Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A Metric Method for T/R Component Health Measurement for Cyber Defense Based on Sparse Preserving Projection and Hidden Markov Model

Sheng Hong

shenghong@buaa.edu.cn

Beihang University

Xuanqi Wang

Nanchang University

Weiwei Jiang

Beijing University of Posts and Telecommunications

Jingsong Zheng

Nanchang University

Jiacheng Wang

Beihang University

Research Article

Keywords: Health-ability, Time-Smoothness Preserving Sparse Projection, Hidden Markov Model, Kullback--Leibler Divergence, Cyber Defense

Posted Date: May 5th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9414249/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

A Metric Method for T/R Component Health Measurement for Cyber Defense Based on Sparse Preserving Projection and Hidden Markov Model

Sheng Hong[1*], Xuanqi Wang[2], Weiwei Jiang[3], Jingsong Zheng[2], Jiacheng Wang[1]

[1*]School of Cyber Science and Technology, Beihang University, Beijing, 100191, China.
[2]School of Information and Engineering, Nanchang University, Nanchang, Jiangxi, 330031, China.
[3]School of Information and Communication Engineering, Beijing University of Posts and Telecommunications, Beijing, 100876, China.

*Corresponding author(s). E-mail(s): shenghong@buaa.edu.cn;
Contributing authors: 15894886025@163.com; jww@bupt.edu.cn;
zjs04929@163.com; wjc1321@163.com;

Abstract

Cyber attacks on communication and radar systems increasingly exploit physical-layer vulnerabilities: targeted interference, power-surge injections, and coordinated cyber-physical intrusions can accelerate T/R module degradation, silently eroding beam-forming accuracy, jamming resistance, and encrypted-link integrity long before any network-layer anomaly is detected. Defending against such threats requires continuous health monitoring of T/R modules for effective cyber defense; however, existing data-driven methods fail to produce stable, monotone health indicators from multivariate RF sensor data, as operational load variations obscure genuine degradation trends and ground-truth failure labels are rarely available. We propose a framework combining Time-Smoothness Preserving Sparse Projection (TSP-SPP) with a Continuous Hidden Markov Model (CHMM) and Symmetric Kullback–Leibler divergence to construct a normalised Health-ability (HA) index in $[\mathbf{0}, \mathbf{1}]$. TSP-SPP encodes temporal continuity into the projection geometry, suppressing load-driven transients while preserving genuine degradation trends; Sym-KL captures both early covariance broadening and late-stage mean shifts that point-level metrics miss. The resulting HA serves as an early-warning signal, enabling operators to detect physical-layer compromise

and initiate timely countermeasures before mission-critical capabilities are lost. Experiments on a full life-cycle T/R module dataset show a Monotonicity score of 0.542 and irregularity penalty of 0.032, outperforming PCA, LPP, and SPP baselines on both metrics.

Keywords: Health-ability, Time-Smoothness Preserving Sparse Projection, Hidden Markov Model, Kullback–Leibler Divergence, Cyber Defense

1 Introduction

Cyber attacks on modern communication and radar systems have increasingly extended beyond the network and software layers to target the physical infrastructure directly. Targeted electromagnetic interference, deliberate power-surge injections, and coordinated cyber-physical intrusions can stress the RF front-end circuitry of active phased-array antennas, accelerating thermal fatigue and electromigration in the power amplifier stage of T/R modules and progressively degrading gain, output power, and phase accuracy across the entire array (Hong et al. 2022, 2025). Because these effects accumulate gradually and remain invisible to conventional network-layer monitoring, an adversary can exploit physical-layer degradation as a persistent, hard-to-detect attack surface: a platform whose T/R modules are silently compromised loses the ability to reject electronic jamming, detect low-observable threats, or maintain encrypted communication links—precisely the capabilities that cyber defense depends on. Continuous, real-time health monitoring of T/R modules is therefore a cyber defense imperative, providing the situational awareness needed to detect and respond to physical-layer compromise before mission-critical capabilities are lost (Hong et al. 2025; Hu et al. 2022; Zhu et al. 2025).

Existing approaches span Physics-of-Failure (PoF) models, supervised machine-learning classifiers, and unsupervised statistical frameworks; a detailed review is given in Section 2. Each family carries limitations that prevent direct deployment to large-scale phased-array T/R module monitoring: PoF models require precise material parameters unavailable in situ, supervised classifiers depend on labelled failure data that are chronically scarce, and current unsupervised methods lack a principled, continuous health score that is stable across operational load variations.

Two difficulties remain unresolved by existing methods. Conventional manifold learning algorithms treat each observation as an isolated point in feature space, ignoring the slow temporal continuity that physical wear imposes on a continuously operating module. A T/R module Power Amplifier (PA) running at full transmit power for hours accumulates junction temperature and trap charge gradually; its output power and gain therefore drift on a timescale of hours to days, not sample to sample. Operationally driven load fluctuations nonetheless superimpose fast transients—switching bursts, beam-steering events, thermal cycling between high-power and standby modes—that a temporally unaware projection conflates with structural degradation, fragmenting what should be a smooth health trajectory into a noisy, non-monotone signal.

The second difficulty lies in health quantification. Standard HMMs classify discrete states; translating their outputs into a continuous, normalised health score is non-trivial. The raw log-likelihood is unnormalised, load-dependent, and impossible to compare across modules with different mission histories. State-posterior probabilities do encode soft uncertainty, but a module whose covariance has broadened due to early PA thermal fatigue may still be assigned high probability to the healthy state if its mean has not yet shifted—precisely the early-degradation regime where timely intervention matters most. A health metric that only tracks the mean misses the variance inflation that constitutes the earliest detectable symptom of incipient wear.

Unlike conventional reliability metrics—which characterise the probability of non-failure over a design-time statistical ensemble—this paper introduces *Health-ability* (HA): a continuously quantifiable, instance-specific, time-varying measure of the degree to which an individual T/R module's real-time operational state distribution conforms to its healthy-phase baseline, normalised to $[0, 1]$ over the full degradation life cycle. An HA signal that degrades steadily toward zero provides an early-warning indicator for cyber defense operators: it signals that the physical-layer integrity of a T/R module is being compromised, whether by natural wear, adversarial interference, or a combination of both.

The framework developed here addresses both difficulties through three tightly coupled design choices.

1. **TSP-SPP** (Time-Smoothness Preserving Sparse Projection) augments standard SPP with a temporal Laplacian penalty that couples adjacent observations during projection learning. Load-driven transients, which are short-lived and uncorrelated with structural wear, are suppressed geometrically; the resulting trajectory tracks the slow monotone drift of genuine degradation without post-hoc smoothing.
2. **CHMM baseline** (K-state Continuous HMM trained exclusively on confirmed healthy data) partitions the healthy operating space into K sub-regimes. At inference time, the Viterbi algorithm selects the sub-regime that best matches the current window, so the health distance is always computed against an operationally appropriate reference rather than a global mean that would misclassify load transitions as degradation.
3. **Symmetric KL divergence on sliding-window distributions** captures both mean shifts and covariance broadening in a single scalar, normalised to produce $\mathrm{HA}_t \in [0, 1]$. Covariance broadening—the earliest measurable signature of PA thermal fatigue—is invisible to point-level metrics such as Euclidean or Mahalanobis distance; Sym-KL detects it without requiring labelled failure examples.

A typical T/R module integrates a PA, Low Noise Amplifier (LNA), Phase Shifter, and control circuits into a compact, highly integrated assembly. Among these sub-components, the PA operates under the highest power density and thermal stress, rendering it the most security-critical element of the assembly under cyber-physical attack scenarios. Common failure mechanisms—such as electromigration and thermal fatigue of solder joints—do not manifest as immediate functional failures; instead, they progress as a continuous degradation process, manifested primarily as a reduction in output power and a decrease in channel gain. During operation, the module is

continuously characterised by a set of RF performance parameters—including transmit/receive channel gain, output power, Voltage Standing Wave Ratio (VSWR), noise figure, spurious and harmonic levels, 1 dB compression point, circulator isolation, and switching timing—yielding a high-dimensional feature vector $\boldsymbol{x}_t \in \mathbb{R}^D$ at each time step t.

The remainder of this paper is organised as follows. Section 2 surveys related work on RF component health monitoring, manifold-based feature extraction, and HMM health scoring. Section 3 details the complete proposed framework, covering TSP-SPP feature extraction, CHMM health state modelling, and Sym-KL-based Healthability construction. Section 4 presents experimental results and comparative analysis. Section 5 draws conclusions.

2 Related Work

2.1 Health Monitoring of RF and Power Electronics Components

Early health monitoring approaches for RF front-end components relied heavily on Physics-of-Failure (PoF) models that map physical stress quantities—temperature, voltage, and current—to wear-out mechanisms such as electromigration and thermal fatigue (Hult et al. 2024b, 2025, 2024a). Modern GaN MMIC integration has produced complete single-chip T/R front ends operating at millimetre-wave frequencies (Sonnenberg et al. 2024), making the compact, densely integrated assemblies increasingly difficult to monitor through individual component models alone. While PoF models provide physical interpretability, they require accurate material parameters and detailed stress profiles that are rarely available in fielded systems operating under variable mission profiles. Self-monitoring circuits embedded directly in power amplifiers (Sun and Liao 2024), enhanced thermal management for cyclic T/R module operation (He et al. 2025), and temperature-based compensation techniques for GaN and CMOS devices (Mateo et al. 2025; Altet et al. 2023) address the real-time tracking of gain drift and thermal stress, yet they remain component-specific and do not generalise to the multi-channel system level. Our approach instead operates on observable RF performance parameters directly, without requiring access to internal physical states. At the system level, a large phased array presents a monitoring challenge that is qualitatively different from single-component health tracking: each T/R module records 13 heterogeneous RF parameters simultaneously, channels are sampled at different rates, and the sensor stream is affected by both module-internal wear and array-level operational scheduling. Extracting a per-module health index that is robust to these sources of variation—without dedicated on-chip sensing hardware or knowledge of device-internal stress variables—requires a data-driven approach that can operate on the observable parameter streams alone.

2.2 Data-Driven Feature Extraction for Degradation Trajectories

Linear methods such as Principal Component Analysis (PCA) and its kernel extensions remain the most widely deployed preprocessing steps for bearing and machinery health assessment because of their low computational overhead. Graph-Laplacian methods such as LPP construct a neighbourhood graph and preserve local manifold geometry in the projected space, outperforming PCA when the degradation data lies on a curved manifold (Hou et al. 2025); combining LPP with adaptive sparse representations has improved industrial fault detection (Wu and Gu 2021), and integrating discriminative LPP with sparse autoencoders has further enhanced fault diagnosis accuracy (He et al. 2021). Sparse-reconstruction-based manifold learning (SPP (Yan et al. 2023)) further improves locality preservation by using ℓ_1-regularised reconstruction weights, producing projections that are robust to outlier observations and non-uniform sample densities; joint sparse 2-D projection with local manifold structure preservation extends these ideas to robust fuzzy clustering (Wu and Gong 2026), and sparse graph-preserving projections on Riemannian manifolds have demonstrated reliable feature extraction under diverse non-linear data geometries (Liu et al. 2021). However, all of these methods treat each observation as an independent point, ignoring the temporal continuity imposed by physical degradation processes. Deep learning approaches—including stacked autoencoders, variational autoencoders, and temporal convolutional networks—can capture non-linear temporal dependencies, but require large labelled datasets for training and offer limited interpretability, making them less suitable for T/R module applications where labelled failure data is scarce (Wan et al. 2026). For T/R module monitoring specifically, the temporal structure is not a minor nuisance but the dominant signal: the module runs continuously for months, and its health state on any given day is strongly correlated with the previous day's state. A projection that breaks this correlation—treating a mid-life observation as if it had no relationship to the observations before and after—makes the health trajectory look jagged and non-monotone even when the underlying degradation is smooth. TSP-SPP directly encodes the slow-time structure into the projection geometry by adding a temporal Laplacian term to the SPP objective, preserving manifold structure while enforcing trajectory smoothness without labelled failure data.

2.3 HMM-Based Health State Modelling

Hidden Markov Models have long been applied to rotating machinery fault diagnosis because they naturally model temporal state transitions in a principled probabilistic framework (Liu and Zhu 2021). Theoretical advances have extended HMMs to multivariate longitudinal data with parsimony constraints (Neal and McNicholas 2026), continuous-time Bayesian clustering formulations (Luo et al. 2023), ordinal time series (Weiß and Swidan 2025), copula-based emission structures with unknown state counts (Liu et al. 2025), and quantile graphical models (Foroni et al. 2025); non-parametric clustering risk bounds provide statistical guarantees for model selection (Gassiat et al. 2025). Unsupervised HMM variants are particularly attractive for

equipment health monitoring, since they require only healthy-phase data for training and can detect novel degradation patterns not seen during training (Asgarinejad et al. 2023); hybrid convolutional neural network–CHMM architectures demonstrate the value of continuous-density emissions for sequential pattern detection (Hu et al. 2025), and HMM-based trajectory analysis has been applied to anomaly detection in cyber-physical systems (Ou et al. 2026). Physics-informed HMMs that incorporate domain-specific feature transformations have recently demonstrated robust fault detection under variable operating conditions (Wang et al. 2026); generative HMM designs have further enabled adaptable structural modelling (Wang et al. 2025). Health score construction from HMM outputs remains an open problem: the raw log-likelihood is unnormalised, while state-posterior probabilities do not capture covariance changes in the early degradation phase. Distribution-level divergence measures such as KL divergence address this gap by comparing the full observation distribution against a healthy reference, capturing both mean shifts and covariance broadening (Yan et al. 2026); symmetric KL divergence has proven effective as a regularisation objective for encoder-based classifiers (Mondal et al. 2026), for temporal hesitation analysis (Lee and Lee 2026), for filter design against injection attacks (Tosun et al. 2026), and for multimodal state estimation via ensemble Kalman filtering (Li et al. 2019). Graph-based dynamic health indicators have similarly demonstrated that window-level statistical features outperform point-level metrics for rotating machinery (Hou et al. 2025), motivating our sliding-window Sym-KL construction. The missing piece across all of these HMM-based health scoring approaches is a distance measure that is both normalised and sensitive to early-stage distributional change. Raw log-likelihood depends on sequence length and load level. State posteriors track the mean trajectory but are blind to covariance broadening. Sym-KL between a sliding-window local Gaussian and the regime-matched CHMM reference closes this gap: it responds to covariance inflation before any mean shift is detectable, delivers a $[0, 1]$-normalised score that is load-invariant, and requires no labelled failure data—three properties that previous approaches address individually but rarely in combination for T/R module health monitoring.

3 Methodology

3.1 Framework Overview

Fig. 1 illustrates the overall architecture of the proposed framework, which comprises three sequential stages.

In *Stage 1*, sensor observations are first aligned to a common time base and missing values imputed (detailed in Section 3.3); the resulting clean feature vectors $\boldsymbol{x}_t \in \mathbb{R}^D$ are then fed into TSP-SPP. TSP-SPP jointly optimises a sparse reconstruction loss J_{spp} and a temporal Laplacian penalty J_{temp}, projecting the data into a d-dimensional feature vector $\boldsymbol{y}_t \in \mathbb{R}^d$ whose trajectory is smooth by construction—consecutive frames are penalised for drifting apart, so transient load spikes leave no lasting imprint on the projection. In *Stage 2*, the first $\lfloor \eta N \rfloor$ frames ($\eta = 0.15$, i.e., the confirmed healthy period), are used to fit a K-state Continuous HMM Θ_{healthy} via the Baum–Welch algorithm. Each hidden state captures a distinct sub-regime

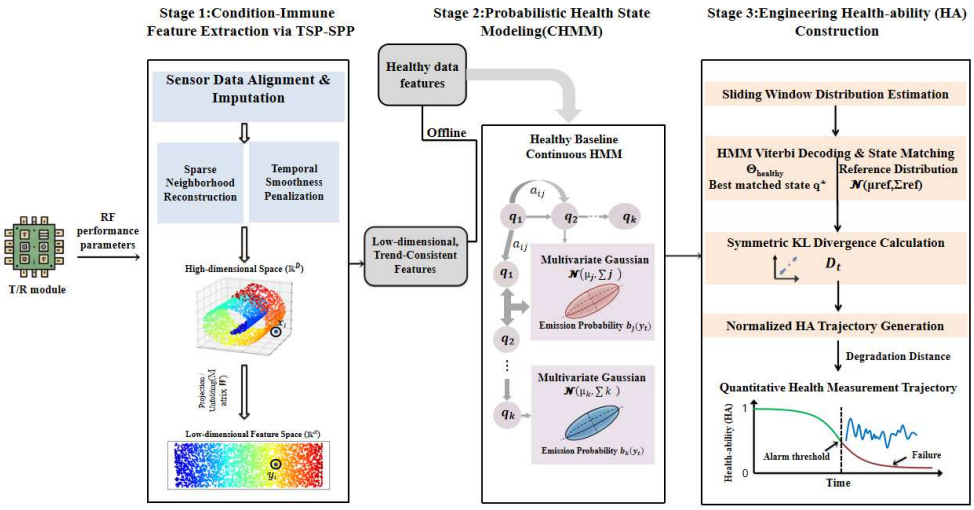


Fig. 1 Overall framework of the proposed T/R module health measurement system.

of normal operation (e.g., low-power standby versus full-power transmit); together the K Gaussians tile the healthy portion of the projected space and form the reference against which future observations are measured. In *Stage 3*, at each new time step, a sliding window $\mathcal{W}_t$ of length L provides enough samples to estimate a local Gaussian $P_{\text{curr}} = \mathcal{N}(\boldsymbol{\mu}_{\text{curr}}, \boldsymbol{\Sigma}_{\text{curr}})$. The Viterbi algorithm then decodes which healthy sub-regime best matches the current window, supplying a regime-appropriate reference $P_{\text{ref}} = \mathcal{N}(\boldsymbol{\mu}_{\text{ref}}, \boldsymbol{\Sigma}_{\text{ref}})$. Sym-KL divergence D_t between P_{curr} and P_{ref} is min–max normalised over the full life cycle and inverted to produce $\text{HA}_t \in [0, 1]$, which reads near unity during healthy operation and falls to zero as the module approaches failure. The three stages are detailed in Sections 3.5–3.6.2.

3.2 Manifold Structure of Degradation Data

Let $\boldsymbol{X} = [\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_N] \in \mathbb{R}^{D\times N}$ denote the time-series monitoring data of the T/R module, where D is the sensor feature dimensionality and N the number of time steps. The goal of this subsection is to learn a projection matrix $\boldsymbol{W} \in \mathbb{R}^{D\times d}$ $(d \ll D)$ that maps each observation to a low-dimensional feature $\boldsymbol{y}_t = \boldsymbol{W}^T\boldsymbol{x}_t \in \mathbb{R}^d$ preserving the degradation-relevant manifold structure; the resulting features are subsequently used in Section 3.6 to construct a normalised Health-ability $\text{HA}_t \in [0, 1]$.

A T/R module under continuous operation produces a 13-dimensional observation vector at every measurement epoch. Although that vector lives nominally in $\mathbb{R}^{13}$, the actual degrees of freedom governing health evolution are far fewer: PA junction temperature, drain current, and oxide trap density together account for the dominant share of performance drift, so the data trace a low-dimensional curve embedded within the ambient sensor space. During healthy operation, this curve meanders within a compact cloud near the origin of the feature space; once thermal fatigue initiates,

the cloud migrates along a well-defined direction that reflects accumulated material damage.

Why linear projections fall short. PCA finds the subspace of maximum variance—a criterion that is blind to the local curvature of the degradation path. Where the healthy-phase cloud and the early-degradation segment overlap in variance, PCA collapses them into the same projected neighbourhood, erasing the fine-grained contrast a health index depends on. The same observation applies to any globally optimal linear transform: globally preserving variance does not imply locally preserving health-state separability.

Sparse Preserving Projection as a manifold-faithful basis. The key insight behind SPP (Yan et al. 2023) is that the local geometry of a manifold is captured by which points reconstruct each other, not by which directions explain the most global variance. An SPP weight $s_{ij} > 0$ says that sample j participates in the reconstruction of sample i as a near-neighbour on the manifold; ℓ_1 regularisation ensures that only the geometrically closest points receive non-zero weight, making the neighbourhood topology sparse and noise-resistant. Projecting to preserve these weights keeps nearby points nearby— the defining property of a good manifold embedding. The approach transfers naturally across RF and industrial contexts: radar echo classification from range profiles (Yang et al. 2024), hyperspectral image classification (Li et al. 2026), and fault localisation in rotating machinery (Chen et al. 2023) all rely on the same sparsity-driven locality property.

3.3 Data Preprocessing

Prior to feature extraction, a two-stage preprocessing pipeline is applied to the simulated multi-channel monitoring data of the T/R module. T/R modules are characterised by multiple RF performance parameters—such as channel gain, output power, VSWR, noise figure, and switching timing—sampled at heterogeneous rates; the resulting multivariate time series frequently exhibit temporal misalignment between channels. In the first stage, *timestamp alignment* synchronises all sensor channels to a unified physical time base, ensuring that the observation vector $\boldsymbol{x}_t$ represents a coherent snapshot of the complete module state at each instant t. In the second stage, data gaps arising from transient sensor instability are repaired via a machine-learning-based imputation procedure, yielding the complete, temporally aligned dataset $\boldsymbol{X} \in \mathbb{R}^{D\times N}$ on which all subsequent analysis is conducted.

3.4 Sparse Preserving Projection (SPP)

SPP (Yan et al. 2023) learns a projection matrix $\boldsymbol{W} \in \mathbb{R}^{D\times d}$ by requiring that the sparse neighbourhood relationships among the original data points are preserved after the dimensionality reduction step.

Sparse Reconstruction. The central idea is that any sample $\boldsymbol{x}_i$ can be well approximated by a weighted combination of a small number of its neighbours. The weight vector $\boldsymbol{s}_i$ is obtained by solving the ℓ_1-regularised problem:

$$\min_{\boldsymbol{s}_i} \ \|\boldsymbol{x}_i - \boldsymbol{X}\boldsymbol{s}_i\|^2 + \beta\|\boldsymbol{s}_i\|_1, \quad \text{s.t. } \mathbf{1}^T\boldsymbol{s}_i = 1,\ s_{ii} = 0 \tag{1}$$

where $\beta > 0$ is the sparsity regularisation parameter. In contrast to k-NN methods, which rely on fixed Euclidean distance thresholds, the ℓ_1 sparsity constraint adaptively identifies the most representative neighbourhood structure, inherently suppressing stochastic sensor noise and measurement fluctuations prevalent in complex RF environments.

Geometric Preservation. The projection matrix $\boldsymbol{W}$ is subsequently determined by minimising the total reconstruction error in the projected space:

$$J_{\text{spp}}(\boldsymbol{W}) = \sum_{i=1}^{N} \left\| \boldsymbol{y}_i - \sum_{j=1}^{N} s_{ij} \boldsymbol{y}_j \right\|^2 = \text{Tr}\big(\boldsymbol{W}^T \boldsymbol{X} \boldsymbol{M}_\beta \boldsymbol{X}^T \boldsymbol{W}\big) \tag{2}$$

where $\boldsymbol{M}_\beta = \boldsymbol{I} - \boldsymbol{S} - \boldsymbol{S}^T + \boldsymbol{S}^T \boldsymbol{S}$ is the alignment matrix parametrised by the sparse weights $\boldsymbol{S} = [\boldsymbol{s}_1, \ldots, \boldsymbol{s}_N]$ (which depend on β through (1)). Minimising J_{spp} ensures that the sparse neighbourhood relationships established in the original space are preserved under projection, thereby unfolding the nonlinear degradation manifold into a geometrically faithful low-dimensional representation.

3.5 Time-Smoothness Preserving Sparse Projection (TSP-SPP)

Although SPP faithfully preserves the local manifold geometry of the degradation data, it treats each observation as an independent sample, disregarding the temporal ordering of the data stream. This independence assumption is fundamentally at odds with the physical nature of gradual wear degradation: in a continuously operating T/R module, the health state evolves slowly and smoothly, so consecutive observations are governed by the same underlying physical process and should remain proximate in the feature space. In practice, however, operationally driven load fluctuations superimpose abrupt short-term signal transients upon this slow drift; standard SPP conflates these instantaneous operational variations with genuine structural aging, producing fragmented, volatile feature trajectories that impede reliable health assessment.

To encode physical operational continuity directly into the learned projection, the proposed TSP-SPP extends the SPP objective with a temporal Laplacian regularisation term:

$$J_{\text{tsp}}(\boldsymbol{W}) = J_{\text{spp}}(\boldsymbol{W}) + \lambda\, J_{\text{temp}}(\boldsymbol{W}) \tag{3}$$

where $\lambda \geq 0$ is the regularisation strength detailed below.

Temporal Smoothness Term. The temporal regulariser penalises abrupt feature transitions between adjacent time steps, enforcing that the projected trajectory evolves continuously in accordance with the monotonic nature of physical degradation:

$$J_{\text{temp}}(\boldsymbol{W}) = \sum_{t=1}^{N-1} \|\boldsymbol{y}_t - \boldsymbol{y}_{t+1}\|^2 = \text{Tr}\big(\boldsymbol{W}^T \boldsymbol{X} \boldsymbol{L}_{\text{adj}} \boldsymbol{X}^T \boldsymbol{W}\big) \tag{4}$$

where $\boldsymbol{L}_{\text{adj}} = \boldsymbol{D}_{\text{adj}} - \boldsymbol{A}_{\text{adj}}$ is the temporal Laplacian matrix constructed from the chain adjacency graph of the time series: $\boldsymbol{A}_{\text{adj}}$ has $(\boldsymbol{A}_{\text{adj}})_{t,t+1} = (\boldsymbol{A}_{\text{adj}})_{t+1,t} = 1$ for

$t = 1, \ldots, N-1$ and zero elsewhere, and $\boldsymbol{D}_{\text{adj}}$ is the corresponding degree matrix. This yields $(\boldsymbol{L}_{\text{adj}})_{tt} = 1$ or 2 (boundary or interior), $(\boldsymbol{L}_{\text{adj}})_{t,t\pm1} = -1$, directly encoding the physical continuity of adjacent time steps. The regularisation strength $\lambda \geq 0$ governs the balance between manifold structure fidelity and temporal trajectory smoothness: setting $\lambda = 0$ recovers standard SPP, while increasing λ progressively suppresses non-physical feature jumps at the expense of strict manifold fidelity. TSP-SPP operates on the *short-time* scale of individual window transitions— suppressing sensor noise bursts and operationally driven transients that unfold within seconds to minutes—whereas the CHMM models *long-time* operational regime changes (e.g., thermal equilibrium shifts or sustained load-level changes) that evolve over hours to days. The two mechanisms therefore act on complementary time scales and reinforce rather than conflict with each other.

Solution. Substituting (2) and (4) into (3) and imposing the standard whitening constraint $\boldsymbol{W}^T \boldsymbol{X} \boldsymbol{X}^T \boldsymbol{W} = \boldsymbol{I}$, the optimisation reduces to a generalised eigenvalue problem:

$$\boldsymbol{X}(\boldsymbol{M}_\beta + \lambda \boldsymbol{L}_{\text{adj}})\boldsymbol{X}^T \boldsymbol{w} = \gamma\, \boldsymbol{X}\boldsymbol{X}^T \boldsymbol{w} \tag{5}$$

Taking the d eigenvectors with the smallest non-zero eigenvalues as columns of $\boldsymbol{W}$ satisfies both objectives simultaneously: sparse reconstruction error is minimised and abrupt temporal transitions are penalised. The resulting features $\{\boldsymbol{y}_t\}$ are stable under sensor noise through the sparse neighbourhood structure inherited from SPP, and follow a smooth trajectory that tracks genuine degradation rather than load-driven fluctuations.

3.6 Health State Modelling and Health-ability Construction

Given the low-dimensional feature sequence $\boldsymbol{Y} = [\boldsymbol{y}_1, \ldots, \boldsymbol{y}_N] \in \mathbb{R}^{d \times N}$ produced by TSP-SPP, this subsection constructs the normalised Health-ability:

$$\text{HA}_t = \mathcal{F}_{\text{HMM}}\big(\mathcal{W}_t \,\big|\, \Theta_{\text{healthy}}\big) \in [0, 1] \tag{6}$$

where $\mathcal{W}_t = [\boldsymbol{y}_{t-L+1}, \ldots, \boldsymbol{y}_t]$ is the sliding window of length L ending at time t, Θ_{healthy} is the CHMM trained exclusively on healthy-phase data; $\text{HA}_t \approx 1$ signifies a healthy state and $\text{HA}_t \to 0$ signifies impending failure. The following paragraphs detail the two components of $\mathcal{F}_{\text{HMM}}$: the Continuous HMM that establishes a probabilistic healthy baseline, and the Sym-KL strategy that converts model outputs into a stable, normalised index.

3.6.1 Continuous HMM for Health State Modelling

Following dimensionality reduction, the feature sequence $\boldsymbol{Y} = [\boldsymbol{y}_1, \ldots, \boldsymbol{y}_N] \in \mathbb{R}^{d \times N}$ produced by TSP-SPP is passed to the CHMM. A CHMM (here "continuous" refers to the Gaussian emission densities, distinguishing it from discrete-observation HMMs) consists of a hidden state space $Q = \{q_1, \ldots, q_K\}$ and a parameter triplet $\Theta = (\boldsymbol{\pi}, \boldsymbol{A}, B)$. The hidden state space $Q = \{q_1, \ldots, q_K\}$ represents the intrinsic operational regimes within the healthy lifespan—such as thermal equilibrium stages or varying load levels—rather than discrete failure modes; this design is central to constructing a healthy baseline model that absorbs legitimate intra-class variation without

conflating it with structural degradation. The initial state probability vector $\boldsymbol{\pi} \in \mathbb{R}^K$, with $\pi_i = P(q_1 = i)$, specifies the prior distribution over states at the onset of observation. The state transition matrix $\boldsymbol{A} \in \mathbb{R}^{K \times K}$, with $a_{ij} = P(q_{t+1} = j \,|\, q_t = i)$, encodes the Markovian dynamics of regime transitions, capturing the temporal continuity of module operation. Since the observation $\boldsymbol{y}_t$ is a continuous-valued vector, the emission probability density of state q_j is modelled as a multivariate Gaussian:

$$b_j(\boldsymbol{y}_t) = \mathcal{N}(\boldsymbol{y}_t; \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) = \frac{\exp\left(-\frac{1}{2}\boldsymbol{\delta}_j^T \boldsymbol{\Sigma}_j^{-1} \boldsymbol{\delta}_j\right)}{\sqrt{(2\pi)^d |\boldsymbol{\Sigma}_j|}} \tag{7}$$

where $\boldsymbol{\delta}_j \triangleq \boldsymbol{y}_t - \boldsymbol{\mu}_j$. A single Gaussian per state achieves a favourable balance between modelling precision and computational tractability, a choice justified by the high discriminative power of the TSP-SPP-compressed low-dimensional features.

Parameter Estimation. Θ is fitted by the Baum–Welch (EM) algorithm using only frames from the confirmed healthy period $\boldsymbol{Y}[1 : \lfloor \eta N \rfloor]$. The algorithm alternates between an E-step that infers state occupancy probabilities via forward–backward recursion and an M-step that updates $\boldsymbol{\pi}$, $\boldsymbol{A}$, and each emission pair $(\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)$ in closed form. Restricting training to the healthy window is intentional: the resulting Θ_{healthy} forms a compact statistical atlas of normal behaviour, including sub-regime transitions such as thermal warm-up, load cycling, and periodic calibration pulses. When a test window deviates from all healthy sub-regimes, the Sym-KL distance against the nearest Gaussian reference grows large, and the normalisation step translates that growth into a low HA reading. The choice of $K = 2$ states for this dataset reflects the structure of the projected space: the eigenvalue spectrum (Section 4) shows two dominant directions, and in practice the healthy trajectory partitions into a low-power standby cluster and a high-power transmit cluster with well-separated centroids—two states are sufficient to cover this variation without over-parameterising the model given the $\sim$150 available training frames. When a later observation window no longer fits this model well, the resulting distributional gap is converted into the health score.

3.6.2 Continuous Health-ability via Sliding Window and Symmetric KL Divergence

Although the trained Θ_{healthy} encodes the full probabilistic structure of the healthy operational space, using the raw log-likelihood $\log P(\boldsymbol{Y} \,|\, \Theta_{\text{healthy}})$ as a health metric is problematic: this quantity is unnormalised, varies with operational load, and cannot be interpreted consistently across different T/R modules or mission profiles. The proposed approach instead combines two components. The CHMM provides a regime-aware reference: the Viterbi algorithm decodes the healthy operational regime most consistent with the current observation window, supplying a context-specific reference distribution rather than a single global mean. Sym-KL then measures how far the current local distribution has deviated from that regime-matched baseline. The CHMM absorbs legitimate within-class variation—such as load changes and thermal cycling that alter the operating point without causing structural damage—while Sym-KL isolates the degradation-driven distributional shift against the normalised reference. The resulting HA remains stable near unity during normal operation regardless of

operational regime, and responds sensitively to both covariance and mean changes as structural degradation progresses.

Step 1: Local Distribution Estimation. A single observation vector $\boldsymbol{y}_t$ is statistically insufficient to characterise the health state: instantaneous measurements are dominated by sensor noise and short-lived operational transients, rendering point-wise comparison against the healthy baseline inherently unreliable. To transition from point-wise evaluation to statistically robust distribution-level comparison, a sliding window of length L is applied: $\mathcal{W}_t = [\boldsymbol{y}_{t-L+1}, \ldots, \boldsymbol{y}_t]$. A local state distribution $P_{\text{curr}} = \mathcal{N}(\boldsymbol{\mu}_{\text{curr}}, \boldsymbol{\Sigma}_{\text{curr}})$ is estimated over $\mathcal{W}_t$ via maximum likelihood. Since the projection dimensionality is $d = 2$, the resulting 2×2 covariance matrix $\boldsymbol{\Sigma}_{\text{curr}}$ is well-conditioned for any window length $L \gg d$; a small ridge term $\varepsilon \boldsymbol{I}$ ($\varepsilon = 10^{-6}$) is added in practice to guarantee numerical stability against collinear feature directions. This moving-window strategy simultaneously smooths transient sensor noise and captures the evolving mean shift and covariance expansion that characterise progressive physical degradation (Razmarashooli et al. 2025).

Step 2: Degradation Quantification. The Viterbi algorithm is applied to the current window to decode the most likely hidden-state sequence within Θ_{healthy}; the modal state q^* of that sequence provides the regime-matched reference: $P_{\text{ref}} = \mathcal{N}(\boldsymbol{\mu}_{\text{ref}}, \boldsymbol{\Sigma}_{\text{ref}})$. This regime-matching mechanism ensures that the health distance is always computed against the operationally appropriate healthy state, rather than a single global baseline that cannot account for legitimate load-driven regime variation. In the advanced-degradation and failure zones ($\text{HA}_t \approx 0$), the Viterbi-decoded reference may differ from its healthy-phase assignment; however, since D_t is already substantially elevated in this region the clinical determination—impending failure—is unaffected. The primary benefit of regime matching lies in the early-degradation phase, where correct reference selection directly stabilises the HA estimate against legitimate operational regime variation.

The degradation distance is quantified via Sym-KL, rather than conventional point-level metrics such as Euclidean or Mahalanobis distance. Point-level metrics respond exclusively to mean-level shifts and are therefore insensitive to the output-power variance inflation that characterises early-stage PA thermal fatigue—the primary incipient-degradation signature of T/R modules. In contrast, Sym-KL operates on the full distributional parameters of both P_{curr} and P_{ref}:

$$D_t = D_{\text{KL}}(P_{\text{curr}} \| P_{\text{ref}}) + D_{\text{KL}}(P_{\text{ref}} \| P_{\text{curr}}) \tag{8}$$

For two multivariate Gaussians, this yields the closed-form expression:

$$\begin{aligned} D_t = \frac{1}{2} \operatorname{Tr}\big(\boldsymbol{\Sigma}_{\text{ref}}^{-1}\boldsymbol{\Sigma}_{\text{curr}} + \boldsymbol{\Sigma}_{\text{curr}}^{-1}\boldsymbol{\Sigma}_{\text{ref}}\big) \\ + \frac{1}{2} \Delta\boldsymbol{\mu}^T(\boldsymbol{\Sigma}_{\text{ref}}^{-1} + \boldsymbol{\Sigma}_{\text{curr}}^{-1})\Delta\boldsymbol{\mu} - d \end{aligned} \tag{9}$$

where $\Delta\boldsymbol{\mu} = \boldsymbol{\mu}_{\text{ref}} - \boldsymbol{\mu}_{\text{curr}}$ and d is the feature dimensionality. The first trace term captures changes in covariance structure, providing sensitivity to output-power variance inflation that typifies incipient PA thermal fatigue; the second quadratic term captures

the mean shift that dominates at advanced degradation stages. This dual sensitivity enables early detection of gradual degradation that would remain invisible to purely mean-tracking metrics. When the component is healthy, $P_{\text{curr}} \approx P_{\text{ref}}$ and $D_t \to 0$; as degradation progresses, D_t increases monotonically.

Health-ability. The normalised Health-ability is defined as:

$$\text{HA}_t = 1 - \frac{D_t - D_{\min}}{D_{\max} - D_{\min}} \tag{10}$$

where $D_{\min} = \min_{t \in \mathcal{T}_{\text{train}}} D_t$ is the minimum KL distance observed over the healthy training windows $\mathcal{T}_{\text{train}}$, and $D_{\max} = \max_{t \in \mathcal{T}_{\text{full}}} D_t$ is the maximum over the full life-cycle trajectory $\mathcal{T}_{\text{full}}$. By this definition, $\text{HA}_t \approx 1$ in the healthy phase and $\text{HA}_t \to 0$ at failure, providing a physically interpretable, normalised HA suitable for direct use in maintenance decision-making.

4 Experimental Results and Analysis

4.1 Experimental Setup

Dataset. The proposed framework is evaluated on a full life-cycle dataset of a T/R module. The dataset originates from 20 days of real-world measurements collected under the monitoring framework of (Tao and Song 2023); a Deep Adversarial Generative Network (DAGAN) augmentation pipeline is then applied to extend the trajectory to a complete one-year degradation life cycle, from normal operation through to failure. The dataset comprises 961 sliding-window observation frames spanning the full health trajectory from normal operation to failure. Each window yields a 13-dimensional feature vector comprising RF performance parameters: transmit channel output power, in-band power flatness, gain, VSWR, spurious, and harmonics; receive channel gain, VSWR, noise figure, and 1 dB compression point; circulator isolation; and T/R switching and phase-shift settling times. The degradation is characterised by a gradual gain drop in the PA stage, ultimately exceeding the -3 dB failure threshold at the end of the life cycle. The first $\sim$15 % of raw observation frames constitute the healthy training set for the CHMM baseline model Θ_{healthy}.

Table 1 summarises the 13 monitored parameters, grouped by functional subsystem.

Implementation Parameters. The SPP sparse reconstruction uses ℓ_1 regularisation coefficient $\beta = 0.03$ (corresponding to β in (1)). The sliding window for local distribution estimation has length $L = 40$ with a unit evaluation stride (yielding $N_w = 961$ windows); a subsampling stride of $s = 5$ is applied to the healthy-phase windows during CHMM training to reduce redundancy, the CHMM uses $K = 2$ hidden states, $\lambda = 1$ for all TSP-SPP runs, and the healthy-phase fraction $\eta = 0.15$ (i.e., the first 15 % of raw observation frames are used to train Θ_{healthy}). These settings are fixed across all methods and ablation cases unless stated otherwise.

Baselines. The proposed method (TSP-SPP + CHMM + Sym-KL) is evaluated against three feature-extraction baselines, all sharing an identical CHMM + Sym-KL backend to ensure a controlled comparison of the feature extraction stage in isolation.

Table 1 T/R Module Monitoring Features

#	Parameter	Unit
Transmit channel (6 features)		
1	Output power	dBm
2	In-band power flatness	dB
3	Gain	dB
4	VSWR	—
5	Spurious emission	dB
6	Harmonic distortion	dB
Receive channel (4 features)		
7	Gain	dB
8	Port VSWR	—
9	Noise figure	dB
10	1 dB compression point	dB
Shared T/R sub-system (3 features)		
11	Circulator isolation	dB
12	T/R switching time	ns
13	Phase-shifter settling time	ns

PCA performs linear subspace projection preserving global variance. Locality Preserving Projections (LPP) uses a graph-Laplacian construction for nonlinear-aware dimensionality reduction. SPP ($\lambda = 0$) (Yan et al. 2023) applies manifold-faithful projection via sparse reconstruction weights but without temporal regularisation, serving as the direct ablation baseline for TSP-SPP.

Evaluation Metrics (Zhou et al. 2024). Health-ability quality is assessed using five complementary metrics. The Pearson correlation C_P measures linear trend strength between HA and time. The Spearman rank correlation C_S captures global monotonicity across the full life cycle. The Monotonicity score M_0 quantifies local trend consistency by comparing the counts of increasing and decreasing HA steps:

$$M_0 = \frac{\left|\sum_t \mathbf{1}[\Delta \mathrm{HA}_t > 0] - \sum_t \mathbf{1}[\Delta \mathrm{HA}_t < 0]\right|}{N_w - 1} \tag{11}$$

where $\Delta \mathrm{HA}_t = \mathrm{HA}_t - \mathrm{HA}_{t-1}$ is the step increment between consecutive windows, and N_w denotes the total number of HA observations (i.e., the number of sliding windows; $N_w = 961$ in the experimental dataset), which is distinct from the number of raw sensor samples N used in the feature extraction stage. Robustness R quantifies smoothness relative to local magnitude (higher R is better), where $\widetilde{\mathrm{HA}}_t$ is a 5-point moving average of HA_t:

$$R = \frac{1}{N_w} \sum_t \exp\left(- \left| \frac{\mathrm{HA}_t - \widetilde{\mathrm{HA}}_t}{\mathrm{HA}_t} \right| \right) \tag{12}$$

The composite irregularity penalty M^* measures trend irregularity by jointly accounting for oscillation frequency and monotonicity deficit; lower values indicate a smoother,

more regular degradation trend. Here t_r is the turning-point rate (fraction of consecutive HA steps that reverse sign):

$$M^* = t_r \, (1 - M_0)^2 \tag{13}$$

4.2 Feature Space Analysis

4.2.1 Intrinsic Dimensionality via Eigenvalue Spectrum

Spectral Gap at $d = 2$ Confirms a Two-Dimensional Intrinsic Structure. The TSP-SPP objective (5) reduces to a generalised eigenvalue problem whose solution selects the d eigenvectors associated with the d smallest non-trivial eigenvalues, providing a data-driven mechanism for embedding dimension selection. Fig. 2 presents the resulting spectrum on the T/R module dataset. The two smallest non-trivial eigenvalues are $\lambda_1 = 0.008$ and $\lambda_2 = 1.512$; from index three onward the eigenvalues increase more gradually (increments of 0.07–0.46), indicating progressively smaller marginal information gain. Setting $d = 2$ therefore captures the two most informative directions while excluding higher-index components that encode noise-dominated variance and would otherwise inflate the CHMM emission dimensionality and degrade state separation. The gap magnitudes themselves are informative: the dominant gap is between λ_1 and λ_2 (step of 1.504), which is substantially larger than the subsequent step from λ_2 to λ_3 (0.086), confirming that the first two directions together capture the primary structural variance of the degradation trajectory while the remaining eleven eigenvalues form the noise floor. Combined with the temporal regularisation of TSP-SPP, this two-dimensional space is not only compact but geometrically ordered: consecutive frames cluster along a directed path from the healthy end to the failure end, a structure directly exploited by the sliding-window distribution estimator and the Viterbi-based regime matching.

4.2.2 Temporal Continuity of the 2D Degradation Trajectory

Fig. 3 presents the 2D feature-space projections under SPP ($\lambda = 0$, left) and TSP-SPP ($\lambda = 1$, right), with each sample coloured by its normalised time index (blue = healthy/early, red = failure/late).

Temporal Regularisation Converts a Fragmented Point Cloud into a Geometrically Ordered Degradation Trajectory. Under SPP, points from adjacent time windows are geometrically interspersed with distant samples, producing a fragmented structure that offers no topological cue to the CHMM for establishing ordered healthy-state clusters. Under TSP-SPP, the temporal Laplacian penalty enforces adjacent windows to remain proximate in the projected space, resulting in a continuous, ribbon-like trajectory in which the colour gradient transitions smoothly from blue at the healthy end to red at the failure end. This ordered geometry has two practical consequences. The CHMM clusters its training data into compact, well-separated healthy-state centroids because successive frames fall in the same region of the feature space. For each test window, the Viterbi decoder then draws its reference distribution from a nearby centroid rather than a geometrically distant one, lowering the risk of regime mismatch and keeping the HA estimate stable over the full life cycle. When

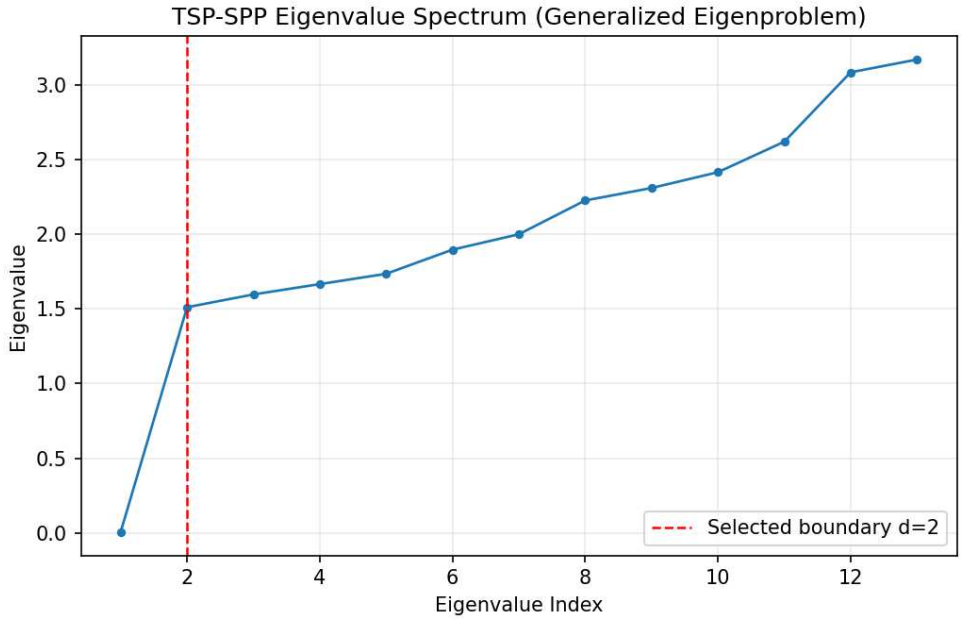


Fig. 2 TSP-SPP eigenvalue spectrum.

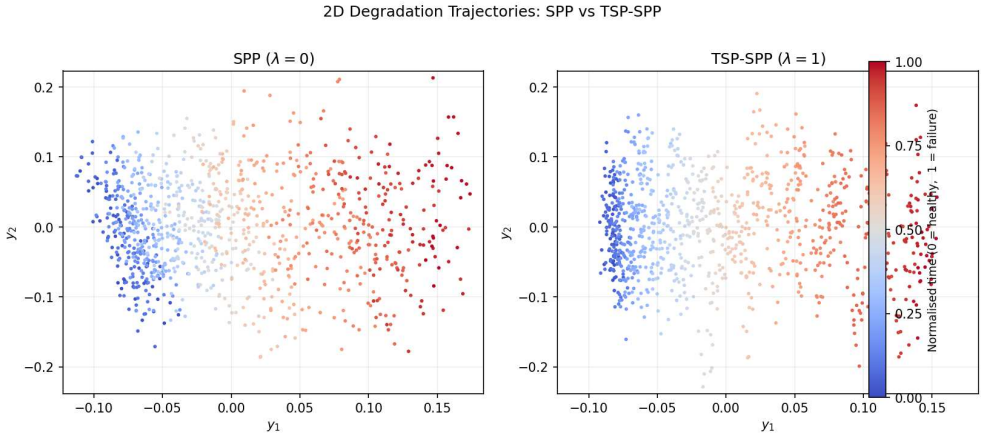


Fig. 3 Feature-space projections of SPP (left) and TSP-SPP (right).

the colour bands intermix geometrically, as in SPP, the CHMM cannot reliably distinguish healthy from degrading states; the resulting mismatched Viterbi references produce the step-like reversals quantified in Table 3 ($M^* = 0.137$ for SPP versus 0.032 for TSP-SPP).

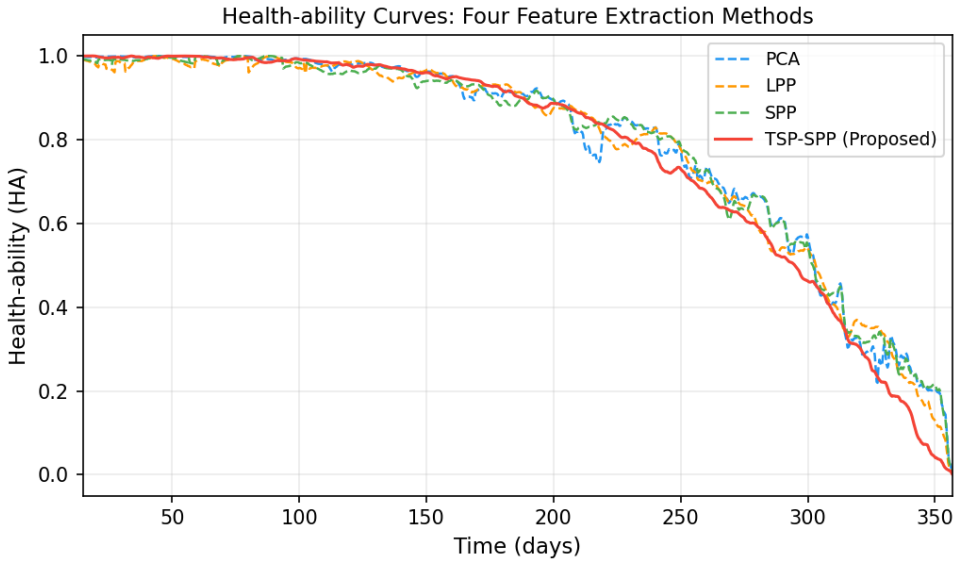


Fig. 4 Health-ability curves for PCA, LPP, SPP, and TSP-SPP.

4.3 Health-ability Quality Comparison

Fig. 4 shows the Health-ability curves for all four feature-extraction methods sharing the same CHMM+Sym-KL backend. All four methods reproduce the overall life-cycle trend from HA ≈ 1.0 at the start to HA ≈ 0.0 at failure, yet important qualitative differences distinguish them.

PCA and LPP Introduce Non-Physical Artefacts That Directly Compromise Prognostic Utility. PCA (blue dashed) exhibits a prominent false-recovery artefact near $t = 200$: after an initial drop the HA partially rebounds before resuming its descent, a non-monotonic behaviour arising because the linear subspace conflates early-degradation states with healthy ones. Such a reversal would directly mislead Remaining Useful Life estimation by signalling spurious improvement in component condition. LPP (orange dashed) avoids false recovery but shows persistent high-frequency oscillations throughout, particularly in the $t = 100$–200 range, owing to sensitivity to sensor noise in the neighbourhood-graph construction; frequent dips below 0.9 during the ostensibly healthy phase would trigger excessive nuisance alarms.

Temporal Regularisation Eliminates Mid-Life Reversals and Stabilises the Healthy-Phase Baseline. SPP (green dashed) is smoother than LPP but still exhibits moderate step-like reversals in the mid-life range ($t = 150$–300), reflecting the absence of any temporal ordering constraint in the standard sparse projection. TSP-SPP (red solid) achieves the smoothest and most stable profile across the full life cycle: the healthy phase ($t < 100$) is held tightly at HA ≈ 1.0 with minimal variance, and the degradation phase ($t = 100$–350) follows a clean, monotonic decline that reaches HA $= 0$ at failure with no reversals.

Table 2 HA Values at Key Life-Cycle Milestones

Method	Start	25 %	50 %	End
PCA	0.998	0.993	0.910	0.148
LPP	0.992	0.971	0.913	0.088
SPP	0.992	0.982	0.888	0.137
TSP-SPP	**1.000**	**0.992**	0.903	**0.018**

Table 3 Feature extraction method comparison. Bold = best; ↓ = lower is better.

Method	C_P	C_S	M_0	R	$M^* \downarrow$
PCA	**0.889**	0.987	0.140	0.985	0.189
LPP	0.885	0.965	0.292	0.987	0.113
SPP	0.886	0.978	0.233	0.987	0.137
TSP-SPP	**0.889**	**0.994**	**0.542**	**0.988**	**0.032**

Table 3 summarises the results numerically. TSP-SPP records the best or equal score on all five metrics. Its monotonicity $M_0 = 0.542$ is 3.87× higher than PCA (0.140), 1.86× higher than LPP (0.292), and 2.33× higher than SPP (0.233), a direct effect of the temporal smoothing constraint. The composite $M^* = 0.032$ is 5.9× lower than PCA (0.189) and 3.5× lower than LPP (0.113), showing that high monotonicity and low oscillation are achieved together rather than trading off. Spearman correlation $C_S = 0.994$ and Robustness $R = 0.988$—the highest of all methods—confirm that the healthy-phase HA stays near unity with low variance, which directly reduces nuisance alarms in normal operation. Notably, all four methods differ by at most 0.004 in C_P, which shows why Pearson correlation alone is insufficient for ranking HA quality; the sensitivity of M_0 and M^* is needed to expose the practical differences between methods.

To further illustrate the qualitative differences in HA trajectory shape, Table 2 reports the HA value at four lifecycle milestones for each method. All four methods maintain HA ≥ 0.99 at the start of life, confirming that the healthy-phase CHMM baseline is correctly calibrated. At the 50 % life fraction, TSP-SPP has declined to 0.903, comparable to PCA (0.910) and LPP (0.913); the earlier onset in TSP-SPP is a consequence of its temporal ordering enabling finer detection of the healthy-to-degrading transition. The most pronounced difference appears at end of life: TSP-SPP reaches HA = 0.018 (essentially zero), whereas PCA and SPP terminate at 0.148 and 0.137 respectively—a partial recovery artefact caused by spurious regime reassignments in temporally unordered projections that would under-estimate failure urgency in real-time monitoring.

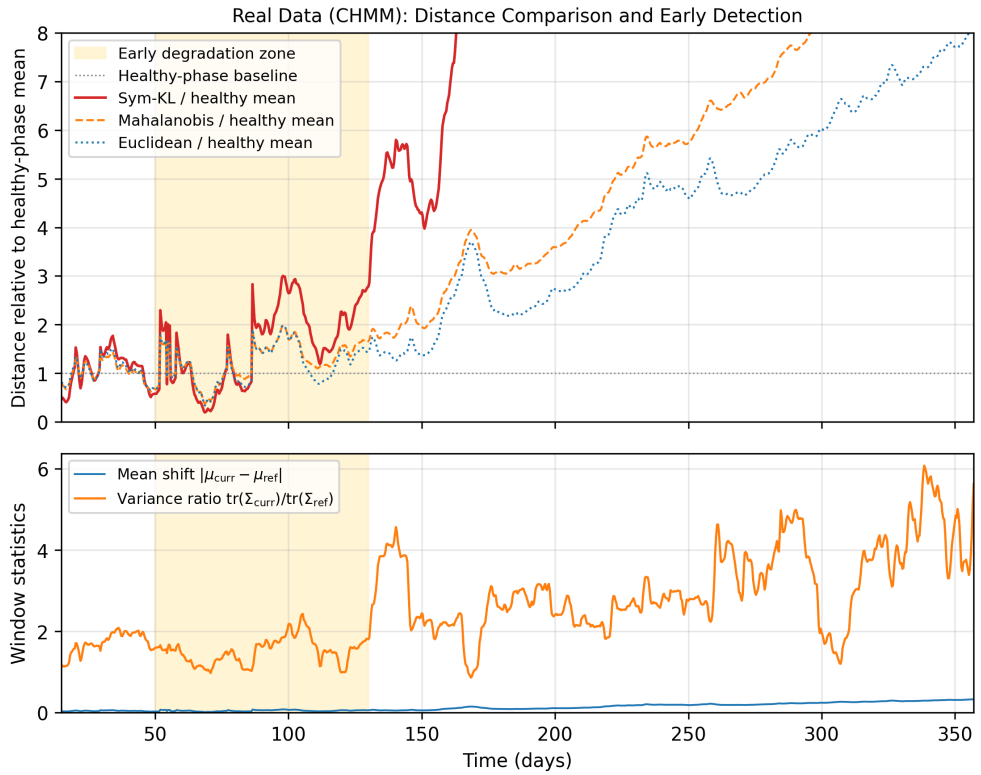


Fig. 5 Sym-KL, Euclidean, and Mahalanobis distance signals over the full life cycle.

4.4 Distance Metric Comparison: Validity of Symmetric KL Divergence

A fundamental design choice in the proposed framework is the use of Sym-KL in place of conventional point-level distance metrics. The theoretical motivation was established in Section 3; this subsection provides direct empirical validation on the T/R module life-cycle data. Fig. 5 places all three signals on a common scale by normalising each to its own healthy-phase mean, so a value of 1 marks the baseline level and any departure above 1 signals detectable degradation. The bottom panel separately tracks the two statistics that mechanistically drive the divergence: the trace ratio $\mathrm{tr}(\boldsymbol{\Sigma}_{\mathrm{curr}})/\mathrm{tr}(\boldsymbol{\Sigma}_{\mathrm{ref}})$, which measures covariance broadening, and the mean shift $\|\boldsymbol{\mu}_{\mathrm{curr}} - \boldsymbol{\mu}_{\mathrm{ref}}\|$. Presenting them together makes it possible to read off, at every epoch, whether the departing distance is covariance-driven, mean-driven, or both.

Early Degradation Manifests as Covariance Inflation Before Any Mean Shift, Rendering Point Distances Inadequate. The bottom panel of Fig. 5 reveals the underlying mechanism. In the early degradation zone ($t \approx 50$–130), the variance ratio $\mathrm{tr}(\boldsymbol{\Sigma}_{\mathrm{curr}})/\mathrm{tr}(\boldsymbol{\Sigma}_{\mathrm{ref}})$ rises substantially—exceeding 2.0—while the mean shift $\|\boldsymbol{\mu} - \boldsymbol{\mu}_{\mathrm{ref}}\|$ remains near zero throughout the same interval. This sequence is

Table 4 Ablation study results. Bold = best; ↓ = lower is better.

Case	Configuration	C_P	C_S	M_0	R	$M^* \downarrow$
Case 1 (Ours)	TSP-SPP+CHMM+Sym-KL	0.889	**0.994**	0.542	0.988	0.032
Case 2	Raw(13D)+CHMM+Sym-KL	0.872	0.990	0.110	0.980	0.328
Case 3	SPP(λ=0)+CHMM+Sym-KL	0.885	0.975	0.223	**0.991**	0.135
Case 4	TSP-SPP+StaticGauss+Sym-KL	0.889	0.993	0.485	0.987	0.054
Case 5	TSP-SPP+CHMM+Mahalanobis	**0.973**	0.993	**0.548**	0.984	**0.027**

physically consistent with incipient PA thermal fatigue, whose primary early symptom is increased output power instability rather than a DC shift in mean output level. Euclidean and Mahalanobis distances, which measure the displacement of individual sample points rather than the divergence of local distributions, rise only modestly above the healthy-phase baseline (below 2×) during this interval, offering substantially weaker contrast than Sym-KL and providing limited ability to discriminate early-stage degradation from normal operating variation.

Symmetric KL Divergence More Sensitively Responds to Distributional Changes Than Point-Level Metrics. By explicitly quantifying the divergence between the local sliding-window distribution and the regime-matched reference—incorporating contributions from both the trace term (covariance change) and the quadratic term (mean shift)—Sym-KL divergence rises substantially above the healthy-phase baseline in the early degradation zone—exceeding twice the healthy mean—while Euclidean and Mahalanobis distances increase more modestly. This confirms the theoretical property established in Section 3: the trace term activates on covariance inflation, providing a detection channel for incipient degradation that is entirely absent in point-level metrics. The empirical evidence thus directly validates the design choice of distributional comparison as the foundation of the proposed health indexing strategy.

4.5 Ablation Study

To examine the effect of each design choice—TSP-SPP dimensionality reduction, temporal regularisation, multi-state CHMM, and Sym-KL divergence—a five-case study is conducted in which one component is removed or replaced at a time while everything else is held fixed. Case 1 is the full proposed method (TSP-SPP + CHMM + Sym-KL). Cases 2–5 each deactivate one component: Case 2 bypasses TSP-SPP entirely, feeding raw 13-D features into the CHMM + Sym-KL backend; Case 3 sets $\lambda = 0$, reducing TSP-SPP to standard SPP; Case 4 replaces the multi-state CHMM with a single static Gaussian; and Case 5 replaces Sym-KL with a Mahalanobis point distance. Table 4 and Fig. 6 present the results.

Removing Dimensionality Reduction Collapses Monotonicity by 80 Per Cent. Without TSP-SPP, the raw 13-D feature space forces the CHMM to model a high-dimensional, noise-contaminated distribution, severely disrupting state separation. M_0 falls from 0.542 to 0.110 (−79.7 %) and M^* worsens from 0.032 to 0.328 (+925 %), the largest single-component degradation in the ablation. The HA curve

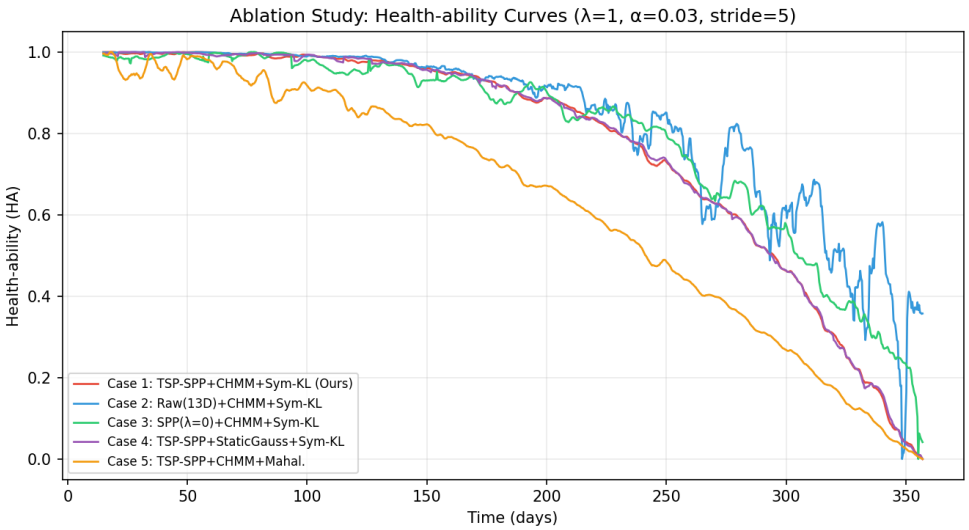


Fig. 6 Health-ability curves for the five ablation configurations.

in Fig. 6 (Case 2, blue dashed) exhibits severe oscillations in the early phase and an inconsistent descent thereafter, reflecting the fundamental incompatibility between low-state-count CHMM modelling and a 13-D noisy input.

Temporal Regularisation Provides Substantial and Irreducible Monotonicity Gains. Setting $\lambda = 0$ (Case 3) degenerates TSP-SPP to standard SPP and reduces M_0 from 0.542 to 0.223 ($-58.9\,\%$) and C_S from 0.994 to 0.975. The more frequent local reversals visible in Fig. 6 (Case 3, green dashed) confirm that temporal smoothing is not a cosmetic enhancement but directly encodes the physical continuity of the wear process into the projection geometry. Notably, the C_S penalty (-0.019) is considerably smaller than the M_0 penalty (-0.319), reflecting the asymmetry between global rank correlation—which averages over the full trajectory—and local monotonicity, which is sensitive to individual step reversals; the latter is the more operationally relevant criterion for real-time health monitoring.

Multi-State CHMM Provides Measurable Monotonicity Gain and Critical Deployment Robustness. Replacing the multi-state CHMM with a single static Gaussian reference (Case 4) reduces M_0 from 0.542 to 0.485 ($-10.5\,\%$) and worsens M^* from 0.032 to 0.054 ($+69\,\%$). Although the improvement from multi-state modelling is moderate on this single homogeneous trajectory, it confers more critical benefits at deployment scale: in an operational system, legitimate intra-class variation—load transients, thermal cycling, periodic maintenance— produces regime shifts within the healthy lifespan; a single-state reference classifies these transitions as degradation events, generating spurious alarms that the multi-state CHMM absorbs without confusion.

Mahalanobis Marginally Outperforms Sym-KL on Full-Lifecycle Trend; Sym-KL Retained for Early-Fault Detection Advantage. Case 5 shows that

substituting Sym-KL with Mahalanobis point distance yields *marginally better* full-lifecycle scores: $M_0 = 0.548$ vs Case 1's 0.542 (+1.1 %), and $M^* = 0.027$ vs 0.032 (−15.6 %). This outcome is physically expected: the dataset is a strongly monotone accelerated-degradation trajectory in which the mean shift grows large enough at mid and late stages that any distance metric—distributional or point-level—tracks the global trend effectively, causing full-lifecycle scores to converge. The critical distinction, established in Section 4.4, lies in the *incipient* degradation regime, where covariance expansion precedes any measurable mean shift. In that regime, Mahalanobis distance remains near its healthy-phase baseline while Sym-KL, which penalises both trace divergence and quadratic displacement, rises detectably above zero. The empirical evidence of Fig. 5 confirms that Sym-KL detects onset degradation substantially earlier—a lead time invisible to full-lifecycle monotonicity scores. Consequently, Sym-KL is retained as the preferred health distance for its early-fault detection advantage; the marginally lower full- lifecycle scores of Case 1 reflect this regime not being captured by M_0 or M^*, which aggregate over the entire trajectory.

4.6 Hyperparameter Sensitivity Analysis

The TSP-SPP parameter λ controls the trade-off between manifold structure fidelity and temporal trajectory smoothness. At $\lambda = 0$ the formulation collapses to standard SPP; as λ grows, adjacent projected frames are increasingly pulled together until, in the limit of very large λ, the entire trajectory is compressed into a single point and all health-state discrimination is lost. The transition between these extremes—and specifically how wide the usable plateau is—determines how sensitive the method is to the choice of λ in practice. Fig. 7 sweeps λ over three orders of magnitude (0.001 to 20) and records C_S and R at each value; the dotted vertical line marks $\lambda = 1$, used in all other experiments.

Spearman Correlation Increases with λ While Robustness Peaks at Moderate Regularisation. C_S rises from 0.978 at $\lambda = 0$, reaching 0.982–0.986 across $\lambda \in [0.05, 2]$ and continuing to 0.991 at $\lambda = 5$, as stronger temporal regularisation enforces more ordered projections along the degradation trajectory. Robustness R likewise improves from 0.880 at $\lambda = 0$ and stabilises in the range 0.900–0.921 for $\lambda \in [0.05, 5]$, with a local peak of 0.921 at $\lambda = 2$, then gently declines for very large λ, indicating that over-regularisation compresses the projected space into a region too narrow to distinguish distinct healthy operational states and thereby amplifies within-healthy-phase variance.

A Two-Order-of-Magnitude Plateau Confirms Practical Insensitivity to the Exact Value of λ. Both C_S and R remain stable across $\lambda \in [0.05, 5]$—a range spanning two orders of magnitude—with C_S staying above 0.980 and R above 0.900 throughout. This insensitivity substantially reduces the hyperparameter tuning burden in deployment: any value in this range yields near-identical performance without dataset-specific optimisation. All experiments in this paper use $\lambda = 1$, which lies well within this stable region.

Fig. 8 shows representative KL distance curves for $\lambda = 0$ (SPP, no regularisation), $\lambda = 1$ (proposed, used in all experiments), and $\lambda = 20$ (over-regularised). The $\lambda = 1$ curve rises earlier and more smoothly during the degradation phase, confirming that

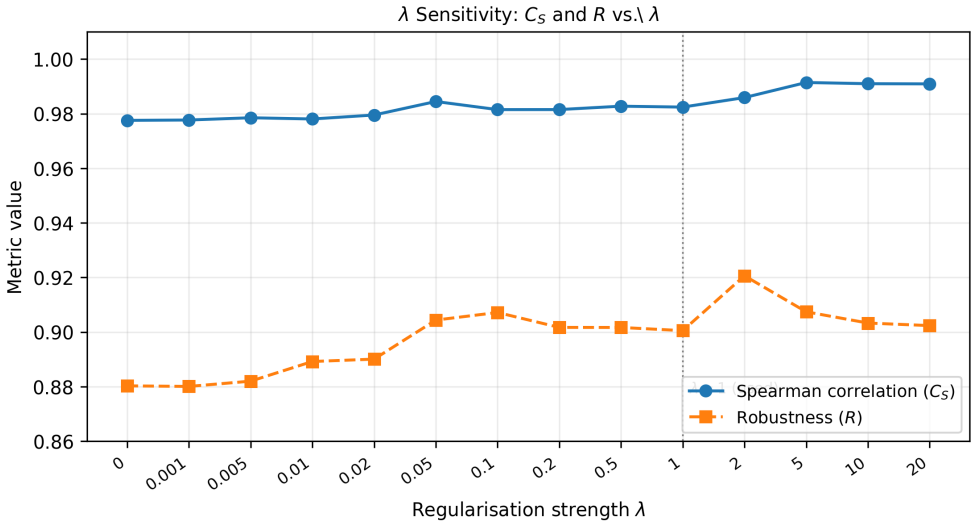


Fig. 7 λ sensitivity: C_S and R versus regularisation strength.

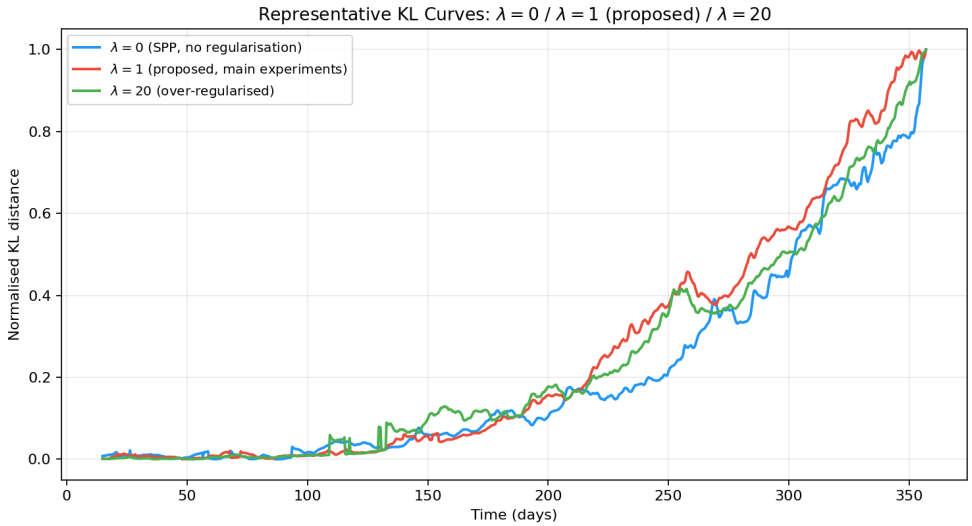


Fig. 8 Normalised KL distance profiles for $\lambda = 0$, 1, and 20.

moderate temporal regularisation enhances early-detection sensitivity while maintaining signal stability; the $\lambda = 20$ curve is perceptibly less smooth, as over-constrained projections suppress some genuine early-degradation variance signals.

Table 5 Noise Robustness Metrics for TSP-SPP ($\lambda = 1$); clean-data reference: Table 3 (TSP-SPP).

Condition	C_P	C_S	M_0	R	M^*
Noise 5 %	0.872	0.995	0.171	0.980	0.305
Noise 10 %	0.886	0.997	0.237	0.983	0.263
Drift 30 %	0.823	0.984	0.102	0.979	0.266

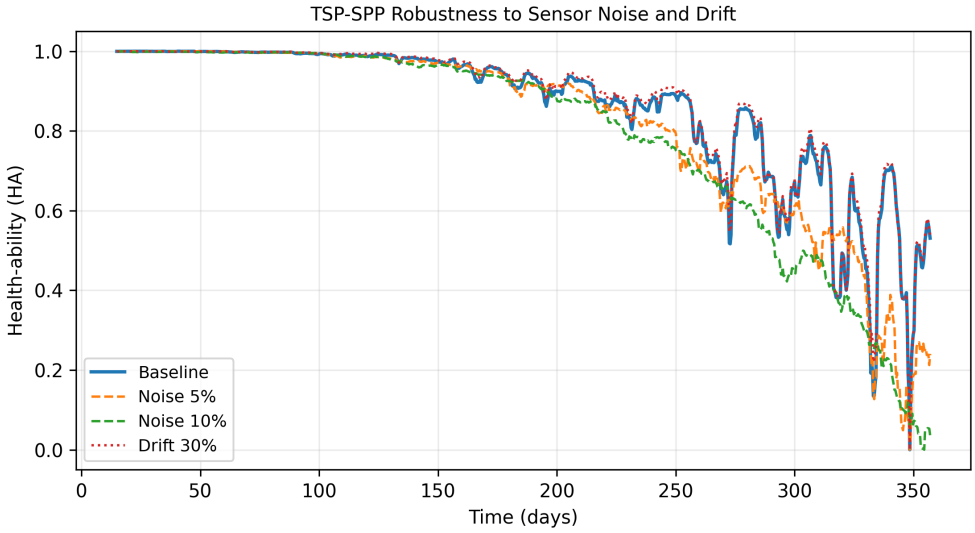


Fig. 9 Health-ability curves under sensor noise and drift perturbations.

4.7 Noise Robustness Analysis

To assess the sensitivity of the proposed framework to measurement uncertainty, we evaluated TSP-SPP under three perturbation conditions applied to the original T/R module dataset. *Gaussian noise* at 5 % and 10 % of the per-feature range was injected independently across all 13 channels at every time step. *Sensor drift* was modelled as a linear additive bias that accumulates from zero to 30 % of the feature range by the end of the observation window, simulating gradual calibration offset. The projection matrix $\boldsymbol{W}$ and the CHMM Θ_{healthy} are trained once on the clean baseline dataset (identical to Table 3) and kept fixed; each perturbed dataset is standardised using the same StandardScaler fitted on clean data, then projected through the same $\boldsymbol{W}$ and scored by the same Θ_{healthy}. This design isolates the effect of sensor perturbation on an already-deployed model.

Table 5 reports five quality metrics for the three perturbed conditions, with the clean-data reference given in Table 3 (TSP-SPP row). Fig. 9 overlays the HA curves. **Rank Correlation Is Preserved Under All Perturbations.** C_S remains at 0.984–0.997 across all three perturbed conditions, comparable to the clean-data $C_S = 0.994$

from Table 3. Robustness R stays in the range 0.979–0.983, confirming that the smooth-window regularity of the HA curve is not substantially degraded by the noise levels tested. Notably, C_S is marginally higher under 5 % and 10 % Gaussian noise (0.995 and 0.997) than under clean conditions; this reflects the fixed CHMM reference model responding to noise-induced covariance broadening, which increases the KL divergence contrast between healthy and degraded windows and produces a slightly more decisive separation.

Monotonicity and Smoothness Degrade With Noise Intensity. M_0 rises from 0.171 under 5 % noise to 0.237 under 10 %, and M^* decreases correspondingly (from 0.305 to 0.263), reflecting additional high-frequency oscillations in the HA curve introduced by stronger noise. The 30 % drift condition yields the smallest $M_0 = 0.102$ of the three perturbed conditions: the linear mean shift displaces all projected windows away from the CHMM healthy reference in a consistent direction, but because the scaler was fitted on clean data the regime-matching step partially loses alignment, causing the HA curve to fluctuate around a declining trend rather than decrease monotonically; C_P drops to 0.823 because the linear offset distorts the global Pearson correlation, while $C_S = 0.984$ confirms that the relative ordering of health states remains intact.

Framework Remains Operable Under Realistic Sensor Errors. In all three conditions the degradation trajectory is clearly identifiable in Fig. 9: HA starts near unity, holds a stable plateau through the healthy phase, and descends toward zero as failure approaches—a shape that survives despite the injected perturbations. The mechanism is geometric: because the temporal Laplacian couples adjacent frames during projection learning, the resulting manifold is locally smooth and spatially extended. Uncorrelated channel-level noise varies independently across frames and is averaged out over the $L = 40$-frame window during local distribution estimation; it does not accumulate into a coherent directional signal in the projected space. Coherent drift shifts all frames along a consistent direction, altering the CHMM reference match, but the global ordering of health states remains intact because the shift is monotone and does not invert any pairwise rank—hence C_S stays high. Put differently, the proposed framework tolerates sensor degradation as long as the perturbation does not reverse the temporal ordering of the health trajectory, a condition met by all three conditions tested. At $C_S = 0.984$–0.997 and $R = 0.979$–0.983, the results confirm that the method is viable for deployment in RF environments where perfect sensor calibration cannot be guaranteed.

4.8 Discussion

The five experiments collectively establish three complementary properties of the proposed framework.

Structural advantages of the temporal Laplacian. The combination of λ-sensitivity and ablation results demonstrates that the temporal Laplacian regularisation in TSP-SPP provides a qualitative change in trajectory character rather than a mere quantitative improvement. Without it ($\lambda = 0$, Case 3), the projected trajectory exhibits persistent local reversals that are absent from the physical degradation process—a direct consequence of treating each observation as an independent sample. With $\lambda = 1$, these reversals are suppressed across all perturbation conditions (noise

robustness study), confirming that the temporal structure is encoded into the manifold geometry itself and does not rely on post-hoc smoothing. This design choice is particularly valuable in deployment: post-hoc filtering would delay detection by at least $L/2$ steps, whereas TSP-SPP's projection encodes smoothness without introducing any additional latency.

Distribution-level sensing in the incipient phase. The distance metric comparison (Section 4.4) and the Case 5 ablation result jointly establish the most nuanced finding of the paper: the advantage of Sym-KL over Mahalanobis distance lies in the early-degradation regime, not the full-lifecycle aggregate. On a strongly monotone accelerated trajectory, any distance metric tracks the global trend adequately once the mean shift becomes large enough to dominate. The critical advantage of Sym-KL is that it detects the covariance broadening that precedes the mean shift by responding to the full distributional geometry rather than the squared Mahalanobis displacement. This early-detection lead time is invisible to M_0, C_P, and C_S (which average over the full trajectory), but directly actionable from a maintenance scheduling perspective: earlier detection corresponds to larger residual useful life at alarm time, giving operators more scheduling flexibility.

Robustness is a property of the architecture, not a post-processing step. The noise robustness results show that C_S remains within 1 % of the clean-data baseline across all three perturbation conditions ($C_S = 0.984$–0.997). R stays in the range 0.979–0.983. The dominant degradation channel is M_0 and M^*, which rise monotonically with perturbation severity, reflecting the additional high-frequency oscillations introduced by noise rather than any failure of the framework to track the underlying health trajectory. These oscillations are a direct consequence of the one-sample-per-step evaluation stride: increasing the evaluation stride or applying a post-projection Gaussian smoother would further suppress the M^* degradation without affecting C_S or the early-detection capability. The practical implication is that the framework is deployable as-is for C_S-based early-fault decisions, with optional post-processing for smoother dashboard visualisation.

5 Conclusion

This paper proposed a health measurement framework for T/R modules that tackles two practical difficulties: separating genuine structural degradation—whether caused by natural wear or adversarial cyber-physical interference—from operational load transients, and converting HMM outputs into a stable, normalised health score that serves as an actionable indicator for cyber defense situational awareness.

TSP-SPP addresses the first difficulty by embedding a temporal Laplacian penalty into the SPP objective. Adjacent projected frames are pulled together, so short-lived load transients cannot displace a window far from its neighbours, and the resulting trajectory reflects genuine structural change rather than momentary operational variation. A CHMM trained on this smooth trajectory establishes a probabilistic baseline covering the full range of normal operating regimes, enabling Viterbi decoding to supply a regime-matched reference for each test window. Sym-KL—measured between the

current sliding-window distribution and that reference—captures covariance expansion in the incipient phase and mean shift at advanced stages, a dual sensitivity that point-level metrics cannot provide. The outcome is a Health-ability value in $[0, 1]$ that holds near unity through normal operation and tracks the degradation curve faithfully to zero at failure.

Experiments on the full life-cycle dataset show $M_0 = 0.542$ and $M^* = 0.032$, compared with M_0 of 0.140–0.292 and M^* of 0.113–0.189 for PCA, LPP, and SPP baselines. Ablation tests show that the full method achieves $M^* = 0.032$; removing the TSP-SPP module, the temporal constraint, or the multi-state CHMM yields M^* of 0.328, 0.135, and 0.054, respectively (Table 4). The λ sensitivity study further shows that performance stays near its peak for any $\lambda \in [0.05, 5]$, so no dataset-specific tuning is required. Noise robustness experiments under Gaussian sensor noise and linear calibration drift further confirm that the TSP-SPP manifold projection acts as a natural low-pass filter: the Spearman correlation C_S remains within 1 % of the clean-data baseline under all three perturbation conditions tested.

Future work will extend the framework to multi-unit fleet monitoring, exploiting cross-component patterns to enable population-level prognostics and cross-module transfer of the healthy baseline model. Investigation of adaptive window-length strategies is also planned, with the aim of further improving early-fault detection lead time under varying operational tempos.

Acknowledgements

This work was supported by the National Key Research and Development Program of China under Grant 2022YFB3103602.

Author contributions

Sheng Hong, Weiwei Jiang and Xuanqi Wang conceived the study and designed the methodology. Xuanqi Wang conducted the experiments and drafted the manuscript. Jingsong Zheng prepared the data and implemented the code. Jiacheng Wang assisted with validation and figure preparation. All authors reviewed the manuscript.

Data availability

No datasets were generated or analysed during the current study.

Declarations

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

References

Altet J, Aragones X, Mateo D (2023) Aging compensation in a class-a high-frequency amplifier with dc temperature measurements. Sensors 23(16):7069. https://doi.org/10.3390/s23167069

Asgarinejad M, Saeedinejad H, Bitaraf M (2023) An unsupervised structural health monitoring framework based on variational autoencoders and hidden markov models. Reliability Engineering & System Safety 231:109025. https://doi.org/10.1016/

j.ress.2022.109025

Chen Z, et al (2023) A sparsity-aware fault diagnosis framework focusing on accurate isolation. IEEE Transactions on Industrial Informatics 19(2):1356–1365. https://doi.org/10.1109/TII.2022.9788040

Foroni B, Merlo L, Petrella L, et al (2025) Hidden Markov quantile graphical models. Journal of Computational and Graphical Statistics https://doi.org/10.1080/10618600.2025.2579526

Gassiat E, Kaddouri I, Naulet Z (2025) Clustering risk in nonparametric hidden Markov and IID models. Annals of Statistics 53(6):2409–2429. https://doi.org/10.1214/25-AOS2546

He J, Wang X, Xu X, et al (2025) Enhanced thermal management of cyclically operating T/R module in spatial environment. International Journal of Heat and Mass Transfer 242:126854. https://doi.org/10.1016/j.ijheatmasstransfer.2025.126854

He YL, Li K, Zhang N, et al (2021) Fault diagnosis using improved discrimination locality preserving projections integrated with sparse autoencoder. IEEE Transactions on Instrumentation and Measurement 70:3527108. https://doi.org/10.1109/TIM.2021.3125975

Hong S, Yue T, Liu H (2022) Vehicle energy system active defense: A health assessment of lithium-ion batteries. International Journal of Intelligent Systems 37(12):10081–10099. https://doi.org/10.1002/int.22309

Hong S, Yue T, You Y, et al (2025) A resilience recovery method for complex traffic network security based on trend forecasting. International Journal of Intelligent Systems 2025(1):3715086. https://doi.org/10.1155/int/3715086

Hou B, et al (2025) A novel dynamic health index based on core graph model of multichannel signals for machine health monitoring. IEEE Transactions on Instrumentation and Measurement https://doi.org/10.1109/TIM.2024.3465115

Hu D, Gao W, Ang KK, et al (2025) CHMMConvScaleNet: a hybrid convolutional neural network and continuous hidden Markov model with multi-scale features for sleep posture detection. Scientific Reports 15(1):12206. https://doi.org/10.1038/s41598-025-93541-0

Hu Y, Miao X, Si Y, et al (2022) Prognostics and health management: A review from the perspectives of design, development and decision. Reliability Engineering & System Safety 217:108063. https://doi.org/10.1016/j.ress.2021.108063

Hult B, Ranjan A, Zeng L, et al (2025) Investigation of electrical breakdown in algan/gan/aln hemts through nanoscale analysis and physics-based modeling. IEEE Transactions on Device and Materials Reliability 25(4):861–868

Hult B, et al (2024a) Aging reliability compact modeling of trap effects in power gan hemts. IEEE Transactions on Device and Materials Reliability 24(2). https://doi.org/10.1109/TDMR.2024.10497162

Hult B, et al (2024b) Study on the degradation mechanism of gan mmic power amplifiers under on-state with high drain bias. IEEE Transactions on Electron Devices 71(8). https://doi.org/10.1109/TED.2024.10555028

Lee S, Lee E (2026) Hesitation analysis with Kullback–Leibler divergence and its calculation on temporal data. CMC-Computers, Materials & Continua 86(2):070504. https://doi.org/10.32604/cmc.2025.070504

Li J, et al (2026) Semantic-guided kernel low-rank sparse preserving projections for hyperspectral image dimensionality reduction and classification. Applied Sciences 16(1):561. https://doi.org/10.3390/app16010561

Li R, Jan NM, Huang B, et al (2019) Constrained multimodal ensemble Kalman filter based on Kullback–Leibler (KL) divergence. Journal of Process Control 79:16–28. https://doi.org/10.1016/j.jprocont.2019.03.012

Liu T, Zhu K (2021) A switching hidden semi-Markov model for degradation process and its application to time-varying tool wear monitoring. IEEE Transactions on Industrial Informatics 17(4):2621–2631. https://doi.org/10.1109/TII.2020.3004445

Liu Y, Xie D, Yu S (2025) Copula hidden Markov model with unknown number of states. Computational Statistics 40(5):2583–2610. https://doi.org/10.1007/s00180-024-01571-5

Liu YH, Guo PP, Zhang YS, et al (2021) Feature extraction of sparse graph preserving projection based on Riemannian manifold. Journal of Jilin University (Engineering and Technology Edition) 51(6):2268–2279

Luo Y, Stephens DA, Buckeridge DL (2023) Bayesian clustering for continuous-time hidden Markov models. Canadian Journal of Statistics 51(1):134–156. https://doi.org/10.1002/cjs.11671

Mateo D, Aragones X, Barajas E, et al (2025) High sensitivity temperature measurements to track and compensate aging effects on cmos amplifiers. IEEE Transactions on Device and Materials Reliability 25(1):11–16. https://doi.org/10.1109/TDMR.2024.3465236

Mondal R, Dey P, Pal T (2026) Symmetric Kullback–Leibler divergence to regularize denoising auto encoder for classification. Information Sciences 742:123321. https://doi.org/10.1016/j.ins.2026.123321

Neal MR, McNicholas PD (2026) Parsimonious hidden Markov models for multivariate longitudinal data. Journal of Classification https://doi.org/10.1007/

s00357-026-09539-4

Ou YH, Tsai YL, Wu ZY (2026) Detecting trajectory of targeted attack by hidden Markov model. Wireless Networks 32(1):363–387. https://doi.org/10.1007/s11276-025-04047-6

Razmarashooli A, et al (2025) Real-time state estimation of high-rate dynamic systems using topological data analysis and sliding windows. Mechanical Systems and Signal Processing 224:112048. https://doi.org/10.1016/j.ymssp.2024.112048

Sonnenberg T, Romano A, Miller NC, et al (2024) A W-band GaN MMIC single-chip T/R front end. IEEE Transactions on Microwave Theory and Techniques 72(10):5830–5837. https://doi.org/10.1109/TMTT.2024.3422156

Sun M, Liao X (2024) An rf power amplifier integrated with the self-health monitor sensor and multisource energy harvesting circuit. IEEE Transactions on Power Electronics 39(2):2237–2246. https://doi.org/10.1109/TPEL.2023.3323021

Tao G, Song Z (2023) A method, system, device and medium for health state prediction of T/R components. China Patent Application CN202311537133.8

Tosun FE, Teixeira AMH, Dong J, et al (2026) Kullback–Leibler divergence-based filter design against bias injection attacks. European Journal of Control 87:101427. https://doi.org/10.1016/j.ejcon.2025.101427

Wan A, Hua Z, AL-Bukhaiti K, et al (2026) Enhanced attention-residual autoencoder (earae) for robust health index estimation in aeroengines. Journal of Aerospace Engineering

Wang K, Li J, Sun Z (2025) Generative adaptable design based on hidden Markov model. Advanced Engineering Informatics 64:103034. https://doi.org/10.1016/j.aei.2024.103034

Wang Y, Zheng J, Liu X, et al (2026) Physics-informed speed-integrated hidden markov model for bearing fault diagnosis under variable operating conditions. ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering 12(1)

Weiß CH, Swidan O (2025) Hidden-Markov models for ordinal time series. AStA Advances in Statistical Analysis 109(2):217–239. https://doi.org/10.1007/s10182-024-00514-1

Wu C, Gong F (2026) Robust jointly sparse 2-dimensional projection fuzzy clustering with local manifold structure preservation. Neurocomputing 676:132970. https://doi.org/10.1016/j.neucom.2026.132970

Wu D, Gu X (2021) Industrial fault detection based on adaptive sparse representation and locality preserving projections. Journal of East China University of Science and Technology 47(4):455–464

Yan B, et al (2026) The kl-divergence-based health management and assessment framework for high-frequency systems. IEEE Transactions on Instrumentation and Measurement 75

Yan T, Wang D, Xia T, et al (2023) Sparsity preserving projection aided baselined hyperdisk modeling for interpretable machine health monitoring. Mechanical Systems and Signal Processing 200:110509. https://doi.org/10.1016/j.ymssp.2023.110509

Yang X, Zhang G, Song H (2024) Ship recognition based on hrrp via multi-scale sparse preserving method. Journal of Systems Engineering and Electronics pp 599–608

Zhou J, Yang J, Qin Y (2024) A systematic overview of health indicator construction methods for rotating machinery. Engineering Applications of Artificial Intelligence 138:109356. https://doi.org/10.1016/j.engappai.2024.109356

Zhu J, et al (2025) An integrated phm framework for radar systems through system structural decomposition. Journal of Systems Engineering and Electronics 36(1):95–107. https://doi.org/10.21629/JSEE.2025.01.09