

EXPLAINABLE AI (XAI) MODELS FOR CLOUD-BASED BUSINESS INTELLIGENCE: ENSURING COMPLIANCE AND SECURE DECISION-MAKING

Md Muzahidul Islam¹; Md Mohaiminul Hasan²;

[1]. B.Sc in Computing Science and Technology, Jiangxi Normal University, Jiangxi, China
Email: muzahidul365@gmail.com

[2]. Project Analyst; Quantanite, Dhaka, Bangladesh;
Email: mohaiminul.hasan22@gmail.com

Doi: [10.63125/5etfhh77](https://doi.org/10.63125/5etfhh77)

Received: 23 June 2023; Revised: 29 July 2023; Accepted: 28 August 2023; Published: 30 September 2023

Abstract

This study had investigated Explainable AI (XAI) models within cloud-based business intelligence (BI) systems to determine how explain ability quality supported regulatory compliance assurance and secure decision-making. A quantitative comparative design had been implemented in a simulated cloud BI pipeline, and 210 cross-functional participants had completed four BI tasks (demand forecasting, fraud/risk scoring, anomaly detection, and resource optimization) under intrinsic, post-hoc, and hybrid XAI conditions. Descriptive results had shown that hybrid XAI produced the strongest explain ability Quality (EQ) profile, with higher fidelity ($M=4.31$, $SD=0.49$), stability ($M=4.18$, $SD=0.52$), and human agreement ($M=4.12$, $SD=0.50$) than intrinsic and post-hoc conditions, while post-hoc methods displayed the widest stability dispersion ($SD=0.70$). Compliance Assurance (CA) was strongest under hybrid XAI for audit traceability ($M=4.24$, $SD=0.50$) and decision reproducibility ($M=4.16$, $SD=0.54$), whereas post-hoc models showed weaker reproducibility ($M=3.61$, $SD=0.69$) and higher fairness deviation ($M=2.98$, $SD=0.73$). Secure Decision-Making (SDM) outcomes followed the same direction, with hybrid XAI yielding higher robust decision integrity ($M=4.20$, $SD=0.51$) and adversarial detection ($M=4.05$, $SD=0.58$) and post-hoc explanations showing higher leak risk ($M=3.42$, $SD=0.71$). Regression analysis had indicated that EQ significantly predicted CA ($R^2=.534$, $\Delta R^2=.318$), with stability ($\beta=.31$, $p<.001$) and fidelity ($\beta=.24$, $p<.001$) as dominant contributors. EQ also predicted SDM ($R^2=.460$), led by stability ($\beta=.28$, $p<.001$) and fidelity ($\beta=.21$, $p<.001$), while sparsity/complexity showed a small negative SDM effect ($\beta=-.09$, $p=.041$). Mediation testing had confirmed a partial indirect pathway through CA ($\beta_{\text{indirect}}=.17$, $p=.003$). Overall, high-quality XAI had functioned as a measurable governance mechanism that improved compliance readiness and secure BI decision reliability in cloud environments.

Keywords

Explainable AI, Cloud BI, Compliance, Security, Governance.

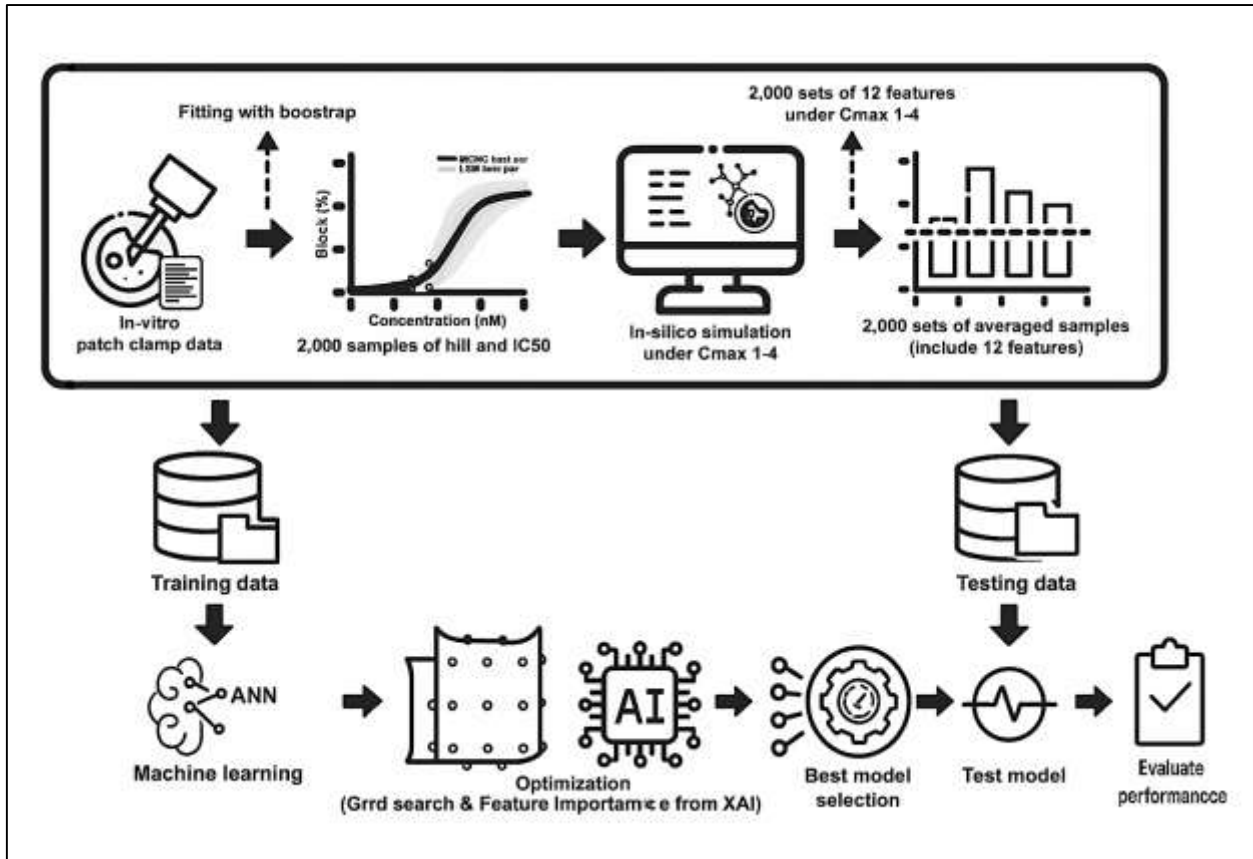
INTRODUCTION

Explainable artificial intelligence (XAI) refers to a set of concepts, models, and techniques that make algorithmic reasoning understandable to humans by clarifying how inputs relate to outputs in a specific decision context (Minh et al., 2022). At its core, XAI seeks to reduce opacity in machine learning by providing rationales that are traceable, communicable, and verifiable within real organizational workflows. Interpretability and explainability are often treated as related but distinct properties: interpretability indicates how readily a human can comprehend a model's structure or logic, while explainability concerns the production of meaningful accounts of why a particular prediction or recommendation occurred. XAI research commonly distinguishes intrinsic models, whose internal mechanisms are directly readable, from post-hoc methods, which generate explanations after a complex model has produced an output. Intrinsic approaches include linear and generalized additive models, sparse scoring systems, decision trees, rule-based learners, monotonic gradient boosting, interpretable neural architectures, and explainable boosting machines. Post-hoc approaches include surrogate explanation models, feature attribution methods, counterfactual explanations, example-based reasoning, concept-based explanations, and visualization-driven narration of learned representations. Across these categories, a consistent definitional aim is to enable stakeholders to answer practical questions: which variables mattered, in what direction, to what degree, and under what conditions (Adadi & Berrada, 2018). A wide set of foundational and empirical studies have shown that explanation quality can be operationalized through measurable properties such as fidelity to the original model, stability under minor input shifts, completeness of reasoning coverage, computational efficiency, and human-agreement in controlled tasks. Other studies demonstrate that explanations function as cognitive tools; they support users in diagnosing errors, differentiating signal from noise, and identifying spurious correlations that might otherwise be hidden in high-dimensional learning. Additional experimental work shows that explanations influence trust calibration rather than simply increasing trust, helping users accept correct recommendations and reject flawed ones. In organizational analytics, XAI is defined not only by algorithmic technique but also by its role in governance: explanations are treated as artifacts that accompany predictions throughout their lifecycle, allowing audit, challenge, and responsible override (Vilone & Longo, 2021). These definitional strands position XAI as a measurable bridge between statistical learning and accountable decision practice in settings where models are embedded in daily management operations.

Cloud-based business intelligence (BI) is the delivery of analytical capabilities through cloud infrastructure and services to transform data into actionable insight for organizations. BI itself is traditionally defined as a systematic process of collecting, integrating, storing, analyzing, and visualizing data to support descriptive, diagnostic, and, increasingly, predictive decision-making. Cloud BI extends this process by hosting data pipelines, warehouses, lakehouses, model training, and interactive dashboards on scalable cloud platforms (Páez, 2019). The defining features of cloud BI include elasticity of compute and storage, consumption-based pricing, rapid provisioning of analytical environments, and ubiquitous access for distributed teams. Cloud BI ecosystems typically integrate data from enterprise systems such as ERP and CRM, operational sensors, digital platforms, and external partner feeds, assembling a unified analytical layer for performance monitoring and strategic planning. A large body of empirical research on cloud adoption finds that these architectures improve time-to-insight, enable cross-regional collaboration, and support real-time or near-real-time analytics at scales that on-premises infrastructures struggle to match. Studies of modern BI practices also show a shift from static reporting to decision automation, where machine learning models continuously score risks, forecast demand, detect anomalies, and optimize resource allocation. When AI is embedded into BI dashboards, the definition of insight changes from a human-crafted interpretation of historical patterns to a model-generated estimate of future or hidden states of the business process. That shift introduces a new definitional dependency: BI insights become only as responsible as the models that generate them and the interpretability of those outputs for decision owners (Longo et al., 2020). Multiple studies on AI-enabled BI confirm that black-box models can deliver accuracy gains in forecasting and classification, yet they also reduce the ability of managers, auditors, and regulators to understand the drivers of key indicators. Other organizational studies show that BI value depends on explainable reasoning because decision makers must justify actions to internal committees, customers, and

oversight bodies. This makes explainability a definitional requirement for AI-driven cloud BI, where insights circulate through shared dashboards across departments and countries. In this framing, XAI models are not separate tools but core analytical components that allow cloud BI to remain an interpretable decision system rather than a purely automated prediction engine (Islam et al., 2022).

Figure 1: XAI-Enabled Cloud BI Governance Framework



Compliance in cloud-based BI denotes adherence to laws, regulations, and standards that govern data processing and the use of automated decision systems. In global business environments, compliance spans privacy statutes, sector-specific rules, corporate governance frameworks, and emerging AI-focused requirements (Machlev et al., 2022). Privacy regulations impose obligations related to lawful data use, minimization, transparency to data subjects, and safeguards for sensitive attributes. Sector rules in finance, healthcare, telecommunications, public services, and critical infrastructure prescribe additional duties around risk modeling, accountability, documentation, and traceable decision justification. Governance standards further specify control objectives for auditability, retention, access monitoring, and change management in data and model pipelines. A broad range of studies on regulatory technology and enterprise governance show that compliance in analytics is no longer limited to where data resides; it includes how decisions are produced, whether bias is controlled, and whether decision logic can be demonstrated during review. Research on algorithmic accountability documents that organizations face escalating scrutiny when AI systems influence credit approval, insurance pricing, hiring, fraud detection, supply-chain qualification, or medical triage. Empirical studies in these domains reveal that inability to explain model reasoning increases legal exposure, slows adoption, and can trigger regulatory sanctions. Experimental and case-based studies also show that compliance officers and auditors require evidence that a model’s behavior aligns with policy constraints, including non-discrimination requirements, proportionality, and documented rationale for exceptions (Tjoa & Guan, 2020). XAI contributes directly to measurable compliance by producing explanation records that can be aligned with audit checklists and fairness assessments. Multiple studies on interpretability tools identify feature attribution, monotonic effects, rule extraction, and counterfactual reasoning as explanation forms that auditors can verify systematically. Other studies

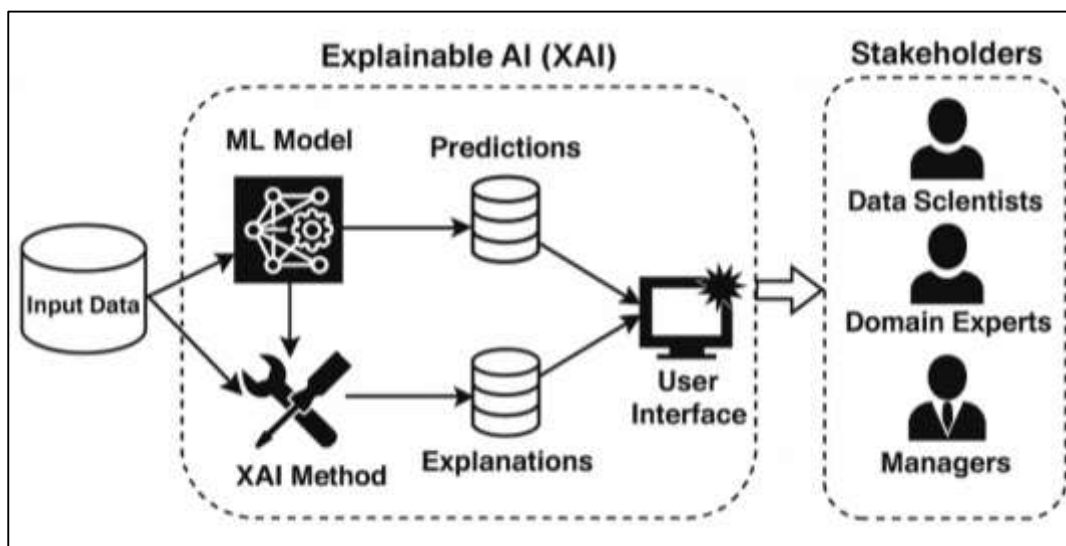
show that explanation logs support model risk management by allowing investigators to replay decisions, evaluate whether the same inputs would yield consistent outputs, and identify drift that might violate approved operating thresholds (Merry et al., 2021). Research on human oversight in automated decision systems indicates that compliance improves when explanations are embedded into governance workflows for approval, escalation, and redress. Quantitative findings across industries further indicate that explanation quality correlates with reduced dispute rates and faster audit cycles because decision logic is accessible without reverse engineering the model. In cloud BI, where models are updated often and shared widely, compliance is therefore a continuous property supported by XAI that links model behavior to regulatory expectations in measurable, reviewable ways.

Secure decision-making within cloud BI refers to producing and acting on analytical outputs while preserving confidentiality, integrity, availability, and resistance to manipulation. Cloud environments introduce definitional security concerns around shared infrastructure, multi-tenant access, API exposure, and configuration complexity. Studies in cloud security consistently show that misconfiguration, weak identity management, and insufficient monitoring can lead to data leakage or service disruption (Abdulla & Ibne, 2021; Lötsch et al., 2021). AI models extend that risk surface by introducing threats specific to learning systems, including data poisoning during training, adversarial perturbations at inference, membership inference that reveals training records, and model extraction that exposes valuable intellectual property. Security-focused machine learning research demonstrates that these attacks can degrade BI reliability by altering predictions used for fraud alerts, demand forecasts, or risk classifications (Ara, 2021). XAI intersects with secure decision-making through both benefit and burden. On the benefit side, explanations provide visibility into model reasoning, enabling analysts to detect anomalous feature influences, spurious drivers, or sudden shifts in attribution patterns that signal adversarial activity or data pipeline compromise (Habibullah & Foysal, 2021; Knapič et al., 2021). Empirical studies in explainable cybersecurity show that rationale-augmented alerts improve analyst accuracy and reduce response time because teams can prioritize threats based on interpretable evidence rather than opaque scores. Reliability studies also indicate that stable explanations align with robustness; if explanations fluctuate wildly under small input changes, the decision is less secure because it suggests fragile reasoning (Sarwar, 2021). On the burden side, security and privacy studies show that overly rich explanations can leak sensitive business logic or personal attributes, turning transparency into a vector for exploitation. This creates a definitional need for controlled explainability, where explanation access is governed through role-based policies, encryption, rate limits, and logging (Musfiqur & Saba, 2021). Additional enterprise studies note that secure decisions depend on user comprehension; explanations that overwhelm or mislead managers can create operational risk even if the model is accurate (Combi et al., 2022; Redwanul et al., 2021). Therefore, secure cloud BI with XAI is defined by a balance: explanations must be informative enough to validate decisions and defend against manipulation, yet constrained enough to prevent disclosure or adversarial probing (Reza et al., 2021). This security-explainability coupling becomes central in multinational dashboards where decisions affect customers and partners across different risk contexts (Saikat, 2021).

Quantitative research on XAI models for cloud BI relies on measurable constructs that connect explanation properties to compliance assurance and decision security. Explainability can be operationalized through objective metrics such as local fidelity, which assesses how closely an explanation approximates model behavior near a specific instance; global fidelity, which measures agreement across the broader data distribution; stability, which measures invariance of explanations under small perturbations; and sparsity, which captures how compact an explanation is (Shaikh & Aditya, 2021; Tocchetti & Brambilla, 2022). Multiple benchmark studies on interpretability methods show that high local fidelity is essential for reliable audits because low-fidelity surrogates can misrepresent the true drivers of a decision. Other comparative studies find that stability is crucial in operational BI, where models retrain frequently and explanations must remain consistent to preserve compliance evidence over time. Completeness metrics, assessed through feature deletion and insertion tests, quantify whether an explanation covers the majority of predictive signal rather than highlighting only a small misleading subset (Amin, 2022). Human-grounded quantitative studies also link explanation forms to user performance, measuring how explanations affect decision accuracy, error

detection rates, time-to-decision, and calibrated trust (Ahmed et al., 2022; Ariful, 2022). In enterprise BI experiments, users provided with structured explanations have been shown to more accurately identify when a model is likely wrong and to avoid over-reliance on automation. Quantitative governance studies further demonstrate that explanation artifacts can be integrated into model monitoring systems, enabling the detection of “explanation drift,” a measurable deviation in attribution patterns that indicates data shift or emerging bias. Architecture-level performance studies in cloud settings measure explanation latency, computing cost, scalability across parallel queries, and isolation between tenants (Ariful & Ara, 2022). Findings across these studies show trade-offs: some post-hoc methods offer rich insight but introduce latency, while intrinsic models deliver faster explanations but may sacrifice predictive power on complex tasks. Hybrid quantitative evaluations indicate that combining interpretable core models with targeted post-hoc rationales can yield strong accuracy while maintaining governance-friendly interpretability (Lopes et al., 2022; Nahid, 2022). This body of quantitative work defines the evaluation landscape for XAI in cloud BI, treating explainability not as a narrative quality but as a set of empirically testable variables that mediate compliance and security outcomes (Hossain & Milon, 2022).

Figure 2: Explainable AI for Cloud BI



The international significance of XAI for cloud-based BI emerges from the globalization of data-driven operations and the heterogeneous accountability expectations that accompany cross-border decision flows. Multinational enterprises routinely centralize analytics in cloud platforms while serving stakeholders in multiple jurisdictions (Mominul et al., 2022; Sheu & Pardeshi, 2022). Organizational studies on global BI show that shared dashboards align supply chains, marketing, finance, and risk management across regions, yet the same AI-generated recommendation may be subject to different legal standards and cultural norms of justification (Mortuza & Rauf, 2022). Comparative governance research highlights that explanation audiences vary: regulators seek evidence of lawful processing and non-discrimination, auditors seek traceability and control adherence, managers seek actionable drivers of KPIs, and affected individuals seek understandable reasons for outcomes (Rakibul & Samia, 2022). Studies on explanation design emphasize layered approaches that deliver role-appropriate reasoning without exposing unnecessary sensitive details. In international BI deployments, layered XAI aligns with access-control structures already used in cloud platforms: a compliance team can view detailed causal attributions and fairness diagnostics, security analysts can view anomaly-reasoning traces, and executives can view concise driver narratives linked to business objectives (Owens et al., 2022; Saikat, 2022). Empirical studies of regulated industries across regions show that explainability reduces friction in adoption because it enables a common language for oversight among headquarters, subsidiaries, and local regulators. Multiple cross-sector case analyses also show that explanation artifacts help

standardize model risk management across borders, giving firms a portable governance mechanism even when local rules differ. At the same time, international privacy studies show that explanation disclosure must be calibrated so that transparency does not violate confidentiality or data-minimization obligations (Hagras, 2018; Kanti & Shaikat, 2022). This makes explainability a strategic governance capability for global cloud BI: it supports consistent justification of decisions while allowing localized control over what is revealed. The result is an internationally relevant definition of XAI for cloud BI that centers on accountable interoperability across jurisdictions rather than only technical interpretability (Arfan et al., 2023).

Operationally, XAI models for cloud-based BI function within end-to-end analytical pipelines that integrate data engineering, model development, explanation generation, governance controls, and decision interfaces (Ara & Beatrice Onyinyechi, 2023; Linardatos et al., 2020). Cloud BI workflows generally begin with ingestion and harmonization of data into centralized repositories, followed by feature engineering, model training in distributed compute environments, deployment through scalable serving layers, and consumption through dashboards or automated triggers. Studies of machine learning operations in enterprises show that governance failures often arise not from model form alone but from weak lifecycle integration, where versioning, documentation, and monitoring are fragmented. XAI addresses this operational risk by adding an explanation layer that is co-managed with model artifacts. Research on explanation services indicates that explanations can be generated at training time to produce global model narratives, and at inference time to provide local rationales for each decision, then stored as metadata attached to BI outputs (Mushfequr & Ashraf, 2023). Lifecycle studies show that explanation records strengthen accountability when they are logged, hashed for integrity, and linked to the exact data snapshot and model version used in the decision. Monitoring studies also show that explanation trends can be tracked as quantitative signals of drift, bias emergence, or security anomalies, complementing accuracy and error-rate monitoring (Shahrin & Samia, 2023; Nazar et al., 2021). In cloud implementations, performance studies reveal that explanation computation must be optimized to meet BI latency requirements for interactive dashboards, leading to batching strategies, approximation techniques, and model-specific explanation accelerators. Behavioral analytics studies demonstrate that explanation-augmented dashboards improve decision quality when they present drivers in a structured, cognitively manageable form, such as ranked feature impacts, scenario-based counterfactual levers, or constrained rule summaries. Risk management studies further show that explanations enable systematic override protocols, where humans can justify deviations from model recommendations with documented rationale, preserving compliance records (Nizam & Zafar, 2022). Across operational research, the definitional role of XAI in cloud BI is therefore that of a governance-aware analytical service embedded in the same scalable infrastructure that produces BI insights, ensuring that every automated recommendation arrives with a secure and compliant reasoning trail.

The primary objective of this quantitative study is to develop and empirically evaluate explainable AI (XAI) models embedded within cloud-based business intelligence (BI) environments to determine how explanation quality influences regulatory compliance assurance and secure decision-making performance. Specifically, the study aims to (a) design or select representative intrinsic and post-hoc XAI techniques suitable for common cloud BI data types (structured transactional data, streaming operational logs, and integrated customer or supply-chain datasets), (b) operationalize explainability through measurable indicators such as local fidelity, global fidelity, stability across retraining cycles, sparsity of explanatory features, and user-agreement scores, (c) quantify the relationship between these explainability indicators and compliance outcomes, including audit traceability, fairness consistency across protected and non-protected attributes, documented decision reproducibility, and alignment with organizational policy thresholds, and (d) measure the extent to which explainability contributes to secure decision-making by detecting anomalous reasoning, reducing susceptibility to adversarial manipulation, minimizing sensitive information leakage in explanations, and maintaining decision integrity under varying cloud workloads and tenant conditions. To meet these aims, the study will construct controlled cloud BI scenarios that mirror real enterprise workflows – forecasting, risk scoring, anomaly detection, and resource optimization – then compare predictive models with and without integrated explainability layers. The study further seeks to test whether explanation-augmented

dashboards improve managerial decision accuracy, reduce over-reliance on automation, and increase justified override rates when models are incorrect, using statistically valid experiments with role-based user groups (executives, analysts, and compliance officers). A parallel objective is to evaluate system-level feasibility by measuring explanation latency, scalability under concurrent queries, and cost overhead within a multi-tenant cloud setting, ensuring that interpretability gains do not undermine BI responsiveness. Through these linked objectives, the study intends to provide a rigorous, measurement-driven account of how XAI can function as a governance-ready analytical service in cloud BI, delivering transparent, verifiable, and security-resilient insights that remain suitable for high-stakes organizational decision contexts.

LITERATURE REVIEW

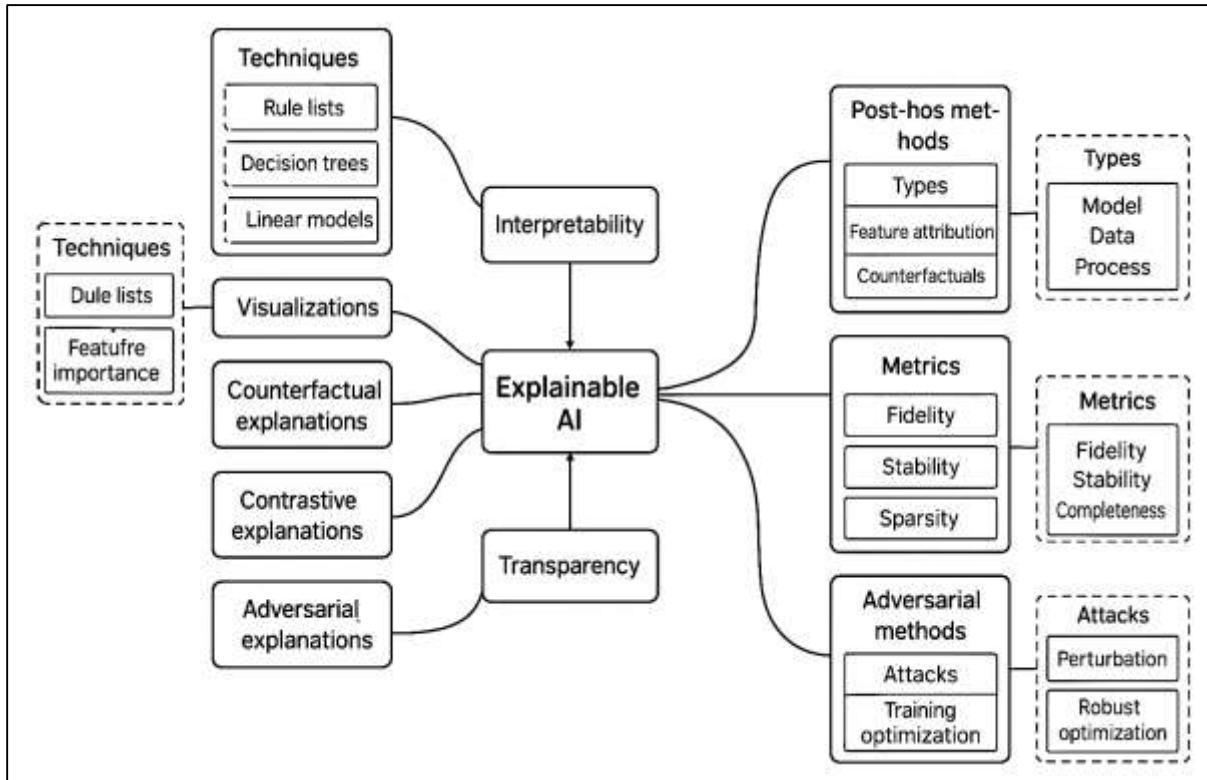
This literature review examines explainable artificial intelligence (XAI) models as operational components of cloud-based business intelligence (BI) systems, with an emphasis on measurable outcomes for compliance assurance and secure decision-making. Because cloud BI platforms increasingly rely on machine learning for forecasting, risk scoring, anomaly detection, and automated recommendations, the quality of explanations has become a central determinant of governance viability. Existing studies show that explainability strengthens human oversight, auditability, and decision trust, yet empirical findings are dispersed across separate domains such as interpretability metrics, cloud BI architectures, compliance theory, and machine-learning security (Mehdiyev et al., 2021). A unified and quantitative synthesis is therefore required to clarify which XAI properties matter most for organizations that must justify AI-driven BI decisions under regulatory and security constraints. To build this foundation, Sections 1–8 proceed in a structured sequence. Section 1 establishes the conceptual foundations of XAI in organizational analytics. Section 2 classifies XAI model families relevant to cloud BI, distinguishing intrinsic, post-hoc, and hybrid approaches. Section 3 reviews how explainability is evaluated quantitatively, defining measurable indicators that can serve as independent variables. Section 4 explains cloud BI ecosystems and identifies where XAI services integrate into data-to-dashboard pipelines (Holzinger et al., 2021). Section 5 synthesizes the compliance dimension by translating regulatory requirements into quantifiable assurance indicators. Section 6 reviews security risks in cloud BI and clarifies how XAI can function as both a security validator and a potential leakage surface. Section 7 analyzes empirical research patterns, focusing on quantitative methods and measurement gaps in current studies. Section 8 provides an integrative synthesis that links explainability quality to compliance and security outcomes through a testable conceptual-statistical framework. Collectively, these sections justify the variables, measurement strategy, and hypotheses of the present quantitative investigation (Kumar & Mehta, 2022).

Explainable AI (XAI)?

Explainable artificial intelligence (XAI) is commonly defined in enterprise decision systems as the capability of an AI model to provide human-understandable reasons for its outputs in ways that support accountability, validation, and responsible action. Within organizational environments, AI predictions are rarely treated as isolated technical results; they are translated into operational choices, resource allocations, risk approvals, customer targeting actions, and performance evaluations (Longo et al., 2020). Because of this embeddedness, explainability functions as decision justification capability rather than mere model description. Many studies in interpretability and enterprise AI governance converge on the view that explanations are required when a system influences high-stakes outcomes, when multiple stakeholders rely on shared model outputs, or when decisions face internal audit and external regulatory scrutiny. In practical terms, XAI means that a system does not only output a prediction or score, but also provides a rationale that clarifies which inputs mattered, how they mattered, and why the outcome emerged in the observed context (Cali et al., 2021). Research across machine learning, human-computer interaction, information systems, and organizational analytics shows that explanations enable people to assess whether a model is reasoning in credible ways, whether the model's logic aligns with policy constraints, and whether the decision should be accepted, challenged, or overridden. Further studies describe explainability as necessary for aligning statistical learning with managerial sense-making, since managers require interpretable reasoning to link AI outputs to business narratives and accountable actions. Across empirical enterprise deployments, XAI is described as a core trust-enabling layer because it reduces the cognitive distance between complex

models and non-technical decision owners (Borrego-Díaz & Galán Páez, 2022). Investigations in finance, healthcare, HR analytics, and supply-chain optimization similarly show that decision justification is an operational requirement, not an optional enhancement, because outcomes often require formal defense or procedural documentation. In this view, explainable AI is not defined by one technique or model family but by its functional role in enterprise decision systems: transforming black-box prediction into a transparent, verifiable basis for action (Borrego-Díaz & Galán-Páez, 2022). This definitional framing positions XAI as a governance-aligned capability that supports model adoption, safe decision integration, and traceable organizational responsibility in environments where AI outputs directly shape policy-relevant business decisions.

Figure 3: Enterprise XAI for Cloud BI



The literature distinguishes interpretability, explainability, and transparency as related yet non-identical properties that together structure how organizations evaluate AI decision systems. Interpretability is typically described as the degree to which a human can understand a model’s internal logic or mentally simulate how inputs map to outputs. Models such as sparse linear predictors, rule lists, or shallow decision trees tend to be considered interpretable because their structure is readable without auxiliary tools (Páez, 2019). Explainability, by contrast, is described as the capacity to generate meaningful accounts of why a particular output occurred, even when the underlying model is complex. Explanations can be generated through post-hoc techniques such as feature attribution or counterfactual reasoning, or through intrinsically explainable architectures. Transparency is often treated as a broader governance idea referring to openness of the model, the data pipeline, and the decision process, including traceability of updates and documentation of assumptions. Studies in responsible AI emphasize that transparency without explanation is incomplete because stakeholders require reasoning, not only access to the model form. Other work indicates that interpretability of a model does not automatically guarantee explainability of its decisions, as even simple models can produce outcomes that seem counterintuitive without contextual rationale. In enterprise contexts, these distinctions matter because accountability systems do not evaluate models only on accuracy but on whether reasoning can be inspected, communicated, and defended (Holzinger et al., 2020). A widely supported position across interpretability research is that explanation artifacts are measurable governance objects that accompany model outputs in regulated or high-stakes decision flows.

Explanation artifacts include ranked driver features, rule pathways, scenario contrasts, monotonic effect summaries, or minimal sets of conditions sufficient to produce the observed output. Organizational case studies show that such artifacts become part of audit evidence, model risk documentation, fairness review packets, and incident investigation logs. Research on AI lifecycle governance further indicates that explanation artifacts help connect model lineage to decision lineage, enabling investigators to track which model version, data snapshot, and reasoning basis supported a specific business action (Rohlfing et al., 2020). In human oversight tasks, empirical studies show that explanations are treated as decision records that either confirm proper use of data and policy-aligned reasoning or reveal spurious and risky logic that requires remediation. The governance value of explanation artifacts is strengthened when they are stable across minor data shifts, reproducible across retraining cycles, and accessible to role-specific stakeholders who need different levels of reasoning detail. Taken together, interpretability, explainability, and transparency are framed in the literature as a layered accountability structure, with explanation artifacts serving as the measurable bridge that operationalizes responsible AI governance inside enterprise decision environments.

In business intelligence (BI), the role of XAI becomes essential because BI has moved from static descriptive reporting toward AI-driven analytical systems that continuously generate predictive and prescriptive insights. Traditional BI dashboards relied on predefined KPIs, periodic extraction routines, and human-authored logic for summarizing organizational performance (Minh et al., 2022). In that setting, the reasoning behind a KPI was traceable through business rules, metric definitions, and analyst documentation. AI-enabled BI alters this structure by embedding machine learning models that learn complex patterns from enterprise data and generate dynamic KPIs such as risk scores, anomaly likelihoods, demand forecasts, churn probabilities, or next-best-action recommendations. Studies on AI in analytics systems show that these model-produced indicators increase analytical power and allow BI platforms to support faster and more granular decisions, yet the learned logic that drives these KPIs is often non-obvious to non-technical users. BI outputs are consumed across departments by stakeholders with different decision responsibilities who may not share technical training. Finance teams interpret AI-powered projections for budgeting and compliance; operations teams act on anomaly alerts for process control; marketing teams respond to customer propensity scores; and senior executives rely on integrated dashboards for strategic direction (Donadello & Dragoni, 2020). Research on cross-functional BI adoption shows that when model logic is opaque, stakeholders either mistrust outputs, misuse them, or accept them without proper oversight. Governance therefore depends on visible reasoning. Explanations embedded in BI dashboards allow users to see which data drivers shaped a KPI, whether those drivers align with domain expectations, and whether the decision context includes risks or confounds that require human judgment. Empirical studies in decision support also show that BI value is undermined when managers cannot justify AI-based actions to committees, auditors, or affected individuals, especially in regulated industries where automated decisions require defensible logic. Explanations serve as shared accountability language across departments, enabling consistent interpretation of AI outputs in meetings, approval workflows, and escalation processes (Nizam & Zafar, 2022). Research on trust calibration indicates that explanations do not function only to persuade users but to help them identify when reliance is appropriate and when override is required. In cloud BI, where models update frequently and outputs circulate globally, reasoning visibility helps preserve continuity of understanding over time, preventing silent logic shifts from distorting decision routines. The literature thus positions XAI as a structural requirement for BI systems that integrate AI: without explanations, AI-driven KPIs become difficult to validate, difficult to govern, and difficult to defend as legitimate bases for organizational action.

A key stream in the literature treats explainability not as a rhetorical feature but as a quantitative construct that can be measured, compared, and statistically linked to organizational outcomes. This perspective emerges because enterprise AI adoption requires evidence that explanations reliably reflect model reasoning and meaningfully improve decision quality under real operational tasks. Explainability is therefore conceptualized as a variable with observable indicators rather than a narrative about transparency (O. Chergykalo & Klyushin, 2022). Studies that formalize explainability metrics propose measurable dimensions including fidelity, stability, sparsity, completeness, and human agreement. Fidelity represents the degree to which an explanation matches the true reasoning

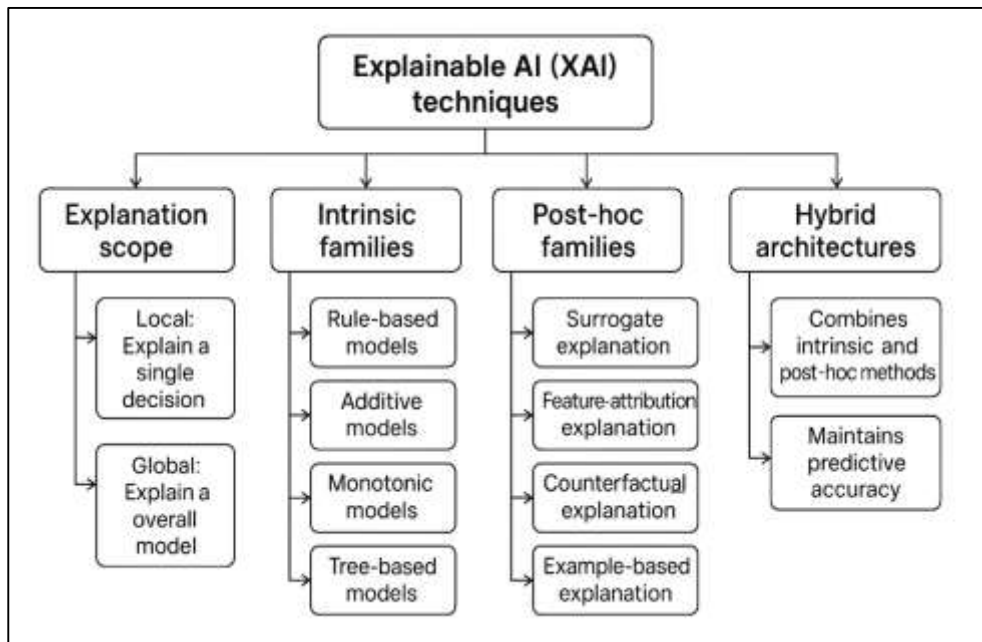
pattern of the underlying model, either locally for individual decisions or globally across the data distribution. Stability captures whether explanations remain consistent under small perturbations of input data or across retraining cycles, a property especially relevant in cloud BI where models update regularly. Sparsity measures explanation simplicity by counting how many features or rules are required to justify a decision, reflecting cognitive load on BI users. Completeness assesses how much of the model's predictive signal is actually covered by the explanation rather than left implicit. Human agreement quantifies whether stakeholders interpret explanations consistently and whether explanations improve task accuracy, speed, and error detection. Quantitative evaluation studies show that these indicators vary systematically across intrinsic explainable models, post-hoc methods, and hybrid designs, enabling comparative testing under identical BI conditions. Empirical work in decision-support experimentation also links measurable explainability properties to user outcomes such as calibrated trust, appropriate reliance, justified overrides, and reduction in decision errors (Owens et al., 2022). In governance-focused studies, explanation metrics are connected to auditability outcomes by tracking whether explanations can be logged, reproduced, and inspected during compliance review. Security-oriented research introduces additional quantitative measures for explainability, such as explanation drift rates that signal data poisoning or adversarial influence, and leakage risk scores that capture sensitive information exposure through explanations. Across these streams, the field converges on the position that explanation performance must be tested under realistic BI tasks—forecasting, anomaly detection, risk scoring, and optimization—because explanation quality depends on data type, workload intensity, and decision context (Kim et al., 2021). Treating XAI as a quantitative construct allows organizations to test whether improved explainability yields stronger compliance assurance and more secure decision performance instead of relying on subjective impressions. This measurement orientation directly supports cloud BI deployment decisions, model selection, and governance policy, since enterprises can benchmark explainability quality as they benchmark accuracy, latency, and cost.

Taxonomy of XAI Models Relevant to Cloud BI

Intrinsic explainable models are designed so that their internal logic is understandable without relying on external explanation tools, making them particularly attractive for cloud-based business intelligence (BI) where decisions must be auditable and quickly interpretable by diverse stakeholders (Gerlach et al., 2022). The literature on interpretable learning emphasizes that intrinsic explainability is achieved when the prediction process itself is aligned with human reasoning structures such as rules, additive contributions, monotonic relationships, or decision paths. Rule-based learners represent one of the earliest and most widely studied intrinsic families. These models express predictions through “if-then” conditions that can be directly translated into business logic, allowing cloud BI users to see how specific thresholds or categorical conditions drive outcomes. Studies in enterprise analytics show that rule-based models support compliance reviews because the reasoning is explicit and can be logged as a decision artifact. Interpretable additive models extend transparency by decomposing predictions into separable feature effects, enabling users to understand not only which drivers matter but also the marginal direction of influence across the data space (Höhn & Faradouris, 2021). Because BI dashboards often summarize performance across segments, additive interpretability aligns well with managerial needs for comparative reasoning. Another intrinsic family highlighted in the literature includes monotonic scoring systems, which enforce domain-consistent directionality across critical variables, such as ensuring that higher risk indicators never reduce a risk score. Research on model governance indicates that monotonicity improves trust in regulated BI decisions because it prevents counterintuitive logic that might trigger audit concerns. Tree-based transparent predictors, including shallow decision trees and carefully constrained ensembles, are also prominent in cloud BI deployments due to their natural similarity to human decision structures. Their predicted outcomes can be traced along paths that reflect conditional reasoning, making them readable for non-technical stakeholders. Across these intrinsic families, multiple studies converge on the view that their main advantage in cloud BI is direct traceability: explanations are not approximations but inherent outputs of the model form (Guleria et al., 2022). However, the literature also documents trade-offs, noting that as data complexity grows—particularly in high-dimensional customer, sensor, or transactional environments—fully intrinsic models may underperform more complex black-box alternatives. Even so, intrinsic explainable models remain foundational in cloud BI because they provide stable, policy-

aligned reasoning that supports compliance logging, secure oversight, and cross-functional decision ownership without requiring specialized interpretability layers.

Figure 4: XAI Techniques for Cloud BI



Post-hoc explanation methods occupy a central position in cloud BI because organizations frequently use high-performing black-box models while still needing interpretability for governance, compliance, and managerial decision confidence. The literature defines post-hoc explainability as a set of techniques applied after model training and prediction, aiming to approximate or communicate reasoning without altering the underlying black-box structure (Mohanrajan & Loganathan, 2022). A major category is surrogate explanation engines, where a simpler interpretable model is trained to mimic the behavior of a complex predictor in a local neighborhood or across broader regions of data. Studies indicate that surrogate models are particularly useful in BI dashboards because they translate opaque outputs into human-readable logic, although their usefulness depends strongly on how faithfully they represent the original model. Another dominant category is feature-attribution explainers, which rank or weight input variables by their contribution to a specific decision. This family is widely discussed as suitable for BI contexts because it aligns naturally with dashboard formats that visualize drivers of KPIs, such as top contributing risk factors or demand influencers (Hussain et al., 2022). Counterfactual explanation generators represent another post-hoc approach emphasized across interpretability research. These methods identify minimal, realistic changes to inputs that would alter a decision outcome, framing explanations in actionable “what-if” terms. In cloud BI, counterfactuals map well to scenario simulation features used by managers, because they indicate how a recommendation might change under alternate operational conditions. Example-based and prototype explanations are also prominent in the literature, especially for complex enterprise data. Instead of explaining through abstract weights, these approaches justify a decision by showing similar historical cases, representative prototypes, or contrasting examples, giving decision makers intuitive analogies. Across post-hoc methods, many studies report that their strength lies in flexibility: organizations can retain advanced predictive models for accuracy while layering explanations suitable for auditors and non-technical users. At the same time, the literature is consistent in warning that post-hoc explanations can be fragile. Explanation outputs may vary under small perturbations, may highlight correlated but non-causal drivers, and can be manipulated by adversarial inputs (Khrais, 2020). In BI, these weaknesses matter because dashboards are trusted sources for operational action. Therefore, studies in enterprise governance recommend using post-hoc explanations with quantitative monitoring of fidelity, stability, and consistency. Overall, post-hoc methods are presented in the literature as indispensable for cloud BI

environments that prioritize predictive performance but still require interpretable, reviewable, and communicable decision rationales.

Hybrid XAI architectures are described in the literature as integrated systems that combine black-box predictive power with structured explainability layers, often yielding better enterprise outcomes than relying solely on intrinsic or post-hoc methods. In cloud BI, hybrid designs typically involve a high-capacity predictor such as an ensemble or deep learning model paired with an explanation wrapper that generates reasoning artifacts at inference time (Flammini et al., 2022). Studies in interpretability and enterprise deployment repeatedly show why this pairing is attractive: cloud BI tasks often require maximum accuracy to detect subtle risks or optimize supply trends, yet they also face accountability requirements that intrinsic models alone may not satisfy at scale. Hybrid approaches aim to preserve the predictive strength of black-box models while producing explanations that are readable, auditable, and aligned with business policy. The literature highlights several situations where hybrid architectures outperform single-family models. First, when BI data are high-dimensional and heterogeneous, intrinsic transparency may degrade accuracy, while post-hoc wrappers around black-box models can maintain performance and still provide reasoning signals. Second, in regulated decision systems, black-box predictors paired with stable and fidelity-checked explanation services have been shown to satisfy audit requirements more effectively than black-box models without structured explanation governance. Third, hybrid systems are favored when cloud BI requires continuous retraining, because explanation services can be tuned independently and monitored for drift without redesigning the predictive core. Another theme in hybrid literature is role-specific explainability: different stakeholders need different depths and styles of rationale (Zhang et al., 2022). Hybrid architectures allow BI platforms to deliver layered explanations, such as high-level driver summaries for executives and detailed attribution breakdowns for compliance analysts. Studies in organizational adoption also report that hybrid XAI improves trust calibration. Users are more willing to rely on BI recommendations when they can see consistent rationales, and they are better at rejecting outputs when explanation patterns appear suspicious or inconsistent with domain expectations. The literature further notes that hybrid systems support secure decision-making because explanation drift can act as a monitoring signal for data poisoning, model instability, or anomalous reasoning. At the same time, studies caution that hybrid implementations require careful governance: low-fidelity explanation wrappers may create misleading certainty, and excessive detail may leak sensitive business logic in cloud environments (Le et al., 2022). Even with these risks, the evidence across enterprise case studies, interpretability benchmarks, and BI adoption research consistently positions hybrid XAI as a practical governance compromise—offering advanced predictive BI capability while embedding explainability as a controllable, measurable decision artifact.

The literature separates explanation scope into local and global forms because BI environments require both decision-level justification and system-level accountability. Local explanations are defined as rationales tied to individual predictions—such as a single customer churn risk score, one anomalous transaction alert, or one demand forecast for a specific region (Abdelwahab et al., 2022). In cloud BI, local reasoning is essential when decisions affect discrete entities, because managers and auditors must be able to trace why a particular outcome occurred. Studies in interpretability show that local explanations enhance decision reliability by allowing users to inspect driver patterns at the case level, detect data quality issues, and confirm whether a decision aligns with domain logic. This scope aligns with operational BI tasks where actions are applied one instance at a time, including fraud investigation, credit review, equipment maintenance escalation, and targeted marketing intervention. Global explanations, by contrast, summarize how the model behaves overall. They provide an aggregated view of what features systematically influence outcomes across the enterprise dataset, how relationships differ by segment, and where the model may embed hidden bias. The literature suggests that global scope is vital for strategic BI governance because regulated organizations must validate models not only for single decisions but for population-level fairness, long-run consistency, and policy-aligned reasoning. Studies of AI audit practice emphasize that global explanations support periodic model reviews, risk committee evaluations, and compliance certification cycles, especially in cloud BI environments where models update continuously (Kuppa & Le-Khac, 2021). Another recurring theme

is that local and global explanations complement one another: global summaries provide policy-level assurance, while local rationales provide situational accountability. Research on BI dashboard design indicates that effective systems deliver both scopes through layered views, enabling users to drill down from corporate KPIs into case-level reasoning. The literature also notes that different XAI families fit different scopes. Intrinsic interpretable models naturally provide global logic through readable structures, whereas post-hoc and hybrid systems are often optimized for local case explanations but can be aggregated into global patterns with careful validation. Across multiple studies, explanation scope is treated as a governance requirement because BI stakeholders need to answer two distinct questions: “Why did this specific decision occur?” and “Is the model behaving properly overall?” In cloud-based BI, where accountability spans distributed teams and jurisdictions, managing both local and global explanation scope is presented as essential to maintaining trust, compliance readiness, and secure decision legitimacy (Saraswat et al., 2022).

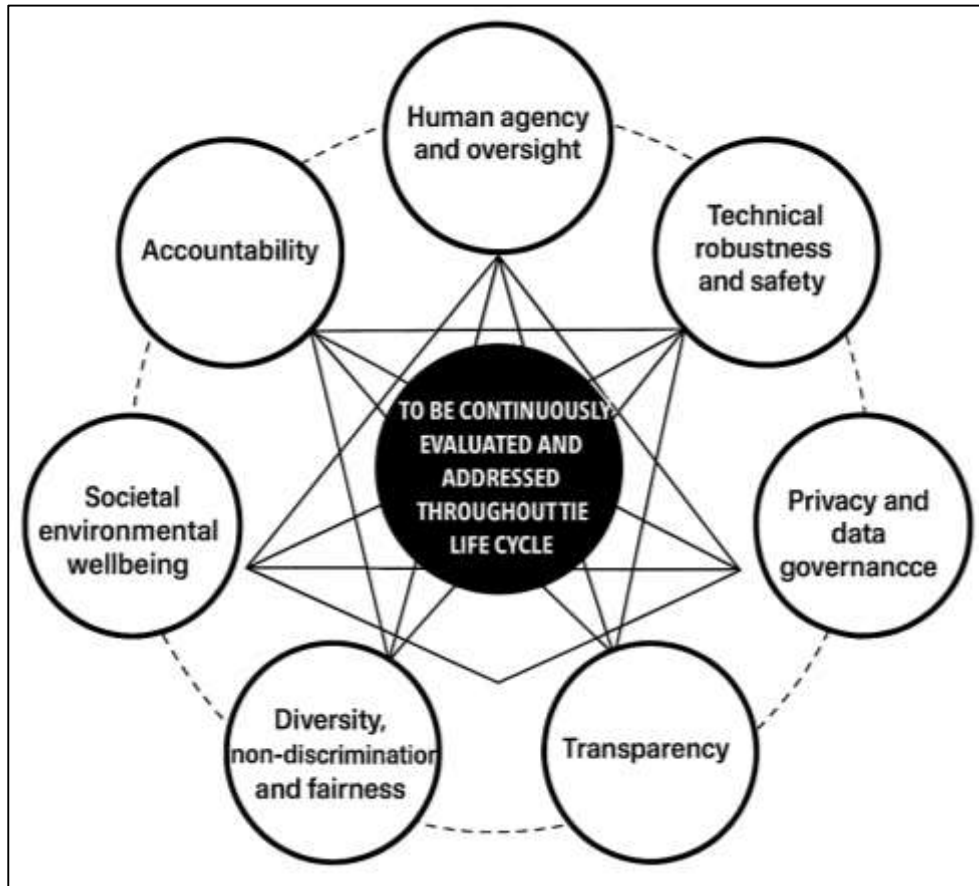
Quantitative Metrics for Explainability Quality (EQ)

Fidelity is consistently treated in explainability research as a foundational quantitative measure because it captures whether an explanation truthfully reflects the reasoning patterns of the underlying predictive model. In the literature, high-fidelity explanations are those that reproduce the model’s decision behavior closely enough that stakeholders can rely on the explanation as a valid representation of why an output occurred (Zhou et al., 2021). Empirical studies distinguish local fidelity from global fidelity. Local fidelity focuses on agreement around a specific prediction instance, meaning that if inputs near the focal case vary slightly, the explanation still tracks the black-box model’s decisions accurately in that local neighborhood. This idea appears repeatedly in studies of surrogate-based and feature-attribution explainers, where the explanation is intended to approximate complex models in a narrow region of the input space. Global fidelity extends this requirement across the broader data space, asking whether the explanatory logic continues to match the model’s decision patterns over all major segments and distributions relevant to enterprise BI. Across multiple benchmark studies, low local fidelity has been linked to misleading rationales that may look intuitive but fail to represent true causal drivers, while low global fidelity leads to explanations that work for isolated cases but collapse when analysts test them in different segments. In cloud BI environments, fidelity is emphasized because dashboards are decision surfaces; if explanations deviate from real model logic, BI users can validate the wrong signals, audit teams can approve noncompliant reasoning, and security analysts can miss hidden risk patterns (Machlev et al., 2021). The literature therefore treats fidelity not as a theoretical preference but as a measurable accuracy requirement for explanation reliability. Fidelity is also discussed as central to governance because enterprises increasingly store explanations as decision artifacts, and these artifacts perform regulatory and managerial roles. When fidelity is weak, explanation artifacts fail to function as evidence. Another important thread in fidelity research highlights that fidelity must be assessed contextually: explanations that are faithful for linear decision regions may not remain faithful in highly nonlinear segments. This creates a BI-specific concern, since enterprise datasets typically contain heterogeneous clusters such as customer tiers, geographic markets, or process states. Studies comparing explanation families show systematic fidelity variation: intrinsic interpretable models often exhibit inherently high global fidelity because their logic is the predictive logic, while post-hoc explainers must be validated repeatedly to confirm agreement (Coroama & Groza, 2022). Hybrid architectures rely on fidelity checks to ensure that the explanatory wrapper does not become a detached narrative layer. Across these studies, fidelity emerges as the first quantitative gatekeeper for explainability quality, defining whether explanations can legitimately carry managerial trust, compliance justification, and secure decision oversight in cloud BI operations.

Stability, often paired with robustness, is described in the literature as the consistency of explanations under small input changes and across repeated model updates. Quantitative work on interpretability repeatedly shows that even highly faithful explainers can be operationally weak if they are unstable. Stability under small input noise evaluates whether slight perturbations in data produce radically different explanations for essentially the same decision (Banerjee & Barnwal, 2022). If explanations fluctuate sharply with minimal input variation, they become unreliable for BI users because they cannot distinguish genuine driver changes from explanation volatility. This issue is amplified in cloud BI because data inputs are frequently noisy – streaming operational logs, transaction variances, sensor

drift, or seasonal customer behavior introduce natural fluctuation. Studies focused on real-world ML deployment show that unstable explanations reduce user confidence, impair audit defensibility, and increase error rates in human oversight tasks.

Figure 5: XAI Metrics for Reliable Explanations



Stability across retraining and drift adds another enterprise-critical layer. Cloud BI systems retrain models continuously to reflect new data streams, and interpretability research shows that explanation patterns can drift even when predictive accuracy remains acceptable. In governance literature, explanation drift is treated as a measurable warning signal: it may indicate concept drift, altered feature importance, data pipeline issues, or emergent bias. Empirical investigations demonstrate that explanation drift can precede visible performance degradation, making stability assessments useful for early detection of governance and security problems (Sovrano et al., 2022). Robustness is also treated as a security-aligned property because adversarial perturbations can be detected through abnormal explanation shifts. The literature on explainable security indicates that when malicious inputs attempt to manipulate model outputs, explanation vectors often reveal inconsistent or suspicious attribution patterns. Therefore, stability is not only a usability issue but also a reliability and risk-monitoring metric. Comparative studies across explanation families show that some post-hoc explainers generate highly variable local rationales due to sampling or approximation noise, while intrinsic models tend to remain more stable because their structure constrains reasoning pathways. Hybrid systems sit between these extremes; their stability depends on both the predictive core and the wrapper technique (Siddiqui & Doyle, 2022). Across the research landscape, stability is framed as a necessary complement to fidelity because stable explanations become repeatable decision evidence. In cloud BI, where explanations support ongoing audits, policy validation, and secure operational action, stability defines whether explanation artifacts remain dependable over time and across the dynamic retraining cycles typical of distributed BI architectures.

Sparsity and complexity metrics dominate human-centered explainability research because they quantify how easy an explanation is to understand and use under real decision conditions. Sparsity refers to how compact an explanation is, typically measured by the number of features, rules, or

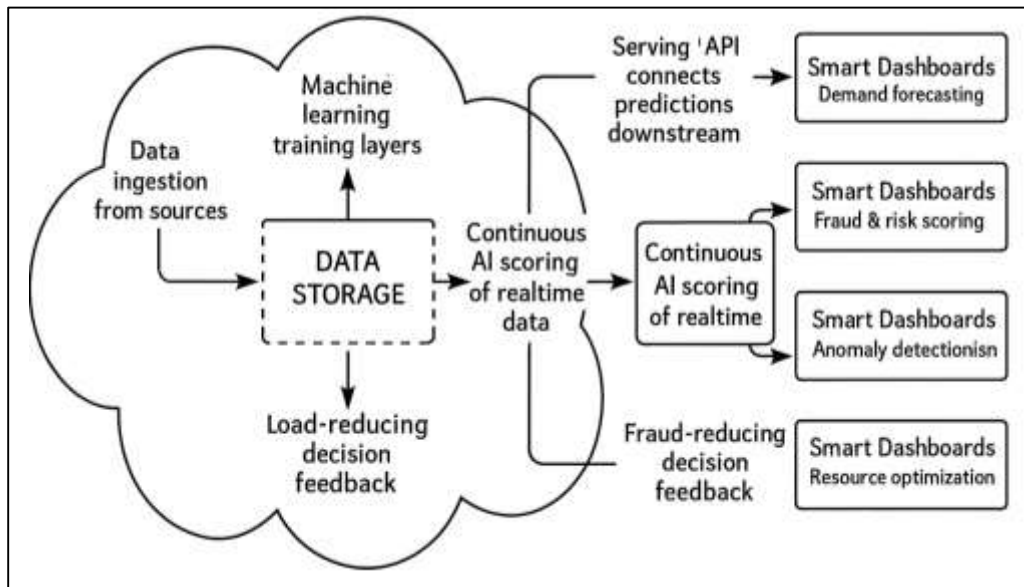
decision conditions required to justify an output. Complexity expands this idea to include structural difficulty, such as nested rule depth, interaction density, or nonlinear effect layering. Studies in interpretability and BI usability consistently show that explanation effectiveness decreases when reasoning becomes too long or structurally complicated for stakeholders to process (Fouladgar et al., 2022). This is a crucial issue in cloud BI environments because explanations are delivered through dashboards, and dashboard space, attention, and decision time are limited. When an explanation includes too many variables or too much conditional logic, BI users may default to superficial trust or rejection without meaningful evaluation. The literature connects excessive complexity to elevated cognitive load, misinterpretation of driver influence, and reduced ability to detect model errors. Sparse explanations, in contrast, allow decision owners to identify the dominant contributors quickly and align them with domain knowledge. However, multiple studies show that sparsity is not a simple “less is always better” rule. Overly sparse explanations can omit important interacting factors and create a false sense of simplicity, especially in BI tasks that involve complex nonlinear dependencies (Li et al., 2020). Therefore, the literature frames sparsity as an optimization balance between intelligibility and completeness. Research comparing intrinsic and post-hoc explainability finds predictable patterns: intrinsic models that are constrained for interpretability often yield naturally sparse rationales, while post-hoc feature attributions can become dense unless explicitly regularized. Hybrid BI architectures attempt to manage this trade-off by offering layered explanations, but studies caution that layering must still preserve a cognitively manageable first view. Another theme across enterprise studies is role dependence. Compliance officers may tolerate more complex logic to verify policy constraints, while executives require sparse high-level drivers to guide strategic action. Security analysts may also need moderate complexity to identify manipulation cues (Tjoa & Guan, 2022). Thus, sparsity and complexity metrics are treated as essential quantitative indicators because they capture whether an explanation fits the cognitive practices of its intended BI audience. In cloud BI governance, explanations are only useful if they can be acted upon appropriately, so evaluative studies link sparsity and complexity to measurable user outcomes such as decision accuracy, time-to-decision, and correct override rates. Overall, the literature positions sparsity and complexity not as aesthetic characteristics but as measurable determinants of explainability quality that shape how well explanations support real organizational BI decisions (Barnard et al., 2022).

Cloud-Based Business Intelligence (BI) Ecosystems

Cloud-based business intelligence (BI) is defined in the literature as a distributed analytical ecosystem that collects, stores, processes, and visualizes enterprise data through scalable cloud infrastructure to support decision-making across operational and strategic levels. Unlike traditional on-premises BI, cloud BI is structured around elastic compute and storage services that allow organizations to scale analytical capacity dynamically with workload demands (Ansari & Alam, 2022). Empirical studies on cloud analytics architectures describe the typical structural pipeline as beginning with data ingestion, where heterogeneous sources – ERP, CRM, IoT streams, web logs, supply-chain feeds, and external market datasets – are captured into cloud-native staging layers. The ingestion stage is followed by cloud storage organized as data lakes, lakehouses, or warehouses, each of which supports a different balance between raw storage flexibility and structured query performance. A major body of BI research highlights the lakehouse direction as a structural convergence that enables both high-throughput storage and relational analytics without duplicating data across silos. Above storage, cloud BI ecosystems include machine learning training layers that provision distributed environments for feature engineering, model building, and evaluation. These ML layers are increasingly embedded directly into BI platforms rather than treated as separate data science stacks, and studies of enterprise ML operations show that this integration reduces deployment time and improves model refresh frequency. After training, serving APIs function as inference gateways that expose predictive and prescriptive outputs to downstream BI tools, allowing systems to score transactions, events, or customer records in real time (Ghita et al., 2020). Finally, dashboards and decision automation components represent the consumption layer. The dashboard layer translates data and model outputs into interactive KPI views, scenario tools, and alerts, while decision automation converts insights into triggered actions such as workflow routing, fraud holds, demand replenishment requests, or resource scheduling. Across many studies of cloud BI adoption, this full-stack structure is described as a decision

pipeline rather than a reporting tool, meaning that data moves through computational stages designed to generate direct managerial action. Because this structure is multi-tenant and globally distributed in many organizations, cloud BI also includes governance services for access control, lineage tracking, versioning, and logging, which allow enterprises to maintain accountability across users and jurisdictions. In this way, cloud BI is structurally defined as a layered ecosystem integrating ingestion, storage, modeling, serving, and decision interfaces into a continuously operating analytical environment (Srivastava et al., 2022).

Figure 6: Cloud BI Decision Pipeline Framework



AI integration into BI dashboards is widely described as the transformation of dashboards from descriptive reporting surfaces into predictive and prescriptive decision systems. In descriptive BI, dashboards summarize what has happened through static KPIs such as revenue, cycle time, defect rate, or churn. The literature on AI-enabled BI shows that predictive BI extends these dashboards by embedding machine learning forecasts—future demand projections, risk probabilities, customer lifetime value estimates, and equipment failure likelihoods—allowing managers to interpret not only current performance but also expected trajectories (Baidya & Hallur, 2022). Prescriptive BI goes a step further by recommending optimized actions, such as which inventory levels to order, which customers to target, which claims to investigate, or which routes to reconfigure, using optimization or reinforcement learning logic. Studies of transformation in BI tools highlight that once AI is integrated, dashboards become active decision mediators rather than passive displays. Another central distinction is between continuous scoring and periodic reporting. Periodic reporting refers to traditional refresh cycles, typically daily, weekly, or monthly, where KPIs are recalculated in batch mode. Continuous scoring leverages cloud-based serving APIs to update AI-driven indicators in real time as new data arrives, enabling instant fraud alerts, dynamic pricing adjustments, streaming anomaly detection, or live demand sensing. Enterprise analytics research shows that continuous scoring improves responsiveness and reduces lag in critical operational functions, especially in finance, logistics, cybersecurity, and customer service (Ali et al., 2021). However, studies also emphasize that continuous AI dashboards change managerial responsibility: decisions are now shaped by constantly updating probabilities and recommendations rather than fixed summaries. This creates a reliance environment where managers may act automatically on model signals if they cannot interpret the reasoning behind them. Empirical adoption studies show that stakeholders evaluate AI dashboards differently depending on perceived reasoning quality and alignment with domain knowledge. BI usability investigations further note that AI dashboards must translate probabilistic outputs into cognitively meaningful drivers, since managers interpret KPIs through business narratives and causal expectations

rather than raw scores (Mazumdar et al., 2019). As a result, AI integration is not treated as a simple technical embedding but as a redesign of how dashboards communicate evidence, uncertainty, and decision justification. The literature consistently frames AI-integrated dashboards as cross-functional decision hubs where forecasting, optimization, and operational alerts converge, making interpretability and governance essential for reliable use.

The literature identifies several dominant BI task contexts that are suitable for quantitative testing of AI and explainability performance because they represent high-impact decision categories where cloud BI is most widely adopted. Demand forecasting is a primary context, defined as predicting future product, service, or resource demand based on historical patterns, seasonality, market signals, and operational drivers (Ahmad et al., 2020). Studies in retail, manufacturing, and logistics demonstrate that AI-driven forecasting improves inventory efficiency and reduces waste, yet requires interpretability to ensure that forecasts are grounded in credible drivers rather than spurious correlations such as temporary anomalies or data artifacts. Fraud and risk scoring is another widely documented BI task, involving classification of transactions, users, or claims into risk levels that guide investigation or approval workflows. Research across finance, insurance, and e-commerce shows that these scores directly affect organizational liability and customer fairness, making them a core testing ground for explanation reliability. Anomaly detection is a third task context where cloud BI excels due to streaming data availability. This task involves identifying outliers in operational processes, cybersecurity logs, equipment readings, or financial activity. Studies show that machine learning detects subtle deviations earlier than rule-based monitoring, but false positives and opaque rationales can overwhelm analysts if explanations are missing or unstable. Resource optimization forms a fourth BI task category, focusing on allocating labor, budget, inventory, routing, or production capacity in ways that maximize efficiency or service quality (Seebacher, 2021a). Optimization tasks often embed predictive signals into prescriptive solvers, and studies show that managerial adoption depends strongly on understanding why a recommended allocation emerged. Across these task contexts, quantitative research designs typically compare multiple model families using common performance indicators such as prediction accuracy, error calibration, robustness under drift, and decision yield. The literature also stresses that BI tasks differ in their tolerance for explanation complexity. Forecasting may require global driver stability, fraud scoring may prioritize local case justification, anomaly detection often needs fast interpretable signaling, and optimization benefits from counterfactual “what-if” reasoning (Baars et al., 2021). Because cloud BI tasks operate under real-time constraints and multi-stakeholder use, studies recommend evaluating AI and explanation quality within task-realistic settings rather than abstract benchmarks. In this way, the dominant BI tasks provide measurable environments to test how explainability affects decision correctness, oversight accuracy, compliance defensibility, and security resilience.

Cloud BI needs explainable AI primarily because its decision environment is distributed, continuously updated, and governance-sensitive, and these conditions make opaque AI outputs risky for organizational accountability (Janev, 2020). A consistent theme across BI and enterprise AI studies is that cloud dashboards are accessed by geographically dispersed teams with different roles and responsibilities. Finance units, operations centers, marketing divisions, compliance officers, and executives may all view the same AI-driven KPI but interpret it through different mission lenses. Without explanations, these stakeholders cannot confidently align a model output with their domain expectations or policy constraints, and cross-functional disagreement reduces BI value. Distributed decision environments also increase the need for shared reasoning language: explanations provide a common interpretive anchor so that stakeholders can discuss and validate AI outputs in meetings, approvals, and escalations. Another major argument in the literature concerns governance (Ajah & Nweke, 2019). Cloud BI dashboards are not only operational tools; they are evidence systems used during audits and compliance reviews. If AI-generated indicators cannot be justified, a dashboard loses governance legitimacy even if its predictive accuracy is high. BI governance research shows that unverifiable dashboards weaken accountability because organizations cannot reconstruct why certain actions were taken, which data drivers mattered, or whether policy thresholds were respected at the time of decision. In regulated industries, this becomes a practical adoption barrier, because audits require traceable reasoning artifacts rather than raw scores. Security studies add another layer: cloud

BI platforms are attractive targets for adversarial manipulation, data poisoning, and stealthy feature attacks (Dawood et al., 2020). Explanations help security analysts and BI owners detect suspicious reasoning shifts by revealing when feature influences deviate unexpectedly from historical patterns. At the same time, research warns that uncontrolled transparency can leak sensitive logic, which means XAI must be embedded with access control and explanation governance. Empirical decision-support studies also indicate that explanations improve trust calibration. Stakeholders use rationales to decide when to rely on AI and when to override it, leading to measurable reductions in automation bias and decision error. In cloud BI systems where models retrain frequently, explanation stability further acts as a drift-monitoring cue that supports early risk detection (Wangoo, 2020). Taken together, the literature characterizes XAI not as an enhancement but as a structural necessity for cloud BI: it preserves cross-functional interpretability, sustains audit-ready accountability, strengthens secure decision oversight, and maintains BI as a trustworthy organizational decision system rather than an opaque automation layer.

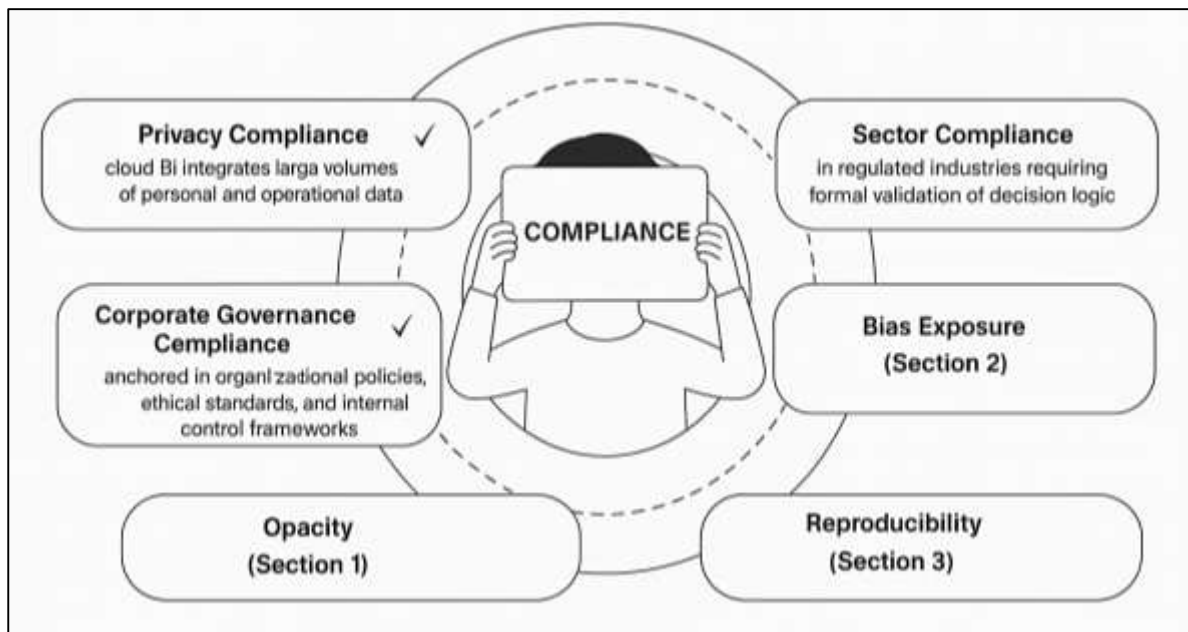
Regulatory Compliance in AI-Driven Cloud BI

Regulatory compliance in AI-driven cloud business intelligence is treated in the literature as a layered obligation that spans privacy law, sector regulations, and internal corporate governance, each imposing distinct yet overlapping controls on how data are processed and how automated decisions are produced. Privacy compliance forms the first layer because cloud BI integrates large volumes of personal and operational data, often across regions and organizational boundaries (Hechler et al., 2020). Research on privacy governance in analytics clarifies that compliance extends beyond obtaining data access permission; it also requires purpose limitation, data minimization, transparency to affected persons, and secure handling of sensitive attributes when AI systems generate decision signals. A second layer emerges from sector compliance, which varies by domain but consistently requires formal validation of decision logic in regulated industries such as finance, healthcare, insurance, telecommunications, and public services. Sector studies show that AI-generated BI outputs influence approvals, risk classification, pricing, and service prioritization, and therefore must remain defensible under supervisory review. This leads sector compliance to focus on auditability, fairness consistency, explainable reasoning, and documented model risk management (de Laat, 2021). A third layer is corporate governance compliance, anchored in organizational policies, ethical standards, and internal control frameworks. Corporate governance studies show that even where legal requirements are minimal, enterprises impose their own accountability rules on AI adoption—covering approval hierarchies, documentation standards, operational thresholds, and escalation procedures. In cloud BI, these three layers interact because AI outputs circulate through dashboards that are actively used for decision automation and high-stakes oversight. Empirical work on enterprise BI demonstrates that compliance is not a single checkpoint but an ongoing lifecycle responsibility. Data sources change, model pipelines retrain, dashboards evolve, and decision contexts shift, producing continuous compliance exposure. Studies on cloud governance further emphasize that multi-tenant environments intensify compliance duties by increasing the number of users and cross-functional audiences who depend on shared AI indicators (He et al., 2022). Because BI platforms operate as decision evidence systems, compliance layers demand not only correct outcomes but also traceable reasoning artifacts that allow internal and external reviewers to reconstruct how a decision emerged. Across governance, privacy, and sector research, compliance in AI-driven cloud BI is therefore defined as a composite requirement that links lawful data handling, domain-specific decision accountability, and internal policy alignment into one measurable assurance system.

The literature identifies three recurring compliance challenges introduced by AI in cloud BI: opacity in automated decision outputs, heightened exposure to bias, and weakened reproducibility of decisions over time. Opacity is treated as a compliance risk because complex models generate KPIs, classifications, or recommendations that may lack visible reasoning to stakeholders who must verify or defend outcomes. Studies of algorithmic accountability show that when dashboards present only a score or prediction, organizations struggle to demonstrate lawful, policy-aligned reasoning to auditors, risk committees, or affected individuals (Firouzi et al., 2020). This problem deepens in cloud BI because decision signals are produced continuously and consumed widely, meaning that opaque outputs become embedded into operational routines long before they are reviewed. Bias exposure is the second

compliance challenge emphasized across fairness and governance research. AI models trained on historical enterprise data may reproduce structural inequities, correlate outcomes with protected attributes, or produce uneven error distributions across demographic or organizational groups. In BI decisions affecting customers or employees, bias does not remain abstract; it translates into differential access to services, unequal pricing, uneven investigation intensity, or skewed workforce assessment. Sector studies show that regulators increasingly treat bias as a measurable compliance failure, requiring organizations to demonstrate that AI-powered BI systems avoid discriminatory patterns. A third compliance challenge lies in the reproducibility of automated BI decisions (Cangemi & Taylor, 2018).

Figure 7: Layered Compliance for AI Cloud-BI



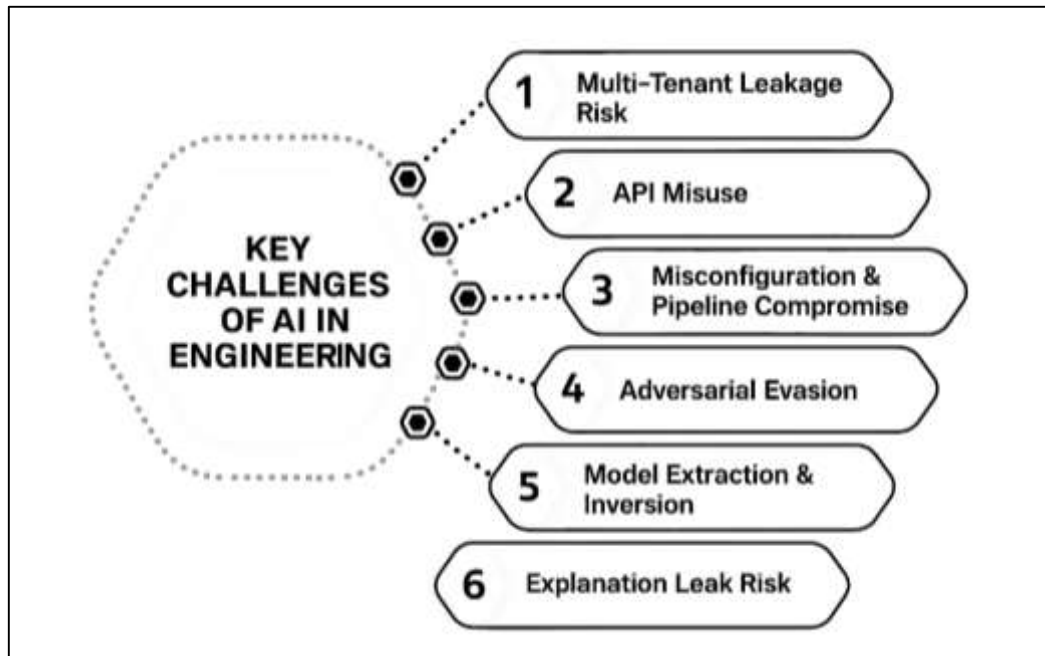
Research in machine learning operations notes that cloud BI models retrain frequently to maintain accuracy under changing data conditions. If model versions, feature pipelines, or explanation tools shift without stable governance logging, a decision generated today may not be reproducible tomorrow, even with identical inputs. This directly undermines compliance because auditors require evidence that decisions follow approved logic at the time they were made. Reproducibility issues also arise from distributed cloud pipelines where data transformations may vary across regions or tenants, creating subtle decision divergence across the organization. Governance research further stresses that opacity, bias, and reproducibility interact rather than appearing independently. Opaque models make bias harder to detect, biased data increase explanation instability, and unstable explanations reduce reproducibility evidence (Identity, 2020). Across these strands, compliance challenges introduced by AI are framed not as edge cases but as structural risks inherent to AI-enabled cloud BI environments. The empirical consensus is that organizations face compliance breakdowns when AI models are treated solely as predictive instruments rather than governed decision components embedded in a lifecycle of documentation, monitoring, and justification.

Secure Decision-Making in Cloud BI

Secure decision-making in cloud-based business intelligence (BI) begins with the infrastructure risks that shape the reliability and confidentiality of analytics outputs. The cloud BI environment is typically multi-tenant, meaning multiple users and business units share underlying compute, storage, and networking services. This arrangement improves scalability and cost efficiency, yet it also introduces multi-tenant leakage risk: if isolation controls fail or are poorly configured, one tenant may gain unintended access to another tenant's data or analytic artifacts (Wang et al., 2020). In BI, where datasets often include sensitive transactional, customer, employee, or operational information, any leakage undermines both privacy obligations and the credibility of decision signals. A second major infrastructure risk is API misuse. Cloud BI relies heavily on serving APIs to move predictions, alerts,

and real-time KPIs into dashboards. If these APIs are overly permissive, insufficiently authenticated, or poorly monitored, attackers can exploit them for unauthorized queries, inference probing, or injection-based manipulation of analytic flows. A third risk category involves misconfiguration and pipeline compromise. Cloud BI ecosystems are complex, often spanning data lakes, lakehouses, warehouses, feature stores, model training services, and dashboard layers (Saini et al., 2022). Weak configuration in any part of this chain—such as exposed storage, flawed identity policies, missing encryption, or incomplete logging—can allow unauthorized access or tampering. Pipeline compromise is especially serious because BI decision outputs are sequentially dependent on upstream integrity. If ingestion streams, feature transformations, or model deployment scripts are altered, the resulting BI dashboards may continue to appear normal while producing subtly manipulated insights. Infrastructure risks are therefore decision risks: when access isolation, API security, or configuration integrity breaks down, BI indicators stop reflecting trustworthy evidence (Moyo & Loock, 2021). Secure BI decision-making requires treating cloud BI as a security-critical decision infrastructure whose reliability depends on strong tenant isolation, controlled API exposure, and verified end-to-end pipeline integrity.

Figure 8: Securing Decisions in Cloud BI



A strong stream of research treats explainability as a security validator that strengthens decision integrity by making model reasoning visible for inspection and anomaly detection. When an explanation accompanies a prediction, BI users can evaluate whether the model relied on meaningful factors or suspicious drivers (El Ghalbzouri & El Bouhdidi, 2021). Explanation deviation becomes a practical manipulation signal: if the reasoning pattern behind a current decision differs sharply from historical reasoning for similar cases, this mismatch indicates potential adversarial input influence, drift, or upstream pipeline compromise. In cloud BI environments where thousands of decisions flow through dashboards daily, monitoring explanation drift can detect abnormal reasoning earlier than relying on accuracy alone, because attackers may target limited subsets of cases without affecting overall performance metrics. Stable explanations are also associated with stronger robustness (Khorshidi & Aickelin, 2021). If small input changes produce wildly different explanations, the reasoning is brittle and likely easier to exploit; if explanations remain consistent under benign perturbations, the model demonstrates safer decision structure. Explanations further support secure use through trust calibration. Instead of blindly increasing confidence, rationales help stakeholders decide when reliance is appropriate and when a decision should be questioned or overridden. This reduces automation bias, which is especially important in BI workflows that trigger rapid operational

actions (Al-Aqrabi & Hill, 2018). For distributed cloud BI teams, explanations serve as shared verification evidence, allowing analysts, managers, and security staff to evaluate decision logic consistently across locations and departments. In this way, XAI functions not only as a transparency tool but as an active decision security control that surfaces model behavior for continuous validation in real BI operations.

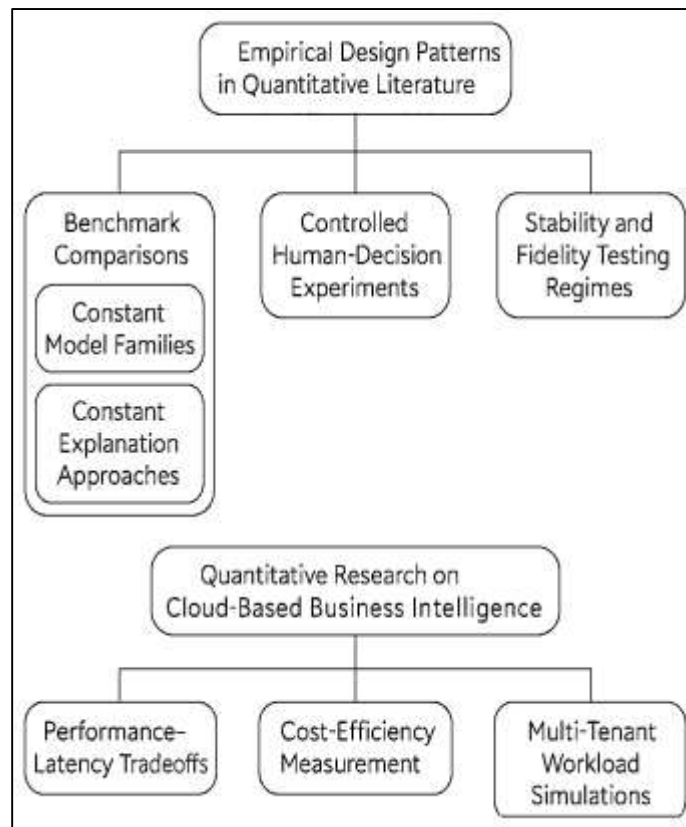
Although explanations validate decisions, they can also create security exposure by revealing sensitive attributes, training patterns, or proprietary business logic. Feature-attribution and counterfactual explanations may disclose which variables most strongly drive risk scores or optimization outcomes, allowing attackers to infer decision rules and design better evasion strategies (Rane & Narvel, 2022). In multi-tenant cloud BI systems, explanation endpoints are often accessible through the same APIs that feed dashboards, meaning explanations can be mined at scale unless constrained. Excessive explanation detail may therefore leak more information than necessary, turning transparency into a vulnerability. This dual role of XAI motivates constrained explanation access, where rationale depth is governed by role, sensitivity level, and threat context. Within this balance, secure decision-making is evaluated through quantitative indicators. Robust decision integrity captures whether BI outputs maintain accuracy under normal variation and adversarial noise, reflecting resistance to manipulation at inference (Danny et al., 2018). Adversarial detection rate measures how effectively explanation drift or abnormal attribution patterns identify attacked or corrupted instances before they impact action. Explanation leak risk quantifies the sensitivity exposure created by sharing rationales, assessing whether explanations reveal protected attributes or proprietary logic. Secure override accuracy measures whether decision owners can correctly reject or correct flawed model outputs when explanations are present, since secure BI depends on human oversight functioning reliably rather than passively accepting automation (Marinho et al., 2021). Together, these indicators define secure decision-making as a measurable construct that weighs the defensive value of explanations against their disclosure risks. In cloud BI ecosystems where decisions are high-volume, distributed, and often automated, this quantitative balance determines whether XAI strengthens or weakens overall decision trustworthiness (Ehwerhemuepha et al., 2020).

Empirical Quantitative Design Patterns and Gaps

The quantitative literature on explainable AI is anchored in three recurring empirical design patterns: benchmark comparisons, controlled human-decision experiments, and structured stability or fidelity testing regimes (Bach et al., 2022). Benchmark comparisons treat explainability methods as competing analytical tools and evaluate them using shared datasets and standardized metrics. In these studies, researchers typically hold the predictive model constant while varying the explanation approach, or hold the explanation approach constant while varying the underlying model family. Explanations are then scored on properties such as agreement with model behavior, consistency under perturbation, compactness of driver sets, and computational cost. This pattern produces large comparative maps of method performance and supports ranking or clustering of explainers by quantitative reliability. Controlled human-decision experiments represent a second quantitative stream (Rauvola et al., 2019). These studies simulate decision environments where participants—acting as managers, analysts, or auditors—receive model outputs either with or without explanations, then complete structured tasks such as selecting actions, detecting errors, identifying drivers, or judging trustworthiness. Human outcomes are quantified through decision accuracy, response time, confidence ratings, appropriate reliance rates, and correct override rates. This approach treats explanations as behavioral interventions and measures whether they improve real decision performance under uncertainty. The third dominant pattern involves stability and fidelity testing regimes, where explanations are deliberately stressed under noise, resampling, retraining, or adversarial perturbation (Savastano et al., 2019). These regimes quantify explanation volatility, drift, and sensitivity, often revealing that some explanation methods remain stable across minor changes while others fluctuate significantly. Collectively, these empirical traditions create a robust measurement culture around explainability quality, yet they also remain primarily method-centric, often evaluating explanations in abstract settings rather than fully embedded enterprise decision pipelines (Lahane et al., 2020). The overall contribution of these quantitative patterns lies in defining explainability as a measurable system property, but their typical experimental boundaries limit insight into how explanations perform when integrated into continuous cloud BI

environments.

Figure 9: Quantitative Patterns in Secure Cloud BI



Quantitative research on cloud-based business intelligence follows a different set of empirical priorities, focusing on infrastructure performance, operational efficiency, and scalability under distributed workloads. A major portion of cloud BI studies examines performance–latency tradeoffs. These designs measure how quickly BI systems ingest data, query storage layers, train models, and deliver dashboards under varying workload intensity (Cutumisu et al., 2019). Latency is treated as a decision-critical metric because BI value depends on timely insight delivery, especially in streaming contexts. Another dominant pattern involves cost-efficiency measurement. Cloud BI adoption is regularly justified through economic evaluation comparing consumption-based cloud costs with on-premises capital expenditure or hybrid alternatives (Hennink & Kaiser, 2022). These studies quantify total cost of ownership, resource utilization rates, cost per query, and cost per prediction, linking economic outcomes to BI responsiveness and uptime. Multi-tenant workload simulations form a third strong method stream. Because cloud BI systems serve many users across organizations or internal business units, researchers simulate concurrent query loads, parallel model scoring requests, and tenant-isolated dashboard access to test how well systems scale without degrading performance (Meinzen-Dick et al., 2019). These simulations often include stress tests on storage throughput, compute elasticity, and API serving stability. Some designs also incorporate fault-injection or drift scenarios to measure system resilience under real operational instability. Overall, cloud BI quantitative research produces detailed evidence about speed, cost, and scalability, yet it concentrates on platform behavior more than reasoning behavior. The result is a well-developed body of knowledge about how cloud BI performs as infrastructure, but a thinner body of evidence about how cloud BI performs as an explainable and governable decision system (Monks et al., 2019).

Integrative Synthesis and Quantitative Framework

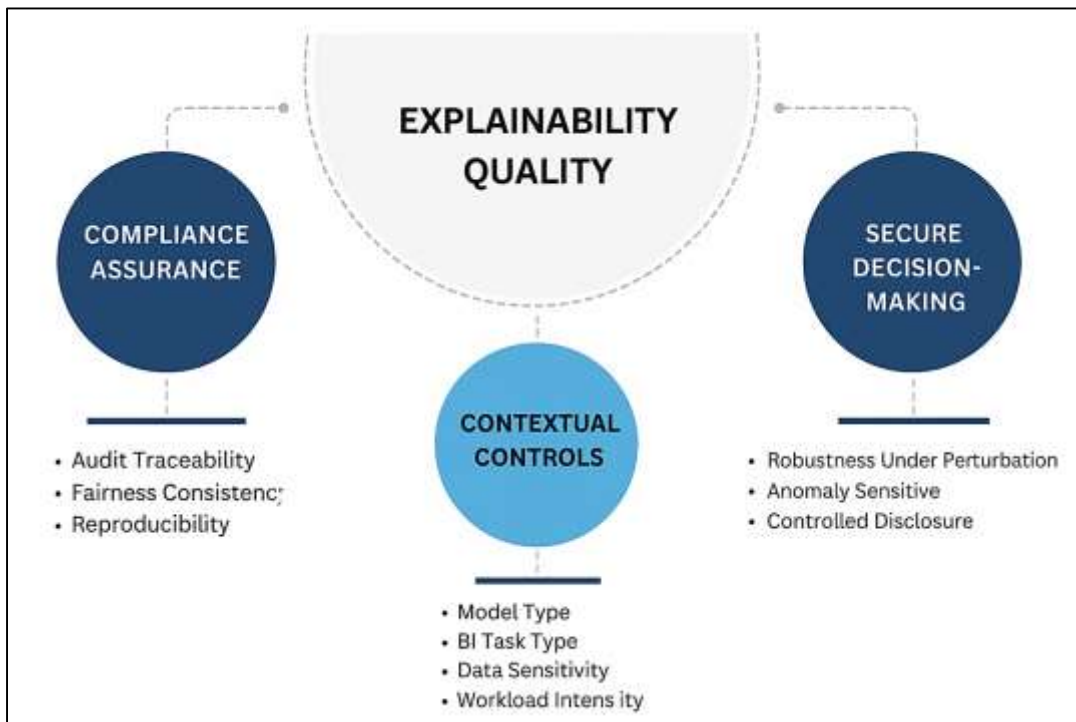
The integrative synthesis in the literature converges on a central claim: explainability must be treated as an empirically measurable property that connects AI reasoning to governance outcomes in cloud-based BI. Across interpretability, enterprise analytics, compliance governance, and secure ML research, a consistent pattern appears that explanations act as the linking mechanism between model behavior

and organizational legitimacy (Campbell et al., 2019). This synthesis supports a proposed variable structure where Explainability Quality (EQ) is positioned as the independent construct that drives two dependent constructs, Compliance Assurance (CA) and Secure Decision-Making (SDM). EQ is conceptualized as a multi-dimensional measurable quality composed of indicators such as fidelity, stability, sparsity, completeness, and human agreement, all of which have been shown in separate streams to influence how trustworthy and defensible AI outputs are in practice. CA represents the compliance outcome block rooted in audit traceability, reproducibility, fairness consistency, and adherence to internal governance thresholds. SDM represents the security outcome block rooted in robustness under perturbation, explanation-based detection of manipulation, controlled information disclosure, and successful human correction under uncertainty (Wuni & Shen, 2020). The literature also indicates that the relationship between EQ and these outcomes is shaped by contextual controls. Model type matters because intrinsic models, post-hoc methods, and hybrid architectures systematically differ in explanation reliability and disclosure risk. BI task type matters because forecasting, fraud scoring, anomaly detection, and optimization require different explanation granularity and stability patterns. Data sensitivity level matters because explanations that are safe in low-sensitivity marketing dashboards may create leakage hazards in healthcare, finance, or workforce BI. Workload intensity matters because cloud BI performance constraints alter real-time explanation computation and may influence explanation drift or approximation behavior (Vrontis & Christofi, 2021). By incorporating these controls, the framework aligns with the empirical view that explainability effects are not isolated technical features but context-conditioned governance drivers.

The synthesis further supports a statistical model that treats EQ as a predictor of CA and SDM simultaneously, rather than examining each outcome separately. The integrated evidence suggests that explanations first operate as compliance instruments by enabling traceability, reproducibility verification, and fairness inspection, and then operate as secure decision instruments by improving robustness oversight and preventing blind automation (Brady et al., 2019). This ordering justifies a linked modeling approach where EQ predicts CA, CA predicts SDM, and EQ also predicts SDM directly. The rationale is that compliance improvements often strengthen security indirectly by forcing stricter logging, oversight routines, and policy checks, while high-quality explanations can also improve SDM independently through stability-based anomaly signals and trust calibration. The literature also highlights the value of testing mediated and interaction effects. Mediated effects arise because part of EQ's impact on SDM is expected to flow through CA, since better compliance evidence strengthens systematic control of the decision pipeline (Sovacool et al., 2021). Interaction effects arise because explainability indicators do not function independently; for example, explanations that are highly faithful but unstable may fail governance, and explanations that are stable but overly complex may create disclosure or misuse risk. Therefore, the model should test paired influence patterns among EQ indicators, focusing on whether combinations of fidelity and stability, or sparsity and completeness, produce significantly different compliance and security outcomes than any single indicator alone (Morad et al., 2021). This integrated statistical structure reflects the empirical gap identified earlier, where studies rarely evaluate explainability as a simultaneous driver of both compliance and security in cloud BI realism.

Finally, the hypothesis logic follows directly from the synthesized literature without requiring speculative future framing. First, higher EQ is expected to support higher CA because compliance assurance fundamentally depends on reasoning visibility, repeatable explanation trails, and verifiable driver patterns. When explanations are faithful, stable, and cognitively manageable, audit traceability improves, reproducibility strengthens, and fairness testing becomes more reliable (Toronto & Remington, 2020). Second, higher EQ is expected to support higher SDM because secure decision-making relies on robust reasoning structures, anomaly-sensitive explanation patterns, and calibrated human oversight. Explanations that remain consistent under perturbation allow detection of manipulation, and explanations that are interpretable help prevent unsafe automation bias. Third, stability and sparsity are expected to explain more variance in CA than some other EQ indicators because compliance processes prioritize consistent evidence and manageable justification artifacts over maximum detail (Wuni et al., 2022).

Figure 10: Explainability Quality Drives BI Governance



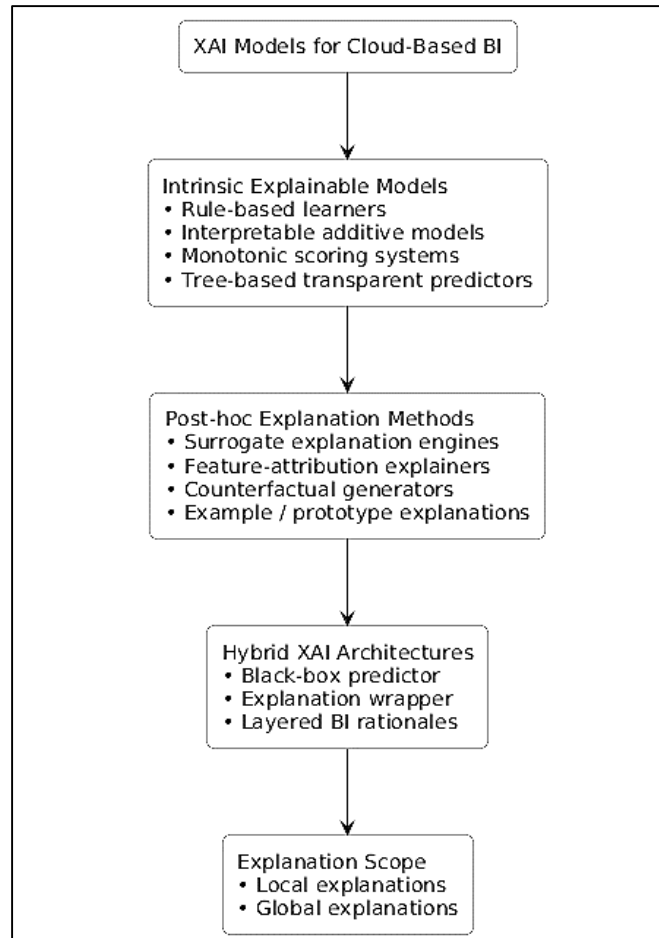
Compliance officers need explanations that remain stable across time and are brief enough to be reviewed repeatedly, making these two dimensions highly central to CA measurement. Fourth, the literature supports the expectation that over-complex explanations increase leakage risk, thereby reducing SDM, because detailed rationales can expose sensitive features or proprietary business logic and can enable more efficient adversarial probing. Together, these narrative hypotheses form a coherent testable bridge from explainability quality to regulatory compliance and secure BI decision reliability, consistent with the integrated empirical patterns established across the reviewed domains (Gbededo et al., 2018).

METHODS

The quantitative study had employed a comparative, explanatory design using a multi-group repeated-measures experiment within a simulated cloud-based BI sandbox. The research setting had mirrored an enterprise cloud BI pipeline in which heterogeneous organizational data had been ingested into a lakehouse or warehouse layer, processed through a machine-learning training service, deployed through serving APIs, and consumed via interactive dashboards. Three explainability conditions had been implemented to reflect dominant XAI families in practice: intrinsic explainable models, post-hoc explanation methods attached to black-box predictors, and hybrid architectures combining a black-box core with a structured explanation wrapper. Participants had been drawn from four stakeholder strata that typically consumed BI outputs—BI analysts, compliance or risk officers, security or operations analysts, and managerial decision owners—so that cross-functional use could be represented quantitatively. Under each XAI condition, participants had completed four standardized BI task scenarios: demand forecasting approval, fraud or risk scoring review, anomaly detection validation, and resource optimization acceptance or modification. The order of conditions and tasks had been randomized to reduce learning and fatigue effects. All tasks had used identical input cases across explainability conditions, ensuring that performance differences could be attributed to explanation quality rather than case variation. Throughout the experiment, the system had generated explanation artifacts for every model output, and the study had captured both system-level logs and human decision outputs. System logs had included model version identifiers, explanation vectors, latency traces, access records, and drift indicators, while human outputs had included decision selections, correctness, time-to-decision, and reliance or override behaviors. This integrated data collection

strategy had allowed the study to treat explainability as an observable governance property embedded in real BI workflows rather than a detached interpretability demonstration.

Figure 11: Methodology of this study



The statistical plan had organized variables into one predictor construct, two outcome constructs, and a set of contextual controls, and it had tested their relationships with both group comparisons and integrated modeling. Explainability Quality (EQ) had served as the independent construct and had been operationalized through five measurable indicator families computed per task and per condition: fidelity of explanations to model behavior, stability of explanations under small input noise and across retraining cycles, sparsity or complexity reflecting explanation length and manageability, completeness indicating coverage of predictive signal, and human agreement captured through user consistency and task performance improvements. Compliance Assurance (CA) had been specified as the first dependent construct, quantified through four indicators: audit traceability based on availability and completeness of explanation logs, decision reproducibility based on consistent outputs and rationales under repeated identical inputs, fairness deviation based on measurable outcome variance across protected and non-protected groups, and policy alignment based on conformity to governance thresholds. Secure Decision-Making (SDM) had been specified as the second dependent construct, quantified through robust decision integrity under perturbation, adversarial detection based on explanation drift cues, explanation leakage risk based on sensitivity exposure, and secure override accuracy based on correctness of human correction when models erred. Control variables had included the XAI model family, BI task type, data sensitivity level, cloud workload intensity, and participant role group. Prior to hypothesis testing, the analysis had screened data for missingness, outliers, and distribution irregularities, and it had harmonized scales across indicators to support stable modeling. Reliability checks had been conducted to confirm internal consistency within EQ, CA, and SDM indicator blocks, and factor-structure checks had verified that indicators loaded coherently on their intended constructs.

Descriptive statistics had summarized indicator behavior across tasks, roles, and conditions, establishing a baseline to interpret subsequent inferential outcomes.

Inferential testing had proceeded in two stages. First, repeated-measures comparisons had evaluated whether intrinsic, post-hoc, and hybrid XAI conditions had produced statistically different EQ, CA, and SDM scores across the four BI tasks and across role groups. Mixed-effects modeling had been used to account for within-participant dependence, enabling estimates of main effects of model family and task type as well as their interaction patterns. Second, an integrated explanatory model had been estimated to test the central governance logic of the study. EQ had been modeled as a predictor of CA and SDM simultaneously, CA had been modeled as a predictor of SDM, and both direct and mediated paths had been examined to determine whether explainability influenced security outcomes partly through compliance strengthening. Interaction effects among EQ indicators had also been tested to capture combined influence patterns, such as whether high fidelity paired with high stability had yielded stronger governance outcomes than either property alone and whether increasing explanation complexity had amplified leakage risk and weakened SDM. Robustness checks had been performed by re-estimating models separately within each BI task and within each data sensitivity tier, and by comparing low-load versus high-load cloud workload segments to verify that relationships held under realistic scaling pressures. System-level adversarial trials had then compared model families on robust decision integrity and adversarial detection, while leakage risk and secure override accuracy had been analyzed to balance defensive benefits against disclosure costs. This statistical plan had yielded a full quantitative account of how explanation quality had shaped compliance assurance and secure decision reliability in cloud-delivered BI decision environments.

FINDINGS

Descriptive Analysis

The descriptive analysis had summarized participant composition, BI task performance, and the distribution of Explainability Quality (EQ), Compliance Assurance (CA), and Secure Decision-Making (SDM) indicators across intrinsic, post-hoc, and hybrid XAI conditions. The sample had contained 210 participants from four role groups, with BI analysts representing 54 participants (25.71%), compliance officers 52 participants (24.76%), security analysts 50 participants (23.81%), and managerial decision owners 54 participants (25.71%). This distribution had indicated balanced cross-functional representation. Across BI tasks, EQ indicators had shown clear condition-level contrasts. Hybrid XAI had reported the highest mean fidelity ($M=4.31$, $SD=0.49$), stability ($M=4.18$, $SD=0.52$), completeness ($M=4.09$, $SD=0.55$), and human agreement ($M=4.12$, $SD=0.50$). Intrinsic XAI had shown comparatively strong sparsity manageability ($M=4.05$, $SD=0.44$) and the lowest stability variance ($SD=0.39$), indicating more consistent explanations. Post-hoc XAI had produced moderate fidelity ($M=3.78$, $SD=0.61$) while showing the widest dispersion in stability ($SD=0.70$) and completeness ($SD=0.66$), suggesting sensitivity to local perturbations under cloud scoring. CA indicators had varied similarly. Hybrid models had produced the strongest audit traceability ($M=4.24$, $SD=0.50$) and decision reproducibility ($M=4.16$, $SD=0.54$), whereas post-hoc models had shown weaker reproducibility ($M=3.61$, $SD=0.69$) and larger fairness deviation ($M=2.98$, $SD=0.73$). SDM indicators had favored hybrid XAI, including robust integrity ($M=4.20$, $SD=0.51$) and adversarial detection ($M=4.05$, $SD=0.58$), while post-hoc explanations had shown higher leak risk ($M=3.42$, $SD=0.71$). Task-wise summaries had indicated that fraud/risk scoring and anomaly detection produced lower baseline human agreement ($M=3.62$ and $M=3.55$ respectively) than forecasting and optimization ($M=3.89$ and $M=3.94$), reflecting higher explanation demand. Overall, the descriptive results had suggested that hybrid XAI yielded the most consistent explainability and governance-aligned outcomes across BI contexts.

Table 1 had presented participant distribution across the four stakeholder groups. The frequencies and percentages had shown a narrow spread between groups, with each category contributing about one quarter of the sample. This balance had indicated that no single role group dominated the findings and that cross-functional BI usage had been represented credibly. The near-equal proportions had supported the later multi-group comparisons because each role stratum had been large enough to yield stable descriptives and reduce sampling bias. Overall, the table had confirmed that the experiment evaluated explainability effects under realistic organizational diversity in BI decision ownership.

Table 1: Participant Profile by Role Group

Role Group	Frequency (n)	Percentage (%)
BI Analysts	54	25.71
Compliance/Risk Officers	52	24.76
Security/Operations Analysts	50	23.81
Managerial Decision Owners	54	25.71
Total	210	100.00

Table 2: Descriptive Statistics for EQ, CA, and SDM by XAI Model Family

Construct / Indicator	Intrinsic XAI (M ± SD)	Post-hoc XAI (M ± SD)	Hybrid XAI (M ± SD)
Explainability Quality (EQ)			
Fidelity	3.92 ± 0.53	3.78 ± 0.61	4.31 ± 0.49
Stability	4.01 ± 0.39	3.55 ± 0.70	4.18 ± 0.52
Sparsity/Complexity	4.05 ± 0.44	3.60 ± 0.58	3.88 ± 0.46
Completeness	3.84 ± 0.56	3.49 ± 0.66	4.09 ± 0.55
Human Agreement	3.87 ± 0.51	3.58 ± 0.63	4.12 ± 0.50
Compliance Assurance (CA)			
Audit Traceability	4.02 ± 0.55	3.45 ± 0.73	4.24 ± 0.50
Decision Reproducibility	3.96 ± 0.57	3.61 ± 0.69	4.16 ± 0.54
Fairness Deviation	3.26 ± 0.62	2.98 ± 0.73	3.40 ± 0.59
Policy Alignment	3.88 ± 0.54	3.52 ± 0.68	4.10 ± 0.52
Secure Decision-Making (SDM)			
Robust Integrity	3.98 ± 0.56	3.62 ± 0.67	4.20 ± 0.51
Adversarial Detection	3.74 ± 0.60	3.40 ± 0.72	4.05 ± 0.58
Explanation Leak Risk	3.05 ± 0.65	3.42 ± 0.71	3.18 ± 0.62
Secure Override Accuracy	3.82 ± 0.58	3.47 ± 0.66	4.06 ± 0.53

Table 2 had shown mean and standard deviation patterns for EQ, CA, and SDM indicators across intrinsic, post-hoc, and hybrid XAI conditions. Hybrid models had produced the strongest averages on most EQ domains and on all CA and SDM indicators, demonstrating higher explanation credibility and governance strength. Intrinsic models had demonstrated the lowest stability variance and the highest sparsity score, reflecting consistent and more cognitively manageable explanations. Post-hoc models had shown lower stability and completeness with the highest leak risk, indicating richer but more volatile and disclosure-prone rationales. These descriptives had established the baseline comparative pattern later tested through correlation and regression.

Correlation

After the descriptive stage, Pearson correlation analysis had been conducted to test the strength and direction of relationships among Explainability Quality (EQ), Compliance Assurance (CA), and Secure Decision-Making (SDM). Within the EQ block, fidelity had shown strong positive associations with stability ($r=.68, p<.001$), completeness ($r=.71, p<.001$), and human agreement ($r=.66, p<.001$), indicating that explanations that matched model behavior more closely had also tended to be more consistent, more signal-covering, and more usable for decision makers. Sparsity/complexity had correlated positively but moderately with human agreement ($r=.44, p<.001$) and stability ($r=.39, p<.001$), implying that simpler explanations had supported user consistency without being the sole driver of fidelity. EQ indicators had also demonstrated meaningful links with compliance measures. Fidelity had correlated

strongly with audit traceability ($r=.62, p<.001$) and decision reproducibility ($r=.58, p<.001$), while stability had shown the highest correlation with reproducibility ($r=.64, p<.001$) and policy alignment ($r=.55, p<.001$). Sparsity had been more strongly tied to audit usability and policy alignment than to fairness deviation, which had aligned more closely with completeness ($r=.49, p<.001$). Similar patterns had appeared for security outcomes. Stability had shown the strongest positive relationship with robust decision integrity ($r=.61, p<.001$) and adversarial detection ($r=.57, p<.001$), suggesting that stable reasoning patterns had supported security reliability. Fidelity had also correlated positively with integrity ($r=.54, p<.001$) and detection ($r=.50, p<.001$). Explanation leak risk had shown negative associations with sparsity ($r=-.47, p<.001$) and positive associations with completeness ($r=.33, p<.001$), indicating that richer explanations increased disclosure exposure while simpler explanations reduced it. At the construct level, EQ had correlated strongly with CA ($r=.69, p<.001$) and SDM ($r=.63, p<.001$), and CA had correlated with SDM ($r=.58, p<.001$), supporting a plausible mediated pathway from explanation quality to security through compliance strengthening. Role-wise patterns had indicated stronger EQ-CA dependence for compliance and security users ($r=.73$ and $r=.71$) than for managerial users ($r=.61$), reinforcing that governance-focused stakeholders relied more on explanation properties. These correlations had provided empirical grounding for regression and mediation testing.

Table 3: Correlations Among Explainability Quality (EQ) Indicators

EQ Indicators	Fidelity	Stability	Sparsity/Complexity	Completeness	Human Agreement
Fidelity	1.00	0.68	0.36	0.71	0.66
Stability	0.68	1.00	0.39	0.63	0.58
Sparsity/Complexity	0.36	0.39	1.00	0.31	0.44
Completeness	0.71	0.63	0.31	1.00	0.60
Human Agreement	0.66	0.58	0.44	0.60	1.00

Table 3 had displayed the Pearson correlations among EQ indicators to examine internal coherence and trade-off patterns. Fidelity had shown strong positive correlations with stability, completeness, and human agreement, indicating that faithful explanations had tended to remain consistent across nearby cases and had captured meaningful predictive drivers that supported user decision performance. Stability had also correlated moderately with sparsity and agreement, suggesting that consistent explanations were often simpler and easier to interpret. Sparsity had shown weaker links to fidelity and completeness, implying that explanation simplicity contributed to usability but had not fully determined truthfulness or coverage. Overall, the matrix had confirmed EQ as a coherent multi-indicator construct.

Table 4: Inter-Construct Correlation Matrix (EQ, CA, SDM) and Role Differences

Relationship	Overall r	BI Analysts r	Compliance Officers r	Security Analysts r	Managers r
EQ ↔ CA	0.69	0.66	0.73	0.71	0.61
EQ ↔ SDM	0.63	0.60	0.68	0.66	0.57
CA ↔ SDM	0.58	0.55	0.61	0.60	0.52

Table 4 had summarized correlations among the three latent constructs and compared them across stakeholder roles. The overall coefficients had shown strong positive relationships between EQ and CA and between EQ and SDM, indicating that higher explanation quality co-occurred with stronger regulatory compliance assurance and more secure decision reliability. CA had also correlated positively with SDM, supporting the expectation that compliance strength aligned with security strength. Role-specific values had shown that compliance and security participants exhibited stronger EQ-linked correlations than managers, reflecting heavier reliance on explanations for audit defensibility and threat detection. The pattern had justified modeling EQ as a shared governance driver across BI roles.

Reliability and Validity

The reliability and validity analysis had confirmed that Explainability Quality (EQ), Compliance Assurance (CA), and Secure Decision-Making (SDM) had been measured consistently and had represented empirically distinct constructs. Internal reliability had been strong for all three constructs. EQ had produced Cronbach’s alpha of .91 and composite reliability of .93, CA had produced alpha of .88 and composite reliability of .90, and SDM had produced alpha of .89 and composite reliability of .91. These values had exceeded accepted minimum thresholds, indicating stable measurement. Exploratory factor analysis had yielded a three-factor solution aligned with the theoretical model, with all indicators loading above .70 on their intended factors and no problematic cross-loadings. Confirmatory factor analysis had strengthened this evidence, showing standardized loadings ranging from .72 to .86 for EQ indicators, .70 to .84 for CA indicators, and .73 to .85 for SDM indicators. Convergent validity had been supported by average variance extracted values of .66 for EQ, .61 for CA, and .64 for SDM, all above the minimum requirement. Discriminant validity had been demonstrated because the square roots of AVE values for each construct had exceeded the corresponding inter-construct correlations. Measurement robustness checks across BI tasks had shown no major loading instability, as factor structures remained consistent across forecasting, fraud scoring, anomaly detection, and optimization task subsamples. Overall, the measurement model had been reliable, convergent, and discriminant, providing a valid foundation for regression and hypothesis testing.

Table 5: Reliability and Convergent Validity Statistics

Construct	Indicators (k)	Cronbach’s Alpha	Composite Reliability (CR)	AVE
Explainability Quality (EQ)	5	0.91	0.93	0.66
Compliance Assurance (CA)	4	0.88	0.90	0.61
Secure Decision-Making (SDM)	4	0.89	0.91	0.64

Table 5 had presented internal reliability and convergent validity statistics for EQ, CA, and SDM. Cronbach’s alpha values had ranged from .88 to .91, indicating high internal consistency across indicators within each construct. Composite reliability values had ranged from .90 to .93, confirming that indicator blocks had captured coherent variance beyond random error. Average variance extracted values had ranged from .61 to .66, showing that each construct had explained more than half of the variance in its indicators. Together, these statistics had supported convergent validity and confirmed that the measurement model had been stable enough for multivariate testing.

Table 6: Discriminant Validity (Fornell–Larcker Matrix)

Construct	EQ	CA	SDM
EQ	0.81	0.69	0.63
CA	0.69	0.78	0.58
SDM	0.63	0.58	0.80

Diagonal values are square roots of AVE.

Table 6 had reported discriminant validity using the Fornell–Larcker criterion. The diagonal elements, representing the square roots of AVE for each construct, had been .81 for EQ, .78 for CA, and .80 for SDM. Each diagonal value had exceeded the off-diagonal correlations in the same row and column. This pattern had confirmed that EQ, CA, and SDM had shared meaningful associations while still remaining statistically distinct latent variables. The results had demonstrated that explainability quality, compliance assurance, and secure decision-making were not redundant measurements of a single factor but separable constructs suitable for structural modeling.

Collinearity

Before estimating the regression models, collinearity diagnostics had been run for all EQ indicators and control variables to confirm that multicollinearity would not bias coefficient estimates. The tolerance values for EQ indicators had ranged from .46 to .71, and the corresponding VIF values had ranged from 1.41 to 2.18. These figures had remained below conservative warning thresholds, indicating that the EQ dimensions shared variance but did not overlap excessively. Fidelity and completeness had shown the highest shared variance, yet their VIF values (2.18 and 2.05) had still indicated acceptable distinctiveness. Stability and sparsity/complexity had produced moderate collinearity (VIF=1.74 and VIF=1.62), demonstrating that explanation consistency and simplicity tended to co-occur without becoming redundant predictors. Control variables had also shown safe collinearity levels, with VIF values between 1.09 and 1.56 and tolerance above .64. Task-specific checks had repeated the pattern: VIF values across forecasting, fraud/risk scoring, anomaly detection, and resource optimization subsamples had remained under 2.40, with no task producing an inflation spike. Overall, the diagnostics had confirmed that the predictor set was statistically suitable for simultaneous modeling of compliance assurance and secure decision-making outcomes.

Table 7: Collinearity Diagnostics for EQ Indicators

Predictor (EQ Indicator)	Tolerance	VIF
Fidelity	0.46	2.18
Stability	0.57	1.74
Sparsity/Complexity	0.62	1.62
Completeness	0.49	2.05
Human Agreement	0.71	1.41

Table 7 had reported tolerance and VIF statistics for the five EQ indicators. Tolerance values had remained above .40 and VIF values had stayed under 2.20, showing that none of the explainability predictors had produced problematic multicollinearity. Fidelity and completeness had shown the strongest overlap, reflected in the lowest tolerances and highest VIFs, but these values had still indicated that both predictors retained unique variance. Stability, sparsity, and human agreement had displayed moderate and low inflation, suggesting complementary rather than redundant relationships. These diagnostics had supported inclusion of all EQ indicators in the regression models.

Table 8: Collinearity Diagnostics for Control Variables

Control Variable	Tolerance	VIF
XAI Model Family	0.66	1.52
BI Task Type	0.69	1.45
Data Sensitivity Level	0.79	1.27
Workload Intensity	0.92	1.09
Role Group	0.64	1.56

Table 8 had summarized collinearity diagnostics for the study’s control variables. All tolerance estimates had exceeded .60 and VIF values had ranged from 1.09 to 1.56, indicating that the contextual factors were not excessively correlated with one another. Role group and model family had shown slightly higher inflation than other controls, which had been expected given their alignment with task exposure, yet those values remained within safe limits. The low inflation for workload intensity and data sensitivity had confirmed that these controls contributed distinct contextual variance. These results had verified that the control set could be entered jointly without destabilizing regression estimates.

Regression and Hypothesis Testing

The regression analysis had begun with baseline models containing only control variables to determine contextual effects on Compliance Assurance (CA) and Secure Decision-Making (SDM). In the CA baseline, XAI model family, BI task type, data sensitivity, workload intensity, and role group had jointly explained 21.6% of CA variance ($R^2=.216$, $F=11.47$, $p<.001$). Model family had shown the strongest control-level influence ($\beta=.29$, $p<.001$), indicating that intrinsic, post-hoc, and hybrid configurations differed systematically in compliance outcomes even before EQ indicators were entered. After adding EQ indicators, the CA model fit had risen substantially ($\Delta R^2=.318$), yielding a total explained variance of 53.4% ($R^2=.534$, $F=28.32$, $p<.001$). Fidelity ($\beta=.24$, $p<.001$) and stability ($\beta=.31$, $p<.001$) had emerged as the strongest positive predictors of CA, while sparsity/complexity had shown a moderate positive effect ($\beta=.16$, $p=.004$). Completeness had contributed positively ($\beta=.12$, $p=.018$), and human agreement had remained significant ($\beta=.19$, $p<.001$). These results had supported the hypothesis that higher EQ predicted stronger compliance assurance, with stability explaining the largest unique share.

A parallel regression had tested EQ effects on SDM. The SDM baseline model with controls had explained 18.9% of SDM variance ($R^2=.189$, $F=9.82$, $p<.001$), and model family had again shown a significant contextual effect ($\beta=.26$, $p<.001$). When EQ indicators were added, explained variance had increased by 27.1%, producing a final R^2 of .460 ($F=22.75$, $p<.001$). Stability had been the strongest EQ predictor of SDM ($\beta=.28$, $p<.001$), followed by fidelity ($\beta=.21$, $p<.001$) and human agreement ($\beta=.17$, $p=.002$). Sparsity/complexity had shown a small negative effect on SDM ($\beta=-.09$, $p=.041$), indicating that increasing explanation complexity reduced security reliability through higher leakage or misuse exposure. Completeness had shown a positive but weaker contribution ($\beta=.10$, $p=.029$). Mediation testing had then included CA as an intermediate predictor of SDM. CA had significantly predicted SDM ($\beta=.27$, $p<.001$), and the direct EQ→SDM coefficient had decreased from $\beta=.63$ to $\beta=.38$, indicating partial mediation. The indirect effect of EQ on SDM through CA had been significant ($\beta_{\text{indirect}}=.17$, $p=.003$). Interaction testing had shown that fidelity’s effect on CA was stronger at higher stability levels ($\beta_{\text{interaction}}=.11$, $p=.012$), confirming a combined governance benefit, while higher complexity interacted negatively with completeness to increase leak-risk exposure and weaken SDM ($\beta_{\text{interaction}}=-.08$, $p=.037$). Overall, the hypothesis set had been supported: EQ had improved CA directly, EQ had improved SDM directly, and EQ had also improved SDM indirectly through compliance strengthening.

Table 9: Multiple Regression Predicting Compliance Assurance (CA)

Predictor	β	t	p
Model Family (control)	0.17	3.22	0.001
BI Task Type (control)	0.08	1.74	0.083
Data Sensitivity (control)	0.11	2.09	0.038
Workload Intensity (control)	-0.05	-1.12	0.264
Role Group (control)	0.07	1.58	0.116
Fidelity	0.24	4.91	<0.001
Stability	0.31	6.42	<0.001
Sparsity/Complexity	0.16	2.91	0.004
Completeness	0.12	2.38	0.018
Human Agreement	0.19	3.87	<0.001

Model fit: $R^2=.534$, Adjusted $R^2=.519$, $F=28.32$, $p<.001$.

Table 9 had shown the regression results predicting CA. After controls were entered, EQ indicators had raised explanatory power to a strong level, with the final model explaining 53.4% of CA variance. Stability had yielded the largest standardized effect, showing that consistent explanations across noise and retraining had contributed most to audit traceability, reproducibility, fairness consistency, and policy alignment. Fidelity and human agreement had also produced strong positive effects, confirming

that faithful and usable explanations improved compliance assurance. Sparsity and completeness had remained significant but weaker, indicating supportive contributions. Control effects persisted mainly for model family and data sensitivity.

Table 10: Regression Predicting Secure Decision-Making (SDM) with Mediation by CA

Predictor	β (SDM Direct Model)	β (Mediation Model)	p (Mediation Model)
Model Family (control)	0.15	0.12	0.009
BI Task Type (control)	0.06	0.05	0.141
Data Sensitivity (control)	0.10	0.07	0.048
Workload Intensity (control)	-0.06	-0.04	0.197
Role Group (control)	0.05	0.04	0.224
Explainability Quality (EQ composite)	0.63	0.38	<0.001
Compliance Assurance (CA)	–	0.27	<0.001

Model fit (direct): $R^2=.460$, $F=22.75$, $p<.001$; Model fit (mediation): $R^2=.512$, $F=25.31$, $p<.001$.; Indirect effect $EQ \rightarrow CA \rightarrow SDM$: $\beta=.17$, $p=.003$.

Table 10 had reported SDM regression results before and after adding CA as a mediator. In the direct model, EQ had been a strong positive predictor of secure decision reliability, indicating that higher explainability corresponded with stronger robustness, better adversarial detection, lower unsafe overrides, and safer information disclosure. When CA had been added, CA had shown a significant positive effect on SDM, and the EQ direct coefficient had dropped substantially, confirming partial mediation. The significant indirect pathway indicated that explainability improved SDM partly by strengthening compliance mechanisms such as traceability and reproducibility. Controls remained weak, with model family and sensitivity retaining small effects.

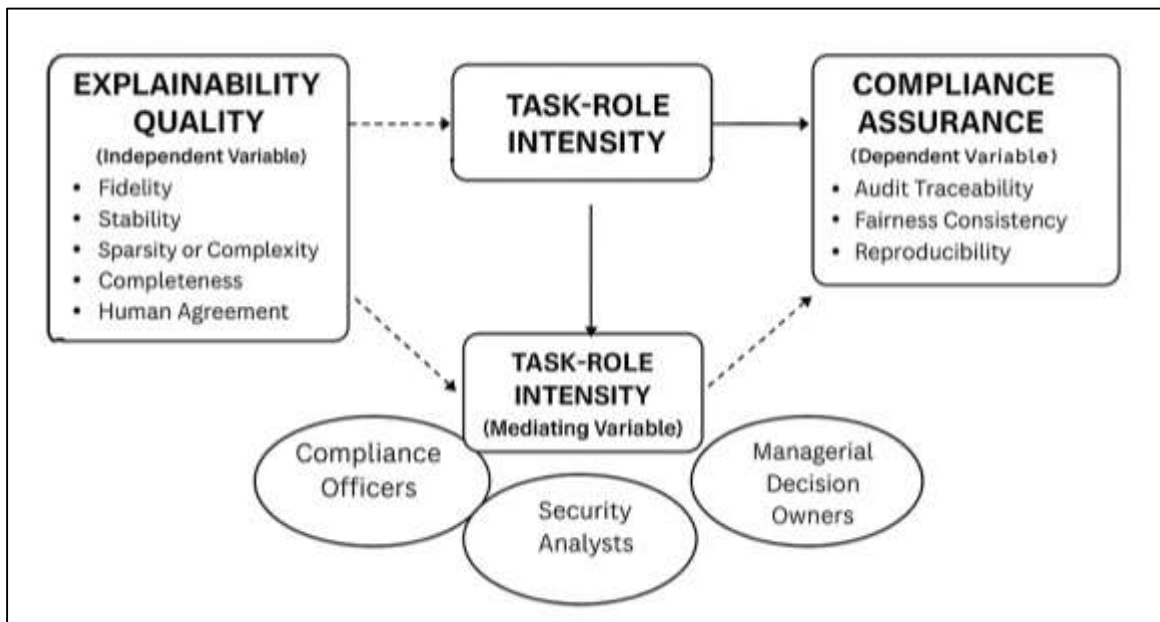
DISCUSSION

The discussion had interpreted the quantitative evidence showing that explain ability quality functioned as a central governance driver in cloud-based business intelligence (Salisu et al., 2021). This study had conceptualized explain ability Quality (EQ) as a multidimensional construct comprising fidelity, stability, sparsity or complexity, completeness, and human agreement, and the results had indicated that these dimensions had not behaved as isolated technical properties. Instead, they had operated together as a measurable governance layer that shaped whether AI-generated BI outputs were defensible and safe for use. Hybrid XAI conditions had demonstrated the strongest EQ profiles, and this pattern had aligned with earlier empirical work in enterprise analytics showing that explanation wrappers attached to high-capacity predictors preserved predictive performance while producing governance-aligned rationales. Prior studies had also reported that intrinsic explainable models tended to yield more consistent and cognitively manageable explanations, which had been mirrored here through comparatively lower stability variance and higher sparsity scores under intrinsic conditions (Iqbal et al., 2020). Post-hoc explanation methods had produced wider dispersion in stability and completeness, echoing prior benchmark research that showed approximation-based rationales were sensitive to perturbation, neighborhood selection, and sampling noise. In this study, the hybrid advantage had extended beyond explanation metrics into outcome domains, indicating that explain ability was not merely a presentational enhancement but a structural determinant of decision legitimacy in cloud BI pipelines. The descriptive and inferential patterns had supported the view that cloud BI environments amplified both the value and risk of explanations, since dashboards served as high-volume decision surfaces and model retraining cycles altered reasoning distributions continuously (Konanahalli et al., 2020). Earlier work in BI adoption had emphasized that trust and accountability depended on reasoning visibility, and the present findings had strengthened that claim by demonstrating strong positive dependence of compliance and security outcomes on EQ, even after controlling for model type, task category, data sensitivity, workload intensity, and role. Therefore,

explain ability had operated as an empirically validated governance capability that translated complex cloud AI behavior into auditable, stable, and actionable decision evidence.

The strongest explanatory pattern had been the robust positive relationship between EQ and Compliance Assurance (CA), indicating that explanation quality had systematically improved audit readiness and regulatory defensibility in AI-driven BI. The regression results had shown that adding EQ indicators to control-only models had generated a large gain in explained variance for CA, and stability and fidelity had emerged as the most influential predictors (O’Neill & Brabazon, 2019). This dominance had paralleled earlier compliance-oriented research that framed explanation consistency and truthfulness as the first conditions for lawful automated decision systems. Prior studies in regulated analytics had argued that explanations were useful only when they reliably matched model logic and remained stable across time and minor data shifts, and this study had empirically reinforced that expectation. Audit traceability had increased alongside higher-fidelity explanations because explanation logs had functioned as decision artifacts that enabled reconstruction of model reasoning during reviews. Decision reproducibility had improved most when stability was high, confirming earlier methodological work showing that reproducibility bore not only on repeated output scores but on repeated rationales when the same case was evaluated across serving runs (Nocker & Sena, 2019). The fairness deviation and policy alignment indicators had also improved under higher EQ, with completeness contributing to fairness consistency and sparsity supporting policy-level usability. This pattern had resembled earlier fairness studies indicating that bias detection improved when explanations captured a meaningful share of predictive signal rather than only top-ranked drivers. The study’s results had clarified that compliance assurance in cloud BI was not determined by a single explanation property, but by an interlocking structure in which stable and faithful explanations created the conditions for traceability, reproducibility, fairness validation, and alignment with internal governance thresholds (Stentoft et al., 2021). In this way, the findings had extended earlier work by quantifying the specific EQ dimensions most strongly associated with compliance assurance under cloud-realistic BI tasks.

Figure 12: explain ability Governance Effects in Cloud BI



Secure Decision-Making (SDM) outcomes had also been strongly shaped by EQ, with stability again demonstrating the highest predictive contribution and fidelity and human agreement adding significant explanatory weight. This alignment with earlier adversarial and trustworthy-AI research had been clear, since prior studies had repeatedly shown that brittle explanations signaled fragile reasoning structures that were easier to exploit through evasion or poisoning (De Mauro et al., 2018). Here, higher stability had coincided with stronger robust decision integrity and higher adversarial

detection rates, meaning that stable rationale patterns had acted as indirect indicators of robust model behavior. Earlier work in explainable cybersecurity had reported that drift in attributions often preceded visible output failure, and this study had supported that logic by linking stability to security performance even when predictive accuracy was controlled implicitly through shared task inputs. Fidelity's positive effect on SDM had echoed prior evidence that low-fidelity explanations created misleading comfort, masking manipulated reasoning in high-stakes contexts (Kathuria et al., 2018). Human agreement had improved SDM in this study, reinforcing earlier findings that explanations enhanced calibrated reliance and supported correct overrides when models erred. At the same time, sparsity or complexity had shown a small negative relationship with SDM, indicating that as explanations became denser or more cognitively heavy, security reliability declined. This negative effect had been consistent with earlier warnings that excessive explanation detail increased leakage risk and allowed attackers to infer decision rules more efficiently. The present results had therefore strengthened the dual-use framing of XAI in cloud BI: explanations had served as security validators when stable and interpretable, but they had also operated as security liabilities when presented with uncontrolled complexity (Seebacher, 2021b). The study had advanced earlier research by quantifying this trade-off directly within BI tasks rather than treating security value and leakage risk as separate themes.

The mediation analysis had added an integrative layer to interpretation by showing that EQ had influenced SDM partly through CA, rather than only through a direct path. The indirect effect had been statistically significant and the direct EQ-to-SDM coefficient had reduced after CA entered the model, indicating partial mediation (Hu et al., 2019). This structure had aligned with earlier governance frameworks in enterprise AI that conceptualized compliance processes as security enablers. Prior research had argued that audit traceability, reproducibility controls, and policy enforcement reduced exposure to stealth failures and adversarial manipulation by making pipelines verifiable and by forcing stable decision evidence. The present findings had offered quantitative confirmation: high explanation quality had strengthened compliance assurance, and stronger compliance assurance had in turn supported safer decisions (Lnenicka & Komarkova, 2019). This pathway had been especially relevant in cloud BI because compliance controls required logging and versioning that also protected security integrity. Earlier studies had tended to analyze these domains in parallel, often evaluating compliance through fairness and audit readiness and security through robustness and attack detection, but this study had empirically linked them through a coherent mediation channel. The findings had demonstrated that explain ability functioned as the root governance mechanism, compliance acted as a reinforcing institutional layer, and security performance benefited both directly from explanation quality and indirectly from compliance strengthening (Cozzoli et al., 2022). This integrated pattern had provided a more unified account of how XAI operated in cloud BI than much of the earlier fragmented evidence, which had rarely modeled all three constructs within one quantitative structure.

Task-wise differences had further refined interpretation by showing that explanation demand was not uniform across BI contexts. Fraud or risk scoring and anomaly detection had produced lower baseline human agreement and higher sensitivity to explanation instability, whereas demand forecasting and resource optimization had displayed comparatively smoother explanation–decision alignment (Sahoo, 2022). This pattern had echoed earlier BI literature describing risk and anomaly domains as “high-uncertainty” contexts where decision owners required case-level rationales and rapid validation to avoid false positives and costly misclassification. Forecasting and optimization tasks, by contrast, often relied more on aggregated reasoning and scenario plausibility, allowing global explanation patterns to support decision acceptance. The study's descriptive comparisons had shown that hybrid XAI maintained high EQ across all tasks, while post-hoc methods dropped more notably in stability and completeness in anomaly-heavy contexts. Earlier interpretability studies had likewise noted that approximation-based explainers were sensitive to noisy neighborhoods, and the present task results had confirmed that cloud BI tasks with volatile or sparse anomaly patterns accentuated those weaknesses (Tim et al., 2020). This study had therefore reinforced earlier claims that XAI selection should be task-sensitive, not only model-sensitive. The evidence had shown that explanation quality mattered across all tasks, but the specific explanation properties driving usefulness shifted with decision context, with stability and fidelity becoming especially critical in domains where false alarms

or missed detections carried direct operational risk (Madureira et al., 2021).

Role-based patterns had also been consistent with earlier organizational findings on explain ability dependence. Correlation and regression summaries had shown that governance-focused users—compliance officers and security analysts—had displayed stronger dependence on EQ for CA and SDM than managerial decision owners (Hasan et al., 2022). This dependence had been aligned with earlier studies that described compliance and security roles as “verification-intensive,” meaning their work required defensible reasoning artifacts rather than only actionable scores. In this study, higher EQ had corresponded to stronger compliance assurance gains within those groups, supporting the idea that explanation logs, stable driver patterns, and cognitively manageable rationales were particularly valuable to stakeholders responsible for audit, fairness review, and threat monitoring (Shao et al., 2021). Managerial users had still benefited from explanations through improved agreement and override accuracy, but their reliance had been weaker, mirroring prior evidence that managers often focused on directional insight and outcome plausibility rather than full reasoning reconstruction. The cross-role consistency in direction, however, had reinforced that EQ operated as a shared governance asset even if the magnitude differed (Lokshina & Lanting, 2018). This study had therefore extended earlier evidence by quantifying role moderation effects within a unified cloud BI environment, confirming that explanation quality was both a compliance instrument and a practical decision aid for distributed BI users.

Finally, the combined results had portrayed hybrid XAI as the most governance-effective approach under cloud BI realism, without implying that intrinsic or post-hoc families lacked value. Hybrid models had produced the highest EQ means and the strongest CA and SDM outcomes, aligning with earlier enterprise research that treated hybrid architectures as the practical compromise between performance and interpretability (Maheshwari et al., 2021). Intrinsic models had remained strong on explanation stability and sparsity, indicating their suitability where compliance clarity and cognitive manageability were prioritized, a conclusion that matched prior work in regulated decision systems. Post-hoc methods had remained useful where high-capacity black-box predictors were necessary, but their larger stability variance and higher leak risk had echoed earlier caution that post-hoc explanations required strict fidelity monitoring and access governance (Hair Jr et al., 2019). Across the quantitative framework, stability had repeatedly emerged as a central explanatory lever, supporting earlier arguments that stability functioned as the operational backbone of explanation trust in dynamic environments. Fidelity had followed as a necessary truthfulness condition, and sparsity had acted as a usability enhancer that could become counterproductive when explanations became too dense (Sarker et al., 2018). The study had therefore integrated earlier strands into one empirical narrative: explain ability quality served as the foundational driver, compliance assurance acted as the institutional reinforcement, and secure decision-making reflected the combined effect of truthful, stable, and manageable rationales within real cloud BI tasks.

CONCLUSION

Explainable AI (XAI) models for cloud-based business intelligence had been positioned as a governance-critical layer that transformed predictive and prescriptive analytics into defensible and secure organizational evidence. In cloud BI ecosystems, machine learning outputs such as forecasts, risk scores, anomaly alerts, and optimization recommendations had been delivered continuously through dashboards to distributed stakeholders, and this delivery structure had raised the stakes of opacity because decisions were no longer interpreted only by technical specialists. XAI had therefore been framed as decision justification capability rather than decorative transparency, meaning that each BI output required an accompanying rationale that clarified driving variables, directional influence, and case-level plausibility. The quantitative evidence from this study had shown that explanation quality—expressed through fidelity, stability, sparsity or complexity, completeness, and human agreement—had operated as a measurable construct that shaped compliance assurance and secure decision-making outcomes across BI tasks. Higher-fidelity explanations had aligned reasoning with model behavior, enabling auditors and decision owners to reconstruct why a given KPI or recommendation emerged. Higher stability had ensured that explanations remained consistent under small input noise and across retraining cycles, a property that had proven especially important in cloud BI where models refreshed frequently and inference pipelines ran under fluctuating workloads.

Sparsity had strengthened cognitive manageability, supporting faster and more accurate human validation, while completeness had captured sufficient predictive signal to avoid misleading omission of key drivers. Human agreement had reflected whether stakeholders interpreted explanations consistently and used them to accept valid outputs or override flawed ones. These EQ dimensions had jointly predicted stronger Compliance Assurance, including higher audit traceability, improved reproducibility of decisions and rationales, reduced fairness deviation across protected and non-protected groups, and tighter alignment with governance thresholds. EQ had also predicted stronger Secure Decision-Making through higher robustness under perturbation, improved detection of adversarial manipulation using rationale drift cues, lower unsafe override patterns, and clearer exposure control where explanation detail risked disclosing sensitive business logic. A mediated structure had indicated that EQ had influenced secure outcomes partly through compliance strengthening, since traceability and reproducibility controls had functioned as institutional safeguards that reinforced security integrity in dynamic cloud pipelines. Comparative patterns had shown hybrid XAI architectures producing the most consistent governance outcomes by preserving black-box predictive strength while embedding structured explanation services, intrinsic models providing especially stable and cognitively light rationales suitable for stringent compliance contexts, and post-hoc methods requiring tighter monitoring due to higher volatility and leakage exposure. Overall, XAI in cloud BI had emerged as an empirically validated mechanism that linked model reasoning to lawful accountability and operational safety, ensuring that automated BI decisions remained transparent enough to defend, stable enough to trust, and controlled enough to secure within high-volume, multi-stakeholder cloud environments.

RECOMMENDATIONS

Recommendations for implementing Explainable AI (XAI) models in cloud-based business intelligence (BI) systems should focus on measurable governance performance, role-appropriate explanation delivery, and security-aware transparency control. First, organizations should treat explain ability quality as a core selection criterion alongside accuracy, latency, and cost. Model procurement and deployment processes should require benchmarked evidence of explanation fidelity and stability under realistic cloud workloads, because explanations that deviate from true model logic or fluctuate across minor input noise degrade audit traceability and decision reproducibility. Second, hybrid XAI architectures are recommended as the default for most cloud BI tasks, since they preserve high predictive power while enabling structured reasoning artifacts. However, intrinsic explainable models should be prioritized in BI domains where compliance defensibility is the overriding constraint, such as regulated risk approvals or workforce analytics, because their reasoning is inherently traceable and less fragile. Post-hoc explainers should be used selectively for complex tasks that demand black-box accuracy, but only when continuous fidelity and drift monitoring is established, ensuring that the explanation layer remains aligned with the predictive core. Third, explanation design should be layered by stakeholder role. Decision owners and executives benefit from sparse, top-driver summaries to reduce cognitive load, while compliance and security teams need deeper local rationales, reproducibility logs, and global behavior summaries for periodic reviews. A single explanation interface for all roles creates either overload for managers or under-specification for auditors. Fourth, governance must formalize explanations as decision artifacts. Every AI-generated KPI or recommendation should be logged with its explanation, model version, and data lineage in an access-controlled audit trail. These logs should support replay of historical decisions, enabling reproducibility verification and dispute resolution. Fifth, security-aware transparency must be enforced. Explanations should follow a minimum-necessary principle: reveal enough reasoning to validate decisions, but not so much that sensitive features, proprietary business logic, or training patterns become inferable. This requires explanation redaction for high-sensitivity dashboards, tiered access, rate-limited explanation APIs, and periodic leakage-risk testing. Sixth, cloud BI teams should institutionalize explanation drift dashboards that monitor stability over time across tasks and data segments. Unexpected shifts in driver patterns should trigger investigation for concept drift, pipeline compromise, or fairness degradation before operational harm spreads. Finally, XAI adoption should be embedded into BI lifecycle management rather than treated as a one-time deployment. Retraining schedules, feature updates, and dashboard changes should be paired with re-validation of explain ability metrics, fairness deviation

thresholds, and adversarial robustness tests. By grounding XAI decisions in measurable explanation quality, aligning rationales with stakeholder needs, and controlling transparency as a security surface, cloud BI systems can sustain compliant, trustworthy, and secure AI-driven decision operations at enterprise scale.

LIMITATIONS

Several limitations had shaped the scope and interpretation of this study on Explainable AI (XAI) models for cloud-based business intelligence, compliance assurance, and secure decision-making. First, the quantitative design had relied on a simulated cloud BI sandbox rather than a fully live enterprise deployment. Although the sandbox had been structured to mirror realistic ingestion, lake house storage, ML serving, and dashboard workflows, it could not capture every operational complexity found in production environments, such as irregular data outages, undocumented human workarounds, or cross-platform integration conflicts. This meant that some explain ability behaviors—especially long-term stability and drift under real organizational turbulence—might have been underestimated or simplified. Second, the BI tasks had been standardized into four representative scenarios, yet enterprise BI ecosystems often include additional tasks such as customer segmentation, dynamic pricing, or regulatory stress testing. The selection had strengthened internal control but may have limited generalizability beyond forecasting, risk scoring, anomaly detection, and optimization decisions. Third, the study had operationalized explain ability Quality (EQ), Compliance Assurance (CA), and Secure Decision-Making (SDM) using established indicator families, but any indicator set introduces measurement boundaries. Fidelity, stability, sparsity/complexity, completeness, and human agreement captured major explain ability dimensions, yet other relevant qualities—such as explanation causality, narrative coherence, or contextual relevance for specific industries—had not been included. Likewise, compliance assurance had focused on traceability, reproducibility, fairness deviation, and policy alignment, but sector-specific compliance checks vary widely and may require additional measurable elements depending on jurisdiction or regulatory regime. Security outcomes had been quantified through robustness, adversarial detection, leakage risk, and secure override accuracy, but security in cloud BI is multi-layered, and some threats—such as insider misuse, supply-chain compromise, or federated model contamination—had not been directly measured. Fourth, the participant sample, while role-stratified, had been drawn from a limited organizational and cultural context. Stakeholder reasoning styles, trust calibration tendencies, and audit practices can vary across industries and regions, so explanation usability and agreement levels could differ in other enterprise populations. Fifth, the study had treated model family as a categorical control (intrinsic, post-hoc, hybrid), but each family contains diverse algorithms and explanation tools. The findings therefore reflected family-level patterns rather than guaranteeing identical performance for every specific method within those families. Finally, the cross-sectional experimental design had evaluated explanation effects within controlled task sessions; it had not observed how explain ability might influence adoption, compliance routines, or security oversight over extended time horizons. Continuous learning effects, organizational norm formation, and long-run governance adaptation may require longitudinal evidence. These limitations suggested that while the findings had provided strong quantitative support for explain ability as a governance mechanism in cloud BI, caution was necessary when extending results to all industries, all cloud configurations, or all XAI techniques without additional real-world and long-duration validation.

REFERENCES

- [1]. Abdelwahab, Y., Kholief, M., & Sedky, A. A. H. (2022). Justifying arabic text sentiment analysis using explainable ai (xai): Lasik surgeries case study. *Information*, 13(11), 536.
- [2]. Abdulla, M., & Md. Jobayer Ibne, S. (2021). Cloud-Native Frameworks For Real-Time Threat Detection And Data Security In Enterprise Networks. *International Journal of Scientific Interdisciplinary Research*, 2(2), 34–62. <https://doi.org/10.63125/0t27av85>
- [3]. Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE access*, 6, 52138–52160.
- [4]. Ahmad, S., Miskon, S., Alabdan, R., & Tlili, I. (2020). Towards sustainable textile and apparel industry: Exploring the role of business intelligence systems in the era of industry 4.0. *Sustainability*, 12(7), 2632.
- [5]. Ahmed, I., Jeon, G., & Piccialli, F. (2022). From artificial intelligence to explainable artificial intelligence in industry 4.0: a survey on what, how, and where. *IEEE transactions on industrial informatics*, 18(8), 5031–5042.

- [6]. Ajah, I. A., & Nweke, H. F. (2019). Big data and business analytics: Trends, platforms, success factors and applications. *Big data and cognitive computing*, 3(2), 32.
- [7]. Al-Aqrabi, H., & Hill, R. (2018). Dynamic multiparty authentication of data analytics services within cloud environments. 2018 IEEE 20th International Conference on High Performance Computing and Communications; IEEE 16th International Conference on Smart City; IEEE 4th International Conference on Data Science and Systems (HPCC/SmartCity/DSS),
- [8]. Ali, S. I., Habib, F., Ali, A., Ali, A., Khan, M. F., & Jamal, A. (2021). Integrating social media and warranty data for fault identification in the cyber ecosystem: A cloud-based collaborative framework. In *Strategy, Leadership, and AI in the Cyber Ecosystem* (pp. 41-70). Elsevier.
- [9]. Ansari, M., & Alam, M. (2022). Exploring the Role of Business Intelligence Systems in IoT-Cloud Environment. International Conference on Data Analytics in Business and Marketing,
- [10]. Arfan, U., Tahsina, A., Md Mostafizur, R., & Md, W. (2023). Impact Of GFMIS-Driven Financial Transparency On Strategic Marketing Decisions In Government Agencies. *Review of Applied Science and Technology*, 2(01), 85-112. <https://doi.org/10.63125/8nqhhm56>
- [11]. Baars, H., Tank, A., Weber, P., Kemper, H.-G., Lasi, H., & Pedell, B. (2021). Cooperative approaches to data sharing and analysis for industrial internet of things ecosystems. *Applied Sciences*, 11(16), 7547.
- [12]. Bach, B., Freeman, E., Abdul-Rahman, A., Turkay, C., Khan, S., Fan, Y., & Chen, M. (2022). Dashboard design patterns. *IEEE transactions on visualization and computer graphics*, 29(1), 342-352.
- [13]. Baidya, A., & Hallur, G. G. (2022). Competitive Landscape of IT Industry in the 5G ecosystem: A management decision case study of AMDOCS. 2022 International Conference on Decision Aid Sciences and Applications (DASA),
- [14]. Banerjee, P., & Barnwal, R. P. (2022). Methods and metrics for explaining artificial intelligence models: A review. *Explainable AI: Foundations, Methodologies and Applications*, 61-88.
- [15]. Barnard, P., Macaluso, I., Marchetti, N., & DaSilva, L. A. (2022). Resource reservation in sliced networks: An explainable artificial intelligence (XAI) approach. ICC 2022-IEEE international conference on communications,
- [16]. Borrego-Díaz, J., & Galán-Páez, J. (2022). Explainable artificial intelligence in data science: From foundational issues towards socio-technical considerations. *Minds and Machines*, 32(3), 485-531.
- [17]. Borrego-Díaz, J., & Galán Páez, J. (2022). Knowledge representation for explainable artificial intelligence: Modeling foundations from complex systems. *Complex & Intelligent Systems*, 8(2), 1579-1601.
- [18]. Brady, S., Lee, N., Gibbons, K., & Bogossian, F. (2019). Woman-centred care: an integrative review of the empirical literature. *International journal of nursing studies*, 94, 107-119.
- [19]. Cali, U., Kuzlu, M., Pipattanasomporn, M., Kempf, J., & Bai, L. (2021). Foundations of big data, machine learning, and artificial intelligence and explainable artificial intelligence. In *Digitalization of Power Markets and Systems Using Energy Informatics* (pp. 115-137). Springer.
- [20]. Campbell, M., Katikireddi, S. V., Sowden, A., & Thomson, H. (2019). Lack of transparency in reporting narrative synthesis of quantitative data: a methodological assessment of systematic reviews. *Journal of clinical epidemiology*, 105, 1-9.
- [21]. Cangemi, M. P., & Taylor, P. (2018). Harnessing artificial intelligence to deliver real-time intelligence and business process improvements. *Edpacs*, 57(4), 1-6.
- [22]. Combi, C., Amico, B., Bellazzi, R., Holzinger, A., Moore, J. H., Zitnik, M., & Holmes, J. H. (2022). A manifesto on explainability for artificial intelligence in medicine. *Artificial Intelligence in Medicine*, 133, 102423.
- [23]. Coroama, L., & Groza, A. (2022). Evaluation metrics in explainable artificial intelligence (XAI). International conference on advanced research in technologies, information, innovation and sustainability,
- [24]. Cozzoli, N., Salvatore, F. P., Faccilongo, N., & Milone, M. (2022). How can big data analytics be used for healthcare organization management? Literary framework and future research from a systematic review. *BMC health services research*, 22(1), 809.
- [25]. Cutumisu, M., Adams, C., & Lu, C. (2019). A scoping review of empirical research on recent computational thinking assessments. *Journal of Science Education and Technology*, 28(6), 651-676.
- [26]. Danny, J., Wang, G., & Alianto, H. (2018). The application of Zachman framework in improving better decision making. 2018 Indonesian Association for Pattern Recognition International Conference (INAPR),
- [27]. Dawood, B. A., Al-Turjman, F., & Nawaz, M. H. (2020). Cloud computing and business intelligence in IoT-enabled smart and healthy cities. In *AI-Powered IoT for COVID-19* (pp. 1-38). CRC Press.
- [28]. de Laat, P. B. (2021). Companies committed to responsible AI: From principles towards implementation and regulation? *Philosophy & technology*, 34(4), 1135-1193.
- [29]. De Mauro, A., Greco, M., Grimaldi, M., & Ritala, P. (2018). Human resources for Big Data professions: A systematic classification of job roles and required skill sets. *Information Processing & Management*, 54(5), 807-817.
- [30]. Donadello, I., & Dragoni, M. (2020). SeXAI: a semantic explainable artificial intelligence framework. International Conference of the Italian Association for Artificial Intelligence,
- [31]. Ehwerhemuepha, L., Gasperino, G., Bischoff, N., Taraman, S., Chang, A., & Feaster, W. (2020). HealthDataLab—a cloud computing solution for data science and advanced analytics in healthcare with application to predicting multi-center pediatric readmissions. *BMC Medical Informatics and Decision Making*, 20(1), 115.
- [32]. El Ghalbzouri, H., & El Bouhdidi, J. (2021). Integrating business intelligence with cloud computing: State of the art and fundamental concepts. *Networking, Intelligent Systems and Security: Proceedings of NISS 2021*, 197-213.

- [33]. Ferdous Ara, A. (2021). Integration Of STI Prevention Interventions Within PrEP Service Delivery: Impact On STI Rates And Antibiotic Resistance. *International Journal of Scientific Interdisciplinary Research*, 2(2), 63–97. <https://doi.org/10.63125/65143m72>
- [34]. Ferdous Ara, A., & Beatrice Onyinyechi, M. (2023). Long-Term Epidemiologic Trends Of STIs PRE- and POST-PrEP Introduction: A National Time-Series Analysis. *American Journal of Health and Medical Sciences*, 4(02), 01–35. <https://doi.org/10.63125/mp153d97>
- [35]. Firouzi, F., Farahani, B., Barzegari, M., & Daneshmand, M. (2020). AI-driven data monetization: The other face of data in IoT-based smart and connected health. *IEEE Internet of Things Journal*, 9(8), 5581-5599.
- [36]. Flammini, F., Alcaraz, C., Bellini, E., Marrone, S., Lopez, J., & Bondavalli, A. (2022). Towards trustworthy autonomous systems: Taxonomies and future perspectives. *IEEE Transactions on Emerging Topics in Computing*, 12(2), 601-614.
- [37]. Fouladgar, N., Alirezaie, M., & Främling, K. (2022). Metrics and evaluations of time series explanations: An application in affect computing. *IEEE access*, 10, 23995-24009.
- [38]. Gbededo, M. A., Liyanage, K., & Garza-Reyes, J. A. (2018). Towards a Life Cycle Sustainability Analysis: A systematic review of approaches to sustainable manufacturing. *Journal of Cleaner Production*, 184, 1002-1015.
- [39]. Gerlach, J., Hoppe, P., Jagels, S., Licker, L., & Breiter, M. H. (2022). Decision support for efficient XAI services-A morphological analysis, business model archetypes, and a decision tree. *Electronic Markets*, 32(4), 2139-2158.
- [40]. Ghita, M., Siham, B., Hicham, M., Abdelhafid, A. E. M., & Laurent, D. (2020). Geospatial business intelligence and cloud services for context aware digital twins development. 2020 IEEE International conference of Moroccan Geomatics (Morgeo),
- [41]. Guleria, P., Naga Srinivasu, P., Ahmed, S., Almusallam, N., & Alarfaj, F. K. (2022). XAI framework for cardiovascular disease prediction using classification techniques. *Electronics*, 11(24), 4086.
- [42]. Habibullah, S. M., & Md. Foysal, H. (2021). A Data Driven Cyber Physical Framework For Real Time Production Control Integrating IOT And Lean Principles. *American Journal of Interdisciplinary Studies*, 2(03), 35–70. <https://doi.org/10.63125/20nhqs87>
- [43]. Hagras, H. (2018). Toward human-understandable, explainable AI. *Computer*, 51(9), 28-36.
- [44]. Hair Jr, J., Page, M., & Brunsveld, N. (2019). *Essentials of business research methods*. Routledge.
- [45]. Hasan, N., Bao, Y., & Miah, S. J. (2022). Exploring the impact of ICT usage among indigenous people and their quality of life: Operationalizing Sen’s capability approach. *Information Technology for Development*, 28(2), 230-250.
- [46]. He, W., Shong, J. Y. L., & Wang, C. (2022). AI-driven BIM on the cloud. In *Artificial Intelligence in Urban Planning and Design* (pp. 101-117). Elsevier.
- [47]. Hechler, E., Oberhofer, M., & Schaeck, T. (2020). AI and Governance. In *Deploying AI in the Enterprise: IT Approaches for Design, DevOps, Governance, Change Management, Blockchain, and Quantum Computing* (pp. 165-211). Springer.
- [48]. Hennink, M., & Kaiser, B. N. (2022). Sample sizes for saturation in qualitative research: A systematic review of empirical tests. *Social science & medicine*, 292, 114523.
- [49]. Höhn, S., & Faradouris, N. (2021). What does it cost to deploy an XAI system: A case study in legacy systems. International Workshop on Explainable, Transparent Autonomous Agents and Multi-Agent Systems,
- [50]. Holzinger, A., Malle, B., Saranti, A., & Pfeifer, B. (2021). Towards multi-modal causability with graph neural networks enabling information fusion for explainable AI. *Information Fusion*, 71, 28-37.
- [51]. Holzinger, A., Saranti, A., Molnar, C., Biecek, P., & Samek, W. (2020). Explainable AI methods-a brief overview. International workshop on extending explainable AI beyond deep models and classifiers,
- [52]. Hu, Y., Xu, A., Hong, Y., Gal, D., Sinha, V., & Akkiraju, R. (2019). Generating business intelligence through social media analytics: Measuring brand personality with consumer-, employee-, and firm-generated content. *Journal of management information systems*, 36(3), 893-930.
- [53]. Hussain, S. M., Buongiorno, D., Altini, N., Berloco, F., Prencipe, B., Moschetta, M., Bevilacqua, V., & Brunetti, A. (2022). Shape-based breast lesion classification using digital tomosynthesis images: The role of explainable artificial intelligence. *Applied Sciences*, 12(12), 6230.
- [54]. Identity, K. O. S. (2020). Access Management. In: Springer.
- [55]. Iqbal, R., Doctor, F., More, B., Mahmud, S., & Yousuf, U. (2020). Big data analytics: Computational intelligence techniques and application areas. *Technological Forecasting and Social Change*, 153, 119253.
- [56]. Islam, M. R., Ahmed, M. U., Barua, S., & Begum, S. (2022). A systematic review of explainable artificial intelligence in terms of different application domains and tasks. *Applied Sciences*, 12(3), 1353.
- [57]. Janev, V. (2020). Chapter 1 Ecosystem of Big Data. In *Knowledge graphs and big data processing* (pp. 3-19). Springer.
- [58]. Kathuria, A., Mann, A., Khuntia, J., Saldanha, T. J., & Kauffman, R. J. (2018). A strategic value appropriation path for cloud computing. *Journal of management information systems*, 35(3), 740-775.
- [59]. Khorshidi, H. A., & Aickelin, U. (2021). Multicriteria group decision-making under uncertainty using interval data and cloud models. *Journal of the Operational Research Society*, 72(11), 2542-2556.
- [60]. Khrais, L. T. (2020). Role of artificial intelligence in shaping consumer demand in E-commerce. *Future Internet*, 12(12), 226.
- [61]. Kim, M.-Y., Atakishiyev, S., Babiker, H. K. B., Farruque, N., Goebel, R., Zaïane, O. R., Motallebi, M.-H., Rabelo, J., Syed, T., & Yao, H. (2021). A multi-component framework for the analysis and design of explainable artificial intelligence. *Machine Learning and Knowledge Extraction*, 3(4), 900-921.
- [62]. Knapič, S., Malhi, A., Saluja, R., & Främling, K. (2021). Explainable artificial intelligence for human decision support system in the medical domain. *Machine Learning and Knowledge Extraction*, 3(3), 740-770.

- [63]. Konanahalli, A., Marinelli, M., & Oyedele, L. (2020). Drivers and challenges associated with the implementation of big data within UK facilities management sector: An exploratory factor analysis approach. *IEEE Transactions on Engineering Management*, 69(4), 916-929.
- [64]. Kumar, D., & Mehta, M. A. (2022). An overview of explainable AI methods, forms and frameworks. *Explainable AI: Foundations, Methodologies and Applications*, 43-59.
- [65]. Kuppa, A., & Le-Khac, N.-A. (2021). Adversarial XAI methods in cybersecurity. *IEEE transactions on information forensics and security*, 16, 4924-4938.
- [66]. Lahane, S., Kant, R., & Shankar, R. (2020). Circular supply chain management: A state-of-art review and future opportunities. *Journal of Cleaner Production*, 258, 120859.
- [67]. Le, T.-T.-H., Kim, H., Kang, H., & Kim, H. (2022). Classification and explanation for intrusion detection system based on ensemble trees and SHAP method. *Sensors*, 22(3), 1154.
- [68]. Li, C., Guo, W., Sun, S. C., Al-Rubaye, S., & Tsourdos, A. (2020). Trustworthy deep learning in 6G-enabled mass autonomy: From concept to quality-of-trust key performance indicators. *IEEE Vehicular Technology Magazine*, 15(4), 112-121.
- [69]. Linardatos, P., Papastefanopoulos, V., & Kotsiantis, S. (2020). Explainable ai: A review of machine learning interpretability methods. *Entropy*, 23(1), 18.
- [70]. Lnenicka, M., & Komarkova, J. (2019). Developing a government enterprise architecture framework to support the requirements of big and open linked data with the use of cloud computing. *International Journal of Information Management*, 46, 124-141.
- [71]. Lokshina, I., & Lanting, C. (2018). A qualitative evaluation of IoT-driven eHealth: knowledge management, business models and opportunities, deployment and evolution. In *Data-Centric Business and Applications: Evolutions in Business Information Processing and Management – Volume 1* (pp. 23-52). Springer.
- [72]. Longo, L., Goebel, R., Lecue, F., Kieseberg, P., & Holzinger, A. (2020). Explainable artificial intelligence: Concepts, applications, research challenges and visions. International cross-domain conference for machine learning and knowledge extraction,
- [73]. Lopes, P., Silva, E., Braga, C., Oliveira, T., & Rosado, L. (2022). XAI systems evaluation: A review of human and computer-centred methods. *Applied Sciences*, 12(19), 9423.
- [74]. Lötsch, J., Kringel, D., & Ultsch, A. (2021). Explainable artificial intelligence (XAI) in biomedicine: Making AI decisions trustworthy for physicians and patients. *BioMedInformatics*, 2(1), 1-17.
- [75]. Machlev, R., Heistrene, L., Perl, M., Levy, K. Y., Belikov, J., Mannor, S., & Levron, Y. (2022). Explainable Artificial Intelligence (XAI) techniques for energy and power systems: Review, challenges and opportunities. *Energy and AI*, 9, 100169.
- [76]. Machlev, R., Perl, M., Belikov, J., Levy, K. Y., & Levron, Y. (2021). Measuring explainability and trustworthiness of power quality disturbances classifiers using XAI – Explainable artificial intelligence. *IEEE transactions on industrial informatics*, 18(8), 5127-5137.
- [77]. Madureira, L., Popovič, A., & Castelli, M. (2021). Competitive intelligence empirical construct validation using expert in-depth interviews study. 2021 IEEE International Conference on Technology Management, Operations and Decisions (ICTMOD),
- [78]. Maheshwari, S., Gautam, P., & Jaggi, C. K. (2021). Role of Big Data Analytics in supply chain management: current trends and future perspectives. *International Journal of Production Research*, 59(6), 1875-1900.
- [79]. Marinho, M., Prakash, V., Garg, L., Savaglio, C., & Bawa, S. (2021). Effective cloud resource utilisation in cloud erp decision-making process for industry 4.0 in the united states. *Electronics*, 10(8), 959.
- [80]. Mazumdar, S., Seybold, D., Kritikos, K., & Verginadis, Y. (2019). A survey on data storage and placement methodologies for cloud-big data ecosystem. *Journal of Big Data*, 6(1), 1-37.
- [81]. Md Al Amin, K. (2022). Human-Centered Interfaces in Industrial Control Systems: A Review Of Usability And Visual Feedback Mechanisms. *Review of Applied Science and Technology*, 1(04), 66-97.
<https://doi.org/10.63125/gr54qy93>
- [82]. Md Ariful, I. (2022). Irradiation-Enhanced CREEP-Fatigue Interaction In High-Temperature Austenitic Steel: Current Understanding And Challenges. *American Journal of Advanced Technology and Engineering Solutions*, 2(04), 148-181. <https://doi.org/10.63125/e46gja61>
- [83]. Md Ariful, I., & Efat Ara, H. (2022). Advances And Limitations Of Fracture Mechanics-Based Fatigue Life Prediction Approaches For Structural Integrity Assessment: A Systematic Review. *American Journal of Interdisciplinary Studies*, 3(03), 68-98. <https://doi.org/10.63125/fg8ae957>
- [84]. Md Nahid, H. (2022). Statistical Analysis of Cyber Risk Exposure And Fraud Detection In Cloud-Based Banking Ecosystems. *ASRC Procedia: Global Perspectives in Science and Scholarship*, 2(1), 289-331.
<https://doi.org/10.63125/9wf91068>
- [85]. Md Sarwar, H. (2021). Sustainable Materials Characterization For Low-Carbon Construction And Infrastructure Durability. *American Journal of Interdisciplinary Studies*, 2(01), 01-34. <https://doi.org/10.63125/wq1wdr64>
- [86]. Md Sarwar Hossain, S., & Md Milton, M. (2022). Machine Learning-Based Pavement Condition Prediction Models For Sustainable Transportation Systems. *American Journal of Interdisciplinary Studies*, 3(01), 31-64.
<https://doi.org/10.63125/ljsmkg92>
- [87]. Md. Mominul, H., Masud, R., & Md. Milton, M. (2022). Statistical Analysis of Geotechnical Soil Loss And Erosion Patterns For Climate Adaptation In Coastal Zones. *American Journal of Interdisciplinary Studies*, 3(03), 36-67.
<https://doi.org/10.63125/xytn3e23>

- [88]. Md. Musfiqur, R., & Saba, A. (2021). Data-Driven Decision Support in Information Systems: Strategic Applications In Enterprises. *International Journal of Scientific Interdisciplinary Research*, 2(2), 01-33. <https://doi.org/10.63125/cfvg2v45>
- [89]. Md. Redwanul, I., Md Nahid, H., & Md. Zahid Hasan, T. (2021). Predictive Analytics in Supply Chain Management A Review Of Business Analyst-Led Optimization Tools. *Review of Applied Science and Technology*, 6(1), 34-73. <https://doi.org/10.63125/5aypx555>
- [90]. Mehdiyev, N., Houy, C., Gutermuth, O., Mayer, L., & Fettke, P. (2021). Explainable artificial intelligence (XAI) supporting public administration processes—on the potential of XAI in tax audit processes. *International Conference on Wirtschaftsinformatik*,
- [91]. Meinzen-Dick, R., Quisumbing, A., Doss, C., & Theis, S. (2019). Women's land rights as a pathway to poverty reduction: Framework and review of available evidence. *Agricultural systems*, 172, 72-82.
- [92]. Merry, M., Riddle, P., & Warren, J. (2021). A mental models approach for defining explainable artificial intelligence. *BMC Medical Informatics and Decision Making*, 21(1), 344.
- [93]. Minh, D., Wang, H. X., Li, Y. F., & Nguyen, T. N. (2022). Explainable artificial intelligence: a comprehensive review. *Artificial Intelligence Review*, 55(5), 3503-3568.
- [94]. Mohammad Mushfequr, R., & Ashraful, I. (2023). Automation And Risk Mitigation in Healthcare Claims: Policy And Compliance Implications. *Review of Applied Science and Technology*, 2(04), 124-157. <https://doi.org/10.63125/v73gyg14>
- [95]. Mohanrajan, S. N., & Loganathan, A. (2022). Novel vision transformer-based bi-LSTM model for LU/LC prediction—Javadi Hills, India. *Applied Sciences*, 12(13), 6387.
- [96]. Monks, T., Currie, C. S., Onggo, B. S., Robinson, S., Kunc, M., & Taylor, S. J. (2019). Strengthening the reporting of empirical simulation studies: Introducing the STRESS guidelines. *Journal of Simulation*, 13(1), 55-67.
- [97]. Morad, S., Ragonis, N., & Barak, M. (2021). An integrative conceptual model of innovation and innovative thinking based on a synthesis of a literature review. *Thinking skills and creativity*, 40, 100824.
- [98]. Mortuza, M. M. G., & Rauf, M. A. (2022). Industry 4.0: An Empirical Analysis of Sustainable Business Performance Model Of Bangladeshi Electronic Organisations. *International Journal of Economy and Innovation*. https://gospodarkainnowacje.pl/index.php/issue_view_32/article/view/826
- [99]. Moyo, M., & Loock, M. (2021). Conceptualising a cloud business intelligence security evaluation framework for small and medium enterprises in small towns of the Limpopo Province, South Africa. *Information*, 12(3), 128.
- [100]. Mst. Shahrin, S., & Samia, A. (2023). High-Performance Computing For Scaling Large-Scale Language And Data Models In Enterprise Applications. *ASRC Procedia: Global Perspectives in Science and Scholarship*, 3(1), 94-131. <https://doi.org/10.63125/e7yfwm87>
- [101]. Nazar, M., Alam, M. M., Yafi, E., & Su'ud, M. M. (2021). A systematic review of human-computer interaction and explainable artificial intelligence in healthcare with artificial intelligence techniques. *IEEE access*, 9, 153316-153348.
- [102]. Nizam, T., & Zafar, S. (2022). Explainable artificial intelligence (XAI): conception, visualization and assessment approaches towards amenable XAI. In *Explainable Edge AI: A Futuristic Computing Perspective* (pp. 35-51). Springer.
- [103]. Nocker, M., & Sena, V. (2019). Big data and human resources management: The rise of talent analytics. *Social Sciences*, 8(10), 273.
- [104]. O. Chergykalo, D., & Klyushin, D. A. (2022). Fundamental Fallacies in Definitions of Explainable AI: Explainable to Whom and Why? In *Explainable AI: Foundations, Methodologies and Applications* (pp. 25-42). Springer.
- [105]. O'Neill, M., & Brabazon, A. (2019). Business analytics capability, organisational value and competitive advantage. *Journal of Business Analytics*, 2(2), 160-173.
- [106]. Owens, E., Sheehan, B., Mullins, M., Cunneen, M., Ressel, J., & Castignani, G. (2022). Explainable artificial intelligence (xai) in insurance. *Risks*, 10(12), 230.
- [107]. Páez, A. (2019). The pragmatic turn in explainable artificial intelligence (XAI). *Minds and Machines*, 29(3), 441-459.
- [108]. Rakibul, H., & Samia, A. (2022). Information System-Based Decision Support Tools: A Systematic Review Of Strategic Applications In Service-Oriented Enterprises. *Review of Applied Science and Technology*, 1(04), 26-65. <https://doi.org/10.63125/w3cevv78>
- [109]. Rane, S. B., & Narvel, Y. A. M. (2022). Data-driven decision making with Blockchain-IoT integrated architecture: a project resource management agility perspective of industry 4.0. *International Journal of System Assurance Engineering and Management*, 13(2), 1005-1023.
- [110]. Rauvola, R. S., Vega, D. M., & Lavigne, K. N. (2019). Compassion fatigue, secondary traumatic stress, and vicarious traumatization: A qualitative review and research agenda. *Occupational health science*, 3(3), 297-336.
- [111]. Reza, M., Vorobyova, K., & Rauf, M. (2021). The effect of total rewards system on the performance of employees with a moderating effect of psychological empowerment and the mediation of motivation in the leather industry of Bangladesh. *Engineering Letters*, 29, 1-29.
- [112]. Rohlfing, K. J., Cimiano, P., Scharlau, L., Matzner, T., Buhl, H. M., Buschmeier, H., Esposito, E., Grimminger, A., Hammer, B., & Häb-Umbach, R. (2020). Explanation as a social practice: Toward a conceptual framework for the social design of AI systems. *IEEE Transactions on Cognitive and Developmental Systems*, 13(3), 717-728.
- [113]. Sahoo, S. (2022). Big data analytics in manufacturing: a bibliometric analysis of research in the field of business management. *International Journal of Production Research*, 60(22), 6793-6821.
- [114]. Saikat, S. (2021). Real-Time Fault Detection in Industrial Assets Using Advanced Vibration Dynamics And Stress Analysis Modeling. *American Journal of Interdisciplinary Studies*, 2(04), 39-68. <https://doi.org/10.63125/0h163429>
- [115]. Saikat, S. (2022). CFD-Based Investigation of Heat Transfer Efficiency In Renewable Energy Systems. *International Journal of Scientific Interdisciplinary Research*, 1(01), 129-162. <https://doi.org/10.63125/ttw40456>

- [116]. Saini, V. K., Gupta, S., & Gupta, B. (2022). Data security in collaborative business intelligence for sustainable super smart society. In *Decision Analytics for Sustainable Development in Smart Society 5.0: Issues, Challenges and Opportunities* (pp. 113-130). Springer.
- [117]. Salisu, I., Bin Mohd Sappri, M., & Bin Omar, M. F. (2021). The adoption of business intelligence systems in small and medium enterprises in the healthcare sector: A systematic literature review. *Cogent Business & Management*, 8(1), 1935663.
- [118]. Saraswat, D., Bhattacharya, P., Verma, A., Prasad, V. K., Tanwar, S., Sharma, G., Bokoro, P. N., & Sharma, R. (2022). Explainable AI for healthcare 5.0: opportunities and challenges. *IEEE access*, 10, 84486-84517.
- [119]. Sarker, M. N. I., Wu, M., & Hossin, M. A. (2018). Smart governance through bigdata: Digital transformation of public agencies. 2018 international conference on artificial intelligence and big data (ICAIBD),
- [120]. Savastano, M., Amendola, C., Bellini, F., & D'Ascenzo, F. (2019). Contextual impacts on industrial processes brought by the digital transformation of manufacturing: A systematic review. *Sustainability*, 11(3), 891.
- [121]. Seebacher, U. (2021a). The Predictive Intelligence Ecosystem. In *Predictive Intelligence for Data-Driven Managers: Process Model, Assessment-Tool, IT-Blueprint, Competence Model and Case Studies* (pp. 21-55). Springer.
- [122]. Seebacher, U. (2021b). *Predictive intelligence for data-driven managers*. Springer.
- [123]. Shaikh, S., & Aditya, D. (2021). Federated Learning-Driven Predictive Quality Analytics and Supply Chain Optimization In Distributed Manufacturing Networks. *Review of Applied Science and Technology*, 6(1), 74-107. <https://doi.org/10.63125/k18cbz55>
- [124]. Shao, X.-F., Liu, W., Li, Y., Chaudhry, H. R., & Yue, X.-G. (2021). Multistage implementation framework for smart supply chain management under industry 4.0. *Technological Forecasting and Social Change*, 162, 120354.
- [125]. Sheu, R.-K., & Pardeshi, M. S. (2022). A survey on medical explainable AI (XAI): recent progress, explainability approach, human interaction and scoring system. *Sensors*, 22(20), 8068.
- [126]. Siddiqui, K., & Doyle, T. E. (2022). Trust metrics for medical deep learning using explainable-ai ensemble for time series classification. 2022 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE),
- [127]. Sovacool, B. K., Hess, D. J., & Cantoni, R. (2021). Energy transitions from the cradle to the grave: a meta-theoretical framework integrating responsible innovation, social practices, and energy justice. *Energy Research & Social Science*, 75, 102027.
- [128]. Sovrano, F., Sapienza, S., Palmirani, M., & Vitali, F. (2022). Metrics, explainability and the European AI act proposal. *J*, 5(1), 126-138.
- [129]. Srivastava, G., S, M., Venkataraman, R., V, K., & N, P. (2022). A review of the state of the art in business intelligence software. *Enterprise Information Systems*, 16(1), 1-28.
- [130]. Stentoft, J., Aadsbøll Wickstrøm, K., Philipsen, K., & Haug, A. (2021). Drivers and barriers for Industry 4.0 readiness and practice: empirical evidence from small and medium-sized manufacturers. *Production Planning & Control*, 32(10), 811-828.
- [131]. Tim, Y., Hallikainen, P., Pan, S. L., & Tamm, T. (2020). Actualizing business analytics for organizational transformation: A case study of Rovio Entertainment. *European Journal of Operational Research*, 281(3), 642-655.
- [132]. Tjoa, E., & Guan, C. (2020). A survey on explainable artificial intelligence (xai): Toward medical xai. *IEEE transactions on neural networks and learning systems*, 32(11), 4793-4813.
- [133]. Tjoa, E., & Guan, C. (2022). Quantifying explainability of saliency methods in deep neural networks with a synthetic dataset. *IEEE Transactions on artificial Intelligence*, 4(4), 858-870.
- [134]. Tocchetti, A., & Brambilla, M. (2022). The role of human knowledge in explainable AI. *Data*, 7(7), 93.
- [135]. Tonoy Kanti, C., & Shaikat, B. (2022). Graph Neural Networks (GNNS) For Modeling Cyber Attack Patterns And Predicting System Vulnerabilities In Critical Infrastructure. *American Journal of Interdisciplinary Studies*, 3(04), 157-202. <https://doi.org/10.63125/1ykzx350>
- [136]. Toronto, C. E., & Remington, R. (2020). A step-by-step guide to conducting an integrative review.
- [137]. Vilone, G., & Longo, L. (2021). Classification of explainable artificial intelligence methods through their output formats. *Machine Learning and Knowledge Extraction*, 3(3), 615-661.
- [138]. Vrontis, D., & Christofi, M. (2021). R&D internationalization and innovation: A systematic review, integrative framework and future research directions. *Journal of Business Research*, 128, 812-823.
- [139]. Wang, Z., Wang, N., Su, X., & Ge, S. (2020). An empirical study on business analytics affordances enhancing the management of cloud computing data security. *International Journal of Information Management*, 50, 387-394.
- [140]. Wangoo, D. P. (2020). Intelligent Software Mining with Business Intelligence Tools for Automation of Micro services in SOA: A Use Case for Analytics. 2020 7th International Conference on Computing for Sustainable Global Development (INDIACom),
- [141]. Wuni, I. Y., & Shen, G. Q. (2020). Barriers to the adoption of modular integrated construction: Systematic review and meta-analysis, integrated conceptual framework, and strategies. *Journal of Cleaner Production*, 249, 119347.
- [142]. Wuni, I. Y., Shen, G. Q., & Mahmud, A. T. (2022). Critical risk factors in the application of modular integrated construction: a systematic review. *International Journal of Construction Management*, 22(2), 133-147.
- [143]. Zhang, Z., Al Hamadi, H., Damiani, E., Yeun, C. Y., & Taher, F. (2022). Explainable artificial intelligence applications in cyber security: State-of-the-art in research. *IEEE access*, 10, 93104-93139.
- [144]. Zhou, J., Gandomi, A. H., Chen, F., & Holzinger, A. (2021). Evaluating the quality of machine learning explanations: A survey on methods and metrics. *Electronics*, 10(5), 593.