

The Business Analyst as an AI-Transformation Broker: A Framework for Stakeholder Alignment, Responsible Automation, And Value Realization in Small and Medium-Sized Enterprises

¹Joseph Idanesi Alieme, ²Aumbur Sule, ³Valentina Ochuko Obukadata

¹ Georgia State University,
USA; aliemejoseph@gmail.com

²Access Bank Plc, Nigeria;
Aumbur.sule@accessbankplc.com

³ University of New Haven, Connecticut United States;
obukadataval@gmail.com

Corresponding Author; aliemejoseph@gmail.com
DOI: 10.56201/ijebm.vol.11.no8.2025.pg299.355

Abstract

Small and medium-sized enterprises adopting artificial intelligence rarely fail because a model does not work. They fail at the seams: a process is automated that nobody agreed on, a model is built on data never fit for the decision it informs, a governance checklist is satisfied without any design decision changing, and a benefit is declared that nobody measured. This paper argues that these are boundary failures that is, knowledge that did not survive transfer between groups who do not share a frame, and that they are best understood through Carlile's syntactic, semantic and pragmatic boundary hierarchy. Large firms answer boundary failure structurally, through data offices, governance functions, transformation offices and model-risk teams; a smaller firm has none of these, yet the brokering work does not disappear, it is either performed by a role or not at all. The paper argues, from role positioning rather than from empirical demonstration, that the business analyst is the role structurally placed to perform it, and develops a framework of four brokering functions consisting of translation, legitimation, constraint injection and benefit tracing, specifying for each what is moved, across which boundary, which boundary objects are produced, what competence is required, and what fails without it. Six tables set out the functions, a boundary-by-phase matrix, the substitutions smaller firms must make, a benefits dependency structure, the pathologies of brokering, and the traditions grounding each function. An applied-literature survey shows the functions surfacing across domains, and eight propositions are offered for testing.

Keywords: *business analysis; boundary spanning; intelligent process automation; small and medium-sized enterprises; AI governance; benefits realization; stakeholder salience*

1. Introduction

An automation initiative in a smaller firm rarely has a single point of failure; it has seams. The sponsor's stated outcome, the operations manager's actual process, the data available at the moment of decision, the compliance requirement that arrives after the design is fixed, and the baseline nobody took are all handovers between groups with different vocabularies, interests and accountabilities. The failure lies not in any of the groups but between them.

This is the capstone of a programme concerned with the *specification gap*, the distance between the way a business states what it wants and the precision an automated system requires to deliver it. Earlier work in this programme has argued that decision logic must be standardized before it can be automated; that readiness in a smaller firm is organisational rather than technological and must be sequenced; that ambiguous problems can be converted through structured gates into AI-ready requirements; and that different gates demand different delivery postures. All of that describes *work*. None of it says who does the work, or what happens when nobody does.

The argument makes three moves. AI transformation fails at boundaries rather than components: the recurring failure patterns are instances of knowledge failing to cross a syntactic, semantic or pragmatic boundary (Carlile, 2004). In a smaller firm that work cannot be housed in a department, because the departments large firms use will not be created; it is carried by a role or abandoned. The role best positioned to carry it is the business analyst, who sits at the intersection of business intent, data reality, delivery capability and governance obligation (International Institute of Business Analysis, 2015; Paul et al., 2020); argued from structural position, not established empirically. And brokering has failure modes of its own, since naming an analyst "the AI broker" does not create brokering competence (Levina & Vaast, 2005).

The contribution is a reframing of these failures as boundary failures, a four-function brokering framework, an account of what smaller firms lack structurally and what a broker substitutes, and eight propositions for testing. Section 2 reviews the literature, Section 3 states what kind of paper this is, and Section 4 develops the framework and its six tables. Section 5 surveys applied and practitioner writing to show how the four functions surface in particular domains, flagging explicitly that this material is illustrative rather than evidential. Section 6 works the framework through a hypothetical initiative, and Sections 7 to 9 discuss implications, limitations and the programme.

2. Background and related literature

2.1 Knowledge at boundaries

Knowledge is not simply transmitted across organisational boundaries; it is at stake in them. It is localised, embedded and invested in practice, so what one function knows is not readily available to another even where both are willing (Carlile, 2002). Carlile (2004) organised this into three boundaries with matching processes: at a *syntactic* boundary the parties share meanings and interests but lack a common lexicon, and *transfer* across a shared syntax suffices; at a *semantic* boundary meanings differ even where words are shared, and *translation* is required; at a *pragmatic* boundary interests themselves differ, and *transformation* is required, altering knowledge and interests on both sides. The hierarchy is cumulative, so an instrument pitched at the wrong level does not merely underperform; it misdiagnoses.

Star and Griesemer (1989) supplied the mechanism: boundary objects inhabit several intersecting social worlds and satisfy the informational requirements of each, being "plastic enough to adapt to local needs... yet robust enough to maintain a common identity across sites," and at pragmatic boundaries making one group's consequences for another negotiable (Carlile, 2002).

The role literature is consistent. Boundary-spanning roles were characterised early as performing two structurally distinct functions: processing information originating outside the unit, and representing the unit to those outside it, with their prominence rising as the environment becomes less certain (Aldrich & Herker, 1977). Cross-boundary communication concentrates in a few special boundary roles (Tushman, 1977); the broker produces "brokered knowledge" from a position of "double peripherality" (Meyer, 2010); and value arises from recombination rather than

speed of transmission (Hargadon & Sutton, 1997). Levina and Vaast (2005) supply the corrective, distinguishing *nominated* boundary spanners from *boundary spanners-in-practice*, whose capability emerges from practice and is recognised as legitimate by both sides, and explaining why designated liaison roles fail.

Two results qualify what such a role can carry: organisational knowledge resides in members, tools and tasks and in the networks among them, so what is embedded in their interaction moves least readily (Argote & Ingram, 2000), and knowing is enacted in practice rather than held and delivered (Orlikowski, 2002). A broker builds the conditions in which each side can enact enough of the other's.

2.2 Boundary objects in use, and how they fail

The concept has been applied widely enough that its originator intervened to correct the usage. Star (2010) objected that later applications kept the appealing half: the object that bridges, while discarding the apparatus that made it work: the interplay between a loosely structured common object and the tightly structured local versions each site produces, and the residual categories neither absorbs. An artefact is therefore not a boundary object because someone designates it as one; it becomes one where information requirements genuinely differ and overlap, and the work of moving between common and local versions is performed.

Empirical work is correspondingly cautionary. Barrett and Oborn (2010), studying cross-cultural software development teams, found that boundary objects can entrench differences as readily as bridge them, the effect depending on how they are used across occupational lines rather than on any property of the artefact. Nicolini, Mengis and Swan (2012) held "boundary object" too coarse a label for what objects do in cross-disciplinary collaboration, distinguishing motivational, epistemic and infrastructural roles performed by different objects at different moments. Bechky (2003) showed on a production floor that cross-occupational understanding was achieved by transforming meaning rather than transmitting it; Carlile's translation observed rather than postulated. And Oborn and Dawson (2010) established that the work is political: some knowledge domains are privileged over others.

2.3 Knowledge transfer, stickiness and situated practice

Beneath the boundary literature sits an older question: why knowledge held in one part of an organisation does not arrive in another. Nonaka (1994) modelled organisational knowledge creation as conversion between tacit and explicit forms, the tacit component being precisely what cannot be handed over as a document. Szulanski (1996) located the impediments empirically, finding internal stickiness driven by the recipient's absorptive capacity, by causal ambiguity about why a practice works, and by an arduous source–recipient relationship rather than, as managerial accounts assume, by unwillingness to share. Brown and Duguid (1991) added that canonical practice as documented diverges systematically from practice as performed, learning occurring through participation in communities of practice rather than through the procedure.

Brown and Duguid's community of practice received its full statement from Wenger (1998), for whom competence is negotiated through participation and reification, that is, abstractions, tools and documents whose meaning depends on the participation they accompany, and was extended to the organisation as a social learning system (Wenger, 2000). A decision table written in a workshop and read a year later by an engineer who was not there has lost half of what made it work. Absorptive capacity supplies the receiving side: the capacity to assimilate and apply external knowledge is a function of prior related knowledge (Cohen & Levinthal, 1990), and potential

absorptive capacity is separable from realised, so a firm may acquire AI knowledge without converting it into use (Zahra & George, 2002). Tacit knowledge and unknown unknowns, meanwhile, are categories of ignorance calling for different mitigations rather than more interviewing (Sutcliffe & Sawyer, 2013).

Together these results specify what translation must accomplish. Where the decisive knowledge is held tacitly, reading the documented procedure recovers the canonical account and misses the practice; where stickiness follows from causal ambiguity, a requirement stating what the system should do without stating why the rule exists transfers a conclusion without its warrant; and where absorptive capacity conditions what a receiving group can take up, the same specification is not equally usable by an engineering team with domain exposure and one without.

2.4 Stakeholders, salience, change recipients and the basis of claims

Freeman (1984) established that management is accountable to those who affect or are affected by the firm's objectives, and Mitchell, Agle and Wood (1997) made this tractable through a salience typology in which attention follows power, legitimacy and urgency. Donaldson and Preston (1995) distinguished descriptive, instrumental and normative uses of stakeholder theory, a distinction used in Section 4.3 to mark what kind of claim this paper makes. Two constraints matter: the association between user involvement and system success is mostly positive but heterogeneously defined and measured (Bano & Zowghi, 2015), and a reorganisation disregarding the social system it displaces degrades the technical result itself, an axis of joint optimisation that technology-led framings erode (Sarker et al., 2019; Trist & Bamforth, 1951).

The people whose work an automation reallocates are what the change literature calls recipients, and that literature carries two cautions. The first is evidentiary: By (2005), reviewing organisational change management critically, found a field resting on prescriptive models with a weak empirical base, the frequently repeated claim that most change programmes fail being contested rather than established, a caution observed here by asserting no failure rate. The second is substantive: Oreg, Vakola and Armenakis (2011), synthesising six decades of quantitative studies of change recipients' reactions, established that what recipients anticipate the change will mean for them is central to whether they support or resist it. Waizenegger and Techatassanasoontorn (2022), interviewing eighteen staff at a New Zealand financial institution, identified four configurations of employee perception and reaction, finding that employees' perceived job consequences influenced their collaboration with the automation team. Asatiani et al. (2020) observed knowledge workers reacting to a *planned* introduction of robotic process automation, before any software was deployed, which makes the announcement itself an intervention. Reaction is therefore not a residue to be managed after design; it is information about the design, available before it.

Two strands sharpen who is heard. Rowley (1997) moved stakeholder influence beyond the dyad to network density and centrality, a correction that bites in a smaller firm, where the same individuals occupy several stakeholder positions at once. Readiness is likewise not one quantity: it combines change commitment and change efficacy as a shared state rather than an aggregate of attitudes (Weiner, 2009), is measured by instruments frequently lacking demonstrated reliability and validity (Weiner et al., 2008), and operates at individual, group and organisational levels at once (Rafferty et al., 2013). Doherty and King (2001) supply the consequence: organisational rather than technical issues are systematically under-addressed in systems development projects, an omission associated with poorer outcomes.

2.5 Governance instruments and the principles-to-practice gap

The governance environment has thickened rapidly. Value-based principles for trustworthy AI were set out in the OECD Recommendation, adopted in 2019 and subsequently amended (OECD, 2019), and elaborated as a lawful, ethical and robust conception with seven key requirements (High-Level Expert Group on Artificial Intelligence, 2019); operational structure followed in a voluntary framework organised around Govern, Map, Measure and Manage (National Institute of Standards and Technology, 2023), guidance on AI risk management, and the first certifiable AI management system standard (International Organization for Standardization, 2023a, 2023b). Regulation (EU) 2024/1689, the Artificial Intelligence Act, was adopted on 13 June 2024, published in the Official Journal on 12 July 2024 and entered into force on 1 August 2024, establishing a risk-tiered structure in which obligations scale with the risk class of the system; prohibitions and AI-literacy provisions applied from 2 February 2025 and general-purpose AI and governance provisions from 2 August 2025 (European Parliament and Council, 2024). This paper engages with that structure and principle, not with the compliance calendar.

Alongside the risk frameworks sits a process standard for ethics itself: IEEE Std 7000-2021 defines a model process for addressing ethical concerns during system design, covering traceability of ethical values into ethical requirements and risk-based design (Institute of Electrical and Electronics Engineers, 2021), ethics by design as a traceable activity with defined outputs. The difficulty is not scarcity of principle but its inertness. Jobin, Ienca and Vayena (2019) analysed 84 documents and found convergence on five principles: transparency, justice and fairness, non-maleficence, responsibility and privacy, alongside substantial divergence in interpretation and implementation, and Mittelstadt (2019) argued that principles alone cannot guarantee ethical AI, the field lacking the professional history, proven methods and accountability structures that make principles operative in medicine. Interviewed practitioners applying a national principle set encountered tensions and trade-offs requiring situated judgement rather than a checklist, their authors concluding that "a gap remains between high-level principles and practical techniques" (Sanderson et al., 2023). Section 4.4 draws on the literature around this gap and on Green's (2022) argument that a principle not converted into a design constraint capable of failing a test is decoration.

2.6 Human oversight, automation bias and algorithm aversion

The governance instruments converge on human oversight as the principal residual control, and a human-factors literature four decades old establishes that this control is weaker than it looks. Bainbridge (1983) named the ironies: automating the tractable parts of a task leaves the operator the hard residue while removing the practice by which the skill to handle it was maintained. Parasuraman and Riley (1997) decomposed reliance failure into four separately diagnosable modes: use, misuse, disuse, and abuse, misuse being over-reliance with inadequate monitoring, disuse under-reliance frequently triggered by false alarms, and abuse the automation of functions without regard for downstream human-performance consequences; a control addressing only over-reliance addresses one of four. Skitka, Mosier and Burdick (1999) established automation bias experimentally as producing both omission errors, where a person fails to act because the automation did not prompt them, and commission errors, where a person acts on an automated recommendation against contrary evidence. Lee and See (2004) reframed the objective accordingly: not maximised but calibrated trust, reliance matched to actual capability, a property of the human–system pair rather than of the system.

The direction of miscalibration is not fixed. Dietvorst, Simmons and Massey (2015) reported algorithm aversion, people abandoning a statistically superior algorithm faster than an equally erring human forecaster after seeing it err once; Logg, Minson and Moore (2019) reported the opposite under other conditions; and Burton, Stein and Jensen (2020), reviewing 61 peer-reviewed articles spanning 1950 to 2018, organised the causes under five themes, including expectations and expertise, decision autonomy, incentivisation, cognitive compatibility, and divergent rationalities, arguing that the contradictory findings reflect incompatible framings rather than failed replication.

Green's (2022) objection therefore has an empirical literature behind it rather than resting on argument alone: a control specified as "a human reviews the output" states neither which failure mode it addresses nor what would count as evidence that discretion was exercised rather than deference performed.

2.7 Model risk, MLOps and the practice of AI audit

Two adjacent practice literatures supply structure a smaller firm can borrow rather than invent. The first is model risk management, a supervisory regime predating the current AI governance wave by more than a decade: the guidance issued jointly by the Board of Governors of the Federal Reserve System and the Office of the Comptroller of the Currency (2011) defines model risk as the potential for adverse consequences from decisions based on incorrect or misused model outputs, and organises control around validation independent of development, effective challenge, and documentation maintained across the model's life. Its central move, attaching risk to the decision taken on the output rather than to the model in isolation, is the one most often lost when governance is organised around the model alone.

The second is the engineering practice grown up around production machine learning. Kreuzberger, Kühl and Hirschl (2023) derived a definition, principle set, component set and role set for MLOps from a combined literature review, tool review and expert interviews, giving a named process spine to which controls attach; and model cards report intended use, out-of-scope use, and evaluation disaggregated across relevant groups (Mitchell et al., 2019). Assurance can also be organised by lifecycle stage rather than by artefact, with stated desiderata for data management, model learning, model verification and deployment (Ashmore et al., 2021), the cleanest structure for an assurance case, and the one from which a reduced version can be cut; documentation extends to the data as well as the model (Geburu et al., 2021), warranted because data work is undervalued inside AI-building organisations and its failures compound downstream as data cascades (Sambasivan et al., 2021), the debts that matter accruing at system rather than model level (Sculley et al., 2015).

Audit practice should not be over-credited: Costanza-Chock, Raji and Buolamwini (2022), scanning the algorithmic auditing ecosystem, found a market without accreditation, agreed standards of practice, or any requirement that findings be acted upon. The auditing literature has begun to bound its own claims, defining ethics-based auditing as "a structured process whereby an entity's present or past behaviour is assessed for consistency with relevant principles or norms" and stating explicit feasibility constraints on it (Mökander et al., 2021). Structure exists and scales down; an audit is not yet an assurance.

2.8 Value, benefits and the productivity question

IT generates value only alongside complementary organisational resources (Melville et al., 2004), attribution to discrete investments is difficult (Kohli & Grover, 2008), and the productivity

shortfall is explained by mismeasurement, lags, redistribution and mismanagement (Brynjolfsson, 1993) and, for AI, by implementation lags and complementary intangible capital (Brynjolfsson et al., 2019). The same result appears from the economics side, identical technology yielding very different returns across firms because of complementary organisational investment rather than the technology itself (Brynjolfsson & Hitt, 2000). The method response is benefits management: benefits arise from business changes that IT enables rather than from the investment, linked to those changes through a Benefits Dependency Network (Ward & Daniel, 2012), each with a named business owner (Peppard et al., 2007). The field began from the observation that appraisal and realisation had come apart (Ward et al., 1996), a diagnosis revisited a decade later (Ward et al., 2007). Benefits management has since been restated as an organisational competence rather than a document. Ashurst, Doherty and Peppard (2008) proposed a benefits realisation capability model organised around four competences consisting of planning, delivery, review and exploitation, of which the last two are routinely omitted, both falling after the delivery organisation has dispersed; Doherty, Ashurst and Peppard (2012) found realisation depending on practices sustained beyond go-live; and Serra and Kunc (2015) positioned benefits realisation management as the mechanism connecting projects to strategy execution. Badewi (2016) separated project management practice from benefits management practice as contributors to different dimensions of success: delivering on time is not delivering benefit. The implication for a smaller firm is uncomfortable: the competences that matter most are furthest from the delivery event, which is the only moment at which a small organisation reliably has anyone's attention.

Three measurement results complete the picture. "Success" is itself multi-dimensional rather than one variable (DeLone & McLean, 1992, 2003); value is more defensibly measured first at business-process level (Elbashir et al., 2008); and a measurement system fails chiefly by drifting out of alignment with a changing strategy rather than by poor initial design (Bourne et al., 2000). Where an indicator must be chosen rather than invented, 140 process-related indicators have been synthesised from 76 publications into eleven categories, cost measurement and organisational flexibility among the thinly covered dimensions (Van Looy & Shafagatova, 2016).

2.9 The business analyst role and the evidence about it

The paper's central claim concerns a role, and the honest position is that the empirical literature on that role is thin. The strongest adjacent result is Bassellier and Benbasat's (2004), which conceptualised the business competence of IT professionals and related it to effective IT–business partnering, establishing that what makes the intersection work is business knowledge held by technical professionals rather than the reverse. Directly on the role, Sonteya and Seymour (2012) analysed the competencies of the business process analyst empirically, distinguishing what practitioners do from what job advertisements describe. Neither study concerns AI, and neither concerns smaller firms.

Two further results bear on the setting. Aranda, Easterbrook and Wilson (2007) found requirements practice in small companies managed successfully through informal, tacit and culture-carried mechanisms rather than the documented processes textbooks prescribe. Nahar, Zhou, Lewis and Kästner (2022), studying 45 practitioners across 28 organisations building machine-learning-enabled systems, found the hardest problems organisational rather than algorithmic, clustering at three collaboration points: requirements and planning, training-data handling, and product–model integration. Those are boundaries between groups, and none is a modelling problem.

Adjacent to the role, the requirements literature has begun to ask what changes when the specified artefact contains a learned component: the work demands familiarity with performance measures, elicitation of quality requirements such as explainability and freedom from discrimination, and accommodation of the shift from coding to training (Vogelsang & Borg, 2019); and across 43 primary studies current approaches "were not adequately adaptable for building AI systems" (Ahmad et al., 2023). Neither names a role; both describe work that has to be someone's. No evidence exists that assigning these boundaries to a business analyst improves outcomes; the claim developed in Section 4 is therefore positional.

2.10 Small and medium-sized enterprises and AI adoption

The adoption base is established: perceived usefulness and ease of use (Davis, 1989), the unified acceptance constructs (Venkatesh et al., 2003) and their consumer extension (Venkatesh, Thong, & Xu, 2012), the technological–organisational–environmental frame (Tornatzky et al., 1990), and diffusion (Rogers, 2003). An older small-firm tradition sits alongside these: adoption modelled through perceived benefits, organisational readiness and external pressure (Iacovou et al., 1995), resource scarcity as the defining boundary condition of implementation (Thong, 2001), and digitalisation antecedents located in managerial and employee competence (Eller et al., 2020); all placing the binding constraint inside the organisation.

Applied to smaller firms and AI, the picture is consistent in direction and thin in depth. Schwaeye et al. (2024), reviewing 106 studies, synthesised the field around eight TOE-grounded clusters: compatibility, infrastructure, knowledge, resources, culture, competition, regulation and ecosystem, and Ayinaddis (2025), in a TOE-framed review of 78 articles, located the smaller-firm gap in infrastructure, finance, skills and management support. Sánchez, Calderón and Herrera (2025) analysed adoption through a TOE–DOI lens; despite "Survey" in its title, that work is a literature-based conceptual analysis rather than a primary firm survey and must not be read as a source of survey statistics. Sagala and Öri (2024), reviewing 106 articles, found incremental implementation matched to the organisational baseline recurring as a success condition. On readiness, Jöhnk, Weißert and Wyrski (2021) derived 5 categories, 18 organisational AI readiness factors and 58 indicators from 25 expert interviews, and Uren and Edwards (2023) established that readiness spans people, process and data as well as technology. Though Sadiq et al. (2021), reviewing 15 primary studies, found most AI maturity models proposed without empirical validation, a caution applying here. Prerequisites and post-development difficulties: data drift, infrastructure standardization, stakeholder expectation management, are documented elsewhere (Baier et al., 2019; Kinkel et al., 2022), while Alsheibani, Cheung and Messom (2018) proposed a firm-level AI-readiness framework that is research-in-progress with no empirical results and is cited here only as a conceptual structure.

On numbers, Eurostat (2025b) reports that for reference year 2025, 17% of small enterprises, 30.36% of medium enterprises and 55.03% of large enterprises used AI, with an EU average of 19.95% for enterprises with ten or more employees, the EU size classes being small (10–49 employees), medium (50–249) and large (250 or more); Eurostat (2025a) reports 20.0% of EU enterprises with ten or more employees using AI technologies in 2025, up 6.5 percentage points from 13.5% in 2024, against 7.7% in 2021 and 8.1% in 2023. The OECD (2025) provides qualitative evidence on adoption barriers and enablers across G7 economies and Brazil. IBM (2023) surveyed 2,342 IT professionals across 20 markets, reporting 42% deployed and 40% exploring among enterprises with more than 1,000 employees against 25% deployed and 46% exploring among companies with 1,000 or fewer; the unit of observation is IT professionals rather

than firms, and the 1,000-employee threshold does not correspond to any standard SME definition, so these figures cannot be read as firm-level small-enterprise adoption rates.

Official firm-level statistics point downward. Bonney et al. (2024), in a US Census Bureau working paper using the Business Trends and Outlook Survey for September 2023 to February 2024, report bi-weekly estimates of the AI use rate rising "from 3.7% to 5.4%, with an expected rate of about 6.6% by early Fall 2024", use being defined as use "for business purposes"; it is a working paper, no sample size is stated on the record consulted, and the measure is not comparable with Eurostat's.

This evidence base is thin and inconsistently defined: Eurostat measures "use of AI technologies" as a broad category including uses far removed from process automation; the IBM index uses an incompatible size threshold and a respondent unit that is not the firm; and the reviews synthesise studies defining "SME" differently across jurisdictions. No defensible point estimate exists for the proportion of smaller firms that have automated a business decision using AI, and none is offered here. The claim that survives is directional: adoption rises steeply with firm size, and the gap is attributed to infrastructure, finance, skills and management support rather than to available technology (Ayinaddis, 2025; Schwaewe et al., 2024). Practitioner-facing work has begun to propose AI-based advisory services as a route by which capabilities previously confined to larger organisations, trade compliance advice among them, might be extended to smaller firms (Jones & Ominyi, 2025a), a proposal in a venue without an established peer-review record, cited for its position and not as evidence that any such extension has occurred.

2.11 Automation implementation evidence

Van der Aalst, Bichler and Heinzl (2018) characterised robotic process automation as "outside-in" automation operating at the user-interface level and leaving underlying systems unchanged, fast to deploy and structurally fragile. What this paper calls intelligent process automation has a peer-reviewed anchor rather than a vendor one, being characterised as a technical extension of that pattern covering learning, unstructured input and dialogue (Chakraborti et al., 2020). The implementation record is case-based: the Telefónica O2 deployment (Lacity & Willcocks, 2016a, 2016b); the Xchanging deployment, reported in an LSE working paper (Willcocks et al., 2015) and revisited as a lever for global business services (Willcocks et al., 2017); 39 evidence-based action principles across 22 organisations (Lacity et al., 2021); thematic and success-factor syntheses (Plattfaut et al., 2022; Syed et al., 2020); an eleven-stage framework from 35 cases, 8 expert interviews and 5 validation workshops (Herm et al., 2023); and the field's open questions (Plattfaut et al., 2024).

Practitioner evidence requires care. A frequently cited industry observation places early RPA project failure at thirty to fifty per cent (EY, 2016); it should be read as practitioner experience rather than as a measured rate, and the academic literature has recirculated it without independent verification. Deloitte's global RPA survey reported that 72% of respondents had already started their RPA journey while only 3% of organizations had scaled their digital workforce, defined as fifty or more robots, on a base of over 400 responses globally (Deloitte, 2018); an earlier wave of the same series, published under the same title, reports different figures, and the two must not be mixed. Wright et al. (2020) reported that 73% of respondents said their organisations had embarked on a path to intelligent automation and 78% were implementing RPA, with 13% deploying 51 or more automations and only 26% of piloting organisations holding an enterprise-wide intelligent automation strategy. Telford et al. (2022), surveying 479 executives across 35 countries, reported that "seventy-four per cent of survey respondents are already implementing

RPA, and 50 per cent are already implementing OCR," naming process fragmentation, lack of clear vision, IT readiness and workforce resistance as barriers to scaling. These are self-selected practitioner samples reported by vendors of advisory services; they indicate the shape of a problem, not its magnitude, which lies in the organisational envelope rather than the automation: task–technology and task–environment fit and supporting conditions (Trkman, 2010), and six core elements of which technology is one (Rosemann & vom Brocke, 2014).

Two cautions follow. The evidence base is unevenly composed, a systematic mapping finding much of the available material industrial rather than scientific (Enríquez et al., 2020); and single-number failure rates are weak for a deeper reason than sampling, since success and failure are enacted judgements that can coexist for one system (Cecez-Kecmanovic et al., 2014). This paper asserts no failure rate anywhere.

What the case literature supplies is where the work goes: adoption driven from below rather than by IT, with practitioners not expecting reduced headcount (Cooper et al., 2019), and deployment creating new human work rather than only removing it (Grønsund & Aanestad, 2020).

2.12 Explanation and fairness: two literatures that do not supply a control

Two technical literatures are routinely offered as the operational answer to governance obligations, and both resist that role.

Explanation is not one technique. Post hoc explanation of an opaque model is the weaker control for high-stakes decisions, an inherently interpretable model being preferable (Rudin, 2019), which relocates explainability from a reporting layer to a design-time modelling choice; and explainability as deployed is used overwhelmingly by internal engineers for debugging rather than by the people the decisions affect (Bhatt et al., 2020), a distinction between explanation-for-development and explanation-as-accountability that a requirement reading "the system must be explainable" fails to make. Fairness is plural in a way that defeats the same phrasing. There are mutually incompatible formal definitions and several pipeline points at which bias enters, so a control not stating which definition it adopts has stated nothing testable (Mehrabi et al., 2021); disparate outcomes arise from ordinary modelling choices such as target-variable definition, training data, feature selection, proxy variables, masking of intentional discrimination (Barocas & Selbst, 2016); and technical interventions fail through abstraction traps when lifted out of the social system they act on (Selbst et al., 2019). Neither literature hands a smaller firm a control it can adopt unexamined; both require the act Section 4.4 describes, of choosing at design time which construal applies and what evidence would show it met.

2.13 The gap

Each literature addresses part of the problem and none joins them. Brokering theory explains why knowledge fails to cross but was developed in large-firm settings; the requirements literature details the artefacts of translation but treats governance and value as out of scope; the responsible AI literature says what ought to constrain design but not who converts a principle into a constraint; and the IT value literature assumes measurement infrastructure a smaller firm lacks. No account joins them in a single role. The additions above sharpen this description without closing the gap. The boundary literature carries its own corrections about when objects bridge and when they entrench (Barrett & Oborn, 2010; Star, 2010); the knowledge-transfer literature specifies why translation sticks (Szulanski, 1996); the human-factors literature specifies why oversight fails (Bainbridge, 1983; Parasuraman & Riley, 1997); the benefits literature specifies which competences are omitted (Ashurst et al., 2008); and the engineering literature on machine-

learning-enabled systems locates the difficulty at collaboration points (Nahar et al., 2022). The explanation and fairness literatures name the choices a control must make without saying who makes them (Mehrabian et al., 2021; Rudin, 2019), and the assurance and readiness literatures supply a structure and a construct without saying whose job either is (Ashmore et al., 2021; Weiner, 2009). Each remains a literature about a phenomenon rather than about a role. The gap is not that the phenomena are unstudied but that nothing assigns them.

3. Method

This paper is a conceptual contribution. The framework was developed by structured narrative synthesis of the published literature identified in Section 2, read against the author's professional experience of business analysis practice. No primary data were collected, and the framework has not been empirically validated. Section 8 sets out the resulting limitations, and Section 7.4 proposes a validation agenda. Failure patterns reported in the implementation literature and practitioner surveys were characterised by what knowledge failed to move and between which groups, classified against Carlile's (2004) hierarchy, grouped into four functions, and stress-tested against the small-firm constraint. The illustration in Section 6 is hypothetical and reports no data. The evidence base is of two kinds, kept apart. The theoretical claims rest on the peer-reviewed literature of Section 2; Section 5 surveys applied and practitioner writing, much of it in venues without an established peer-review record, solely to show how the four functions surface in particular domains. No proposition or quantity depends on that second body, and Section 8 restates the distinction as a limitation.

4. The framework: the business analyst as AI-transformation broker

4.1 The premise: failure at boundaries, not at components

The four failure patterns named at the outset sort onto the hierarchy. Automating a process nobody agreed on and satisfying a checklist that changes no design decision are both pragmatic: the first encodes interests never surfaced, the second builds a lexicon without transforming any interest. Deploying a model whose data was never fit for the decision is semantic, since terms such as "complete application" mean different things in the operational and data worlds while the words are shared. Declaring a benefit nobody measured compounds both, since no party held an interest in a baseline that could later contradict them. The diagnostic value is that this forecloses answering a higher-order boundary with a lower-order instrument: a glossary does not resolve conflicting interests, and nor does a workshop unless it changes what one party will give up. The work therefore needs a broker, that is, someone who can make each side's consequences for the other visible and has standing in both to hold the negotiation open, producing brokered knowledge (Meyer, 2010), creating value by recombinant access (Hargadon & Sutton, 1997), and holding authority earned in practice rather than conferred by an organisation chart (Levina & Vaast, 2005).

4.2 F1 Translation

What is moved. Intent, operational constraint and evidentiary expectation outward; feasibility, data reality and error behaviour back. The characteristic error is treating this as unidirectional elicitation.

Boundary type. Predominantly semantic on a syntactic substrate.

Boundary objects. The requirement specification is primary, and must hold the Star and Griesemer (1989) dual property: plastic enough that operations staff recognise their work in it,

robust enough to remain one object rather than two readings. Its normative process and quality characteristics come from ISO/IEC/IEEE 29148 (International Organization for Standardization, 2018); a specification is a state of agreement, not a document (Pohl, 2010); the practical apparatus is standard (Robertson et al., 2024; Wiegers & Beatty, 2013); goal-oriented refinement preserves the link to the intent served (van Lamsweerde, 2009); and requirements smells give automatable defect indicators grounded in those criteria (Femmer et al., 2017). For automation the decisive objects are decision-side: DMN expresses decision logic separately from process flow (Object Management Group, 2024) and SBVR formalises the vocabulary and rules it depends on (Object Management Group, 2019), the two integrated on stated principles (Hasić et al., 2018) once decision logic is extracted from the control flow in which real process models bury it (Batoulis et al., 2015). The decision point, not the process step, is the unit of translation.

Why it is hard. The difficulty is not notation. The decisive knowledge at a decision point is frequently tacit, held in the performance rather than the procedure (Brown & Duguid, 1991; Nonaka, 1994), and what obstructs its movement is absorptive capacity and causal ambiguity about why a rule exists rather than reluctance to explain it (Szulanski, 1996). So a requirement should carry the rationale for a rule and not the rule alone; the specification must be tested for usability by the receiving field rather than by its author, the boundary object's dual property being a claim about use and not about drafting (Star, 2010); and it should not be asked to motivate, adjudicate and instruct at once (Nicolini et al., 2012). Tooling does not yet substitute: across 404 studies over 36 years and 170 venues, 67.08% were solution proposals validated only by laboratory experiment or example, approximately 7% were assessed industrially, and of 130 identified tools only 17 (13.08%) remained available (Zhao et al., 2021).

Three results specify what goes missing. Rationale is the first: the traceability literature distinguishes traceability after a specification exists from the pre-specification trail of where a requirement came from and why, and it is the latter that is characteristically lost (Gotel & Finkelstein, 1994), Szulanski's causal ambiguity from another direction. Second, elicitation has a floor thoroughness does not raise (Sutcliffe & Sawyer, 2013), knowledge lodged in the interaction among people, tools and tasks moving least readily (Argote & Ingram, 2000). Third, a specification is a reification whose meaning depends on the participation that produced it (Wenger, 1998) and on knowing enacted in the receiving practice (Orlikowski, 2002), which is why an artefact that reads as complete can fail on arrival.

Competence. Modelling and abstraction; elicitation under tacit knowledge; the discipline to specify decision, data-at-decision, tolerated error, escalation and evidence.

Failure mode when absent. The specification gap persists undetected until acceptance.

4.3 F2 Legitimation

What is moved. Standing: claims into the design conversation, and consequences out to those who bear them.

Boundary type. Pragmatic: vocabularies may be shared while interests diverge, and clarification dissolves nothing.

Boundary objects. A stakeholder registers with explicit salience rationale; an impact statement naming who bears each option's consequences; a decision record showing which claims were admitted and why.

Substance. The obligation is Freeman's (1984); the power–legitimacy–urgency typology explains why claims receive attention (Mitchell et al., 1997) and predicts a systematic bias, since the workforce being automated holds legitimacy and sometimes urgency but little power, and such stakeholders are attended to only where an advocate supplies the missing attribute, which in a smaller firm no other function will. Donaldson and Preston's (1995) distinction disciplines the claim: instrumentally, it is argued here that automations designed without operators' exception knowledge are likely to perform worse, because exceptions are where value and risk both live, which is joint optimisation restated (Sarker et al., 2019; Trist & Bamforth, 1951); normatively, people whose work is reallocated have standing to be consulted whether or not that improves outcomes, which is a commitment rather than a finding. Because the involvement–success relationship is heterogeneously defined (Bano & Zowghi, 2015), legitimation is proposed not as a lever with a known effect size but as a condition without which the pragmatic boundary cannot be crossed.

Timing. Recipients' reactions are structured and largely anticipatory (Oreg et al., 2011), and in automation settings shape whether staff collaborate with the automation team at all (Waizenegger & Techatassanasoontorn, 2022). Legitimation is therefore not consultation after a design is fixed: by then the operators' exception knowledge has been priced out, and what remains is a negotiation over compliance. The same literature counsels claiming less rather than more, since change management rests on prescriptive models with a weak empirical base and its most-repeated failure statistic is contested (By, 2005). And because a cross-boundary settlement reflects the relative standing of the domains involved rather than their evidential merit (Oborn & Dawson, 2010), an advocate for the least-resourced domain changes it only if credible in the better-resourced one.

What legitimation is building. Salience is structural as well as dyadic, influence following from network density and position within it (Rowley, 1997), so in a smaller firm, where one person may be sponsor in one relation and affected operator in another, part of the act is choosing the relation in which a matter is raised. What is built is not agreement but a shared state of change commitment and efficacy (Weiner, 2009), forming across levels at once (Rafferty et al., 2013) and measured by instruments of frequently undemonstrated reliability (Weiner et al., 2008). Ethnography shows the negotiation: building an operational AI system is a contest between developers' and domain experts' knowledge claims, so what is deployed encodes an organisational settlement rather than a data-driven optimum (van den Broek et al., 2021). Where that settlement is never held, organisational issues go unaddressed, an omission associated with poorer outcomes independently of technical quality (Doherty & King, 2001).

Competence. Stakeholder analysis; facilitation of conflict rather than avoidance; standing to convene people who do not report to the analyst.

Failure mode when absent. Design follows the best-resourced interest, and deployment resistance is misread as change fatigue.

4.4 F3 Constraint injection

What is moved. Regulatory, ethical and risk constraints, moved into the design space while it can still change.

Boundary type. Pragmatic, because constraints cost delivery interests, and semantic, because governance and engineering vocabularies do not map cleanly.

Boundary objects. A risk register in design terms; acceptance criteria that encode a constraint as a test that can fail; documentation of intended use, data provenance and known limitations; and an oversight specification stating not that a human reviews but what they see, what they can change, and what evidence shows discretion was exercised.

Substance. The instruments of Section 2.5 supply the structure: Govern, Map, Measure and Manage as a sequence (National Institute of Standards and Technology, 2023); risk-management guidance and a certifiable management system (International Organization for Standardization, 2023a, 2023b); the framing-stage question of which risk tier applies (European Parliament and Council, 2024); and the principle set (High-Level Expert Group on Artificial Intelligence, 2019; OECD, 2019), whose non-bindingness is documented elsewhere (Floridi et al., 2018; Hagendorff, 2020; Jobin et al., 2019; Kaur et al., 2022; Mittelstadt, 2019). The governance literature distinguishes team, organisation and industry layers (Shneiderman, 2020), of which it is argued here that a smaller firm reaches only the team layer; supplies the SMACTR internal audit framework (Raji et al., 2020), argued here to be usable as scaffolding that scales down; asks, across 242 synthesised articles, to whom accountability is owed and what stays legible outside (Diakopoulos, 2015; Wieringa, 2020), identifies whether a constraint attaches to data, model or system (Mäntymäki et al., 2022; Schneider et al., 2023), and locates the enablers in organisational rather than technical conditions (Rakova et al., 2021). Green (2022) supplies the test: specifying human oversight is not achieving it, and a nominal reviewer can legitimise a flawed system while protecting no one. The brokering act is to convert a principle into a constraint that can fail, a threshold, an acceptance criterion, a required evidence artefact, a stop condition. A principle that cannot fail anything has not been injected; it has been filed.

The conversion requires choices the principle does not make. "Explainable" must become either a commitment to an interpretable model or a specified post hoc artefact with a named recipient, since those are different controls with different failure modes (Bhatt et al., 2020; Rudin, 2019). "Fair" must name which incompatible formal definition is adopted and at which pipeline stage bias is addressed (Mehrabian et al., 2021), knowing that disparate outcomes follow from ordinary modelling choices (Barocas & Selbst, 2016) and that a remedy detached from the social system meets the abstraction traps Selbst et al. (2019) name; IEEE Std 7000-2021 supplies a normative process for the same conversion on the values side (Institute of Electrical and Electronics Engineers, 2021). Documentation, where it is the control, must reach the data and not only the model (Geburu et al., 2021), since data work is undervalued and its failures compound downstream (Sambasivan et al., 2021); the assurance case can be laid out by lifecycle stage and cut down deliberately (Ashmore et al., 2021); and whatever audit follows should be bounded by its own feasibility constraints (Mökander et al., 2021).

Why oversight must be specified rather than asserted. That test has an empirical literature behind it (Section 2.6). Reliance fails in at least four separable modes (Parasuraman & Riley, 1997), through omission where the system does not prompt and commission where it prompts wrongly (Skitka et al., 1999), against a background in which automation has taken the tractable

part of the task and left the hard residue (Bainbridge, 1983). The target is calibrated rather than maximal reliance (Lee & See, 2004), and the direction of miscalibration cannot be assumed (Burton et al., 2020; Dietvorst et al., 2015; Logg et al., 2019). An oversight specification surviving this literature states which failure mode it addresses, what the reviewer sees, what they may change, and what evidence distinguishes exercised discretion from performed deference. One caution belongs to the loop: human–AI collaboration can degrade the reviewer's judgement, removing the independence the control was meant to supply (Fügener et al., 2021).

Borrowable structure. Model risk management supplies the older precedent, organising control around independent validation, effective challenge and lifecycle documentation (Board of Governors of the Federal Reserve System & Office of the Comptroller of the Currency, 2011); MLOps supplies a named process spine (Kreuzberger et al., 2023); and model cards supply a documentation form covering intended use, out-of-scope use and disaggregated evaluation (Mitchell et al., 2019) close in shape to the one-page artefact Table 3 proposes. The caution is that audit is not assurance (Costanza-Chock et al., 2022). Applied governance work makes several of these points from the practitioner side: that model risk management is the practice through which algorithmic accountability is operationalised (Ominyi & Anichukwueze, 2023); that AI-enabled financial systems raise auditability problems distinct from conventional ones (Fobellah, 2025b); that human-in-the-loop is a design pattern to be specified rather than invoked (Ladapo et al., 2022); and that readiness for a specific AI regime can be assessed systematically, though the enterprises so assessed were multinational rather than small (Ominyi et al., 2024b). These are conceptual contributions in venues without established peer-review records, cited for their positions and never for a figure.

Competence. Regulatory literacy sufficient to identify which instruments apply; ability to render an obligation as a testable criterion; willingness to raise a constraint when it is expensive.

Failure mode when absent. Governance arrives as retrospective documentation: a compliance artefact and no additional safety.

4.5 F4 Benefit tracing

What is moved. Causal reasoning about value forward into design, and evidence back from operation.

Boundary type. Semantic between finance, operations and delivery; pragmatic, because a measured baseline can later contradict a claim someone has an interest in making.

Boundary objects. A benefits dependency structure (Ward & Daniel, 2012); a baseline captured *before* implementation; a measurement specification naming metric, source, owner and cadence.

Substance. Because benefits arise from business changes enabled by IT rather than from the IT itself, the business case becomes a dependency chain whose links can be checked, each benefit carrying a named business owner (Peppard et al., 2007; Ward & Daniel, 2012), necessary because value emerges only alongside complementary organisational resources (Melville et al., 2004), attribution is difficult (Kohli & Grover, 2008), mismeasurement and mismanagement are the broker's business (Brynjolfsson, 1993), and a lag before measurable return is expected (Brynjolfsson et al., 2019). Measurement design rests on the insistence that financial measures alone "can give misleading signals for continuous improvement and innovation" (Kaplan & Norton, 1992), later elaborated as a management method (Kaplan & Norton, 1996), and on stated requirements for a coherent measurement system (Neely et al., 1995), but has a precondition often binding in smaller firms: the reporting infrastructure must exist, usability and decision impact depending on organisational and cultural factors as well as technical ones (Popović et al., 2012;

Wixom & Watson, 2001). Where reporting is a set of spreadsheets a baseline can be asserted but not defended, and the brokering act may reduce to instrumenting one quantity before build.

Competence, not documentation. Benefits realisation is better read as an organisational competence than an artefact, and its competences are unevenly available in a smaller firm. Planning and delivery attach to the project; review and exploitation fall after it, when the delivery organisation has dispersed and nobody's objectives contain them (Ashurst et al., 2008), and realisation depends on practices sustained past go-live (Doherty et al., 2012) that connect an initiative to strategy (Serra & Kunc, 2015). The brokering implication is a scheduling one: those competences must be bought at framing, while approval is still being sought, by fixing a date, an owner and a metric rather than proposing them at closure. The metric should be selected from a catalogued space rather than invented (Van Looy & Shafagatova, 2016) and pitched at process level before any organisational claim is made on it (Elbashir et al., 2008). The claim it supports must stay modest: project and benefits management practice bear on different dimensions of success, so an initiative delivered on time may deliver nothing (Badewi, 2016), a gap documented before benefits management existed as a method and still open a decade later (Ward et al., 1996, 2007). Practitioner-facing work proposes instrumenting value measurement continuously inside the data pipeline rather than reconstructing it afterwards (Ogundairo & Ayivi-Donkor, 2023b), a proposal in a venue without an established peer-review record, noted as a direction rather than as evidence, but locating the binding constraint correctly in whether the quantity is instrumented at all.

Competence. Business-case construction; baseline discipline; measurement literacy sufficient to know when a metric moves for unrelated reasons.

Failure mode when absent. Benefit is declared from output rather than outcome, and an automation that worked cannot be distinguished from one that shifted work elsewhere.

4.6 Summary of the four functions

Table 1. The four brokering functions.

Function	What is moved	Boundary type (Carlile, 2004)	Boundary objects produced	Competence required	Failure mode when absent
F1 Translation	Intent, constraint and evidentiary expectation outward; feasibility, data reality and error behaviour back	Semantic, on a syntactic substrate	AI-ready requirement specification; decision model (DMN); vocabulary and rules (SBVR); goal refinement structure; glossary	Modelling and abstraction; elicitation under tacit knowledge; specification of decision, data-at-decision, tolerated error,	Specification gap persists undetected; divergence found at acceptance

Function	What is moved	Boundary type (Carlile, 2004)	Boundary objects produced	Competence required	Failure mode when absent
				escalation, evidence	
F2 Legitimation	Standing: claims into the design conversation, consequences out to those who bear them	Pragmatic	Stakeholder register with salience rationale; impact statement; record of admitted and rejected claims	Stakeholder analysis; facilitation of conflict; convening authority; judgement on which claims bind	Design follows the best-resourced interest; resistance misread as change fatigue
F3 Constraint injection	Regulatory, ethical and risk constraints into the design space early enough to alter design	Pragmatic and semantic	Risk register in design terms; acceptance criteria that can fail; intended-use and limitations record; oversight specification	Regulatory literacy; rendering obligation as testable criterion; willingness to raise costly constraints early	Governance becomes retrospective documentation; nominal oversight legitimises without protecting
F4 Benefit tracing	Causal reasoning about value forward; evidence back from operation	Semantic and pragmatic	Benefits dependency structure; pre-implementation baseline; measurement specification (metric, source, owner, cadence)	Business-case construction; baseline discipline; measurement literacy	Benefit declared from output not outcome; next investment decision made blind

4.7 Brokering across the initiative lifecycle

The functions are not phases: each recurs at every stage, but the dominant boundary shifts. Table 2 states the act required at each intersection.

Table 2. Boundary type by initiative phase: the brokering act required.

	Framing	Specification	Build	Deployment	Operation
Syntactic	Establish common	Fix notation: one process	Maintain one glossary	Ensure runbooks and	Keep the metric

	Framing	Specification	Build	Deployment	Operation
	vocabulary for the process and its decision points	model, decision logic in DMN, terms and rules in SBVR	binding data field names to business terms	procedures use the specification's terms	dictionary current as fields and systems change
Semantic	Surface the difference between the process as performed, as documented and as desired	Translate intent into an AI-ready requirement; apply requirements-smell checks	Reconcile what the data represents with what the requirement assumed	Translate model behaviour into operational language: what it does, cannot do, when to distrust it	Translate drift, error-rate change and exception patterns into business consequence
Pragmatic	Establish stakeholder salience, including those whose work is automated; identify risk tier and obligations	Convert principles into acceptance criteria that can fail; agree the baseline before build	Hold constraints against delivery pressure; force explicit rather than silent trade-offs	Negotiate reallocation of work, authority and accountability; make oversight exercisable	Sustain accountability: who answers to whom, on what evidence, with what consequence

4.8 What changes in a small or medium-sized enterprise

A large firm meets each boundary problem with a structure; a smaller firm has none and cannot buy it at the scale required. The functions must be substituted at a scale one role can sustain, smaller artefacts, tighter scope, and acceptance that some large-firm assurance is not reproduced, a genuine reduction rather than an equivalent alternative. Three conditions shape Table 3: the resource constraint, since the adoption gap lies in infrastructure, finance, skills and management support (Ayinaddis, 2025; Schwaeye et al., 2024); the sequencing constraint (Jöhnk et al., 2021; Sagala & Öri, 2024; Uren & Edwards, 2023); and the proximity advantage, since where sponsor, operations lead and analyst share a room the pragmatic boundary is shorter than in a matrixed organisation. The substitutions in Table 3 should not be read as impoverished imitations. Requirements practice in small companies has been observed to work through informal, tacit and culture-carried mechanisms rather than documented processes (Aranda et al., 2007), which suggests both that the small-firm form is a distinct mode rather than a deficient one and that formalising it risks destroying what made it effective. The minimum viable artefact set exists to make specific things inspectable: which claims were admitted, which constraint can fail, what the baseline was, and not to reproduce documentation.

Table 3. SME-specific conditions for each brokering function.

Function	What a large firm does structurally	What an SME lacks	What the broker substitutes
F1 Translation	Dedicated requirements or product function; enterprise architecture; data governance office maintaining canonical definitions; specialist decision-modelling capability	Any definitional authority; architecture function; time to maintain models not immediately used	One small set of living artefacts one decision model per automated decision point, one glossary, one requirement set, used in the meetings that make decisions, so maintenance is a by-product of use
F2 Legitimation	HR business partnering; organisational development; formal consultation machinery; programme-level stakeholder management	All of it; and no forum in which affected staff can raise consequences without addressing their own manager	Explicit convening: operators brought into design sessions as the source of exception knowledge, with admitted and rejected claims recorded so the reasoning stays inspectable; proximity as the compensating advantage
F3 Constraint injection	AI governance function; model risk management; legal and compliance gates; internal audit; second-line challenge	Second-line challenge of any kind; regulatory specialists; capacity to run an audit framework	A minimum viable control set: risk tier identified at framing; three to five constraints as acceptance criteria that can fail; one page on intended use and limitations; a named human decision owner, with the reduction in assurance stated to the sponsor
F4 Benefit tracing	Benefits realisation office; portfolio management; finance business partnering; established reporting infrastructure and metric ownership	Reporting infrastructure beyond operational systems and spreadsheets; anyone whose job is post-implementation review	One benefit, one metric, one baseline captured before build, one named business owner, one review date; where instrumentation is absent, building the measurement is part of scope

4.9 A benefits dependency structure

Table 4 renders benefit tracing concretely in the structural form of the Benefits Dependency Network (Ward & Daniel, 2012). Its entries are worked for the programme's illustrative casework domain and are entirely hypothetical, containing no data and describing no organisation.

Table 4. Illustrative benefits dependency structure for intelligent process automation in a document-heavy casework operation (structural form after Ward & Daniel, 2012). All entries are hypothetical.

Business driver	Investment objective	Business benefit	Business change	Enabling change	IT enabler
Rising case volume without proportionate headcount growth	Increase cases completed per analyst-day without raising error rates	Reduced elapsed time from intake to decision issuance	Reallocate analyst time from document checking to substantive assessment; revise team targets	Redesign the intake queue; retrain staff on exception handling; agree a new definition of "complete file"	Document classification and data extraction over the intake channel
Regulator expects consistent application of published eligibility criteria	Reduce unexplained variation in outcomes for materially similar cases	Improved defensibility of decisions under external scrutiny	Adopt the decision model as the authoritative statement of criteria; require recorded reasons for deviations	Standardize decision logic; establish a rule-change process; instrument a deviation log	Decision engine implementing the modelled logic, with rule versioning
Rework caused by incomplete files reaching assessment	Reduce the proportion of cases returned for missing or invalid documents	Lower rework cost; shorter cycle time for applicants	Move completeness checking to the point of submission; change applicant communication at that point	Define completeness rules; agree tolerance for borderline documents; establish escalation for contested rejections	Automated completeness validation at intake with a human escalation path
Regulator reporting assembled manually each period	Reduce effort and error in periodic regulatory reporting	Reporting produced from operational data rather than reconstructed	Retire the parallel manual record; accept the operational system as reporting source	Agree metric definitions against the regulator's categories; assign metric ownership;	Reporting layer built directly on case-management data

Business driver	Investment objective	Business benefit	Business change	Enabling change	IT enabler
				establish data quality checks	

No benefit is attributable to the IT enabler alone: each requires a business change someone must decide to make and own (Peppard et al., 2007), and each carries an enabling change that is neither technical nor free.

4.10 Pathologies of brokering

Nomination does not create competence: a designated liaison not recognised as legitimate in both fields produces artefacts neither adopts (Levina & Vaast, 2005), and double peripherality makes the position precarious (Meyer, 2010). Table 5 is the counterweight to the paper's main claim: the role can fail, predictably.

Table 5. Pathologies of the broker role.

Pathology	Observable signal	Mitigating condition
Capture by one side	The analyst reports the delivery team's constraints to the business, or the sponsor's demands to the team, but never the reverse; artefacts read as advocacy	Structural dual accountability; artefacts reviewed by both fields before they bind; rejected claims recorded from both directions
Nominal brokering	The role exists on the organisation chart; specifications are produced but neither operations nor engineering treats them as authoritative	Competence recognised in practice rather than designated: the analyst participates in the work of both fields sufficiently to be credible in each (Levina & Vaast, 2005)
Throughput bottleneck	All cross-boundary communication routes through one person; decisions queue behind their availability; direct contact between fields has ceased	The objective is boundary objects that reduce the need for mediation, not remaining the channel; direct contact convened rather than intercepted
Responsibility without authority	The analyst is accountable for benefits, compliance and scope simultaneously while unable to stop, defer or resource anything	Explicit escalation rights and a named decision owner for each constraint and benefit; undecidable matters visibly transferred, not silently absorbed
Illegitimate legitimization	Sign-off given by the analyst on matters such as clinical, legal, actuarial, on which they have no standing; a template substitutes for competent judgement	Constraints stated as testable criteria whose satisfaction is evidenced rather than attested; matters requiring professional standing routed outward (Green, 2022)

Pathology	Observable signal	Mitigating condition
Peripheral drift	The broker's vocabulary and artefacts become intelligible only to other brokers; neither field can read the outputs unaided	Boundary objects retain the Star and Griesemer (1989) dual property; periodic test that each field can act on the artefact without mediation

4.11 What grounds each function, and what it does not

The four functions are not equally supported by the same literature, and Table 6 states which tradition does the work in each case and where it stops. A framework assembled from several literatures can appear better evidenced than it is, because each borrowed result carries the authority of its own field into a setting that field never studied.

Table 6. Theoretical grounding of the four brokering functions, and its limits.

Function	Primary tradition	Supporting traditions	What the tradition supplies	What it does not supply
F1 Translation	Knowledge transfer and situated practice (Brown & Duguid, 1991; Nonaka, 1994; Szulanski, 1996)	Boundary objects and their critique (Nicolini et al., 2012; Star, 2010; Star & Griesemer, 1989); requirements engineering (International Organization for Standardization, 2018; Pohl, 2010)	Why documented procedure diverges from practice; why transfer sticks even where all parties are willing; what an artefact must do to be usable in both fields	Any account of who performs the work in a firm with no requirements function, or of how much artefact is enough
F2 Legitimation	Stakeholder theory (Freeman, 1984; Mitchell et al., 1997)	Change recipients' reactions (By, 2005; Oreg et al., 2011); socio-technical joint optimisation (Sarker et al., 2019; Trist & Bamforth, 1951); politics of cross-boundary work (Oborn & Dawson, 2010)	Why claims without power are ignored; why reactions form before deployment; why exception knowledge is a design input rather than a courtesy	An effect size for involvement, the underlying relationship being heterogeneously defined and measured (Bano & Zowghi, 2015)

Function	Primary tradition	Supporting traditions	What the tradition supplies	What it does not supply
F3 Constraint injection	Human oversight, automation bias and reliance (Bainbridge, 1983; Lee & See, 2004; Parasuraman & Riley, 1997; Skitka et al., 1999)	Governance instruments (European Parliament and Council, 2024; National Institute of Standards and Technology, 2023); model risk management (Board of Governors of the Federal Reserve System & Office of the Comptroller of the Currency, 2011); audit practice (Costanza-Chock et al., 2022; Raji et al., 2020)	Why a nominal reviewer is not a control; what a constraint must state to be testable; control structure that can be scaled down	Evidence that a reduced control set delivers proportionate assurance, or that laboratory findings transfer to an operator whose job is being reallocated
F4 Benefit tracing	Benefits management and realisation competence (Ashurst et al., 2008; Ward & Daniel, 2012)	IT business value (Kohli & Grover, 2008; Melville et al., 2004); benefits realisation and strategy execution (Serra & Kunc, 2015); performance measurement (Neely et al., 1995)	Why benefit follows business change rather than IT; which competences are systematically omitted and when; what a coherent measurement system requires	The measurement infrastructure the method presupposes, which smaller firms frequently lack (Popovič et al., 2012; Wixom & Watson, 2001)

4.12 Propositions

P1. In smaller firms, failures of intelligent process automation initiatives are more frequently attributable to unresolved semantic and pragmatic boundaries than to technical inadequacy of the components.

P2. Initiatives in which one identifiable role performs all four brokering functions realise more of their stated benefits than those distributing the functions without explicit ownership, controlling for scope and firm size.

P3. Formal nomination as boundary spanner is not associated with improved outcomes unless the role is recognised as legitimate by both the business and the technical field.

P4. Governance constraints expressed as acceptance criteria capable of failing a test are more frequently reflected in delivered system behaviour than those expressed as principles, checklists or attestations.

P5. Initiatives capturing a baseline before implementation are more likely to produce a defensible benefit assessment, and more likely to report a smaller benefit, than those measuring after deployment.

P6. The salience of the workforce whose tasks are automated is lower at framing than at deployment, and initiatives raising it earlier encounter less deployment-stage resistance.

P7. The substitutions available to smaller firms deliver lower assurance than the corresponding large-firm structures, but the marginal benefit of the first minimal control set exceeds that of subsequent additions.

P8. Brokering pathologies are more prevalent where the analyst holds accountability without escalation rights, and the throughput-bottleneck pathology increases with the number of concurrent initiatives assigned to one broker.

5. Brokering in applied settings

This section surveys applied and practitioner literature of varying peer-review provenance, and it now constitutes a substantial share of this paper's reference base rather than a marginal addition to it. Many of the contributions cited below appear in venues without an established peer-review record; several are conceptual frameworks or narrative reviews rather than reports of primary investigation; and none has been independently verified beyond the position its title asserts. It is offered for one purpose: to show how the four brokering functions manifest across specific settings, in the vocabulary practitioners there use. No statistic, percentage, sample size, effect size or validated result is taken from any of it, and no proposition, table or argument in Section 4 depends on it. The four-function framework and every theoretical claim advanced in this paper rest on the peer-reviewed organisation-theory, requirements-engineering, governance and benefits literature reviewed in Section 2; were this section removed entirely, nothing in Sections 4, 6 or 7 would need restating. What follows illustrates scope, not effect. Each subsection closes by naming which functions the domain evidence exercises and which of Carlile's (2004) boundary types is in play, so the material reads as a map of where the functions appear rather than as support for any claim about how well they are performed.

5.1 Translation between business and technical functions in enterprise system and platform initiatives

The boundary here is not the general one between "business" and "technology" but a narrower one: between a stated business intention and the configuration vocabulary of a packaged platform, which supplies a third language: objects, roles, permission models, flows, integration patterns, that neither sponsor nor operations lead speaks.

Applied writing on enterprise digital transformation treats operational fixes and stakeholder engagement as objectives of a single initiative rather than as separable workstreams (Sanni & Attah, 2023a); in platform work the translation becomes a design constraint, privacy obligations restated as requirements on a platform's role and access model (Badmus et al., 2024) and generative components in data preparation given an explicit human review point (Ladapo et al., 2024). A second strand treats the delivery apparatus itself as requiring specification rather than assumption: pipeline strategies for moving change through a platform are reviewed as practices bearing on deployment outcomes rather than as tooling detail (Badmus et al., 2018); the choice among competing deployment tools is presented as varying with enterprise context rather than settled on merit (Badmus et al., 2020); data movement is modelled as part of the platform architecture rather than left to implementation (Badmus et al., 2019a); and integration with surrounding systems is described through recurring patterns rather than as bespoke work repeated each time (Badmus et al., 2022b). Where the business interface is a service-management platform the translation is visible in the artefact itself, a request having to be classified in the platform's categories before it can be routed (Ladapo et al., 2023a); migration of applications and information systems to cloud infrastructure is reported as carrying organisational and not only technical lessons (Ladapo et al., 2025); and integrating systems engineering with an operational function is framed as needing a framework rather than an interface (Okonkwo et al., 2023). Where an initiative spans previously separate technology estates the problem is unmistakable: alignment between information and operational technology is treated as a design principle rather than an outcome (Adegbite et al., 2022), a semantic boundary problem before a security one, since the estates name the same assets differently. The shape recurs where guardrails and an integration layer are prerequisites for agentic orchestration rather than later hardening (Ogundairo, 2025), where a digitalised operational chain needs an explicit rather than an inherited governance framework (Okonkwo et al., 2024a), and where transformations are described through staged change models (Ilodigwe & Adesemoye, 2025b). The function exercised is F1, and the boundary is semantic on a syntactic substrate in Carlile's (2004) sense: the platform's object names must first be bound to business terms that transfer across a shared syntax before anyone can ask whether the configured object means what the sponsor meant, which is translation.

5.2 Legitimation and stakeholder alignment in regulated and multi-party programmes

The boundary in multi-party programmes is one of standing rather than understanding. Where no single organisation can decide, a claim has no automatic route into the design conversation: each party has its own mandate, regulator and accountability, and vocabulary is rarely the obstacle. Applied work treats coordination across stakeholders in regulated joint ventures as requiring an explicit governance alignment mechanism (Eyetsesmitan et al., 2020), models accountability between public and private partners rather than leaving it to contract (Aminu-Ibrahim et al., 2024), and holds cross-border data sharing to require a jointly agreed framework (Akomolafe et al., 2025a). The allocation of exposure receives the same treatment, risk in public-private arrangements being represented as a layered architecture assigning who bears what rather than as a schedule appended to an agreement (Ike et al., 2024a). Collaboration across institutional boundaries is modelled deliberately rather than left to emerge (Isiekwu et al., 2025), and supplier relationships are described as a designed practice supporting collaboration rather than as transactional administration (Ike et al., 2021). Two contributions on multinational infrastructure programmes treat project governance as the thing to be optimised rather than as an overhead on delivery (Amayo et al., 2023a, 2024b), and coordination across sites in a clinical implementation

is presented as a project-management problem in its own right (Omo Enabulele et al., 2025). What these share with Section 4.3 is that a claim without an admitted channel is not a claim; what they lack is any account of why some claims are attended to and others not (Mitchell et al., 1997; Rowley, 1997). Two threads bear on the smaller-firm case. One positions a functional leader as the coordinating point across partner organisations (Isiekwu et al., 2021), structurally this paper's move for the analyst, and open to the same objection, that a position does not confer competence (Levina & Vaast, 2005). The other treats acceptance as something to be operationalised, framing social licence as a measurable variable (Komi, 2024) and naming user trust, privacy and regulatory governance as recurring conditions on acceptance of AI-enabled monitoring (Akintola et al., 2025), the question Donaldson and Preston's (1995) distinction poses.

The function exercised is F2 and the boundary is pragmatic: parties can state each other's positions accurately and still hold incompatible interests, so clarification dissolves nothing and only a transformation of what each will give up moves the work forward.

5.3 Constraint injection in security, data governance and compliance settings

The boundary here is between an obligation as written; in a statute, a standard, a contract or a policy, and a design decision that can be evaluated against it. Security and compliance functions have been crossing it for far longer than AI governance has existed, and it is unsurprising that this is by a wide margin the largest applied literature of the four.

Applied work describes regulatory constraints embodied in the data architecture feeding analytics rather than checked against it afterwards (Atakpa & Abolaji, 2022), data classified by sensitivity and traced to the obligation governing it (Aliliele et al., 2024a), impact assessment performed before deployment (Mbonu et al., 2022a), and privacy requirements translated into access-control policy (Ladapo et al., 2023b). The same conversion appears where a control weakness in a platform is detected by continuous automated monitoring rather than periodic inspection (Aliliele et al., 2023a), where integration points between systems are given a governance regime and a means of prioritising risk across them (Aliliele et al., 2023b), where legal and ethical exposures from enterprise data handling are modelled explicitly rather than treated as residual (Mbonu et al., 2018), and where divergent obligations across jurisdictions are held to bear on implementation choices rather than on documentation alone (Mbonu et al., 2019). Access is treated as an architecture rather than an administrative task (Ogbole et al., 2021), published frameworks are credited with a structuring role in how organisations enact governance (Jimoh et al., 2023), and control effectiveness is subjected to continuous validation rather than attested periodically (Dosunmu & Ogundele, 2024a); the practitioner form of the requirement that a constraint be capable of failing something. The AI-specific writing is younger and points the same way. Adversarial exposure is modelled within a system's governance rather than discovered in operation (Adebayo et al., 2022); agentic systems are described as introducing failure modes including the granting of more agency than intended (Nwakamma et al., 2024b); oversight is proposed as continuous rather than periodic (Aliliele et al., 2025a); and, in a domain vocabulary, Section 2.12's point recurs that a model may be accurate in aggregate while distributing error unevenly across groups (Komi, 2023; cf. Barocas & Selbst, 2016; Mehrabi et al., 2021). Governance structures purpose-built for automated decision systems in regulated industries are argued to be necessary rather than adaptable from general practice (Ominyi & Anichukwueze, 2024), quality control over multi-agent output in a regulated review task is itself designed as a framework (Ominyi et al., 2025a), and confidentiality exposure is named as a distinct barrier to enterprise deployment of large language models (Ominyi et al., 2025c). Three legal contributions locate constraints the technical literature does not supply: that

automating a decision creates compliance obligations attaching to the decision logic itself (Annan, 2025a), that demands for algorithmic explanation can conflict with proprietary protection of the system (Annan, 2024), and that ownership of generated output is unsettled in a way that concerns any organisation deploying such systems (Annan, 2023).

The function exercised is F3, and the boundary is pragmatic and semantic, together pragmatic because constraints cost delivery interests, semantic because governance and engineering vocabularies do not map cleanly. The domain's value here lies in the artefact form it has settled on: a classification that can be wrong, an assessment that can be adverse, a policy that can deny, a validation that can fail. That is what Section 4.4 asks of a principle and what an attestation cannot provide. It does not follow that the artefacts are adequate, since a documented process is not an assurance (Costanza-Chock et al., 2022; Mökander et al., 2021).

5.4 Benefit tracing in analytics, automation and service-delivery initiatives

The boundary here runs between a claim about value and evidence that could contradict it: semantic where finance, operations and delivery use the same words for different quantities, pragmatic wherever a measured baseline could later embarrass someone with an interest in the claim. The applied signature of F4 is accordingly the recurrence of attribution as an acknowledged difficulty rather than a solved practice.

Return on a specific investment is treated as a measurement problem in its own right (Sanni et al., 2020a), an analytic capability's contribution as something to be measured rather than asserted (Oduleye & Medon, 2023b), process optimisation as having a financial impact to be assessed rather than presumed (Tonoyan et al., 2022b), and outcome measurement in large service organisations as a distinct modelling problem (Adelanwa et al., 2024), restating in domain language the attribution difficulty the IT value literature established long ago (Kohli & Grover, 2008; Melville et al., 2004). The link between an investment and its intended outcome is elsewhere treated as something to be made measurable rather than assumed visible (Ogbete & Aminu-Ibrahim, 2024), and long-term value from an asset is tracked through a defined set of performance metrics rather than reconstructed at review (Aminu-Ibrahim et al., 2025c). Two patterns bear on Table 3, the first connecting a measurement layer to a reporting audience: performance frameworks are designed with a board in view (Dosunmu & Ogundele, 2025b), a measurement layer is joined explicitly to an executive reporting layer within one governance architecture (Adegbite et al., 2025), indicators are organised into a decision framework for service delivery and portfolio performance rather than left as a dashboard (Edivri & Oteri, 2024), monitoring is automated so that indicators are available continuously rather than assembled periodically (Tonoyan et al., 2024b), and dashboards are built for real-time strategic oversight (Walawalkar et al., 2025), the practitioner form of the insistence that a metric needs an owner and a cadence, not only a definition (Peppard et al., 2007). The second uses quantified risk to prioritise investment (Dosunmu & Ogundele, 2024b), which presupposes the baseline discipline F4 requires; governance and compliance are treated as one integrated strand (Akomolafe & Agu, 2019c), the same argument as instrumenting value inside the pipeline rather than reconstructing it later (Ogundairo & Ayivi-Donkor, 2023b).

The function exercised is F4 and the boundary is semantic and pragmatic. What none of this establishes is the point the benefits literature is clear about: the competences deciding realisation fall after delivery and are the ones organisations omit (Ashurst et al., 2008; Doherty et al., 2012). The applied writing describes the instrument, not the moment at which it must be installed, which is the gap the broker's timing discipline exists to fill.

5.5 Brokering in financial services, audit and risk functions

The boundary in audit and risk functions is between first-line delivery, which owns the process, and second-line challenge, which owns the question. Both speak finance; the divergence is over what counts as evidence, and it is pragmatic, because a finding costs the first line something. It is also the one domain here in which the structure a smaller firm lacks is present and describable.

Applied writing treats internal audit quality as improvable by technology-enabled risk assessment rather than by more inspection (Akomolafe & Agu, 2018), risk-based auditing as a model to be implemented within financial management rather than as an audit-department preference (Akomolafe et al., 2023a), and audit quality as strengthened by analytics applied across institutions (Ozowara et al., 2025). Compliance frameworks are compared across jurisdictions rather than assumed uniform (Agu et al., 2023a), and the contribution closest to F2 in this domain, corrective action planning and the liaison role that carries it are described as a distinct organisational control process rather than as follow-up correspondence (Ebhojie et al., 2023a). On the AI side, compliance monitoring and anomaly detection in high-volume transaction environments are framed as requiring a purpose-built conceptual framework rather than an extension of existing sampling (Ominyi et al., 2024a).

Two functions are exercised rather than one. F3 dominates, the domain's whole business being to convert an obligation into a test that can fail; but the audit-liaison material is F2 in a domain vocabulary, since a corrective action has no force unless someone with standing in the first line admits the finding as binding. The boundary is pragmatic in both cases, and the instructive feature for a smaller firm is that this assurance rests on an organisational separation it cannot afford, which is why Section 4.8 requires the reduction to be stated to the sponsor.

5.6 Brokering in health and clinical service settings

The boundary in clinical settings is unusually sharp because one side holds professional standing the other cannot acquire: a data scientist may understand a clinical pathway and may not adjudicate it. Terms are shared and mean different things, and the interests behind them differ, so the domain runs semantic and pragmatic boundaries at once, with a legitimation problem no amount of added analyst knowledge resolves.

Applied work treats artificial intelligence in clinical decision making as carrying implementation challenges and risks alongside its efficiency case rather than as a straightforward capability gain (Ilodigwe & Adesemoye, 2025a), and organisational readiness for generative AI in healthcare operations as a management capability that varies by context rather than a technical threshold (Afrihyia et al., 2025). Underneath both sits a data problem stated as a governance and architecture problem: interoperability and data quality are framed as requiring an explicit framework rather than emerging from integration (Ilodigwe & Adesemoye, 2021), and analytics over community health data are built on an architecture designed against a specific compliance regime rather than secured afterwards (Aliliele et al., 2024b). Smaller clinical organisations appear too, workflow automation in small practices being described as a process-efficiency framework rather than a technology deployment (Eyetsemitan et al., 2024a).

Three functions are exercised. F1 appears in the interoperability and data-quality material, which is translation between clinical recording practice and analytic representation; F2 is unavoidable, because the professional whose judgement the system informs holds standing the initiative cannot override, and Table 5's pathology of illegitimate legitimation is nowhere more available than here; F3 follows from the regulatory envelope. The boundary is semantic and pragmatic together, and

the domain supplies the clearest instance of the caution that some matters must be routed outward rather than absorbed by the broker.

5.7 Brokering in procurement, supply chain and logistics programmes

The boundary in procurement is first syntactic. Item identifiers, vendor records, unit-of-measure conventions and contract references must denote the same things across systems and organisations before any question of meaning arises; where they do not, an automation reproduces the divergence at speed. Above that substrate sit the semantic boundary, what "delivered," "compliant" or "approved" means to buyer, supplier and auditor, and the pragmatic one.

Applied writing treats process automation in procurement as pursuing transparency alongside efficiency rather than efficiency alone (Akinleye & Adeyoyin, 2021), builds integrated digital platforms so that a supply chain becomes visible to parties who could not previously see it (Okoruwa et al., 2024), and describes decision support as an analytic capability applied to procurement rather than as reporting (Akinleye et al., 2023). The compliance strand is explicit: supplier risk assessment is organised around a regulatory framework in pharmaceutical procurement (Akin-Oluyomi et al., 2023b), and procurement in a high-risk operating environment is framed as needing a regulatory-compliance framework of its own (Okonkwo et al., 2021). Automation is elsewhere described as running on service-management platforms, so that supply-chain work inherits a request-and-routing vocabulary from another discipline (Okonkwo et al., 2024f). The functions exercised are F1 and F3, and the boundary sequence is this section's clearest illustration of Carlile's (2004) cumulative hierarchy: master-data reconciliation is transfer across a shared syntax, agreeing what "compliant delivery" means is translation, and settling whose tolerance governs a borderline case is transformation, since it changes what one party will accept. An initiative answering the third with an instrument built for the first misdiagnoses in precisely the way Section 4.1 predicts.

5.8 Brokering in public sector and infrastructure initiatives

The boundary in public-sector and infrastructure work is between the delivery organisation and a body it neither controls nor can ignore: a regulator, an inspectorate, a funder, or a statutory scheme. The obligation is externally defined and non-negotiable in substance while its application to a design is negotiable, which is why the work concentrates on evidence rather than argument.

Applied writing applies AI to an audit process organised around a defined metric set rather than narrative assessment (Okojie et al., 2023), and analytics to risk detection in public financial systems (Adelanwa et al., 2023b). A statutory-entitlement strand is unusually explicit about the artefact: procedural safeguards and regulatory compliance in special education are reviewed as the substance of service delivery rather than its administration (Oloto & Yeboah, 2024), and risk assessment models are developed for documentation and accountability systems specifically (Oloto & Yeboah, 2025), a domain in which the record is the compliance, which is the condition Section 4.4 describes when documentation is the control. A maturity model for anticipating audit outcomes in resource-constrained regulatory environments is proposed (Adeyelu & Dagodzo, 2024), outcome measurement in large public services is treated as a modelling problem distinct from performance reporting (Adelanwa et al., 2024); and governance between public and private partners appears as in Section 5.2.

F3 and F4 are exercised jointly and, unusually, in that order: the external body defines both constraint and measure, so the benefit definition arrives from outside rather than being negotiated within. The boundary is pragmatic; the regulator's interest is not the delivery organisation's, with

a semantic layer wherever statutory and operational categories diverge. The resource-constrained variant of this domain is the closest applied analogue to the smaller firm's position, describing an organisation obliged to satisfy a regime designed for better-resourced entities.

5.9 Brokering in customer-facing and marketing technology initiatives

The boundary in customer-facing initiatives is between a commercial judgement and a modelled proxy for it. A segment, a propensity or a next-best action restates something a person previously decided, and the restatement is lossy in ways the commercial owner is rarely shown.

Applied work applies predictive models to a routine commercial targeting decision inside a regulated service business (Atima et al., 2022) and argues that the arrival of autonomous agents changes how a business function must be architected rather than merely how it executes (Ogundairo & Ayivi-Donkor, 2024b). The constraint strand is unusually clear in this domain: automation operating under compliance constraints is held to require an explicit control layer rather than a policy statement (Sanni et al., 2022), and lifecycle control over an automation is made auditable through process mining rather than asserted (Sanni et al., 2025). The measurement strand runs alongside, dashboard frameworks being addressed to executive visibility rather than to analyst use (Sanni & Atima, 2021b).

The functions exercised are F1 and F4, with F3 appearing wherever a regulated product is involved. The boundary is predominantly semantic; the gap between commercial construct and modelled proxy is a difference in meaning, not vocabulary, but becomes pragmatic at the point Section 4.3 identifies, when a control layer imposes a cost on the function whose targets the automation was meant to serve. That the domain's own writing reaches for an explicit control layer rather than for better models is what is worth carrying forward.

5.10 Capability building and workforce transition alongside automation

The final boundary is internal and temporal: between the organisation before an automation and the organisation that must operate afterwards. It is the boundary Section 4.3 argues is crossed too late, the people whose work is reallocated being consulted after the design is fixed.

Standard operating procedures are treated as strategic assets in small business operations rather than as compliance documents (Eyetssemitan et al., 2022), a claim about codified knowledge holding value in the organisation rather than in whoever happens to hold it. Improvement method is adapted downward for resource-constrained organisations rather than imported whole (Ambali et al., 2021), and a later synthesis joins procedure design, automation, compliance and change management into one framework for scaling small operations (Eyetssemitan et al., 2025), which is the domain's own recognition that these arrive together rather than in sequence. Regulatory obligations are translated into workable compliance routines for smaller firms specifically (Eyetssemitan et al., 2021), and acceptance testing is framed as structured validation suited to lean operational environments rather than as a scaled-down enterprise practice (Eyetssemitan et al., 2023b), the point at which those who will operate a system are given a defined role in judging it, and the nearest thing in this literature to legitimisation with a named artefact attached.

The functions exercised are F1 and F2: F1 because a procedure is a boundary object in the ordinary sense, carrying practice from those who hold it to those who must enact it and vulnerable to exactly the failure Section 4.2 describes when it records the rule without the rationale; F2 because acceptance testing, done seriously, admits the operator's judgement as binding on the design rather than as feedback. The boundary is semantic where practice must be codified and pragmatic where

the reallocation of work is decided; the domain's most useful signal is that its own literature places procedure design, automation and change management in one frame.

5.11 What the survey does and does not establish

The survey does three things and fails at a fourth. It shows that all four functions have domain vocabularies and recognisable artefacts across settings as different as clinical operations, procurement, statutory entitlement administration and marketing technology; that they arrive unevenly, F3 far better developed than F4, consistent with Section 4.5's argument that benefit tracing is most often left unowned; and that they are described separately, by different specialisms in different venues, with no contribution treating them as one role's work, the gap of Section 2.13, seen from the practitioner side. A further observation follows from the range: the same four functions recur under domain-specific names, offered as no more than weak evidence that the categories are joints in the work rather than impositions of the theory they came from. What the survey does not supply is evidence: each contribution asserts a position, none reports a controlled comparison, the venues carry no established peer-review record, and nothing here bears on whether concentrating the functions in one role improves any outcome.

6. Illustration

This section is illustrative and hypothetical: it describes no organisation and reports no data. The domain is the programme's shared hypothetical case: a mid-sized services organisation handling high-volume, document-heavy casework like intake, eligibility assessment, document verification, decision issuance and regulator reporting. A sponsor proposes to "use AI to speed up assessments", a technology and a direction, not a decision, a data condition, an error tolerance, or an outcome.

F1. The broker decomposes the process into decision points; of those that emerge, such as file completeness, document type and legibility, and eligibility, only the first two suit a first increment. Modelling completeness separately from process flow (Hasić et al., 2018) reproduces the pattern Batoulis et al. (2015) describe: the rules sit in workflow and in staff habits, and terms that seemed unambiguous; "complete," "valid," "current", carry three operational meanings. The requirement then states the decision, the data at the moment of decision, the tolerated error behaviour, (here false rejections are worse than false acceptances, since a rejection reaches the applicant), the escalation path and the evidence of performance, with a requirements-smell pass before circulation (Femmer et al., 2017).

F2. The framing contains sponsor, operations manager and vendor; on the salience typology (Mitchell et al., 1997) verification clerks hold legitimacy but little power, and applicants legitimacy and urgency but no channel. Convening the clerks is justified instrumentally, since they hold the exception knowledge determining whether the automation degrades gracefully, and normatively, since their work is being reallocated. In this hypothetical initiative that changes the design: a borderline document category clerks accept with a note is unrepresented in the proposed rules and would have been rejected outright, joint optimisation in the sense of Trist and Bamforth (1951).

F3. At framing the broker asks the risk-tier question implied by the Act's structure (European Parliament and Council, 2024) and works through Map and Measure (National Institute of Standards and Technology, 2023). The requirement of meaningful human review over adverse

decisions becomes constraints that can fail: no automated rejection without review by a named role; the reviewer sees the extracted evidence, the rule that fired and the confidence indicator, and may override without justification burden; and above a logged, stated override threshold the rule is suspended. That answers Green's (2022) objection, since a reviewer who cannot see what drove the decision supplies legitimation rather than oversight.

F4. Before build a baseline is captured, the proportion of files returned for missing or invalid documents over a defined prior period, with a named owner and a review date. Table 4 shows why: the rework reduction depends on moving completeness checking to submission and on changing applicant communication, which the automation does not perform. Where reporting cannot support the baseline, instrumenting it enters scope (Popovič et al., 2012; Wixom & Watson, 2001), and return is expected to lag deployment (Brynjolfsson et al., 2019).

Two failures are equally available in this hypothetical initiative: an analyst nominated but not credible to operations would not have surfaced the clerks' exception knowledge, leaving a specification formally correct and operationally wrong (Levina & Vaast, 2005); and one attesting to the oversight arrangement rather than specifying it as testable behaviour would have gained a governance artefact and no protection (Green, 2022).

7. Discussion

7.1 Implications for practice

Five implications follow. The brokering functions should be assigned explicitly: the common state is not that they are done badly but that nobody was asked to do them, though naming them is not sufficient (Levina & Vaast, 2005). The analyst's equipment must extend beyond requirements technique: the bodies of knowledge already define a discipline spanning elicitation, analysis, solution evaluation and stakeholder collaboration (International Institute of Business Analysis, 2015; Paul et al., 2020), empirical work on analyst competencies distinguishes what practitioners do from what job descriptions claim (Sonteya & Seymour, 2012), and the strongest adjacent evidence concerns business competence held by technical professionals rather than technical fluency held by business ones (Bassellier & Benbasat, 2004); what changes for AI work is depth in decision modelling and governance literacy sufficient to convert an obligation into a test. Minimum viable control is the right target, since a smaller firm reproducing a large firm's apparatus produces documentation instead of control, and the reduction in assurance should be stated to the sponsor rather than concealed. Measurement precedes automation: the broker's most consequential act is often temporal, insisting on a baseline before the intervention, because it cannot be recovered afterwards (Peppard et al., 2007; Ward & Daniel, 2012). Premature declaration of victory is among the errors that fail transformation efforts (Kotter, 1995, 1996), benefit tracing prevents it, and Trkman (2010) and Rosemann and vom Brocke (2014) locate the point structurally.

The fifth cuts against normal practice: because benefits review and exploitation are the competences most often omitted, and both fall after the delivery organisation has dispersed (Ashurst et al., 2008; Doherty et al., 2012), they must be bought at framing, a review date and a named owner agreed while approval is still being sought. Applied governance work observes compatibly that placing AI assistance inside an existing workflow raises governance, risk and readiness questions at once rather than in a queue (Badmus et al., 2025), a contribution in a venue without an established peer-review record, cited for that observation alone.

7.2 When brokering roles succeed and when they fail

The boundary-spanning literature says more about when such a role works than the framework has so far used. Because the two structural functions can come apart (Aldrich & Herker, 1977), a broker strong on representation and weak on information processing produces advocacy dressed as analysis, and the reverse produces intelligence nobody acts on. Levina and Vaast (2005) add the decisive condition, that the competence emerges in practice and must be recognised as legitimate in both fields, which is why Table 5's mitigation for nominal brokering is participation in the work of both fields rather than a clearer job description.

At the artefact level the conditions are equally specific. Boundary objects work where information requirements genuinely overlap and the movement between common and local versions is performed (Star, 2010); they entrench differences where use across occupational lines is asymmetric (Barrett & Oborn, 2010); and they perform roles that generally belong to different objects (Nicolini et al., 2012), so a specification asked at once to persuade a sponsor, settle what "complete" means and serve as the build instruction will likely fail at two of the three. Failure is also political: settlements reflect the unequal privilege of the domains involved (Oborn & Dawson, 2010), and understanding is reached only where each side rebuilds enough of the other's context for a shared referent (Bechky, 2003). In AI-enabled work, the points of difficulty have been identified empirically as requirements and planning, training-data handling, and product-model integration (Nahar et al., 2022), a map corresponding to F1 and F3.

One qualification bears on all of this: the literature was developed almost entirely in large-firm settings; new product development (Carlile, 2002), a production floor (Bechky, 2003), multidisciplinary and cross-cultural professional teams (Barrett & Oborn, 2010; Oborn & Dawson, 2010), research and development units (Tushman, 1977), organisations large enough for brokering between specialised functions to be necessary at all. Its transfer to a firm where one person writes the requirement, convenes the workshop and takes the baseline is an assumption of this paper, not a finding of that literature; the argument for transfer is that the boundaries persist where the specialised functions do not.

Three conditions can now be stated more precisely. The first is participative: if artefacts carry meaning only against the participation that produced them (Wenger, 1998, 2000), a broker attending neither field's working meetings produces artefacts that read correctly and act on nothing, so Table 5's mitigation for nominal brokering is attendance rather than a clearer mandate. The second is a capacity condition on the receiving side, since what a group takes up depends on what it already knows (Cohen & Levinthal, 1990) and acquisition is not conversion into use (Zahra & George, 2002). The third is temporal: reactions form before deployment and, in automation, on announcement (Asatiani et al., 2020; Oreg et al., 2011); readiness is built across levels rather than crossed once (Rafferty et al., 2013; Weiner, 2009); and the review competences fall after everyone disperses (Ashurst et al., 2008). A role engaging at specification and disengaging at go-live is mistimed at both ends. Brokering also fails more visibly than it succeeds, which should bias retrospective study toward the pathologies.

7.3 The applied settings and the four functions

Section 5 earns two conclusions and forecloses a third. The first is about scope: each function has an applied vocabulary and artefact set of its own configuration and access design, alignment mechanisms, classification and assessment instruments, measurement and reporting frameworks, so the categories are joining the field already uses rather than artefacts of the organisation theory they came from.

The second concerns distribution. The applied material is heavily weighted toward F3, security and compliance practice having converted obligations into failable artefacts for decades; F4 is thinnest, and appears as an acknowledged difficulty rather than a solved practice. That matches Section 4.5's prediction that benefits review and exploitation are omitted because they fall after the delivery organisation disperses (Ashurst et al., 2008; Doherty et al., 2012), and it locates a smaller firm's marginal effort: where a compliance function exists, F3 has a home and the brokering act is one of timing; where it does not, F4 is the function nobody holds and the only one whose absence cannot later be repaired.

The third conclusion the survey cannot support: that the functions are separately described does not show that concentrating them in one role improves outcomes. P2 proposes concentration as a testable claim on which this literature is silent, consistent both with the diagnosis that nobody owns the four and with the alternative that the consolidation proposed here is unnecessary.

7.4 Implications for research and a validation agenda

The framework is unvalidated and offered as claims to be attacked. A multiple-case study across smaller firms that have completed at least one AI-enabled automation would test P1 and P2, coding which functions were performed, by whom and with which artefacts, against independently assessed outcomes; the dependent variable requires care, since self-reported success is unreliable, success and failure are enacted judgements that can coexist for the same system (Cecez-Kecmanovic et al., 2014), and the practitioner surveys of Section 2.11 are self-selected samples reported by advisory firms (Deloitte, 2018; EY, 2016; Telford et al., 2022; Wright et al., 2020), so outcomes should be assessed against each initiative's own pre-stated benefit, measured at process level before any organisational claim is made (Elbashir et al., 2008). That design is what makes P5 testable. A document-analytic study classifying constraints as principle-, checklist- or criterion-stated and tracing each to observable system behaviour would test P4, the empirical form of Green's (2022) argument and Wieringa's (2020) accountability dimensions. A role-level study following Levina and Vaast (2005) would test P3 and P8 by examining where boundary-spanning competence actually emerges, which may not be in the person holding the analyst title; the competency instruments developed for the analyst role provide a starting measurement frame (Bassellier & Benbasat, 2004; Sonteya & Seymour, 2012), and the collaboration points identified for machine-learning-enabled systems provide observation sites (Nahar et al., 2022). A separate study should ask whether the oversight specifications a broker produces alter reviewer behaviour as the human-factors literature predicts (Lee & See, 2004; Skitka et al., 1999), and whether they mitigate or aggravate the degradation of independent judgement observed in human–AI collaboration (Fügener et al., 2021). Comparative work on Table 3 would test P7, and extension work should establish whether the framework transfers across organisational baselines (Sagala & Óri, 2024), using existing readiness instruments (Alsheibani et al., 2018; Baier et al., 2019; Jöhnk et al., 2021; Kinkel et al., 2022; Uren & Edwards, 2023), with the caution that most maturity models here remain empirically unvalidated (Sadiq et al., 2021) and that readiness instruments more broadly have been found conceptually heterogeneous and frequently without demonstrated reliability and validity (Weiner et al., 2008), so an instrument should be selected against the multi-level construct it claims to measure (Rafferty et al., 2013; Weiner, 2009) rather than by availability. One further study follows from Section 5: a venue-stratified review would test whether the distribution observed there, with F3 densely developed and F4 thin, survives direct comparison of the peer-reviewed and applied literatures.

8. Limitations

Eight limitations bear on how the framework should be read. It is conceptual and unvalidated: it reports no primary data and demonstrates nothing, and the appropriate response is empirical challenge rather than adoption.

The central claim that the business analyst is the appropriate broker is argued from role positioning rather than demonstrated. The argument is that the role already sits at the intersection of business intent, data reality, delivery capability and governance obligation, and that no alternative role in a smaller firm sits at all four; that is an argument about structural position, not evidence about performance. The functions may in practice be better performed by an operations manager, a technically fluent sponsor or an external adviser, and Levina and Vaast (2005) give reason to expect competence to emerge where practice puts it rather than where a title places it.

The theoretical base was developed elsewhere. Boundary objects, boundary roles, technology brokering and the boundary hierarchy itself all originate in large-firm research, development and product-development settings with specialised functions and substantial resources (Aldrich & Herker, 1977; Carlile, 2002, 2004; Hargadon & Sutton, 1997; Star & Griesemer, 1989; Tushman, 1977), as was the empirical work extending them (Barrett & Oborn, 2010; Bechky, 2003; Oborn & Dawson, 2010) and the knowledge-transfer results underpinning F1 (Argote & Ingram, 2000; Szulanski, 1996; Wenger, 1998). Their transfer to AI adoption in resource-constrained smaller firms is assumed here, not tested, and the defence that the boundaries persist even where the functions are informal is itself conceptual.

The widened evidence base carries limitations of its own. The human-factors results on which Section 4.4 rests were produced largely in aviation, process control and laboratory judgement tasks (Bainbridge, 1983; Parasuraman & Riley, 1997; Skitka et al., 1999), with participants facing stakes unlike those of an operator whose own work is being reallocated (Dietvorst et al., 2015; Logg et al., 2019), and the review literature holds these results to be framed incompatibly rather than settled (Burton et al., 2020). They establish that oversight has failure modes worth specifying against, not how any particular reviewer will behave. The model risk regime cited was written for regulated financial institutions possessing independent validation functions (Board of Governors of the Federal Reserve System & Office of the Comptroller of the Currency, 2011), which a smaller firm cannot reproduce; and the MLOps and documentation artefacts cited (Kreuzberger et al., 2023; Mitchell et al., 2019) assume organisations operating models continuously at scale, so their downward scaling is proposed rather than observed.

The composition of the evidence base. Somewhat more than a third of the reference base, the whole of Section 5, together with a small number of sources elsewhere, consists of applied and practitioner contributions in venues without established peer-review records, and no theoretical, comparative or quantitative claim in this paper depends on any of it. The expansion of Section 5 into a ten-domain survey has raised that proportion considerably, which sharpens rather than relaxes the discipline: the material carries illustration only. The applied-governance and SME-practice sources outside Section 5 sit in Sections 2.10, 4.4, 4.5 and 7.1 (Badmus et al., 2025; Fobellah, 2025b; Jones & Ominyi, 2025a; Ladapo et al., 2022; Ogundairo & Ayivi-Donkor, 2023b; Ominyi & Anichukwueze, 2023; Ominyi et al., 2024b), none independently verified beyond the position each title asserts. They are cited for illustration rather than evidential weight: no quantitative claim, sample, effect size, proposition or theoretical commitment rests on any of them, and every argument in Section 4 would stand unchanged if Section 5 were removed. The risk is of framing, since writing about a practice can make it look better established than it is. A related

caution applies on the peer-reviewed side, where several bodies of evidence used here are reviews rather than primary studies and inherit the heterogeneity of their primary literatures (Ahmad et al., 2023; Enríquez et al., 2020; Weiner et al., 2008).

The governance instruments cited are recent and moving, being revised and reinterpreted continuously, so specific obligations and thresholds will change; the framework's dependence on any particular instrument is therefore deliberately shallow, resting on the structural insight that a constraint must be expressed so that it can fail a test.

Definitions of "small and medium-sized enterprise" differ across the evidence cited: employment bands of 10–49, 50–249 and 250 or more (Eurostat, 2025b); a 1,000-employee division with IT professionals rather than firms as the unit of observation (IBM, 2023); a firm-level official survey using a different instrument again and reported in a working paper rather than a peer-reviewed article (Bonney et al., 2024); and jurisdiction-specific definitions inside the systematic reviews (Ayinaddis, 2025; Schwaeye et al., 2024). Comparisons across these sources are directional at best. Finally, the framework may describe an idealised role. Performing all four functions competently requires modelling skill, facilitation authority, regulatory literacy and measurement discipline, together with standing to raise costly constraints and time to maintain artefacts. Few analysts in smaller firms are resourced or mandated to do all of this, and the pathologies of responsibility without authority and throughput bottleneck are the predictable result of asking one under-resourced role to carry it. The framework should be read as a specification of what the work requires, against which an organisation can decide how much to fund, rather than a description of what analysts currently do.

9. Conclusion: the programme as a single argument

This paper completes a programme whose parts can now be stated as one claim: organisations do not fail at intelligent process automation because the technology does not work, but because the business problem was never specified in a form a machine could act on, and because nobody owns the translation.

Standardization came first, because a decision point whose logic varies without record cannot be expressed as a requirement, and automating it encodes the variation rather than resolving it. *Readiness* came next, because the constraint in smaller firms is organisational rather than technological, spanning people, process and data alongside technology (Jöhnk et al., 2021; Uren & Edwards, 2023), and must be sequenced against a baseline (Sagala & Öri, 2024). *Specification* came third as the core method, the structured conversion of an ambiguous problem into an AI-ready requirement, and *lifecycle* came fourth, because different gates in that conversion are best carried by different delivery postures. This paper is the fifth, and asks what the others leave open: who does this, and what happens when nobody does.

The answer proposed is that the work is brokering, decomposing into translation, legitimation, constraint injection and benefit tracing. In a large firm each has a structural home; in a smaller firm none does, and the functions are either carried by a role or left undone. Each function rests on a different tradition, and Table 6 records what those traditions supply and what they do not. The value of business analysis in AI transformation, on this account, does not lie in producing documents. It lies in the fact that the specification gap is a boundary phenomenon, that boundary phenomena do not resolve themselves, and that the artefacts a competent analyst produces; a decision model both sides recognise, a record of which claims were admitted and why, a constraint written so it can fail, a baseline captured before anyone had an interest in its value, are boundary

objects in the precise sense: plastic enough to adapt to local needs, robust enough to maintain a common identity across sites (Star & Griesemer, 1989).

What would have to be true for the argument to hold? At minimum: that failures cluster at semantic and pragmatic boundaries rather than technical components (P1); that explicit ownership of the four functions is associated with better realised outcomes than diffuse ownership (P2); that the association depends on competence recognised in practice rather than formal nomination (P3); and that constraints expressed as failable criteria change delivered system behaviour in ways principles and attestations do not (P4). Each is testable; none is tested here. If the first is false, the framework diagnoses the wrong problem; if the second, the functions may be real but their concentration in one role unnecessary; if the third, the apparatus reduces to an organisation-chart change; and if the fourth, the governance function collapses into the documentation it was designed to prevent. The programme's contribution, if it has one, is to have named the gap precisely enough that these questions can be asked.

References

- Adebayo, A., Adegbite, M. P., & Ahmed, M. O. (2022). Adversarial machine learning in critical infrastructure: A conceptual framework for threat modeling AI enabled OT systems. *World Journal of Innovation and Modern Technology*, 6(1), 184–234. <https://doi.org/10.56201/wjimt.v6.no1.2022.pg184.234>
- Adegbite, M. P., Adebayo, A., & Ahmed, M. O. (2022). A security architecture model for IT and OT convergence in regulated energy networks: Design principles and governance alignment. *International Journal of Computer Science and Mathematical Theory*, 8(2), 81–132. <https://doi.org/10.56201/ijcsmt.v8.no2.2022.pg81.132>
- Adegbite, M. P., Adebayo, A., & Ahmed, M. O. (2025). A cyber risk quantification and governance architecture for critical infrastructure: From posture measurement to executive reporting. *International Journal of Computer Science and Mathematical Theory*, 11(12), 232–293. <https://doi.org/10.56201/ijcsmt.vol.11.no12.2025.pg232.293>
- Adelanwa, A., Basnet, A., & Anene, U. N. (2023b). Predictive analytics models for financial risk detection and fraud prevention in public systems. *International Journal of Advanced Multidisciplinary Research and Studies*, 3(6). <https://doi.org/10.62225/2583049X.2023.3.6.5969>
- Adelanwa, A., Basnet, A., & Anene, U. N. (2024). Performance intelligence models for optimization and outcome measurement in large scale public services. *Shodhshauryam, International Scientific Refereed Research Journal*, 6(1).
- Adeyelu, O. O., & Dagodzo, D. (2024). A maturity model for predicting airport safety audit outcomes in resource-constrained regulatory environments. *International Journal of Scientific Research in Civil Engineering*, 8(4), 132–170. <https://doi.org/10.32628/IJSRCE248423>
- Afrihyia, E., Akinse, S. G., & Ojukwu, P. U. (2025). Organizational readiness for generative AI integration in healthcare operations: Comparative management capabilities between the U.S. and low- and middle-income countries. *Iconic Research and Engineering Journals*, 8(10), 1673–1697. <https://doi.org/10.64388/IREV8I10-1714670>
- Agu, M. U., Akomolafe, O., & Bello, A. (2023a). A comparative review of SOX compliance frameworks in cross-border financial auditing. *International Journal of Advanced Multidisciplinary Research and Studies*, 3(6), 2297–2306. <https://doi.org/10.62225/2583049X.2023.3.6.5362>
- Ahmad, K., Abdelrazek, M., Arora, C., Bano, M., & Grundy, J. (2023). Requirements engineering for artificial intelligence systems: A systematic mapping study. *Information and Software Technology*, 158, Article 107176. <https://doi.org/10.1016/j.infsof.2023.107176>
- Akin-Oluyomi, O. T., Atima, M. E., & Akinleye, O. K. (2023b). Regulatory compliance and supplier risk assessment frameworks in international pharmaceutical procurement. *International Journal of Advanced Multidisciplinary Research and Studies*, 3(6), 2194–2204.
- Akinleye, O. K., & Adeyoyin, O. (2021). Process automation framework for enhancing procurement efficiency and transparency. *Shodhshauryam, International Scientific Refereed Research Journal*, 4(4), 356–387.
- Akinleye, O. K., Okoruwa, P. O., Babatope, O. M., & Akokodaripon, D. A. (2023). Leveraging big data and business intelligence for optimization of manufacturing sector procurement. *International Journal of Advanced Multidisciplinary Research and Studies*, 3(6), 2164–2172.

- Akintola, A. S., Fawehinmi, Y. O., Aiyenitaju, O., Chinemerem, B., & Emmanuella, O. (2025). Digital transformation in telehealth: A systematic review of user trust, privacy, and regulatory governance in AI-powered remote monitoring systems. *Journal of Scientific Research and Reports*, 31(11), 163–182. <https://doi.org/10.9734/jsrr/2025/v31i113658>
- Akomolafe, O., & Agu, M. U. (2018). A conceptual model for enhancing internal audit quality through technology-enabled risk assessment frameworks. *Iconic Research and Engineering Journals*, 1(9), 458–475.
- Akomolafe, O., & Agu, M. U. (2019c). Advances in financial resilience through integrated governance and compliance strategies. *Iconic Research and Engineering Journals*, 2(10), 607–620.
- Akomolafe, O., Agu, M. U., & Bello, A. (2023a). A conceptual model for implementing risk-based auditing in strategic financial management. *International Journal of Advanced Multidisciplinary Research and Studies*, 3(6), 2274–2286. <https://doi.org/10.62225/2583049X.2023.3.6.5359>
- Akomolafe, O., Olaogun, B. O., Adesuyi, M. O., Ndukwe, V. U., & Sakyi, J. K. (2025a). Collaborative governance framework for secure cross-border payment data sharing. *International Journal of Advanced Multidisciplinary Research and Studies*, 5(6), 849–865. <https://doi.org/10.62225/2583049X.2025.5.6.5283>
- Aldrich, H., & Herker, D. (1977). Boundary spanning roles and organization structure. *The Academy of Management Review*, 2(2), 217. <https://doi.org/10.2307/257905>
- Aliliele, C., Mbonu, I. S., & Iwuanyanwu, U. (2023a). A conceptual framework for continuous cloud misconfiguration monitoring and enterprise risk mitigation strategies. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 9(10), 373–394. <https://doi.org/10.32628/CSEIT2361071>
- Aliliele, C., Mbonu, I. S., & Iwuanyanwu, U. (2023b). A review of API governance and risk prioritization frameworks in modern financial institutions. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 9(10), 395–433. <https://doi.org/10.32628/CSEIT2361072>
- Aliliele, C., Mbonu, I. S., & Iwuanyanwu, U. (2024a). A conceptual framework for enterprise data sensitivity classification and regulatory traceability mechanisms. *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), 3103–3124. <https://doi.org/10.62225/2583049X.2024.4.6.5991>
- Aliliele, C., Mbonu, I. S., & Iwuanyanwu, U. (2024b). Advances in HIPAA compliant data architecture and secure analytics frameworks for community healthcare organizations. *Shodhshauryam, International Scientific Refereed Research Journal*, 7(2), 277–324. <https://doi.org/10.32628/SHISRRJ2472163>
- Aliliele, C., Mbonu, I. S., Uzoka, E., & Iwuanyanwu, U. (2025a). A review of AI assisted continuous auditing systems in technology risk and cybersecurity oversight. *Gyanshauryam, International Scientific Refereed Research Journal*, 8(4), 210–250. <https://doi.org/10.32628/GISRRJ258369>
- Alsheibani, S., Cheung, Y., & Messom, C. (2018). Artificial intelligence adoption: AI-readiness at firm-level. In *Proceedings of the 22nd Pacific Asia Conference on Information Systems (PACIS 2018)*, Paper 37.
- Amayo, E. B., Owulade, O. A., & Isi, L. R. (2023a). Optimizing project governance in multinational infrastructure projects: Insights from General Electric's global operations.

- International Journal of Multidisciplinary Research and Growth Evaluation*, 4(1), 975–983. <https://doi.org/10.54660/IJMRGE.2023.4.1.975-983>
- Amayo, E. B., Owulade, O. A., & Isi, L. R. (2024b). Effective project governance in multinational infrastructure projects: A case study from General Electric. *Iconic Research and Engineering Journals*, 8(5), 1305–1324.
- Ambali, K. B., Eyetsemitan, R. A., Oyeleye, A. O., & Fadayomi, O. (2021). Lean Six Sigma for small enterprises: A systematic review and Lite-DMAIC adaptation framework for resource-constrained organizations. *Iconic Research and Engineering Journals*, 5(5), 562–583. <https://doi.org/10.64388/IREV5I5-1716957>
- Aminu-Ibrahim, A. Y., Ogbete, J. C., & Ambali, K. B. (2024). Governance and accountability models for public private partnerships in healthcare infrastructure development. *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), 2943–2960. <https://doi.org/10.62225/2583049X.2024.4.6.5699>
- Aminu-Ibrahim, A. Y., Ogbete, J. C., & Iwuanyanwu, O. C. (2025c). Sustainable healthcare infrastructure performance metrics for long-term asset management and value creation. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 11(4), 566–601. <https://doi.org/10.32628/CSEIT251116277>
- Annan, A. O. (2023). Intellectual property ownership and authorship in AI-generated works. *Gyanshauryam, International Scientific Refereed Research Journal*, 6(3), 425–450.
- Annan, A. O. (2024). Algorithmic accountability and trade secret protection in artificial intelligence. *Shodhshauryam, International Scientific Refereed Research Journal*, 7(5), 315–347.
- Annan, A. O. (2025a). Automated decision-making and anti-discrimination compliance under U.S. law. *International Journal of Advanced Multidisciplinary Research and Studies*, 5(2), 2522–2540.
- Aranda, J., Easterbrook, S., & Wilson, G. (2007). Requirements in the wild: How small companies do it. In *15th IEEE International Requirements Engineering Conference (RE 2007)* (pp. 39–48). IEEE. <https://doi.org/10.1109/RE.2007.54>
- Argote, L., & Ingram, P. (2000). Knowledge transfer: A basis for competitive advantage in firms. *Organizational Behavior and Human Decision Processes*, 82(1), 150–169. <https://doi.org/10.1006/obhd.2000.2893>
- Asatiani, A., Penttinen, E., Ruissalo, J., & Salovaara, A. (2020). Knowledge workers' reactions to a planned introduction of robotic process automation: Empirical evidence from an accounting firm. In *Information systems outsourcing* (Progress in IS, pp. 413–452). Springer. https://doi.org/10.1007/978-3-030-45819-5_17
- Ashmore, R., Calinescu, R., & Paterson, C. (2021). Assuring the machine learning lifecycle: Desiderata, methods, and challenges. *ACM Computing Surveys*, 54(5), 1–39. <https://doi.org/10.1145/3453444>
- Ashurst, C., Doherty, N. F., & Peppard, J. (2008). Improving the impact of IT development projects: The benefits realization capability model. *European Journal of Information Systems*, 17(4), 352–370. <https://doi.org/10.1057/ejis.2008.33>
- Atakpa, M. I., & Abolaji, T. O. (2022). A privacy-preserving data architecture model for regulated industry analytics under GDPR and HIPAA compliance. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 8(3), 781–808. <https://doi.org/10.32628/CSEIT22558>

- Atima, M. E., Sanni, J. O., & Attah, A. (2022). Predictive audience segmentation models resolving targeting inefficiencies in regulated professional service enterprises. *Shodhshauryam, International Scientific Refereed Research Journal*, 5(1), 271–303. <https://doi.org/10.32628/SHISRRJ247134>
- Ayinaddis, S. G. (2025). Artificial intelligence adoption dynamics and knowledge in SMEs and large firms: A systematic review and bibliometric analysis. *Journal of Innovation & Knowledge*, 10(3), 100682. <https://doi.org/10.1016/j.jik.2025.100682>
- Badewi, A. (2016). The impact of project management (PM) and benefits management (BM) practices on project success: Towards developing a project benefits governance framework. *International Journal of Project Management*, 34(4), 761–778. <https://doi.org/10.1016/j.ijproman.2015.05.005>
- Badmus, O., Dosunmu, A. A., & Anunagba, C. O. (2025). A governance framework for AI-assisted CRM workflows: Agentforce deployment, risk management, and organizational readiness in enterprise Salesforce environments. *International Journal of Scientific Research in Humanities and Social Sciences*, 2(1), 59–80.
- Badmus, O., Dosunmu, A. A., & Ozowara, D. E. (2018). A systematic review of CI/CD pipeline strategies in Salesforce DevOps: Tools, practices, and deployment outcomes. *Iconic Research and Engineering Journals*, 2(6).
- Badmus, O., Dosunmu, A. A., & Ozowara, D. E. (2019a). A conceptual model for ETL design and data integration in Salesforce-centric enterprise architectures. *Iconic Research and Engineering Journals*, 3(5).
- Badmus, O., Dosunmu, A. A., Ozowara, D. E., & Anunagba, C. O. (2020). A comparative framework for Salesforce DevOps tooling: Evaluating Copado, Flosun, and Salesforce DX across enterprise deployment contexts. *International Journal of Multidisciplinary Research and Growth Evaluation*, 1(5), 1021–1031. <https://doi.org/10.54660/IJMRGE.2020.1.5.1021-1031>
- Badmus, O., Jooda, D., & Anunagba, C. O. (2024). A privacy-by-design framework for role-based security architecture in Salesforce environments handling sensitive personal data. *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), 3244–3254. <https://doi.org/10.62225/2583049X.2024.4.6.6243>
- Badmus, O., Jooda, D., Ozowara, D. E., & Anunagba, C. O. (2022b). A systematic review of MuleSoft ESB integration patterns in multi-cloud Salesforce architectures. *Gyanshauryam, International Scientific Refereed Research Journal*, 5(3), 462–481.
- Baier, L., Jöhren, F., & Seebacher, S. (2019). Challenges in the deployment and operation of machine learning in practice. In *Proceedings of the 27th European Conference on Information Systems (ECIS 2019)*, Research Paper 163.
- Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775–779. [https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8)
- Bano, M., & Zowghi, D. (2015). A systematic review on the relationship between user involvement and system success. *Information and Software Technology*, 58, 148–169. <https://doi.org/10.1016/j.infsof.2014.06.011>
- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671–732. <https://doi.org/10.15779/Z38BG31>
- Barrett, M., & Oborn, E. (2010). Boundary object use in cross-cultural software development teams. *Human Relations*, 63(8), 1199–1221. <https://doi.org/10.1177/0018726709355657>

- Bassellier, G., & Benbasat, I. (2004). Business competence of information technology professionals: Conceptual development and influence on IT–business partnerships. *MIS Quarterly*, 28(4), 673–694. <https://doi.org/10.2307/25148659>
- Batoulis, K., Meyer, A., Bazhenova, E., Decker, G., & Weske, M. (2015). Extracting decision logic from process models. In *Advanced Information Systems Engineering (CAiSE 2015)* (LNCS Vol. 9097, pp. 349–366). Cham: Springer. https://doi.org/10.1007/978-3-319-19069-3_22
- Bechky, B. A. (2003). Sharing meaning across occupational communities: The transformation of understanding on a production floor. *Organization Science*, 14(3), 312–330. <https://doi.org/10.1287/orsc.14.3.312.15162>
- Bhatt, U., Xiang, A., Sharma, S., Weller, A., Taly, A., Jia, Y., Ghosh, J., Puri, R., Moura, J. M. F., & Eckersley, P. (2020). Explainable machine learning in deployment. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 648–657). Association for Computing Machinery. <https://doi.org/10.1145/3351095.3375624>
- Board of Governors of the Federal Reserve System & Office of the Comptroller of the Currency. (2011). *Supervisory guidance on model risk management* (SR Letter 11-7). Board of Governors of the Federal Reserve System. <https://www.federalreserve.gov/boarddocs/srletters/2011/sr1107.htm>
- Bonney, K., Breaux, C., Buffington, C., Dinlersoz, E., Foster, L., Goldschlag, N., Haltiwanger, J., Kroff, Z., & Savage, K. (2024). *Tracking firm use of AI in real time: A snapshot from the Business Trends and Outlook Survey* (CES Working Paper No. CES-24-16). US Census Bureau, Center for Economic Studies. <https://www.census.gov/library/working-papers/2024/adrm/CES-WP-24-16.html>
- Bourne, M., Mills, J., Wilcox, M., Neely, A., & Platts, K. (2000). Designing, implementing and updating performance measurement systems. *International Journal of Operations & Production Management*, 20(7), 754–771. <https://doi.org/10.1108/01443570010330739>
- Brown, J. S., & Duguid, P. (1991). Organizational learning and communities-of-practice: Toward a unified view of working, learning, and innovation. *Organization Science*, 2(1), 40–57. <https://doi.org/10.1287/orsc.2.1.40>
- Brynjolfsson, E. (1993). The productivity paradox of information technology. *Communications of the ACM*, 36(12), 66–77. <https://doi.org/10.1145/163298.163309>
- Brynjolfsson, E., & Hitt, L. M. (2000). Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4), 23–48. <https://doi.org/10.1257/jep.14.4.23>
- Brynjolfsson, E., Rock, D., & Syverson, C. (2019). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 23–60). Chicago: University of Chicago Press. <https://doi.org/10.7208/chicago/9780226613475.003.0001>
- Burton, J. W., Stein, M.-K., & Jensen, T. B. (2020). A systematic review of algorithm aversion in augmented decision making. *Journal of Behavioral Decision Making*, 33(2), 220–239. <https://doi.org/10.1002/bdm.2155>
- By, R. T. (2005). Organisational change management: A critical review. *Journal of Change Management*, 5(4), 369–380. <https://doi.org/10.1080/14697010500359250>
- Carlile, P. R. (2002). A pragmatic view of knowledge and boundaries: Boundary objects in new product development. *Organization Science*, 13(4), 442–455. <https://doi.org/10.1287/orsc.13.4.442.2953>

- Carlile, P. R. (2004). Transferring, translating, and transforming: An integrative framework for managing knowledge across boundaries. *Organization Science*, 15(5), 555–568. <https://doi.org/10.1287/orsc.1040.0094>
- Cecez-Kecmanovic, D., Kautz, K., & Abrahall, R. (2014). Reframing success and failure of information systems: A performative perspective. *MIS Quarterly*, 38(2), 561–588. <https://doi.org/10.25300/misq/2014/38.2.11>
- Chakraborti, T., Isahagian, V., Khalaf, R., Khazaeni, Y., Muthusamy, V., Rizk, Y., & Unuvar, M. (2020). From robotic process automation to intelligent process automation: Emerging trends. In *Business process management: Blockchain and robotic process automation forum* (Lecture Notes in Business Information Processing, pp. 215–228). Springer. https://doi.org/10.1007/978-3-030-58779-6_15
- Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128–152. <https://doi.org/10.2307/2393553>
- Cooper, L. A., Holderness, D. K., Jr., Sorensen, T. L., & Wood, D. A. (2019). Robotic process automation in public accounting. *Accounting Horizons*, 33(4), 15–35. <https://doi.org/10.2308/acch-52466>
- Costanza-Chock, S., Raji, I. D., & Buolamwini, J. (2022). Who audits the auditors? Recommendations from a field scan of the algorithmic auditing ecosystem. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1571–1583). Association for Computing Machinery. <https://doi.org/10.1145/3531146.3533213>
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340. <https://doi.org/10.2307/249008>
- Deloitte. (2018). *The robots are ready. Are you? Untapped advantage in your digital workforce* (Deloitte Global RPA Survey). Deloitte LLP.
- DeLone, W. H., & McLean, E. R. (1992). Information systems success: The quest for the dependent variable. *Information Systems Research*, 3(1), 60–95. <https://doi.org/10.1287/isre.3.1.60>
- DeLone, W. H., & McLean, E. R. (2003). The DeLone and McLean model of information systems success: A ten-year update. *Journal of Management Information Systems*, 19(4), 9–30. <https://doi.org/10.1080/07421222.2003.11045748>
- Diakopoulos, N. (2015). Algorithmic accountability: Journalistic investigation of computational power structures. *Digital Journalism*, 3(3), 398–415. <https://doi.org/10.1080/21670811.2014.976411>
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126. <https://doi.org/10.1037/xge0000033>
- Doherty, N. F., & King, M. (2001). An investigation of the factors affecting the successful treatment of organisational issues in systems development projects. *European Journal of Information Systems*, 10(3), 147–160. <https://doi.org/10.1057/palgrave.ejis.3000401>
- Doherty, N. F., Ashurst, C., & Peppard, J. (2012). Factors affecting the successful realisation of benefits from systems development projects: Findings from three case studies. *Journal of Information Technology*, 27(1), 1–16. <https://doi.org/10.1057/jit.2011.8>
- Donaldson, T., & Preston, L. E. (1995). The stakeholder theory of the corporation: Concepts, evidence, and implications. *Academy of Management Review*, 20(1), 65–91. <https://doi.org/10.5465/amr.1995.9503271992>

- Dosunmu, A. A., & Ogundele, P. O. (2024a). Breach and attack simulation frameworks for continuous validation of enterprise security controls. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 10(3), 1100–1119.
- Dosunmu, A. A., & Ogundele, P. O. (2024b). Cyber risk quantification models for prioritizing enterprise security investment decisions. *International Journal of Multidisciplinary Research and Growth Evaluation*, 5(6), 1777–1785. <https://doi.org/10.54660/IJMRGE.2024.5.6.1777-1785>
- Dosunmu, A. A., & Ogundele, P. O. (2025b). Cyber defense performance measurement frameworks for executive and board level governance. *Computer Science and IT Research Journal*, 6(11), 895–913. <https://doi.org/10.51594/csitrj.v6i11.2164>
- Ebhojie, O., Dogbatsey, E. A., & Oyeleye, A. O. (2023a). Audit liaison, corrective action planning, and control compliance in multinational organisations: A systematic literature review. *International Journal of Advanced Multidisciplinary Research and Studies*, 3(6), 2839–2850. <https://doi.org/10.62225/2583049X.2023.3.6.6200>
- Edivri, J., & Oteri, O. (2024). A conceptual KPI-driven decision and optimization framework for IT service delivery, portfolio performance, and adoption. *Gyanshauryam, International Scientific Refereed Research Journal*, 7(1), 167–187. <https://doi.org/10.32628/GISRRJ246643>
- Elbashir, M. Z., Collier, P. A., & Davern, M. J. (2008). Measuring the effects of business intelligence systems: The relationship between business process and organizational performance. *International Journal of Accounting Information Systems*, 9(3), 135–153. <https://doi.org/10.1016/j.accinf.2008.03.001>
- Eller, R., Alford, P., Kallmünzer, A., & Peters, M. (2020). Antecedents, consequences, and challenges of small and medium-sized enterprise digitalization. *Journal of Business Research*, 112, 119–127. <https://doi.org/10.1016/j.jbusres.2020.03.004>
- Enríquez, J. G., Jiménez-Ramírez, A., Domínguez-Mayo, F. J., & García-García, J. A. (2020). Robotic process automation: A scientific and industrial systematic mapping study. *IEEE Access*, 8, 39113–39129. <https://doi.org/10.1109/access.2020.2974934>
- European Parliament and Council. (2024). Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L, 12 July 2024. <http://data.europa.eu/eli/reg/2024/1689/oj>
- Eurostat. (2025a). 20% of EU enterprises use AI technologies. Eurostat News, 11 December 2025. <https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20251211-2>
- Eurostat. (2025b). Use of artificial intelligence in enterprises. Statistics Explained. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Useofartificialintelligencein_enterprises
- EY. (2016). *Get ready for robots: Why planning makes the difference between success and disappointment*. EYGM Limited.
- Eyetsemitan, R. A., Ambali, K. B., Oyeleye, A. O., & Fadayomi, O. (2020). Multi-stakeholder governance alignment in joint venture operations: A conceptual framework for coordinating business processes in highly regulated environments. *Iconic Research and Engineering Journals*, 4(4), 418–441. <https://doi.org/10.64388/IREV4I4-1716955>
- Eyetsemitan, R. A., Ambali, K. B., Oyeleye, A. O., & Fadayomi, O. (2021). Translating tax and regulatory requirements into SME compliance workflows: A conceptual framework for

- implementing IAS 12, VAT, PAYE, and withholding tax. *Iconic Research and Engineering Journals*, 4(11), 621–641. <https://doi.org/10.64388/IREV4I11-1716956>
- Eyetsemitan, R. A., Ambali, K. B., Oyeleye, A. O., & Fadayomi, O. (2023b). User acceptance testing in small business technology deployment: A structured validation framework for lean operational environments. *International Journal of Multidisciplinary Research and Growth Evaluation*, 4(6), 1512–1531. <https://doi.org/10.54660/IJMRGE.2023.4.6.1512-1531>
- Eyetsemitan, R. A., Ambali, K. B., Oyeleye, A. O., & Fadayomi, O. (2025). An integrated lean-digital framework for scaling small business operations: Synthesizing SOP design, automation, compliance, and change management. *International Journal of Multidisciplinary Research and Growth Evaluation*, 6(6), 1341–1360. <https://doi.org/10.54660/IJMRGE.2025.6.6.1341-1360>
- Eyetsemitan, R. A., Oyeleye, A. O., Ambali, K. B., & Fadayomi, O. (2022). Standard operating procedures as strategic assets in small business operations: A systematic review and implementation framework. *Gyanshauryam, International Scientific Refereed Research Journal*, 5(2), 438–465. <https://doi.org/10.32628/GISRRJ225356>
- Eyetsemitan, R. A., Oyeleye, A. O., Ambali, K. B., & Fadayomi, O. (2024a). CRM and workflow automation in small healthcare practices: A process efficiency framework for scalable patient engagement. *International Journal of Multidisciplinary Research and Growth Evaluation*, 5(6), 1931–1949. <https://doi.org/10.54660/IJMRGE.2024.5.6.1931-1949>
- Femmer, H., Méndez Fernández, D., Wagner, S., & Eder, S. (2017). Rapid quality assurance with requirements smells. *Journal of Systems and Software*, 123, 190–213. <https://doi.org/10.1016/j.jss.2016.02.047>
- Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ... Vayena, E. (2018). AI4People: An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- Fobellah, A. N. (2025b). Navigating digital transformation: Auditing artificial intelligence-powered financial systems: A conceptual review. *International Journal of Science and Research Archive*, 16(2), 23–28. <https://doi.org/10.30574/ijsra.2025.16.2.2274>
- Freeman, R. E. (1984). *Strategic management: A stakeholder approach*. Boston: Pitman.
- Fügener, A., Grahl, J., Gupta, A., & Ketter, W. (2021). Will humans-in-the-loop become borgs? Merits and pitfalls of working with AI. *MIS Quarterly*, 45(3), 1527–1556. <https://doi.org/10.25300/misq/2021/16553>
- Geburu, T., Morgenstern, J., Vecchione, B., Wortman Vaughan, J., Wallach, H., Daumé, H., III, & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92. <https://doi.org/10.1145/3458723>
- Gotel, O. C. Z., & Finkelstein, A. C. W. (1994). An analysis of the requirements traceability problem. In *Proceedings of the IEEE International Conference on Requirements Engineering (ICRE '94)* (pp. 94–101). IEEE. <https://doi.org/10.1109/ICRE.1994.292398>
- Green, B. (2022). The flaws of policies requiring human oversight of government algorithms. *Computer Law & Security Review*, 45, 105681. <https://doi.org/10.1016/j.clsr.2022.105681>
- Grønsund, T., & Aanestad, M. (2020). Augmenting the algorithm: Emerging human-in-the-loop work configurations. *The Journal of Strategic Information Systems*, 29(2), 101614. <https://doi.org/10.1016/j.jsis.2020.101614>

- Hagendorff, T. (2020). The ethics of AI ethics: An evaluation of guidelines. *Minds and Machines*, 30(1), 99–120. <https://doi.org/10.1007/s11023-020-09517-8>
- Hargadon, A., & Sutton, R. I. (1997). Technology brokering and innovation in a product development firm. *Administrative Science Quarterly*, 42(4), 716–749. <https://doi.org/10.2307/2393655>
- Hasić, F., De Smedt, J., & Vanthienen, J. (2018). Augmenting processes with decision intelligence: Principles for integrated modelling. *Decision Support Systems*, 107, 1–12. <https://doi.org/10.1016/j.dss.2017.12.008>
- Herm, L.-V., Janiesch, C., Helm, A., Imgrund, F., Hofmann, A., & Winkelmann, A. (2023). A framework for implementing robotic process automation projects. *Information Systems and e-Business Management*, 21(1), 1–35. <https://doi.org/10.1007/s10257-022-00553-8>
- High-Level Expert Group on Artificial Intelligence. (2019). *Ethics Guidelines for Trustworthy AI*. Brussels: European Commission. <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>
- Iacovou, C. L., Benbasat, I., & Dexter, A. S. (1995). Electronic data interchange and small organizations: Adoption and impact of technology. *MIS Quarterly*, 19(4), 465–485. <https://doi.org/10.2307/249629>
- IBM. (2023). *IBM Global AI Adoption Index 2023: Enterprise report*. Armonk, NY: IBM.
- Ike, P. N., Aifuwa, S. E., Nnabueze, S. B., Olatunde-Thorpe, J., Ogbuefi, E., Oshoba, T. O., & Akokodaripon, D. (2024a). Quantitative risk architecture for public-private partnerships: A multi-layered model for allocating public and private risk. *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), 2669–2682. <https://doi.org/10.62225/2583049X.2024.4.6.5021>
- Ike, P. N., Ogbuefi, E., Nnabueze, S. B., Olatunde-Thorpe, J., Aifuwa, S. E., Oshoba, T. O., & Akokodaripon, D. (2021). Supplier relationship management strategies fostering innovation, collaboration, and resilience in global supply chain ecosystems. *International Journal of Multidisciplinary Evolutionary Research*, 2(2), 52–62. <https://doi.org/10.54660/IJMER.2021.2.2.52-62>
- Ilodigwe, L., & Adesemoye, A. C. (2021). The data backbone of health system transformation: A governance and architecture framework for interoperability, data quality, and advanced analytics at national scale. *International Journal of Health and Pharmaceutical Research*, 6(2), 52–76. <https://doi.org/10.56201/ijhpr.vol.6.no2.2021.pg52.76>
- Ilodigwe, L., & Adesemoye, A. C. (2025a). Advances, risks, and implementation challenges of artificial intelligence as a force multiplier for clinical decision making and health system efficiency. *International Journal of Medical Evaluation and Physical Report*, 9(7), 165–185. <https://doi.org/10.56201/ijmepr.v9.no7.2025.pg165.185>
- Ilodigwe, L., & Adesemoye, A. C. (2025b). Executing healthcare transformation at scale: A strategic change model derived from large programs across payers, providers, and integrated health systems. *International Journal of Health and Pharmaceutical Research*, 10(12), 253–272. <https://doi.org/10.56201/ijhpr.vol.10.no12.2025.pg253.272>
- Institute of Electrical and Electronics Engineers. (2021). *IEEE standard model process for addressing ethical concerns during system design* (IEEE Std 7000-2021). IEEE. <https://standards.ieee.org/ieee/7000/6781/>
- International Institute of Business Analysis. (2015). *A Guide to the Business Analysis Body of Knowledge (BABOK Guide)* (Version 3). Toronto: International Institute of Business Analysis.

- International Organization for Standardization. (2018). *ISO/IEC/IEEE 29148:2018 Systems and software engineering Life cycle processes: Requirements engineering* (2nd ed.). Geneva: ISO. <https://www.iso.org/standard/72089.html>
- International Organization for Standardization. (2023a). *ISO/IEC 23894:2023 Information technology Artificial intelligence Guidance on risk management*. Geneva: ISO. <https://www.iso.org/standard/77304.html>
- International Organization for Standardization. (2023b). *ISO/IEC 42001:2023 Information technology Artificial intelligence Management system*. Geneva: ISO. <https://www.iso.org/standard/42001>
- Isiekwu, C. P., Oluwo, K., & Dada, T. (2021). A conceptual model for CFO-led strategic finance in joint venture and cross-border partnerships. *Iconic Research and Engineering Journals*, 4(7), 383–400. <https://doi.org/10.64388/IREV4I7-1714350>
- Isiekwu, C. P., Oluwo, K., & Dada, T. (2025). A post-pandemic strategic collaboration model for banks and capital markets: Addressing resilience, recovery, and growth challenges. *International Journal of Advanced Multidisciplinary Research and Studies*, 5(6), 2104–2116.
- Jimoh, H. O., Ahmed, M. O., & Fagbade, M. O. (2023). The role of frameworks in cybersecurity governance. *Journal of Behavioural Informatics, Digital Humanities and Development Research*, 9(4), 7–16. <https://doi.org/10.22624/AIMS/BHI/V9N4P2>
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- Jöhnk, J., Weißert, M., & Wyrski, K. (2021). Ready or not, AI comes: An interview study of organizational AI readiness factors. *Business & Information Systems Engineering*, 63(1), 5–20. <https://doi.org/10.1007/s12599-020-00676-7>
- Jones, B. C., & Ominyin, M. (2025a). Conceptualizing AI-driven advisory services for expanding trade compliance access to small and medium enterprises. *IIARD International Journal of Economics and Business Management*, 11(12), 441–474. <https://doi.org/10.56201/ijebm.vol.11.no12.2025.pg441.474>
- Kaplan, R. S., & Norton, D. P. (1992). The balanced scorecard: Measures that drive performance. *Harvard Business Review*, January–February 1992.
- Kaplan, R. S., & Norton, D. P. (1996). *The balanced scorecard: Translating strategy into action*. Boston, MA: Harvard Business School Press.
- Kaur, D., Uslu, S., Rittichier, K. J., & Durresi, A. (2022). Trustworthy artificial intelligence: A review. *ACM Computing Surveys*, 55(2), 1–38. <https://doi.org/10.1145/3491209>
- Kinkel, S., Baumgartner, M., & Cherubini, E. (2022). Prerequisites for the adoption of AI technologies in manufacturing: Evidence from a worldwide sample of manufacturing companies. *Technovation*, 110, 102375. <https://doi.org/10.1016/j.technovation.2021.102375>
- Kohli, R., & Grover, V. (2008). Business value of IT: An essay on expanding research directions to keep up with the times. *Journal of the Association for Information Systems*, 9(1), Article 1. <https://doi.org/10.17705/1jais.00147>
- Komi, N. M. (2023). When accurate forecasts are not fair forecasts: Deep learning load prediction and equity across districts. *Shodhshauryam, International Scientific Refereed Research Journal*, 6(6), 476–546. <https://doi.org/10.32628/SHISRRJ247102>

- Komi, N. M. (2024). Turning social license into a measurable variable: Predicting community trust in renewable energy development. *International Journal of Engineering and Modern Technology*, 10(11), 239–296. <https://doi.org/10.56201/ijemt.v10.no11.2024.pg239.296>
- Kotter, J. P. (1995). Leading change: Why transformation efforts fail. *Harvard Business Review*, March–April 1995.
- Kotter, J. P. (1996). *Leading change*. Boston, MA: Harvard Business School Press.
- Kreuzberger, D., Köhl, N., & Hirschl, S. (2023). Machine learning operations (MLOps): Overview, definition, and architecture. *IEEE Access*, 11, 31866–31879. <https://doi.org/10.1109/ACCESS.2023.3262138>
- Lacity, M. C., & Willcocks, L. P. (2016a). A new approach to automating services. *MIT Sloan Management Review*, 58(1), 41–49.
- Lacity, M., & Willcocks, L. (2016b). Robotic process automation at Telefónica O2. *MIS Quarterly Executive*, 15(1), 21–37.
- Lacity, M., Willcocks, L., & Gozman, D. (2021). Influencing information systems practice: The action principles approach applied to robotic process and cognitive automation. *Journal of Information Technology*, 36(3), 216–240. <https://doi.org/10.1177/0268396221990778>
- Ladapo, O. O., Dosunmu, A. A., Jooda, D., & Abolaji, T. O. (2022). Human-in-the-loop machine learning: A state of the art. *Journal of Frontiers in Multidisciplinary Research*, 3(1), 656–669. <https://doi.org/10.54660/JFMR.2022.3.1.656-669>
- Ladapo, O. O., Dosunmu, A. A., Jooda, D., & Abolaji, T. O. (2025). Migration of applications and information systems to cloud computing infrastructure: Lessons from a South African retail bank. *International Journal of Multidisciplinary Research and Growth Evaluation*, 6(6), 1361–1375. <https://doi.org/10.54660/IJMRGE.2025.6.6.1361-1375>
- Ladapo, O. O., Jooda, D., Dosunmu, A. A., & Abolaji, T. O. (2023a). A comprehensive survey on ServiceNow for IT service management. *International Journal of Multidisciplinary Research and Growth Evaluation*, 4(6), 1532–1545. <https://doi.org/10.54660/IJMRGE.2023.4.6.1532-1545>
- Ladapo, O. O., Jooda, D., Dosunmu, A. A., & Abolaji, T. O. (2023b). Systematic literature review on security access control policies and techniques based on privacy requirements in a BYOD environment. *International Journal of Advanced Multidisciplinary Research and Studies*, 3(6), 2890–2904. <https://doi.org/10.62225/2583049X.2023.3.6.6206>
- Ladapo, O. O., Jooda, D., Dosunmu, A. A., & Abolaji, T. O. (2024). Keeping humans in the loop: Human-centered automated annotation with generative AI. *International Journal of Multidisciplinary Futuristic Development*, 5(1), 81–95. <https://doi.org/10.54660/IJMFD.2024.5.1.81-95>
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392
- Levina, N., & Vaast, E. (2005). The emergence of boundary spanning competence in practice: Implications for implementation and use of information systems. *MIS Quarterly*, 29(2), 335–363.
- Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90–103. <https://doi.org/10.1016/j.obhdp.2018.12.005>
- Mäntymäki, M., Minkkinen, M., Birkstedt, T., & Viljanen, M. (2022). Defining organizational AI governance. *AI and Ethics*, 2(4), 603–609. <https://doi.org/10.1007/s43681-022-00143-x>

- Mbonu, I. S., Aliliele, C., Iwuanyanwu, U., & Oluoha, O. M. (2018). A conceptual framework for legal and ethical risk modeling in enterprise data protection governance systems. *Iconic Research and Engineering Journals*, 2(2), 207–226. <https://doi.org/10.64388/IREV212-1714911>
- Mbonu, I. S., Aliliele, C., Uzoka, E., & Oluoha, O. M. (2019). A review of comparative data protection regulations and secure cloud implementation strategies across jurisdictions. *Iconic Research and Engineering Journals*, 2(9), 482–501. <https://doi.org/10.64388/IREV219-1714912>
- Mbonu, I. S., Iwuanyanwu, U., Aliliele, C., & Uzoka, E. (2022a). A review of data protection impact assessment models in multi cloud financial infrastructure systems. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 8(1), 589–623. <https://doi.org/10.32628/CSEIT25442>
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35. <https://doi.org/10.1145/3457607>
- Melville, N., Kraemer, K., & Gurbaxani, V. (2004). Review: Information technology and organizational performance: An integrative model of IT business value. *MIS Quarterly*, 28(2), 283–322.
- Meyer, M. (2010). The rise of the knowledge broker. *Science Communication*, 32(1), 118–127. <https://doi.org/10.1177/1075547009359797>
- Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 220–229). Association for Computing Machinery. <https://doi.org/10.1145/3287560.3287596>
- Mitchell, R. K., Agle, B. R., & Wood, D. J. (1997). Toward a theory of stakeholder identification and salience: Defining the principle of who and what really counts. *Academy of Management Review*, 22(4), 853–886. <https://doi.org/10.5465/amr.1997.9711022105>
- Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1(11), 501–507. <https://doi.org/10.1038/s42256-019-0114-4>
- Mökander, J., Morley, J., Taddeo, M., & Floridi, L. (2021). Ethics-based auditing of automated decision-making systems: Nature, scope, and limitations. *Science and Engineering Ethics*, 27(4), Article 44. <https://doi.org/10.1007/s11948-021-00319-4>
- Nahar, N., Zhou, S., Lewis, G. A., & Kästner, C. (2022). Collaboration challenges in building ML-enabled systems: Communication, documentation, engineering, and process. In *Proceedings of the 44th International Conference on Software Engineering (ICSE '22)* (pp. 413–425). ACM. <https://doi.org/10.1145/3510003.3510209>
- National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (NIST AI 100-1). Gaithersburg, MD: NIST. <https://doi.org/10.6028/NIST.AI.100-1>
- Neely, A., Gregory, M., & Platts, K. (1995). Performance measurement system design: A literature review and research agenda. *International Journal of Operations & Production Management*, 15(4), 80–116. <https://doi.org/10.1108/01443579510083622>
- Nicolini, D., Mengis, J., & Swan, J. (2012). Understanding the role of objects in cross-disciplinary collaboration. *Organization Science*, 23(3), 612–629. <https://doi.org/10.1287/orsc.1110.0664>

- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37. <https://doi.org/10.1287/orsc.5.1.14>
- Nwakamma, S., Ojukwu, J., & Oyesiji, S. O. (2024b). Securing agentic AI enterprise workflows against prompt injection, tool poisoning, memory manipulation, and excessive agency threats. *World Journal of Innovation and Modern Technology*, 8(6), 217–248. <https://doi.org/10.56201/wjimt.v8.no6.2024.pg217.248>
- Object Management Group. (2019). *Semantics of Business Vocabulary and Business Rules (SBVR), Version 1.5* (OMG Document formal/19-10-02). <https://www.omg.org/spec/SBVR/>
- Object Management Group. (2024). *Decision Model and Notation (DMN), Version 1.5* (OMG Document formal/24-01-01). <https://www.omg.org/spec/DMN/1.5/>
- Oborn, E., & Dawson, S. (2010). Knowledge and practice in multidisciplinary teams: Struggle, accommodation and privilege. *Human Relations*, 63(12), 1835–1857. <https://doi.org/10.1177/0018726710371237>
- Oduleye, T. E., & Medon, J. J. (2023b). A quantitative model for measuring the strategic impact of financial analysis on enterprise growth. *International Journal of Advanced Multidisciplinary Research and Studies*, 3(1), 1663–1672. <https://doi.org/10.62225/2583049X.2023.3.1.5334>
- OECD. (2019). *Recommendation of the Council on Artificial Intelligence* (OECD/LEGAL/0449). Paris: OECD. <https://legalinstruments.oecd.org/en/instruments/oecd-legal-0449>
- OECD. (2025). *The adoption of artificial intelligence in firms: New evidence for policymaking*. Paris: OECD Publishing. <https://doi.org/10.1787/f9ef33c3-en>
- Ogbete, J. C., & Aminu-Ibrahim, A. Y. (2024). Translating healthcare infrastructure investment into measurable population health and diagnostic outcomes. *International Journal of Scientific Research in Humanities and Social Sciences*, 1(2), 955–985. <https://doi.org/10.32628/IJSRSSH242775>
- Ogbole, J. I., Okoruwa, P. O., Fadayomi, O., Abolaji, T. O., Edivri, J., & Akeju, B. (2021). Conceptual model for identity-centric zero trust architecture in enterprise security governance. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 7(5), 393–415. <https://doi.org/10.32628/IJSRCSEIT217562>
- Ogundairo, K. M. (2025). Hardening the autonomous value chain: Lean Six Sigma guardrails and multi-enterprise ETL architecture for agentic supply chain orchestration. *IIARD International Journal of Economics and Business Management*, 11(12), 441–478. <https://doi.org/10.56201/ijebm.vol.11.no12.2025.pg441.478>
- Ogundairo, K. M., & Ayivi-Donkor, S. S. (2023b). Telemetry-driven value realization: Infusing economic evaluation algorithms directly into enterprise data pipelines. *International Journal of Economics and Financial Management*, 8(8), 195–232. <https://doi.org/10.56201/ijefm.v8.no8.2023.pg195.232>
- Ogundairo, K. M., & Ayivi-Donkor, S. S. (2024b). The vanishing interface: Why the coming era of autonomous agents demands a paradigm shift in product marketing architecture. *IIARD International Journal of Economics and Business Management*, 10(11), 350–385. <https://doi.org/10.56201/ijebm.v10.no11.2024.pg350.385>
- Okojie, J. S., Filani, O. M., Ike, P. N., Okojokwu-Idu, J. O., Nnabueze, S. B., & Ihwughwavwe, S. I. (2023). Integrating AI with ESG metrics in smart infrastructure auditing for high-impact

- urban development projects. *International Journal of Multidisciplinary Futuristic Development*, 4(1), 32–44. <https://doi.org/10.54660/IJMFD.2023.4.1.32-44>
- Okonkwo, C. S., Agbabiaka, J., Mayo, W., & Okeke, O. T. (2024a). Conceptual framework for digital supply chain governance in energy and infrastructure sectors. *Gyanshauryam, International Scientific Refereed Research Journal*, 7(4), 335–356. <https://doi.org/10.32628/GISRRJ247423>
- Okonkwo, C. S., Agbabiaka, J., Mayo, W., & Okeke, O. T. (2024f). Supply chain automation framework using service management platforms. *Shodhshauryam, International Scientific Refereed Research Journal*, 7(2), 157–177. <https://doi.org/10.32628/SHISRRJ2472157>
- Okonkwo, C. S., Ahiaekwe Patrick, M. C., Okeke, O. T., & Mayo, W. (2023). Framework for integrating IT systems engineering with supply chain operations. *International Journal of Advanced Multidisciplinary Research and Studies*, 3(6), 2580–2589. <https://doi.org/10.62225/2583049X.2023.3.6.5500>
- Okonkwo, C. S., Ogunwale, O., Mayo, W., & Okeke, O. T. (2021). Framework for regulatory-compliant procurement in high-risk energy environments. *International Journal of Multidisciplinary Research and Growth Evaluation*, 2(6), 595–605. <https://doi.org/10.54660/IJMRGE.2021.2.6.595-605>
- Okoruwa, P. O., Babatope, O. M., Akokodaripon, D. A., & Akinleye, O. K. (2024). Developing integrated digital platforms for enhancing transparency in procurement and supply chain management. *International Journal of Multidisciplinary Research and Growth Evaluation*, 5(6), 1719–1729.
- Oloto, A. M., & Yeboah, T. J. (2024). A critical review of procedural safeguards and regulatory compliance in special education programs. *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), 3125–3137.
- Oloto, A. M., & Yeboah, T. J. (2025). Developing conceptual risk assessment models for special education documentation and accountability systems. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 11(4), 691–734. <https://doi.org/10.32628/CSEIT251116283>
- Ominyi, M., & Anichukwueze, C. C. (2023). A review of algorithmic accountability and model risk management approaches in financial services. *Journal of Accounting and Financial Management*, 9(12), 218–238. <https://doi.org/10.56201/jafm.v9.no12.2023.pg218.238>
- Ominyi, M., & Anichukwueze, C. C. (2024). Conceptualizing responsible AI governance structures for automated decision systems in regulated industries. *International Journal of Social Sciences and Management Research*, 10(11), 403–430. <https://doi.org/10.56201/ijssmr.v10.no11.2024.pg.403.430>
- Ominyi, M., Anichukwueze, C. C., & Uzougbo, N. S. (2024a). Developing a conceptual framework for AI-assisted compliance monitoring and anomaly detection in high-volume transaction environments. *World Journal of Innovation and Modern Technology*, 8(6), 204–228. <https://doi.org/10.56201/wjimt.v8.no6.2024.pg204.228>
- Ominyi, M., Anichukwueze, C. C., & Uzougbo, N. S. (2024b). A systematic review of organizational readiness for the EU Artificial Intelligence Act in multinational enterprises. *World Journal of Innovation and Modern Technology*, 8(6), 185–203. <https://doi.org/10.56201/wjimt.v8.no6.2024.pg185.203>
- Ominyi, M., Anichukwueze, C. C., & Uzougbo, N. S. (2025a). A conceptual framework for multi-agent AI quality control in the review of regulated documents. *International Journal of*

- Social Sciences and Management Research*, 11(8), 534–557.
<https://doi.org/10.56201/ijssmr.vol.11no8.2025.pg534.557>
- Ominyi, M., Anichukwueze, C. C., & Uzougbo, N. S. (2025c). Reviewing legal privilege and confidentiality challenges in enterprise deployment of large language models. *Journal of Law and Global Policy*, 10(3), 150–176.
<https://doi.org/10.56201/jlgp.vol.10.no3.2025.pg150.176>
- Omo Enabulele, A. B., Oyebamiji, I., Ikubanni, O., Okwah, M. N., Babalola, W. S., Azeez, O. M., & Olaniyi, A. O. (2025). A coordinated project-management approach to multisite implementation of motor-rehabilitation programs for children with autism spectrum disorder in United States healthcare systems: A narrative review. *Cureus*, 17(12), e100513.
<https://doi.org/10.7759/cureus.100513>
- Oreg, S., Vakola, M., & Armenakis, A. (2011). Change recipients' reactions to organizational change: A 60-year review of quantitative studies. *The Journal of Applied Behavioral Science*, 47(4), 461–524. <https://doi.org/10.1177/0021886310396550>
- Orlikowski, W. J. (2002). Knowing in practice: Enacting a collective capability in distributed organizing. *Organization Science*, 13(3), 249–273.
<https://doi.org/10.1287/orsc.13.3.249.2776>
- Ozowara, D. E., Adebayo, A., & Anunagba, C. O. (2025). Advanced conceptual model for strengthening audit quality using data analytics across financial institutions. *International Journal of Advanced Multidisciplinary Research and Studies*, 5(6), 2246–2268.
<https://doi.org/10.62225/2583049X.2025.5.6.6050>
- Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2), 230–253. <https://doi.org/10.1518/001872097778543886>
- Paul, D., Cadle, J., Eva, M., Rollason, C., & Hunsley, J. (2020). *Business Analysis* (4th ed.). Swindon: BCS, The Chartered Institute for IT.
- Peppard, J., Ward, J., & Daniel, E. (2007). Managing the realization of business benefits from IT investments. *MIS Quarterly Executive*, 6(1), Article 3.
- Plattfaut, R., Borghoff, V., Godefroid, M., Koch, J., Trampler, M., & Coners, A. (2022). The critical success factors for robotic process automation. *Computers in Industry*, 138, 103646. <https://doi.org/10.1016/j.compind.2022.103646>
- Plattfaut, R., Rehse, J.-R., Jans, C., Schulte, M., & van Wendel de Joode, J. (2024). Robotic process automation: Research impulses from the BPM 2023 panel discussion. *Process Science*, 1, Article 5. <https://doi.org/10.1007/s44311-024-00005-1>
- Pohl, K. (2010). *Requirements Engineering: Fundamentals, Principles, and Techniques*. Berlin: Springer.
- Popović, A., Hackney, R., Coelho, P. S., & Jaklič, J. (2012). Towards business intelligence systems success: Effects of maturity and culture on analytical decision making. *Decision Support Systems*, 54(1), 729–739. <https://doi.org/10.1016/j.dss.2012.08.017>
- Rafferty, A. E., Jimmieson, N. L., & Armenakis, A. A. (2013). Change readiness: A multilevel review. *Journal of Management*, 39(1), 110–135.
<https://doi.org/10.1177/0149206312457417>
- Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., ... Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 33–44). New York: ACM.
<https://doi.org/10.1145/3351095.3372873>

- Rakova, B., Yang, J., Cramer, H., & Chowdhury, R. (2021). Where responsible AI meets reality: Practitioner perspectives on enablers for shifting organizational practices. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), 1–23. <https://doi.org/10.1145/3449081>
- Robertson, J., Robertson, S., & Reed, A. (2024). *Mastering the Requirements Process* (4th ed.). Boston: Addison-Wesley Professional.
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). New York: Free Press.
- Rosemann, M., & vom Brocke, J. (2014). The six core elements of business process management. In J. vom Brocke & M. Rosemann (Eds.), *Handbook on business process management 1* (2nd ed., pp. 105–122). Berlin: Springer. https://doi.org/10.1007/978-3-642-45100-3_5
- Rowley, T. J. (1997). Moving beyond dyadic ties: A network theory of stakeholder influences. *The Academy of Management Review*, 22(4), 887. <https://doi.org/10.2307/259248>
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- Sadiq, R. B., Safie, N., Abd Rahman, A. H., & Goudarzi, S. (2021). Artificial intelligence maturity model: A systematic literature review. *PeerJ Computer Science*, 7, e661. <https://doi.org/10.7717/peerj-cs.661>
- Sagala, G. H., & Őri, D. (2024). Toward SMEs digital transformation success: A systematic literature review. *Information Systems and e-Business Management*, 22(4), 667–719. <https://doi.org/10.1007/s10257-024-00682-2>
- Sambasivan, N., Kapania, S., Highfill, H., Akrong, D., Paritosh, P., & Aroyo, L. (2021). "Everyone wants to do the model work, not the data work": Data cascades in high-stakes AI. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (pp. 1–15). Association for Computing Machinery. <https://doi.org/10.1145/3411764.3445518>
- Sánchez, E., Calderón, R., & Herrera, F. (2025). Artificial intelligence adoption in SMEs: Survey based on TOE–DOI framework, primary methodology and challenges. *Applied Sciences*, 15(12), 6465. <https://doi.org/10.3390/app15126465>
- Sanderson, C., Douglas, D., Lu, Q., Schleiger, E., Whittle, J., Lacey, J., Newnham, G., Hajkowicz, S., Robinson, C., & Hansen, D. (2023). AI ethics principles in practice: Perspectives of designers and developers. *IEEE Transactions on Technology and Society*, 4(2), 171–187. <https://doi.org/10.1109/TTS.2023.3257303>
- Sanni, J. O., & Atima, M. E. (2021b). Business intelligence dashboard frameworks resolving executive visibility gaps in strategic marketing governance. *International Journal of Multidisciplinary Research and Growth Evaluation*, 2(6), 633–646. <https://doi.org/10.54660/IJMRGE.2021.2.6.633-646>
- Sanni, J. O., & Attah, A. (2023a). A comprehensive framework for digital transformation in capital markets: Solving operational challenges and enhancing stakeholder engagement. *Gyanshauryam, International Scientific Refereed Research Journal*, 6(6), 275–302. <https://doi.org/10.32628/GISRRJ236640>
- Sanni, J. O., Ajiga, D., & Atima, M. E. (2020a). Analytical models addressing measurement challenges of marketing return on investment in regulated services. *International Journal of Multidisciplinary Research and Growth Evaluation*, 1(5), 636–648. <https://doi.org/10.54660/IJMRGE.2020.1.5.636-648>
- Sanni, J. O., Iwuanyanwu, U. A., & Essien, M. A. (2025). Problem-oriented process mining for auditable marketing automation lifecycle control. *International Journal of Advanced*

- Multidisciplinary Research and Studies*, 5(6), 1933–1947.
<https://doi.org/10.62225/2583049X.2025.5.6.5650>
- Sanni, J. O., Iwuanyanwu, U. A., Essien, M. A., Atima, M. E., & Attah, A. (2022). Adaptive control models for AI-driven marketing automation in financial compliance environments. *Shodhshauryam, International Scientific Refereed Research Journal*, 5(1), 243–270. <https://doi.org/10.32628/SHISRRJ247133>
- Sarker, S., Chatterjee, S., Xiao, X., & Elbanna, A. (2019). The sociotechnical axis of cohesion for the IS discipline: Its historical legacy and its continued relevance. *MIS Quarterly*, 43(3), 695–719. <https://doi.org/10.25300/MISQ/2019/13747>
- Schneider, J., Abraham, R., Meske, C., & vom Brocke, J. (2023). Artificial intelligence governance for businesses. *Information Systems Management*, 40(3), 229–249. <https://doi.org/10.1080/10580530.2022.2085825>
- Schwaeke, J., Peters, A., Kanbach, D. K., Kraus, S., & Jones, P. (2024). The new normal: The status quo of AI adoption in SMEs. *Journal of Small Business Management*. <https://doi.org/10.1080/00472778.2024.2379999>
- Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. In *Advances in Neural Information Processing Systems 28 (NIPS 2015)*. Curran Associates.
- Selbst, A. D., boyd, d., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 59–68). Association for Computing Machinery. <https://doi.org/10.1145/3287560.3287598>
- Serra, C. E. M., & Kunc, M. (2015). Benefits realisation management and its influence on project success and on the execution of business strategies. *International Journal of Project Management*, 33(1), 53–66. <https://doi.org/10.1016/j.ijproman.2014.03.011>
- Shneiderman, B. (2020). Bridging the gap between ethics and practice: Guidelines for reliable, safe, and trustworthy human-centered AI systems. *ACM Transactions on Interactive Intelligent Systems*, 10(4), 1–31. <https://doi.org/10.1145/3419764>
- Skitka, L. J., Mosier, K. L., & Burdick, M. (1999). Does automation bias decision-making? *International Journal of Human-Computer Studies*, 51(5), 991–1006. <https://doi.org/10.1006/ijhc.1999.0252>
- Sonteya, T., & Seymour, L. F. (2012). Towards an understanding of the business process analyst: An analysis of competencies. *Journal of Information Technology Education: Research*, 11, 43–63. <https://doi.org/10.28945/1568>
- Star, S. L. (2010). This is not a boundary object: Reflections on the origin of a concept. *Science, Technology, & Human Values*, 35(5), 601–617. <https://doi.org/10.1177/0162243910377624>
- Star, S. L., & Griesemer, J. R. (1989). Institutional ecology, "translations" and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907–39. *Social Studies of Science*, 19(3), 387–420. <https://doi.org/10.1177/030631289019003001>
- Sutcliffe, A., & Sawyer, P. (2013). Requirements elicitation: Towards the unknown unknowns. In *2013 21st IEEE International Requirements Engineering Conference (RE)* (pp. 92–104). IEEE. <https://doi.org/10.1109/RE.2013.6636709>

- Syed, R., Suriadi, S., Adams, M., Bandara, W., Leemans, S. J. J., Ouyang, C., ... Reijers, H. A. (2020). Robotic process automation: Contemporary themes and challenges. *Computers in Industry*, 115, 103162. <https://doi.org/10.1016/j.compind.2019.103162>
- Szulanski, G. (1996). Exploring internal stickiness: Impediments to the transfer of best practice within the firm. *Strategic Management Journal*, 17(S2), 27–43. <https://doi.org/10.1002/smj.4250171105>
- Telford, T., Thopalli, K., Schaefer, G., Polner, A., Urbaniak, T., & Wright, D. (2022). *Automation with intelligence* (2022 survey). Deloitte Insights.
- Thong, J. Y. L. (2001). Resource constraints and information systems implementation in Singaporean small businesses. *Omega*, 29(2), 143–156. [https://doi.org/10.1016/S0305-0483\(00\)00035-9](https://doi.org/10.1016/S0305-0483(00)00035-9)
- Tonoyan, A., Dada, O., & Ayivi-Donkor, S. S. (2022b). Algorithmic process optimization and cost reduction: A review of simulation modeling and financial impact assessment. *Journal of Accounting and Financial Management*, 8(8), 139–169. <https://doi.org/10.56201/jafm.vol.8.no8.2022.p139.169>
- Tonoyan, A., Dada, O., & Ayivi-Donkor, S. S. (2024b). Real-time KPI tracking systems: A review of automated performance monitoring and data-driven decision making. *World Journal of Innovation and Modern Technology*, 8(6), 247–281. <https://doi.org/10.56201/wjimt.v8.no6.2024.pg247.281>
- Tornatzky, L. G., Fleischer, M., & Chakrabarti, A. K. (1990). *The processes of technological innovation*. Lexington, MA: Lexington Books.
- Trist, E. L., & Bamforth, K. W. (1951). Some social and psychological consequences of the longwall method of coal-getting. *Human Relations*, 4(1), 3–38. <https://doi.org/10.1177/001872675100400101>
- Trkman, P. (2010). The critical success factors of business process management. *International Journal of Information Management*, 30(2), 125–134. <https://doi.org/10.1016/j.ijinfomgt.2009.07.003>
- Tushman, M. L. (1977). Special boundary roles in the innovation process. *Administrative Science Quarterly*, 22(4), 587–605. <https://doi.org/10.2307/2392402>
- Uren, V., & Edwards, J. S. (2023). Technology readiness and the organizational journey towards AI adoption: An empirical study. *International Journal of Information Management*, 68, 102588. <https://doi.org/10.1016/j.ijinfomgt.2022.102588>
- van den Broek, E., Sergeeva, A., & Huysman, M. (2021). When the machine meets the expert: An ethnography of developing AI for hiring. *MIS Quarterly*, 45(3), 1557–1580. <https://doi.org/10.25300/misq/2021/16559>
- van der Aalst, W. M. P., Bichler, M., & Heinzl, A. (2018). Robotic process automation. *Business & Information Systems Engineering*, 60(4), 269–272. <https://doi.org/10.1007/s12599-018-0542-4>
- van Lamsweerde, A. (2009). *Requirements Engineering: From System Goals to UML Models to Software Specifications*. Chichester: John Wiley & Sons.
- Van Looy, A., & Shafagatova, A. (2016). Business process performance measurement: A structured literature review of indicators, measures and metrics. *SpringerPlus*, 5, Article 1797. <https://doi.org/10.1186/s40064-016-3498-1>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>

- Venkatesh, V., Thong, J. Y. L., & Xu, X. (2012). Consumer acceptance and use of information technology: Extending the unified theory of acceptance and use of technology. *MIS Quarterly*, 36(1), 157–178. <https://doi.org/10.2307/41410412>
- Vogelsang, A., & Borg, M. (2019). Requirements engineering for machine learning: Perspectives from data scientists. In *2019 IEEE 27th International Requirements Engineering Conference Workshops (REW)* (pp. 245–251). IEEE. <https://doi.org/10.1109/REW.2019.00050>
- Waizenegger, L., & Techatassanasoontorn, A. A. (2022). When robots join our team: A configuration theory of employees' perceptions of and reactions to robotic process automation. *Australasian Journal of Information Systems*, 26. <https://doi.org/10.3127/ajis.v26i0.3833>
- Walawalkar, G., Adesuyi, M. O., Kalu, A., & Oduleye, T. E. (2025). Executive financial dashboards for real-time strategic oversight. *International Journal of Advanced Multidisciplinary Research and Studies*, 5(6), 2042–2054.
- Ward, J., & Daniel, E. (2012). *Benefits management: How to increase the business value of your IT projects* (2nd ed.). Chichester: John Wiley & Sons.
- Ward, J., De Hertogh, S., & Viaene, S. (2007). Managing benefits from IS/IT investments: An empirical investigation into current practice. In *Proceedings of the 40th Annual Hawaii International Conference on System Sciences (HICSS 2007)* (p. 206a). IEEE. <https://doi.org/10.1109/HICSS.2007.330>
- Ward, J., Taylor, P., & Bond, P. (1996). Evaluation and realisation of IS/IT benefits: An empirical study of current practice. *European Journal of Information Systems*, 4(4), 214–225. <https://doi.org/10.1057/ejis.1996.3>
- Weiner, B. J. (2009). A theory of organizational readiness for change. *Implementation Science*, 4(1), 67. <https://doi.org/10.1186/1748-5908-4-67>
- Weiner, B. J., Amick, H., & Lee, S.-Y. D. (2008). Review: Conceptualization and measurement of organizational readiness for change. *Medical Care Research and Review*, 65(4), 379–436. <https://doi.org/10.1177/1077558708317802>
- Wenger, E. (1998). *Communities of practice*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511803932>
- Wenger, E. (2000). Communities of practice and social learning systems. *Organization*, 7(2), 225–246. <https://doi.org/10.1177/135050840072002>
- Wieggers, K., & Beatty, J. (2013). *Software Requirements* (3rd ed.). Redmond, WA: Microsoft Press.
- Wieringa, M. (2020). What to account for when accounting for algorithms: A systematic literature review on algorithmic accountability. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 1–18). New York: ACM. <https://doi.org/10.1145/3351095.3372833>
- Willcocks, L. P., Lacity, M., & Craig, A. (2015). *Robotic process automation at Xchanging* (The Outsourcing Unit Working Research Paper Series, Paper 15/03). London: London School of Economics and Political Science.
- Willcocks, L., Lacity, M., & Craig, A. (2017). Robotic process automation: Strategic transformation lever for global business services? *Journal of Information Technology Teaching Cases*, 7(1), 17–28. <https://doi.org/10.1057/s41266-016-0016-9>
- Wixom, B. H., & Watson, H. J. (2001). An empirical investigation of the factors affecting data warehousing success. *MIS Quarterly*, 25(1), 17–41.

- Wright, D., Telford, T., Schaefer, G., Watson, J., Polner, A., & Witherick, D. (2020). *Automation with intelligence: Pursuing organisation-wide reimagination*. Deloitte Insights.
- Zahra, S. A., & George, G. (2002). Absorptive capacity: A review, reconceptualization, and extension. *The Academy of Management Review*, 27(2), 185–203. <https://doi.org/10.2307/4134351>
- Zhao, L., Alhoshan, W., Ferrari, A., Letsholo, K. J., Ajagbe, M. A., Chioasca, E.-V., & Batista-Navarro, R. T. (2021). Natural language processing for requirements engineering: A systematic mapping study. *ACM Computing Surveys*, 54(3), 1–41. <https://doi.org/10.1145/3444689>