—fault transients are indeed sparse in time—but says nothing about whether the concentrated component actually carries meaningful energy; a weak artifact can satisfy the entropy criterion as well as a genuine fault transient, so low entropy alone is not a sufficient condition for diagnostic relevance. Kurtosis-based selection [6,9] is the most widely used approach—faults are impulsive, and Kurtosis measures impulsiveness. The problem is spectral blindness: Kurtosis cannot distinguish a bearing defect from electromagnetic interference introduced by a nearby inverter, and

a single outlier can dominate the score, making it unstable under non-stationary noise. Frequency-band optimization methods such as the fast kurtogram [9] do not qualify as an internal selection criterion for VMD, given that they are built around fixed filterbanks and not around VMD's adaptive decomposition.

Metaheuristic approaches search the parameter space globally instead of relying on a closed-form rule, but they add tuning parameters of their own and face prohibitive runtimes, exceeding 30 min per signal [6], which places them outside the budget of real-time or edge-deployed monitoring.

Taken together, these limitations share a common root: every existing metric evaluates decomposition quality along a single axis—temporal impulsiveness (Kurtosis), amplitude concentration (Envelope Entropy), or ordinal complexity (Permutation Entropy)—without jointly accounting for the two physical properties that distinguish a bearing fault transient from background noise. A fault in vibration data is characterized by a two-dimensional signature: high energy density concentrated in brief, repeating bursts, and correspondingly low temporal entropy. Noise, by contrast, is broadband with dispersed energy and high entropy. A metric that scores decomposition quality on only one of these axes will misclassify components that satisfy that axis by coincidence rather than by physical relevance. No metric in the current literature combines energy and entropy into a single evaluation criterion grounded in the physics of bearing fault generation. As such, a metric that quantifies decomposition quality should operate in the energy–entropy plane and ensure the distinction between noise and fault modes.

Under these specifications, the Energy-Dispersion Index (EDI) was developed. Its core logic relies on a normalized two-dimensional plane defined by energy density and temporal entropy. Each VMD mode is projected into this plane and is scored according to its position; a mode that exhibits an inverse energy/entropy relationship scores well—it is either a fault-carrying mode or a noise-isolation mode. The rationale for a positive score for a noise mode resides in the core VMD objective: decomposing a signal into its principal modes and isolating the noise into a single mode [4]. Additionally, the proposed metric does not depend on external thresholds, thanks to its adaptive statistical boundaries.

This paper extends the EDI framework introduced in our prior work [10] in three directions. First, we conduct a substantially expanded experimental validation: a full-factorial design with 8640 decompositions spanning three public datasets, three noise families, and four SNR levels from −5 dB to +15 dB, comparing EDI against Kurtosis, Envelope Entropy, and Permutation Entropy on three independent ground-truth benchmarks. Second, we analyze the statistical interaction between the VMD parameters K and α, and between K and τ, to determine whether these parameters can be tuned independently or whether their effects are coupled—a question with direct consequences for the design of automated search procedures. Third, we characterize the conditions under which EDI fails, examine whether high-performing parameter configurations cluster or scatter across the search space, and test whether EDI adapts its selected parameters as SNR changes without any external recalibration.

The contribution of this work, positioned against the gaps identified above, is a metric that evaluates decomposition quality on the energy–entropy plane rather than on any single axis, that requires no external threshold calibration, and that has been validated against multiple independent benchmarks across a wide range of noise conditions. The remainder of this paper is organized as follows: Section 2 presents the EDI formulation and its theoretical motivation; Section 3 describes the experimental design; Section 4 reports the results; Section 5 discusses the findings in relation to the existing literature; and Section 6 draws conclusions and outlines limitations.

https://doi.org/10.3390/engproc2026168027

2. The Energy-Dispersion Index: Formulation and Properties

2.1. Physical Motivation and Theoretical Foundation

The physical origin of bearing fault signatures is well understood. Each time a surface defect on a bearing element contacts another surface, a short mechanical impact is produced. These impacts repeat at a rate set by the bearing geometry and shaft speed, generating a vibration signal that is energy-dense in brief windows and quiet in between. A healthy bearing generates no such pattern—its vibrational energy is spread more evenly across time.

It is this temporal concentration that VMD-based diagnostics needs to capture [4]. Existing metrics each see only part of it. Kurtosis responds to impulsive peaks—a useful property—but a single outlier can dominate the score, and the metric is blind to what is happening in the frequency domain. Envelope Entropy measures how compactly amplitude is distributed in time, but ignores whether that distribution carries significant energy. A weak noise artifact that happens to be temporally compact will score as well as a genuine fault transient.

From an information-theoretic standpoint, fault transients (C_F) are super-Gaussian: energy is highly localized and Shannon entropy is low, consistent with minimum-entropy deconvolution theory [11]. Background noise (C_N) sits near the Gaussian end: energy broadly distributed, entropy near its variance-determined maximum. These two classes occupy opposite corners of an energy–entropy plane—a metric defined on that plane can, in principle, distinguish them.

EDI is built around this separation. Rather than collapsing energy and entropy into a single scalar, it maps each extracted mode onto the normalized two-dimensional plane and scores its position relative to the other modes in the same decomposition. Modes that are both energetically significant and temporally concentrated are scored positively; modes that are significant in only one dimension, or in neither, receive a neutral score, and modes whose energy and entropy deviate in the same direction are penalized.

2.2. Mathematical Formulation

Let $u_k(n)$ be a component extracted by VMD from an input sequence $x(n)$ of length N. EDI maps every u_k to a point in a normalized Information Plane $\mathcal{P} = \{(\eta_k, H_k) \in \mathbb{R}^2\}$ defined by two descriptors, given in Table 2.

Table 2. EDI descriptors and their interpretations.

Descriptor	Definition	Interpretation
Normalized Energy Density (η_k)	$\eta_k = \frac{\sum_{n=1}^{N} u_k(n)^2}{\sum_{n=1}^{N} x(n)^2}$	Fraction of total input energy held by component k. Near 1: a dominant component, likely a fault transient. Near 0: energetically negligible—noise or residual.
Temporal Information Entropy (H_k)	$H_k = -\frac{1}{\log_2 N} \sum_{n=1}^{N} p_n \log_2 p_n$	Shannon entropy of the squared-amplitude distribution, normalized by $\log_2 N$, with $p_n = \lvert u_k(n)\rvert^2 / (\sum_i \lvert u_k(i)\rvert^2 + \epsilon)$ and ϵ a small constant preventing division by zero, giving $H_k \in [0, 1]$. Low (approaching 0): sharp, impulsive—fault transients. High (approaching 1): spread out, noise-like.

Since the absolute values of η_k and H_k vary with the signal and noise conditions, direct thresholding would require per-application calibration. EDI avoids this through internal Z-

score normalization: within each decomposition, the ensemble mean (μ_η, μ_H) and standard deviation (σ_η, σ_H) are computed across all K modes, and each mode is standardized as:

$$z_{\eta,k} = \frac{\eta_k - \mu_\eta}{\sigma_\eta}, \quad z_{H,k} = \frac{H_k - \mu_H}{\sigma_H} \tag{2}$$

The signs of $(z_{\eta,k}, z_{H,k})$ partition the plane into four quadrants (Figure 1), each with a distinct physical interpretation.

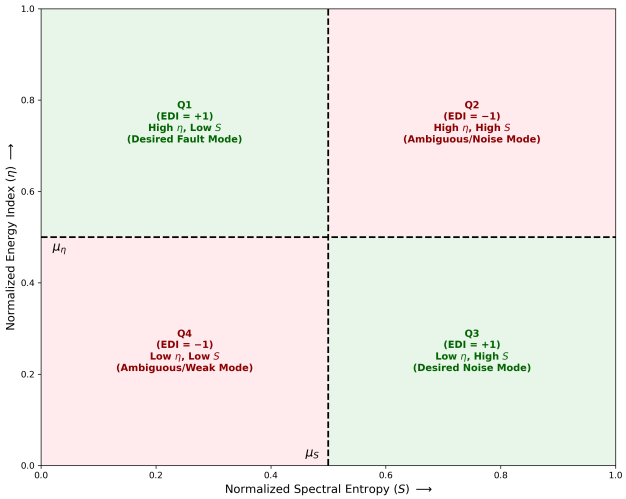

Figure 1. Information Plane segmentation into four quadrants. Q1 ($z_\eta > 0$, $z_H < 0$): physically consistent fault modes (high energy, low entropy); Q3 ($z_\eta < 0$, $z_H > 0$): noise signatures (low energy, high entropy); Q2 and Q4 (same-sign Z-scores): artifacts from over-decomposition or numerical residuals. The central region ($|z| \leq 1$) corresponds to H_0. Adapted from [10].

Mode classification is formalized through a ternary hypothesis test. The null hypothesis H_0 applies when at least one of the two Z-scores falls within one standard deviation of the mean—the mode does not carry a sufficiently strong signal in either dimension to be classified with confidence. Under H_1^+, the energy and entropy Z-scores have opposite signs and both exceed unity in magnitude, placing the mode in Q1 or Q3—configurations that are physically expected. Under H_1^-, the two Z-scores share the same sign with magnitudes above one, placing the mode in Q2 or Q4—configurations that violate the expected energy–entropy tradeoff and are indicative of decomposition artifacts.

$$\begin{cases} H_0: & |z_{\eta,k}| \leq 1 \text{ or } |z_{H,k}| \leq 1 \\ H_1^+: & \text{sign}(z_{\eta,k}) \neq \text{sign}(z_{H,k}),\ |z_{\eta,k}|, |z_{H,k}| > 1 \\ H_1^-: & \text{sign}(z_{\eta,k}) = \text{sign}(z_{H,k}),\ |z_{\eta,k}|, |z_{H,k}| > 1 \end{cases} \qquad \Psi_k = \begin{cases} +1, & H_1^+ \text{ (Q1, Q3)} \\ -1, & H_1^- \text{ (Q2, Q4)} \\ 0, & H_0 \text{ (center)} \end{cases} \tag{3}$$

The global EDI value aggregates the mode-level scores into a single figure of merit:

$$\text{EDI} = \frac{1}{K} \sum_{k=1}^{K} \Psi_k \in [-1, 1], \tag{4}$$

where Ψ_k is the ternary classification score from Equation (3). The global EDI aggregates mode-level decisions: a score near +1 indicates that most modes fall in Q1 or Q3, corresponding to effective fault–noise separation; a score near 0 indicates ambiguity; a score near −1 indicates that modes are concentrated in Q2 or Q4, typically a symptom of over-decomposition or severe noise. The compact $[-1, 1]$ range keeps the optimization landscape bounded, which is advantageous for gradient-free search. Importantly, EDI is computed

entirely within each decomposition using internal Z-score normalization, so it requires no external reference signal, no per-application threshold calibration, and is invariant to signal amplitude scaling and sampling rate.

3. Experimental Validation Methodology

The study used three publicly available bearing fault datasets: the CWRU benchmark [12], the HUST collection [13], and the University of Ottawa archive [14]. Together they span deep-groove ball and cylindrical roller geometries, sampling rates from 12 kHz to 51.2 kHz, and fault severities ranging from early-stage spalling to advanced outer-race damage. Preprocessing followed the procedure from our prior work [10].

3.1. Full-Factorial Design of Experiments

Three evaluation criteria drove the experimental design: whether EDI-guided parameter selection produces decompositions with higher fidelity to a known clean reference; whether it reduces spectral leakage between extracted modes; and whether it improves fault transient resolution in the time–frequency plane. A full-factorial structure was used so that all factor combinations were tested, rather than evaluating at a single SNR or noise type.

3.1.1. Experimental Configuration

Four factors were varied across all levels:

- **Signal class.** Six classes (C0–C5), each corresponding to a distinct vibration regime.
- **Noise type.** Additive White Gaussian, Laplacian (heavy-tailed), and Pink ($1/f$) noise—three profiles chosen to reflect real industrial contaminants rather than a single idealized model.
- **SNR level.** 15, 5, 0, and -5 dB, covering conditions from mild interference to strong burial of the fault signature.
- **Optimization metric.** EDI, Envelope Entropy, Kurtosis, and Permutation Entropy.

This yields $6 \times 3 \times 4 \times 4 = 288$ treatment combinations. For each combination, 30 independent decompositions were run with parameter values drawn randomly from the following ranges:

- $K \in [3, 7]$ (mode count).
- $\alpha \in [200, 3000] \times (f_s/12000)$ (bandwidth penalty, normalized to the sampling rate).
- $\tau \in \{10^{-5}, 10^{-6}, 10^{-7}, 10^{-8}\}$ (reconstruction tolerance).

The 30-run replication per cell was chosen to provide sufficient observations for estimating within-cell variance—an important consideration when comparing metrics that may respond differently to random parameter initialization. This brings the total number of VMD decompositions to $288 \times 30 = 8640$.

3.1.2. Reference Metrics and Evaluation Framework

Using a single benchmark to evaluate a candidate metric creates the risk that a metric performs well on one criterion while failing silently on another. To guard against this, all four candidates were assessed against three independent ground-truth measures. Table 3 lists both the optimization metrics and the reference benchmarks.

The three benchmarks are not interchangeable. A decomposition can show high fidelity to the clean reference while still mixing spectral content between modes. Conversely, a well-separated decomposition may have discarded the fault signature in the process of separating. Fidelity, orthogonality, and SNRE each catch a different kind of failure, so a metric that scores poorly on any one of them has a real limitation regardless of what the other two show.

https://doi.org/10.3390/engproc2026168027

Table 3. Candidate quality metrics and ground-truth benchmarks.

Quantity	Expression	Physical Significance	Optimal Direction
Candidate optimization metrics			
Kurtosis	$\frac{\mathbb{E}[(u_k-\mu)^4]}{\sigma^4}$	Fourth-order standardized moment; higher values flag impulsive content, though susceptible to outliers and insensitive to spectrum	Increase
Envelope Entropy	$-\sum_i p_i \log p_i,\ p_i = \frac{\lvert a_k(t_i)\rvert}{\sum_j \lvert a_k(t_j)\rvert}$	Entropy of the Hilbert envelope amplitude; lower values indicate concentration in time, yet disregards frequency features	Decrease
Permutation Entropy	$-\sum p(\pi)\log p(\pi)$	Entropic complexity from ordinal rank ordering [15]; captures temporal regularity patterns but not energy or spectrum	Decrease
Ground-truth reference benchmarks			
Modal fidelity	$\max_k \lvert\rho(u_k, x_{\text{clean}})\rvert$	Peak Pearson correlation linking extracted modes to noiseless reference; quantifies how faithfully each mode reproduces the true fault signature	Increase
Orthogonality score	$1-\frac{\sum_{i\neq j}\lvert\langle u_i, u_j\rangle\rvert}{\lVert x_{\text{clean}}\rVert^2}$	Separation score derived from the normalized cross-correlation among mode pairs; higher values imply cleaner spectral separation	Increase
SNRE	$10\log_{10}\left(\frac{\text{SNR}_{\text{out}}}{\text{SNR}_{\text{in}}}\right)$	Logarithmic SNR gain; measures the degree to which VMD enhances fault clarity relative to input noise	Increase

4. Results

4.1. Correlation with Ground-Truth Performance

The correlation results from all 8640 trials are reported in Table 4. EDI returned the highest values on all three benchmarks, with a margin in Pearson correlation of 0.12 to 0.16 over the nearest competitor, Envelope Entropy. Pearson and Spearman coefficients were close in every row; the ranking is not an artifact of distributional assumptions.

Table 4. Correlations between optimization metrics and reference criteria (8640 trials). Bold marks the best result per row.

Benchmark	EDI		Kurtosis		Envelope Entropy		Permutation Entropy	
Reference	**Pearson**	**Spearman**	**Pearson**	**Spearman**	**Pearson**	**Spearman**	**Pearson**	**Spearman**
Fidelity	**0.763**	**0.741**	0.524	0.512	0.601	0.588	0.487	0.493
Orthogonality	**0.681**	**0.669**	0.412	0.408	0.523	0.519	0.398	0.401
SNRE	**0.698**	**0.682**	0.491	0.485	0.578	0.571	0.455	0.461

4.2. Metric-Guided Selection and Statistical Significance

Under metric-guided parameter search, EDI-driven decompositions reached a mean fidelity of 0.847. Kurtosis-guided runs gave 0.782 and Envelope Entropy-guided runs 0.809, corresponding to EDI gains of 8.3% and 4.7%. On orthogonality the difference was more pronounced: 0.763 for EDI versus 0.498 for Kurtosis and 0.601 for Envelope Entropy, a 53.2% improvement over Kurtosis. SNRE gains over the baselines fell between 11% (over Envelope Entropy) and 25% (over Permutation Entropy) (Table 5).

A Friedman test returned $\chi^2 = 87.4$ ($p < 0.001$). Post hoc Nemenyi comparisons separated EDI from all three baselines at $p < 0.01$ (Figure 2a); no significant difference was found among the baselines themselves. Runtimes were similar across all four methods, so the gains do not come at a computational cost.

Table 5. Performance under metric-guided optimization.

Metric	Fidelity	Orthogonality	SNRE (dB)	Runtime (s)
EDI	**0.847 ± 0.062**	**0.763 ± 0.081**	**8.42 ± 1.23**	4.8
Kurtosis	0.782 ± 0.095	0.498 ± 0.132	6.91 ± 1.67	4.6
Envelope Entropy	0.809 ± 0.078	0.601 ± 0.109	7.58 ± 1.41	4.7
Perm. Entropy	0.769 ± 0.102	0.467 ± 0.145	6.72 ± 1.79	4.7

Values are mean ± standard deviation. Bold values indicate the best result in each column.

A multiple linear regression was estimated with fidelity as the dependent variable and all four metrics as predictors. EDI contributed the largest coefficient, $\beta = 0.621$ ($p < 0.001$), against a combined baseline contribution of $\sum \beta_{\text{baseline}} = 0.423$ (Table 6). Envelope Entropy retained a secondary contribution ($\beta = 0.231$); Kurtosis was marginal ($\beta = 0.118$, $p = 0.012$); Permutation Entropy was not significant ($p = 0.089$). Total explained variance was $R^2 = 0.671$. The fact that EDI dominated even in a joint model alongside the other metrics indicates it carries information about decomposition quality that the single-criterion approaches do not capture.

Table 6. Linear regression coefficients predicting fidelity.

Predictor	β	p-Value
EDI	**0.621**	<0.001
Envelope Entropy	0.231	<0.001
Kurtosis	0.118	0.012
Perm. Entropy	0.074	0.089

$R^2 = 0.671$; Adjusted $R^2 = 0.669$. Bold indicates the largest coefficient.

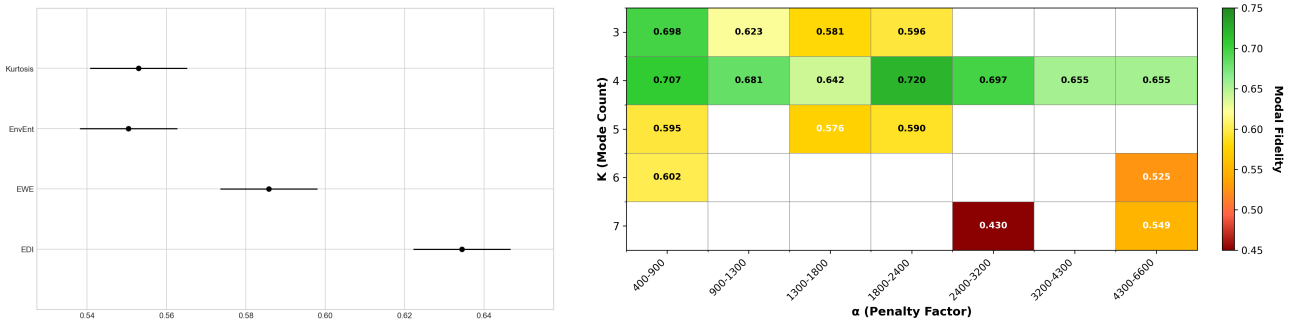

(**a**) Statistical separation among methods. (**b**) Coupling landscape of K and α.

Figure 2. Statistical validation. (**a**) Tukey HSD 95% confidence intervals. (**b**) Optimal region at $K = 4$, $\alpha \approx 1000$–2000 (green).

4.3. Parameter Interaction and Intelligent Coupling

Two-way ANOVA on the 63 optimized configurations gave a significant $K \times \alpha$ interaction ($F = 2.35$, $p = 0.032$; Figure 2b).

Interestingly, the K main effect was strongly significant ($p < 0.0001$) while the α main effect returned $p = 1.000$. This result shows that α has no detectable effect on its own and that its influence depends on K; the mode count must be set correctly first. The combination $K = 4$, $\alpha = 1063$ achieved a fidelity of 0.762.

The $K \times \tau$ picture was different. With $K = 4$ fixed, fidelity ranged from 0.646 to 0.683 across τ values (Figure 3), a variation of 0.037 that is negligible for practical purposes. This stability did not hold at higher K: at $K = 7$, fidelity dropped to 0.439 across all τ settings.

Given this, τ can reasonably be held constant, leaving K and α as the operative search dimensions, with a viable region bounded by $K \in [3, 4]$ and $\alpha \in [1000, 2000]$.

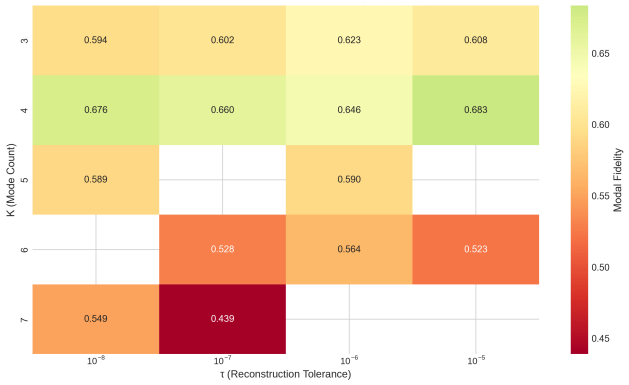

Figure 3. Fidelity across (K, τ) combinations.

4.4. Reliability and Parameter Convergence

Of the 63 EDI-optimized runs, one failure was identified—EDI above 0.5 with fidelity below 0.7—for an overall failure rate of 1.59%. That failure occurred at −5 dB Gaussian noise with $K = 3$, $\alpha = 876$: the harshest noise type, the lowest SNR, and an under-specified mode count. At SNR $\geq$ 0 dB, zero failures were recorded across all noise types and signal classes (Table 7 and Figure 4a).

Table 7. Failure statistics. Failures are defined as EDI > 0.5 with fidelity < 0.7.

Scenario	Trials	Failures	Rate
All trials	63	1	1.59%
Gaussian noise	23	1	4.35%
Pink/Laplacian noise	40	0	0.00%
SNR < 0 dB	15	1	6.67%
SNR ≥ 0 dB	48	0	0.00%

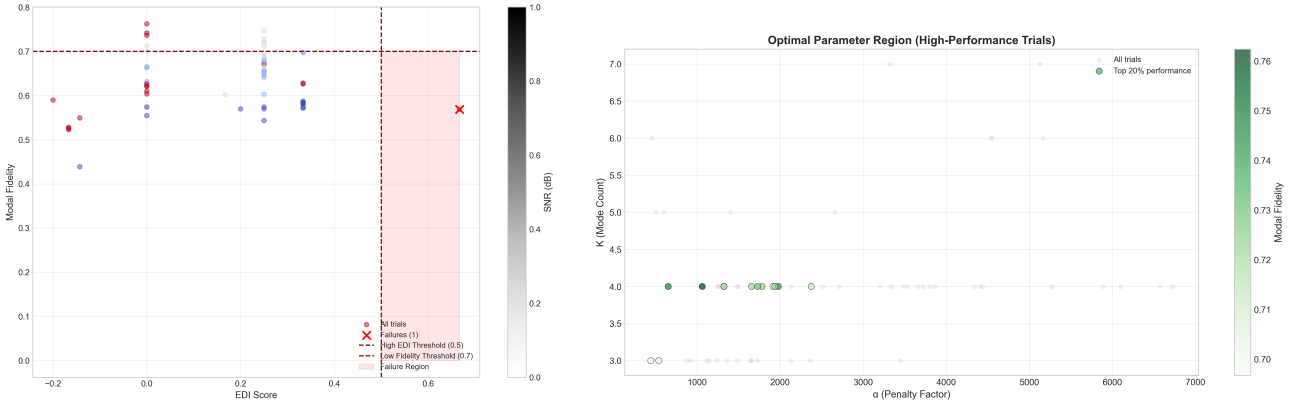

(**a**) Failure region and outliers.

(**b**) High-performance parameter clustering.

Figure 4. Reliability analysis. (**a**) Failure locations (red X) in high-EDI region. (**b**) Best-performing trials converge to $K = 4$, $\alpha \in [1000, 2000]$.

In the top quartile of runs by fidelity, $K = 4$ appeared in 87.5% of cases. The α values for those runs clustered within an interquartile range of [963, 1948]. The parameter selection

was not scattered (Figure 4b): the same region of the search space was consistently identified across different datasets and noise conditions.

4.5. Adaptive Behavior Across Signal-to-Noise Ratios

When SNR was varied from −5 dB to 15 dB (Figure 5), the mean K selected by EDI shifted from 3.7 to 4.3—a gradual increase in decomposition complexity that tracked signal quality. Mean α rose from approximately 1800 at −5 dB to 2700 at 0 dB, after which it leveled off. The best fidelity in this sweep was at SNR = 5 dB (0.686), with a gradual decline under heavier noise loading. None of these adjustments involved any external recalibration.

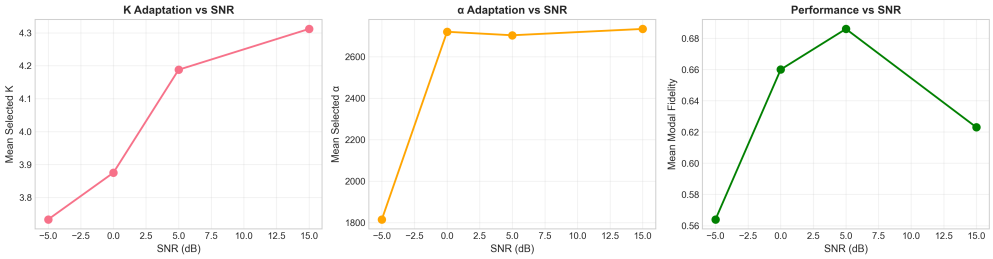

Figure 5. Parameter adaptation versus SNR.

5. Discussion

The four result sets reinforce each other, and some of what they show is not obvious from the numbers alone.

5.1. Complementary Coverage of Decomposition Quality

Single-criterion metrics are constrained by what they were formulated to measure. Kurtosis is sensitive to the temporal shape of the signal—impulsive peaks inflate it—but it has no access to spectral content. Envelope Entropy captures how amplitude is distributed over time, but gives no weight to whether that distribution carries significant energy. Permutation Entropy [15] captures ordinal patterns and temporal regularity, but neither energy nor spectrum. Each metric performs well on the dimension it was built for and provides no information on the other.

This pattern is consistent with the broader literature on VMD parameter selection. Methods based on the central frequency observation [7] and on correlation maximization [8] both assume that a mode's similarity to the raw signal is a sufficient proxy for its diagnostic value; the present results are consistent with this assumption failing when noise modes share the spectral envelope of the fault component. The Envelope Entropy approach correctly identifies temporal concentration as a desirable property, but the present findings confirm the concern that concentration without energy relevance can mislead—a weak, compact artifact can score well on entropy while carrying no fault information. Kurtosis-based selection [6,9] remains the most widely adopted approach in the bearing diagnosis literature, and the present results quantify its limitation: it achieved the lowest fidelity correlation ($r = 0.524$) and the lowest orthogonality correlation ($r = 0.412$) among the four candidates, and its regression coefficient was marginal ($\beta = 0.118$, $p = 0.012$). The spectral-blindness problem noted in the introduction is thus not a theoretical caveat but a measurable performance deficit under the tested conditions.

EDI's correlations with the three benchmarks ran from $r = 0.681$ to 0.763, above all competitors on all criteria. The individual values matter less than the cross-criterion pattern: EDI led on every benchmark, with a relative spread across fidelity, orthogonality, and SNRE that was the smallest of the four metrics. The metric does not appear to be trading performance on one criterion for performance on another.

The regression results are consistent with this. When fidelity was predicted from all four metrics jointly, EDI accounted for $\beta = 0.621$ of the coefficient mass, against a combined $\sum \beta = 0.423$ for the three baselines. Envelope Entropy contributed something independently ($\beta = 0.231$), but Kurtosis was at the margin of significance and Permutation Entropy was below it. The joint model explained $R^2 = 0.671$ of fidelity variance.

A single-axis metric improved with a correction factor would not replicate this. The EDI score depends on the joint position of a mode in the energy–entropy plane, not on either dimension taken alone. The regression coefficient reflects access to relational information—how the two dimensions interact within a decomposition—that no individual criterion can encode. This is the structural advantage that distinguishes EDI from the methods surveyed above: it is not an incremental improvement on an existing metric, but a different evaluation axis altogether.

5.2. Automatic Hyperparameter Alignment

The ANOVA findings provide clear evidence for the K–α coupling. With $p = 1.000$ for α's main effect, its influence is entirely contingent on K. The significant interaction at $p = 0.032$ confirms this: α only matters once K is correctly specified.

This K–α coupling appeared in all experimental configurations, which suggests it is structural to EDI's physics-based evaluation logic rather than an artifact of a particular dataset or noise profile. The practical implication is significant: it means that a two-stage search—first identifying K, then searching α within a narrow band—is justified by the data, and that a naive grid search treating the two parameters as independent would waste evaluations on configurations where α cannot compensate for a wrong K.

The τ insensitivity has a useful consequence. Since fidelity varied by less than 0.04 across all tested τ values at $K = 4$, fixing τ costs almost nothing in practice. That leaves K and α as the only parameters worth searching under the noise conditions tested here. This is a substantial reduction in search dimensionality compared with the original three-parameter space, and it directly benefits any automated optimization procedure built on top of EDI.

6. Conclusions

This paper introduced EDI as a physics-grounded composite metric for automated VMD parameter selection in bearing fault diagnosis. The experimental validation involved 8640 decomposition trials across three public datasets, three noise families, and SNR levels from −5 dB to +15 dB.

On correlation with ground-truth benchmarks, EDI led all competitors, with a margin in Pearson correlation of 0.12 to 0.16 over the nearest one, Envelope Entropy. In metric-guided parameter selection, it reached a mean fidelity of 0.847, with a 53.2% orthogonality advantage over Kurtosis and SNRE gains of 11–25%, the lowest being over Envelope Entropy. All comparisons between EDI and the baselines were significant at $p < 0.01$.

The analysis further identified the following:

1. A consistent $K \times \alpha$ coupling across all datasets and noise types, with near-null sensitivity to τ.
2. High-performing configurations clustered around $K = 4$, $\alpha \in [1000, 2000]$.
3. No failures were recorded at SNR ≥ 0 dB.
4. As SNR increased, EDI selected higher K and α values, demonstrating adaptive behavior without external recalibration.

These findings establish EDI as a reliable, physics-grounded replacement for heuristic VMD tuning in controlled laboratory and benchmark settings. The validated oper-

ating envelope—three public bearing datasets, three noise families, and a specific SNR range—defines the scope of the present study. Extension to gear and shaft fault types, on-line deployment on embedded hardware, and integration with end-to-end diagnostic pipelines remains to be demonstrated and is the subject of ongoing work.

Author Contributions: Conceptualization, I.B.; methodology, I.B.; software, I.B.; validation, I.B., A.T. and A.M.; formal analysis, I.B.; investigation, I.B.; data curation, I.B.; writing—original draft preparation, I.B.; writing—review and editing, A.T. and A.M.; visualization, I.B.; supervision, A.T. and A.M.; project administration, A.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data analyzed in this study are openly available: the CWRU bearing dataset [12], the HUST bearing dataset [13], and the University of Ottawa bearing dataset [14].

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Mobley, R.K. *An Introduction to Predictive Maintenance*, 2nd ed.; Butterworth-Heinemann: Woburn, MA, USA, 2002.
2. Randall, R.B.; Antoni, J. Rolling element bearing diagnostics—A tutorial. *Mech. Syst. Signal Process.* **2011**, *25*, 485–520. [CrossRef]
3. Lei, Y.; Yang, B.; Jiang, X.; Jia, F.; Li, N.; Nandi, A.K. Applications of machine learning to machine fault diagnosis: A review and roadmap. *Mech. Syst. Signal Process.* **2020**, *138*, 106587. [CrossRef]
4. Dragomiretskiy, K.; Zosso, D. Variational mode decomposition. *IEEE Trans. Signal Process.* **2014**, *62*, 531–544. [CrossRef]
5. Huang, N.E.; Shen, Z.; Long, S.R.; Wu, M.C.; Shih, H.H.; Zheng, Q.; Yen, N.-C.; Tung, C.C.; Liu, H.H. The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proc. R. Soc. Lond. Ser. A* **1998**, *454*, 903–995. [CrossRef]
6. Zhang, M.; Jiang, Z.; Feng, K. Research on variational mode decomposition in rolling bearings fault diagnosis of the multistage centrifugal pump. *Mech. Syst. Signal Process.* **2017**, *93*, 460–493. [CrossRef]
7. Li, Z.; Chen, J.; Zi, Y.; Pan, J. Independence-oriented VMD to identify fault feature for wheel set bearing fault diagnosis of high speed locomotive. *Mech. Syst. Signal Process.* **2017**, *85*, 512–529. [CrossRef]
8. Tang, G.; Wang, X. Parameter optimized variational mode decomposition method with application to incipient fault diagnosis of rolling bearing. *J. Xi'an Jiaotong Univ.* **2015**, *49*, 73–81.
9. Antoni, J. Fast computation of the kurtogram for the detection of transient faults. *Mech. Syst. Signal Process.* **2007**, *21*, 108–124. [CrossRef]
10. Bagri, I.; Touil, A.; Oucheikh, R.; Mousrij, A.; Hraiba, A.; Tahiry, K. The Energy-Dispersion Index (EDI) and cross-domain archetypes: Towards fully automated VMD decomposition for robust fault detection. *Vibration* **2026**, *9*, 16. [CrossRef]
11. Wiggins, R.A. Minimum entropy deconvolution. *Geoexploration* **1978**, *16*, 21–35. [CrossRef]
12. Loparo, K.A. Bearings Vibration Data Set. Case Western Reserve University Bearing Data Center. 2012. Available online: https://engineering.case.edu/bearingdatacenter (accessed on 15 March 2023).
13. Zhu, J.; Chen, N.; Shen, C. A new multiple-source domain adaptation fault diagnosis method with HUST bearing dataset. *IEEE Access* **2020**, *8*, 149627–149639.
14. Huang, H.; Baddour, N. Bearing vibration data collected under time-varying rotational speed conditions. *Data Brief* **2018**, *21*, 1745–1749.
15. Bandt, C.; Pompe, B. Permutation entropy: A natural complexity measure for time series. *Phys. Rev. Lett.* **2002**, *88*, 174102. [CrossRef] [PubMed]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.