

RESEARCH ARTICLE

Operationalising Consulting Methodology Through Workflow Orchestration: A Design Science Study of CWORT with Governed Agentic AI Support

Oluwaseun Iyiola^{1,*} ¹Lapace Services UK Ltd, UK

Abstract: Consulting engagements play a critical role in organisational decision-making, yet delivery often remains highly manual, fragmented, and dependent on individual practitioner experience. While consulting methodologies provide conceptual guidance, they are rarely operationalised as executable systems, leading to variability in delivery quality, limited traceability between inputs and outputs, and challenges in scaling consistent advisory practices. The growing use of digital and artificial intelligence (AI)-enabled tools in consulting has tended to focus on isolated tasks rather than the orchestration of end-to-end workflows, raising concerns around governance and accountability. This paper introduces CWORT (Consulting Workflow Orchestration Tool), a socio-technical system that operationalises consulting methodology through a digitally governed, workflow-based approach. CWORT represents consulting delivery as a sequence of explicit workflow states, integrating role-based governance, artefact traceability, and constrained agentic reasoning components to support analysis and synthesis activities. AI capabilities are embedded within predefined workflow stages and operate under strict human-in-the-loop control, ensuring that professional judgement and accountability remain central to advisory outcomes. The study is framed within a design science research paradigm and presents the conceptual model, system architecture, and a structured case-based evaluation in a real consulting engagement. The evaluation adopts a within-context comparative field design, demonstrating that workflow orchestration improves transparency, consistency, and traceability of consulting outputs, while enabling substantial reductions in discovery and assessment effort relative to prior approaches. AI-supported components are shown to augment interpretation and insight generation without altering deterministic analytical outcomes.

Keywords: consulting, agentic AI, human-in-the-loop AI, artificial intelligence, decision-support system

1. Introduction

Consulting engagements play a central role in supporting organisations as they navigate digital transformation, operational change, and complex technology decisions [1]. Professional advisory work typically involves discovery, assessment, interpretation, and the synthesis of findings into recommendations that inform strategic and operational action. Despite the maturity of consulting as a profession, prior research has consistently observed that the delivery of consulting engagements remains heavily dependent on manual coordination, tacit knowledge, and individual practitioner experience [2]. As a result, consulting delivery often exhibits variability in quality, limited transparency in decision rationale, and challenges in scaling consistent practices across engagements [3]. From an information systems perspective, consulting work can be characterised as a complex knowledge-intensive activity, where inputs, interpretations, and outcomes are distributed across multiple artefacts and tools [4]. Discovery data, assessment

responses, analytical notes, and recommendations are frequently captured in disconnected documents, spreadsheets, and presentation materials [5]. Research on work systems and knowledge work has shown that such fragmentation weakens coordination, reduces traceability between inputs and outputs, and obscures the rationale underpinning decisions [6, 7]. In consulting contexts, these limitations can undermine client confidence, complicate assurance and audit activities, and increase delivery risk in regulated or high-stakes environments [8].

Consulting firms have long sought to address these challenges through the development of formal methodologies and frameworks [9]. These methodologies typically define phases, deliverables, and guiding principles intended to standardise engagement delivery while preserving professional judgement [2]. However, the literature suggests that such methodologies are predominantly conceptual in nature and rely on consultants to interpret and enact them through manual processes [10]. The absence of executable mechanisms to operationalise methodological intent means that consistency, traceability, and accountability remain dependent on individual discipline rather than systemic support [2].

*Corresponding author: Oluwaseun Iyiola, Lapace Services UK Ltd, UK. Email: olu.iyiola@cwort.com

In parallel, the consulting industry has adopted an expanding range of digital tools to support aspects of advisory delivery, including survey platforms, maturity assessment tools, analytics systems, and collaboration environments [11]. While these tools can improve efficiency at specific stages, research in information systems has shown that optimising isolated tasks rarely addresses broader issues of workflow coordination and accountability [6]. More recently, artificial intelligence (AI) has been introduced to augment consulting activities such as analysis and content generation [12]. Studies of AI in professional services suggest that such technologies can not only enhance speed and pattern recognition but also introduce risks related to opacity, accountability, and overreliance on automated outputs if not carefully governed [13]. This paper introduces CWORT (Consulting Workflow Orchestration Tool) as a response to these limitations. CWORT is positioned as a consulting methodology operationalised through a digital platform, rather than as a generic software solution to which advisory practices are subsequently attached. Its core aim is to encode established consulting practices into structured, executable workflows that improve consistency, traceability, and defensibility while preserving the central role of human judgement. CWORT incorporates constrained, agentic reasoning components to support analysis and synthesis activities within predefined workflow stages, while maintaining explicit human-in-the-loop control over all consulting decisions [8].

Conceptually, CWORT is designed as a socio-technical system in which people, processes, and technology are explicitly aligned. Consultants retain responsibility for interpretation, validation, and decision-making, while the platform provides workflow governance, artefact traceability, and controlled automation. Clients interact with the system in a guided and role-appropriate manner, contributing inputs and reviewing outputs without undermining consultant-led methodological control. This design reflects established principles in professional service delivery, where collaboration is essential but accountability for recommendations must remain clearly defined [2]. The proposed approach is grounded in a conceptual model that links consulting methodology, workflow execution, and governed agentic AI support to accountable and transparent consulting outcomes. This paper makes four primary contributions to the study of digitally enabled consulting and AI-supported professional work. It operationalises consulting methodology as an executable system by representing consulting delivery as a sequence of governed workflow states, enabling consistent execution, traceability, and scalability. It introduces a workflow orchestration approach that integrates role-based governance and artefact traceability within a socio-technical system, improving coordination, transparency, and accountability in advisory processes. It further proposes a governed model of AI integration through the architectural separation of deterministic control-based analysis and AI-supported augmentation, demonstrating how AI can support summarisation, explanation, and insight generation while preserving human judgement and accountability. Finally, the paper provides a structured, case-based evaluation of the artefact through a real-world consulting engagement using a within-context comparative field design, offering empirically grounded evidence of improvements in workflow coherence, traceability, and delivery efficiency [14].

2. Literature Review

2.1. Problem statement and motivation

Consulting delivery constitutes a form of knowledge-intensive professional work characterised by high levels of

ambiguity, contextual dependence, and reliance on expert judgement. Such work typically involves the collection of qualitative and quantitative inputs, iterative interpretation, and the synthesis of findings into recommendations that inform organisational decision-making. Research on professional service firms has consistently highlighted that the quality and credibility of consulting outcomes depend not only on analytical rigour but also on the processes through which evidence is gathered, interpreted, and justified [11]. Despite the widespread use of formal consulting methodologies, the operational execution of consulting engagements remains largely manual [2]. Methodologies typically define phases, deliverables, and guiding principles but leave the coordination of activities, artefacts, and decisions to individual consultants [15]. Studies of professional knowledge work suggest that this reliance on tacit coordination introduces variability and makes delivery outcomes sensitive to individual experience and working practices [7]. As a result, engagements following the same nominal methodology may differ substantially in structure, depth, and defensibility.

A central challenge arising from this mode of delivery is the lack of traceability between engagement inputs and outputs. Discovery data, stakeholder interviews, assessment responses, and analytical insights are frequently captured in disparate artefacts, such as documents, spreadsheets, and slide decks [11]. Research in information systems and work system theory indicates that when information, reasoning, and decision-making are distributed across disconnected tools, the rationale underpinning recommendations becomes difficult to reconstruct [6]. In consulting contexts, this opacity weakens accountability and poses risks in environments where recommendations must be justified to regulators, auditors, or senior executives [16]. Closely related to traceability is the issue of delivery variability. While professional judgement necessarily introduces some degree of variation, excessive inconsistency can undermine the perceived reliability of consulting services and limit opportunities for organisational learning and reuse. Prior research on professional services highlights a persistent tension between standardisation and expertise-driven discretion, noting that the absence of structured execution mechanisms often leads to unpredictable outcomes and inefficiencies in scaling advisory practices [2]. The digital tools commonly employed in consulting engagements have not fully addressed these challenges. Survey platforms, maturity models, analytics tools, and collaboration systems typically support isolated tasks rather than the orchestration of the consulting lifecycle as an integrated whole. Research on workflow and process management has shown that task-level optimisation, when not embedded within an end-to-end process model, offers limited benefits for complex, interdependent work [17]. Consequently, consultants must manually bridge gaps between tools, increasing cognitive load and the likelihood of undocumented assumptions, inconsistent execution, or missed dependencies [18].

The increasing introduction of AI into consulting workflows further accentuates these limitations [19]. AI-enabled tools can assist with analysis, pattern recognition, and content generation, but their unstructured adoption risks amplifying opacity and undermining accountability. Studies of AI in organisational decision-making caution that without explicit governance and human oversight, AI systems may obscure responsibility and erode trust in professional outcomes (Davenport & Kirby, 2016). These concerns are particularly salient in consulting, where recommendations often carry significant organisational, financial, or regulatory implications. From a design science perspective, these limitations indicate the need for a new class of artefact that moves beyond static methodologies and point solutions. Such an artefact

must operationalise consulting methodology through structured, executable workflows that explicitly align people, processes, and technology. It must support traceability across engagement stages, reduce unnecessary delivery variability, and enable the governed use of digital and AI-assisted capabilities without displacing professional judgement. Design science research (DSR) emphasises that effective artefacts should be grounded in real-world problems and evaluated through practical use rather than theoretical abstraction alone [14].

The motivation for CWORT arises directly from this gap. CWORT is proposed as a workflow orchestration approach that treats consulting delivery as a socio-technical system rather than a collection of ad hoc tasks. Embedding consulting methodology into structured workflows, supporting artefact traceability, and constraining AI-assisted reasoning within explicit human validation, CWORT seeks to address long-standing challenges in transparency, consistency, and defensibility that are not adequately resolved by existing consulting frameworks or digital tools.

3. Research Objective and Method

3.1. Research question

This study addresses the following research question:

How can consulting methodologies be operationalised through a digitally orchestrated workflow incorporating governed, agentic AI support while preserving professional judgement and accountability?

This research focuses on the design and feasibility of a socio-technical artefact rather than on predictive performance or statistical effectiveness. Specifically, it examines how consulting methodology can be translated into structured, executable workflow constructs and how agentic AI capabilities can be embedded as governed, decision-support mechanisms within those constructs without displacing human responsibility for interpretation and decision-making. The CWORT conceptual model provides the structural basis through which this research question is addressed. It defines the relationships between consulting methodology, workflow execution, orchestration mechanisms, artefact traceability, constrained agentic reasoning, and human validation, thereby guiding both artefact design and evaluation.

3.1.1. Research method

This study adopts a DSR approach to investigate how consulting methodologies can be operationalised through digitally orchestrated workflows incorporating governed agentic AI support. The research objective is to design and evaluate a purposeful socio-technical artefact that improves consistency, traceability, and accountability in consulting delivery while preserving professional judgement. Consistent with established design science guidance, the study follows an artefact-centred research process comprising artefact construction, demonstration, and formative evaluation [14]. The CWORT artefact is developed iteratively, informed by the conceptual model presented in Section 4 and realised through a bounded application architecture described in Section 5.

The design of CWORT is informed by kernel theories drawn from work system theory [6], socio-technical systems theory, and decision-support system (DSS) design principles. These theories inform the meta-requirements of structured workflow execution,

explicit governance, artefact traceability, and human-centred augmentation, which collectively guide the artefact design.

Evaluation is qualitative and exploratory in nature, focusing on feasibility, coherence with consulting practice, and the extent to which the artefact supports structured workflow execution, traceability of decisions, and governed use of agentic AI. The study does not seek statistical generalisation but rather to provide analytically grounded insights into the applicability and limitations of workflow-orchestrated consulting delivery supported by AI under human governance.

4. Related Work

CWORT sits at the intersection of multiple research streams, including consulting methodologies and professional service delivery, maturity and assessment models, workflow and process management systems, DSS, and emerging research on AI-assisted knowledge work. This section reviews relevant literature across these areas to situate CWORT within existing scholarship and to clarify its contribution relative to prior work.

4.1. Consulting methodologies and professional service delivery

Consulting methodologies have long been used to structure advisory engagements, providing phased models for discovery, analysis, synthesis, and recommendation development [15]. Within professional service firms, such methodologies function as coordination mechanisms intended to promote consistency, knowledge reuse, and quality assurance across engagements [15]. However, research on professional services consistently observes that these methodologies are primarily articulated as conceptual frameworks rather than operational systems, leaving execution highly dependent on individual consultant judgement and informal practices.

Studies of knowledge-intensive work emphasise that professional expertise is enacted through tacit coordination, situated interpretation, and emergent practices rather than rigid process adherence [7]. While this flexibility is central to professional value creation, it also introduces delivery variability and limits the scalability and repeatability of consulting practices. Prior work highlights a persistent structural tension between the need for standardisation to ensure reliability and the need for discretion to accommodate contextual complexity [2]. CWORT builds on this literature by seeking to operationalise consulting methodology through structured workflows that preserve professional judgement while reducing unnecessary variability through orchestration and governance.

4.2. Maturity models and assessment-based approaches

Maturity models and capability assessments are widely employed in consulting to evaluate organisational practices and inform improvement roadmaps. Research documents the prevalence of staged maturity models across domains such as IT management, process improvement, and digital capability, where they provide a shared assessment language and comparative benchmarks [20]. Despite their widespread use, maturity models are often applied as stand-alone diagnostic instruments rather than as integrated components of an end-to-end consulting workflow.

Scholars have raised concerns regarding the transparency, validity, and interpretability of maturity assessments, particularly when assessment outcomes are presented without explicit linkage to underlying evidence, interpretive reasoning, or methodological context [20]. When assessments are detached from broader delivery processes, the connection between responses, analysis, and recommendations can become opaque. CWORT addresses this limitation by embedding assessment activities within an orchestrated workflow that explicitly links inputs, interpretations, and outputs, thereby preserving assessment value while improving traceability and defensibility.

4.3. Workflow management and process-oriented systems

Workflow and business process management systems have been extensively studied as mechanisms for coordinating complex organisational activities. Research in this area emphasises the importance of explicit process models, state management, and dependency handling in improving transparency, control, and repeatability [17]. Such systems have been successfully applied to transactional and operational processes, where activities are predictable and outcomes are well defined.

However, consulting work differs fundamentally from operational processes in that it involves interpretive reasoning, iterative sense-making, and negotiated meaning. Prior research on knowledge work and emergent processes highlights that traditional workflow systems often struggle to accommodate the flexibility required for such activities [7]. CWORT extends workflow-oriented thinking into the consulting domain by combining explicit state management with human-in-the-loop control, enabling structure and coordination without imposing rigid automation or eliminating professional discretion.

4.4. Decision-support systems and socio-technical perspectives

DSS research provides important foundations for understanding how digital systems can augment human judgement in complex decision contexts. Early DSS literature emphasised that such systems should support, rather than replace, human decision-making, particularly in situations characterised by uncertainty, ambiguity, and competing objectives [6]. Subsequent socio-technical perspectives reinforce this view, highlighting the importance of aligning technological capabilities with organisational roles, responsibilities, and accountability structures [21].

CWORT aligns with these principles by positioning digital and AI-assisted capabilities as supportive artefacts embedded within a governed workflow. Rather than generating decisions autonomously, CWORT requires explicit human validation at key points, preserving accountability and professional responsibility. This design reflects long-standing guidance in DSS research regarding transparency, controllability, and trust in DSS [6].

4.5. Artificial intelligence and agentic support in professional work

The introduction of AI into professional and advisory contexts has generated significant scholarly interest. Research suggests that AI can enhance professional work by supporting pattern recognition, analysis, and information synthesis, while also raising concerns related to opacity, bias, and accountability

(Davenport & Kirby, 2016). Studies of human–AI collaboration emphasise the importance of maintaining human oversight and contextual interpretation, particularly in high-stakes decision environments [22]. Within AI literature, agentic systems are commonly defined as entities capable of goal-directed behaviour within an environment [13]. While much agent research focuses on autonomy and coordination, there is increasing recognition that constrained, human-governed agentic components may be more appropriate for professional domains where accountability must remain explicit. CWORT adopts this constrained interpretation by incorporating agentic reasoning components that operate within predefined workflow stages and cannot act independently of human approval.

4.6. Positioning of CWORT

Existing research provides insights into consulting methodologies, assessment models, Computer-Supported Cooperative Work, human–computer interaction (HCI), workflow and Business Process Modelling systems, DSS, and AI-assisted knowledge work [2, 6, 17, 20, 23, 24], as reflected in widely adopted frameworks and platforms (e.g., COBIT, ITIL; ServiceNow, Camunda; Power BI; ChatGPT). However, while existing approaches address specific aspects of consulting activity, the literature does not fully operationalise consulting methodology as an end-to-end, digitally orchestrated, and governed socio-technical workflow that integrates these elements coherently.

CWORT contributes to this gap by integrating workflow orchestration, role-based governance, artefact traceability, and constrained agentic support within a single execution model. Rather than replacing existing methodologies or tools, CWORT provides an enabling infrastructure that aligns them within a governed consulting workflow, thereby extending prior work across multiple research streams.

Prior research has addressed elements of collaboration, workflow management, and decision support; these approaches typically operate at the level of task coordination, user interaction, or process automation. Collaboration systems focus on enabling communication and shared artefact access, while workflow and process management systems are primarily designed for structured, repeatable processes with predefined execution logic. Similarly, HCI research has emphasised usability and interaction with intelligent systems, including AI-assisted tools. However, these approaches do not operationalise consulting methodology itself as an executable, governed workflow. CWORT differs in that it treats consulting methodology as the primary object of system design, embedding methodological phases, roles, artefacts, and validation logic into a unified orchestration layer. This enables the execution of knowledge-intensive consulting work as a structured yet flexible workflow, while preserving interpretive reasoning, professional judgement, and accountability. In doing so, CWORT extends beyond existing systems by integrating workflow orchestration, governance, traceability, and constrained AI support within a methodology-native execution model.

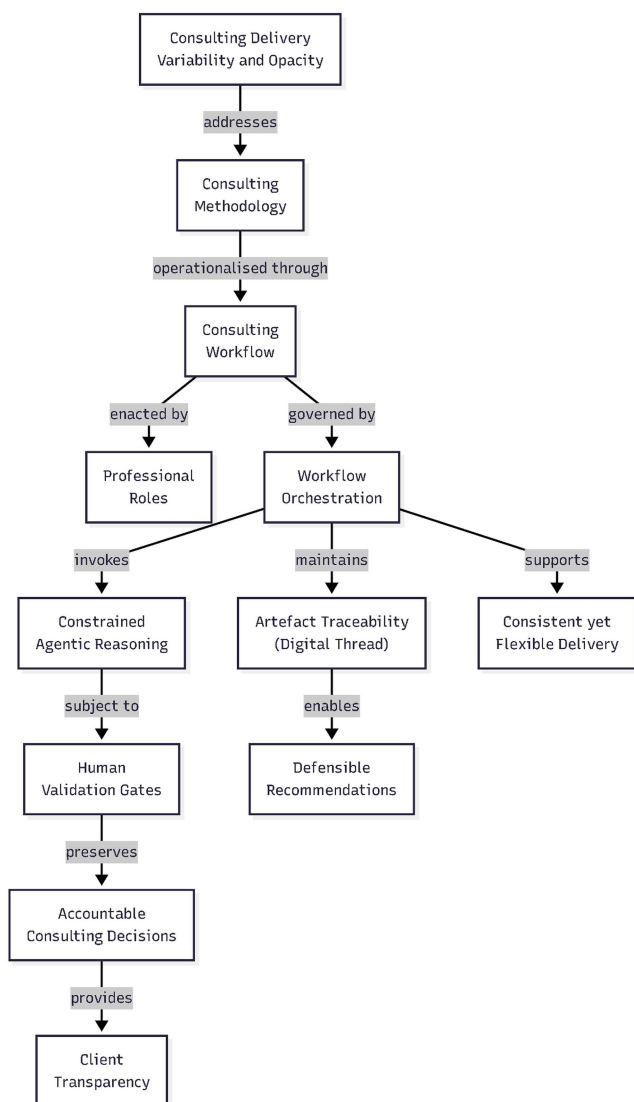
5. CWORT Conceptual Model

The CWORT conceptual model defines how consulting methodology can be operationalised as a workflow-governed socio-technical system that integrates human expertise, structured process, and constrained agentic reasoning [25]. Rather than treating consulting delivery as a loosely connected sequence

of activities, CWORT conceptualises it as an orchestrated system in which workflow stages, professional roles, artefacts, and decision-support mechanisms are explicitly aligned. As illustrated in Figure 1, the model addresses the problem of variability and opacity in consulting delivery by linking methodological intent to executable workflow structures and accountable outcomes. The model captures the relationships between consulting methodology, workflow execution, orchestration mechanisms, artefact traceability, constrained agentic reasoning, and explicit human validation. Together, these elements form a coherent structure that supports consistency, transparency, and defensibility while preserving professional judgement.

The conceptual model is intentionally abstract and technology-agnostic. It does not prescribe specific implementation technologies or deployment architectures but instead provides a unifying abstraction that guides system design, demonstration, and evaluation. In the remainder of this paper, the model serves as the conceptual foundation for the CWORT application architecture (Section 6), the illustrative use case (Section 6), and the evaluation of observed outcomes (Section 7).

Figure 1
CWORT conceptual model



5.1. Consulting methodology as an executable workflow

Traditional consulting methodologies describe phases such as discovery, assessment, synthesis, and recommendation but typically rely on consultants to manually interpret and enact these phases in practice. CWORT reframes consulting methodology as an executable workflow, in which each methodological phase is represented as a defined workflow state with explicit inputs, outputs, and validation conditions. This approach aligns with work system theory, which emphasises the importance of modelling work as an integrated system of interacting components rather than as isolated tasks [6]. Progression between workflow states is governed by methodological criteria rather than by tool availability or individual preference. While consultants retain discretion over how activities are performed within each phase, the workflow ensures that key methodological steps are neither omitted nor bypassed. In doing so, CWORT introduces systemic support for consistency and traceability while preserving the interpretive nature of professional consulting work.

5.2. Workflow orchestration and professional roles

Workflow orchestration in CWORT provides the coordination logic that binds methodology, roles, and artefacts into a coherent delivery structure. The orchestration layer governs state transitions, enforces methodological sequencing, and ensures that required inputs and validations are satisfied before progression occurs. Professional roles are embedded directly into this orchestration logic. Consultants retain authority over workflow configuration, interpretation, and decision-making, while clients participate through structured mechanisms for input provision and output review. This separation of roles reflects established principles in professional service delivery, where collaboration is essential but accountability for recommendations must remain clearly defined [2]. Embedding role distinctions into the conceptual model rather than leaving them implicit, CWORT supports collaboration without eroding methodological control. The model thus balances coordination and autonomy, enabling shared engagement while maintaining professional responsibility.

5.3. Constrained agentic reasoning and human validation

A foundational principle of the CWORT conceptual model is that professional judgement remains human-led. Agentic reasoning components are incorporated to support analytical and synthetic activities, but they operate strictly within predefined workflow stages and cannot act independently of human oversight. These components may generate candidate interpretations, summaries, or pattern indicators, but they do not advance workflow state or produce binding conclusions. Progression and acceptance of outcomes require explicit human validation. This design reflects established decision-support research, which argues that systems supporting complex, non-routine work should augment human reasoning rather than replace it [6]. Constraining agentic behaviour within workflow logic and requiring explicit consultant approval, CWORT mitigates risks associated with opaque or overconfident automation. The model preserves accountability, ensures traceability of reasoning, and supports trust in AI-assisted analysis without delegating responsibility to automated agents [26].

5.4. Artefact traceability and digital thread

CWORD maintains a continuous digital thread that links engagement inputs, agent-supported analysis, human interpretation, and final recommendations across the consulting lifecycle. Each artefact is associated with its originating workflow state and validation history, enabling traceability between evidence, reasoning, and outcomes. This explicit linkage addresses documented limitations of assessment-based and document-centric consulting approaches, where outputs are often presented without sufficient evidential grounding or explanation of interpretive steps [20]. Embedding traceability into the conceptual model, CWORD supports defensibility, auditability, and learning across engagements, while reducing reliance on post hoc reconstruction of decision rationale.

5.5. Conceptual model summary and design science alignment

From a design science perspective, the CWORD conceptual model represents a purposeful artefact designed to address a real-world problem in consulting delivery. It integrates socio-technical principles, workflow orchestration, role-based governance, artefact traceability, and constrained agentic reasoning into a coherent abstraction that can be instantiated and evaluated through practical use. Consistent with design science guidance, the model emphasises relevance to professional practice, conceptual rigour, and suitability for formative evaluation rather than optimisation or prediction [14]. The conceptual model provides the foundation for the bounded application architecture described in the following section, which illustrates how these mechanisms are realised within a web-based system.

6. CWORD System Architecture and Workflow Orchestration

The CWORD application architecture in Figure 2 realises the conceptual model presented in Figure 1 by operationalising consulting methodology as a governed, stateful workflow supported by role-based interaction, artefact traceability, and constrained agentic reasoning. The architecture is intentionally bounded and illustrative, focusing on the core mechanisms required to support coordination, transparency, and accountability in consulting delivery rather than on deployment or infrastructure concerns.

6.1. Architectural overview

CWORD adopts a layered application architecture comprising workflow orchestration, role-based interaction, artefact management, and decision-support services. The workflow orchestration layer provides the central coordination mechanism, governing state transitions, validation checkpoints, and role-mediated access to analytical support. Agentic reasoning components are implemented as subordinate services that operate strictly within workflow constraints and are invoked only at defined stages. They do not control workflow progression or produce binding outcomes. This separation of concerns supports modularity and traceability and aligns with design science principles for artefact construction, abstraction, and evaluation [14].

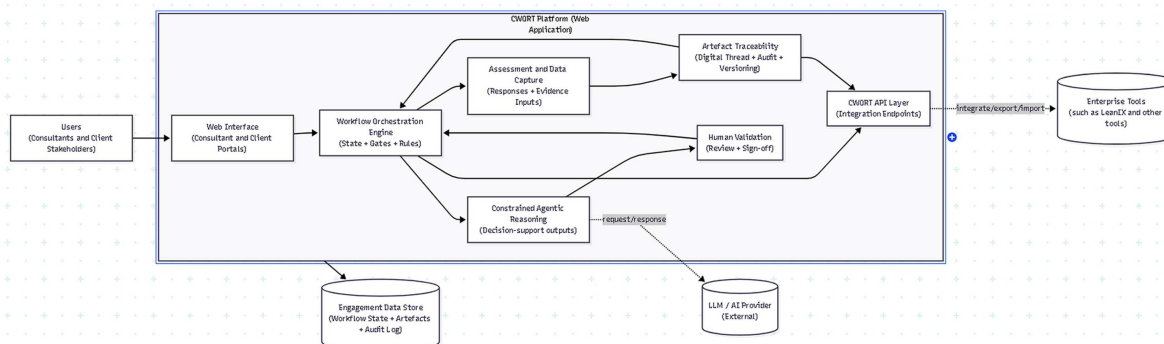
6.2. Workflow state management

Consulting delivery within CWORD is represented as a sequence of explicit workflow states corresponding to methodological phases such as discovery, assessment, synthesis, and recommendation development. Each state defines permissible activities, required artefacts, and validation conditions that must be satisfied before progression occurs. Agentic reasoning components may be invoked within a workflow state to assist with analytical or synthetic tasks, but they cannot initiate state transitions. Progression between states requires explicit consultant validation, ensuring that professional judgement remains central to delivery decisions. This approach extends workflow management concepts traditionally applied to operational processes into the domain of knowledge-intensive consulting work, where interpretive flexibility must be preserved [17].

6.3. Role-based interaction and governance

Role-based interaction is enforced at the architectural level. Consultants are responsible for configuring workflows, invoking agentic reasoning support, interpreting analytical outputs, and validating progression between workflow states. Clients participate through constrained interfaces that allow structured input provision and review of outputs without altering methodological flow or decision logic. Agentic components do not interact directly with clients and cannot bypass consultant mediation. This architectural governance model reflects established practices

Figure 2
CWORD application architecture



in professional service delivery, where responsibility for advisory outcomes must remain clearly assigned and auditable [2].

6.4. Artefact and evidence management

All artefacts produced during a CWORT engagement, including assessment responses, consultant interpretations, and agent-supported analytical outputs, are treated as first-class entities within the architecture. Artefacts are explicitly linked to their originating workflow states and associated validation events, enabling traceability across the consulting lifecycle. This approach supports reconstruction of how conclusions were reached and facilitates assurance and audit activities. Prior research on maturity assessments and diagnostic tools has highlighted the importance of such linkage for transparency, credibility, and defensibility of consulting outcomes [20].

6.5. Decision support and agentic components

Agentic reasoning components within CWORT provide decision-support capabilities by identifying patterns, summarising inputs, or highlighting potential inconsistencies across artefacts. Outputs generated by these components are explicitly labelled as decision-support artefacts and require consultant interpretation and approval before they can inform recommendations or workflow progression. This design aligns with socio-technical perspectives on DSS, which emphasise transparency, controllability, and trust in systems supporting complex professional judgement [6]. Constraining agentic behaviour within workflow logic and embedding human validation, CWORT enables AI-assisted analysis without displacing professional accountability [13]. Within CWORT, agentic AI components are applied to bounded analytical tasks, including summarisation of qualitative inputs, identification of recurring themes, and detection of potential inconsistencies for further review. Their operation is constrained by the workflow architecture: outputs are generated only within predefined stages, presented as decision-support artefacts, and require explicit consultant validation before influencing interpretation or outcomes. Agentic components do not interact directly with clients or control workflow progression, ensuring that AI remains assistive rather than autonomous while preserving accountability [27].

7. Demonstration

This section demonstrates the applicability of CWORT by illustrating how the conceptual model (Figure 1) and application architecture (Figure 2) are instantiated in a representative consulting engagement. The purpose of the demonstration is to show how CWORT can be applied in practice under workflow constraints, rather than to evaluate performance or effectiveness.

7.1. Demonstration context

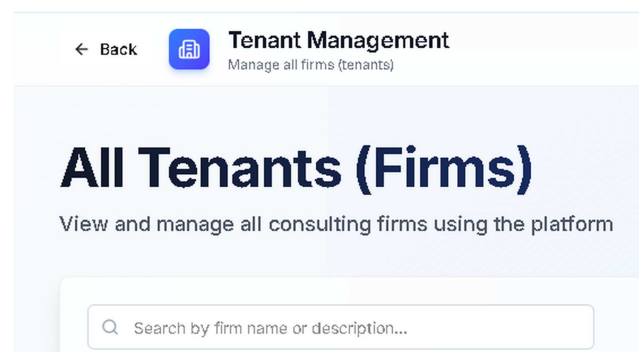
The demonstration is based on a representative consulting engagement involving a mid-sized organisation seeking to assess its digital and operational capabilities. The consulting objective was to identify capability gaps and prioritised improvement actions through structured assessment and expert interpretation. This context was selected as it reflects common advisory scenarios in which qualitative and quantitative inputs must be synthesised into defensible recommendations.

This demonstration is illustrative and is intended to show system instantiation and workflow behaviour; it is distinct from the structured, real-world case-based evaluation presented in Section 8.

7.2. Instantiation of the CWORT conceptual model

For the purposes of the demonstration, the consulting methodology was instantiated as a workflow comprising discovery, assessment, synthesis, and recommendation states as seen in Figures 3 and 4. Each state corresponded to a defined phase of the methodology, with explicit inputs, artefacts, and validation conditions, consistent with the CWORT conceptual model shown in Figure 1. Professional roles were instantiated such that consultants retained authority over workflow configuration, interpretation, and progression, while clients participated through structured input and review activities. Agentic reasoning support was enabled only within selected workflow states and was explicitly excluded from workflow control and recommendation finalisation.

Figure 3
Demonstration of a representative consulting engagement page




7.3. Application of the CWORT architecture

The CWORT application architecture in Figure 2 supported the instantiated workflow by providing orchestration, role-based interaction, artefact management, and decision-support services. Consultants configured the workflow and assessment structure prior to client engagement and controlled invocation of agentic reasoning components. Clients interacted with the system through guided interfaces to provide structured responses and supporting evidence. Agentic components were not exposed to clients and did not generate client-facing outputs, ensuring that all AI-supported reasoning remained mediated by consultants and subject to governance constraints.

7.4. Human-in-the-loop agent-supported analysis

During the synthesis stage, agentic reasoning components were invoked to support analytical activities such as summarising qualitative responses and highlighting recurring themes across artefacts. These outputs were presented as candidate interpretations rather than authoritative conclusions. Consultants reviewed agent-generated outputs and explicitly validated, modified, or discarded them prior to progressing the workflow. At no point did agentic components initiate workflow transitions or

Figure 4
Assessment page



generate recommendations independently. This interaction pattern reflects established principles in human–AI collaboration, where AI augments professional reasoning without substituting human judgement [16].

7.5. Artefact traceability and recommendation formation

Throughout the engagement, CWORT maintained explicit linkage between assessment inputs, agent-supported analysis, consultant interpretations, and final recommendations. Final recommendations were authored by consultants and explicitly connected to underlying evidence and validated reasoning. Agent-generated output remained part of the evidential record but was clearly distinguished from human-authored conclusions as illustrated in Figure 5. This traceability enabled reconstruction of how recommendations were derived and supported transparency and accountability across the consulting lifecycle.

7.6. Demonstration scope and transition

This demonstration illustrates the feasibility of applying CWORT in a representative consulting context and shows how workflow orchestration, governed agentic reasoning, and artefact traceability operate together in practice. This section does not evaluate effectiveness, efficiency, or comparative performance. These aspects are examined separately through a structured, real-world case-based evaluation in Section 8.

8. Evaluation and Observed Outcomes

In DSR, evaluation focuses on determining whether an artefact addresses the identified problem and performs as intended within its application context [14]. Given the early-stage maturity of CWORT, the evaluation presented in this paper is formative in nature and combines qualitative observations with quantitative indicators drawn from a structured case-based evaluation.

Its purpose is to assess feasibility, internal coherence, and alignment with consulting practice, rather than to provide statistical generalisation or performance benchmarking.

8.1. Evaluation approach

CWORT was evaluated through a structured observational case study, consistent with approaches for nascent design artefacts [14]. The case involves an energy distribution network operator subject to the NIS2 Directive, operating fragmented Information Technology (IT) and Operational Technology (OT) security environments.

The evaluation focused on four criteria:

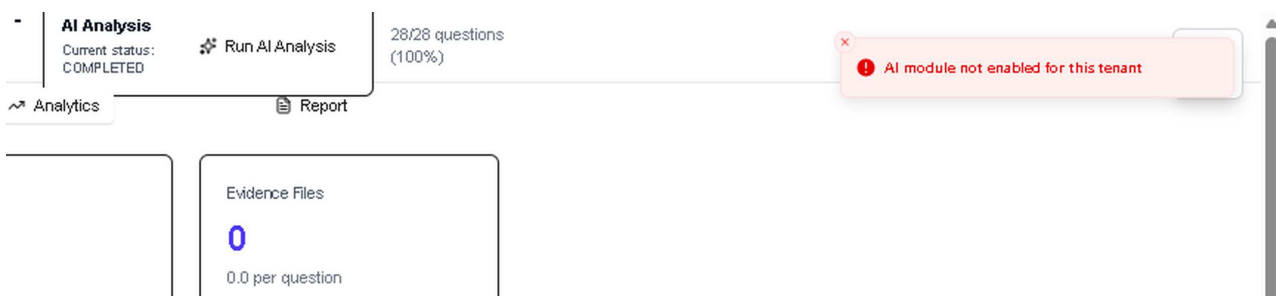
- 1) Workflow coherence—structured execution of consulting activities
- 2) Traceability—explicit linkage between inputs, analysis, and outputs
- 3) Governance and accountability—preservation of human responsibility
- 4) Practical usability—support for, rather than disruption of, consulting work

Data sources included CWORT-generated artefacts (control mappings, maturity scores, remediation plans), workflow execution records, consultant observations from Projectic Solution Ltd, and organisational reporting outputs. Where relevant, findings were compared against historical delivery baselines from prior engagements.

8.2. Observations on AI-supported analysis

CWORT incorporates constrained, agentic reasoning components to support analytical and synthesis activities within defined workflow stages. In the evaluated case, these components were primarily used to summarise qualitative assessment inputs, identify recurring themes, and surface potential inconsistencies for consultant review. Consistent with the system design, all

Figure 5
AI module disabled, requiring human-in-the-loop oversight



AI-supported outputs were treated as decision-support artefacts requiring explicit human validation. Consultants retained full discretion to accept, modify, or reject outputs, and no instances were observed in which agentic outputs were adopted without scrutiny. This reinforces the effectiveness of workflow constraints and validation checkpoints in maintaining accountability.

The usefulness of AI-supported analysis was dependent on the structure and quality of input data, with better-defined inputs producing more relevant summaries and insights. This observation aligns with prior research indicating that AI performance in professional contexts is sensitive to task framing and data quality (Davenport & Kirby, 2016). Overall, AI functioned as an augmentation mechanism for sense-making rather than as an autonomous decision-maker.

8.3. Workflow consistency and traceability

The explicit representation of consulting delivery as a sequence of workflow states enabled structured execution across the engagement. All requirements from Article 21 of the NIS2 Directive were mapped to CWORT's control taxonomy, resulting in the identification of 23 critical gaps across IT and OT environments. Previously unassessed OT systems, including SCADA and ICS environments, were incorporated using IEC 62443-aligned modelling. CWORT maintained a digital thread linking assessment inputs, AI-supported analysis, consultant interpretation, and resulting recommendations. This enabled retrospective reconstruction of decision rationale and strengthened the defensibility of outputs. Consultants reported improved clarity regarding engagement progress, outstanding activities, and required artefacts. While professional judgement remained essential, the structured workflow improved consistency and reduced ambiguity in execution [20].

8.4. Efficiency, governance, and professional accountability

A key outcome of the evaluation was the improvement in delivery efficiency. Prior to CWORT deployment, initial discovery and assessment phases in comparable engagements conducted by Projectic Solution Ltd typically required 8–12 weeks, based on historical delivery timelines. In the evaluated case, the same phase was completed in 3 weeks, representing an estimated 62–75% reduction in delivery time. This reduction was enabled by structured workflow orchestration, predefined control taxonomy, and coordinated Small and Medium-sized Enterprise (SME) input within a single system [6]. In addition, CWORT's dependency intelligence engine prioritised 8 high-impact controls, resolving dependencies for 14 downstream controls, and was associated with an estimated 35% reduction in remediation effort compared to prior manual planning approaches. Governance mechanisms ensured that no workflow transition or recommendation could occur without explicit human validation. Agentic components did not initiate actions or produce binding outputs, preserving consultant accountability. The presence of structured validation checkpoints encouraged more deliberate reflection on assumptions and interpretations, consistent with research on DSS and human-AI collaboration [7].

8.5. Limitations of the evaluation design

The evaluation is based on a single organisational case within a regulated energy-sector context, providing in-depth insight into

CWORT's application but not, on its own, supporting generalisation across diverse settings. The study adopts a within-context comparative field design, in which the same consulting team operated in the same organisational environment using both prior approaches and CWORT. This enables a robust, context-controlled comparison of outcomes while preserving ecological validity. However, the evaluation does not involve controlled experimental conditions, repeated trials, or statistical analysis, and the observed efficiency gains should therefore be interpreted as contextually grounded rather than universally generalisable.

CWORT's architecture separates deterministic control-based analysis from AI-supported augmentation, enabling partial comparison of system behaviour with and without AI support. While the underlying analytical outputs remained consistent, AI contributed to summarisation, explanation, and insight generation. As such, the evaluation supports an assessment of AI's augmentative role but does not isolate causal effects through experimental design. While the present study reports a single structured case, it forms part of a broader, ongoing multi-context evaluation programme across both academic and industry settings. These studies apply DSR evaluation methods and incorporate systematic data collection through system logs, observation, questionnaires, and interviews, designed to enable systematic comparison across contexts and strengthen empirical generalisation. Consistent with DSR methodology, this study represents a formative evaluation of the artefact. The limitations identified reflect the current stage of the research and define a clear pathway for continued empirical validation [14].

9. Design Implications, Limitations of the Artefact, and Future Research Directions

The design and evaluation of CWORT is a set of generalised design implications and principles for operationalising consulting methodology as a socio-technical system can be identified. The findings indicate that effective digital support for consulting delivery requires a shift from loosely coordinated activities to structured, workflow-oriented execution, balancing standardisation with flexibility to preserve professional judgement while ensuring consistent delivery. Governance mechanisms are essential to maintain explicit and auditable accountability for decisions, particularly in regulated or high-stakes contexts, and traceability must be sustained across inputs, analysis, and outcomes throughout the workflow lifecycle. Furthermore, AI capabilities are most effective when embedded as decision-support mechanisms that augment, rather than replace, human interpretation. Collectively, these principles suggest that workflow systems for knowledge-intensive professional work should provide structured progression, embedded governance, end-to-end traceability, and controlled flexibility, offering a conceptual foundation for designing transparent, accountable, and human-centred consulting systems without prescribing specific implementation approaches.

As an early-stage design artefact, CWORT is subject to limitations that warrant acknowledgment. The current evaluation is formative in nature and based on a structured case study of a real consulting engagement, rather than controlled experimentation or large-scale deployment. While this is consistent with early phases of DSR, it constrains the generalisability of findings and precludes definitive claims regarding universal performance improvements or outcome quality [14]. In addition, CWORT has been evaluated primarily in consulting engagements focused on assessment and advisory outcomes; further validation is required to assess

transferability across other advisory domains, organisational contexts, and regulatory environments.

CWORT also incorporates deliberately constrained agentic reasoning components. AI-supported analysis depends on the quality and structure of underlying inputs, and the system does not currently support adaptive learning across engagements or dynamic workflow optimisation. These limitations reflect intentional design choices aimed at preserving accountability and methodological control, but they also represent opportunities for future research. Potential directions include longitudinal and comparative studies examining the impact of workflow orchestration and AI-assisted reasoning on delivery efficiency, consistency, and user experience, as well as technical extensions such as multi-agent coordination, adaptive reasoning informed by prior engagements, and enhanced validation and provenance mechanisms. Any such developments would require careful consideration of governance, transparency, and accountability, consistent with research on human–AI collaboration and socio-technical system design [25]. Future research may also explore the broader organisational implications of operationalising consulting methodology through digital workflow orchestration, including effects on capability development, knowledge reuse, and professional identity in AI-supported consulting environments.

10. Conclusion

This paper introduced CWORT as an innovative approach to operationalising consulting methodology through a digitally orchestrated socio-technical system. Drawing on research in professional service delivery, information systems, workflow management, and DSS, CWORT addresses long-standing challenges in consulting practice related to consistency, traceability, and accountability. The core contribution of CWORT lies in its reframing of consulting delivery as an executable workflow governed by methodological intent rather than as a collection of loosely coordinated activities. Integrating workflow orchestration, role-based governance, artefact traceability, and constrained agentic reasoning within a unified execution model, CWORT enables structured and transparent consulting delivery while preserving the central role of professional judgement.

The study presented the CWORT conceptual model and application architecture and evaluated the artefact through a structured case-based evaluation in a real consulting engagement. Hence, using a within-context comparative field design, the findings demonstrate that workflow orchestration can improve the coherence, traceability, and transparency of consulting activities, while enabling significant reductions in discovery and assessment effort relative to prior approaches. The results also show that AI-supported components can effectively augment analysis through summarisation and insight generation without altering deterministic analytical outcomes or displacing human accountability. These findings provide initial, empirically grounded evidence of the feasibility and utility of workflow-orchestrated consulting supported by governed AI. At the same time, the evaluation reflects a formative stage of the research, and further validation across multiple organisational contexts is required to strengthen generalisability. Ongoing evaluation across academic and industry settings is expected to extend these findings and support broader empirical validation.

In summary, CWORT contributes to both research and practice by offering an integrative model for digitally enabled consulting delivery that operationalises methodology, structures workflow execution, and embeds AI within accountable,

human-governed processes. As organisations increasingly seek to scale advisory capabilities while maintaining trust and transparency, approaches such as CWORT provide a robust foundation for the responsible integration of AI into professional consulting work.

Ethical Statement

This study did not require formal ethical approval. The research concerns the design, demonstration, and case-based evaluation of a socio-technical artefact in a professional consulting setting. No personal data, sensitive organisational data, or identifiable participant information were collected or reported in this paper. All findings are presented at an abstracted level for the purpose of evaluating artefact utility, workflow behaviour, and governance characteristics, rather than assessing individual or organisational performance.

Conflicts of Interest

The author declares that he has no conflicts of interest to this work.

Data Availability Statement

Data sharing is not applicable to this article as no new data were created or analysed in this study.

Author Contribution Statement

Oluwaseun Iyiola: Conceptualisation, Methodology, Software, Writing – original draft, Writing – review & editing, Visualisation, Supervision.

References

- [1] Crişan, E. L., & Stanca, L. (2021). The digital transformation of management consulting companies: A qualitative comparative analysis of Romanian industry. *Information Systems and e-Business Management*, 19(4), 1143–1173. <https://doi.org/10.1007/s10257-021-00536-1>
- [2] Crişan, E. L., & Marincean, A. (2023). The digital transformation of management consulting companies: A review. *Information Systems and e-Business Management*, 21(2), 415–436. <https://doi.org/10.1007/s10257-023-00624-4>
- [3] Mirafzal, M., Fhal, S., Wadhwa, P., & Stal-Le Cardinal, J. (2024). Intelligent competency mapping for improving knowledge management in consulting firms. *Proceedings of the Design Society*, 4, 433–442. <https://doi.org/10.1017/pds.2024.46>
- [4] Zbucnea, A., Dinu, E., Iliescu, A. N., Stăneiu, R. M., & Salageanu, B. R. (2023). Managing knowledge in Romanian KIBS during the COVID-19 pandemic. *Knowledge*, 3(1), 18–39. <https://doi.org/10.3390/knowledge3010002>
- [5] Chapela-Campa, D., & Dumas, M. (2023). From process mining to augmented process execution. *Software and Systems Modeling*, 22(6), 1977–1986. <https://doi.org/10.1007/s10270-023-01132-2>
- [6] Alter, S. (2013). Work system theory: Overview of core concepts, extensions, and challenges for the future. *Journal of the Association for Information Systems*, 14(2), 72–121. <https://doi.org/10.17705/1jais.00323>

- [7] Berente, N., Seidel, S., & Safadi, H. (2019). Research commentary—Data-driven computationally intensive theory development. *Information Systems Research*, 30(1), 50–64. <https://doi.org/10.1287/isre.2018.0774>
- [8] Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., . . . , & Williams, M. D. (2021). Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, 57, 101994. <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>
- [9] Verhoef, P. C., Broekhuizen, T., Bart, Y., Bhattacharya, A., Dong, J. Q., Fabian, N., & Haenlein, M. (2021). Digital transformation: A multidisciplinary reflection and research agenda. *Journal of Business Research*, 122, 889–901. <https://doi.org/10.1016/j.jbusres.2019.09.022>
- [10] Sinkovics, N., Vieira, L. M., & van Tulder, R. (2022). Working toward the sustainable development goals in earnest—critical international business perspectives on designing and implementing better interventions. *Critical Perspectives on International Business*, 18(4), 445–456. <https://doi.org/10.1108/cpoib-05-2022-0059>
- [11] Parelusson, B., Æsøy, V., & Giskeødegård, M. F. (2022). Professions, work, and digitalization: Technology as means to connective professionalism. *Journal of Professions and Organisation*, 9(1), 100–114. <https://doi.org/10.1093/jpo/joab023>
- [12] Norton, L. W., Howell, A. W., DiGirolamo, J. A., & Hayes, T. L. (2024). Using artificial intelligence in consulting psychology. *Consulting Psychology Journal*, 76(2), 137–162. <https://doi.org/10.1037/cpb0000274>
- [13] Feng, F., Wang, J., Yang, X., & Geng, X. (2025). LearnGene: Inheritable “genes” in intelligent agents. *Artificial Intelligence*, 348, 104421. <https://doi.org/10.1016/j.artint.2025.104421>
- [14] Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research I. *MIS Quarterly*, 28(1), 75–106. <https://doi.org/10.2307/25148625>
- [15] Block, P. (2024). *The flawless consulting fieldbook & companion: A guide to understanding your expertise*. USA: John Wiley & Sons.
- [16] Upase, P. B. (2026). Reframing organisational decision-making in the age of artificial intelligence: A conceptual review of human–AI augmentation. *International Journal of Scientific Research in Engineering and Management*, 10, 2279–2290. <https://doi.org/10.55041/ijrsrem.ibfe071>
- [17] Van der Aalst, W. M. (2013). Decomposing Petri nets for process mining: A generic approach. *Distributed and Parallel Databases*, 31(4), 471–507. <https://doi.org/10.1007/s10619-013-7127-5>
- [18] McLeod, S., & Schapper, J. H. (2024). Understanding quality in planning consultancy: A review and synthesis of the literature. *Journal of Planning Education and Research*, 44(1), 39–51. <https://doi.org/10.1177/0739456X20943928>
- [19] Financieras, C. N. (2025). *Accenture signs partner responsible for artificial intelligence in consulting at KPMG [Data set]*. ProQuest. <https://www.proquest.com/docview/3162077473>
- [20] Aras, A., & Büyüközkan, G. (2023). Digital transformation journey guidance: A holistic digital maturity model based on a systematic literature review. *Systems*, 11(4), 213. <https://doi.org/10.3390/systems11040213>
- [21] Benbya, H., Pachidi, S., & Jarvenpaa, S. (2021). Special issue editorial: Artificial intelligence in organisations: Implications for information systems research. *Journal of the Association for Information Systems*, 22(2), 10. <https://doi.org/10.17705/1jais.00662>
- [22] Bresnahan, T. F. (2024). Power and prediction: The disruptive economics of artificial intelligence. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *Business economics* (pp. 120–122). Harvard Business Review Press. <https://doi.org/10.1057/s11369-023-00337-z>
- [23] Davenport, T. H., & Kirby, J. (2022). Just how smart are smart machines? *MIT Sloan Management Review*, 57(3), 21–35.
- [24] Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organisational decision-making. *Business Horizons*, 61(4), 577–586. <https://doi.org/10.1016/j.bushor.2018.03.007>
- [25] Qiu, J., Lam, K., Li, G., Acharya, A., Wong, T. Y., Darzi, A., . . . , & Topol, E. J. (2024). LLM-based agentic systems in medicine and healthcare. *Nature Machine Intelligence*, 6(12), 1418–1420. <https://doi.org/10.1038/s42256-024-00944-1>
- [26] Tyagi, A. K. (2023). *Automated secure computing for next-generation systems*. USA: John Wiley & Sons. <https://doi.org/10.1002/9781394213948>
- [27] Fragiadakis, G., Diou, C., Kousiouris, G., & Nikolaidou, M. (2024). Evaluating human-ai collaboration: A review and methodological framework. arXiv. <https://doi.org/10.48550/arXiv.2407.19098>

How to Cite: Iyiola, O. (2026). Operationalising Consulting Methodology Through Workflow Orchestration: A Design Science Study of CWORT with Governed Agentic AI Support. *Artificial Intelligence and Applications*. <https://doi.org/10.47852/bonviewAIA62028988>