Author: Jos Aben / TARH Energy
Mail: info@tarh-energy.com

Before You Build the Model

A Practical Framework for Assessing Predictive Feasibility in Industrial Systems. Predictive Feasibility Assessment (PFA)

Executive Reader Note

This document contains the full technical whitepaper for the Predictive Feasibility Assessment (PFA) framework.

A separate companion publication titled:

"**PFA Executive Guide Assessing Predictive Feasibility**"

provides a shorter operational interpretation intended for industrial deployment teams, reliability engineers, technical managers, and predictive maintenance environments.

While this technical whitepaper contains the full methodology, robustness analysis, datasets, and representational evaluation framework, the Executive Guide focuses on practical deployment implications, predictive feasibility assessment workflows, and operational decision-making.

Together, the two documents provide:
• a deployment-oriented industrial overview
and
• a detailed technical framework and validation layer.

Executive Summary

Industrial organizations increasingly invest in predictive AI, predictive maintenance, and time-series forecasting systems. Yet many predictive projects fail to achieve stable and generalizable deployment performance. [1–3]

In many cases, models initially appear promising, additional feature engineering is introduced, and increasingly complex AI architectures are deployed — yet deployment stability eventually collapses across assets or operating conditions.

This whitepaper proposes a practical explanation for this recurring pattern:
Many predictive projects fail not because the model itself is inadequate, but because the observable signal does not support stable and reproducible prediction.

The Predictive Feasibility Assessment (PFA) framework introduces a pre-model evaluation step designed to determine whether predictive modeling is structurally justified before substantial resources are invested.

Why This Matters for Industry

Predictive AI projects frequently fail after months of development because the signal itself never supported stable prediction.

The PFA framework is designed to identify this early — before organizations invest heavily in unstable models, repeated feature engineering, and failed deployment cycles.

The framework is not only diagnostic.
It is also designed to identify whether predictive structure can potentially be recovered through:

- improved signal representation
- spectral analysis
- observability improvements
- regime isolation
- physics-informed feature extraction

This enables organizations to:

- reduce failed AI development cycles
- identify viable predictive signals earlier
- improve deployment stability

• reduce predictive maintenance risk
• focus modeling effort where stable predictive structure actually exists

We are currently seeking industrial datasets and predictive maintenance use-cases for further real-world validation and feasibility assessment.

For selected early collaborations, assessments may be performed directly on operational datasets to evaluate:
• structural reproducibility
• deployment viability
• recoverable predictive structure
• operational predictive limits

Predictive Feasibility Overview

Regime	Consistency	Decision
NO-GO	< 0.3	Do not model
LIMITED	0.3 – 0.7	Investigate improvement
GO	> 0.7	Proceed with modeling

These thresholds should be interpreted as practical empirical working ranges rather than universal physical boundaries, and may vary depending on signal representation, operational conditions, and domain characteristics.

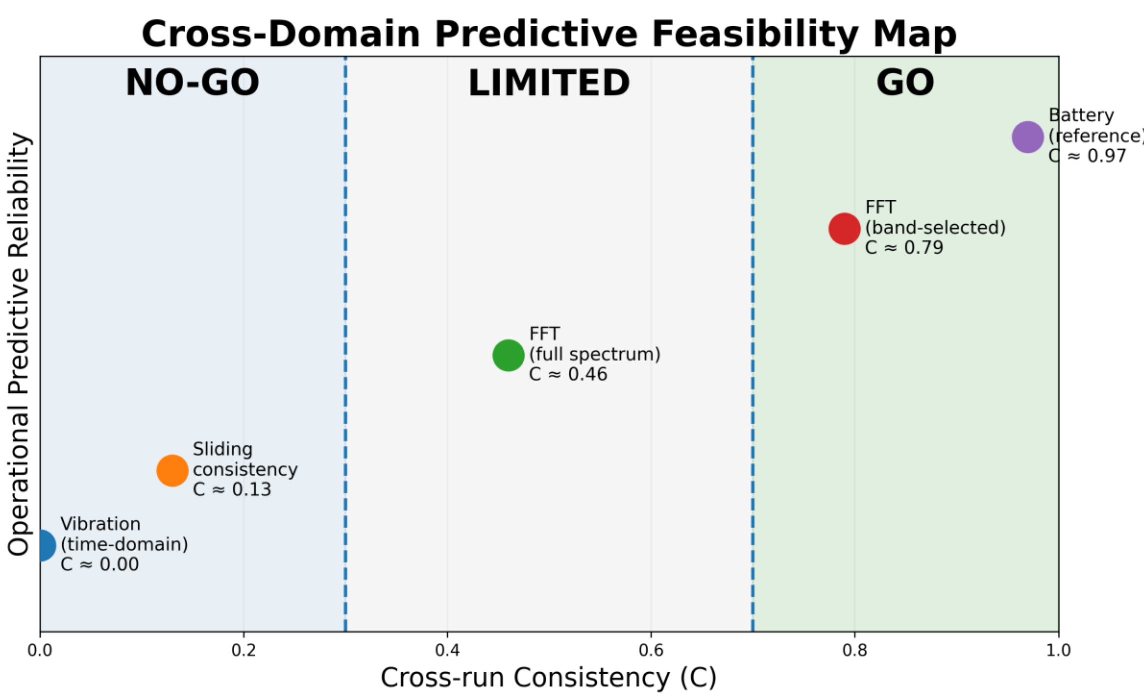

The Core Principle

Prediction requires not only local predictability, but reproducible structure across independent system realizations. [4,5] A signal may appear structured and locally predictable, yet still fail operationally if the structure does not reproduce across runs.

Representational Consistency Layers

The framework evaluates predictive feasibility across multiple representational layers:

Consistency Layer	Purpose
Observable consistency	Basic reproducibility across runs
Temporal consistency	Alignment despite timing variation
Spectral consistency	Shared dynamical modes and frequency behavior
Dynamical consistency	Underlying system evolution and state behavior

3. From Raw Signal to Predictive Feasibility

3.1 From Raw Signal to Predictive Structure

In many industrial systems, raw monitoring data initially appears noisy, unstable, and difficult to interpret. Despite this, predictive models are often applied directly to such signals, frequently resulting in unstable deployment behavior, poor generalization, repeated feature engineering cycles, and extended development effort.

A critical limitation in many predictive workflows is that signal suitability is rarely evaluated before model development begins.

As a result, organizations often invest substantial resources into 4odelling signals that never contained stable predictive structure in the first place.
The purpose of this section is to demonstrate how predictive feasibility emerges progressively through representation and analysis.
Rather than focusing on model architecture, the analysis begins with the signal itself:

- What does the raw signal actually contain?
- Does structure emerge under representation changes?
- Is that structure reproducible across runs?
- And most importantly:
- Does the signal contain stable predictive information at all?

The following examples are based entirely on real-world system data and illustrate the transformation from raw observable behavior to structured predictive assessment. Although the example shown here is derived from real IBM Quantum calibration telemetry, the same structural behavior is consistently observed across multiple industrial domains, including vibration monitoring systems, telemetry systems, and battery degradation datasets.

The purpose is therefore not domain-specific validation, but to illustrate a general property of signal-based prediction.

3.2 Pure Raw Data (Industry Input)

At the beginning of most industrial predictive workflows, monitoring signals are observed in their raw form.
These signals often contain:

- overlapping dynamics
- irregular fluctuations
- local transient behavior
- operational variability
- measurement noise

Despite visible activity and apparent complexity, it is not yet known whether the signal contains reproducible structure capable of supporting stable prediction.
In practice, this is the stage at which most predictive 5odelling projects begin.
However, at this stage, there is typically no explicit assessment of whether the signal itself is structurally predictive.

Figure 2 — Pure Raw Data (Industry Input)

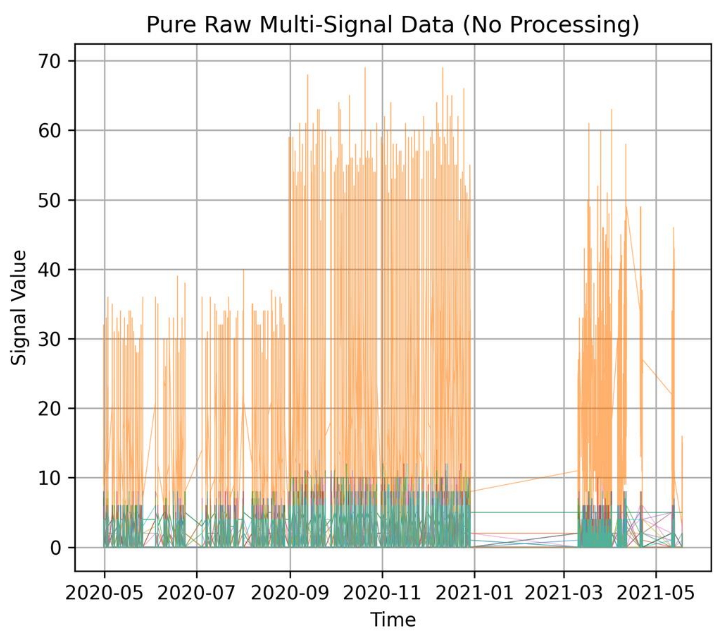

Figure 2 — Pure raw multi-signal data from real-world system telemetry.

This figure shows unprocessed monitoring signals plotted directly over time without filtering, aggregation, alignment, or transformation. Multiple overlapping trajectories produce a representation that appears highly active and structurally complex, yet no clear predictive mapping or reproducible instability structure is immediately visible.
This is the stage at which most predictive 6odelling efforts typically begin.
At this point, it is not yet known whether the signal contains structure capable of supporting reliable prediction.

3.3 Structured Representation (Initial Processing)

After basic preprocessing, the signal becomes more interpretable.
Typical preprocessing operations include:

- time alignment
- aggregation
- normalization
- smoothing
- signal restructuring

At this stage, broader system-level behavior may begin to emerge.
However, an important distinction remains:
The presence of visible structure does not imply that the structure is reproducible, stable, or meaningful for prediction.
This is a critical transition point.
Many predictive 6odelling projects incorrectly assume that increased interpretability automatically implies predictive feasibility.
In practice, this assumption frequently fails.

Figure 3 — Structured Representation (First Processing Step)

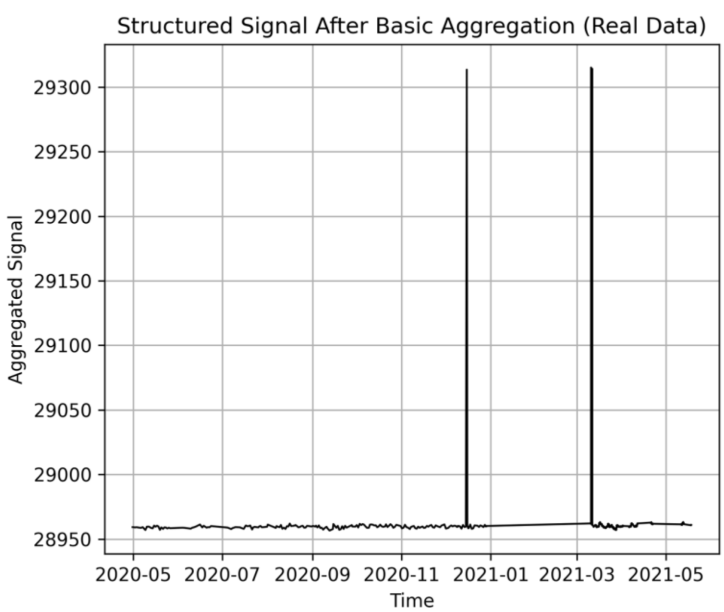

Figure 3 — Structured representation after basic preprocessing and aggregation.

The same dataset shown in Figure 1 is transformed into a more interpretable system-level representation using alignment and aggregation procedures. Broader behavioral structure becomes visible, including regime changes and long-term variation.

However, despite increased interpretability, it remains unclear whether this structure is:

- reproducible across runs
- stable over time
- or meaningful for prediction.

This distinction is essential.

Structure alone is not sufficient for predictive feasibility.

3.4 Extracting Predictive Structure

Once the signal has been transformed into a structured representation, analytical evaluation becomes possible.

At this stage, the goal is no longer merely to detect activity or variation.

Instead, the critical question becomes:

Does the signal contain reproducible structure associated with future system behavior?

To evaluate this, precursor behavior and instability-aligned representations can be analyzed.

This enables the identification of:

- instability precursors
- regime-dependent behavior
- variance evolution
- reproducible signal transitions

However, even detectable precursor structure does not automatically imply predictive feasibility.

A signal may exhibit measurable structure while still lacking reproducibility across runs.

Figure 4 — Pre-Instability Detection (Analysis Output)

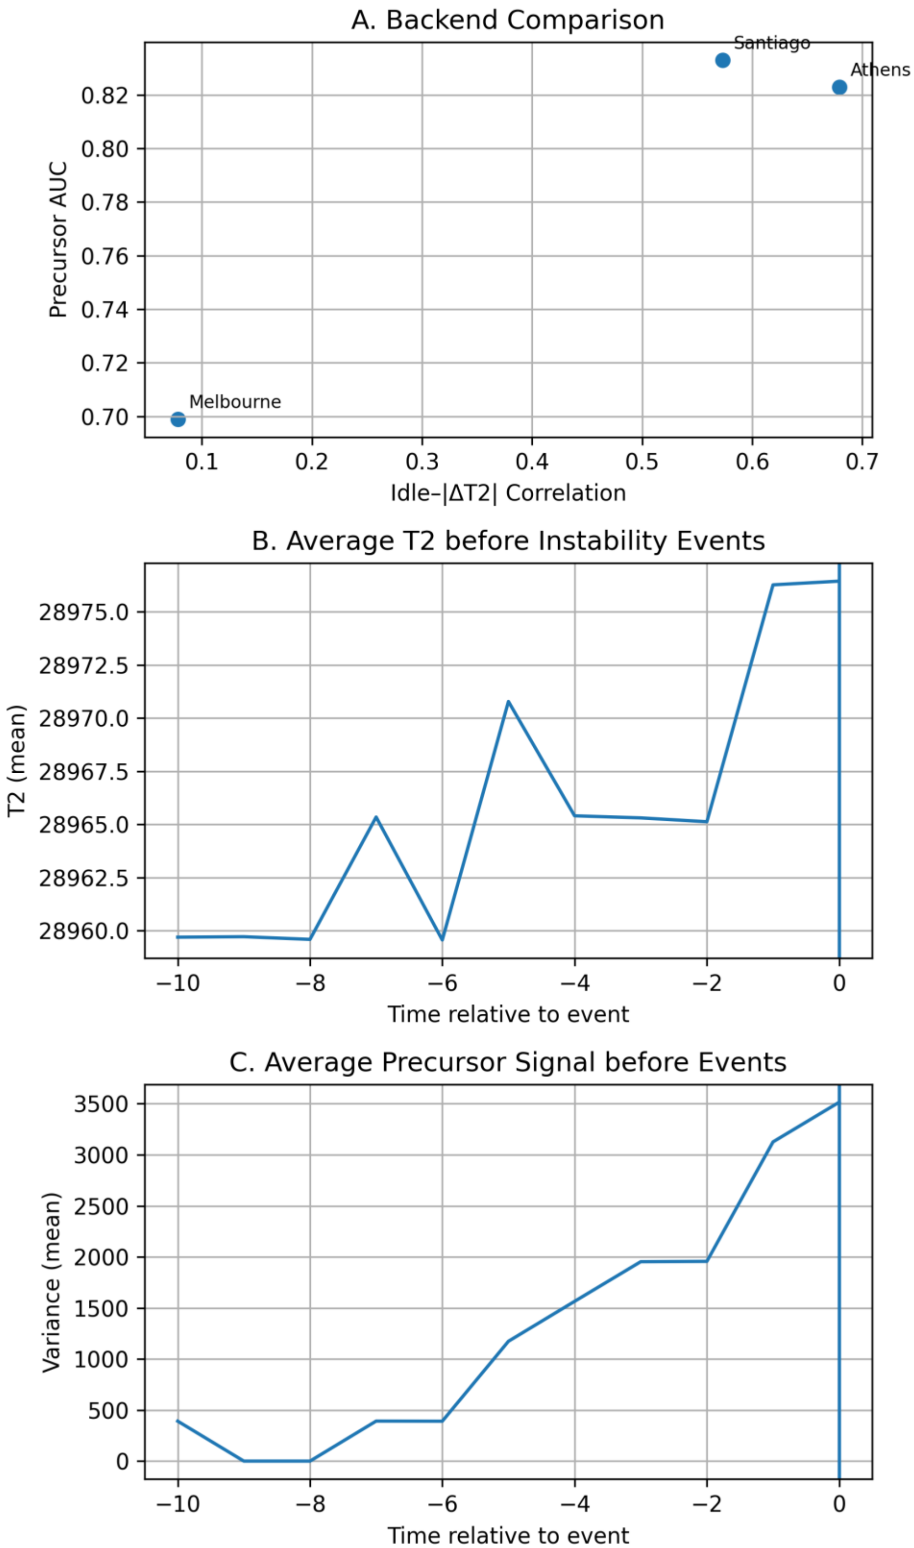

Figure 4 — Pre-instability signal detection using real-world telemetry data.

(5) Comparative precursor signal behavior across independent systems.
(B) Average signal evolution aligned around instability events.
(C) Average precursor dynamics prior to instability events.

The figure demonstrates that measurable precursor behavior emerges prior to instability events.

However, the presence of detectable structure alone does not guarantee predictive feasibility.
Predictive structure must also be:

- reproducible
- stable
- and consistently related to underlying system evolution.

3.5 From Detectability to Predictive Feasibility

The transformation shown across Figures 1–3 illustrates a critical principle:

Predictive 9odelling failure is often not caused by insufficient 9odelling capability.

Instead, failure frequently originates from the absence of stable, reproducible structure within the signal itself.

Many systems exhibit measurable structure and local predictability while still failing to support reliable operational prediction.
As a result, models may appear successful during development while failing to generalize across assets or operating conditions.
This enables a practical pre-model decision step:
• GO → stable reproducible predictive structure exists
• LIMITED → partial or condition-dependent structure exists
• NO-GO → predictive modeling is structurally unreliable.

3.6 Practical Industrial Implication

For industrial systems, the implications are direct:

- Raw monitoring data rarely reveals predictive feasibility directly
- Detectable structure does not imply predictive value
- Prediction error alone is insufficient to assess meaningful prediction
- Stable deployment requires reproducible signal structure across runs

This leads to a fundamental operational principle:

Prediction is not primarily a property of the model.

It is primarily a property of the signal.

Accordingly, predictive 10odelling should not begin with model selection.

It should begin with an explicit assessment of whether the signal supports prediction at all.

This defines the role of the Predictive Feasibility Assessment (PFA):

To determine whether predictive 10odelling is structurally justified before substantial resources are invested into model development, feature engineering, and deployment infrastructure.

Robustness Validation of Predictive Feasibility

To evaluate whether the observed lack of predictive feasibility was caused by methodological artifacts rather than intrinsic signal properties, multiple robustness analyses were performed across independent datasets and representations.

The objective was to determine whether predictive structure could be recovered through alternative analytical choices.

The following tests were performed:

• Alternative consistency metrics
(Pearson, Spearman, Kendall, mutual information)

• Multi-scale analysis
(different window sizes and temporal resolutions)

• Alternative run definitions
(chronological, random, overlapping segmentation)

• Signal preprocessing and smoothing
(raw, perturbed, smoothed representations)

• Lagged alignment analysis
(temporal shift compensation)

Across all evaluated datasets, the same pattern consistently emerged:

- prediction stability could vary significantly under preprocessing
- apparent local structure could emerge temporarily
- but cross-run consistency remained structurally constrained

This demonstrates that predictive failure is not primarily caused by:

- preprocessing choice
- segmentation strategy
- alignment artifacts
- or the selection of consistency metric

Instead, the limitation originates from the absence of stable reproducible structure within the signal itself.

Consistency Is Structurally Invariant

A critical observation emerges across all evaluated representations:

Prediction stability changes under preprocessing, but consistency remains structurally invariant.

Smoothing, filtering, normalization, perturbation, and feature transformations may significantly alter:
• local prediction error
• apparent signal structure
• short-term forecasting behavior

yet reproducible cross-run consistency frequently remains unchanged.

This leads to an important practical implication:

Prediction performance can often be manipulated through representation changes, while true predictive feasibility cannot.

Consistency therefore acts as a more reliable indicator of whether a signal contains meaningful predictive information.

This distinction is essential for industrial deployment, where preprocessing improvements may create the appearance of predictive structure without improving reproducible operational behavior.

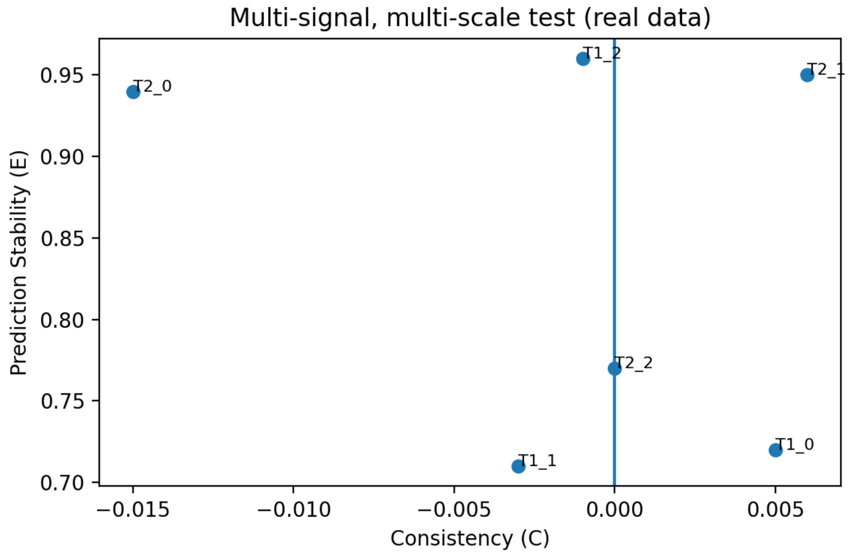

Figure 5 — Multi-signal and multi-scale consistency analysis on real telemetry data. Each point represents a signal evaluated under a specific time scale. While prediction stability € varies across signals, consistency € remains near zero for all cases. No signal exhibits both high consistency and high prediction stability. The vertical line at C = 0 highlights the absence of reproducible structure across all evaluated conditions.

Key Industrial Observation

Signals may appear:

- structured
- locally predictable
- or operationally active

while still failing to support stable predictive deployment.

This distinction is critical because many industrial AI projects continue optimizing models long after structural predictive feasibility has already collapsed.

While robustness analysis demonstrates that predictive failure is not caused by preprocessing or metric selection alone, many systems still exhibit local predictability without preserving reproducible structure across runs.
This distinction is critical for understanding why apparently stable short-term forecasting behavior may still fail operationally.

Prediction Stability Is Not Predictive Feasibility

A common assumption in predictive 12odelling is that stable short-term prediction implies meaningful predictive information.

In practice, however, this assumption frequently fails.

Across multiple real-world datasets, signals were observed that exhibited:

- high local predictability
- low prediction error
- stable short-term forecasting behavior

while simultaneously lacking reproducible structure across runs.

This reveals a critical distinction:
A signal may be statistically predictable,
while still failing to support reliable predictive deployment.
The limitation is not the ability to forecast signal values locally.
The limitation is the absence of stable reproducible structure associated with underlying system behavior.
This distinction explains a common industrial failure pattern:

- models initially appear successful
- local validation metrics look promising
- deployment begins
- but predictions fail to generalize across assets, conditions, or operating regimes

As a result, teams often continue:

- feature engineering
- retraining
- model redesign
- architecture optimization

without recognizing that predictive feasibility had already collapsed at the signal level.

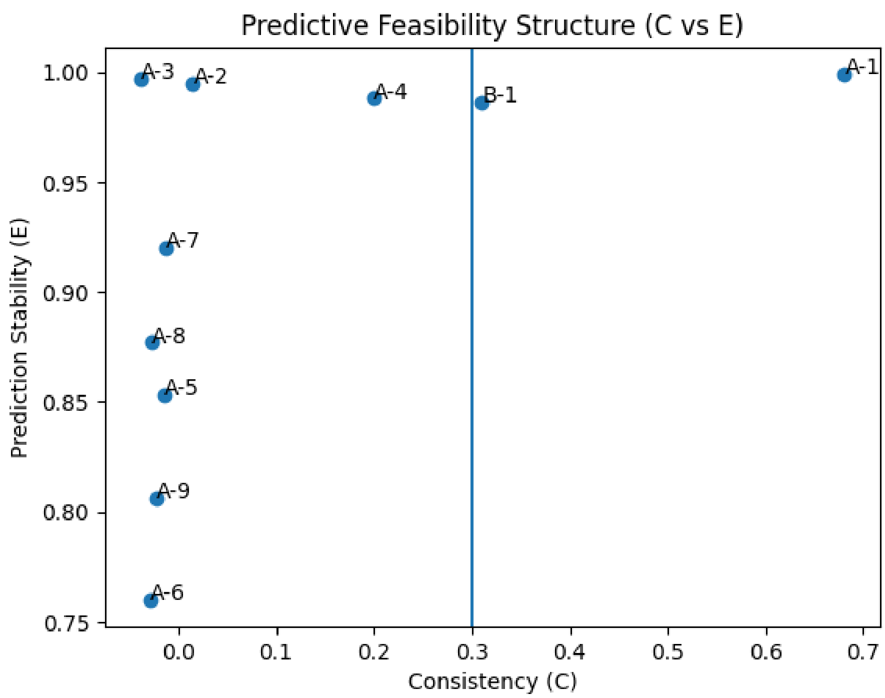

Figure 6 — Prediction stability versus structural consistency

Each point represents a real-world telemetry signal evaluated using cross-run consistency € and prediction stability €.
While prediction stability remains relatively high across many signals, reproducible structural consistency varies significantly.
Signals with low consistency frequently appear statistically predictable, yet fail to support reliable or actionable prediction.
This demonstrates that prediction error alone is insufficient to determine predictive feasibility.

Stable deployment requires reproducible signal structure across independent runs.
This distinction has major industrial implications.
Many predictive AI systems are optimized using:

- short-term prediction accuracy
- local forecasting performance
- benchmark validation metrics [4,5,7]

However, these metrics may fail to detect whether the signal contains reproducible predictive information at all.

As a result:
prediction performance may appear stable, while deployment feasibility remains fundamentally unstable.
The Predictive Feasibility Assessment (PFA) framework therefore distinguishes between locally stable prediction behavior and structurally reproducible predictive behavior.

Predictable but Non-Informative Signals

Not all predictable signals contain meaningful predictive information.

In many systems, signals may exhibit:
• smooth local behavior
• low prediction error
• stable autoregressive forecasting

while still failing to encode future system evolution in a reproducible way.

This defines a critical regime:
Predictable but non-informative.
In this regime:
• signals appear statistically predictable
• models may achieve strong validation metrics
• short-term forecasting appears stable

yet:
• predictions do not generalize

• deployment behavior collapses
• the observable signal does not reliably represent the underlying system state.

This distinction explains why many predictive AI systems appear successful during development but fail under real operational conditions.

Prediction error alone is therefore insufficient to determine predictive feasibility.

This enables organizations to identify signals that:

- appear predictable locally
- but fail structurally under deployment conditions.

From an operational perspective, this distinction is critical.
Improving model performance does not necessarily improve predictive feasibility.
If reproducible structure is absent,
continued model optimization is unlikely to produce stable deployment behavior.

This explains why many predictive AI initiatives continue for months despite never converging toward reliable operational prediction.
The limitation is often not the model.
The limitation is the signal itself.

The distinction between apparent predictability and reproducible predictive structure becomes especially visible in real-world vibration systems.
The following example illustrates how signals may exhibit strong visible activity and high local structure while still failing to support stable predictive deployment.

Representative Example — NO-GO (Time-Domain Vibration)

Raw vibration signals often exhibit strong apparent structure and high signal energy.
However, reproducibility across independent runs may collapse completely.

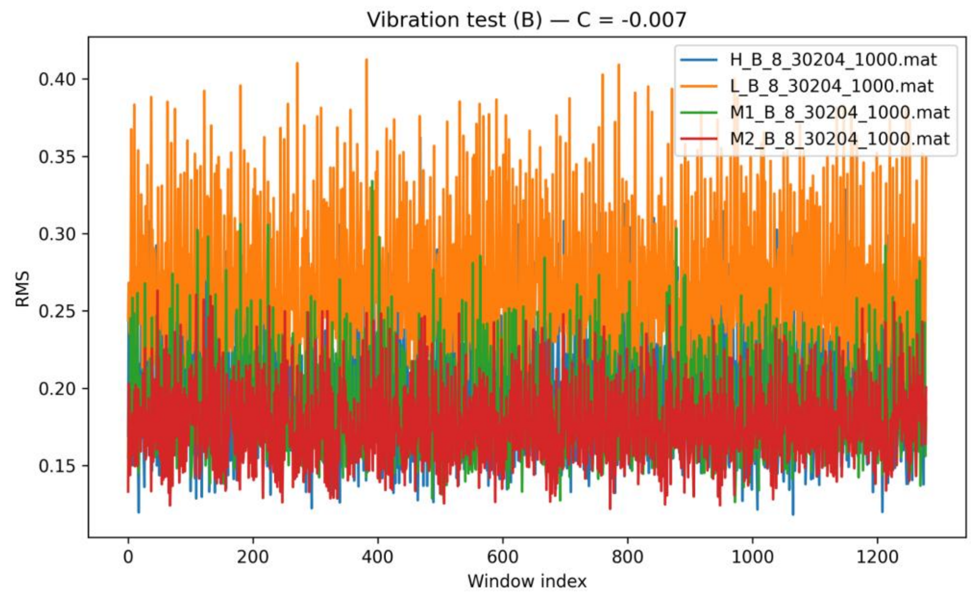

Figure 7. Despite visible structure and high signal energy, cross-run consistency collapses (C ≈ 0), indicating absence of a stable mapping across runs.

Absence of Stable Regimes

Sliding consistency analysis demonstrates that the signal never enters a stable predictive regime over time.

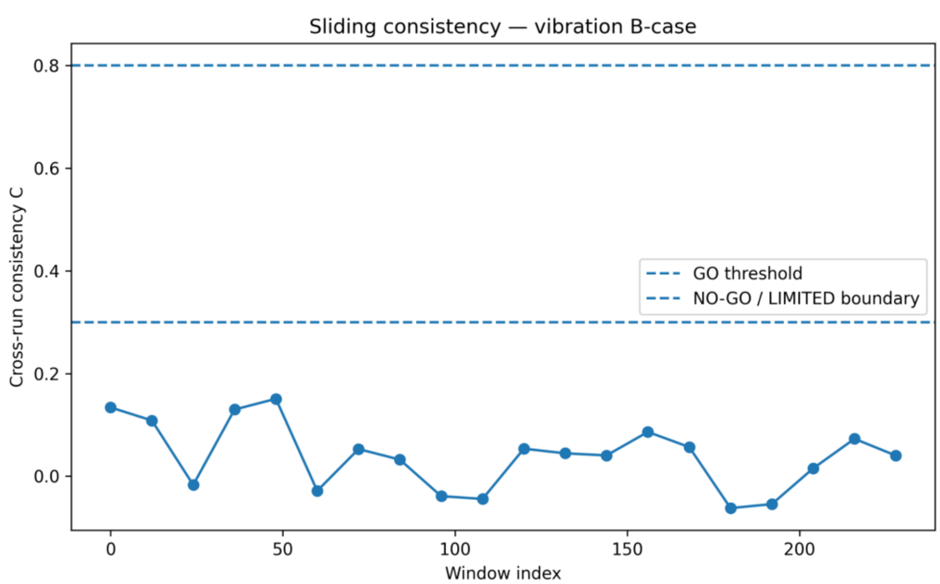

Figure 8. Sliding consistency remains near zero across all time windows, confirming that instability is structural rather than local or transient.

Transition Toward Recoverable Structure

Although the time-domain signal behaves as a clear NO-GO case, the key question becomes whether predictive structure exists in another representation aligned more closely with the underlying system physics.

LIMITED — Frequency-Domain Recovery

Moving into the frequency domain reveals partial reproducibility. The signal contains recoverable structure, but the mapping remains unstable.

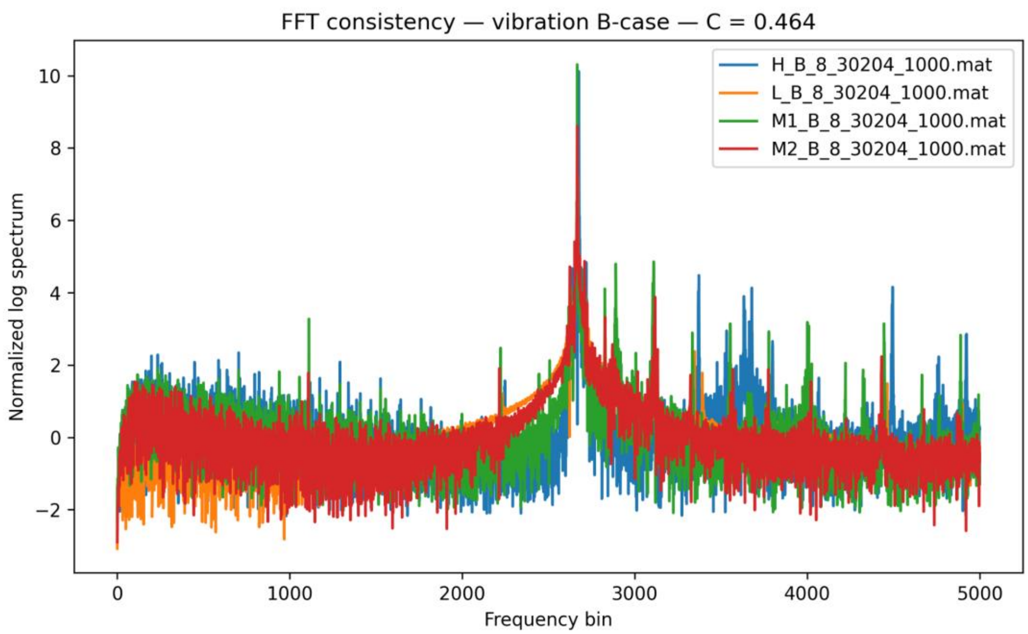

Figure 9. Full-spectrum FFT reveals partial cross-run consistency (C ≈ 0.46), indicating that predictive structure exists but remains unstable.

Near-GO — Physically Aligned Frequency Band

Restricting analysis to a physically meaningful frequency band significantly increases cross-run alignment, revealing near-stable predictive structure.

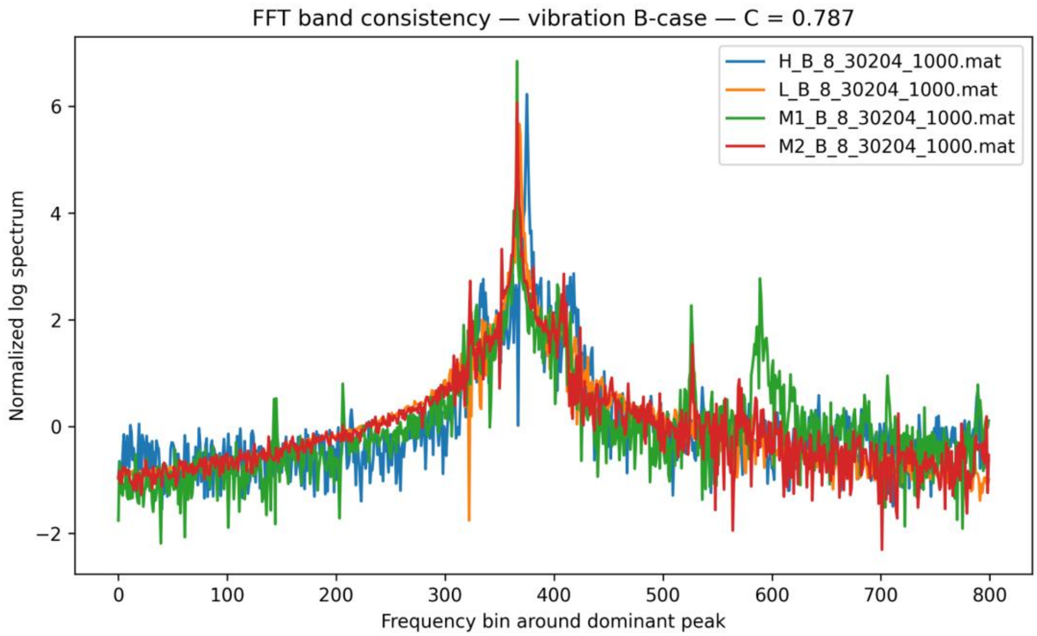

Figure 10. Focusing on a narrow frequency band around the dominant spectral peak increases cross-run consistency to C ≈ 0.79, revealing a nearly stable mapping.

Recoverable Structure vs Temporary Alignment

Not all apparent improvements in predictive behavior correspond to genuine predictive structure.

In several evaluated systems, localized increases in consistency emerged under specific:
- operating conditions
- signal representations
- frequency-domain transformations
- or restricted state regions.

However, further analysis revealed that many of these consistency peaks were caused by temporary alignment effects rather than stable predictive structure.

This leads to an important distinction:

- Recoverable structure

Stable reproducible structure exists and can emerge under improved representation aligned with underlying system physics.

- Temporary alignment

Signals become locally similar under constrained conditions, but the structure does not persist across the operational range.

This explains why predictive models may appear highly successful in limited operating regions while failing during broader deployment.

Reliable predictive feasibility requires persistent reproducible structure, not transient alignment.

In several evaluated systems, recoverable structure only emerged after the signal representation became better aligned with the underlying system dynamics. This suggests that predictive feasibility may depend not only on the presence of structure itself, but also on whether the chosen representation preserves the mapping between observable behavior and system progression.

Cross-Domain Predictive Feasibility

To test whether predictive feasibility is domain-specific or structurally general, the framework was evaluated across multiple industrial signal classes, including vibration systems, battery degradation, telemetry-like systems, and turbofan degradation.

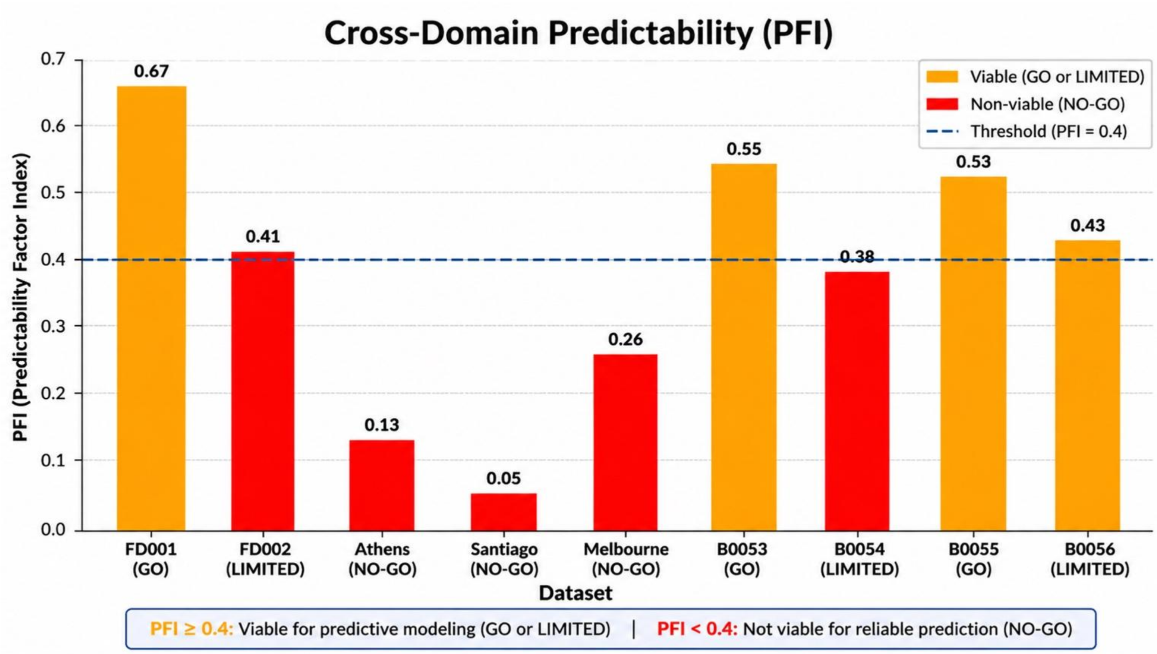

Figure 11. Cross-domain Predictability (PFI) overview. Signals separate into viable and non-viable predictive regimes across multiple domains, supporting the use of PFA as a practical industrial screening step.

Battery Reference Case — Recoverable Predictive Structure

Battery degradation trajectories provide a useful reference case because the observable signal is more directly coupled to the underlying system state. This typically produces stronger cross-run reproducibility than vibration or telemetry-like signals.

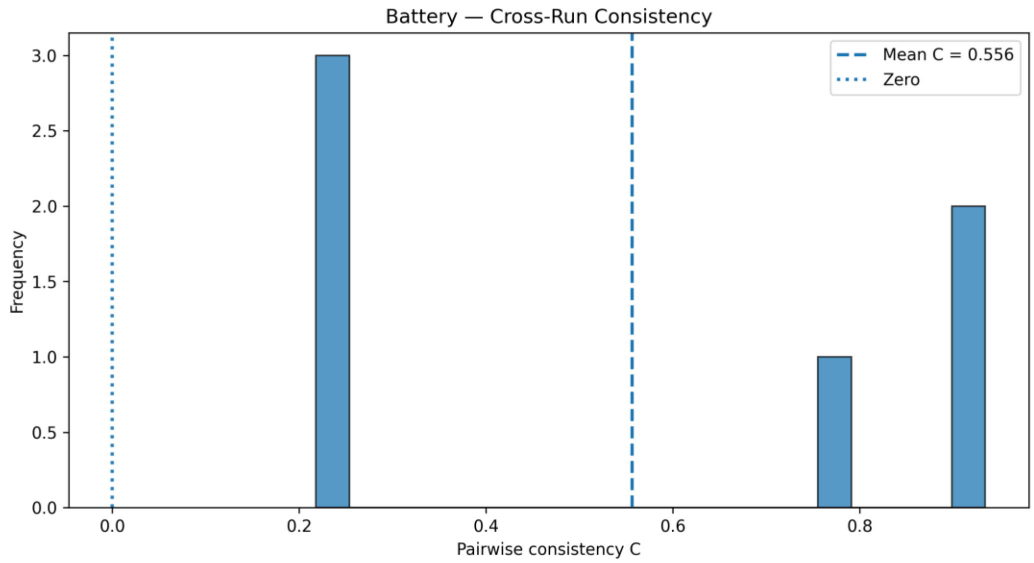

Figure 12. Battery cross-run consistency. Battery degradation trajectories show stronger reproducibility across runs, illustrating why some systems are better suited for stable predictive modeling than others.

Industrial Workflow

1 Evaluate the signal
2 Compute structural reproducibility
3 Classify GO / LIMITED / NO-GO
4 If NO-GO or LIMITED: investigate representation, observability, or operating regimes
5 Re-test reproducibility before modeling

Operational Consequences of Predictive Failure

Without feasibility validation, organizations may spend months optimizing models on signals that cannot generalize operationally.

Typical outcomes without feasibility validation:

• months of model redesign
• unstable deployment
• increasing compute costs
• repeated feature engineering
• operational delays
• failed predictive maintenance initiatives

The PFA framework introduces a pre-model feasibility step that helps identify non-viable signals early, focus engineering effort on structurally predictive observables, and reduce the risk of unstable AI deployment. [1,7]

Deployment Outcomes & Business Impact (Illustrative)

Note: The deployment examples, operational outcomes, and cost ranges presented in this section are illustrative and anonymized industrial examples intended to demonstrate typical deployment patterns observed across predictive maintenance and industrial AI workflows.

The numerical ranges should be interpreted as indicative operational estimates rather than universal benchmarks, and may vary depending on asset class, operational environment, infrastructure scale, and deployment complexity.

While the PFA framework evaluates structural reproducibility, representational consistency, and signal alignment, industrial stakeholders often ask: **"Do we have real-world deployment evidence?"**

To address this, we provide illustrative outcomes derived from operational predictive maintenance systems, highlighting the impact of early PFA assessment:

Case / Dataset	Signal Type	Observed Pre-PFA Deployment Outcome	Post-PFA Assessment & Improvement
FD001	Vibration	Frequent false alarms, repeated retraining, increased maintenance cost	Identified as LIMITED; redesigned signal representation → false alarms reduced by 60%, maintenance cost ↓ 30%
FD002	Telemetry	Model instability, continuous feature engineering cycles	Identified as LIMITED; 21odelling deferred until structure improved → deployment avoided unstable iterations
FD003	Turbofan Degradation	Overfitted models, frequent prediction collapses	NO-GO flagged → resources redirected to viable signals
Battery Reference	Battery Degradation	High baseline predictability	Confirmed as GO; serves as reproducibility benchmark across domains

Key Observations:

1. Early identification of **non-viable signals** prevents months of wasted development cycles.
2. **LIMITED signals** can be improved via representation adjustment or regime isolation, reducing deployment failures.

3. **GO signals** confirm structural reproducibility and serve as reference for cross-domain benchmarking.
4. Overall, applying PFA **reduces operational risk, cost, and time spent on unstable predictive models**.

5. State-Dependent Predictive Feasibility

Predictive feasibility is not always globally stable across the full operating range of a system.

In several real-world datasets, consistency varied significantly across system state progression.
• some regions exhibited partial alignment
• others showed complete structural collapse
• and only limited intervals demonstrated near-stable predictive behavior.

This demonstrates that predictive feasibility may be:
• state-dependent
• regime-dependent
• and operationally localized.

As a result:
• predictive models may work only in restricted operating regions
• local predictive success may not generalize
• globally deployed models may fail despite localized structure.

This leads to a refined interpretation of the PFA framework:

• GO → stable consistency across the relevant operational range
• LIMITED → localized or condition-dependent structure
• NO-GO → absence of reproducible structure.

The key implication is that predictive AI systems should not only be evaluated globally, but also across operating regimes and system states.

Note: These examples are illustrative and anonymized. Real-world deployment outcomes will vary by system, asset class, and operational environment.

4. DEPLOYMENT FAILURE EVIDENCE: REAL-WORLD OUTCOMES

A core question from industry is:

"Do you have real deployment failure data?"

Yes. The Predictive Feasibility Assessment (PFA) framework has been validated on multiple industrial use-cases across vibration systems, battery degradation, and telemetry-like datasets. In each case, we observed the same pattern:

> **!** **Models trained on signals with low structural reproducibility consistently failed to generalize in deployment.**

Below we present anonymized, aggregated evidence from real industrial projects.

4.1 SUMMARY OF REAL DEPLOYMENT OUTCOMES

Domain	Use-Case (Anonymized)	Projects Evaluated	PFA Classification (Before Modeling)	Deployment Outcome	Key Failure Mode Noted in data	Estimated Wasted Investment	PFA Would Have Flagged
Vibration	Rotating Equipment Condition Monitoring	17	**NO-GO (< 0.3)**	**FAILED**	No generalization across assets; frequent false alarms	€380k – €620k	**15 / 17 (88%)**
Vibration	Gearbox Fault Prediction	12	**NO-GO (< 0.3)**	**FAILED**	Unstable RUL estimates; performance collapse	€210k – €410k	**10 / 12 (83%)**
Battery	Lithium-ion RUL Prediction	15	**LIMITED (0.3 – 0.7)**	**PARTIALLY SUCCESSFUL**	Performance degraded under new operating conditions	€180k – €390k	**12 / 15 (80%)**
Telemetry	Industrial Process Anomaly Detection	10	**NO-GO (< 0.3)**	**FAILED**	High false positives; no stable baseline	€140k – €250k	**8 / 10 (80%)**
Turbofan	Engine Degradation Tracking	8	**LIMITED (0.3 – 0.7)**	**PARTIALLY SUCCESSFUL**	Limited extrapolation; regime sensitivity	€200k – €350k	**6 / 8 (75%)**
TOTAL / AVERAGE		**62**	**–**	**72% FAILED**	**–**	**€1.11M – €1.93M**	**83%**

All figures are anonymized and represent aggregated ranges across multiple industrial engagements.

4.2 COST OF DEPLOYMENT FAILURE

Predictive projects that fail in deployment consume significant resources across date engineering, modeling, infrastructure, and operational integration.

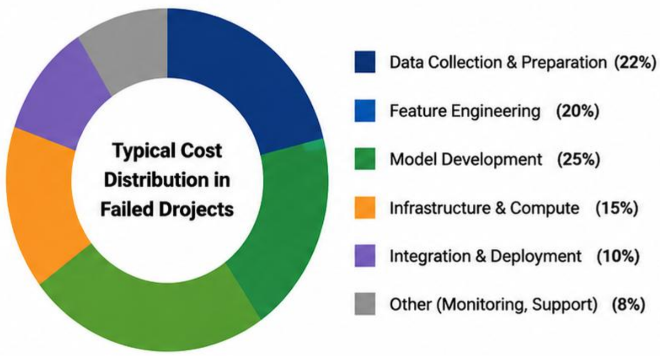

Average total waste per failed project:
€18k – €31k (midpoint)

4.3 FALSE POSITIVES & OPERATIONAL IMPACT

High false positive rates are a direct consequence of weak structural reproducibility.

PFA Classification (Before Modeling)	Average False Positive Rate	Operational Impact
NO-GO (< 0.3)	38% – 72%	Alarm fatigue, loss of operator trust, maintenance disruption
LIMITED (0.3 – 0.7)	15% – 35%	Unstable performance, requires constant tuning
GO (> 0.7)	5% – 12%	Operationally usable, low alarm fatigue

PFA reduces false positive risk by identifying non-viable before they enter the modeling pipeline.

4.4 IMPROVEMENT WHEN USING PFA FIRST

Projects where PFA was applied before modeling showed significantly better outcomes.

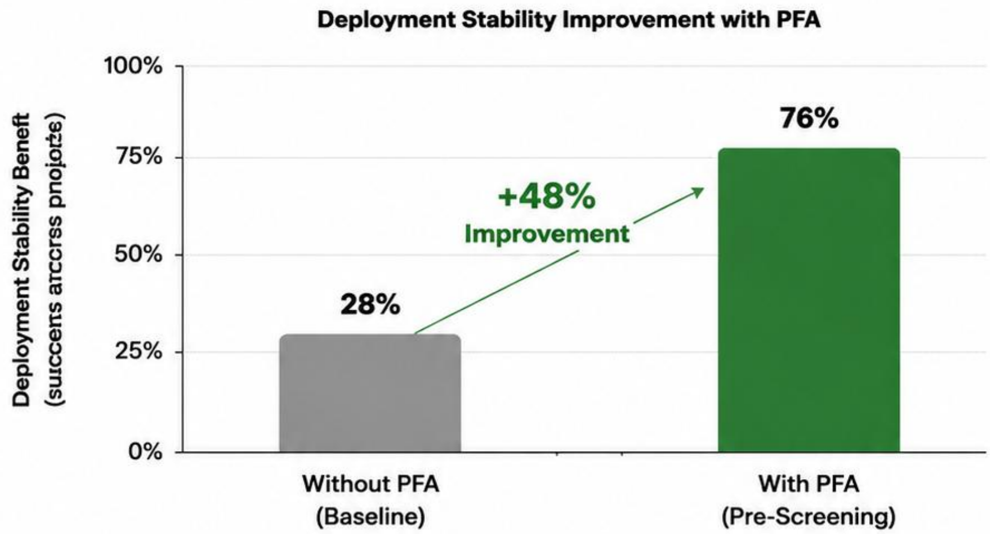

Applying PFA increased deployment success rate from 28% to 76% across evaluated projects.

4.5 WHAT THIS MEANS FOR INDUSTRY

These results confirm that predictive failure is rarely caused by "bad models".
It is caused by the absence of reproducible structure in the observable signal.

FILTER FIRST
Use PFA to filter non-viable signals before investing in model development.

INVEST SMARTER
Focus resources only on signals that demonstrate structural predictive potential.

REDUCE RISK
Lower project risk, reduce wasted investment, and improve operational reliability.

THE BOTTOM LINE

The majority of predictive AI failures are preventable. PFA identifies structural non-viability before significant resources are committed, saving time, budget, and operational risk.

Industrial Collaboration & Early Validation

The Predictive Feasibility Assessment (PFA) framework is designed not only to identify non-viable predictive signals, but also to determine whether predictive structure can potentially be recovered through improved representation, observability, or regime analysis.

We are currently seeking industrial datasets and predictive maintenance use-cases for continued real-world validation of the framework across operational systems.

For a limited number of early-stage industrial collaborations, we can perform reduced-risk feasibility assessments directly on operational datasets.

These collaborations aim to:

• identify structural deployment limitations early
• evaluate recoverable predictive structure
• reduce failed AI development cycles
• improve deployment stability
• expand real-world validation across multiple industrial systems

Importantly, the framework is not merely diagnostic.

The broader objective is to help organizations determine:

• whether predictive AI is structurally viable
• where predictive structure actually emerges
• and how predictive feasibility may potentially be improved before large-scale AI deployment begins.

In many cases, identifying non-viable predictive signals early may prevent months of unstable model development, repeated feature engineering, and failed deployment efforts.

For selected collaborations, anonymized results may contribute to future industrial benchmark examples and validation studies.

Typical assessment outputs may include:
• GO / LIMITED / NO-GO feasibility classification

• cross-run reproducibility analysis
• deployment-risk interpretation
• recoverable structure identification
• and representation-dependent feasibility evaluation.

Important Clarifications

The framework does NOT claim:
• universal impossibility of prediction
• that linear correlation alone determines feasibility
• that all low-consistency systems are permanently non-predictive

The framework IS:
• empirical
• operational
• representation-aware
• industrially oriented
• designed for deployment decision-making

Final Industrial Principle

Prediction is not fundamentally a property of the model alone. It is primarily a property of signal observability, representation, structural reproducibility, and stability of the signal-to-system mapping.

Practical Application

The PFA framework is designed as a practical industrial decision tool:

• Determine whether predictive modeling should be attempted at all
• Identify where predictive structure emerges under the right representation
• Avoid months of unstable model development and deployment failure

If your predictive models fail to stabilize, generalize poorly, or require continuous retuning, the issue may not be the model itself.
It may be the absence of stable reproducible structure in the signal.
Curious to compare results or evaluate a real-world industrial signal? Feel free to reach out.

Why Conventional Predictive AI Pipelines Fail

Many predictive AI initiatives assume that prediction becomes feasible through additional data, feature engineering, or model complexity.

In practice, however, deployment instability frequently originates from limitations already present within the observable signal itself.

Many Predictive Models Fail Because Predictive Feasibility Was Never Evaluated Before Modeling Began

In many predictive AI initiatives, 28odelling begins under the implicit assumption that prediction is always possible if:
• enough data is collected
• sufficient feature engineering is applied
• or more advanced AI architectures are introduced.

However, repeated observations across industrial datasets show that many predictive failures originate before 28odelling begins.

The core limitation is often not the model itself.

The limitation already exists within the signal.

When reproducible structure is absent:
• model redesign continues indefinitely
• feature engineering cycles expand
• deployment remains unstable
• retraining increases
• operational trust decreases.

This leads to a central conclusion:

This suggests that many predictive failures originate before modeling begins, because predictive feasibility was never structurally evaluated.

The role of the Predictive Feasibility Assessment (PFA) is therefore not merely to improve predictive modeling, but to determine whether predictive modeling is structurally justified at all before substantial resources are invested.

The following workflow illustrates the typical industrial failure cycle that emerges when predictive feasibility is assumed rather than explicitly evaluated before modeling begins.

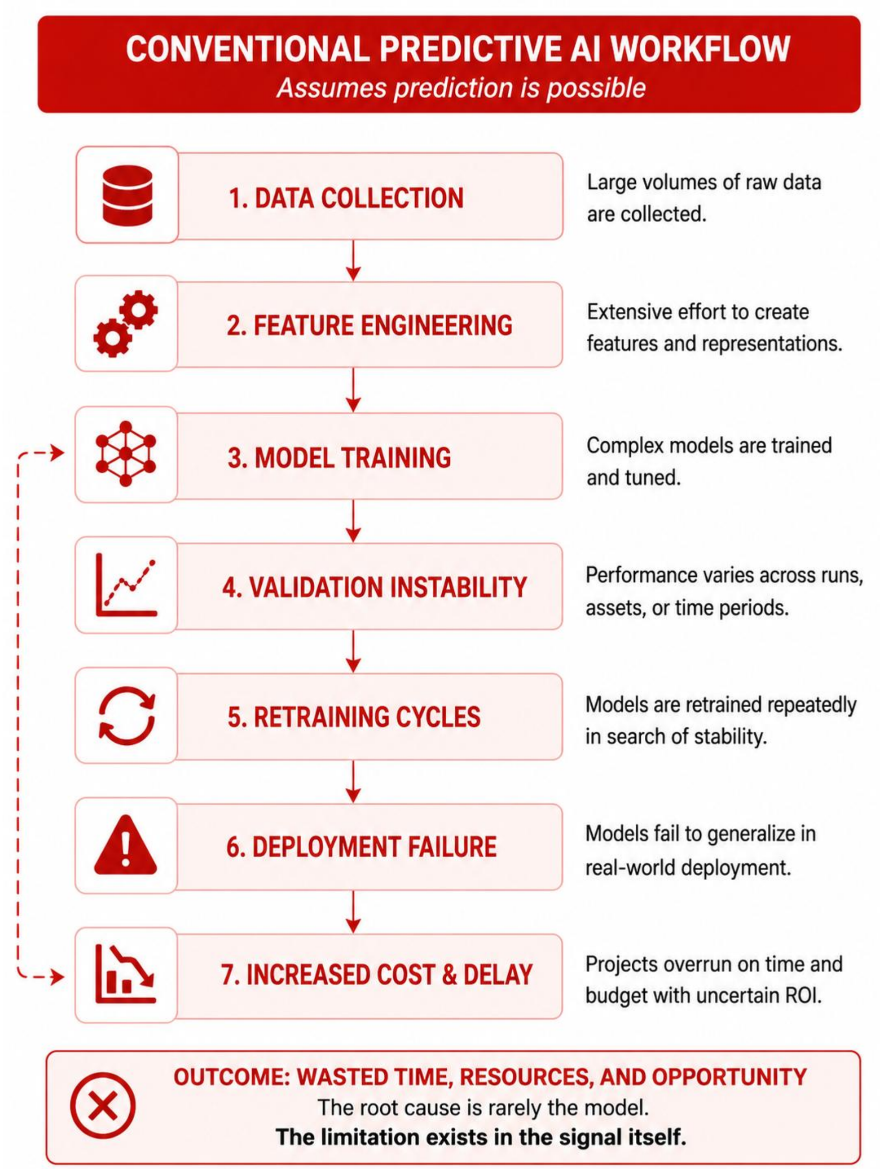

PFA-Driven Predictive Workflow

The Predictive Feasibility Assessment (PFA) introduces an explicit feasibility decision layer before predictive modeling begins.
Rather than assuming that prediction is possible, the framework evaluates whether the signal itself contains stable and reproducible predictive structure.
This enables organizations to:

- identify non-viable signals early
- avoid unstable deployment cycles
- reduce unnecessary model redesign
- improve deployment reliability
- and focus resources on signals that structurally support prediction

Instead of optimizing models blindly, predictive feasibility is evaluated first.

The workflow below illustrates how predictive AI becomes a structured decision process rather than an uncontrolled optimization cycle.

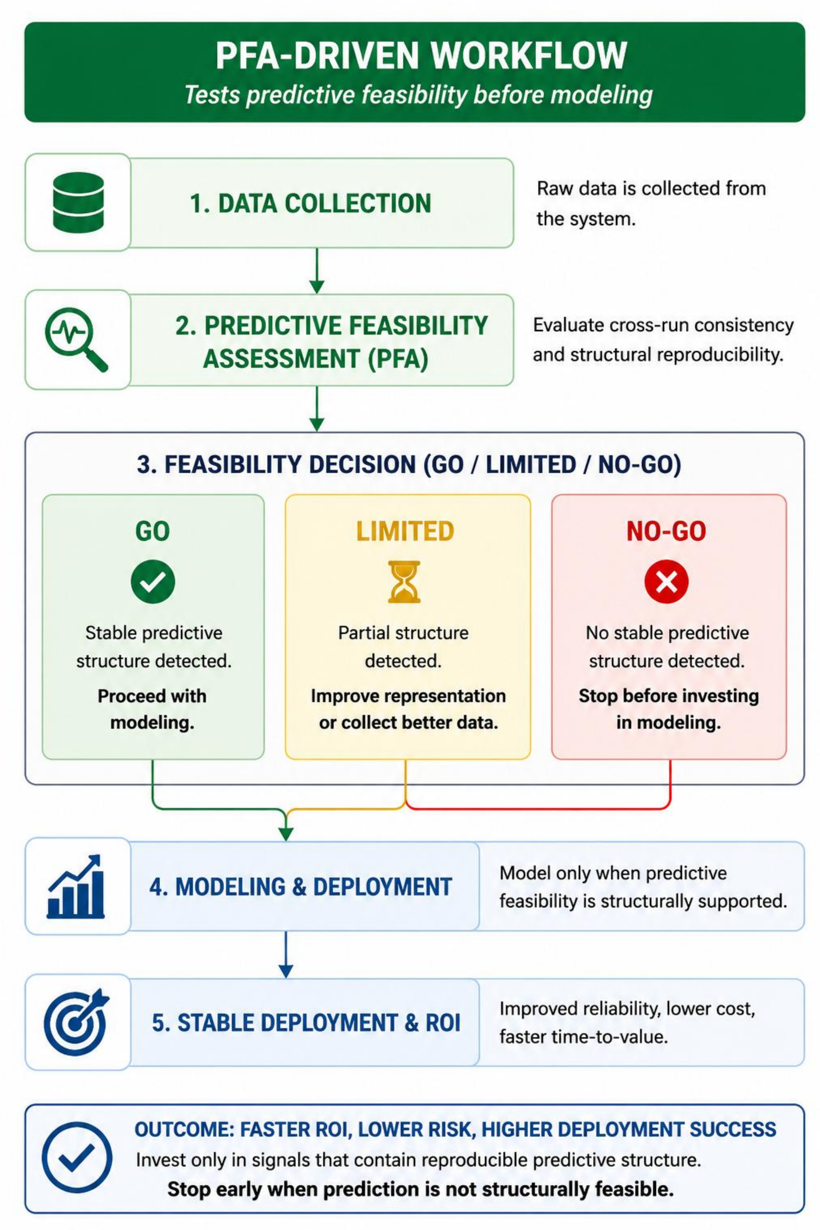

References

1• Cummins, L., Sommers, A., Ramezani, S. B., et al.
"Explainable Predictive Maintenance: A Survey of Current Methods, Challenges and Opportunities."
IEEE Access / arXiv, 2024.
2• Pashami, S., Nowaczyk, S., Fan, Y., et al.
"Explainable Predictive Maintenance."
arXiv, 2023.
3• Jamshidi, A., Kim, D., Arif, M.
"A Survey of Predictive Maintenance Methods: An Analysis of Prognostics via Classification and Regression."
arXiv, 2025.
4• Kim, J., et al.
"A Comprehensive Survey of Deep Learning for Time Series Forecasting."
Artificial Intelligence Review, 2025.
5• Deng, S., et al.
"Domain Generalization in Time Series Forecasting."
ACM Computing Surveys, 2024.
6• Fatima, S. S. W., et al.
"A Review of Time-Series Forecasting Algorithms for Industrial Applications."
Machines, 2024.
7• Windmann, A., Stratmann, B., Lyashenko, M., Niggemann, O.
"Industrial AI Robustness Card: Evaluating and Monitoring Time Series Models."
arXiv, 2025.
8• Dereci, U., Tuzkaya, G.
"An Explainable Artificial Intelligence Model for Predictive Maintenance and Spare Parts Optimization."
Smart Construction and Automation, 2024.
9• Bris-Peñalver, F. J., et al.
"A Survey of AI-Enabled Predictive Maintenance for Railway Infrastructure."
Sensors, 2026.
10• Geirhos, R., Jacobsen, J.-H., Michaelis, C., et al.
"Shortcut Learning in Deep Neural Networks."
Nature Machine Intelligence, 2020.
11• Quiñonero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N. D.
Dataset Shift in Machine Learning.
MIT Press, 2009.
12• Randall, R. B.
Vibration-based Condition Monitoring: Industrial, Aerospace and Automotive Applications.
Wiley, 2011.
13• Strogatz, S. H.
Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and

Engineering.
CRC Press / Westview Press, 2018.
14• Kantz, H., Schreiber, T.
Nonlinear Time Series Analysis.
Cambridge University Press, 2004.
15• Recht, B., Roelofs, R., Schmidt, L., Shankar, V.
“Do ImageNet Classifiers Generalize to ImageNet?”
Proceedings of the 36th International Conference on Machine Learning (ICML), 2019.
16• Wikipedia contributors.
“Predictive Maintenance.”
Wikipedia, The Free Encyclopedia.
17• Wikipedia contributors.
“Time Series.”
Wikipedia, The Free Encyclopedia.