



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

PFA Executive Guide

Assessing Predictive Feasibility Before Large-Scale AI Deployment

This Executive Guide provides a concise operational overview of the Predictive Feasibility Assessment (PFA) framework.

A full technical whitepaper containing methodology, robustness analysis, datasets, and supporting technical validation is available separately.

<https://doi.org/10.5281/zenodo.20159329>

Part 1 — Why Predictive AI Systems Fail

Industrial organizations increasingly invest in predictive AI, predictive maintenance, anomaly detection, and time-series forecasting systems.

Yet despite major advances in AI architectures, many predictive systems still fail to achieve stable operational deployment.

A recurring industrial pattern is frequently observed:

- validation metrics initially appear promising
- additional feature engineering is introduced
- increasingly complex AI models are deployed
- retraining cycles continue
- yet deployment stability eventually collapses.

In practice, this often leads to:

- unstable predictions across assets
- excessive false positives
- unreliable Remaining Useful Life (RUL) estimates
- repeated tuning cycles
- operational distrust in predictive systems
- increasing engineering and maintenance cost.

A central observation emerging across multiple industrial datasets is that these failures are often not caused by insufficient model complexity.

The limitation frequently already exists within the observable signal itself. Many industrial signals may appear:

- highly active
- statistically predictable
- locally structured
- or operationally meaningful

while still failing to support stable predictive deployment across independent runs, assets, or operating conditions.

This creates a major industrial problem: Organizations may spend months optimizing models on signals that never contained stable reproducible predictive structure. The Predictive Feasibility Assessment (PFA) framework was developed to identify this limitation before large-scale deployment begins.

Figure 1 — Raw Industrial Signal Behavior Before Predictive Feasibility Assessment




Figure 1. Pure raw multi-signal industrial telemetry data prior to predictive feasibility assessment. The figure shows unprocessed monitoring signals plotted directly over time without filtering, aggregation, alignment, or transformation. Multiple overlapping trajectories create a visually active and structurally complex representation, yet it remains unclear whether stable reproducible predictive structure is actually present. This represents the stage at which many predictive modeling initiatives typically begin before predictive feasibility is explicitly evaluated.

Related Technical Whitepaper Sections

- Pure Raw Data (Industry Input) → Whitepaper pp. 4–5
- Structured Representation → Whitepaper pp. 6–7
- Prediction Stability vs Predictive Feasibility → Whitepaper pp. 13–16
- Operational Consequences of Predictive Failure → Whitepaper pp. 21–22

Part 2 — Prediction Stability Does Not Guarantee Predictive Feasibility

One of the most common misconceptions in predictive AI is the assumption that stable short-term prediction automatically implies reliable deployment behavior.

In practice, however, this assumption frequently fails.

Across multiple industrial datasets, signals were observed that exhibited:

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

while simultaneously lacking reproducible predictive structure across independent runs.

This creates a critical industrial problem:

Predictive AI systems may appear statistically successful during validation while remaining structurally unstable during deployment.

In many operational environments, validation metrics remain strong while deployment stability collapses across assets or operating conditions.

The limitation is often not the AI model itself, but the absence of reproducible predictive structure within the observable signal.

The Predictive Feasibility Assessment (PFA) framework therefore distinguishes between:

- apparent predictability
- and
- structurally reproducible predictability.

This distinction is operationally critical because stable deployment requires reproducible predictive structure across operating conditions, assets, and independent system realizations. Prediction stability alone is therefore insufficient to determine predictive feasibility.

The framework evaluates whether predictive structure remains reproducible under real deployment conditions before large-scale engineering effort is invested.

Figure 2 — Prediction Stability vs Operational Predictive Feasibility




Figure 2. Structural distinction between local prediction stability and operational predictive feasibility. Many industrial AI systems achieve stable validation metrics and low short-term prediction error while still failing during deployment because the observable signal lacks reproducible predictive structure across assets, operating conditions, or independent runs. The PFA framework evaluates this structural reproducibility before large-scale deployment effort begins.

Related Technical Whitepaper Sections

- Prediction Stability Is Not Predictive Feasibility → Whitepaper pp. 13–15
- Predictable but Non-Informative Signals → Whitepaper pp. 15–16
- Why Conventional Predictive AI Pipelines Fail → Whitepaper pp. 28–29
- Why Predictive Feasibility Must Be Evaluated Before Modeling → Whitepaper pp. 29–30

Part 3 — Recoverable Structure vs Structurally Non-Viable Signals

A critical industrial question is not only whether predictive structure currently exists, but whether predictive structure can potentially be recovered through improved representation and analysis.

Across multiple industrial datasets, an important distinction emerges between:

- structurally non-viable predictive signals
and
- signals containing recoverable predictive structure.

Many industrial systems initially appear unstable in their raw representation.

These signals may exhibit:

- strong variability
- unstable alignment
- overlapping dynamics
- high signal activity
- and low reproducibility across runs.

Under these conditions, predictive deployment frequently becomes unstable.

However, additional analysis demonstrates that some systems contain latent predictive structure that only becomes visible when the signal representation becomes better aligned with the underlying system physics.

This may occur through:

- frequency-domain analysis
- narrow-band spectral isolation
- observability improvements
- regime isolation
- or alternative signal representations.

A key operational finding of the PFA framework is that predictive feasibility may improve significantly when analysis focuses on physically meaningful signal regions rather than the raw signal alone.

This creates two fundamentally different deployment scenarios:

Structurally Non-Viable Signals

These signals:

- remain inconsistent across representations
- fail to recover reproducible structure
- continue producing unstable deployment behavior.

In these cases, increasing model complexity is unlikely to solve the deployment problem.

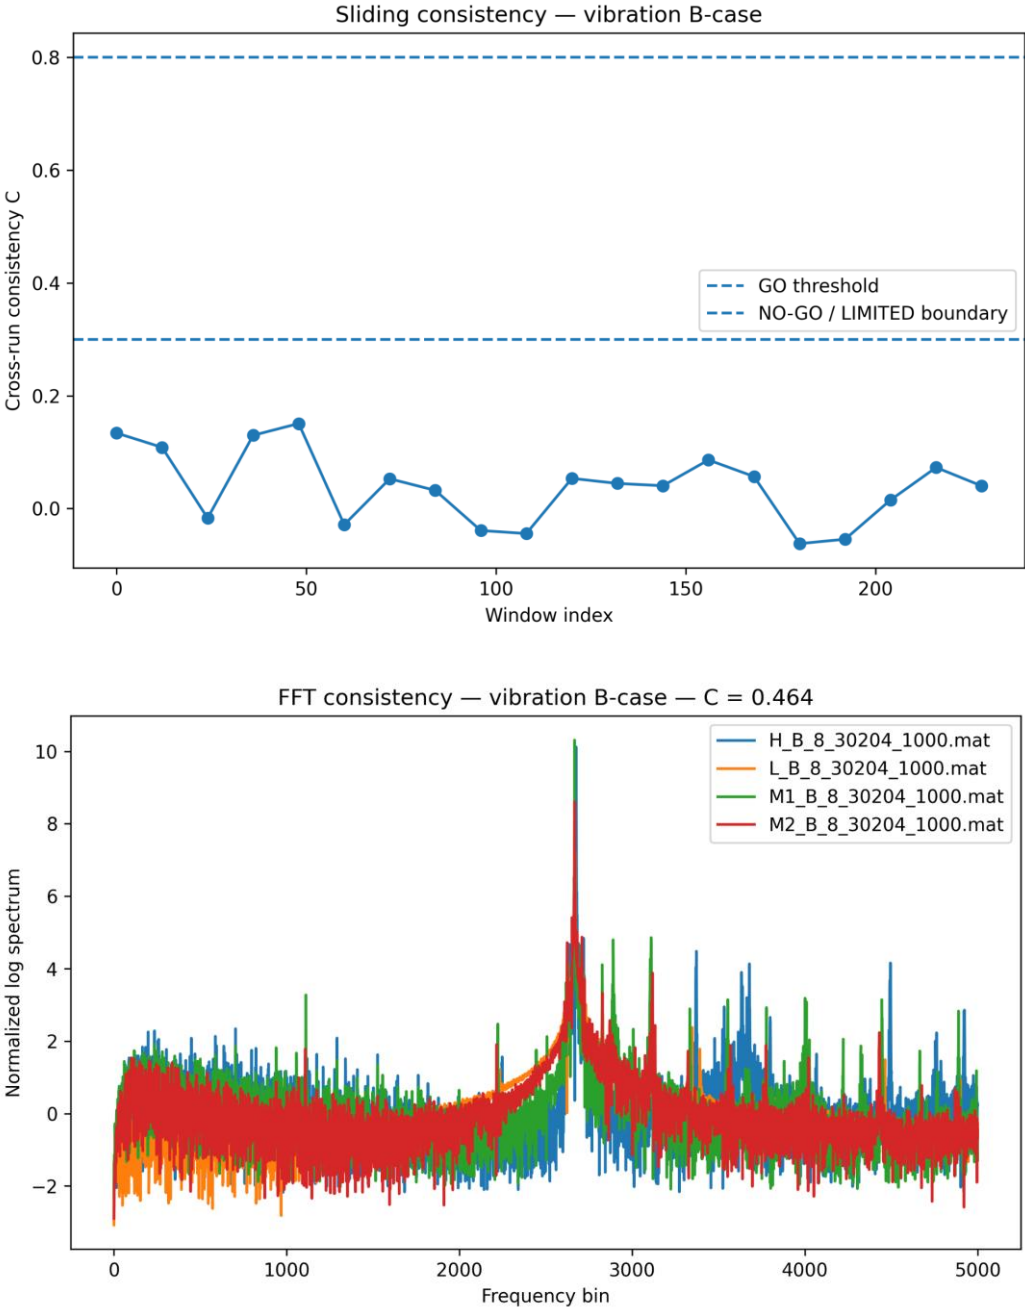
Recoverable Predictive Structure

These signals initially appear unstable, but reproducible predictive structure begins to emerge under improved representation.

In these cases, predictive deployment may become operationally viable after representation recovery.

The practical objective of the framework is therefore not only to identify non-viable predictive systems, but also to determine whether predictive structure can become operationally viable before large-scale deployment effort begins.

Figure 3 — Recovery of Predictive Structure Through Representation Alignment






Figure 3. Multi-stage representation-dependent recovery of predictive structure in an industrial vibration signal. The raw time-domain representation exhibits near-zero reproducibility (NO-GO), while progressively improved signal representations reveal increasing structural consistency. This illustrates how predictive feasibility may become operationally viable when signal representation becomes better aligned with the underlying system physics.

Related Technical Whitepaper Sections

- Frequency-Domain Recovery → Whitepaper pp. 17–18
- Recoverable Structure vs Temporary Alignment → Whitepaper pp. 19–20
- Cross-Domain Predictive Feasibility → Whitepaper pp. 20–21
- Industrial Workflow → Whitepaper pp. 21–22

Part 4 — Operational Consequences of Ignoring Predictive Feasibility

Many industrial predictive AI initiatives begin with the implicit assumption that prediction becomes feasible if:

- enough data is collected
- additional features are engineered
- larger AI models are deployed
- or retraining cycles are increased.

In practice, however, many predictive systems continue failing despite increasing technical complexity.

Across multiple industrial environments, a recurring operational pattern emerges:

- models initially appear promising
- validation metrics improve

- additional engineering effort is invested
- deployment begins
- yet operational instability continues.

This frequently leads to repeated retraining cycles, false positives, growing engineering effort, and unstable deployment behavior.

A central observation of the PFA framework is that these failures often originate before deployment begins, because reproducible predictive structure is absent within the observable signal itself.

Rather than continuously optimizing unstable models, the framework evaluates whether predictive structure exists, whether it is recoverable, and whether deployment viability is operationally realistic.

Figure 4 — Operational Consequences of Structurally Unstable Predictive Signals

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.

Figure 4. Illustrative operational consequences observed when predictive AI systems are deployed on signals lacking stable reproducible predictive structure. Typical outcomes include unstable deployment behavior, repeated retraining cycles, false positives, increasing engineering effort, and operational inefficiency. The PFA framework is designed to identify these limitations before large-scale deployment investment begins.

Related Technical Whitepaper Sections

- Operational Consequences of Predictive Failure → Whitepaper pp. 21–22
- Deployment Outcomes & Business Impact → Whitepaper pp. 22–24
- Why Conventional Predictive AI Pipelines Fail → Whitepaper pp. 28–29
- Most Predictive Models Fail Because They Should Never Have Been Built → Whitepaper pp. 29–30

Part 5 — Practical Predictive Feasibility Assessment Workflow

The Predictive Feasibility Assessment (PFA) framework is designed as a practical pre-model evaluation process for industrial predictive systems.

The objective is not merely to improve predictive AI models.

The objective is to determine whether predictive modeling is structurally viable before organizations invest heavily in:

- model development
- feature engineering
- deployment infrastructure
- retraining workflows
- or large-scale predictive maintenance integration.

A typical PFA assessment begins with operational industrial signal data such as:

- vibration telemetry
- sensor streams
- degradation trajectories
- industrial monitoring systems
- predictive maintenance datasets
- process telemetry
- or anomaly-related operational signals.

The framework then evaluates whether the observable signal contains reproducible predictive structure across:

- independent runs
- assets
- operating conditions
- temporal regions
- and alternative signal representations.

The practical assessment process typically includes:

1. Structural Reproducibility Analysis

Evaluation of whether signal structure remains reproducible across independent realizations.

2. Representation Recovery Analysis

Evaluation of whether predictive structure improves under alternative representations such as:

- spectral analysis
- frequency-domain isolation
- observability enhancement
- or regime separation.

3. Deployment Feasibility Assessment

Evaluation of whether predictive structure remains sufficiently stable for operational deployment.

4. GO / LIMITED / NO-GO Classification

Operational classification of predictive feasibility based on structural reproducibility and deployment stability.

This process allows organizations to determine where engineering effort should be focused before large-scale deployment begins.

The practical objective is simple:

Avoid investing months of engineering effort into signals that were never structurally predictive to begin with.

At the same time, the framework helps identify systems where predictive feasibility may become operationally viable after representation recovery or observability improvement.

The framework therefore acts as both:

- a deployment-risk assessment methodology
and
- a predictive structure recovery methodology.

Figure 5 — Practical Industrial Predictive Feasibility Assessment Process

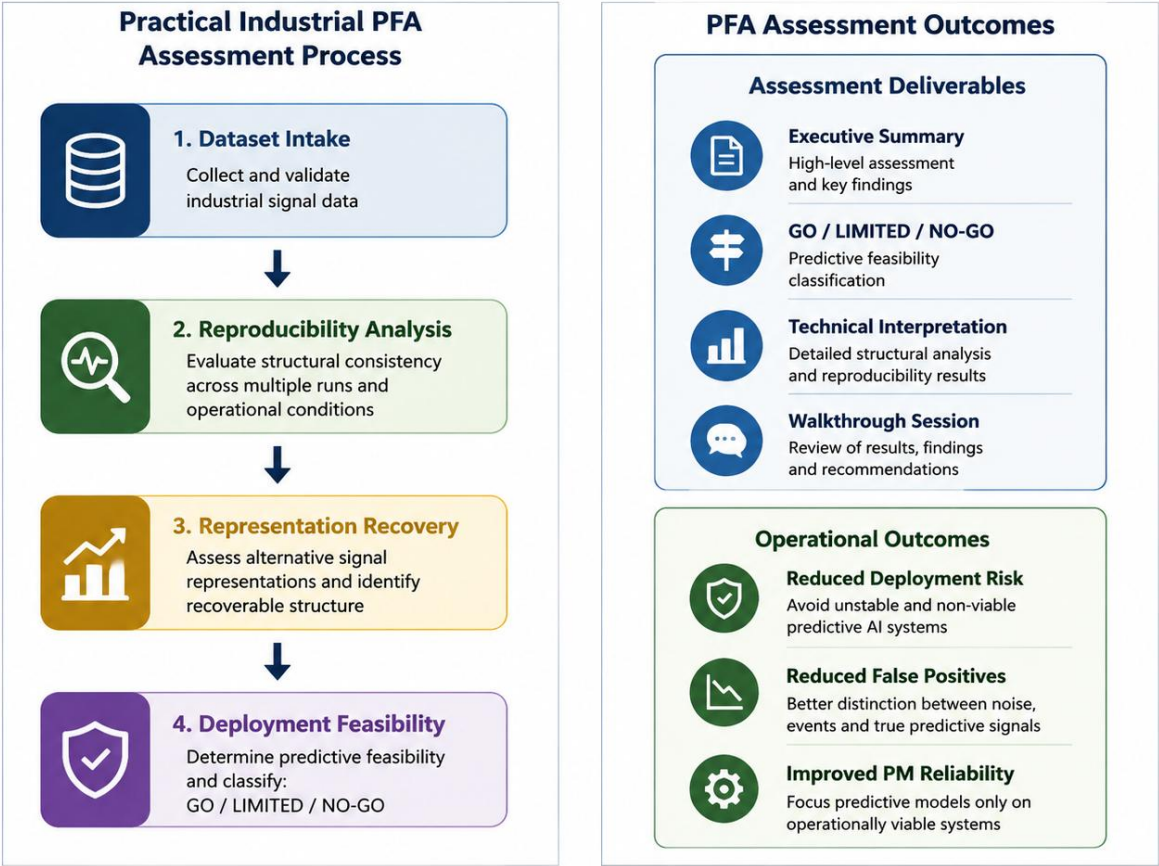


Figure A5 (Part 1). Practical Industrial PFA Assessment Process.

Figure A5 (Part 2). PFA Assessment Deliverables and Outcomes.

Figure 5. Practical industrial Predictive Feasibility Assessment (PFA) process. The framework evaluates whether operational signals contain stable reproducible predictive structure before substantial deployment effort is invested. The assessment process identifies operationally viable predictive systems, recoverable predictive structure, and deployment risks across industrial environments.

Related Technical Whitepaper Sections

- Industrial Workflow → Whitepaper pp. 21–22
- Deployment Outcomes & Business Impact → Whitepaper pp. 22–24
- Industrial Collaboration & Early Validation → Whitepaper pp. 27–28
- Practical Application → Whitepaper pp. 28–29
- PFA-Driven Predictive Workflow → Whitepaper pp. 30–31

Part 6 — Predictive Feasibility Across Industrial Domains

To evaluate whether predictive feasibility is domain-specific or structurally general, the Predictive Feasibility Assessment (PFA) framework was evaluated across multiple industrial signal classes.

These include:

- vibration monitoring systems
- telemetry systems
- battery degradation trajectories
- turbofan degradation systems
- industrial monitoring streams
- and real-world operational time-series datasets.

Across these domains, a consistent structural pattern emerges:

- some signals support stable reproducible prediction
- some signals contain partially recoverable predictive structure
- and some signals remain structurally unstable regardless of additional model complexity.

Predictive feasibility depends fundamentally on the reproducibility of observable signal structure rather than model complexity alone.

Some systems exhibit stable short-term forecasting behavior while still failing operationally because reproducible predictive structure is absent.

Other systems become significantly more reproducible when representation becomes better aligned with the underlying system physics.

This enables the framework to distinguish between:

- structurally non-viable predictive systems
- recoverable predictive structure
- and operationally viable predictive signals.

Figure 6 — Cross-Domain Predictive Feasibility Overview

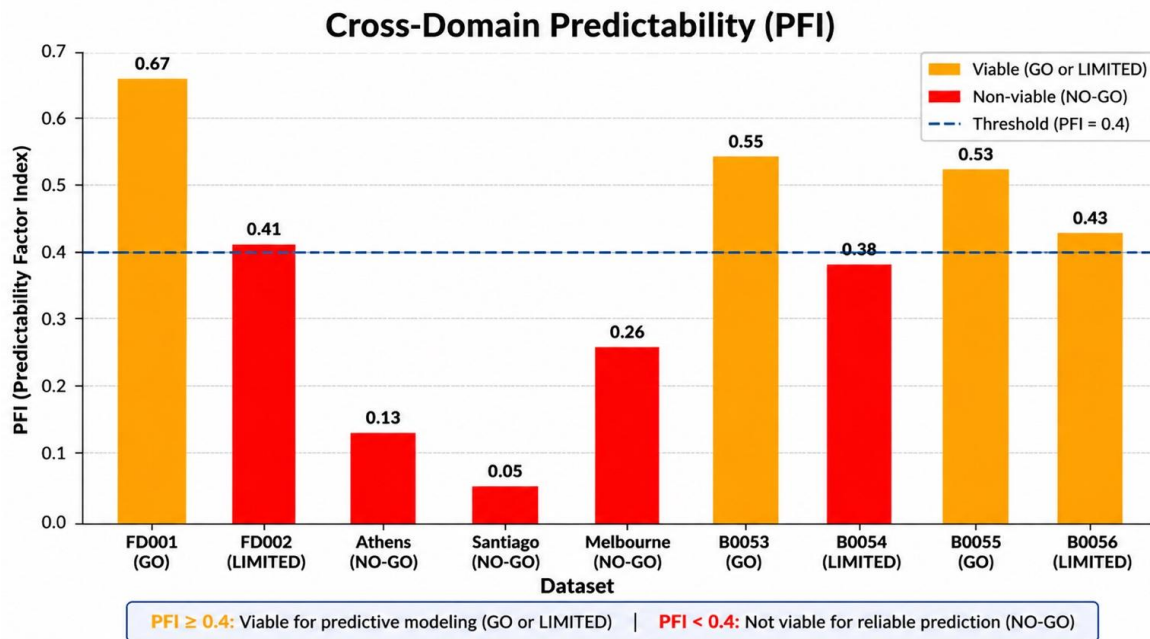


Figure 6. Cross-domain Predictive Feasibility Assessment (PFA) overview across multiple industrial signal classes. Signals separate into operationally viable, partially recoverable, and structurally non-viable predictive regimes across multiple domains, demonstrating that predictive feasibility is fundamentally linked to reproducible signal structure rather than model complexity alone.

Related Technical Whitepaper Sections

- Cross-Domain Predictive Feasibility → Whitepaper pp. 19–20
- Battery Reference Case — Recoverable Predictive Structure → Whitepaper pp. 20–21
- Industrial Workflow → Whitepaper pp. 21–22
- Deployment Outcomes & Business Impact → Whitepaper pp. 22–24

Part 7 — Collaboration & Assessment Process

The Predictive Feasibility Assessment (PFA) framework is designed as a practical industrial assessment process that can be applied before large-scale predictive AI deployment begins.

The purpose of the assessment is not merely to evaluate model performance.

The objective is to determine:

- whether predictive deployment is structurally viable
- whether predictive structure is reproducible
- whether predictive feasibility may be recoverable under improved representation
- and whether operational deployment risk is acceptable before substantial engineering effort is invested.

A typical industrial PFA collaboration may include:

- cross-run reproducibility analysis
- representation-dependent consistency evaluation
- deployment stability analysis
- recoverable structure investigation
- operational regime evaluation
- and signal-to-system alignment analysis.

The framework can be applied across multiple industrial domains, including:

- predictive maintenance
- vibration monitoring
- telemetry systems
- degradation modeling
- Remaining Useful Life (RUL) prediction
- anomaly detection
- industrial monitoring systems
- and sensor-based operational forecasting.

Typical operational outputs may include:

- GO / LIMITED / NO-GO classification
- reproducibility interpretation
- deployment-risk assessment
- representation recovery recommendations
- operational stability interpretation
- and consistency analysis across assets and conditions.

For suitable industrial collaborations, assessments may be performed directly on operational datasets under NDA or confidential conditions.

What a Predictive Feasibility Assessment (PFA) May Include

A Predictive Feasibility Assessment (PFA) may include:

- evaluation of raw operational telemetry or monitoring data
- cross-run reproducibility analysis
- GO / LIMITED / NO-GO feasibility classification
- identification of structurally unstable observables
- representation-dependent feasibility evaluation
- representation-dependent frequency-domain and regime analysis
- assessment of deployment stability risk
- identification of partially recoverable predictive structure
- evaluation of whether additional signals improve or reduce ambiguity
- interpretation of predictive limitations under varying operational conditions

Typical Deliverables

Depending on the industrial use-case and signal characteristics, a typical assessment may provide:

- GO / LIMITED / NO-GO decision per signal or representation
- structural reproducibility report
- deployment-risk interpretation
- signal viability assessment before model development
- identification of misleading or non-generalizable observables
- recommendations for improved representation or observability
- practical interpretation of predictive feasibility limitations
- optional comparative analysis across multiple signals or operating regimes

Collaboration Format

Typical exploratory collaborations may involve:

- limited operational datasets
- anonymized telemetry or vibration signals
- short exploratory feasibility studies (typically 1–2 weeks)
- NDA-based collaboration if required
- reduced-risk pilot assessments for early-stage validation

All operational data remains the property of the collaborating organization.

No public sharing or publication occurs without explicit approval.

To ensure technical clarity, reproducibility, and structured interpretation throughout the assessment process, detailed feasibility interpretation is typically delivered using:

- structured walkthrough presentations
- annotated figures and slides
- deployment-oriented interpretation
- and documented technical summaries.

This allows engineering and operational teams to:

- revisit technical conclusions internally
- reduce risk of miscommunication
- maintain traceable analytical interpretation
- and retain a reusable technical reference throughout predictive AI deployment evaluation.

For selected early-stage industrial collaborations, reduced-cost pilot assessments may be performed in exchange for permission to use fully anonymized operational findings as future benchmark or case-study material.

The objective of these early collaborations is to further validate the framework across real operational environments while simultaneously providing industrial partners with low-risk predictive feasibility assessment prior to large-scale deployment investment.

The broader industrial objective of the framework is simple:

Evaluate predictive feasibility before large-scale predictive AI deployment begins.

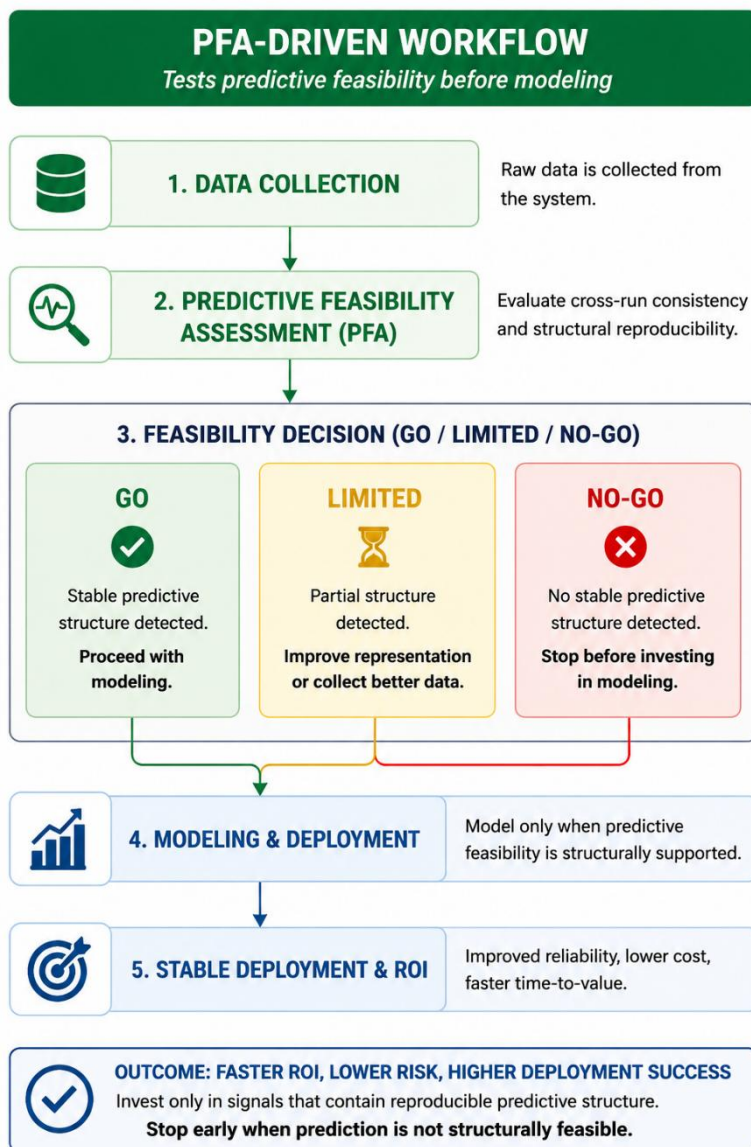


Figure 7. PFA-driven predictive AI deployment workflow. Rather than assuming that prediction is automatically feasible, the framework introduces an explicit predictive feasibility

decision layer before large-scale model development begins. The workflow helps organizations identify non-viable predictive systems early, investigate recoverable predictive structure, and reduce unstable deployment cycles before substantial engineering effort is invested.

Related Technical Whitepaper Sections

- Industrial Collaboration & Early Validation → Whitepaper pp. 27–28
- Practical Application → Whitepaper pp. 28–29
- Why Conventional Predictive AI Pipelines Fail → Whitepaper pp. 28–29
- PFA-Driven Predictive Workflow → Whitepaper pp. 30–31