

Anomaly Detection as an Inverse Problem on Quotient Spaces

Equivalence Relations and Functional Manifolds as a General Diagnostic Framework

Technical Note

Eduardo Di Santi¹

¹Department of Applied Mathematics, University of Colorado Boulder

June 28, 2026

Abstract

This work proposes a quotient-space formulation of diagnostic monitoring in physical systems. The central claim is that anomaly detection is not fundamentally a supervised classification problem in raw observation space, nor does it require identifying the full forward physical model, which may be unknown, nonlinear, high-dimensional, or impractical to learn. Instead, the objective is to recover a representation in which observations that are physically equivalent with respect to the latent state are mapped to the same equivalence class.

The key step is to separate physical-state information from nuisance variability. Measurements are affected by transformations such as amplitude scaling, phase shift, speed variation, sensor offset, installation effects, and additive noise. These transformations modify the observed signal while preserving the underlying physical state. They therefore induce an equivalence relation on the observation space. The resulting quotient space is the natural geometric domain for diagnosis: the inverse problem is no longer posed as explicit recovery of the full generative model, but as recovery of membership in physically meaningful equivalence classes.

Under this formulation, a nominal operating condition is represented not by a single waveform or label, but by a connected quotient-space structure. Operationally, and in the sense of the companion manifold-saturation report, this structure is interpreted as a functional manifold. Anomaly detection becomes the problem of deciding whether an observation belongs to this nominal quotient manifold. Fault classification, when required, becomes a secondary operation: once an observation is detected outside the nominal manifold, it may be assigned to another quotient-space region or fault manifold. This provides a geometric mechanism for the emergence of functional manifolds: they arise by collapsing nuisance orbits that are irrelevant to the physical state.

After establishing the mathematical object, we introduce one computational approximation. Nominal trajectories are generated by Monte Carlo simulation over nuisance variables, a simple one-dimensional convolutional autoencoder is trained only on the synthetic nominal orbit, and the reconstruction error is used as an inverse-problem residual. An adaptive threshold estimates the local boundary of the nominal manifold for each physical asset. Synthetic validation on four parallel bearing signals provides preliminary evidence that quotient-space recovery is operationally feasible under controlled conditions. The experiments do not prove the mathematical formulation itself; they demonstrate that one computational approximation can recover sufficient quotient geometry to support anomaly detection and subsequent signal classification.

1 Introduction

Diagnostic monitoring in physical systems is often reduced to a pattern recognition problem: given an observation window, assign a semantic label such as nominal, degraded, or faulty. In rotating machinery, this usually becomes a bearing fault classification task, where vibration windows are mapped to labels such as healthy operation, outer-race fault, inner-race fault, or ball fault. This formulation is operationally convenient, and it is well represented in the machinery diagnostics and condition-monitoring literature [10, 11, 13]. However, it hides a more fundamental structure.

The physical state of a system is not observed directly. Instead, one observes a signal produced by a measurement process that mixes the latent physical state with nuisance factors such as operating speed, amplitude, phase, sensor gain, installation conditions, measurement offsets, and noise. Consequently, two observations can look different in raw signal space while representing the same physical condition. Conversely, two observations can appear superficially similar while belonging to different physical states.

This is the signature of an inverse problem. The observation is not the object of interest; it is the output of a forward process. However, in many industrial settings the full forward model may be unknown, nonlinear, high-dimensional, or impractical to identify. The objective is therefore not necessarily to learn or invert the complete generative model F . The objective is to recover a representation in which physically equivalent observations are identified, while physically meaningful deviations remain detectable. This distinction is related to, but not identical with, classical anomaly and novelty detection formulations [3, 4, 5], and to the broader theory of ill-posed inverse problems [1, 2].

The central claim of this work is the following:

Diagnostic monitoring can be formulated as inverse recovery over a quotient space induced by nuisance transformations of the observation process.

The quotient space appears because nuisance transformations modify the observed signal without changing the underlying physical state. These transformations induce an equivalence relation on the observation space. The diagnostic task can therefore be reformulated as recovery of membership in physically meaningful equivalence classes, rather than direct classification in raw observation space or explicit reconstruction of the full forward physical model. The emphasis on invariance, transformations, and geometry is consistent with the broader geometric learning perspective [9].

Under this view, a nominal operating condition is not represented by a single waveform or by a predefined label. It is represented by a connected functional manifold in quotient space. Anomaly detection is then a membership test: an observation is nominal if it belongs to the nominal quotient manifold and anomalous if it does not. Classification, when required, becomes a secondary operation: once an observation is detected outside the nominal manifold, it may be assigned to another quotient-space region or fault manifold.

The quotient space is not introduced because an autoencoder is used. The autoencoder is introduced because the diagnostic problem naturally lives on a quotient space. This distinction is essential. The mathematical formulation should remain valid if the computational operator is replaced by another representation learner, such as a contrastive encoder, a predictive latent model, a diffusion-based score model, or a physics-informed residual operator.

This work separates three layers that are often conflated:

- (i) the physical observation model;
- (ii) the quotient-space inverse problem induced by nuisance equivalence;

(iii) one practical computational approximation.

The first two layers define the scientific contribution. The third layer demonstrates that the formulation can be approximated in a concrete setting. In this work, the concrete setting is a synthetic bearing fault detection experiment, used not as the center of the theory but as a controlled validation of the proposed quotient-space inverse recovery framework.

2 Observation Model

Let $\mathcal{X} \subset \mathbb{R}^n$ denote the space of fixed-length observation windows. In the experimental section these observations are vibration windows, but the formulation does not depend on this particular sensing modality.

Let $\theta \in \Theta$ denote the latent physical state of the system. In a bearing application, examples of θ include nominal operation, outer-race degradation, inner-race degradation, ball damage, or transitional states not yet assigned a semantic label.

Let $\eta \in \mathcal{E}$ denote nuisance variables that modify the observed signal without changing the underlying physical state. Typical nuisance variables include:

- amplitude scaling due to gain or load variation;
- phase shift due to sampling alignment;
- speed perturbation;
- DC offset or baseline drift;
- sensor placement effects;
- additive measurement noise.

Physically, nuisance variables are typically constrained to compact ranges (bounded amplitude scaling, bounded phase or speed perturbation, bounded offset). This compactness of \mathcal{E} is used in Section 4 to establish boundedness of the nominal functional manifold.

The observation process is written as

$$x = F(\theta, \eta), \quad x \in \mathcal{X}. \tag{1}$$

The diagnostic objective is not to identify η . Nor is it necessarily to reconstruct the full forward model F . The diagnostic objective is to recover information about θ from x , while remaining invariant to nuisance degrees of freedom. This produces the inverse problem

$$x \mapsto \theta, \tag{2}$$

or, more precisely, recovery of the physical state modulo nuisance transformations of the observation process.

3 Diagnostic Monitoring as Inverse Recovery

Equation (1) describes a forward observation map. The inverse diagnostic problem is ill-posed in the classical sense: multiple nuisance configurations may produce different observations corresponding to the same latent physical state. The goal is therefore not to invert F exactly, nor to identify all nuisance variables explicitly. Instead, the goal is to construct a diagnostic representation

$$\Phi : \mathcal{X} \rightarrow \mathcal{Z} \tag{3}$$

such that $\Phi(x)$ preserves information about the latent physical state while suppressing nuisance variability.

A direct classifier attempts to approximate the map

$$x \mapsto y,$$

where y is a semantic or operational label. However, the label is not the primitive object. The primitive object is the latent physical state θ , and semantic labels are assigned only after a stable diagnostic structure is recovered.

This motivates a two-step view of diagnostic monitoring:

- (a) recover the geometric structure associated with latent-state equivalence classes;
- (b) assign operational semantics to the recovered structures.

In this view, classification is downstream of inverse recovery. Anomaly detection is not the prediction of a predefined class label. It is the decision that an observation does not belong to the recovered nominal structure. If additional non-nominal structures are recovered, subsequent classification may be performed by assigning the observation to another quotient-space region or fault manifold.

The experimental bearing case studied later in this work should therefore be read as one concrete instantiation of this general formulation. The theory concerns diagnostic monitoring under nuisance transformations; the bearing simulator provides a controlled physical example in which the proposed quotient-space recovery can be evaluated.

4 Quotient-Space Formulation

The key mathematical object is the equivalence relation induced by nuisance transformations. Let \mathcal{T} denote the physically motivated family of nuisance transformations introduced in Section 2: amplitude scaling, phase shift, speed perturbation, offset, and additive noise.

For the formal quotient argument, we let G_{nuis} denote the transformation group generated by \mathcal{T} , together with the identity transformation and the inverses of its elements. This is a closure construction: even if the physically motivated transformations in \mathcal{T} are sampled only over bounded operational ranges, the group generated by transformations of the same type is mathematically well defined.

The formal quotient-space construction is therefore stated in terms of the group action of G_{nuis} on \mathcal{X} . The compactness argument below does not require the full group G_{nuis} to be compact. It uses only the compact nuisance parameter set \mathcal{E} sampled in the Monte Carlo construction. Thus, the mathematical quotient is defined using an honest group action, while the experiment samples a compact physically admissible subset of nuisance configurations.

Definition 1 (Nuisance equivalence). *Two observations $x_1, x_2 \in \mathcal{X}$ are nuisance-equivalent, written $x_1 \sim x_2$, if there exists $g \in G_{\text{nuis}}$ such that*

$$x_2 = g \cdot x_1, \quad (4)$$

and both observations correspond to the same latent physical state θ .

The orbit of an observation x under nuisance transformations is

$$[x] = \{g \cdot x : g \in G_{\text{nuis}}\}. \quad (5)$$

The quotient space is the set of such equivalence classes:

$$\mathcal{X}/G_{\text{nuis}} = \{[x] : x \in \mathcal{X}\}. \quad (6)$$

Definition 2 (Nominal quotient manifold). *The nominal quotient manifold \mathcal{M}_{N} is the connected subset of $\mathcal{X}/G_{\text{nuis}}$ generated by observations sharing the nominal latent physical state:*

$$\mathcal{M}_{\text{N}} = \{[x] \in \mathcal{X}/G_{\text{nuis}} : x = F(\theta_{\text{N}}, \eta), \eta \in \mathcal{E}\}. \quad (7)$$

Proposition 1 (Compactness of the nominal nuisance orbit). *Suppose the nuisance parameter space \mathcal{E} is compact, and the generative map*

$$\eta \mapsto F(\theta_{\text{N}}, \eta)$$

viewed as a map from \mathcal{E} into the space of fixed-length observation windows equipped with the uniform, or sup-norm, topology, is uniformly bounded and equicontinuous over \mathcal{E} . Then the nominal nuisance orbit

$$\mathcal{O}_{\text{N}} = \{F(\theta_{\text{N}}, \eta) : \eta \in \mathcal{E}\} \subset \mathcal{X}$$

is relatively compact: its closure $\overline{\mathcal{O}_{\text{N}}}$ is a compact subset of \mathcal{X} .

Proof. By hypothesis, the family

$$\{F(\theta_{\text{N}}, \eta)\}_{\eta \in \mathcal{E}}$$

is uniformly bounded and equicontinuous in the relevant sup-norm topology. By the Arzelà–Ascoli theorem, this family is relatively compact in the sup-norm topology. Therefore its closure $\overline{\mathcal{O}_{\text{N}}}$ is compact.

In the bearing instantiation of Section 6.2, \mathcal{E} is a product of intervals and is therefore connected. If the map $\eta \mapsto F(\theta_{\text{N}}, \eta)$ is continuous, then its image \mathcal{O}_{N} is connected. \square

Corollary 1 (Boundedness of the nominal quotient structure). *Assume the hypotheses of Proposition 1. Since the quotient projection $\pi : \mathcal{X} \rightarrow \mathcal{X}/G_{\text{nuis}}$ is continuous, the image*

$$\pi\left(\overline{\mathcal{O}_{\text{N}}}\right)$$

is compact in $\mathcal{X}/G_{\text{nuis}}$. Moreover, in the connected setting described above, $\pi(\mathcal{O}_{\text{N}})$ is connected. Since $\mathcal{M}_{\text{N}} = \pi(\mathcal{O}_{\text{N}})$ by Definition 2, the nominal quotient manifold is contained in a compact quotient structure. This provides the geometric justification for describing \mathcal{M}_{N} as a bounded, connected functional structure, and operationally as a functional manifold in the sense clarified below.

Remark 1 (On the use of the term “manifold”). *Proposition 1 and Corollary 1 establish compactness and connectedness properties of the nominal structure under explicit assumptions, but they do not establish a smooth or topological manifold structure in the differential-geometric sense, such as local Euclidean charts, well-defined dimension, or curvature. The term functional manifold is used here, consistently with the companion report on manifold saturation [15] and with common usage in the representation-learning literature under the manifold hypothesis [9], to denote a bounded, connected region of the quotient space generated by a fixed latent regime under nuisance variability. A fully differential-geometric treatment is left to future work.*

Proposition 2 (Quotient factorization of diagnostic maps). *Assume that G_{nuis} acts on \mathcal{X} , that the action preserves the latent physical state, and that a diagnostic map $h : \mathcal{X} \rightarrow \mathcal{Y}$ depends only on that latent physical state. Then h is constant on G_{nuis} -orbits and therefore factors through the quotient projection $\pi : \mathcal{X} \rightarrow \mathcal{X}/G_{\text{nuis}}$. That is, there exists a map $\bar{h} : \mathcal{X}/G_{\text{nuis}} \rightarrow \mathcal{Y}$ such that*

$$h = \bar{h} \circ \pi.$$

Proof. Let $x_1, x_2 \in \mathcal{X}$ lie in the same G_{nuis} -orbit. Then $x_2 = g \cdot x_1$ for some $g \in G_{\text{nuis}}$. By assumption, the action of g preserves the latent physical state, and the diagnostic map h depends only on that state. Hence $h(x_1) = h(x_2)$. Therefore h is constant on equivalence classes. By the universal property of quotient maps, h factors through π , so there exists $\bar{h} : \mathcal{X}/G_{\text{nuis}} \rightarrow \mathcal{Y}$ such that $h = \bar{h} \circ \pi$. \square

Corollary 2 (Equivalence-class diagnosis). *If two observations lie in the same nuisance orbit, no physically meaningful diagnostic decision should distinguish them on the basis of nuisance variation alone. Consequently, nominal observations should be identified as elements of a nominal equivalence class, or a connected region of such classes, in the quotient space.*

Under this definition, anomaly detection becomes a membership problem:

$$[x] \in \mathcal{M}_N \quad \text{or} \quad [x] \notin \mathcal{M}_N. \tag{8}$$

This is the point at which anomaly detection enters the formulation. An anomaly is not primarily a label. It is a failure of membership in the nominal quotient manifold. Classification, when required, is a subsequent step that assigns non-nominal observations to other quotient-space regions or fault manifolds.

5 Functional Manifolds as Quotient Geometry

The quotient-space formulation provides a geometric mechanism through which functional manifolds emerge. A functional manifold is not postulated as an arbitrary structure discovered by an algorithm. It arises because nuisance transformations generate orbits in observation space, and the quotient projection collapses each orbit into a single equivalence class.

Thus, the observed manifold structure has two sources:

- (i) physical continuity of the underlying system;
- (ii) identification of observations that differ only by nuisance transformations.

In raw signal space, nominal operation may appear as a broad, noisy cloud. In quotient space, the nuisance directions are collapsed. The remaining geometry is closer to the latent physical-state

structure. This is why learning a representation that is invariant to nuisance transformations can reveal compact functional manifolds.

This mechanism is illustrated conceptually in Figure 1. The quotient projection does not create physical structure *ex nihilo*; it removes directions of variability that are irrelevant to the latent state. What remains is a lower-complexity geometry in which nominal and non-nominal regimes can be represented as functional manifolds.

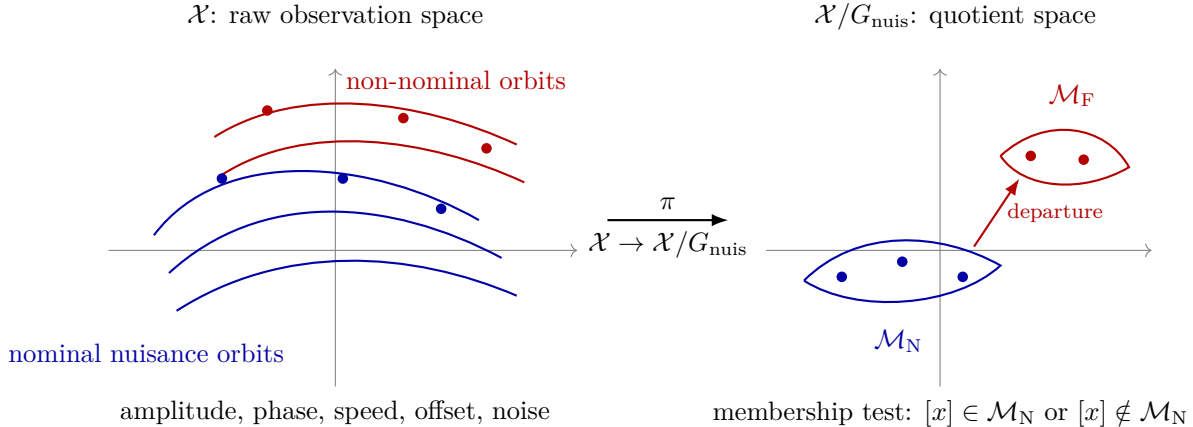


Figure 1: Conceptual illustration of functional manifolds as quotient geometry. In raw observation space, nuisance transformations generate extended orbits of observations that preserve the same latent physical state. The quotient projection π collapses these orbits into equivalence classes, revealing compact nominal and non-nominal functional manifolds.

This also clarifies the relation between anomaly detection and classification. A classifier learns a partition of labeled observations. A quotient-space diagnostic operator learns a representation in which physically equivalent observations are identified. Operational labels may then be attached to the recovered structures, but the structures themselves precede the labels.

6 One Computational Approximation

The previous sections establish the mathematical object. We now introduce one computational approximation.

The approximation has four components:

- (1) Monte Carlo generation of nominal trajectories over nuisance variables;
- (2) self-supervised representation learning on nominal trajectories only;
- (3) reconstruction error as an inverse-problem residual;
- (4) per-asset adaptive threshold estimation.

This section should be interpreted as one realization of the quotient-space formulation, not as the formulation itself. The same mathematical framework could be implemented with alternative representation operators, residual functions, or domain simulators.

6.1 Computational Strategy

Let θ_N denote a nominal latent physical state. The computational objective is to approximate the nominal quotient manifold \mathcal{M}_N without learning the full forward model F . Instead of identifying all physical parameters and nuisance variables explicitly, we generate nominal observations under nuisance variability and train a representation operator that is stable on this nominal orbit.

Synthetic nominal samples are generated as

$$x_i = F(\theta_N, \eta_i), \quad i = 1, \dots, N, \quad (9)$$

where each $\eta_i \in \mathcal{E}$ represents a nuisance configuration. These samples approximate the orbit of the nominal state under nuisance transformations.

The purpose of Monte Carlo generation is not to model every possible non-nominal condition. The purpose is to sample the nominal nuisance orbit densely enough that a representation learner can approximate the quotient geometry of nominal operation. Non-nominal regimes are then detected by their failure to belong to the recovered nominal structure.

6.2 Controlled Bearing Instantiation

The controlled validation uses a synthetic bearing signal emulator designed to expose the geometric structure of the proposed formulation rather than to reproduce a specific industrial asset. The emulator should not be interpreted as a high-fidelity digital twin of a particular bearing assembly. Its purpose is to generate regimes with known latent structure, controlled nuisance variability, and controlled non-nominal departures. This makes it suitable for testing whether the proposed quotient-space approximation can recover membership structure under known ground-truth conditions.

In this instantiation, the generic observation $x \in \mathcal{X}$ is a vibration window. Each window has length $n = 1200$, corresponding to 0.1 seconds sampled at $f_s = 12,000$ Hz. The nominal shaft frequency is 30 Hz.

The nominal bearing signal is generated as a finite harmonic expansion of the shaft frequency:

$$x_N(t) = \sum_{k=1}^5 a \rho^{k-1} \sin(2\pi k f_0 s t + \phi/k),$$

where a is an amplitude factor, $\rho = 0.55$ is the harmonic decay, $f_0 = 30$ Hz is the nominal shaft frequency, s is a speed-scaling factor, and ϕ is a phase perturbation.

The nuisance variables are sampled by Monte Carlo simulation:

$$a \sim U(0.85, 1.15), \quad \phi \sim U(-0.3, 0.3), \quad s \sim U(0.95, 1.05),$$

together with additive offset, additive noise, and a small local dropout perturbation. These variables modify the observed waveform without changing the intended physical regime. They therefore instantiate the nuisance orbit discussed in the quotient-space formulation.

Non-nominal regimes are generated by adding impulsive components to the nominal base signal. The outer-race fault uses an impulse train near 108s Hz, the inner-race fault uses an impulse train near 162s Hz, and the ball-fault regime uses an impulse train near 72s Hz with a low-frequency amplitude modulation. These non-nominal regimes are not used to train the representation operator; they are reserved for evaluating whether departure from the nominal quotient manifold can be detected and subsequently classified.

The emulator therefore provides three ingredients required by the validation: a nominal latent regime, nuisance transformations that preserve that regime, and non-nominal regimes that depart

from it through structured impulsive components. This is sufficient for testing the proposed inverse-recovery mechanism under controlled conditions, although it does not constitute field validation.

The controlled bearing instantiation is situated within the broader literature on machinery diagnostics, condition-based maintenance, and bearing vibration analysis [10, 11, 12, 13, 14].




Figure 2: Controlled bearing signal instantiation. The nominal signal and the non-nominal fault regimes are generated from distinct physical templates, while nuisance transformations modify the observed waveform without changing the intended latent regime.

6.3 Representation Operator

Let

$$E : \mathcal{X} \rightarrow \mathcal{Z}$$

be an encoder and

$$D : \mathcal{Z} \rightarrow \mathcal{X}$$

a decoder. In the present implementation, E and D are realised by a Conv1D autoencoder trained exclusively on nominal synthetic windows. Autoencoders have long been used as nonlinear representation learners [7] and as reconstruction-based anomaly detectors [8].

The learned reconstruction is

$$\hat{x} = D(E(x)). \tag{10}$$

The reconstruction residual is

$$r(x) = \frac{1}{n} \|x - \hat{x}\|_2^2. \tag{11}$$

Assumption 1 (Residual–distance proxy). *The reconstruction residual $r(x)$ is treated as a computable proxy, monotonically related to a notion of distance to the nominal quotient manifold:*

$$r(x) \approx d([x], \mathcal{M}_N). \tag{12}$$

This relation is not derived from the autoencoder’s training objective: minimizing mean-squared reconstruction error does not, by itself, guarantee that r is isometric, or even monotonic, with respect to any particular metric on the quotient space. Assumption 1 should be read as an operational hypothesis, evaluated empirically in Section 7, not as a property established by construction.

This interpretation is important. The contribution is not the use of mean-squared error as such. The contribution is the use of a residual as an operational test for membership in the nominal quotient manifold. In this implementation, the residual is a reconstruction error; in another implementation, it could be replaced by a latent likelihood, an energy score, a predictive error, or a physics-informed residual.

The autoencoder used in the controlled experiment is a one-dimensional convolutional reconstruction model operating directly on fixed-length vibration windows. The encoder applies successive

temporal convolution and compression operations to map each window to a finite-dimensional latent representation. The decoder mirrors this operation through upsampling and convolutional reconstruction layers. The model is trained only on nominal synthetic windows by minimizing mean-squared reconstruction error. No non-nominal samples are used during training.

Table 1: Computational role of the autoencoder in the controlled experiment.

Component	Role
Input	Fixed-length vibration window $x \in \mathcal{X}$
Encoder E	Maps x to latent representation $z \in \mathcal{Z}$
Decoder D	Reconstructs $\hat{x} = D(E(x))$
Training data	Nominal synthetic windows only
Loss	Mean-squared reconstruction error
Diagnostic score	Residual $r(x) = n^{-1} \ x - \hat{x}\ _2^2$

The architectural choice is intentionally not part of the mathematical formulation. It is used here as a simple self-supervised operator capable of learning a compact reconstruction model for the nominal nuisance orbit. Other representation operators could replace it without changing the quotient-space formulation.

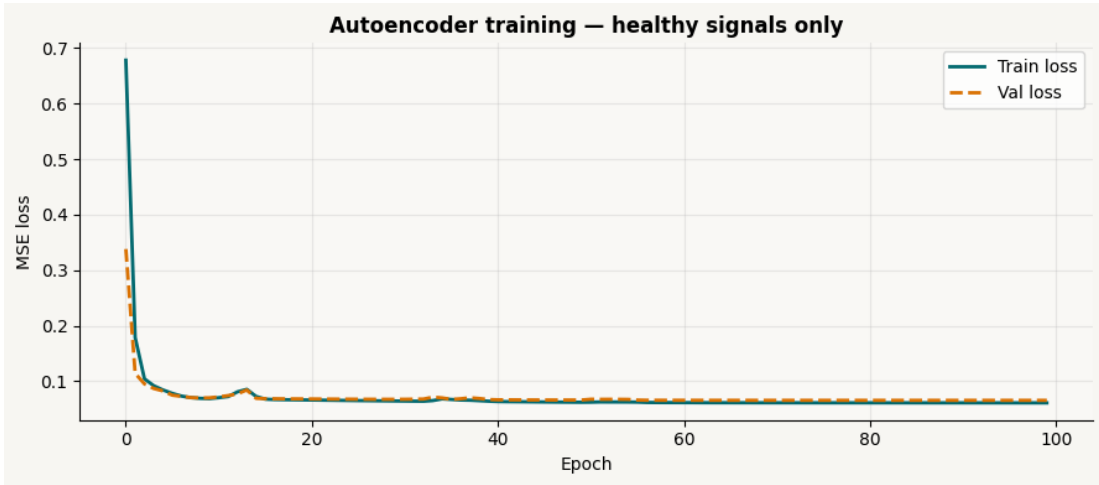


Figure 3: Self-supervised learning of a nominal reconstruction operator from synthetic nuisance-orbit samples. The learned operator is used as one computational approximation of the nominal quotient geometry.

6.4 Monte Carlo Calibration

A Monte Carlo reference threshold is estimated from residuals on held-out nominal synthetic samples. Let

$$\mu_{\text{MC}} = \mathbb{E}_{\text{MC}}[r(x)], \quad \sigma_{\text{MC}} = \text{Std}_{\text{MC}}[r(x)].$$

More generally, let $q \in (0, 1)$ denote an operational quantile controlling the trade-off between sensitivity and false-alarm tolerance. A parametric calibration can be written as

$$\tau_{\text{MC}}(q) = \mu_{\text{MC}} + z_q \sigma_{\text{MC}}, \tag{13}$$

where z_q is the corresponding standard-normal quantile or an empirically estimated equivalent. Alternatively, $\tau_{\text{MC}}(q)$ can be estimated directly as the empirical q -quantile of the nominal residual distribution.

The value of q is not part of the mathematical formulation. It is an operational risk parameter. Lower values increase sensitivity but may increase false alarms; higher values reduce false alarms but may delay detection of weak or incipient deviations.

In the synthetic experiment reported below, q is treated as a calibration parameter. The resulting $\tau_{\text{MC}}(q)$ defines a reference boundary for the learned nominal structure.

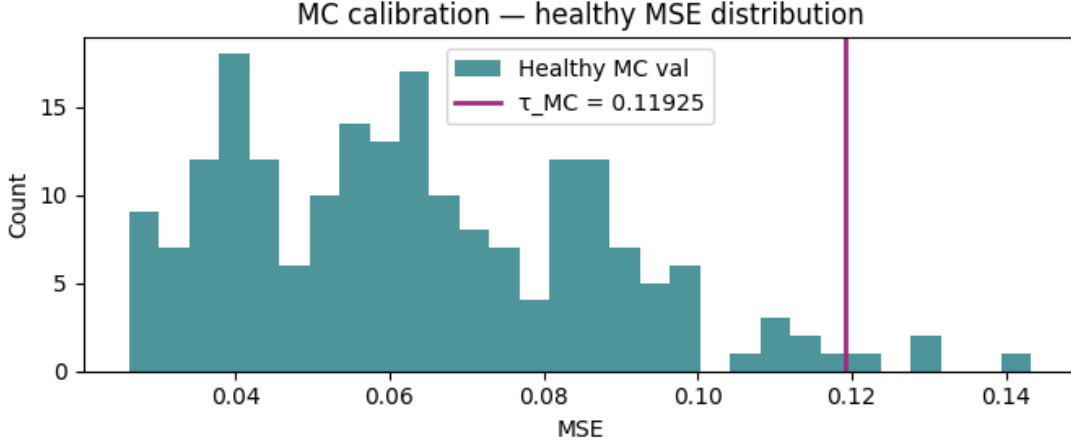


Figure 4: Monte Carlo calibration of the nominal residual distribution. The threshold defines a synthetic reference boundary for the learned nominal quotient structure.

6.5 Per-Asset Adaptive Threshold

A global threshold is insufficient when multiple physical assets have different local baselines. Even if they share the same representation operator, each asset may have a distinct residual scale due to installation, operating conditions, sensor coupling, or local environmental effects.

For each asset b , an adaptive threshold is estimated from a sliding buffer of recent non-anomalous residuals:

$$\tau_t^{(b)}(q) = \mu_t^{(b)} + z_q \sigma_t^{(b)}. \quad (14)$$

Here $\mu_t^{(b)}$ and $\sigma_t^{(b)}$ are computed locally for asset b . The scalar z_q preserves the statistical meaning of the selected operational risk level, while the local mean and standard deviation adapt the threshold to each physical asset.

The anomaly decision is

$$a_t^{(b)} = \begin{cases} 1, & r(x_t^{(b)}) > \tau_t^{(b)}(q), \\ 0, & r(x_t^{(b)}) \leq \tau_t^{(b)}(q). \end{cases} \quad (15)$$

This separates the representation problem from the operational decision problem. The representation is shared across assets. The decision boundary is local. In the bearing experiment, each physical asset b corresponds to one simulated bearing; in the general formulation, b may represent any monitored component, subsystem, or sensorized physical unit.

7 Synthetic Validation

The objective of the synthetic experiment is not to prove the quotient-space formulation. Proposition 2 is a structural statement about diagnostic geometry under nuisance-preserving transformations. The experiment addresses a different question: whether one computational approximation can recover a residual that behaves as a useful proxy for distance to the nominal quotient manifold.

The validation has four purposes:

- (i) verify that the simulated nuisance orbits exhibit saturation toward bounded functional regimes;
- (ii) verify that a representation trained only on nominal trajectories produces low residuals on nominal observations;
- (iii) verify that non-nominal regimes produce systematically larger residuals without being used during training;
- (iv) verify that local adaptive thresholds convert the residual into stable per-asset membership decisions.

The controlled instantiation uses four simultaneous synthetic bearings. Each bearing is treated as one physical asset:

- Asset A: nominal throughout the sequence;
- Asset B: outer-race fault onset;
- Asset C: inner-race fault onset;
- Asset D: ball-fault onset.

The representation operator is trained only on nominal synthetic data. Non-nominal regimes are used only for evaluation. The calibration quantile q is treated as an operational risk parameter, as defined in Equation (13). In the experiment reported here, adaptive thresholds are estimated using a sliding window of 200 samples unless otherwise stated.

7.1 Blueprint Saturation

The term *blueprint saturation* follows the companion technical report on self-supervision, manifold saturation, and emergent operational classes [15]. In that context, a blueprint denotes the finite sampled approximation of a functional regime, and saturation refers to the stabilization of geometric descriptors as the number of samples increases. Here we reuse the same diagnostic to test whether Monte Carlo samples from each simulated regime approach a bounded functional structure.

Before evaluating anomaly decisions, we first examine whether the simulated regimes exhibit blueprint saturation in the sense of [15]. Figure 5 reports a blueprint-saturation diagnostic: the p98 manifold radius in a two-dimensional PCA projection as a function of the number of Monte Carlo samples.

The purpose of this diagnostic is not to define the quotient space through PCA. Rather, it tests whether the sampled nuisance orbits approach bounded functional structures as the number of samples increases. The stabilization of the p98 radius across regimes suggests that the simulated observations are not merely isolated waveform examples, but samples from compact functional regimes. This is consistent with the quotient-space interpretation: nuisance variability expands

the raw observation cloud, but the orbit generated by a fixed latent regime remains geometrically bounded.




Figure 5: Blueprint saturation of the simulated functional regimes, following the manifold-saturation diagnostic introduced in [15]. For each regime, the p98 manifold radius is estimated in a two-dimensional PCA projection as the number of Monte Carlo samples increases. The stabilization of the radius indicates that the sampled nuisance orbit approaches a bounded functional structure. This supports the interpretation that the simulator generates compact quotient-space regimes rather than isolated waveform examples.

7.2 Residual Separation

Figure 6 shows the residual distributions obtained after training the representation operator on nominal trajectories only. The important point is not that the residual is mean-squared error in this implementation. The important point is that the residual behaves as a proxy for membership in the learned nominal structure. Nominal observations concentrate at low residual values, while non-nominal regimes shift to higher residual regions.



Figure 6: Residual distributions produced by the learned representation operator. The operator is trained only on nominal trajectories. The separation between nominal and non-nominal regimes indicates that the learned residual is diagnostically meaningful as a proxy for distance to the nominal quotient manifold.

7.3 Adaptive Window Robustness

The adaptive boundary depends on a sliding window of recent residuals. To test whether the result is sensitive to this choice, we evaluate false-alarm and detection rates across several window sizes. Figure 7 shows that the detection rate remains high while the false-alarm rate remains low across a broad range of window sizes. Thus, the adaptive decision rule is not an artefact of a single hand-tuned window length.

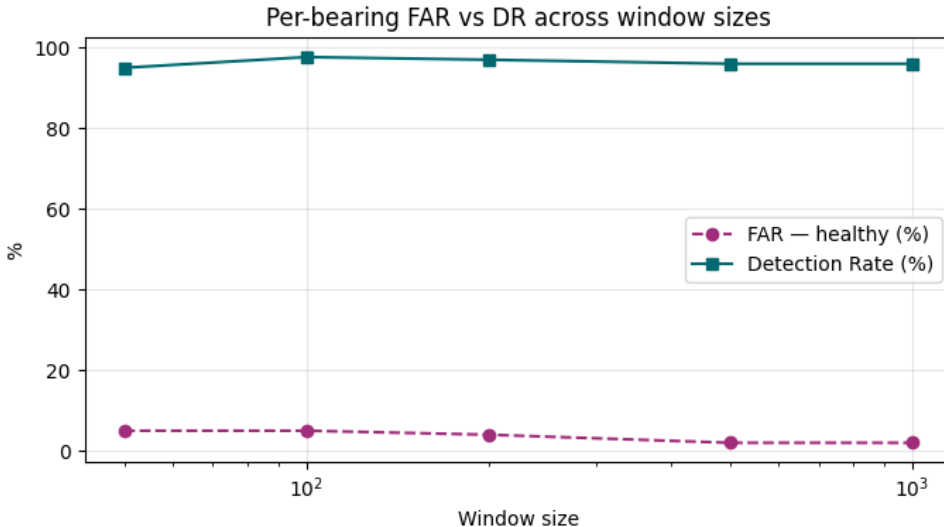


Figure 7: Sensitivity of the adaptive decision rule to the sliding-window size. The detection rate remains high across a wide range of window sizes, while the false-alarm rate remains low. This indicates that the per-asset adaptive boundary is not an artefact of a single window choice.

7.4 Per-Asset Detection Results

Table 2 reports the resulting membership decisions after applying the adaptive threshold. FAR denotes the false-alarm rate during nominal operation. DR denotes the detection rate after non-nominal onset.

Table 2: Synthetic per-asset detection metrics. The representation operator is shared across all assets, while the adaptive decision boundary is estimated locally for each asset.

Asset	Scenario	FAR during nominal period	DR after onset
A	Nominal throughout	1.5%	N/A
B	Outer-race fault onset	2.0%	100%
C	Inner-race fault onset	2.0%	100%
D	Ball-fault onset	2.0%	86%

The results show that the learned residual remains stable for the asset that remains nominal and increases after non-nominal onset for the remaining assets. The outer-race and inner-race regimes are detected completely in the evaluation interval. The ball-fault regime is detected less strongly, with an 86% detection rate, which is consistent with a weaker or less separable deviation from the learned nominal structure.

These numbers should be interpreted as preliminary evidence of operational feasibility, not as a claim of benchmark-level performance. The scientific claim of this note is the quotient-space inverse formulation. The synthetic experiment demonstrates that one computational realization of that formulation can produce useful membership decisions.

8 Interpretation

The experiment supports four claims.

First, the simulated regimes exhibit blueprint saturation in the sense of the companion report [15]: as the number of Monte Carlo samples increases, the p98 radius stabilizes, suggesting that the nuisance orbits approach bounded functional structures. This supports the use of the synthetic emulator as a controlled setting for studying quotient-space recovery.

Second, a representation trained only on nominal nuisance-orbit samples can produce residuals that distinguish nominal and non-nominal regimes. This is consistent with the interpretation of the residual as a proxy for distance to the nominal quotient manifold.

Third, the adaptive boundary should be estimated locally for each physical asset. A shared representation can be global, but the decision boundary must account for local residual baselines. This reflects the distinction between learning a common quotient geometry and estimating an asset-local operational boundary.

Fourth, non-nominal regimes need not be learned as semantic labels. In the controlled bearing experiment, outer-race, inner-race, and ball-fault regimes are detected by departure from the nominal structure. This is the inverse-problem view: first recover membership in the nominal functional structure; assign semantics later.

9 Discussion

The proposed formulation differs from standard supervised fault classification in three ways.

Diagnosis precedes classification. The primary decision is whether an observation belongs to the nominal quotient manifold. Classification of the type of non-nominal condition is a subsequent semantic step.

The representation is geometric. The goal is not to memorize labels, but to collapse nuisance orbits while preserving latent physical-state information. This is why the quotient space is the natural object.

The threshold is asset-local. A physical asset may have a local residual baseline even when the representation operator is shared. Per-asset adaptive thresholding estimates this local boundary without retraining the representation.

This work also clarifies the relationship between quotient spaces and functional manifolds. The quotient-space formulation provides the geometric mechanism through which functional manifolds emerge. They are connected regions of the quotient space after nuisance directions have been collapsed. This gives a mathematical origin to the compact structures observed in self-supervised diagnostic representations.

10 Limitations

This work should not be read as a complete industrial validation.

The current experiment is synthetic. It demonstrates the feasibility of one computational approximation under controlled conditions, but it does not establish robustness under field conditions.

The synthetic emulator is designed to test the geometry of the proposed formulation under controlled ground-truth conditions; it is not intended to reproduce the full mechanical, tribological, or environmental complexity of field bearing data.

The nuisance transformation family is approximate. As anticipated in Section 4, the formal quotient construction uses an honest group action G_{nuis} , while any empirical instantiation—synthetic or field-collected—can only sample a compact, physically admissible subset of that action. In real machines, nuisance effects may not even compose exactly within \mathcal{T} itself, so the closure that defines G_{nuis} should be read as a mathematical idealization of physical nuisance mechanisms rather than as an empirically verified algebraic property. Nevertheless, the quotient-space formulation remains useful as an organizing principle when transformations approximately preserve the physical state.

The autoencoder is only one possible approximation. Future work should compare reconstruction-based residuals with alternative representation operators and distance functions.

Finally, fault classification is not fully addressed here. The present work focuses on anomaly detection as failure of membership in the nominal quotient manifold. Fault classification can be formulated as identifying the destination manifold after departure from the nominal structure.

11 Conclusion

This work establishes a mathematical formulation of diagnostic monitoring as inverse recovery over quotient-space geometry. The formulation begins with the physical observation model, identifies nuisance transformations that preserve the latent physical state, and shows that diagnosis naturally factors through the quotient space induced by those transformations.

The proposed learning architecture constitutes one practical approximation of that formulation. Monte Carlo simulation samples the nominal nuisance orbit, a self-supervised representation operator approximates the nominal quotient geometry, and a per-asset adaptive threshold estimates the local decision boundary. Synthetic experiments provide preliminary empirical evidence that quotient-space recovery is feasible under controlled operating conditions.

The main contribution is therefore not a particular autoencoder architecture or a bearing-specific detector. The main contribution is the geometric formulation: anomaly detection can be posed as a membership problem on quotient-space functional manifolds induced by nuisance-invariant observations.

Future work will investigate field validation, alternative representation operators, fault-manifold classification after anomaly detection, and transfer to industrial rotating machinery datasets.

Accompanying Source Material

The synthetic bearing emulator, Monte Carlo generation scripts, Conv1D autoencoder training notebook, adaptive-threshold calibration code, and figure-generation scripts are deposited as accompanying source material for this work. The source material is provided to make the controlled validation reproducible and to distinguish the general quotient-space formulation from the particular computational approximation used in the experiment.

References

- [1] J. Hadamard. *Lectures on Cauchy's Problem in Linear Partial Differential Equations*. Yale University Press, 1923.
- [2] A. N. Tikhonov and V. Y. Arsenin. *Solutions of Ill-Posed Problems*. Winston, 1977.
- [3] V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), Article 15, 2009.
- [4] B. Schölkopf, R. C. Williamson, A. J. Smola, J. Shawe-Taylor, and J. C. Platt. Support vector method for novelty detection. In *Advances in Neural Information Processing Systems*, 1999.
- [5] C. M. Bishop. Novelty detection and neural network validation. *IEE Proceedings — Vision, Image and Signal Processing*, 141(4), 1994.
- [6] T. Kohonen. Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43, 59–69, 1982.
- [7] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786), 504–507, 2006.
- [8] M. Sakurada and T. Yairi. Anomaly detection using autoencoders with nonlinear dimensionality reduction. In *Proceedings of the 2nd Workshop on Machine Learning for Sensory Data Analysis*, 2014.
- [9] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv:2104.13478*, 2021.
- [10] A. K. S. Jardine, D. Lin, and D. Banjevic. A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing*, 20(7), 1483–1510, 2006.
- [11] R. B. Randall and J. Antoni. Rolling element bearing diagnostics—A tutorial. *Mechanical Systems and Signal Processing*, 25(2), 485–520, 2011.
- [12] W. A. Smith and R. B. Randall. Rolling element bearing diagnostics using the Case Western Reserve University data: A benchmark study. *Mechanical Systems and Signal Processing*, 64–65, 100–131, 2015.
- [13] R. Zhao, R. Yan, Z. Chen, K. Mao, P. Wang, and R. X. Gao. Deep learning and its applications to machine health monitoring: A survey. *Mechanical Systems and Signal Processing*, 115, 213–237, 2019.
- [14] Y. Lei, J. Lin, Z. He, and M. J. Zuo. A review on empirical mode decomposition in fault diagnosis of rotating machinery. *Mechanical Systems and Signal Processing*, 35(1–2), 108–126, 2013.
- [15] E. Di Santi. *Self-Supervision, Manifold Saturation, and Emergent Operational Classes: Learning from Scarce Observations as Inverse Recovery of Compact Functional Structure*. Technical report, University of Colorado Boulder, 2026.