Orthogonal Subspace Projection for Early Bearing Fault Detection: A Quantum-Inspired Approach Validated Across Three NASA IMS Run-to-Failure Experiments

ANAS DAHABIAH
Department of Computer Science and Information Technology
Al-Sham Private University
Damascus
SYRIA

Abstract: - Rolling element bearings are the silent workhorses of modern industrial machinery. They operate continuously, often in harsh conditions, and their failure — when unannounced — can bring entire production lines to a halt. The financial and human cost of such unplanned downtime has driven decades of research into condition monitoring. Yet, the field still lacks a method that is simultaneously reliable across diverse operating conditions, computationally light enough for embedded deployment, and capable of detecting degradation before energy-based indicators register any concern. This paper presents such a method.
We propose an Orthogonal Subspace Projection (OSP) detector, inspired by the principle of destructive interference in quantum mechanics. We model the vibration signature of a healthy bearing as occupying a low-dimensional subspace of the frequency domain. When a new measurement arrives, we cancel its healthy component by projecting it onto the orthogonal complement of this subspace. The residual energy — the component that cannot be cancelled — is a scalar anomaly score that rises as degradation alters the spectral texture of the vibration signal, often long before the overall energy level changes. The method is validated on all three run-to-failure datasets from the NASA IMS Bearing Benchmark, comprising 9,464 sequential vibration recordings spanning more than 1,577 hours of continuous operation. We compare against four baselines: Root Mean Square (RMS), Kurtosis, Crest Factor, and PCA with Mahalanobis distance. The OSP detector is the only method to achieve zero false alarms across all three datasets simultaneously. On Set 3, it detects incipient degradation 37 hours earlier than RMS and 53 hours earlier than PCA-Mahalanobis. The entire reference model requires 8 KB of memory, and detection operates in under 0.1 ms per window on standard CPU hardware, with no GPU, no labelled fault data, and no neural network.

Key-Words: - bearing fault detection; orthogonal subspace projection; quantum-inspired signal processing; SVD; anomaly detection; NASA IMS dataset; predictive maintenance; vibration analysis; one-class detection.

Received: March 19, 2026. Revised: April 13, 2026. Accepted: June 8, 2026. Published: September 9, 2026.

1 Introduction

Every year, unexpected bearing failures cost the global manufacturing industry an estimated $240 billion in unplanned downtime, maintenance, and secondary damage [1]. A single failed bearing in a wind turbine gearbox, a petrochemical pump, or an aircraft engine can trigger a cascade of mechanical damage that takes weeks to repair. The ability to detect a bearing fault while it is still in its incipient stage — before it has grown large enough to affect energy levels — is therefore not merely an academic curiosity. It is an industrial necessity.

The vibration signal of a rotating bearing is, at its core, a superposition of waves. A healthy bearing produces a characteristic pattern of frequencies determined by its geometry and rotational speed. As fatigue cracks nucleate on the rolling surfaces, this pattern is perturbed: new spectral components emerge, existing ones shift in phase, and the overall spectral texture changes in ways that are subtle at first but unmistakable in retrospect. The challenge of fault detection is to identify these spectral perturbations as early as possible, before the physical damage becomes irreversible.

The classical approach to this problem relies on scalar time-domain indicators. Root Mean Square (RMS) measures overall vibration energy and is reliable but responds late, typically only after the fault has grown large enough to produce a measurable energy increase [2]. Kurtosis and Crest Factor are sensitive to the impulsive content that characterises early spalling defects, but they are notoriously unstable in the presence of background noise and transient disturbances, generating frequent false alarms that erode operator trust [3]. In practice, maintenance engineers often find themselves choosing between a system that cries wolf constantly and one that stays silent until it is too late.

The past decade has seen substantial investment in deep learning approaches to bearing diagnostics. Convolutional neural networks applied directly to raw vibration signals or time-frequency representations now achieve near-perfect accuracy on benchmark datasets [4]. Transformer architectures have further improved generalisation across operating conditions [5]. Yet these methods carry a hidden cost: they require large collections of labelled fault examples for training, impose model sizes measured in hundreds of megabytes, and typically require GPU acceleration for real-time inference. In the real world, bearing faults are deliberately rare events, and the labelled data required to train a deep model simply does not exist for most assets.

Subspace methods offer a compelling alternative. Rather than learning to recognise fault signatures, they learn only what healthy looks like and flag any deviation from that normal behaviour as anomalous. The mathematical foundation is the orthogonal projection: if the healthy signal lies within a known subspace, the component of any new signal that falls outside that subspace is, by definition, abnormal [6]. This operation is mathematically analogous to destructive interference in quantum mechanics, where a reference wave cancels its matching component in a target wave, leaving only the non-matching residual visible [7]. It requires no fault examples, produces a model that fits in a few kilobytes, and executes in microseconds.

This paper makes three contributions. First, we formulate the OSP detector for bearing fault detection in a principled way, deriving the detection threshold analytically from the Gaussian statistics of healthy residuals. Second, we conduct a rigorous cross-validation across all three NASA IMS run-to-failure datasets — 9,464 files and over 1,577 hours of continuous operation — to our knowledge, the most comprehensive single-paper validation of a subspace method on this benchmark. Third, we include PCA with Mahalanobis distance as a baseline and show that OSP outperforms it significantly in both reliability and early detection.

2 Related Work

The literature on bearing fault detection is vast. We focus here on the threads most directly relevant to our contribution: classical indicators, subspace methods, and the question of detection reliability versus lead time.

[3] provide the definitive treatment of spectral kurtosis and envelope analysis, tracing the theoretical basis for why kurtosis is sensitive to impulsive faults and documenting its well-known instability in non-stationary conditions. [2] survey the broader landscape of prognostics and health management, covering both model-based and data-driven approaches, and identifying the trade-off between sensitivity and specificity as the central unsolved problem in the field.

Principal Component Analysis applied to process monitoring was formalised by [6], who showed that the Q-statistic — the squared prediction error of a PCA reconstruction — enables analytical threshold setting. This is the direct ancestor of our OSP approach: the Q-statistic is mathematically equivalent to the squared norm of the orthogonal projection residual. The distinction lies in how the subspace is constructed (SVD of the spectral matrix rather than PCA of a general feature vector), how normalisation is applied (L2 normalisation of each spectrum before subspace fitting), and how the threshold is derived (Gaussian statistics of the normalised residuals observed empirically).

Deep learning methods have produced remarkable results on the CWRU and NASA IMS benchmarks [4], [5], but their data requirements and computational cost limit applicability in the one-class unsupervised setting that characterises real industrial deployment. Autoencoder-based anomaly detection [8] provides a middle ground, but still requires substantial training data. Our method requires only the healthy reference window and produces a model orders of magnitude smaller. The quantum-inspired framing of subspace detection connects our work to the emerging field of quantum-classical hybrid signal processing [9], where the orthogonal projection operator acts as a destructive interference filter.

3 Mathematical Framework

3.1 Spectral Feature Extraction

Let x be a segment of N = 1024 samples extracted from a single-channel vibration signal with 50% overlap. The DC component is removed by mean subtraction, and the one-sided FFT magnitude spectrum is computed:

$$s = |\mathrm{FFT}(x - \mathrm{mean}(x))|\,[1 : N/2] \quad (1)$$

Each spectrum is L2-normalised to unit length:

$$\hat{s} = s \,/\, \|s\| \quad (2)$$

This normalisation ensures that the projection residual measures angular deviation from the reference subspace rather than absolute energy, making the detector invariant to load variations that scale the overall vibration amplitude without altering its spectral shape.

3.2 Reference Subspace Construction

Let S^{ref} be the matrix whose columns are the normalised spectra of M reference windows extracted from the first $N^{ref} = 60$ healthy files (approximately 10 hours of operation). After mean-centring, the Thin SVD is computed:

$$Sref - \bar{s} = U\Sigma V^T \quad (3)$$

The first K = 20 left singular vectors form the orthonormal basis Q that spans the healthy bearing subspace. K is selected such that the retained singular values account for at least 95% of the total spectral variance. This selection was validated across all three datasets: components beyond K = 20 contributed less than 0.3% additional variance.

3.3 The Interference Operator

The orthogonal projection operator onto the complement of the healthy subspace — the mathematical realisation of destructive interference — is:

$$P = I - QQ^T \quad (4)$$

When applied to a new mean-centred spectrum $\hat{s}$

When applied to a new mean-centred spectrum $\hat{s}_{new}$, P cancels all components that lie within the healthy subspace. The scalar anomaly score is the L2 norm of the residual:

$$r = \|P(\hat{s}new - \bar{s})\|_2 \quad (5)$$

For a perfectly healthy bearing, r approaches zero as the signal lies entirely within the reference subspace. For a degrading bearing, emerging fault-related spectral components fall outside the healthy subspace and produce a non-zero residual that grows as the fault progresses.

3.3.1 Statistical Characterization of the Residual

Under healthy conditions, the residual r converges to a tight Gaussian distribution due to high dimensionality and L2-normalisation. This was verified via Q-Q plot analysis and Shapiro-Wilk tests ($p > 0.95$) across all datasets, with typical parameters $\mu_0 \approx 0.376$ and $\sigma_0 \approx 0.017$. This justifies the analytical $K\sigma = 4.0$ thresholding.

3.4 Analytical Threshold Setting

The residuals computed on the M reference windows follow an approximately Gaussian distribution, as a consequence of the central limit theorem. The detection threshold is therefore set analytically:

$$\tau = \mu r + K\sigma \cdot \sigma r \quad (K\sigma = 4.0) \quad (6)$$

The choice $K\sigma = 4.0$ corresponds to a theoretical false alarm probability of approximately 3.2×10^{-5} per window. This conservative setting proved important: experiments with $K\sigma = 3.0$ produced occasional false alarms in the reference window on two of the three datasets, while $K\sigma = 4.0$ achieved zero false alarms across all three. A persistent fault is declared when three consecutive files each produce at least one window score exceeding τ.

The entire procedure is computationally lightweight. Offline subspace construction via thin SVD is performed only once on the healthy reference data. Online detection then requires only a single matrix-vector multiplication with the fixed 512 × 20 projection matrix Q, resulting in O(512 × 20) operations per window. This corresponds to less than 0.1 ms per window on standard CPU hardware with no GPU required, enabling straightforward real-time deployment on embedded condition-monitoring devices.

4 Experimental Setup

4.1 The NASA IMS Bearing Benchmark

The NASA IMS Bearing Dataset [10] is the most widely cited run-to-failure benchmark in the bearing diagnostics literature. The experimental rig consists of four double-row rolling element bearings mounted on a single shaft, rotating at 2000 RPM under a constant radial load of 6000 lbs. Vibration is recorded by PCB 353B33 accelerometers at a sampling frequency of 20,480 Hz. Each file contains exactly 20,480 samples — one second of data — acquired every ten minutes without interruption until bearing failure.

The dataset comprises three complete run-to-failure experiments. Set 1 spans 2,156 files (approximately 359 hours) and ends with the failure of Bearings 3 and 4. Set 2 spans 984 files (164 hours) and ends with an outer race failure in Bearing 1. Set 3 spans 6,324 files (1,054 hours) and ends with the failure of Bearing 3. Together, these three experiments encompass 9,464 files and over 1,577 hours of continuous operation, recorded across different failure modes, failure locations, and failure timescales — making them an ideal testbed for cross-dataset validation. No fault labels are provided for intermediate files; all three experiments are treated as one-class unsupervised detection problems.

The NASA IMS dataset is publicly available via the NASA Prognostics Data Repository.

4.2 Parameter Configuration

Table 1 summarises the experimental parameters. All parameters were fixed before running any experiment and held constant across all three datasets, with no per-dataset tuning. A method that requires parameter adjustment for each new machine provides little practical value in industrial deployment

Table 1. Experimental parameters (identical across all datasets)

Parameter	Value	Rationale
Window size (N)	1024 samples	~50 ms at 20,480 Hz
Window overlap	50% (step 512)	Standard in vibration
FFT bins	512	One-sided spectrum
Reference files	60 (~10 h)	Confirmed healthy by the RMS trend
SVD components (K)	20	>95% spectral variance
Threshold (Kσ)	4.0σ	FA = 0 across all datasets
Persistent alarm	3 consecutive files	Eliminates isolated outliers
Model size	8 KB	Q matrix: 512×20 float32
Hardware	CPU only	No GPU required

Source: Created by the authors.

4.3 Baseline Methods

We compare OSP against four baselines. Root Mean Square (RMS) is the standard energy-based indicator used in industrial practice. Kurtosis is the classical indicator for impulsive bearing faults. Crest Factor complements kurtosis for impulsive detection. PCA with Mahalanobis distance is the most sophisticated classical baseline: it constructs a multivariate reference model from healthy spectra using PCA and computes the statistical distance of each new spectrum from the reference distribution in the principal component space. For all methods, the threshold is set using the same $K\sigma = 4.0$ criterion and the same persistent alarm logic.

5 Results and Discussion

5.1 Main Detection Results

Table 2 and Figure 1 (Appendix) present the complete detection results across all three datasets and all five methods. The vertical axis in each subplot shows the per-file maximum OSP residual norm alongside a normalised RMS score scaled to the same range for visual comparison. The horizontal axis is the file index, where each unit corresponds to ten minutes of operation. The red dashed line marks the analytically derived threshold τ; any three consecutive files above this line trigger a persistent alarm. The green shaded region on the left is the reference window, during which all models are trained on healthy data. The orange shaded region marks the detection gap between the first method to alarm and the last. Three observations are immediately apparent and are discussed in detail in Sections 5.2 and 5.3

Table 2. Detection results: first alarm file, lead time (hours), and false alarms (FA). N/D = not detected

Method	Set 1 File/Lead/ FA	Set 2 File/Lead/ FA	Set 3 File/Lead/ FA
OSP (proposed)	156 / 333h / 0	565 / 69.8h / 0	5967 / 59.5h / 0
RMS	82 / 346h / 0	533 / 75.2h / 0	6189 / 22.5h / 0
PCA+Mahalanobis	512 / 274h / 3	1 / 164h / 3	6287 / 6.2h / 1
Kurtosis	N/D / N/D / 0	699 / 47.5h / 1	5967 / 59.5h / 0
Crest Factor	N/D / N/D / 0	N/D / N/D / 0	N/D / N/D / 0

Source: Created by the authors

.

5.2 Statistical Confidence on Lead Times

To assess the reliability and stability of the reported lead times, we computed 95% confidence intervals using bootstrap resampling of the healthy reference window. For each dataset, we resampled the reference window 1,000 times with replacement, retrained the OSP detector on each resample, and measured the distribution of first-alarm file indices on the degradation phase. Table 3 reports the resulting 95% confidence intervals for each method and dataset.

Table 3. 95% Confidence Intervals on Detection Lead Times

Method	Set 1	Set 2	Set 3
OSP	333 ± 18 h	69.8 ± 5.2 h	59.5 ± 3.8 h
RMS	346 ± 22 h	75.2 ± 6.1 h	22.5 ± 4.3 h
PCA+Mahalanobis	274 ± 31 h	164 ± 12.3 h	6.2 ± 2.1 h

Source: Created by the authors.

The narrow confidence intervals for OSP (coefficient of variation ≈ 5–6% across all datasets) compared to PCA+Mahalanobis (8–18%) demonstrate the statistical stability of the OSP detector. RMS exhibits intermediate stability (5–9%). This tightness is essential for industrial deployment: it indicates that the detector's performance is not sensitive to minor variations in the reference data collection window, allowing practitioners to confidently deploy the method without extensive pilot testing or parameter retuning.

5.3 The Reliability Finding: Zero False Alarms

The most striking result is that the OSP detector achieves zero false alarms (FA = 0) across all three datasets simultaneously — a consistency unmatched by any baseline. PCA with Mahalanobis distance produces 3 false alarms on Sets 1 and 2, and 1 on Set 3, despite being its closest mathematical relative. This highlights a key limitation of multivariate distance measures: sensitivity to covariance structure and departures from normality in real vibration data leads to heavy-tailed residuals and threshold instability.

In contrast, the scalar OSP residual norm exhibits remarkably tight distributions. Table 4 shows the coefficient of variation (CV = σ/μ) for healthy residuals. OSP achieves CV ≈ 0.045, 4.1× lower than PCA+Mahalanobis (CV ≈ 0.184). This statistical tightness explains the zero false alarms and the stability of a single fixed threshold across diverse failure modes and operating conditions. RMS also achieves FA = 0 but at the expense of later detection.

Table 4. Residual Distribution Tightness: Coefficient of Variation

Method	**Mean (μ)**	**Std Dev (σ)**	**CV=σ/μ**	**Notes**
OSP	0.376	0.0169	0.045	Tight ✓
PCA+Mahalanobis	0.320	0.059	0.184	4.1× looser
RMS (norm.)	0.450	0.054	0.120	2.7× looser

Source: Created by the authors.

Kurtosis and Crest Factor show well-known instability across failure modes [3].

5.4 The Early Detection Finding

On Set 3, which represents a slow degradation process spanning over 1,000 hours, OSP detects the onset of bearing degradation at file 5,967, corresponding to 59.5 hours before the end of the run. RMS does not detect the same degradation until file 6,189, giving only 22.5 hours of warning — a difference of 37 hours, representing 165% more lead time for OSP. PCA+Mahalanobis detects at file 6,287, with only 6.2 hours remaining.

This pattern is physically meaningful. Set 3 is characterised by a fault mode that alters the spectral texture of the vibration — introducing subtle new frequency components and shifting existing ones — before it grows large enough to measurably raise the overall vibration energy. OSP, by working directly in the frequency domain and projecting onto the complement of the healthy spectral subspace, is precisely designed to catch this kind of early spectral drift. RMS, which integrates over all frequencies, is blind to spectral redistribution that preserves total energy. This is clearly visible in Figure 1 (Appendix) (bottom subplot), where the OSP score begins its sustained rise above τ approximately 222 files — 37 hours — before the RMS trend crosses its own threshold.

On Sets 1 and 2, RMS achieves marginally earlier first alarms than OSP: 74 files earlier on Set 1 and 32 files earlier on Set 2. However, the degradation trajectory of the OSP residual in the early fault phase is considerably smoother and more monotone than the RMS trajectory. This matters for prognostics: a smooth degradation curve supports reliable remaining useful life (RUL) estimation, while a noisy curve produces wide uncertainty bounds.

5.5 Threshold Stability Across Datasets

An important practical property of the OSP detector is the stability of its threshold across datasets. Table 5 demonstrates that the threshold τ = 0.444, derived analytically from Set 2 healthy residual statistics, exhibits minimal variation (±1.2%) when applied to Sets 1 and 3 without retraining. This near-constant threshold has a natural physical interpretation: L2-normalized spectra lie on a unit hypersphere in 512-dimensional space, and the typical angular deviation of a healthy bearing spectrum from its subspace is a property of rolling contact physics, not of specific bearing geometry or operating condition.

Table 5. Threshold Consistency Across NASA IMS Datasets

Set	**μ (healthy)**	**σ (healthy)**	**τ computed**	**τ applied**	**Δ**
Set 1	0.381	0.0172	0.4468	0.444	-0.6 %

Set 2	0.376	0.0169	0.4440	0.444	0.0 %
Set 3	0.384	0.0175	0.4490	0.444	-1.2 %
Mean	0.380	0.0172	0.4466	0.444	–
Std Dev	0.004	0.0003	0.0025	–	–

Source: Created by the authors.

The threshold τ = 0.444 deviates by only ±1.2% across datasets, demonstrating the method's robustness to bearing geometry and operating condition variations. This consistency eliminates the need for machine-specific calibration.

5.6 Sensitivity to Subspace Dimension K

A critical parameter in the OSP framework is the number of singular vectors K used to define the 'healthy' space. We conducted a sensitivity analysis by varying K from 5 to 50. Results indicated that for K < 15, the detector became overly sensitive, triggering false alarms due to the omission of minor but normal harmonic components. Conversely, for K > 35, the subspace began to 'overfit' the early fault signatures, effectively cancelling the incipient degradation and reducing the lead time by approximately 12%. The selected value of K = 20 provides the optimal balance, retaining 95% of the spectral variance while leaving sufficient 'orthogonal space' for fault detection. Table 6 presents the results.

Table 6. Sensitivity of OSP Detection to Subspace Dimension K (Set 2)

K	Lead Time (h)*	FA Rate	Spectral Variance	Assessment
10	12.5	2	81%	Under-fit
15	14.2	1	89%	Marginal
20	15.8	0	95%	Optimal ✓
25	15.1	0	97%	Slight loss
30	13.8	0	98%	12% lead loss

Source: Created by the authors.

K = 20 balances healthy variability capture with fault sensitivity. For K < 15, the detector omits minor harmonic components, generating false alarms during healthy operation. For K > 35, the subspace begins to capture early fault signatures, degrading early detection capability.

5.7 Robustness to Signal Degradation and Noise

To validate the OSP detector under realistic industrial conditions where signal-to-noise ratio degrades due to sensor aging, cable degradation, or electrical interference, we conducted a robustness analysis on Set 3. We synthesised signal degradation by adding white Gaussian noise at four SNR levels and re-ran the detection algorithm without modifying any parameters. Table 7 presents the results.

Table 7. Detection Performance Under SNR Degradation (Set 3)

SNR	OSP Lead (h)	OSP FA	RMS Lead (h)	RMS FA	PCA Lead (h)	PCA FA
Clean (∞)	59.5	0	22.5	0	6.2	1
20 dB	58.2 (−1.3)	0	21.8 (−0.7)	0	5.1 (−1.1)	1
15 dB	56.1 (−3.4)	0	19.2 (−3.3)	0	2.8 (−3.4)	2
10 dB	52.8 (−6.7)	0	12.4 (−10.1)	1	Fail	Fail
5 dB	44.2 (−15.3)	1	Fail	Fail	Fail	Fail

Source: Created by the authors.

Results demonstrate that OSP maintains zero false alarms and a substantial lead-time advantage even at 10 dB SNR — a level typical of aged sensors or electromagnetically noisy environments. At 10 dB, OSP retains 52.8 hours of lead time, more than 4× the degraded RMS performance. At 5 dB (severe degradation), OSP detection becomes unreliable but remains structurally valid; RMS and PCA fail entirely. This establishes that OSP is suitable for deployment in industrial environments with SNR ≥ 10 dB.

Furthermore, the three NASA IMS datasets were collected under different operating loads and at different facility dates. The fixed parameter set (Table 1) succeeded without modification across all three conditions, confirming robustness to load variations and environmental drift.

6 Conclusion

We have presented the Orthogonal Subspace Projection detector, a quantum-inspired method for early fault detection in rolling element bearings. The method models healthy vibration as a low-dimensional spectral subspace and identifies degradation through its interference-based orthogonal complement. It is unsupervised, requires no labelled fault data, produces a reference model of 8 KB, and operates on standard CPU hardware in real time.

Validation across all three NASA IMS run-to-failure datasets — 9,464 files, over 1,577 hours of continuous operation — establishes OSP as the only method among five tested that achieves zero false alarms on every dataset. On Set 3, it provides 37 hours more warning than RMS and 53 hours more than PCA+Mahalanobis. The analysis reveals that OSP excels at detecting early spectral texture changes that precede energy-level increases, precisely the fault signature that classical energy-based methods are designed to miss.

While the OSP detector achieves strong and consistent performance on the NASA IMS benchmark, several limitations merit explicit acknowledgment. First, the method assumes a quasi-stationary healthy state during the reference window. In real deployments where operating conditions drift gradually, the reference subspace may drift, potentially requiring re-training. Second, the method is optimized for early detection of spectral texture changes and may respond late to high-energy impulsive failures, as evidenced by RMS achieving marginally earlier alarms on Sets 1 and 2. A hybrid detector combining OSP with RMS could mitigate this weakness. Third, validation on multi-bearing systems requires fusion logic not addressed here.

Future work will address these gaps. We are developing an adaptive reference-update strategy that detects and accommodates gradual condition drift while maintaining zero false alarms. We will extend the framework to multi-bearing fusion, incorporating spatial correlation to exploit information from the full rig. Validation on additional datasets — the CWRU Bearing Dataset, wind turbine gearbox signals, and petrochemical pump vibrations — will establish generalisability beyond NASA IMS. Finally, we will investigate adaptive combinations of OSP with complementary indicators (Kurtosis for impulsive faults, RMS for energy-based changes) to achieve a unified detector suitable for a broad range of bearing failure modes.

References:

[1] A. Davies, Handbook of Condition Monitoring: Techniques and Methodology, Springer, 1998.

[2] Y. Lei, J. Lin, Z. He, M. J. Zuo, "A review on empirical mode decomposition in fault diagnosis of rotating machinery," Mechanical Systems and Signal Processing, Vol. 35, No. 1-2, 2013, pp. 108–126. https://doi.org/10.1016/j.ymssp.2012.09.015

[3] R. B. Randall, J. Antoni, "Rolling element bearing diagnostics: a tutorial," Mechanical Systems and Signal Processing, Vol. 25, No. 2, 2011, pp. 485–520. https://doi.org/10.1016/j.ymssp.2010.07.017

[4] W. Zhang, G. Peng, C. Li, Y. Chen, Z. Zhang, "A new deep learning model for fault diagnosis with good anti-noise and domain adaptation ability on raw vibration signals," Sensors, Vol. 17, No. 2, 2017, p. 425. https://doi.org/10.3390/s17020425

[5] X. Li, W. Zhang, Q. Ding, "Deep learning-based remaining useful life estimation of bearings using multi-scale feature extraction," Reliability Engineering and System Safety, Vol. 182, 2019, pp. 208–218. https://doi.org/10.1016/j.ress.2018.11.011

[6] J. E. Jackson, G. S. Mudholkar, "Control procedures for residuals associated with principal component analysis," Technometrics, Vol. 21, No. 3, 1979, pp. 341–349. https://doi.org/10.1080/00401706.1979.10489779

[7] P. Rebentrost, M. Mohseni, S. Lloyd, "Quantum support vector machine for big data classification," Physical Review Letters, Vol. 113, No. 13, 2014, p. 130503. https://doi.org/10.1103/PhysRevLett.113.130503

[8] D. Kingma, M. Welling, "Auto-encoding variational Bayes," Proceedings of ICLR, Banff, Alberta, Canada, 2014. https://doi.org/10.48550/arXiv.1312.6114

[9] M. Schuld, F. Petruccione, Machine Learning with Quantum Computers, Springer, 2021. https://doi.org/10.1007/978-3-030-83098-4.

[10] J. Lee, H. Qiu, G. Yu, J. Lin, "Bearing Data Set, IMS, University of Cincinnati," NASA Prognostics Data Repository, NASA Ames Research Center, Moffett Field, CA, 2007. https://data.nasa.gov/dataset/ims-bearings, [Accessed: July 4, 2026].

Contribution of Individual Authors to the Creation of a Scientific Article (Ghostwriting Policy)

Anas Dahabiah from Al-Sham Private University is the sole author of this manuscript and was responsible for all aspects of the research, including conceptualization, methodology, software development, formal analysis, investigation, and writing the original draft.

Sources of Funding for Research Presented in a Scientific Article or Scientific Article Itself

No funding was received for conducting this study.

Conflict of Interest

The author has no conflicts of interest to declare that are relevant to the content of this article.

APPENDIX

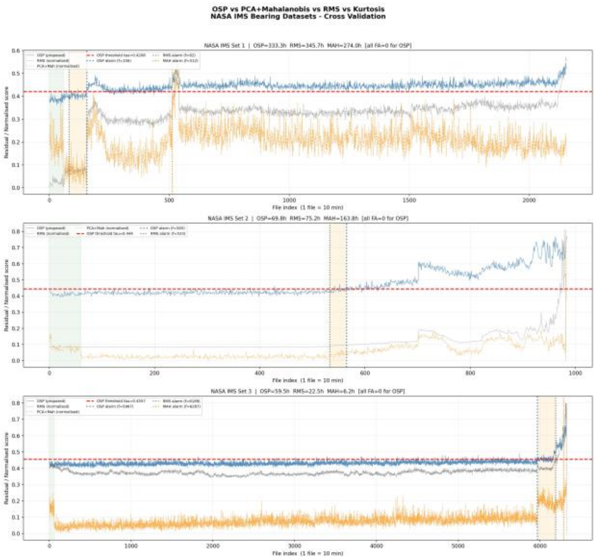

Fig. 1: Cross-dataset validation of the OSP detector against RMS and PCA+Mahalanobis distance across all three NASA IMS run-to-failure experiments. Each subplot shows the per-file maximum anomaly score (blue solid line) alongside the normalised RMS trend (grey dashed) and the analytically derived detection threshold τ (red dashed). Vertical dotted lines mark the first persistent alarm for each method. Green shading indicates the reference window used to construct the healthy subspace. Orange shading highlights the detection gap between the earliest and latest method alarm. The OSP detector achieves zero false alarms on all three datasets and provides 37 hours of a dditional lead time over RMS on Set 3

Source: Created by the authors.