SACJ 38(2) Aug 2026
Research Article

Process Intelligence Maturity as a Leading Indicator of Enterprise AI Return on Investment

Mihir Shah

Independent Researcher, USA

ABSTRACT
Enterprise artificial intelligence investment decisions are frequently made under uncertainty about which initia-tives will realize a measurable return. Existing AI readiness frameworks address data, infrastructure, talent, and governance, yet they tend to treat process intelligence as one factor among many rather than as a distinct leading indicator. This paper proposes that process intelligence maturity, measured across four specific dimensions, may serve as a leading indicator of enterprise AI return-on-investment realization. The four dimensions are catalog completeness and hygiene, process model conformance, repository organization and discoverability, and process telemetry quality. Drawing on practitioner observation from large-scale process-architecture migrations, the pa-per develops five testable propositions and proposes a retrospective empirical methodology spanning 30 to 50 enterprises across at least four sectors. Measurement approaches for each dimension are specified, along with a regression-based analysis plan that controls for sector, organizational size, and use-case category. The paper discusses anticipated findings, validity threats inherent to retrospective designs, and implications for AI invest-ment decisions, readiness-framework development, and vendor platform strategy. The contribution is a testable reframing that positions process intelligence as measurable groundwork preceding AI value.
Keywords: Process Intelligence, Process Mining, AI Return On Investment, AI Readiness, Business Process Management, Conformance Checking, Process Telemetry, Enterprise Architecture.

Article history
Received: 05 July 2026
Accepted: 02 August 2026
Online: 13 August 2026

1. INTRODUCTION

Enterprise AI investment now spans core operations, customer-facing services, and internal au-tomation, and the capital committed continues to climb. Alongside that growth sits a persistent problem. A large share of initiatives reportedly stall before production or fail to demonstrate a return, a pattern documented across reviews of AI business value [1, 2]. Practitioner observation from large-scale process-architecture engagements is consistent with this picture. Programs with similar budgets and comparable models often diverge sharply in the value they realize. The ques-tion that motivates this paper is narrow and practical: why do enterprises with broadly similar

AI ambitions produce such uneven returns, and can any measurable property of the organization predict the difference in advance? One common answer points to readiness. Frameworks for AI readiness and organizational maturity have proliferated, assessing data quality, infrastructure, talent, funding, and governance as predictors of successful adoption [3, 4, 5]. These frameworks are useful, and they capture real constraints on value creation [3]. Their treatment of process knowledge, though, tends to be thin. Where process appears, it is usually folded into a broad process-readiness or operational category and weighted equally with a dozen other factors. The specific quality of an enterprise's process documentation, its conformance to reality, and its

Shah, M. (2026). Process Intelligence Maturity as a Leading Indicator of Enterprise AI Return on Investment. *South African Computer Journal* 38(2), 105–117.

ISSN 1015-7999 (print) ISSN 2313-7835 (online)

telemetry rarely receive separate measurement. This paper argues that this specific quality deserves isolation and study in its own right.

The argument draws on process mining and business process management, which provide tools to model processes, discover them from event logs, and check conformance between the two [6, 7]. That literature has matured considerably over the past decade. Vendors of process intelligence platforms now routinely claim that strong process foundations enable AI value, yet the claim is largely asserted rather than tested. Practitioner observation from structuring and delivering enterprise process-architecture migrations suggests the claim has merit but needs precision. Not all process foundations matter equally. The dimensions that matter can be named, and they can be measured before investment rather than diagnosed afterward.

This paper proposes process intelligence maturity as a distinct leading indicator of AI return on investment, defined across four dimensions: catalog completeness and hygiene, process model conformance, repository organization and discoverability, and process telemetry quality. Section 2 reviews AI readiness frameworks and process intelligence research. Section 3 develops the the-oretical framework and the four dimensions. Section 4 states five testable propositions. Section 5 proposes an empirical methodology. Sections 6 through 8 discuss anticipated findings, valid-ity threats, and practical implications, and Section 9 concludes. The framework's grounding is practitioner observation, not a completed study, and the propositions are offered deliberately for empirical test.

2. RELATED WORK

2.1 AI Readiness and Maturity Frameworks

Organizational readiness for AI has become a well-populated research area. Interview-based work identifies readiness factors spanning strategic alignment, data availability, skills, and cultural con-ditions, offering a catalog of what organizations should have in place before committing to AI [4]. Systematic reviews of AI maturity models find dozens of instruments, most structured as staged ladders from ad hoc experimentation to optimized enterprise capability [5]. Analyses of readiness in multinational corporations add dimensions around governance and ecosystem partnerships [8]. These contributions share a common shape. They enumerate many dimensions and treat them as broadly additive, which spreads analytical attention across the full list rather than concentrating it on the factors that most separate success from failure.

Comparative studies of digital maturity assessment reinforce the point. A systematic compar-ison across models finds substantial overlap in high-level dimensions and considerable divergence in how those dimensions are operationalized [9]. Process capability appears in many models, but usually as a single coarse construct. Recent work integrating AI and business process manage-ment proposes a maturity assessment framework for strategic adoption, mapping AI capabilities onto BPM stages [10]. This is closer to the present concern. It still treats process maturity as a general backdrop rather than decomposing the specific process-intelligence properties that might predict AI value. The gap is one of resolution, not of intent.

A parallel stream examines BPM maturity directly, arguing that assessment methods must evolve as digital transformation changes how processes are executed and sensed [11]. Work on creating AI business value through BPM capabilities suggests that process management capabil-ities mediate the relationship between AI adoption and realized value [12]. That claim aligns with practitioner observation from process-architecture migrations, where downstream automa-tion outcomes tracked closely with the state of the process estate inherited at the start. What remains open is measurement. If BPM capability shapes AI value, which measurable facets carry the effect, and can they be assessed before investment rather than after the fact?

The framework proposed here responds to that opening. Rather than adding another multi-dimensional readiness ladder, it isolates four process-intelligence properties and asks whether

Shah, M. (2026). Process Intelligence Maturity as a Leading Indicator of Enterprise AI Return on Investment. *South African Computer Journal* 38(2), 105–117.

ISSN 1015-7999 (print) ISSN 2313-7835 (online)

they, specifically, predict AI return. This narrows the target on purpose. The persona behind the framework is a solutions architect who has repeatedly scoped AI-assisted automation against enterprise process repositories. The recurring pattern in that work was that a handful of process-estate properties, not the full readiness checklist, separated smooth deployments from stalled ones. The next subsection turns to the process intelligence and process mining research that supplies the measurement tools for those four properties.

Table 1: Comparison of This Framework vs. Existing AI Readiness Frameworks

Aspect	**Typical AI Readiness Frameworks**	**Process Intelligence Maturity Framework**
Treatment of process	Single coarse process-readiness dimension among many	Four decomposed, separately measured dimensions
Primary data source	Self-report surveys and staged maturity ladders	Primary artifacts: catalogs, event logs, repository, telemetry
Predictive target	General AI adoption success	AI ROI realization within 24 months
Conformance and telemetry	Rarely measured explicitly	Measured directly; hypothesized strongest predictors
Intended use	Broad readiness benchmarking	Pre-investment diagnostic and remediation sequencing

1.1 Process Intelligence and Process Mining

Process mining provides methods to discover process models from event logs, check conformance between documented models and observed behavior, and analyze performance [6]. Business process management supplies the surrounding lifecycle of modeling, implementation, and monitoring [7]. Together they define what it means to know a process well enough to act on it. Reviews of process auditing show process mining moving from a niche technique into mainstream assurance and control work [13]. This maturation matters for AI, because the same event data and conformance signals that support auditing also define the substrate on which automation and AI agents operate in production.

A recurring theme in this literature is data quality. Work framing process-data quality as the true frontier of process mining argues that results are only as trustworthy as the event logs beneath them [14]. Detailed treatments of event log quality issues catalog problems of incompleteness, imprecise timestamps, and inconsistent case identifiers, along with remediation strategies [15]. These findings map directly onto the telemetry dimension of the proposed framework. Practitioner observation from delivering process-architecture migrations is consistent with them. Where telemetry was sparse or inconsistently captured, both process improvement and later automation efforts lost the measurement baseline needed to prove value.

Conformance checking has itself become a site of AI research, with recent surveys mapping how learning-based methods detect and explain deviations between models and logs [16]. This intersects the framework twice. Conformance quality is one of the four dimensions, and AI techniques increasingly depend on conformance signals to operate safely. Empirical study of business process models in public repositories, including large-scale analysis of models and clones on GitHub, shows wide variation in modeling quality, duplication, and maintenance [17]. That variation is precisely what catalog hygiene and repository discoverability attempt to measure. The research base therefore supplies both motivation and instruments for the four dimensions.

ISSN 1015-7999 (print) ISSN 2313-7835 (online)

What the literature does not yet offer is a test of whether these process-intelligence properties predict AI return on investment specifically. Process mining research establishes that model and log quality can be measured. BPM research establishes that process capability relates to AI value in general terms [12]. The connecting proposition, that specific measurable process-intelligence dimensions assessed before investment forecast AI ROI, remains unstated and untested. Practi-tioner observation from enterprise engagements motivates that proposition but cannot substitute for a study. The remainder of the paper formalizes the framework, states propositions, and proposes a methodology to close the gap.

3 THEORETICAL FRAMEWORK

The framework defines process intelligence maturity as the degree to which an enterprise can accurately describe, locate, verify, and measure its own processes. It comprises four dimensions, each with sub-measures and a theoretical link to AI return. The dimensions are not offered as exhaustive. They are the properties that practitioner observation from large-scale process- architecture migrations repeatedly identified as separating AI-ready process estates from unready ones. Table 2 summarizes the four dimensions, their sub-measures, and their hypothesized links to AI ROI. Each subsection below develops one dimension and the mechanism by which it may influence realized return.

Table 2: The Four Dimensions of Process Intelligence Maturity

Dimension	Sub-Measures	Theoretical Link to AI ROI
Catalog complete-ness and hygiene	Coverage ratio; metadata com-pleteness; version currency	Scoping accuracy; reduced build-time rework; rational portfolio se-lection
Process model confor-mance	Fitness scores; deviation fre-quency; conformance-evidence coverage	Deployment reliability; fewer pro-duction failures; realistic automa-tion targets
Repository organiza-tion and discoverabil-ity	Taxonomic consistency; search ca-pability; time-to-relevant-process	Scoping efficiency; dependency awareness; lower cumulative nav-igation cost
Process telemetry quality	Telemetry coverage; granularity; integrity	Outcome measurement baseline; training signal for learning-based components

a. Catalog Completeness and Hygiene

Catalog completeness and hygiene describes how fully an enterprise's processes are documented and how well that documentation is maintained. Sub-measures include coverage ratio, the share of operating processes with current documentation, metadata completeness covering ownership, inputs, outputs, and systems, and version currency, how recently documentation was reconciled with practice. In one large process-architecture migration, roughly 30,000 process maps were moved from iGrafx to ARIS, and the exercise made catalog condition unavoidably visible. Cover-age and metadata varied widely across domains. That variation, observed at scale, is the empirical seed of this dimension.

The theoretical link to AI return runs through scoping accuracy. AI and automation initia-tives begin by defining which process, or process fragment, the system will act on. When the catalog is incomplete or stale, scoping proceeds on assumptions, and teams discover missing steps, undocumented exceptions, or retired systems only during build. Each discovery forces rework,

ISSN 1015-7999 (print) ISSN 2313-7835 (online)

and rework erodes the return the initiative was projected to deliver. A complete, current catalog lets scoping proceed against reality. That shortens technical discovery and reduces the risk of building against a process that no longer runs the way its documentation claims.

Catalog hygiene also shapes portfolio decisions. Enterprises choose among candidate processes for AI investment, and that choice depends on comparable, trustworthy descriptions. Where meta-data is sparse, comparison relies on tacit knowledge held by a few staff, which biases selection to-ward familiar rather than valuable targets. Practitioner observation suggests that well-catalogued estates support more rational prioritization. The remediation sequencing that migration work re-quired, cleaning and standardizing before loading into an ARIS repository, mirrors the sequencing an AI program benefits from: establish catalog condition first, then scope the automation against it.

b. Process Model Conformance

Conformance describes the degree to which documented process models match how the process actually runs, as revealed in event logs. Process mining formalizes this through fitness and related measures that quantify how well a model reproduces observed behavior [6]. Sub-measures include fitness scores where logs exist, deviation frequency, and the share of processes with any conformance evidence at all. Documentation that looks authoritative but diverges from execution is a specific and common hazard. It stays invisible until measured, and it is most damaging when a team trusts it.

The link to AI return runs through deployment reliability. An AI or automation system built to a documented model inherits that model's inaccuracies. If the real process contains undocumented branches, the deployed system meets cases it was never designed for, and it fails, escalates, or produces wrong outputs. These failures surface in production, where they are most expensive to remediate. High conformance means the model the team builds against behaves like the process the system will meet, which raises the odds that a pilot survives contact with live operations and reaches sustained use.

Conformance is also where AI increasingly operates as both consumer and producer of sig-nal, since learning-based conformance methods now detect and explain deviations [16]. For an enterprise, conformance evidence assembled before investment offers a realistic read on which processes are stable enough to automate and which need remediation first. Practitioner obser-vation from migration work, where documented maps frequently diverged from executed reality, suggests conformance may carry more predictive weight than catalog completeness. A complete but inaccurate catalog can be more misleading than an incomplete one, because it invites false confidence in the scope.

c. Repository Organization and Discoverability

Repository organization describes how well documented processes can be found, navigated, and related to one another. Sub-measures include taxonomic consistency, search capability, and time-to-relevant-process, how long a knowledgeable user needs to locate the authoritative model for a given process. Configuring an ARIS repository to an enterprise BPMN standard was, in large part, an exercise in discoverability: naming conventions, hierarchy, and cross-references that let thousands of maps function as a navigable estate rather than a pile of files. Discoverability is easy to underrate until an estate grows past human memory.

The link to AI return runs through scoping efficiency and dependency awareness. AI initiatives rarely touch an isolated process. They intersect upstream and downstream steps, shared services, and controls. A well-organized repository lets teams trace those dependencies quickly, which reduces the chance that an automation optimizes one step while breaking an adjacent one. Where discoverability is poor, teams either spend technical discovery time reconstructing the map or

ISSN 1015-7999 (print) ISSN 2313-7835 (online)

proceed without full dependency awareness. Both outcomes raise cost and risk, and both depress the return an initiative ultimately records against its business case.

Discoverability compounds across an AI portfolio. Empirical analysis of process model repositories shows wide variation in structure, duplication, and maintenance, with clones and inconsistent organization common even in curated collections [17]. Duplicated or conflicting models force teams to guess which version is authoritative. Practitioner observation suggests the effect is cumulative. The first initiative in a disorganized estate pays a large navigation cost, and subse-quent initiatives keep paying it until the repository is remediated. Organized repositories let that cost be paid once, which is why discoverability behaves as an enabling condition for a program rather than a single project.

d. Process Telemetry Quality

Telemetry quality describes how well process execution is instrumented and captured as usable event data. Sub-measures include telemetry coverage, the share of process steps that emit events, granularity, the level of detail captured, and integrity, the accuracy and consistency of timestamps, case identifiers, and outcomes. Process mining research treats this substrate as decisive, framing process-data quality as the frontier that governs whether any downstream analysis can be trusted [14], and cataloging the log-quality defects that undermine it [15]. Poor telemetry does not announce itself. It quietly invalidates the analysis built on top of it.

The link to AI return runs through outcome measurement and learning. AI value must be measured to be defended, and measurement requires a telemetry baseline captured before and after deployment. Without it, an initiative cannot show whether cycle time fell, exceptions dropped, or quality improved, so even genuine gains go unproven and unfunded at the next budget cycle. Good telemetry supplies that baseline. It also feeds the models, since many AI-assisted automations learn from historical execution data, and sparse or corrupt telemetry starves them of the training signal they need.

Telemetry quality may be the dimension most tightly coupled to modern AI methods. Techniques that rely on historical event streams degrade quietly when integrity is poor, producing plausible but wrong outputs. Practitioner observation from process-architecture engagements suggests telemetry is often the weakest of the four dimensions in practice, because instrumentation is treated as an operational afterthought rather than an architectural requirement. This combination, high importance and low typical maturity, is why the framework anticipates teleme-try quality and conformance carrying the strongest predictive weight of the four dimensions.

4 TESTABLE PROPOSITIONS

The framework yields propositions that connect each dimension to AI return on investment and one proposition about their combination. Return is defined as realized value from an AI initiative within 24 months of investment, measured through financial, operational, and adoption indicators specified in Section 5. Each proposition is directional, predicting a positive association. The propositions are stated so that they can be supported or contradicted by the proposed study. Table 3 summarizes them alongside their constructs and hypothesized directions, including the ordering the framework expects among the four dimensions.

Propositions P1 through P4 address the dimensions individually. P1 holds that catalog completeness and hygiene at investment time positively predicts AI ROI realization within 24 months. P2 holds the same for process model conformance. P3 holds the same for repository organization and discoverability. P4 holds the same for process telemetry quality. Each proposition isolates one dimension while the others are held constant statistically, which lets the study estimate whether a dimension contributes independent predictive power rather than

Shah, M. (2026). Process Intelligence Maturity as a Leading Indicator of Enterprise AI Return on Investment. *South African Computer Journal* 38(2), 105–117.

ISSN 1015-7999 (print) ISSN 2313-7835 (online)

merely co-varying with the rest of the process estate.

Proposition P5 concerns combination. It holds that the four dimensions together predict AI ROI more strongly than any single dimension alone, and more strongly than a general AI readiness framework score applied to the same enterprises. This proposition tests the paper's central claim, that process intelligence deserves isolation as a leading indicator. If P5 fails while some of P1 through P4 hold, the framework would survive in weaker form as a set of individually useful measures without joint advantage. That distinction matters for how practitioners would actually use the framework in an investment decision.

The propositions carry an implicit ordering hypothesis drawn from practitioner observation. Conformance and telemetry quality are expected to show stronger associations than catalog com-pleteness and repository organization, because the former govern whether a deployed system behaves correctly and whether its value can be measured, while the latter govern the efficiency of getting there. This ordering is not asserted as established. It is a prediction offered for test, and its confirmation or reversal would itself inform how enterprises sequence remediation before AI investment.

5 EMPIRICAL METHODOLOGY

The propositions call for a study that the framework's author, as a solutions architect, proposes rather than has executed. The design below is retrospective and observational, chosen because randomized assignment of process-intelligence maturity across enterprises is infeasible. It trades causal strength for realism and access to real programs with known outcomes. Table 4 summarizes the proposed sample design. The subsections that follow specify the sample, the measurement of both maturity and return, and the statistical analysis plan that connects them to the propositions.

Table 3: Testable Propositions Summary

Proposition	Construct	Hypothesized Direction
P1	Catalog completeness and hygiene to AI ROI within 24 months	Positive
P2	Process model conformance to AI ROI	Positive (expected stronger)
P3	Repository organization and discover-ability to AI ROI	Positive
P4	Process telemetry quality to AI ROI	Positive (expected stronger)
P5	Four dimensions combined vs. any single dimension or general readiness score	Positive; combined exceeds any single and a general readiness score

Table 4: Proposed Sample Design

Sector	Target Enter-prise Count	Inclusion Criteria
Financial services	At least 5	AI program at least 2 years old; process artifacts available
Healthcare	At least 5	AI program at least 2 years old; process artifacts available
Manufacturing	At least 5	AI program at least 2 years old; process artifacts available
Government	At least 5	AI program at least 2 years old; process artifacts available

Shah, M. (2026). Process Intelligence Maturity as a Leading Indicator of Enterprise AI Return on Investment. *South African Computer Journal* 38(2), 105–117.

ISSN 1015-7999 (print) ISSN 2313-7835 (online)

Additional cros s- sector	To reach 30 to 50 total	Deliberate variation in outcome, including under-performing programs

a. Sample Design

The proposed sample comprises 30 to 50 enterprises with documented AI investment programs at least two years old, recruited through industry partnerships and advisory relationships. The two-year minimum ensures that return, or its absence, has had time to materialize. Enterprises would be selected to span at least four sectors, financial services, healthcare, manufacturing, and government, with at least five enterprises per sector. That spread lets sector-specific effects be distinguished from the process-intelligence effects of interest, which is essential if the framework is to generalize beyond a single industry.

Sample size reflects a balance between statistical power and the depth each case demands. Assessing four maturity dimensions from primary artifacts is labor-intensive, which caps feasible sample size, while regression with several predictors and controls needs enough cases to estimate coefficients stably. Thirty to fifty enterprises sits at the lower bound of comfortable regression practice, so the analysis plan favors parsimonious models and treats findings as indicative. Cross-sector recruitment guards against a result that reflects one industry's idiosyncrasies rather than a general pattern across process estates.

Recruitment through industry partnerships is pragmatic and introduces selection bias, addressed in Section 7. Enterprises that agree to expose process-architecture artifacts and AI outcomes tend to be those with something worth showing. The design mitigates this by seeking variation in outcome, deliberately including programs that underperformed, and by measuring maturity from artifacts rather than self-report where possible. Comparable investment-appraisal studies of process automation demonstrate that outcome data of this kind can be assembled from enterprise records with care [18, 19].

b. Measurement

Each maturity dimension is measured at investment time, reconstructed from artifacts that pre-date or coincide with the AI program: process catalogs, event logs, repository exports, and telemetry configurations. Catalog completeness is scored from coverage and metadata against the enterprise's own process inventory. Conformance is scored from available fitness evidence and deviation records. Repository organization is scored from taxonomy and retrieval tests. Telemetry quality is scored from log coverage, granularity, and integrity checks, following defect taxonomies established in the process-data-quality literature [14, 15].

Scoring uses a documented rubric applied by two independent coders, with inter-rater reliabil-ity reported and disagreements reconciled. Dual coding is necessary because several sub-measures require judgment, and single-coder scores would leave reliability unknown. The rubric converts qualitative artifact review into ordinal dimension scores, which the analysis treats as interval-approximate, a common and acknowledged compromise. Anchoring scores to enterprise-specific baselines, rather than an absolute scale, keeps a large and complex estate from being penalized relative to a small and simple one during comparison.

AI ROI is measured two to three years post-investment through three families of indicator: financial, covering cost reduction and revenue contribution, operational, covering cycle time, throughput, and error rate, and adoption, covering usage and sustained deployment. Multiple indicators are used because single-metric ROI invites gaming and misses value that appears in one register but not another. Where enterprises report figures, conservative adjustments and dis-counts are applied, and the study privileges operational and adoption measures that are

Shah, M. (2026). Process Intelligence Maturity as a Leading Indicator of Enterprise AI Return on Investment. *South African Computer Journal* 38(2), 105–117.

ISSN 1015-7999 (print) ISSN 2313-7835 (online)

harder to inflate than self-reported financials [18].

c. Statistical Analysis

The primary analysis is multiple regression with AI ROI as the dependent variable and the four maturity dimensions as independent variables, controlling for sector, organizational size, and use-case category. This specification tests P1 through P4 as individual coefficients and supports P5 through comparison of the full model against single-predictor models and against a general readiness-score model fitted to the same cases. Effect sizes and confidence intervals are reported rather than significance alone, given the modest sample and the risk of over-reading a single threshold.

Subgroup analysis examines whether associations hold within sectors and within use-case categories, since a pooled result could mask heterogeneity. The use-case categories distinguish document-centric automation from decision-support and from agentic AI workflows, because the dependence on process intelligence may differ across them. Sample size limits subgroup power, so subgroup findings are framed as hypothesis-generating. Regularized regression is held in reserve as a contingency if predictor collinearity proves high, which is plausible given that mature enterprises may score well across several dimensions at once.

The analysis plan anticipates collinearity among the four dimensions and treats it as substantively informative rather than merely a nuisance. If dimensions co-vary strongly, that pattern supports treating process intelligence maturity as a coherent construct, consistent with proposition P5. Diagnostic checks cover influential cases, given the small sample, and robustness is assessed by refitting the model without each sector in turn. The plan is deliberately conservative, reporting a testable structure rather than promising a definitive causal estimate the design cannot support.

6 ANTICIPATED FINDINGS

The study is expected to find process intelligence maturity a significant predictor of AI return, with the composite of four dimensions outperforming any single dimension and outperforming a general readiness score. This anticipation follows both from the theoretical mechanisms in Section 3 and from practitioner observation across process-architecture engagements. The prediction is falsifiable. A null result on P5 would meaningfully constrain the framework, and the design is intended to make such a result visible rather than to obscure it behind favorable specification choices.

Within the composite, conformance and telemetry quality are expected to carry stronger predictive weight than catalog completeness and repository organization. The reasoning is mechanis-tic. Conformance governs whether a deployed system behaves correctly against the live process, and telemetry governs whether its value can be measured and whether learning-based components have signal to work from. Research on generative and learning-based automation underscores how heavily such systems depend on execution data of adequate quality [20, 21]. Catalog and repos-itory quality matter, but chiefly by improving efficiency rather than correctness of the deployed system.

If the anticipated ordering holds, the practical implication is a sequencing rule. Enterprises would gain more, per unit of remediation effort, by improving conformance evidence and telemetry instrumentation before or alongside AI investment than by first perfecting catalogs and reposito-ries. This inverts a common instinct to begin with visible documentation cleanup. Practitioner observation is consistent with the inversion. The estates that automated smoothly were not al-ways the best documented, but they were the ones whose documentation matched reality and whose execution was well instrumented.

For researchers, confirmation would justify reframing AI readiness assessment to weight process-

ISSN 1015-7999 (print) ISSN 2313-7835 (online)

intelligence dimensions more heavily and to measure them specifically rather than folding them into a general category [12]. For vendors, the dimensional structure offers a basis for differenti-ating platforms by how well they support conformance evidence and telemetry capture, not only modeling and storage. For advisory practice, the framework offers a defensible way to tell a client which process foundations to fix first. These implications are developed further in Section 8.

7 LIMITATIONS AND VALIDITY THREATS

The most consequential limitation is causal. A retrospective, observational design can establish association between process-intelligence maturity and AI return, but it cannot by itself establish that maturity caused the return. Unmeasured common causes are plausible. A disciplined oper-ating culture might produce both mature process intelligence and successful AI programs. The study can reduce this threat through controls and robustness checks, and it can report associa-tion honestly, but it cannot fully resolve causation. Claims are therefore framed as predictive and associational rather than causal.

Selection bias is a second threat. Enterprises that volunteer process-architecture artifacts and AI outcomes are not a random sample, and they may skew toward more capable organizations. This can compress variance in the maturity measures and attenuate observed associations, biasing the study toward a null rather than toward false confirmation, which is the safer direction for a first test. Deliberate inclusion of underperforming programs and cross-sector recruitment mitigate the threat without eliminating it, and the sampling frame and its limits would be reported in full. Outcome measurement depends partly on enterprise-reported data, which is vulnerable to inconsistent definitions and favorable framing. The design counters this by using multiple indi-cator families, privileging operational and adoption measures that resist inflation, and applying conservative discounts to self-reported financials [18]. Reconstructing maturity from historical artifacts also risks recall and survivorship effects, since artifacts from failed programs may be incomplete. Dual coding, documented rubrics, and sensitivity analyses address measurement reliability, though residual error remains and would be acknowledged plainly in any reporting.

Construct validity deserves scrutiny as well. The four dimensions are proposed as the process-intelligence properties that matter most, but the set is grounded in practitioner observation rather than prior confirmatory work, so it may omit a relevant dimension or over-weight another. The proposed study can partially test this by examining whether the four dimensions cohere and whether they collectively outperform alternatives. Confirming the framework's completeness would require replication across independent samples and analysts, which a single study cannot provide on its own.

8 PRACTICAL IMPLICATIONS

For enterprises weighing AI investment, the framework offers a pre-investment diagnostic. A quantitative maturity assessment across the four dimensions could inform which initiatives to fund, which to defer pending remediation, and where remediation effort would raise the odds of return most. This reframes a decision often made on vendor promises and strategic enthusiasm into one informed by a measurable property of the enterprise's own process estate. The diagnostic is most useful before commitment, when its findings can still change the plan rather than explain a disappointment.

For advisory and consulting practice, the framework structures work that is often done by intuition. Practitioner observation from process-architecture migrations suggests that experienced architects already sense which estates are AI-ready, but that judgment is hard to transfer or defend. Turning it into scored dimensions with explicit sub-measures makes the assessment repeatable and

Shah, M. (2026). Process Intelligence Maturity as a Leading Indicator of Enterprise AI Return on Investment. *South African Computer Journal* 38(2), 105–117.

ISSN 1015-7999 (print) ISSN 2313-7835 (online)

reviewable. It also supports remediation sequencing, giving a client a defensible order of operations: address conformance and telemetry gaps first if the anticipated findings hold, then catalog and repository quality.

For vendors of process intelligence and automation platforms, the dimensional structure suggests a positioning and roadmap logic. Platforms could be evaluated and differentiated by how well they support each dimension, particularly conformance evidence and telemetry capture, which are less commodified than modeling and storage. A vendor that helps customers raise their weak-est dimension addresses a bottleneck the framework identifies as consequential for AI value. This aligns product strategy with the property most likely to gate customer success rather than with feature breadth alone.

For the enterprise architecture function specifically, the framework connects process-architecture work to AI outcomes in language executives fund. Repository migrations and BPMN standard-ization are frequently justified on governance grounds and struggle to show return. Positioning them as raising process intelligence maturity, a measurable leading indicator of AI ROI, reframes that investment as groundwork for the AI agenda the enterprise is already funding. Practitioner observation suggests this framing resonates precisely because it ties unglamorous foundational work to the outcomes leadership is already watching closely.

9 CONCLUSION

Enterprise AI investment continues to outpace the ability to predict which initiatives will return value, and existing readiness frameworks, useful as they are, spread their attention across many dimensions and rarely isolate process intelligence. This paper has argued that process intelligence maturity, defined across catalog completeness, conformance, repository discoverability, and telemetry quality, may function as a distinct leading indicator of AI return. The argument is grounded in process mining and BPM research and in practitioner observation from large-scale process-architecture migrations, and it is offered as a structure to be tested rather than a finding to be accepted.

The contribution is a testable reframing. Four dimensions with explicit sub-measures, five directional propositions, and a retrospective methodology give researchers a concrete way to confirm or refute the claim, including the prediction that conformance and telemetry quality carry more weight than catalog and repository properties. For practitioners, the same structure supplies a prioritization basis: assess the process estate before investing, and sequence remediation toward the dimensions most likely to gate return. The framework's value depends on the study being done, and the design is specified so that it can be carried out.

The stakes justify the effort. Capital committed to enterprise AI is large and rising, and the gap between committed and realized value represents a substantial economic loss that a reliable leading indicator could help close. Measuring process intelligence maturity before investment will not guarantee return, and no single indicator should be asked to. It may, however, move some decisions from hope toward evidence, and it may give the unglamorous work of process-architecture stewardship the standing it has lacked. That reframing, if the propositions survive test, is the paper's intended mark.

References

[1] Ida Merete Enholm, Emmanouil Papagiannidis, Patrick Mikalef, and John Krogstie. Artificial intelligence and business value: A literature review. Information Systems Frontiers, 24(5): 1709–1734, 2022. doi: 10.1007/s10796-021-10186-w.

[2] Nikolaos-Alexandros Perifanis and Fotis Kitsios. Investigating the influence of artificial intelligence on business value in the digital era of strategy: A literature review. Information, 14(2):85, 2023. doi: 10.3390/info14020085.

ISSN 1015-7999 (print) ISSN 2313-7835 (online)

[3] Amit Mitra, Anantha Seetharaman, and Kongath Maddulety. Adoption of artificial in-telligence enabling value creation in business. In 2024 4th International Conference on Robotics, Automation, and Artificial Intelligence (RAAI), pages 213–220, Singapore, 2024. doi: 10.1109/raai64504.2024.10949520.

[4] Jan Jo¨hnk, Marius Weißert, and Katrin Wyrtki. Ready or not, AI comes: An interview study of organizational AI readiness factors. Business & Information Systems Engineering, 63(1):5–20, 2021. doi: 10.1007/s12599-020-00676-7.

[5] Ramla Bint Sadiq, Nurhizam Safie, Ahmad Hanis Abd Rahman, and Shidrokh Goudarzi. Ar-tificial intelligence maturity model: A systematic literature review. PeerJ Computer Science, 7:e661, 2021. doi: 10.7717/peerj-cs.661.

[6] Wil M. P. van der Aalst. Process Mining: Data Science in Action. Springer, Berlin, 2nd edition, 2016. doi: 10.1007/978-3-662-49851-4.

[7] Marlon Dumas, Marcello La Rosa, Jan Mendling, and Hajo A. Reijers. Fundamentals of Business Process Management. Springer, Berlin, 2nd edition, 2018. doi: 10.1007/978-3-662-56509-4.

[8] Aida Nikoumaram Tehrani, Sourav Ray, Sanjit Kumar Roy, Richard L. Gruner, and Francesco Appio. Decoding AI readiness: An in-depth analysis of key dimensions in multi-national corporations. Technovation, 131:102948, 2024. doi: 10.1016/j.technovation.2023. 102948.

[9] Baptiste Cognet, Jean-Philippe Pernot, Louis Rivest, and Christophe Danjou. Systematic comparison of digital maturity assessment models. Journal of Industrial and Production Engineering, 40(7):519–537, 2023. doi: 10.1080/21681015.2023.2242340.

[10]Ouazzani Ayyoub, Sabri Hanae, Zellou Ahmed, and Doumi Karim. Integrating AI and busi-ness process management: A maturity assessment framework for strategic adoption. Procedia Computer Science, 278:264–272, 2026. doi: 10.1016/j.procs.2026.02.462.

[11]Marek Szelagowski and Justyna Berniak-Wo´zny. How to improve the assessment of BPM ma-turity in the era of digital transformation. Information Systems and e-Business Management, 20(1):171–198, 2022. doi: 10.1007/s10257-021-00549-w.

[12]Andrej Zebec and Mojca Indihar Sˇtemberger. Creating AI business value through BPM capabilities. Business Process Management Journal, 30(8):1–26, 2024. doi: 10.1108/BPMJ-07-2023-

[13]Muhammad Imran, Suriayati Hamid, and Maizatul Akmar Ismail. Advancing process audits with process mining: A systematic review of trends, challenges, and opportunities. IEEE Access, 11:68340–68357, 2023. doi: 10.1109/ACCESS.2023.3292117.

[14]Arthur H. M. ter Hofstede et al. Process-data quality: The true frontier of process mining.

[15]Journal of Data and Information Quality, 15(3):29, 2023. doi: 10.1145/3613247.

[16]Duˇsan Daki´c, Tanja Vukovi´c, Milica Zˇiˇzakov, and Bojan Stevanov. Event log data quality issues and solutions. Mathematics, 11(13):2858, 2023. doi: 10.3390/math11132858.

[17]Laura Genga and Karolin Winter. Artificial intelligence in conformance checking: State of the art and research agenda. Process Science, 2:9, 2025. doi: 10.1007/s44311-025-00015-7.

[18]Mahsa Saeedi Nikoo, Shalini Kochanthara, O¨ zgu¨r Babur, and Mark van den Brand. An

[19]empirical study of business process models and model clones on GitHub. Empirical Software Engineering, 30(2):48, 2025. doi: 10.1007/s10664-024-10584-z.

[20]Antti Yla¨-Kujala, Dariusz Kedziora, Lauri Metso, Timo Ka¨rri, Ari Happonen, and Wojciech Piotrowicz. Robotic process automation deployments: A step-by-step method to investment appraisal. Business Process Management Journal, 29(8):163–187, 2023. doi: 10.1108/BPMJ-08-2022-0418.

[21]Muhammad Chandra Al Ayyubi et al. Optimizing core banking operation's ROI with robotic process automation: A case study from a leading southeast asian bank. In 2024 IEEE International Conference on Cybernetics and Innovations in Intelligent Digital Ecosystems (CERIA), pages 1–6, Bandung, Indonesia, 2024. doi: 10.1109/ceria64726.2024.10915031.

Shah, M. (2026). Process Intelligence Maturity as a Leading Indicator of Enterprise AI Return on Investment. *South African Computer Journal* 38(2), 105–117.

ISSN 1015-7999 (print) ISSN 2313-7835 (online)

[22]Stefan Feuerriegel, Jochen Hartmann, Christian Janiesch, and Patrick Zschech. Generative AI. Business & Information Systems Engineering, 66(2):111–126, 2024. doi: 10.1007/s12599-023-00834-7.

[23]Sadia Afrin, Shermin Roksana, and Rakib Akram. AI-enhanced robotic process automation: A review of intelligent automation innovations. IEEE Access, 13:173–197, 2025. doi: 10. 1109/ACCESS.2024.3513279.

[24]Nessrine Omrani, Narjes Rejeb, Adnane Maalaoui, Marina Dabi´c, and Sascha Kraus. Drivers of digital transformation in SMEs. IEEE Transactions on Engineering Management, 71:5030–5043, 2022. doi: 10.1109/TEM.2022.3215727.

ISSN 1015-7999 (print) ISSN 2313-7835 (online)