

Of (AI) Machine and Human (Labour): An Integrated Nexus of Work Operating System Architecture for Orchestrating Human–AI Collaborations

Bin Mai | Department of Engineering Technology and Industrial Distribution, Texas A&M University, College Station, TX, USA | ORCID: 0000-0002-2392-1502

Marc Walker | School of Computing, Newcastle University, Newcastle upon Tyne, UK | ORCID: 0009-0004-7395-3723

David Sweeney | Department of Engineering Technology and Industrial Distribution, Texas A&M University, College Station, TX, USA | ORCID: 0000-0002-6530-7471

Abstract

Rapid advances in artificial intelligence (AI) have significantly impacted virtually all aspects of the global society, including especially the labour markets. AI has the potential to transform almost all occupations to various degrees, and in certain labour market segments, has already started to drastically disrupt and reshape the labour market. The optimal approach to the division of labour between humans and AI has become one of the most fundamental focuses to achieve ideal human–AI collaborations. The existing literature in the domain investigates usually at the level of work process maps or job descriptions. In this paper, we present the system architecture of Nexus of Work Operating System (NOW OS), the first integrated end-to-end operation system that analyses human–AI collaboration from the atomic task (the smallest unit of real work tasks) level, and provides comprehensive functionalities of structuring, governing, and optimising the task redesign, task

Received: 28.09.2025

Accepted: 03.02.2026

Published: 26.03.2026

Cite this article as:

B. Mai, M. Walker, D. Sweeney "Of (AI) machine and human (labour): An integrated nexus of work operating system architecture for orchestrating human–AI collaborations," ACIG, vol. 5, no. 1. doi: 10.60097/ACIG/217714

Corresponding author:

Bin Mai, Department of Engineering Technology and Industrial Distribution, Texas A&M University, College Station, TX, US; E-mail: binmai@tamu.edu

 0000-0002-2392-1502

Copyright:

Some rights reserved

(CC-BY):

Bin Mai
Marc Walker
David Sweeney
Publisher NASK



assignment, and task process re-engineering, to achieve the optimal human–AI collaboration for the organisation. We present a conceptual model for analysis of human–AI collaboration that serves as the theoretic foundation for our proposed system architecture. We detail the technical design of the proposed system functionalities, with emphases on cybersecurity and governance functionalities, and our approach to system performance evaluation. In addition, we describe a sample case to illustrate how our proposed system operates in a healthcare organisation, and how our proposed conceptual metrics for performance evaluation can be implemented. The future directions of the technical development and implementation of the proposed system are also discussed.

Keywords

integrated system, division of labour, nexus of work, human–AI collaboration

1. Introduction

The significance of artificial intelligence (AI) in today's global society is vast and fast-growing [1–3]. The rapid advances in AI have produced tremendous evidence that with proper implementation of human–AI collaborations, AI is capable of driving greater efficiency, generating enhanced productivity, enabling advanced decision-making, sparking novel innovation, and, as a result, reshaping a wide array of industries [4, 5]. Rapid advances in AI have significantly impacted almost all aspects of the global society [6], including the labour markets [7, 8]. AI has the potential to transform almost all occupations to various degrees, and in certain labour market segments, has already started to drastically disrupt and reshape the employment market [9–11]. Industry practitioners have increasingly advocated the urgent need of an optimal approach to leveraging the unique strengths of AI by integrating AI into their job task performances and business processes implementation, and optimising the human–AI collaborations in terms of value generation for the organisation [12–14].

In the meantime, people have started identifying obstacles to the optimisation of human–AI collaborations. Academic researchers have long been interested in investigating the impacts of emerging technologies on the nature of work and labour market, and have applied their analyses to the case of AI [15]. Specifically for AI technology, however, due to the technology's rapid, constant, and continuous advancement, and its potential to fundamentally change

the nature of business processes and work tasks, to fully understand AI's impacts on nature of work and labour market we must address some unique challenges [16, 17]. Frank et al. have laid out the main barriers facing the researchers in measuring the impacts of AI on the future of work [18]: sparse skills data; limited modelling of resilience; and places in isolation.

Sparse skills data refers to the current situation where sufficient granularity of the job task data and job skills data is difficult to obtain, and the high-resolution mapping between the task and skill data is lacking. This barrier makes it difficult to produce a concise mapping between job tasks and skill requirements, hinders the analysis of the impacts of AI on specific job tasks, and thus fundamentally impedes the derivation of optimal division of labour between human and AI. Limited modelling of resilience refers to the situation where there lacks a comprehensive model that illustrates the dynamics between human labour and AI labour in terms of business process implementation and value generation for the organisation. This barrier makes it difficult in assessing the various approaches to division of labour between human and AI. Places in isolation refers to the lack of spatial as well as temporal data and analytical model to investigate how AI impacts human labour over time or how it impacts different geographic locations on a global scale differently. It is worth noting that within an organisation context, the first two main barriers (sparse skills data and limited model of resilience) are the most directly relevant. They are also the focal points of this paper.

It has become increasingly obvious that to optimally implement human-AI collaborations within the organisation context, we need to develop a system that addresses the two main barriers, and integrates AI at both job task performance and business process implementation levels. In addition, throughout the implementation of human-AI collaborations, cybersecurity and governance have been identified as two most significant elements that determine the outcome of the human-AI collaboration [19–21]. Therefore, we intend to address the following research questions in this paper: (1) what is a good system architecture for orchestrating labour between humans and AI agents that overcome the identified barriers?; (2) how does the system implement cybersecurity and AI governance to ensure successful human-AI collaboration?; and (3) what would be the proper approach to evaluating the performance of the proposed system?

In this paper, we propose Nexus of Work Operating System (NOW OS), a modular, execution-grade operating system architecture for

orchestrating labour between humans and AI agents within the context of organisation, as solutions to the research questions. We describe in detail the NOW OS architecture, components of NOW OS, their functionalities, their performance evaluation approach, and their mechanisms of implementing cybersecurity and AI governance. In this paper, we follow the design science/system-oriented research paradigm, and aim to provide technical design details of the NOW OS, including its functionalities (with emphasis on the cybersecurity and governance related features), its performance evaluation approach, and its implementation methods.

The rest of the paper proceed as follows: in Section 2, we present our conceptual model for the analysis of human–AI collaborations; followed by Section 3, where we address all our research questions. We address research question 1 by detailing the components of NOW OS and illustrating how those components are aligned with our conceptual model and achieve the goal of optimal human–AI collaborations; we address research question 2 by specifying the NOW OS component features that address the cybersecurity and AI governance issues; we address research question 3 by explaining the proper approach to NOW OS performance evaluation. In Section 4, we describe a sample case of implementation of NOW OS in a healthcare organisation, with explanations of how NOW OS functions to improve the operation of the organisation, how system performance will be evaluated, and how cybersecurity and AI governance have been implemented. We conclude in Section 5.

2. Conceptual Framework

We suggest that to develop a comprehensive framework to analyse the human–AI collaborations in today’s global labour market, we need to specifically address the human–AI collaboration issues from three levels: micro-, meso-, and macro-level. We propose our conceptual model of the analysis of human–AI collaboration, as illustrated in Figure 1.

In Figure 1, we notice that there are three levels of analyses for us to obtain a comprehensive understanding of human–AI collaboration and overcome the identified barriers to optimal human–AI collaborations:

- At the micro-level, we need to systematically analyse the relationships among specific job tasks, specific skill requirements, and business process compositions. The job tasks are closely related to the organisation’s business processes and corresponding

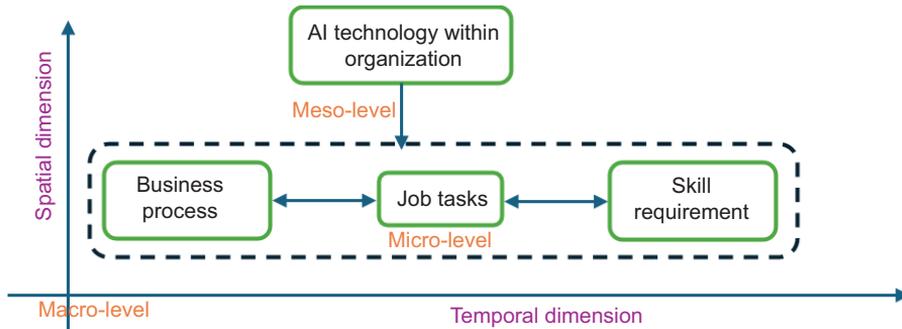


Figure 1. Conceptual model of analysis of human-AI collaboration.

position structure. We need to adopt a systematic approach to connecting the specific job tasks to specific skill requirements, and produce the outcome of task-skill relationship mapping. In addition, the systematic approach should also allow us to determine how the individual job tasks would comprise the business processes. To achieve this outcome, we need to adopt a methodology to break down the job tasks to its smallest possible level as well as a framework for matching the task to required skills. Atomic task decomposition provides the foundational building blocks for the subsequent analyses, and precise mapping between the atomic tasks and skill requirements provides foundations for modelling human-AI collaborations. This approach addresses the sparse skill data barrier. Such analyses provide the foundation for business process re-engineering based on the skill-task mapping. More specifically, based on micro-level analyses, we need to systematically derive the following outcomes:

- For the organisation's existing business processes and job tasks, we determine which ones should be modified (and how). For the transformed job tasks, we then determine which ones should be assigned to human labour, and which ones should be delegated to AI agents. This analysis allows us to re-configure the business processes to enhance alignment with the organisation's value system. We provide detailed rationale for our proposed approach.
- Based on the properties of the AI technology, we analyse if and how new job tasks and, subsequently, new business processes should be created to enhance the organisation's overall value generation. Based on the analyses, we are then able to describe the properties of the proposed new job tasks; and for each of the new job tasks, we are able to determine whether human labour or AI agents should be assigned the specific task.

- At the meso-level, we need to analyse how the AI technology modifies the relationships among job tasks, skill requirements, and business processes within the organisation context, focusing on alignment with the organisation's value system. The analyses at the meso-level critically depend on the analyses conducted and insights obtained from the micro-level. We would identify the properties of AI technology, the characteristics of the organisation's value system, and the properties of the business process-task-skill relationship through the micro-level analyses. We would then adopt a systematic model to illustrate the dynamics among AI, organisation, and human factors to analyse how the AI technology could impact the business process-task-skill relationships within the organisation context to maximise value generation for the organisation. This approach enables us to address the limited modelling of resilience barrier.
 - Not only is it important for us to be able to derive the optimal strategies based on the modelling and analyses we conduct at the meso-level but it is crucial also that we can design and develop a system to operationalise the strategies we derive, to convert our thorough understanding of human-AI collaboration into realistic opportunities to optimise human-AI collaborations for real organisations.
 - It is critical to note that the meso- and micro-level analyses of human-AI collaboration are closely integrated with each other. The micro-level analyses provide fundamental data and insights to inform the meso-level analyses. In the meantime, the meso-level analyses incorporate strategic level of guidance for the micro-level analyses.
- At the macro-level, we need to analyse the impacts of AI on job composition and labour market from these two dimensions:
 - Temporal dimension: We analyse how such impacts of AI evolve over time. To achieve this goal, we need to analyse the evolution of AI technology in the long term, and how our micro- and meso-level analyses results can be adapted and evolve together with the AI technology evolution.
 - Spatial dimension: We analyse how to expand our analyses to multiple geographic locations and study the global dynamics of how AI technologies may impact job task and labour market at different locations in different ways.
 - As described earlier, in this paper, we focus on the impact of AI on human-AI collaborations within the context of individual organisation. Therefore, this macro-level analysis is beyond this paper's scope. We, however, discuss this dimension in Section 5, where we describe possible directions for further research.

Based on the conceptual model, we note that to develop an effective system for an organisation to orchestrate optimal human–AI collaboration, the system must provide functionalities that satisfy the requirements and address the challenges as indicated in the micro- and meso-levels. In this paper, we propose NOW OS, an enterprise-wide system architecture designed for optimisation of division of labour between human and AI to achieve maximum value generations for an organisation.

Next, we describe in detail the basic architecture and technical design of our proposed NOW OS, and illustrate how it achieves the goal of orchestrating the optimal human–AI collaborations and integrates the analyses at micro- and meso-level of our conceptual model.

3. Nexus of Work Operating System

3.1. Basic structure and components of NOW OS

To address research question 1, we propose NOW OS. At its core, NOW OS is an end-to-end operating system for structuring, governing, and optimising the division of labour between humans and AI to obtain optimal business processes configuration for the organisation to maximise its value generation. The design of NOW OS architecture is grounded in the general human–AI teaming (HAT) framework [e.g. 22, 23], which posits that human–AI collaborative systems can be characterised as a joint cognitive system [24] in which two cognitive agents collaborate. Human labour intelligence and AI machine intelligence (MI) achieve intelligent complementarity through their collaboration in performing work tasks that constitute business processes for the organisation, and achieving performance goals in alignment with the organisation’s objectives.

This overall goal is achieved by the two core components of NOW OS: the human–AI division of labour blueprint, and the Unified Operating Model and Management System (UOMAMS). At a high-level abstraction, the labour blueprint and the UOMAMS integrate as illustrated in Figure 2.

3.1.1. The human–AI division of labour blueprint component

This component is located at the lower level of the NOW OS structure, and performs the detailed functionalities of analysing tasks, skills, and business processes. This component directly addresses the challenges identified in the micro-level analyses as

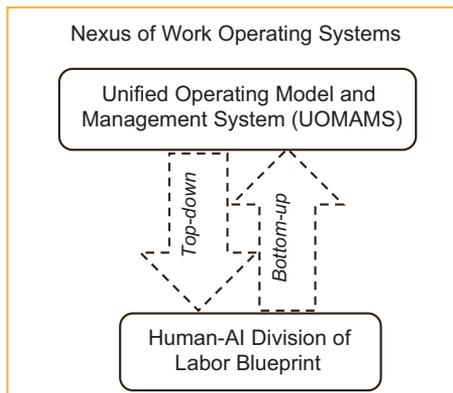


Figure 2. Core components of nexus of work operating system (NOW OS).

shown in our conceptual model. Some unique characteristics in the design of the blueprint components include the following:

- Unlike traditional frameworks that stop at process maps or job descriptions, this component breaks down the job tasks to the level of ‘atomic tasks’, which are defined as the smallest actionable units of work. The critical advantage of this atomic task deconstruction is that now we obtain structured, validated, and unambiguous definitions of job tasks, which provide the foundation for accurate task-skill mapping and optimal business process reconfiguration.
- With the blueprint component, once atomic tasks are captured, each is analysed and scored across a configurable set of ‘logic domains’ (e.g. risk exposure, data sensitivity, business impact, and AI capability match) to generate a task assignment recommendation. These scores and analyses determine whether the specific task is still necessary, whether a new task is needed (and how the new task should be defined), and how – and by whom – each task should be executed, with human oversight layers and escalation paths baked into the output.

This component breaks down work into the smallest meaningful, actionable units (atomic tasks) that cannot be further divided without losing clarity or purpose [25]. Each atomic task would be: clear (one objective, no ambiguity), self-contained (not dependent on multiple other tasks to start), actionable (someone, either a human or an AI agent, can do it), and measurable (performance metrics can be developed for its evaluation) [26, 27].

In designing the functionality of atomic task breakdown, we follow the principles of process mining [28] that focus on end-to-end business processes, extract event, and job task data throughout the organisation's business operations, and perform process discovery, conformance, and enhancement. We develop the overall ontology for the blueprint component that integrates: (1) task ontologies; (2) skill ontologies; and (3) learning and development ontologies. The formal ontology schema is implemented to create the organisation's knowledge graphs, which illustrate relationships among specific work tasks (obtained from process and task mining), skill requirements (obtained from established skill framework, such as Skills Framework for the Information Age [SFIA]), and job positions (obtained from organisation sources). The atomic task breakdown provides foundation for visualising tasks at various granularity levels, and reflects the real-time business process composition for the organisation.

In designing the functionality of task assignment recommendations, we follow the standard human-automation allocation framework such as described by Parasuraman et al. [29], which proposes four broad classes of function automations at various levels (information acquisition, information analysis, decision and action selection, and action implementation) and Horvitz [30], which suggests the mixed-initiative approach to explicitly support an efficient, natural interleaving of contributions by human users and automated services aimed at converging on solutions to problems. We construct an analytical model based on these frameworks, and formulate the task assignment recommendations decision as a constrained optimisation problem. We incorporate cost-benefit parameters (derived from the organisation's value-generation objectives) in our objective function, and optimise the function subject to a series of internal and external constraints (including technology constraints, organisational constraints, governance and compliance constraints, and resource constraints). The solution of this optimisation problem forms the basis of the algorithm for the assist/augment/automate (AAA) decision-making process of the task assignment. The algorithm allows the organisation to develop human-AI collaboration policy functions that incorporate thresholds, confidence intervals, and human-in-the-loop escalation rules.

In summary, the blueprint component goes far beyond a reference architecture or strategy, it's a fully operationalised labour orchestration system for organisations to 'know, govern, and optimise' every unit of work in the AI era. By turning work into structured, scored, and versioned 'atomic tasks', that act as 'nodes' and assets

just as other assets are tracked and managed, the blueprint makes it possible to blend human judgement with machine efficiency in a manner that is auditable, traceable, and continuously self-improving. We are basically building the ‘digital twin of works’ for the organisation. This blueprint provides the organisation with the ability to visualise work tasks with various granularity, laying the foundation for higher executive decision-making support.

3.1.2. The UOMAMS component

This component resides at a higher level of NOW OS, establishing the multi-domain, strategy-to-operations framework that governs all major aspects of an enterprise’s overall structure and value system. This component directly addresses the challenges identified in the meso-level analyses as shown in our conceptual model. UOMAMS aligns an organisation’s strategic ‘why’ and ‘when’ with hands-on ‘how’ and ‘who’ across horizontally multiple critical organisation domains and across vertically multiple organisation implementation layers. UOMAMS transforms abstract strategy into executable, measurable operations.

In designing the UOMAMS layer, we emphasise the alignment between NOW OS and the major AI governance/compliance frameworks, including National Institute of Standards and Technology (NIST) AI risk management framework, European Union (EU) AI Act, and ISO 42001. Such alignments ensure NOW OS to achieve data protection, accountability, risk management, transparency, auditability, and human oversight as required by the frameworks. This design allows us to address research question 2, especially the AI governance aspect. UOMAMS is divided into distinct domains, so that each can define and determine the appropriate response to an AI governance requirement. The business domain has primary responsibility for assessing the impact of legislation and standards, determining the business risk of an incident and deciding upon the appropriate mitigation across some or all of the other domains.

For example, consider the impact of the EU AI Act. The Act introduces a mandatory risk-based governance framework. Each domain must now permanently include ‘AI compliance’ as a function, similar to general data protection regulation (GDPR) privacy compliance. Instructions for strategic and tactical inclusion must be communicated through each UOMAMS layer. Organisations developing AI products need a dedicated AI Governance Committee or officer to oversee compliance and maintain the ‘system inventory’. The software development lifecycle (SDLC) is no longer sufficient for AI. The operating model must shift to an AI lifecycle approach

where technologies, processes, services, and information are assessed for compliance and participation in the AI lifecycle. The leadership of each domain is responsible for particular compliance requirements and should be audited periodically for reassessment.

Another example is NIST's AI Risk Management Framework (AI RMF), which is a voluntary standard emphasising AI system's trustworthiness and cultural aspects, including the governance of AI's human impacts (e.g. deskilling, surveillance, fairness in task routing, and labour relations). NOW OS' governance structure is aligned with NIST AI RMF's four core functions: govern, map, measure, and manage. The business domain determines how specific governance requirements, such as social sustainability (e.g. equitable task assignment, workload fairness, explainability access for workers, and grievance processes), are integrated across all areas of the business processes.

UOMAMS component enables policy-aware orchestration [31], as well as role-based operating views [32], providing a direct interface for NOW OS to allow the organisation's decision makers to customise their 'views' of the organisation's operations. This component empowers leadership to cascade strategy into every domain, equips designers and architects with proven patterns and reference catalogs, gives managers real-time performance visibility, and provides operators with precise, day-to-day execution guides. In summary, UOMAMS transforms organisational complexity into a transparent, governed, and continuously optimised system, driving faster decisions, higher quality outcomes, and sustainable agility.

3.1.3. The Dynamic interaction between the two components

At a high level, the blueprint and UOMAMS form two interlocking layers of a single cohesive organisational architecture:

- The upper layer of UOMAMS component establishes a multi-domain, strategy-to-operations framework that governs every aspect of an enterprise's structure across multiple organisation implementation layers. It provides guidance for the lower layer of blueprint component to enhance its analyses. It generates role-specific 'views' for various organisation decision-makers, including executives, designers, managers, and frontline operators.
- The lower layer of blueprint component works as the granular 'engine' that dissects, scores, routes, and continuously optimises every unit of work (atomic task) across humans and AI agents. It populates and feeds the unified model component's various

domains with real-time task-level data, AAA decisions, and compliance metadata.

Integrating atomic task modelling, policy-aware orchestration, and role-based operating views creates a highly efficient, compliant, and adaptable organisational operating model. Atomic task modelling breaks work into precise, measurable units; policy-aware orchestration ensures that these tasks follow rules, constraints, and compliance requirements; and role-based views ensure that each person sees and executes only the tasks relevant to their responsibilities. Together, they increase clarity, accountability, and auditability; enable safe and scalable automation; improve performance monitoring; reduce operational risk; and make processes more agile and easier to update as policies, organisational structures, or regulatory requirements change [33–35]. Synergised, the two components produce the NOW OS that operates dynamically: it integrates new data continuously, conducts real-time analyses, and adjusts configurations simultaneously. These functionalities enable NOW OS to serve as the organisational backbone for operations in the AI era.

In addition, while designing the NOW OS architecture and its components, we have emphasised the dimension of information security throughout the whole process. To address the cybersecurity aspect of research question 2, we firmly believe security is not a separated and isolated component or a layer in the system. Instead, information security must be woven into all elements of the system. To ensure security features are built-in at both the component level and the system level, we adopt a structured combination of secure design principles, engineering practices, and governance frameworks.

We adopt the secure-by-design architecture's principles of least privilege (each component gets only the access it needs); defence-in-depth (multiple layers of security); fail-safe defaults (deny by default); separation of duties (break up responsibilities); minimise attack surface (reduce inputs, services, and permissions); and zero-trust architecture (never trust, always verify). Specifically, at the system component level, we define component-specific security requirements, conduct threat modelling per component, implement secure coding and automated scanning, and conduct testing and dependency checks. At system level, we adopt the formal framework of NIST SP 800-160 to define organisational and system-level security controls, define enterprise-wide security architecture patterns, specify system-wide identity, logging, encryption, and implement end-to-end pen-tests and monitoring. This combined strategy

guarantees that not only is each component secured individually but the system as a whole also has integrated, layered, and consistent security.

3.1.4. Comparison of NOW OS architecture with other RPA/IDP orchestration and governance solutions

Robotic process automation (RPA) and intelligent document processing (IDP) technologies have been integrated and managed within various existing process orchestration platforms (e.g. UiPath, Appian). Empirical studies and quantitative models [36, 37] have been proposed and conducted to evaluate the performance of the human-automation teaming approaches, including the failure modes in automation governance. NOW OS represents a paradigm shift in orchestration. It adopts many proven elements of RPA frameworks – centralised control, audit logging, and role security – but improves upon them by making the granularity finer, the intelligence higher, and the governance tighter. Crucially, it also rejects the limitations of first-generation automation (rigid processes, siloed bots, and neglect of human factors). The result is an orchestration framework where every task is optimally routed, fully traceable, policy-compliant, and aligned with organisational strategy – a vision that addresses the key shortcomings of today’s RPA platforms. More specifically, NOW OS provides unique advantages from the following perspectives:

- **Task granularity:** NOW OS operates at the level of atomic tasks, treating the smallest actionable units of work as first-class objects that can be independently observed, governed, and re-composed dynamically across humans and AI. In contrast to RPA/IDP systems that orchestrate coarse-grained workflows or transactions, this task-level decomposition enables fine-grained delegation, adaptive reconfiguration, and partial automation within a single process without brittle end-to-end rewrites.
- **Provenance:** NOW OS embeds end-to-end provenance directly into the task lifecycle, capturing lineage, evidence sources, confidence signals, and actor attribution for every atomic task via the lens and immutable ledger. Unlike traditional RPA/IDP audit logs that primarily record execution events, NOW OS enables explainable decision-level traceability that supports regulatory scrutiny, model accountability, and *post hoc* justification of human-AI allocation decisions.
- **Governance hooks:** To address the governance aspect of research question 3, we enforce governance in NOW OS by designing at the task level, using a customer-governed AAA allocation framework that encodes compliance constraints, risk thresholds,

jurisdictional rules, and organisational policy directly into orchestration logic. This contrasts with RPA/IDP platforms where governance is largely externalised to process design conventions or post-execution controls, limiting real-time policy enforcement during task routing.

- **Organisational change:** NOW OS explicitly couples task orchestration with organisational change governance, using a strategic advisory board to continuously monitor and review task classifications, delegation thresholds, and workforce impacts as conditions evolve. Rather than treating change management as an external program layered onto automation initiatives, NOW OS operationalises change as a persistent decision process grounded in task evidence, workforce readiness, and accountable human oversight.

With the basic understanding of the NOW OS infrastructure, its core components, and their respective functionalities, we now describe the technical details of the NOW OS core component design and explain how each component performs its specific functionalities.

3.2. NOW OS core component: Human-AI division of labour blueprint

The blueprint component is the reference architecture for NOW OS. It is an infrastructure-grade framework that defines how organisations can decompose, govern, and orchestrate every unit of work – down to its most atomic actions – across humans and AI agents. At its core, it treats ‘tasks’ as the foundational currency of labour: each task is mapped, scored, and routed according to policy, risk, and capability metadata. It provides foundation for NOW OS to transform it from a conceptual model into a live, governable ‘labour operating system’ that drives real-time execution, auditing, and continuous optimisation. The technical architecture of the blueprint component is summarised in Figure 3. We notice that the Human-AI division of labour blueprint component consists of three lanes: observe, orient, and decide & act, as visualised from top to bottom in the figure.

3.2.1. The ‘Observe’ lane

This lane contains an atomic task model. Instead of relying on static job descriptions or process flows, this lane breaks every role down into nested ‘clusters’ of responsibilities, tasks, sub-tasks, and then ‘atomic tasks’ down to individual atomic actions.

On this lane, data is collected from the organisation about their business processes, job tasks, and skill requirements. The data sources include organisation documentations, interviews, surveys,

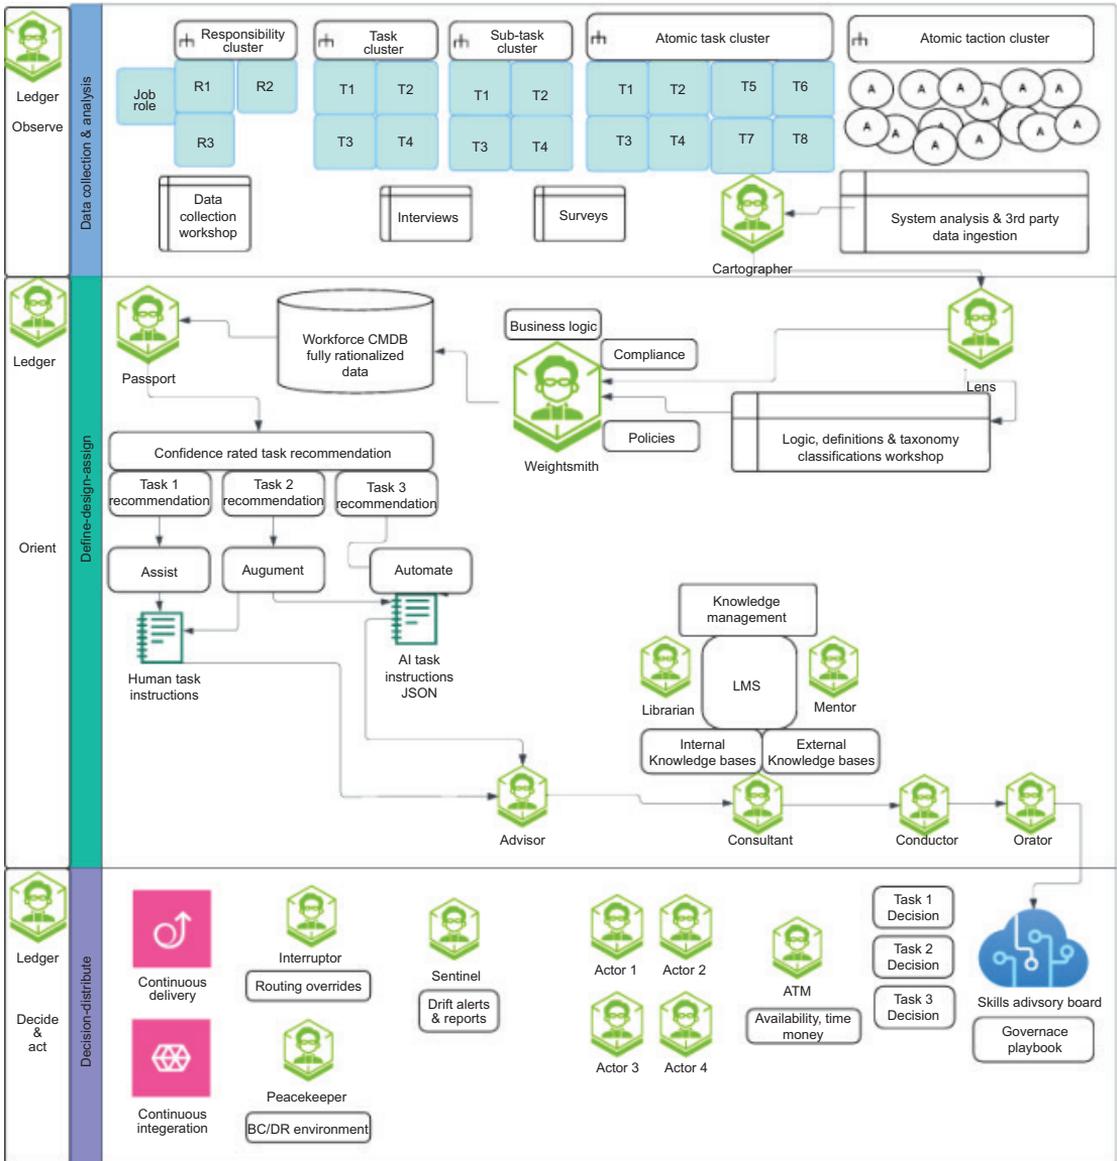


Figure 3. Technical architecture of human-AI division of labour blueprint.

and the third party data sources. The collected data is analysed to generate the atomic action cluster. Obviously, this functionality is labour-intensive. Therefore, in this lane, we design an AI agent named ‘cartographer’ (represented by a green hexagon icon in the ‘observe’ lane) to perform this crucial functionality.

- ‘Cartographer’: The cartographer AI agent solves the critical business problem of organisational workflow invisibility by

Table 1. Cybersecurity features built in in NOW OS AI agents.

Agent feature	Security specifications
Audit logging with cryptographic integrity	Creates tamper-proof records using hashing, blockchain-style hash chaining, and digital signatures that provide legally defensible audit evidence and enable immediate detection of unauthorised modifications.
Role-based access control (RBAC) with multi-factor authentication	Implements authentication with attribute-based permissions, ensuring only authorised personnel can access sensitive functions while maintaining compliance with enterprise security policies.
End-to-end encryption for data at rest and in transit	Applies encryption for sensitive stored data and encryption for all system communications, protecting confidential information throughout its lifecycle, and preventing unauthorised interception.
Zero-trust architecture with continuous identity verification	Implements continuous verification for all system interactions and validates every access attempt against identity governance policies, eliminating implicit trust and reducing attack surface. This feature addresses data-minimisation and purpose-limitation in workflow mining.
Emergency lockdown and automated security incident response	Provides rapid response capabilities to security incidents, including unauthorised access attempts, immediate access lockdown, forensic evidence preservation, and coordinated incident investigation while maintaining service availability for authorised users.

continuously mining diverse data streams (e.g. organisation documentation, such as email, calendar, system logs, and endpoint activities) to automatically map actual work patterns, decompose complex processes into atomic tasks, and identify automation opportunities that remain hidden from traditional process discovery methods. Cartographer integrates with other layers, serving as the foundational intelligence engine that transforms raw organisational data into actionable workflow insights.¹

In the technical design of the cartographer agent (as well as of all other agents we will describe later), we emphasise the secure-by-design architecture's principles, and incorporate a series of cyber-security features consistently and prominently throughout the AI agent ecosystem, as summarised in Table 1.

These features are implemented throughout the agent ecosystem, especially for those directly interacting with sensitive data, such as lens, ledger, sentinel, interruptor, and PeaceKeeper. They support regulatory frameworks such as SOX, GDPR, HIPAA, and ISO.

3.2.2. The 'Orient' lane

This lane obtains the feed of atomic tasks captured from the 'observe' lane, and analyse them to generate an AAA recommendation. Each atomic task is scored across a configurable set of

¹—In Appendix 1, we provide the detailed requirement report of "Cartographer" agent technical specifications. The requirement reports for all other agents discussed in the paper are available upon request.

logic domains (e.g. risk exposure, data sensitivity, business impact, and AI capability match). These scores determine how, and by whom, each task should be executed, with human oversight layers and escalation paths baked into the output.

The hierarchy of feeds from the 'observe' lane are aggregated into a centralised configuration repository called the workforce Configuration Management Database (CMDB), which resides on this lane. Workforce CMDB becomes the single source of truth for who does what, how, why, and under which conditional rules. This lane is also where the workforce CMDB data is analysed to determine recommendations. This sophisticated process requires collaborations from multiple AI agents. We explain the functionalities of each one of the AI agents deployed in this lane in details below.

- 'Lens': The lens AI agent maps every decision or insight to its underlying data sources – tagging which tables, application programming interfaces (APIs), or documents informed that decision. It solves the critical business problem of data lineage and audit compliance by automatically attaching provenance metadata and quality scores to every atomic task within the organisation, ensuring that both downstream agents and external auditors can definitively answer 'where did this data come from and how trustworthy is it?' with measurable outcomes and providing complete audit trails for regulatory compliance. Lens is positioned immediately after the cartographer agent, where it adds the unique capability of enriching every task with standard provenance and data quality assessments that enable transparent, auditable decision-making across the entire agent ecosystem while feeding trustworthy metadata to risk assessment and downstream processing agents.
- 'Weightsmith': The weightsmith AI agent applies scoring weights and AAA thresholds per persona or regulatory zone. It solves the critical problem of organisational bias, imbalance, and inconsistency in how decisions, tasks, and data inputs are prioritised across the enterprise. Weightsmith introduces a standardised, auditable weighting framework that dynamically calibrates the relative importance of signals, tasks, risks, and outcomes against enterprise objectives, ensuring decisions are data-driven, consistent, and explainable. This delivers measurable outcomes in the form of improved alignment between strategy and execution, reduced decision latency, and quantifiable governance assurance by tracing how weights impact actions and outputs. It acts as a balancing and calibration engine across the AI agent ecosystem, ensuring that other agents' analyses and decision-making are all

- grounded in consistent, transparent, and adjustable weighting models aligned with enterprise governance and business drivers.
- ‘Passport’: The passport AI agent generates persona passports and individual human work profiles, linking organisational roles (e.g. ‘AP clerk’) to required capabilities and preferred AAA modes. It solves the core business problem of ensuring that the right tasks are always assigned to the right performer – human or AI – by maintaining a dynamic, enterprise-wide identity, skills, and eligibility registry. Passport can eliminate inefficiencies by continuously validating credentials, mapping skills to task requirements, and enforcing access and eligibility policies in real time. The measurable outcome it delivers is a reduction in task misrouting, improved compliance with governance standards, and optimised utilisation of human and AI resources, directly translating to faster task turnaround times, reduced operational risk, and higher workforce productivity. Passport acts as the authoritative skills, credentials, and permissions registry that feeds into task orchestration decisions. Its unique capability is to serve as the identity and eligibility backbone of the agent ecosystem, ensuring that every action taken by humans or agents is properly validated, compliant, and aligned with enterprise governance requirements.
 - ‘Advisor’: The advisor AI agent solves the critical business problem of translating complex risk assessments and compliance data into clear, actionable recommendations that non-technical stakeholders can understand and implement confidently within their organisational workflows. Advisor serves as an intelligent bridge between Weightsmith’s quantitative risk scores and passport maker’s compliance constraints by generating human-readable and machine-executable action plans that specify which steps should be automated, augmented, or handled by humans with complete rationale and governance transparency.
 - ‘Consultant’: The consultant AI agent addresses the business problem of fragmented, slow, and inconsistent strategic and operational decision-making by transforming disparate enterprise data, governance constraints, and stakeholder inputs into consulting-grade recommendations, business cases, and playbooks that are transparent, risk-aware, and directly tied to measurable outcomes, such as reduced decision cycle time, improved return on investment–net present value (ROI/NPV) accuracy, and higher alignment between initiatives and enterprise strategy. It ensures all recommendations are evidence-linked and governance-compliant. Its unique capability is serving as the decision foundry of the ecosystem – synthesising option sets, modelling trade-offs, and generating implementation blueprints – so that

- humans and downstream agents operate with clear, traceable, and optimised direction.
- ‘Librarian’: The librarian AI agent aggregates, indexes, and curates best-practice entries and lessons learned, availing them to agents (mentor) and human users (process owners). It solves the critical organisational challenge of knowledge fragmentation and decay, where vital business knowledge, including standard operating procedures, best practices, training materials, and domain expertise, becomes siloed, outdated, or inaccessible, leading to duplicated effort, compliance risks, and suboptimal decision-making across the enterprise. The agent delivers measurable outcomes through automated knowledge curation, reduces knowledge discovery time through semantic search capabilities, and ensures 100% audit compliance through immutable provenance tracking and version control of all organisational knowledge assets.
 - ‘Mentor’: The mentor AI agent provides just-in-time tutoring, rationale explanations, and micro-learning nudges whenever a human is in ‘assist’ or ‘augment’ mode. It solves the business problem of fragmented employee development and inconsistent upskilling by providing a structured, AI-driven guidance system that continuously aligns individual learning paths with organisational priorities, compliance requirements, and emerging role demands. It delivers measurable outcomes in the form of improved workforce readiness, accelerated time-to-competency for new skills, higher employee engagement, and reduced training waste through personalised, just-in-time recommendations. Mentor enables context-aware skill guidance and adaptive coaching tied directly to tasks, projects, and workflows, while feeding progress and performance data into the other layers for governance and measurable impact tracking.
 - ‘Conductor’: The conductor AI agent routes each atomic task (and its AAA modes) to the appropriate handler, whether human or AI. It solves the critical organisational challenge of optimal task routing and resource allocation by intelligently analysing risk scores, persona passports, and real-time capacity metrics to make data-driven execution decisions that maximise resource utilisation, ensure policy compliance, and minimise bottlenecks, delivering measurable outcomes, including improved task completion rates, enhanced resource utilisation, and maintaining policy compliance. Conductor serves as the ‘brain of execution routing’ that bridges the analytical insights with actual task execution, adding the unique capability of intelligent route simulation and policy-aware decision making that transforms risk assessments and recommendations into actionable execution plans while maintaining continuous oversight integration with other agents.

- ‘Orator’: The orator AI agent solves the critical business problem of executive decision-making bottlenecks by transforming complex, multi-source task recommendations and analytical outputs into high-integrity, data-dense yet comprehensible presentation formats that enable the skills advisory board to rapidly consume and evaluate wide swaths of organisational intelligence for informed allocation decisions. Orator serves as the final presentation and explainability engine that packages outputs from consultant and other upstream agents into traditional executive formats, including PowerPoint presentations, executive dashboards, and comprehensive reports while posting summarised narratives to ledger for immutable audit compliance.

3.2.3. The ‘Decide & Act’ lane

This lane’s functionalities are built on the foundation of the previous two lanes, and convert the data and analyses from the other lanes into actionable business decision-making. This lane interfaces with the workforce CMDB, which implements the decisions/recommendations generated in this lane. This lane is also where critical and supporting functionalities, such as risk management and governance, are performed. The AI agents deployed in this lane include the following:

- ‘Automated task manager (ATM)’: The ATM AI agent is the automated dispatcher for all tasks within NOW OS, solving the business problem of inconsistent, manual, and error-prone task distribution across human and AI performers. Its core purpose is not to automate the tasks themselves but to automate the process of routing every task – routine, exception-based, or flagged – to the correct executor. ATM evaluates incoming tasks using passport’s authorisation and eligibility rules, the AAA logic, and real-time availability and workload data. It then dispatches each task to an appropriate human operator (credentialed and approved), an AI agent (responsible for a particular domain), or a hybrid path that ensures oversight and escalation options remain available. By functioning as the single automated dispatcher, ATM ensures that every task is routed correctly from the moment it enters the NOW OS ecosystem, removing reliance on *ad hoc* manual assignment and guaranteeing that governance, auditability, and accountability are always applied. Its unique capability is being the enterprise’s automated routing engine, ensuring that the flow of work is optimised, governed, and consistently executed across both humans and agents, making it the control point that underpins scalable and secure human–AI collaboration.

- ‘Sentinel’: The sentinel AI agent serves as the comprehensive performance monitoring and drift detection system responsible for ensuring the reliability, accuracy, and operational integrity of the entire NOW OS multi-model agentic solution by continuously evaluating model performance, detecting data drift, and identifying anomalies before they impact business operations. The sentinel agent provides cross-cutting observability across all main architectural layers and adds the unique capability of predictive performance monitoring with statistical drift detection that enables proactive system optimisation and prevents performance degradation before it affects service delivery.
- ‘Interruptor’: The interruptor AI agent solves the critical business problem of organisational risk escalation by serving as the real-time guardian that monitors service level agreement (SLA) clocks, anomaly alerts, and policy breaches to automatically freeze tasks, halt data mutations, and alert PeaceKeeper (explained next) and on-call humans when safety thresholds are crossed, preventing operational failures from cascading into organisational crises. Interruptor serves as an autonomous safety brake system that provides immediate risk containment and emergency response coordination when tasks exceed acceptable risk parameters or violate governance constraints, ensuring organisational resilience through proactive intervention and systematic failure prevention.
- ‘PeaceKeeper’: The PeaceKeeper AI agent solves the critical business problem of organisational continuity during infrastructure failures and system emergencies by automatically taking over when tasks are paused or systems fail, implementing intelligent workload rerouting to standby actors, alternate regions, or human fall-back queues following business continuity-disaster recovery (BC/DR) runbooks to maintain business operations and minimise organisational disruption. PeaceKeeper serves as the autonomous BC/DR orchestrator that ensures organisational resilience through systematic recovery procedures, alternative resource allocation, and comprehensive disaster recovery coordination while providing real-time recovery metrics to leadership dashboards for informed crisis management and strategic decision-making.

Finally, Figure 3 shows that there is another AI agent (‘ledger’) deployed on all three lanes:

- ‘Ledger’: The ledger AI Agent solves the critical business problem of ensuring financial and operational accountability across all automated and human-executed tasks within the customer’s organisation. Ledger provides a single source of truth for

financial-grade task logging, capturing every event, decision, and outcome with immutable audit trails, thereby reducing compliance costs, accelerating audit readiness, and eliminating manual reconciliation errors. By providing consistent, tamper-proof traceability of actions across the enterprise, it delivers measurable outcomes, such as reduction of audit preparation time, error-free compliance reporting, and demonstrable governance alignment with internal and external standards. Its unique capability in the agent ecosystem is acting as the immutable compliance backbone, ensuring that all agent and human actions are consistently captured, reconciled, and available for regulatory, financial, and governance needs, thereby giving enterprises the confidence to scale automation under full transparency and accountability.

3.3. NOW OS core component: unified operating model and management system

The UOMAMS component is a comprehensive, multi-dimensional framework that aligns an organisation's strategic vision with business operations across eight critical domains: business, agentic AI, service, process, security, people and organisation, information, and technology. By structuring each domain in five layers (definition, design pattern, reference architecture, capability performance, and in-life purpose) and presenting tailored 'views' for executive management, C-level/directors, designers, managers, and operations teams, UOMAMS transforms abstract strategy into executable and measurable operations. UOMAMS is illustrated in Figure 4.

In Figure 4, the eight boxes under the 'Domains Overview' title box represent eight architecture domains, which intersect with the five architecture layers (represented by the five grey boxes under the 'Design Pattern' title box), illustrating the major functionalities of UOMAMS. The five boxes below the 'Definition' title box define each of the five layers UOMAMS operates on. The five boxes under the 'Views' title box represent each of the five role-specific lenses through which the different decision makers consume the output of UOMAMS.

We now illustrate in detail how each of the eight architecture domains are defined and governed at each of the five layers to ensure consistency from strategy down to operations.

Business domain

- Definition layer: Captures vision, mission, and strategic goals (e.g. 'Increases market share by 15% in government AI services').

Unified operating model architecture and management system									
	Why & when	How, what & when	How, what & when	How, what & when	How, what & when	Who, what & when	How, what & when	How, what & when	
	Domain overview								
Design pattern	Business	Agentic AI management	Service management	Process	Security	People & organisation	Information	Technology	Views
Vision and business design pattern	Business vision, strategy and goals	Agentic AI strategy and goals	Service strategy and goals	Process strategy and goals	Security strategy and goals	Organisation and people strategy	Information strategy and goals	Technology strategy and goals	Exec management
Strategic design pattern	Business operating model canvas	Agentic capability model	Service capability model	Process capability model and interfaces	Organisation security capability model	Organisation structure and locations	Information design data & application	Technology design (building blocks)	C Layer/director
Assets/reference architecture pattern	Business capability - units and function	AI agentic CMDB/CMMS catalogue	Service catalogue	Process catalogue	Security controls catalogue	People capability & actor & role catalogue (people capability, work role, task)	Information & knowledge capability information schema catalogue	Technology capability - H & S/W component catalogue	Designer
Capability performance	Business performance	AI agentic service/role performance	Service performance	Process performance	Security performance	Organisation/workforce learning & performance	Information performance	Technology performance	Manager
In life purpose	Business operations	AI agentic service/role operations	Service operations	Process management	Security management	Operational roles	Information management	Technology management	Operations view

Figure 4. Unified operating model and management system (UOMAMS).

- Design pattern: Uses a business operating model canvas to convert strategy into value streams.
- Reference architecture: Maintains a catalogue of business assets and reference architectures (e.g. ‘strategic initiatives portfolio,’ and ‘value-stream map library’).
- Capability performance: Tracks business performance metrics (e.g. revenue growth, margin, and customer satisfaction) against targets.
- In-life purpose: Governs ongoing business operations (quarterly objectives and key results [OKR] reviews, and strategic program management offices).

Agentic AI management domain

- Definition layer: Articulates AI strategy and goals (e.g. ‘implements AAA triage for 80% of service tickets by Q3’).
- Design pattern: Adopts an AI capability model defining how AI agents (e.g. ‘lens’, ‘weightsmith’, and ‘cartographer’) integrate with workflows.
- Reference architecture: Maintains an ‘AI agent CMDB catalogue’ listing each agent’s inputs, outputs, training data sources, and lifecycle stage.
- Capability performance: Measures AI-agent performance (accuracy, confidence scores, uptime, and error rate).

- In-life purpose: Manages live AI agent operations (data-pipeline health checks, continuous retraining, and model drift monitoring).

Service management domain

- Definition layer: Establishes service strategy and goals (e.g. 'offers 24x7 incident response with 90% Service Level Agreement (SLA) compliance').
- Design pattern: Uses a service capability model (service blueprints, SLAs, and service-value chains).
- Reference architecture: Houses a service catalogue (service definitions, dependencies, and routings).
- Capability performance: Tracks service performance key performance indicators (KPIs) (Mean Time to Repair [MTTR], customer satisfaction, and incident volumes).
- In-life purpose: Executes service operations (incident management, problem management, and continuous service improvement).

Process management domain

- Definition layer: Lays out process strategy and goals (e.g. 'reduces end-to-end procurement cycle by 25%').
- Design pattern: Adopts a process capability model (process flow templates, and decision-gate frameworks).
- Reference architecture: Contains a process catalogue (standardised process maps, Responsible, Accountable, Consulted, and Informed (RACI) charts, and workflow engines).
- Capability performance: Monitors process performance (cycle time, defect rates, and resource utilisation).
- In-life purpose: Oversees process management (continuous process orchestration, and business process management (BPM) tool governance).

Security management domain

- Definition layer: Captures security strategy and goals (e.g. 'achieves zero-day breach detection within 10 minutes').
- Design pattern: Implements a security control capability model (NIST, and ISO 27001 patterns).
- Reference architecture: Maintains a security controls catalogue (firewalls, identity and access management [IAM] frameworks, and encryption standards).
- Capability performance: Gauges security performance (vulnerability metrics, compliance scorecards, and mean time to detect/resolve).

- In-life purpose: Runs security management operations (security information and event management (SIEM) monitoring, SOC-as-a-service, and incident response drills).

People and organisation domain

- Definition layer: Specifies organisation and people strategy (e.g. 'shifts 20% of legacy roles to AI-augmented roles by year's end').
- Design pattern: Uses an org-design capability model (organisational charts, skills taxonomies, and responsible, accountable, consulted, informed (RACI) matrices).
- Reference architecture: Contains an org structure and roles catalogue (job families, role descriptions, and skill profiles).
- Capability performance: Tracks people performance (productivity metrics, skill-gap indices, and employee engagement).
- In-life purpose: Manages organisational and human capital operations (recruiting, learning and development programs, and performance reviews).

Information management domain

- Definition layer: Defines information and data strategy (e.g. 'achieves a single source of truth for citizen records by Q4').
- Design pattern: Uses an information and knowledge capability model (data architecture blueprints, and taxonomy frameworks).
- Reference architecture: Houses an information and data catalogue (data dictionaries, metadata repositories, data lineage).
- Capability performance: Measures information performance (data quality scores, usage analytics, and data latency).
- In-life purpose: Operates information lifecycle management (extract, transform, load [ETL] jobs, data governance councils, and Mobile Device Management [MDM] platforms).

Technology management domain

- Definition layer: States technology strategy and goals (e.g. 'migrates 70% of workloads to cloud by Q2 next year').
- Design pattern: Deploys a technology capability model (architectural principles, platform standards, and reference stacks).
- Reference architecture: Maintains a technology components catalogue (system inventories, integration patterns, and container registries).
- Capability performance: Monitors technology performance (availability metrics, latency, and capacity utilisation).

- In-life purpose: Conducts technology and infrastructure operations (DevOps pipelines, change management, and release orchestration).

The role-based visibility and accountability UOMAMS is able to provide five views, which can be summarised as follows:

- Executive management view: High-level dashboards showing aggregated KPI heatmaps per domain, major strategic initiatives, and risk exposure across security, AI, or regulatory domains.
- C-level/director view: Detailed roadmaps and 'how-to-guide' patterns (e.g. service strategy maps, AI agent capability timelines) that show how each domain delivers on Executive objectives.
- Designer view: Domain-specific patterns and reference (P&Rs) architectures, such as the 'security controls pattern', 'process capability model blueprint', or 'Org & roles taxonomy'.
- Manager view: Real-time operational metrics and task assignments: 'who owns which tasks in this service or process?' and 'which AI agents are live vs. deprecated?'
- Operations view: Day-to-day dashboards (e.g. incident queues, AI model performance logs, and daily stand-up checklists) enabling frontline teams to execute precisely against defined patterns and SLAs.

By bridging the gap between 'why/when' (strategic intent and timing) and 'how/who' (detailed execution), UOMAMS ensures every part of the organisation, whether a frontline operator, service-designer, or AI agent, shares a common blueprint. There are no siloed process diagrams or orphaned technology roadmaps; instead, every element is traced back to business objectives, risk parameters, and performance KPIs.

3.4. The Synergy between Human–AI Division of labour blueprint and UOMAMS

We now provide a detailed exploration of the synergy between human–AI division of labour blueprint and UOMAMS, illustrating how these two core components of NOW OS complement each other to deliver a unified, continuously evolving, audit-ready operating system for the organisation.

At strategic level, to achieve the organisation's strategic alignment and traceability, the unified model's role is to provide top-sown strategy. At business definition layer, the unified model captures

high-level vision, mission, strategic goals, and key performance indicators (KPIs); at service/process/AI definition layers, it translates those business objectives into domain-specific strategies (e.g. 'Shift 30% of customer-support tasks to AI-assist agents by Q4'); at People & Organisation definition layer, it defines desired organisational states, skill requirements, and cultural goals (e.g. 'build AI-savvy support teams that can supervise intelligent bots').

In the meantime, the blueprint's role is to generate bottom-up real-world and real-time validation. Through automated data ingestion (email/calendar mining, workflow logs, etc.) and consultant-led workshops, the blueprint uncovers exactly which tasks exist, mapping them back to the unified model's service and process definitions. In addition, each atomic task is scored against the unified model's strategic guardrails (risk, compliance, thresholds, confidence intervals, human-in-the-loop escalation rules, etc.) to produce an AAA recommendation that feeds back into unified model's AI definition and people & organization layers, ensuring that operational reality remains in lockstep with strategic intent.

The synergy between the two components produces a mutual feedback loop. Whenever unified model's stakeholders (e.g. C-level, directors) revise strategic targets, such as new regulatory requirements or shifts in service-delivery goals, the blueprint's scoring engine recalibrates, re-scoring all affected atomic tasks, updating the AAA assignments, and highlighting revised skill-gap reports for People & Organisation. Conversely, real-time signals from the blueprint (e.g. '30% of grant-approval tasks now qualify for full automation') trigger the unified model's capability performance dashboards to surface potential organisational adjustments (e.g. 'phase out three manual review roles, and retrain staff on automated-oversight models').

The interactions between the two components ensure continuous bi-directional optimisation:

- In blueprint to unified model direction, operational signals travel upwards, which achieves:
 - Real-time analytics: Blueprint continuously streams granular metrics (task durations, AI confidence levels, error rates, and compliance flags) into unified model's capability performance layers across all domains.
 - Automated alerting: The moment a metric drifts beyond a predefined threshold (e.g. 'AI model error rate > 5%' or

'Procurement cycle time > 48 hours'), the unified model dashboards raise tickets, assign action items, and reallocate resources in service and process domains.

- Governance escalations: If an atomic task marked 'high sensitivity (GDPR)' is executed by an AI agent with insufficient data-privacy controls, unified model's security domain automatically locks down that agent and notifies the chief information security officer (CISO) for remediation.
- In unified model to blueprint direction, strategic directions travel downward, which achieves:
 - Revised objectives: When executive management revises quarterly goals (e.g. 'achieves 80% cost reduction in back-office operations'), the unified model's business and service definition layers broadcast new targets down to the blueprint's scoring engine, which automatically re-scores all back-office atomic tasks against updated cost-benefit logic and reprioritises automation candidates.
 - Policy and control updates: A new regulatory mandate (e.g. 'all citizen data must remain within EU borders') input into unified model's security definition layer immediately tells the blueprint to tag every related atomic task with a 'Data Residency = EU' flag, forcing those tasks to be routed only to EU-hosted systems or human agents.
 - Capability roadmap changes: If unified model's technology definition layer declares 'Decommission Legacy System Z by year-end', the blueprint identifies all atomic tasks dependent on 'System Z' and kicks off mitigation plans (e.g. 'assign to interim manual process' or 'build new AI-assistant integration by Q3').

In summary, the human-AI division of labour blueprint and UOMAMS together provide the continuum for NOW OS, where:

1. Strategy meets reality: the unified model's top-down targets flow into the blueprint's task-level scoring engine, and the blueprint's bottom-up performance metrics feed back into the unified model's dashboards – creating a continuous, bi-directional feedback loop.
2. Governance is built on actual work: unified model's policies and controls are enforced at the atomic-task level by the blueprint, ensuring that every action, from a simple support ticket to a high-stakes security decision, is governed, auditable, and optimised in real time.
3. Agility and resilience scale together: with the NOW OS in place, executives can adjust strategic priorities on the fly, and the blueprint instantly recalibrates task assignments; similarly, operational signals (e.g. SLA breaches, AI drift alerts, and compliance

flags) are surfaced immediately in the unified model, enabling rapid course correction.

By weaving strategic intent, domain governance, and atomic-level task orchestration into one coherent operating system, the NOW OS accelerates decision cycles, elevates outcome quality, and ensures that every piece of work – no matter how small – advances the organisation’s overarching mission.

3.5. The evaluation approach to NOW OS performance

To address research question 3, we now describe our approach to system performance evaluation. A strong performance-evaluation framework for NOW OS should measure not only technical efficiency but also business value, user experience, governance, and long-term scalability. In order to develop a proper model for NOW OS performance evaluation, we decide to adopt a hybrid approach. For operational performance evaluation, we adopt ITIL Continual Service Improvement model [38], which focuses on availability and reliability, incident and problem resolution, change management success rates, service-level compliance, and user support efficiency. For governance and strategic alignment evaluation, we adopt COBIT framework [39], which uses process capability levels and metrics across alignment with enterprise objectives, information security performance, compliance and governance, IT process maturity, and risk management effectiveness. For system’s overall business performance, we adopt value realisation framework [40], which measures value along three paths: cost reduction (process automation, and resource optimisation), business growth (sales lift, and customer retention), and risk reduction (compliance improvements, and error reduction).

Integrating the three models together, the layered evaluation approach allows us to effectively assess the NOW OS system performance from the dimension of technical performance, user adoption/satisfaction, business impact, governance and risk assurance, and long-term strategic value. Specifically, NOW OS can be evaluated by (i) task-level latency and completion fidelity, (ii) rework and exception rates, (iii) human workload redistribution across AAA modes, and (iv) governance outcomes, including auditability, policy adherence, and escalation accuracy.

Next, we describe a sample use case of implementation of NOW OS in a real-world scenario, accompanied by performance evaluation metrics based on our layered evaluation approach.

4. A Sample Case of NOW OS Implementation

In this section, we describe a sample use case of NOW OS implementation for improving patient experiences and business outcomes for general practice (GP) providers.

4.1. Introduction

A mid-sized GP group operating across 42 suburban clinics found itself at a critical inflection point. Though they had implemented a widely used electronic health records (EHR), integrated a scheduling app, and recently trialed a medical scribe AI plugin, their leadership team was still grappling with a deep structural problem: both care quality and operational performance were trending downward. Their metrics showed increased patient churn, low first-call resolution rates, rising no-show volume, and a troubling decline in post-visit satisfaction scores. They come to us for consultation.

The NOW OS consulting team analyse the situation, and notice the following: (1) each digital solution added its own siloed interface, generated additional tasks, and introduced new workflow dependencies without providing system-level logic for accountability, prioritisation, or resolution; (2) nurses and medical assistants were toggling between inboxes, sticky notes, whiteboards, and EHR alerts, attempting to stitch together processes that fragmented across people, roles, and tools; (3) administrative teams were fire-fighting; clinical teams were fatigued; and (4) patients sensed the chaos, and many simply stopped engaging.

The clinical operations team at this GP identified three patient-centred metrics as their most critical indicators of success and sustainability: communication quality, access to care, and the overall patient satisfaction.

In the next section, we focus on the first KPI, and describe how NOW OS can be implemented to help improve communication quality, together with the business performance gains that followed.

4.2. How NOW OS improves messaging integrity, clinical coordination, and patient trust

NOW OS addresses the communication quality challenge faced by GP not by layering in more messages but by transforming task handoffs into system-orchestrated commitments.

Consider the common task of managing a laboratory order. Traditionally, the process begins when the provider places an order in the EHR, but from that point forward, the actions required to validate, submit, and follow through on that laboratory span more than two dozen steps across five different platforms. The medical assistant (MA) must identify the correct current procedural terminology (CPT) code for the ordered test, verify that the patient's insurance covers it, and determine whether a prior authorisation (PA) is required. This typically involves logging into payer portals like Availity or UnitedHealthcare, navigating benefit search trees, and interpreting ambiguous coverage language. If a coverage condition requires diagnosis code justification, the MA must then reconcile that with the provider's problem list in the EHR – or, more commonly, call the payer to clarify policy interpretations. This means spending time on hold, authenticating the patient, verbally verifying codes, and then jotting notes on paper before transferring the outcome back into the EHR comment field or messaging the provider. This isn't a bug; it's the legacy design of modern healthcare technology stacks. And it is one of the core reasons that patients experience unresponsiveness, delays, or mixed signals about their care.

NOW OS changes this by treating each laboratory order as an orchestrated task string, not a one-time entry in a chart. The system decomposes the workflow into atomic-level steps, each tagged with the most appropriate execution mode – Human-only, Assistive AI, Augmented performance, or full Automation. For instance, the process of checking a laboratory code against payer policies is handled automatically via real-time API integrations. If multiple coverage tiers or diagnosis rules apply, a context-aware decision engine prompts the MA with the most relevant options based on the patient's chart – augmenting judgement not replacing it. Prior authorisation requirements are detected in the background, and if the criteria are met, the authorisation request is automatically compiled and submitted. Status updates are tracked continuously and synchronised back to the patient communication module and the provider's task view. Every step is audited. Every handoff is visible. No hallway conversation or sticky note is needed to 'remember' what comes next.

What once required over 30 scattered human-led steps now resolves in 13 coordinated actions: most of them invisible to the patient, many of them delegated to the right tier of automation, and all of them traceable.

This laboratory order example is emblematic of the silent, repetitive frictions that general practice groups encounter every day – frictions that aren't solved by better messaging templates or stricter checklists but by shifting from messaging to managed completion. If a single order coordination task string can be deconstructed, mapped, and optimised in this way, imagine the downstream potential of doing the same for every recurring action across the practice – Referrals. Imaging requests. Medication changes. Claims reconciliation. Patient check-ins. Every one of these contains dozens of hidden steps that today depend on human memory and multi-system navigation.

NOW OS makes those tasks visible, manageable, and governable. In doing so, it redefines what communication means in clinical care – not just the exchange of information but the assurance of follow-through.

The result is more than efficiency. It is the restoration of patient confidence, the empowerment of frontline teams, and the emergence of trust as a design property of the system itself.

Similar analyses and implementation of NOW OS can be conducted to address other KPIs, and applicable to other scenarios.

4.3. NOW OS performance evaluation metrics for general practice

In this use case of NOW OS implementation, system performance was evaluated across efficiency, quality, governance, and human workload impacts. Metrics were collected from system logs, internal audits, operational reports, and patient experience indicators during comparable observation windows before and after deployment. The system performance metrics and variables are summarised in Table 2.

Pre- and post-implementation metrics were compared using relative percentage change to assess the directional and magnitude impact of NOW OS on system performance. Special attention was given to high-friction task strings, including laboratory orders, prior authorisations, and post-discharge follow-up, where coordination failures had previously driven patient dissatisfaction and staff overload.

Task latency, rework, and escalation rates were evaluated as indicators of orchestration effectiveness, while AAA distribution shifts

Table 2. System performance metrics and variables.

Metric category	Metric	Operational definition	Unit of measurement	Primary data source
Task efficiency	Task execution latency	Mean elapsed time between task trigger and verified completion	Minutes per task	NOW OS task logs
	Task steps per workflow	Number of atomic actions required to complete a defined task string	Count	Workflow deconstruction analysis
	Multi-touch rate	Percentage of tasks requiring more than two human interactions	% Of tasks	Internal audit
Quality and rework	Administrative rework rate	Percentage of tasks requiring correction, resubmission, or repetition	% Of tasks	Internal audit
	Exception escalation rate	Frequency of tasks triggering human escalation due to ambiguity or risk	Escalations per 100 tasks	NOW OS escalation logs
Governance and compliance	Audit completeness	Proportion of tasks with full provenance (actor, evidence, and timestamps)	% Of tasks	Ledger records
	Policy violation incidents	Number of detected deviations from payer, clinic, or compliance rules	Count per period	Compliance review
Human workload impact	Human touch time	Average active human time spent per task string	Minutes	Time sampling
	AAA distribution	Proportion of tasks classified as assist, augment, or automate	% By category	NOW OS allocation logs
Patient experience (outcome)	Communication-related callbacks	Patient-initiated calls due to missing or delayed follow-up	Calls per 1000 visits	Call centre reports
	Patient satisfaction index	Composite post-visit satisfaction score	Index score	Patient surveys

were analysed to understand how work was redistributed between humans and AI. Patient-facing outcomes were interpreted as downstream effects of improved task completion reliability, rather than as isolated satisfaction interventions.

Through this systematic performance evaluation approach, we can quantify the positive impact of NOW OS on the GP group operation efficiency in terms of task execution latency, administrative rework, and payer-related tasks requiring multiple touches etc. These efficiency gains coincided with measurable improvements in patient satisfaction, reduced inbound clarification calls, and improved staff capacity allocation, supporting the hypothesis that task-level orchestration materially improves both operational and experiential outcomes in primary care settings.

5. Conclusions

One of today's most significant technology advancements is AI. Vention estimates that there are over 56,000 active AI companies worldwide [41], with about 15,000 specialised AI tools designed for specific tasks. Rankmy AI tracks more than 45,000 AI tools worldwide [42]. However, several recent surveys and reports (e.g. [43, 44]) suggest that the proportion of AI tools that are genuinely ready for end-to-end deployment at enterprise scale is low. Those who industrialise AI labour and optimise human-AI collaborations are now gaining exponential speed and cost advantages. Those who lag will face loss of competitiveness, talent loss, and inability to deliver cost-effective customer experiences. In the meantime, people have started recognising the obstacles that must be overcome to have a thorough understanding of human-AI collaboration.

In this paper, we present NOW OS, an integrated end-to-end operating system that analyses human-AI collaboration from the atomic task (the smallest unit of real work tasks) level, and provides comprehensive functionalities of structuring, governing, and optimising the task redesign, task assignment, and task process re-engineering to achieve the optimal human-AI collaboration for the organisation. We propose a conceptual model for human-AI collaboration analysis that addresses the identified obstacles to understand human-AI collaborations, and serves as the theoretic foundation for our proposed NOW OS architecture. We also detail the technical design of the proposed NOW OS system functionalities, including how it addresses the security and governance aspects of human-AI collaboration throughout the system components and the system integration, and how the system performance can be evaluated in a systematic approach. For industry practitioners, NOW OS is a model for building a 'digital twin of works' for organisations, and subsequently orchestrating human-AI collaborations to optimise value generation for organisations. For academic research community, the proposed NOW OS architecture directly addresses the major obstacles to human-AI collaborations in organisations, and provides a prototype and testbed for further investigations on the optimisation of human-AI collaboration.

There are several promising research directions that can enhance the proposed NOW OS that we are currently actively pursuing. We are currently developing an ontology for NOW OS that will be implemented as knowledge graphs. The ontology serves as the foundation of 'digital twin of works' for organisations, and allows the organisations to visualise the relationships among atomic tasks, skill requirement, and organisation job positions. We are also

developing an AI capability evaluation model (together with survey instruments) to systematically assess the organisation's current and target AI capacity. The capacity model incorporates the dimensions of people, process, AI technology, and time, and allows organisations to develop their human resource and AI capabilities accordingly. In addition, we are now actively engaging in the technical design of the AI agents and algorithms to be deployed and implemented in NOW OS, especially those directly related to the atomic task breakdown and optimal task assignments, two core IPs for NOW OS.

NOW OS is developed for individual organisations, and its structure and functionalities are perfectly aligned with the micro- and meso-level analyses presented in our conceptual model. The macro-level analyses go beyond the current scope of NOW OS design. Another interesting direction for future research would be to study how the NOW OS architecture can evolve over time to capture the rapidly advancing AI technology landscape.

Acknowledgements

The authors are grateful to Dain Belyeu, who provided valuable insights that are instrumental to the completion of the paper.

References

- [1] L. Rouhiainen, *Artificial intelligence: 101 things you must know today about our future*. Alicante: Lasse Rouhiainen, 2018.
- [2] R. Li, *Artificial intelligence revolution: How AI will change our society, economy, and culture*. New York, NY: Simon and Schuster, 2020.
- [3] J. Bughin, E. Hazan, S. Ramaswamy, M. Chui, T. Allas, P. Dahlström, N. Henke, M. Trench, *Artificial intelligence: The next digital frontier*, vol. 47, no. 3.6. San Francisco, CA: McKinsey Global Institute, Discussion Paper, pp. 6–8, 2017.
- [4] X. Sun, Y. Song, "Unlocking the synergy: Increasing productivity through human-AI collaboration in the Industry 5.0 era," *Computers & Industrial Engineering*, vol. 200, Art. no. 110657, 2025, doi: [10.1016/j.cie.2024.110657](https://doi.org/10.1016/j.cie.2024.110657).
- [5] D. Loske, M. Klumpp, "Intelligent and efficient? An empirical analysis of human-AI collaboration for truck drivers in retail logistics," *The International Journal of Logistics Management*, vol. 32, no. 4, pp. 1356–1383, 2021, doi: [10.1108/IJLM-03-2020-0149](https://doi.org/10.1108/IJLM-03-2020-0149).
- [6] S. Makridakis, "The forthcoming artificial intelligence (AI) revolution: Its impact on society and firms," *Futures*, vol. 90, pp. 46–60, 2017, doi: [10.1016/j.futures.2017.03.006](https://doi.org/10.1016/j.futures.2017.03.006).

- [7] M. Lane, A. Saint-Martin, "The impact of artificial intelligence on the labour market: What do we know so far?," *OECD Social, Employment and Migration Working Papers*, no. 256, pp. 1–60, 2021.
- [8] M. Webb (2019). "The impact of artificial intelligence on the labor market," SSRN. , doi: [10.2139/ssrn.3482150](https://doi.org/10.2139/ssrn.3482150).
- [9] S.E. Faluyi, "AI and job market: Analysing the potential impact of AI on employment, skills, and job displacement," *African Journal of Marketing Management*, vol. 17, no. 1, Art. no. 37DCEDE72936, 2025, doi: [10.5897/AJMM2024.0747](https://doi.org/10.5897/AJMM2024.0747).
- [10] S. Chhibber, S.R. Rajkumar, S. Dassanayake, "Will artificial intelligence reshape the global workforce by 2030? A cross-sectoral analysis of job displacement and transformation," *Blockchain, Artificial Intelligence, and Future Research*, vol. 1, no. 1, pp. 35–51, 2025, doi: [10.70211/bafr.v1i1.178](https://doi.org/10.70211/bafr.v1i1.178).
- [11] C. Challoumis, AI and the economy: How technology is redefining employment and income distribution. ResearchGate. [Online]. Available: https://www.researchgate.net/publication/386282404_AI_And_The_Economy_-_How_Technology_Is_Redefining_Employment_And_Income_Distribution#full-text. [Accessed: Sep. 15, 2025].
- [12] N. Shmatko, V. Ivchuk, "Unleashing the Capabilities of Artificial Intelligence in Managing Businesses," in *Ai-driven transformation: Mapping the course for future business landscapes: monograph*, O. Prokopenko, M. Jarvis, Ed. Tallinn: Teadmus OÜ, 2024, pp. 17–32.
- [13] D. Anny, (2023), "Leveraging artificial intelligence to optimize business processes in enterprise architecture," ResearchGate. [Online]. Available: https://www.researchgate.net/publication/389918517_Leveraging_Artificial_Intelligence_to_Optimize_Business_Processes_in_Enterprise_Architecture. [Accessed: Sep. 15, 2025].
- [14] O.A. Farayola, A.A. Abdul, B.O. Irabor, E.C. Okeleke, "Innovative business models driven by AI technologies: A review," *Computer Science & IT Research Journal*, vol. 4, no. 2, pp. 85–110, 2023, doi: [10.51594/csitrj.v4i2.608](https://doi.org/10.51594/csitrj.v4i2.608).
- [15] G. Fragiadakis, C. Diou, G. Kousiouris, M. Nikolaidou, "Evaluating human-AI collaboration: A review and methodological framework," *arXiv preprint*, arXiv: 2407.19098, 2024, doi: [10.48550/arXiv.2407.19098](https://doi.org/10.48550/arXiv.2407.19098).
- [16] J.A. Jiang, K. Wade, C. Fiesler, J.R. Brubaker, "Supporting serendipity: Opportunities and challenges for human-AI collaboration in qualitative analysis," *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, no. CSCW1, Art. no. 94, 2021, doi: [10.1145/3449168](https://doi.org/10.1145/3449168).
- [17] Q. Yang, A. Steinfeld, C. Rosé, J. Zimmerman, "Re-examining whether, why, and how human-AI interaction is uniquely difficult to design," *Proceedings of 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–13, doi: [10.1145/3313831.3376301](https://doi.org/10.1145/3313831.3376301).
- [18] M.R. Frank, D. Autor, J.E. Bessen, E. Brynjolfsson., M. Cebrian, D.J. Deming, M. Feldman, M. Groh, J. Lobo, E. Moro, D. Wang, H. Youn, I. Rahwan, "Toward understanding the impact of artificial intelligence on labor," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 116, no. 14, pp. 6531–6539, 2019, doi: [10.1073/pnas.1900949116](https://doi.org/10.1073/pnas.1900949116).

- [19] T. Szadeczky, Z. Bederna, "Risk, regulation, and governance: Evaluating artificial intelligence across diverse application scenarios," *Security Journal*, vol. 38, no. 1, Art. no. 35, 2025, doi: [10.1057/s41284-025-00495-z](https://doi.org/10.1057/s41284-025-00495-z).
- [20] G. Fragiadakis, C. Diou, G. Kousiouris, M. Nikolaidou, "Evaluating human-AI collaboration: A review and methodological framework," *arXiv preprint*, arXiv: 2407.19098, 2024, doi: [10.48550/arXiv.2407.19098](https://doi.org/10.48550/arXiv.2407.19098).
- [21] M. Puerta-Beldarrain, O. Gomez-Carmona, R. Sanchez-Corcuera, D. Casado-Mansilla, D. Lopez-de-Ipina, L. Chen, "A multifaceted vision of the human-AI collaboration: a comprehensive review," *IEEE Access*, vol. 13, pp. 29375–29405, 2025, doi: [10.1109/ACCESS.2025.3536095](https://doi.org/10.1109/ACCESS.2025.3536095).
- [22] K. Meimandi, M. Bolton, P. Beling, "Human-agent teaming: A system-theoretic overview," *TechRxiv*, 2024, doi: [10.36227/techrxiv.170629550.02125312/v1](https://doi.org/10.36227/techrxiv.170629550.02125312/v1).
- [23] B. Schelble, C. Flathmann, N. McNeese, G. Freeman, R. Mallick, "Let's think together! Assessing shared mental models, performance, and trust in human-agent teams," *Proceedings of the ACM on Human-Computer Interaction*, vol. 6, no. GROUP, Art. no. 13, 2022, doi: [10.1145/3492832](https://doi.org/10.1145/3492832).
- [24] W. Xu, Z. Gao, "Applying HCAI in developing effective human-AI teaming: A perspective from human-AI joint cognitive systems," *Interactions*, vol. 31, no. 1, pp. 32–37, 2024, doi: [10.1145/3635116](https://doi.org/10.1145/3635116).
- [25] B. Yang, S. Xie, L. Li, "An integrated task decomposition framework considering knowledge reuse and resource availability for complex task crowdsourcing," *Systems*, vol. 13, no. 9, Art. no. 728, 2025, doi: [10.3390/systems13090728](https://doi.org/10.3390/systems13090728).
- [26] A. Caetano, A. Silva, J. Tribolet, "Business process decomposition – An approach based on the principle of separation of concerns," *Enterprise Modelling and Information Systems Architectures*, vol. 5, no. 1, pp. 44–57, 2010, doi: [10.18417/emisa.5.1.3](https://doi.org/10.18417/emisa.5.1.3).
- [27] K. Suri, A. Mos, "Human task monitoring and contextual analysis for domain specific business processes," *arXiv preprint*, arXiv: 1610.05788, 2016, doi: [10.48550/arXiv.1610.05788](https://doi.org/10.48550/arXiv.1610.05788).
- [28] W. Van Der Aalst, "Process mining: Overview and opportunities," *ACM Transactions on Management Information Systems*, vol. 3, no. 2, Art. no. 7, 2012, doi: [10.1145/2229156.2229157](https://doi.org/10.1145/2229156.2229157).
- [29] R. Parasuraman, T. Sheridan, C. Wickens, "A model for types and levels of human interaction with automation," *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans*, vol. 30, no. 3, pp. 286–297, 2000, doi: [10.1109/3468.844354](https://doi.org/10.1109/3468.844354).
- [30] E. Horvitz, "Uncertainty, action, and interaction: In pursuit of mixed-initiative computing," *IEEE Intelligent Systems*, vol. 14, no. 5, pp. 17–20, 1999.
- [31] Y. Zhang, H. Jacobsen, "Decentralized and policy-aware serverless orchestration for the federated web," *Companion proceedings of the ACM on web conference 2025*, 2025, pp. 1539–1543, doi: [10.1145/3701716.3715573](https://doi.org/10.1145/3701716.3715573).
- [32] I. Huvila, "Work and work roles: A context of tasks," *Journal of Documentation*, vol. 64, no. 6, pp. 797–815, 2008, doi: [10.1108/00220410810912406](https://doi.org/10.1108/00220410810912406).

- [33] T. Skouti, R. Seiger, F. Furrer, S. Strahringer, "RBPMN: The value of roles for business process modeling," *Software and Systems Modeling*, vol. 23, no. 6, pp. 1375–1406, 2024, doi: [10.1007/s10270-024-01202-z](https://doi.org/10.1007/s10270-024-01202-z).
- [34] F. Smith, M. Proietti, "Ontology-based representation and reasoning on process models: A logic programming approach," *arXiv preprint*, arXiv: 1410.1776, 2014, doi: [10.48550/arXiv.1410.1776](https://doi.org/10.48550/arXiv.1410.1776).
- [35] R. Singer, "An ontological analysis of business process modeling and execution," *arXiv preprint*, arXiv: 1905.00499, 2019, doi: [10.48550/arXiv.1905.00499](https://doi.org/10.48550/arXiv.1905.00499).
- [36] R. Parasuraman, "Designing automation for human use: Empirical studies and quantitative models," *Ergonomics*, vol. 43, no. 7, pp. 931–951, 2000, doi: [10.1080/001401300409125](https://doi.org/10.1080/001401300409125).
- [37] Y. Du, B. Prébot, T. Malloy, F. Fang, C. Gonzalez, "Experimental evaluation of cognitive agents for collaboration in human-autonomy cyber defense teams," *Computers in Human Behavior: Artificial Humans*, vol. 4, Art. no. 100148, 2025, doi: [10.1016/j.chbah.2025.100148](https://doi.org/10.1016/j.chbah.2025.100148).
- [38] S. Heikkinen, M. Jääntti, M. Tukiainen, "Continual service improvement: A systematic literature review," in *International conference on the quality of information and communications technology*. Cham: Springer, 2023, pp. 30–44, doi: [10.1007/978-3-031-43703-8_3](https://doi.org/10.1007/978-3-031-43703-8_3).
- [39] R. Khther, M. Othman, "COBIT framework as a guideline of effective IT governance in higher education: A review," *International Journal of Information Technology Convergence and Services*, vol. 3, no. 1, p. 21, 2013, doi: [10.5121/ijitcs.2013.3102](https://doi.org/10.5121/ijitcs.2013.3102).
- [40] R. Sheikh, S. Bhatia, S. Metre, A. Faqih, "Strategic value realization framework from learning analytics: a practical approach," *Journal of Applied Research in Higher Education*, vol. 14, no. 2, pp. 693–713, 2022, doi: [10.1108/JARHE-10-2020-0379](https://doi.org/10.1108/JARHE-10-2020-0379).
- [41] Vention. (2025). *State of artificial intelligence*. [Online]. Available: <https://ventionteams.com/solutions/ai/report>. [Accessed: Sep. 15, 2025].
- [42] Rankmy AI. (2025). *Explore and track more than 45,000 AI tools*. [Online]. Available: <https://www.rankmyai.com/>. [Accessed: Sep. 15, 2025].
- [43] Infosys Research. (2024). *While only 2% are ready, most companies expect productivity gains of 10–40% with Enterprise AI*. [Online]. Available: <https://www.infosys.com/newsroom/press-releases/2024/enterprise-ai-readiness.html>. [Accessed: Sep. 15, 2025].
- [44] C. Wang, A. Sullivan. (2025). *Only 7% of enterprises are AI-ready – Here's why*. [Online]. Available: <https://www.fivetran.com/blog/only-7-of-enterprises-are-ai-ready-heres-why>. [Accessed: Sep. 15, 2025].