

Designing A Human-Governed Agentic Project Management Model: A Closed-Loop Architecture for Autonomous Project Intelligence, Decision Control and Human–AI Teaming

Felixta Chioma Onyenwe, Onyedikachi Williams Onwuso, Anumo Ruth Francis, Jennifer Oghenetajiri Ekpiti, Mark Ogechi Munachimso, Mowang Martins Mogbu, Anjolaoluwa Mumeen Sulyman, Maryann Uzoramaka Onyenwe

Department of Theatre and Film Studies,

Nnamdi Azikiwe University, Awka, Nigeria.

Department of social science education, faculty of education,

Imo State University, Owerri, IMO State, Nigeria.

Department of Social work, Chengdu University of Technology, China.

Department of Mass Communications,

University of Nigeria, Nsukka.

Department of Accountancy, Federal Polytechnic Nekede Owerri.

Department of Food Science and Engineering,

South China University of Technology, China.

Editorial Department of Orijo reporter online Newspaper, Nigeria.

Department of Linguistic and Communications,

University of Port Harcourt, Nigeria.

Corresponding Author: ofelixta@gmail.com

<https://doi.org/10.56201/ijssmr.vol.12no9.2026.pg544.576>

Abstract

Purpose: This article develops a Human Governed Agentic Project Management (HG-APM) model for integrating autonomous AI capabilities with human project management authority.

Design/methodology/approach: An integrative conceptual synthesis combines four literature streams: AI-enabled project management, autonomous Manager Agents, human oversight governance, and human–AI teaming. Source supported findings are separated from conceptual inferences, and the resulting architecture is expressed as a closed-loop model.

Findings: The model comprises Sense, Understand, Diagnose, Predict, Plan, Evaluate, Authorize, Act, Monitor, Learn and Replan. It proposes a consequence sensitive, graduated autonomy rather than a single autonomy level. Routine and reversible actions may be automated within predefined boundaries, whereas consequential financial, contractual, scope, safety, stakeholder, reputational and strategic decisions require stronger authority from the designated human or organizational decisionmaker.

Originality: The model integrates project-management intelligence, human–AI teaming and governance into one operational architecture. It contributes to a project specific decision boundary and testable propositions for evaluating whether governed agentic AI improves project outcomes without weakening human control.

Keywords: *Agentic AI; Project management; Human–AI teaming; Autonomous agents; Project governance; Decision autonomy*

1. Introduction

Artificial intelligence has moved from a largely analytical support technology toward systems capable of generating content, reasoning over heterogeneous information, using tools and, increasingly, taking sequences of actions. This transition is particularly consequential for project management because projects are dynamic socio-technical systems in which decisions are distributed across schedules, budgets, resources, risks, dependencies, contracts, quality requirements and stakeholder relationships. Recent reviews show that AI in project management now spans prediction, estimation, optimization, risk management, monitoring, decision support, and increasingly generative and agentic applications. The evidence remains especially concentrated in data intensive processes, while newer work is beginning to examine human guided orchestration and role-based agents (Nenni et al., 2024; Adamantiadou and Tsironis, 2025; Salimimoghdam et al., 2025; Anghel and Cioara, 2026). The emergence of generative AI expands this role. Generative systems can summarize project information, draft documents, retrieve and synthesize knowledge, support communication, and assist managerial reasoning. More recent work proposes a further transition toward agentic systems that can pursue goals, decompose tasks, use tools, coordinate with other agents and adapt their actions. A 2026 systematic review of generative AI in IT project management describes the academic literature as still largely focused on prompt-based applications while identifying agentic architectures as a promising direction for autonomous task execution and human–AI orchestration (Generative AI for IT Project Management, 2026). Kalluri (2026) similarly develops an evidence-based implementation framework and maturity model for agentic AI adoption in project management. These developments indicate that the field is moving beyond the question of whether AI can assist project managers toward the more difficult question of how an AI agent should participate in project execution. The distinction matters because autonomy introduces a governance problem that is different from conventional software automation. An autonomous agent can generate a sequence of actions whose cumulative consequences were not explicitly reviewed by a human. Montgomery, Gallagher, and Mercer (2025) therefore argue that meaningful human oversight requires more than placing a human somewhere in a workflow: the responsible organization must provide decision rights, evidence, intervention capability, accountability and auditability. At the same time, Masters et al. (2025) show through the Manager Agent paradigm that autonomous orchestration of human–AI workflows remains technically difficult. Their MA-Gym experiments demonstrate that even GPT5 based Manager Agents struggle to jointly optimize goal achievement, constraint adherence and workflow runtime across diverse workflows. The implication is important for project management: an agent may be able to act autonomously without yet being reliably capable of managing the competing objectives and constraints that define real projects.

Human–AI teaming research adds another necessary perspective. Gonzalez et al. (2026) define human–AI complementarity as the conditions under which human–AI teams outperform humans alone or AI systems alone. Their framework emphasizes role partitioning, trust calibration, shared mental models, training, attention orchestration, knowledge infrastructure, and meta coordination. This suggests that the design problem is not simply how much authority to grant an AI. It is how human and AI capabilities should be combined so that autonomy improves project performance without displacing legitimate human judgment.

Against this background, this article proposes a Human Governed Agentic Project Management

(HG-APM) model. The model is deliberately positioned between two extremes. At one extreme, AI remains a passive tool that responds only when prompted. At the other, an AI system becomes an unrestricted autonomous manager. The proposed architecture instead treats the AI as an agentic project-management teammate operating inside explicit governance boundaries. The AI can continuously sense project conditions, understand project state, diagnose problems, predict consequences, plan alternatives, evaluate those alternatives, and where authorized, execute actions. The human project manager remains continuously informed and retains decision authority for consequential decisions. The system then monitors outcomes, learns from feedback and replants.

Research problem

Recent literature now provides substantial pieces of this problem, including AI-enabled project management capabilities, autonomous managerial orchestration, organizational implementation and maturity, human–AI teaming, and general governance mechanisms. The remaining problem is narrower: these strands do not yet provide a common project management specific decision control architecture that connects persistent project state reasoning and dependency aware evaluation to differentiated decision rights, action level authorization and post-action replanning. The contribution is therefore an integrative translation and architectural specification, not a claim that governance, dynamic autonomy or agentic execution are individually absent from prior research. The problem addressed in this paper is therefore not whether AI can perform individual project-management tasks. It is how an agentic AI project manager can operate continuously across the project lifecycle while respecting project constraints, organizational governance, human authority and the socio-technical realities of project work.

1.1 Research questions

RQ1. What capabilities are required for an AI system to move from project management assistance toward agentic project management participation?

RQ2. How should human authority and AI autonomy be structured across different project management decisions?

RQ3. What governance mechanisms are required to make autonomous project management action controllable, auditable and reversible where possible?

RQ4. How can project management intelligence, human–AI teaming and AI governance be integrated into a closed-loop agentic architecture?

1.2 Contributions

First, the article synthesizes four previously separated literature streams—AI-enabled project management, autonomous Manager Agents, human oversight governance, and human–AI complementarity.

Second, it proposes the HG-APM closed loop: Sense → Understand → Diagnose → Predict → Plan → Evaluate → Authorize → Act → Monitor → Learn → Re-plan.

Third, it specifies a decision level, consequence sensitive authorization logic that links permissible autonomy to project consequence, uncertainty, reversibility, novelty, data quality and organizational authority, while acknowledging related context sensitive governance work.

Fourth, it formalizes the human and organizational authority boundary within the project decision loop, distinguishing AI capability and recommendation from legitimate authorization to execute consequential actions.

Fifth, it identifies testable research propositions for future empirical validation.

2. Literature Review

2.1 Theoretical foundations: agency, autonomy and human control

The proposed architecture rests on a distinction between an AI-enabled tool and an autonomous agent. Classical agent research treats agency in terms of systems that perceive an environment, maintain internal representations and act in pursuit of objectives rather than merely execute a fixed response to a single input (Wooldridge and Jennings, 1995; Franklin and Graesser, 1996). Contemporary agentic AI extends this logic through language-based reasoning, tool use and multi-step action, but the underlying conceptual issue remains the allocation of initiative and control between the computational system and its human environment. The present study therefore uses ‘agentic’ in an operational rather than promotional sense: an agentic PM system must be able to maintain project state, pursue delegated objectives, select or sequence actions, interact with project tools and adapt its behaviour from feedback within defined boundaries.

Automation research provides an important counterweight to an assumption that greater autonomy is inherently better. Parasuraman, Sheridan and Wickens (2000) distinguish stages of information acquisition, analysis, decision selection and action implementation, showing that automation can be distributed across different functions rather than treated as a single continuum. This supports the HG-APM distinction between AI sensing and analysis, AI recommendation, pre-authorized execution, and human authorization. Shneiderman (2020) similarly argues for designs that combine high levels of computer automation with meaningful human control. These foundations support the paper’s central proposition that autonomy should be assigned to particular actions and decision functions rather than granted to an AI system as an undifferentiated property.

The human control problem is also inseparable from calibrated reliance. Lee and See (2004) show that trust in automation influences whether users rely on automation appropriately, while Amershi et al. (2019) identify design principles for helping people understand, interact with and recover from AI system behaviour. For agentic PM, this implies that the interface cannot simply display recommendations. It must communicate the evidence supporting an action, the system’s uncertainty, the applicable autonomy boundary, the consequences of approval, the status of execution, and the opportunity for intervention. The resulting human role is therefore neither continuous manual control nor nominal oversight but informed and differentiated control.

Project governance, complexity and decision context

Project management provides a distinctive decision environment because projects combine temporary organizational structures, interdependent activities, resource constraints, uncertainty and multiple stakeholders. Baccarini (1996) conceptualizes project complexity through differentiation and interdependency, with integration required to manage those relationships. Ward and Chapman (2003) further argue that project risk management should be understood more broadly as management of uncertainty rather than only discrete risk events. These perspectives matter for agentic PM because an apparently local action can propagate through dependencies, resources and uncertainty. A project agent that optimizes an isolated task without representing these relationships may produce locally rational but systemically harmful decisions.

Project success is likewise multidimensional. Shenhar et al. (2001) show that project success

extends beyond the traditional time-cost-performance triangle to include customer impact, business and organizational success and preparation for the future. Shenhar and Dvir (2007) further frame project management through strategic/business, operational/process and team/leadership perspectives. Consequently, an agentic PM should not evaluate an action solely through schedule improvement or task completion. Evaluation should consider competing project objectives and stakeholder consequences, which provides a theoretical basis for the Evaluate stage of HG-APM and for Proposition P7.

2.2 Human–AI teaming and collective decision-making

This literature also sharpens the meaning of complementarity. Gonzalez et al. (2026) conceptualize human–AI complementarity as conditions under which combined human–AI teams can outperform either humans or AI alone and emphasize role partitioning, meta-coordination, trust calibration, shared mental models and continuous evaluation. Recent project management-specific HAT research similarly identifies human centered AI, agile team transformation and AI-supported risk analysis as emerging themes while calling for stronger empirical testing and evaluation standards (Tian et al., 2026). The HG-APM model builds on this foundation by translating team complementarity into project decision rights and operational dependencies: the AI is responsible for computationally intensive sensing, synthesis, forecasting and coordination support, while the human PM retains authority over decisions whose consequences require organizational judgment, legitimacy or accountability.

2.3 Governance, accountability and continuous control

Accountability research establishes that decision systems require identifiable answerability relationships rather than technical transparency. Bovens (2007) conceptualizes accountability as a relationship in which an actor must explain and justify conduct to a forum capable of questioning and judging it. Kroll et al. (2017) argue that algorithmic decision systems require mechanisms capable of verifying compliance with relevant standards, while Raji et al. (2020) propose end-to-end internal algorithmic auditing to make system development and deployment traceable. These traditions support the HG-APM requirement for action logs, explicit decision ownership, traceability and post-action review. Recent research makes the delegation problem increasingly explicit. Srivastava (2026) models the choice among human control, full AI autonomy and hybrid governance under uncertainty, while Pathak and Panteli (2026) illustrate practical governance dilemmas involved in extending agent autonomy in a regulated project-team context. Armanios and Tucci (2025) argue that organizations and projects require readiness mechanisms before AI can be productively incorporated into project work, including problem scoping, delivery validation scaffolding, task separability and multi-agent governance. Pathak and Panteli (2026) illustrate the practical leadership dilemma of deciding when agents should move from recommendation to limited autonomous action in a project team setting. Srivastava (2026) similarly frames agentic AI as an organizational delegation problem involving the choice between human control and delegated decision authority. These studies reinforce the relevance of HG-APM's authorization gate while leaving open the more specific problem of how project consequences should determine the allocation of authority at each decision point.

2.4 AI-enabled project management

Recent systematic research demonstrates that AI is relevant across multiple project management functions. Nenni et al. (2024) identify strong concentrations in risk management, forecasting,

monitoring and decision support, while Adamantiadou and Tsironis (2025) synthesize applications across cost estimation, duration forecasting, risk assessment and other project processes. Salimimoghadam et al. (2025) further identify organizational, technological and human conditions shaping AI adoption. The 2026 review by Anghel and Cioara extends this literature toward generative AI and identifies process-specific agents, role-based agents and hybrid human-guided orchestration as emerging directions. Collectively, these recent studies establish a substantial evidence base for AI as an analytical and decision support capability, while also showing that the transition toward continuously acting agentic systems remains an emerging research area (Nenni et al., 2024; Adamantiadou and Tsironis, 2025; Salimimoghadam et al., 2025; Anghel and Cioara, 2026).

2.5 From generative AI assistance to agentic project management

Generative AI introduces contextual synthesis and natural-language interaction into project work. The emerging literature shows applications in documentation, communication, knowledge retrieval, planning support, and decision assistance. However, the transition to agentic systems requires the AI to operate over time rather than only respond to isolated prompts. The 2026 systematic review of GenAI for IT project management explicitly identifies agentic architecture as a forward-looking direction for autonomous task execution, collaborative decisionmaking, and human–AI orchestration. (Anghel and Cioara, 2026). Kalluri (2026) advances the field by proposing an evidence-based implementation framework and maturity model for organizations adopting agentic AI in project management. This work is useful because it shifts attention from isolated use cases toward organizational readiness, governance, phase gates, and performance measurement. However, a maturity model primarily answers how organizations can adopt agentic AI; it does not fully specify the cognitive and operational loop through which an agent performs project management reasoning and action at the level proposed here.

2.6 Human oversight and governance

Montgomery et al. (2025) develops the Human Oversight Governance Framework for Autonomous AI Agents (HOGF-AI). Their central argument is that meaningful human oversight requires actual oversight capability. A nominal human-in-the-loop arrangement is insufficient if the reviewer lacks decision rights, adequate evidence, intervention tools, organizational authority, or the practical ability to stop or modify an AI system. The framework addresses autonomy envelopes, risk gates, monitoring, validation, auditability, escalation, accountability and organizational learning. Recent project-governance work also argues that AI governance in projects must address purpose, judgement and accountability rather than readiness and performance alone (Dacre, 2026). This reinforces the need for HG-APM to treat authorization and evaluation as project-governance questions, while not implying that this normative orientation is unique to the proposed model. South et al. (2025) provide a complementary infrastructure perspective by arguing for authenticated and auditable delegation of authority to AI agents. In HG-APM, this supports the distinction between an agent’s technical capacity and the legitimacy of a particular project action but does not by itself establish the project-specific decision cycle proposed here.

2.7 Human–AI teaming and complementarity




Figure 1. Four literature streams and the unresolved project management design problem.

Literature stream	Established contribution	Unresolved issue for agentic PM	HG-APM response
AI-enabled PM	Prediction, estimation, optimization, risk analysis and decision support (Taboada et al., 2023; Nenni et al., 2024; Adamantiadou & Tsironis, 2025).	Mostly function specific; limited continuous autonomous managerial action.	Persistent project state plus closed loop reasoning and action.
Agent theory / automation	Agency, autonomy and distributed automation across information, decision and action functions (Wooldridge & Jennings, 1995; Franklin & Graesser, 1996; Parasuraman et al., 2000).	Does not specify PM-specific decision rights or project consequences.	Action level autonomy and explicit Evaluate/Authorize separation.

Manager Agents	Goal decomposition, task allocation, workflow monitoring and adaptation; multi-objective coordination remains difficult (Masters et al., 2025).	Technical orchestration does not fully specify project governance and authority.	Project-specific consequence-sensitive authorization and governance.
----------------	---	--	--

Literature stream	Established contribution	Unresolved issue for agentic PM	HG-APM response
Human oversight / AI governance	Decision rights, intervention, auditability, accountability, risk management and traceability (Bovens, 2007; Raji et al., 2020; Montgomery et al., 2025; Tabassi, 2023).	Generic governance does not map controls to PM decisions.	Translate governance into scope, cost, schedule, resource, contract, safety and stakeholder boundaries.
Human–AI teaming	Interdependence, shared goals, role clarity, trust, communication and complementarity (Mathieu et al., 2000; O'Neill et al., 2022; Gonzalez et al., 2026).	Limited operationalization of project execution dependencies and decision authority.	Formal human/AI role partitioning within the project decision loop.
Project governance / complexity	Projects involve interdependency, uncertainty, governance structures and multidimensional success (Baccarini, 1996; Ward & Chapman, 2003; Müller, 2011; Shenhar et al., 2001).	AI-agent literature rarely integrates these PM-specific conditions into autonomy design.	Consequence-sensitive autonomy based on project context and governance.

Table 1. Literature synthesis and the unresolved project-management design problem.

2.8 Evidence-to-concept traceability

The conceptual claims in HG-APM are separated into three evidentiary classes. First, literature established premises are propositions or capabilities directly described in prior research. Second, integrative inferences are relationships produced by combining established findings across project management, autonomous-agent governance, and human–AI teaming literature. Third, HG-APM

propositions are forward looking claims that require empirical testing. The model does not treat the third class as established evidence.

Table 2 summarizes this traceability. The table is deliberately conservative: a source is listed as supporting a component only where the literature establishes the underlying premise, not where it merely resembles the proposed architecture.

Table 2. Evidence-to-concept traceability across the HG-APM synthesis.

Conceptual component	What prior literature establishes	Evidence status in HG-APM	Principal sources
AI-enabled capabilities	Prediction, estimation, risk analysis, monitoring and decision support	Literature established premise	Taboada et al., 2023; Nenni et al., 2024; Adamantiadou & Tsironis, 2025; Salimimoghadam et al., 2025
Conceptual component	What prior literature establishes	Evidence status in HG-APM	Principal sources
Agentic orchestration	Goal decomposition, worker allocation, monitoring, adaptation and coordination	Literature established premise	Masters et al., 2025
Governance, delegation and authority	Authenticated delegation, accountability, decision authority and process autonomy	Literature established premise	South et al., 2025; Engin, 2026; Montgomery et al., 2025
Human-AI complementarity	Role partitioning, task structure, coordination and calibrated reliance	Literature established premise	Gonzalez et al., 2026; Tian et al., 2026; O'Neill et al., 2022, 2023
PM-specific implementation and maturity	Organizational readiness, implementation pathways, autonomy levels and governance	Literature established premise	Kalluri, 2026
Project-specific integration	Connecting project state, dependencies, competing constraints, decision rights and post-action feedbacks in one decision cycle	Integrative conceptual inference / HG-APM synthesis	Synthesized from the above literatures and project governance research
State-contingent authorization	Permitted autonomy varies with consequence, uncertainty, reversibility, novelty and data quality	HG-APM theoretical proposition	Proposed for empirical testing; informed by governance/delegation literature but not treated as an established finding

2.9 Boundary of inference

The traceability distinction is consequential for the manuscript's claims. The existence of autonomous planning, governance gates, dynamic autonomy, human–AI teaming and feedback mechanisms is supported by prior work; their combination into HG-APM is a conceptual synthesis. The specific proposition that autonomy should be recalibrated at the level of the governed project decision is proposed here and is not presented as empirically established. Likewise, the eleven-stage sequence is an architectural specification, not a finding that the sequence has already been shown to outperform alternative architectures. The recent literature substantially narrows the space for claiming a general gap in agentic project management. AI-enabled project-management research has established applications in prediction, risk analysis, estimation, monitoring and decision support, while newer work has moved toward generative and agentic systems (Nenni et al., 2024; Khalil et al., 2025; Adamantiadou and Tsironis, 2025; Salimimoghadam et al., 2025; Anghel and Cioara, 2026). Masters et al. (2025) further demonstrate that autonomous Manager Agents can decompose goals, allocate work, monitor progress and adapt workflows, while also showing that multi-objective coordination remains difficult. Kalluri (2026) now provides an explicit implementation framework and maturity model for organizations adopting agentic AI in project management, including autonomy levels, implementation pathways, governance mechanisms, and performance measurement. Human–AI teaming research likewise provides a growing basis for role complementarity and project-specific evaluation (Gonzalez et al., 2026; Tian et al., 2026).

The remaining gap is therefore narrower and should be stated as an integration and translation problem rather than as the absence of governance architectures. Recent work already provides several closely related mechanisms. HAIG conceptualizes decision authority, process autonomy and accountability as continuous governance dimensions; Cohen and Sadeh (2026) propose an Agentic Cognitive Workflow for project management that combines context profiling, graduated autonomy gates, explicit decision rights, audit trails and adaptive reconfiguration; and LATTICE demonstrates a governance-first architecture in which planning, execution and governance are separated and sensitive tool actions pass through explicit approval controls (Engin, 2026; Cohen and Sadeh, 2026; Calboreanu, 2026). These studies materially narrow the novelty space. The unresolved problem addressed here is more specific: how to translate such governance principles into the semantics of project decisions by connecting continuous project-state reasoning and dependency analysis to a differentiated decision boundary in which an intervention is evaluated against multiple project constraints before the relevant authority authorizes execution. HG-APM therefore positions itself as a project management specific decision control synthesis, not as the first framework for dynamic autonomy, human approval or agent governance with its distinctive analytical emphasis on the governed project decision and the explicit Evaluate → Authorize → Act transition within a dependency-aware, closed-loop PM architecture.

3. Methodology

3.1 Research design

This study uses an integrative conceptual synthesis for theory development rather than a new empirical dataset. It integrates findings from four adjacent research streams: AI-enabled project management, autonomous Manager Agents, human oversight/governance and human–AI teaming. The purpose is to derive a domain specific conceptual model, not to claim an exhaustive systematic review.

3.2 Literature identification and selection

The synthesis used four analytical domains: AI-enabled project management capability; autonomous managerial orchestration; human oversight, governance and delegation; and human–AI teaming. Source identification was iterative rather than exhaustive, with priority given to peer-reviewed and authoritative scholarly work through 2026, especially recent research on agentic AI, Manager Agents, delegation, governance and project teaming. Sources were included when they addressed at least one of four substantive criteria: an AI capability relevant to PM; autonomous goal pursuit, planning, orchestration or workflow adaptation; authority, delegation, oversight, intervention, accountability or auditability; or human–AI teaming, complementarity, trust or workload. Sources without a substantive connection to these domains were excluded from the synthesis.

3.3 Analytical extraction and comparison

Each retained source was examined for its specific contribution rather than treated as evidence for the whole model. The analysis extracted the capability or construct addressed, unit of analysis, governance or teaming mechanism, and implication for project decision making. Sources were then compared for convergence, omission and tension, with particular attention to whether authority was separated from capability, whether project dependencies were represented, and whether post-action monitoring or replanning was addressed.

3.4 Conceptual synthesis and model derivation

HG-APM was derived through progressive integration rather than by adopting one existing framework as its template. AI-PM research supplied analytical capabilities; Manager Agent research supplied planning, orchestration and adaptation; governance/delegation research supplied authority, intervention, accountability and auditability; and human–AI teaming research supplied role differentiation and complementarity. These strands were then translated into a project-decision architecture organized around persistent project state, consequence-sensitive evaluation, authorization and feedback.

3.5 From synthesis to propositions

Propositions were generated only after the architecture, and construct boundaries had been specified. Each proposition links a proposed HG-APM mechanism to an observable outcome or boundary condition and is explicitly framed as an empirical expectation rather than an established finding. The propositions were subsequently audited for construct distinctiveness, plausible comparison conditions and potential falsifiability. This procedure is intended to prevent design principles from being presented as empirical propositions.

3.6 Auditability, researcher judgement and limitations

The synthesis involves researcher judgement because the literature spans disciplines, methods and levels of analysis. The audit trail is represented by the four analytical domains, source to concept traceability, explicit separation of evidence from inference, construct definitions, and proposition mapping. The approach does not claim exhaustive retrieval, formal inter-rater coding or statistical aggregation; the model therefore requires subsequent systematic and empirical validation.

3.7 Analytical principle

A strict distinction was maintained between what a source establishes, what it proposes conceptually, and what is inferred for project management. Because several recent studies are conceptual or benchmark oriented rather than longitudinal demonstrations of deployed autonomous project managers, HG-APM is presented as a theory-development contribution requiring empirical testing.

4. The Proposed Framework

4.1 The Human-Governed Agentic Project Management model

The proposed Human Governed Agentic Project Management model treats project management as a closed-loop decision system. The AI agent maintains a continuously updated representation of project state and moves through eleven stages: Sense, Understand, Diagnose, Predict, Plan, Evaluate, Authorize, Act, Monitor, Learn and Replan. The cycle is not strictly linear. Monitoring can trigger a return to diagnosis; new information can invalidate a plan; authorization can be withheld; and learning can modify future decision policies.




Figure 2. Human Governed Agentic Project Management closed-loop architecture.

Stage	Agentic PM function
Sense	Continuously collect and detect changes in schedule, cost, resources, risks, issues, dependencies, quality, communications, stakeholder signals and project artifacts.
Understand	Integrate heterogeneous evidence into a current project-state representation, including objectives, constraints, dependencies, resource availability, decisions and stakeholder context.
Diagnose	Determine likely causes and causal relationships rather than merely identifying deviations; distinguish symptoms from root causes and identify affected downstream elements.
Predict	Forecast consequences, probabilities and trade-offs under alternative responses, including time, cost, resource, quality, risk and stakeholder effects.
Plan	Generate feasible intervention options, sequence activities, account for dependencies, and identify required resources, permissions, and communications.
Evaluate	Test alternatives against project objectives, data quality, constraints, ethics, reversibility, uncertainty, financial, contractual, safety, stakeholder and strategic consequences.
Authorize	Apply the autonomy policy. Low-risk actions may be pre-authorized within an autonomy envelope; consequential actions require the designated human decisionmaker.
Act	Execute authorized actions through connected project management systems, communication channels, or approved tools while preserving traceability.
Monitor	Measure whether the action produced the intended result and detect new deviations, side effects, exceptions or escalation conditions.
Learn	Update the project knowledge state and decision policy from outcomes, human feedback, errors, exceptions and changing conditions.
Re-plan	Modify the project plan, dependencies, priorities, or recommended actions when new evidence changes the feasible or desirable course of action.

Table 3. HG-APM stages and their agentic project management functions.

Sense and Understand: persistent project state. An agentic PM requires more than access to isolated project documents. It needs a persistent project state that integrates structured and unstructured information. Relevant inputs include the baseline schedule, actual progress, resource availability, budgets and forecasts, risk and issue registers, dependencies, change requests, contracts and obligations, quality indicators, meeting records, stakeholder communications and project artifacts. The state must also record provenance and confidence so that the agent can distinguish verified facts from assumptions or stale information.

This requirement responds directly to the partial observability problem identified by Masters et al. (2025). The agent should be capable of recognizing not only what it knows but also what it does not know. Missing or contradictory information should therefore trigger an evidence request or escalation rather than silent inference. Data quality becomes an operational governance condition, not merely a technical concern.

Diagnose: from deviation detection to causal project reasoning. A distinctive element of the proposed model is the separation of understanding from diagnosis. A conventional dashboard may detect that an activity is late. An agentic PM should go further by reasoning why the activity is late, which dependencies are affected, how the delay propagates, and which intervention points are available. For example, a delayed engineering task may be caused by unavailable specialist resources, which may itself result from another workstream's schedule change. The intervention may therefore require cross-workstream resource negotiation rather than simply moving the delayed activity.

Predict and Plan: Multi-objective project reasoning. Project decisions are rarely single objective. A schedule recovery recommendation can affect cost, quality, staff workload, procurement, contractual commitments and stakeholder expectations. The model therefore requires the AI to generate alternatives rather than a single supposedly optimal action. For each alternative, the agent should identify assumptions, dependencies, resource requirements, expected benefits, adverse effects, and reversibility. Consider a schedule that has become tight. The AI might recommend moving work to another engineering workstream. Before recommending execution, it should determine whether two suitably skilled engineers are available, whether reallocating them creates a new bottleneck, whether contractors must be notified, whether additional procurement is required, whether overtime is necessary, whether overtime is permissible and compensated, and whether the change affects contractual milestones. This is the kind of multi-constraint reasoning that distinguishes agentic project management from automated scheduling.

Evaluate: the decision gate before action. Evaluation is deliberately positioned before authorization. The agent's recommendation should be evaluated against the project's decision criteria and the quality of the evidence supporting it. At minimum, the evaluation should consider: data recency; data accuracy and completeness; project context fit; uncertainty; reversibility; novelty; financial consequence; schedule consequence; scope consequence; quality consequence; resource consequence; contractual consequence; stakeholder consequence; safety consequence; reputational consequence; and strategic consequence. This stage also incorporates ethical and organizational constraints. An AI should not infer that overtime is acceptable simply because it improves schedule performance. It should surface the requirement, identify policy constraints, and ask the authorized human decisionmaker where the matter involves worker

welfare, employment rules or material cost. Similarly, contractor notification is not sufficient: the agent should identify what information or commitment is required and whether the contractor is available before an implementation path is chosen.

Authorize: human authority as an explicit control layer. Authorization is the defining governance feature of the proposed architecture. The AI may recommend actions and may autonomously execute predefined low risk actions, but it does not silently cross authority boundaries. Every autonomous action must be traceable to an authorization rule or explicit approval. The human PM remains continuously informed through notifications and decision logs rather than being placed outside the loop. The model therefore rejects a binary distinction between 'human-in-the-loop' and 'human-out-of-the-loop'. Instead, it proposes graduated autonomy with human visibility at all levels. Routine actions can be executed automatically while the PM receives an event notification. Medium risk actions can be recommended and held for approval. High consequence actions require explicit human authorization, potentially involving functional managers, sponsors, legal officers, safety personnel or other authorized stakeholders.

Act, Monitor, Learn and Re-plan. After authorization, the AI may execute the approved action through project-management software, communication tools or other connected systems. Execution must generate an auditable action record showing what was done, when, using which data and under which authorization. Monitoring then tests whether the intervention worked. If it did not, the agent should not simply repeat the same action. It should diagnose the new state, learn from the outcome and re-plan.

This creates a true closed-loop architecture. The endpoint of action is not completion of the AI's task; it is verification of the project's response to the intervention. Learning therefore concerns not only model performance but also project-specific knowledge, organizational policies, human feedback, successful and unsuccessful interventions and changes in stakeholder priorities.

4.2 Dynamic autonomy and human authority

The proposed model does not assign autonomy based solely on financial value or a single confidence score. Financial considerations are important, but a low monetary value can still create contractual, safety, reputational, or ethical consequences. Autonomy is therefore determined by a multidimensional consequence profile. Let an action a be evaluated on consequence dimensions $C = \{\text{financial, schedule, scope, quality, resource, contractual, stakeholder, safety, reputational, strategic}\}$. Let U represent uncertainty, R reversibility, N novelty and D data quality. The autonomy level $A(a)$ is conceptually determined by the combined consequence and evidence profile rather than by a fixed monetary threshold:

$A(a) = f(C, U, R, N, D, \text{policy constraints, stakeholder authority})$ This is a conceptual decision rule, not a validated mathematical formula. Its purpose is to make the governance logic explicit and testable. High consequence, high uncertainty, low reversibility, high novelty, or poor data quality should reduce permitted autonomy. Low consequence, highly reversible, routine actions supported by reliable data may receive greater autonomy.




Figure 3. Graduated autonomy: Five tiers of AI-human decision authority.

Autonomy tier	Typical examples	Human role	AI role
Tier 1: Autonomous routine	Status synchronization, reminders, report refreshes, low-risk notifications	Informed; may review audit trail	Execute within predefined rules
Tier 2: Supervised recommendation	Routine scheduling adjustments, task sequencing, low-impact coordination	Review/approve when policy requires	Analyze, recommend, prepare action
Tier 3: Collaborative decision	Resource trade-offs, complex risk responses, major replanning	Joint decision; PM leads where authority applies	Model alternatives, forecast effects, document trade-offs
Tier 4: Human authorized execution	Material schedule, scope, procurement or stakeholder commitments	Explicit authorization by designated authority	Prepare and execute only after authorization

Autonomy tier	Typical examples	Human role	AI role
---------------	------------------	------------	---------

Tier 5: Human-only / restricted	Budget authority, major contractual commitments, safety critical decisions, termination, sensitive HR decisions	Human/authorized governance body decides	Advise, evidence, scenario analysis only
---------------------------------	---	--	--

Table 4. Autonomy tiers, typical examples, and role allocation.

Human PM role under agentic augmentation. The model does not redefine the project manager as a passive supervisor. It redistributes managerial effort. Routine monitoring, repetitive follow-ups, information synthesis, status preparation and low-risk coordination can increasingly be handled by the agent. The human PM's relative emphasis shifts toward strategic decision-making, stakeholder management, team leadership, negotiation, governance, risk judgment, ethical judgment, conflict resolution and high-consequence problem solving.

This is a proposition rather than an established empirical finding. The literature supports the possibility that AI can absorb routine cognitive and coordination work, but whether this releases human PM capacity for strategic work must be tested empirically. The appropriate research question is therefore not whether AI replaces PMs, but whether governed delegation increases the strategic and relational capacity of PMs while maintaining project control.

Human visibility without human overload. Continuous human notification creates a potential contradiction. If the AI alerts the PM for every action, the PM becomes an approval bottleneck and may experience alert fatigue. If the AI suppresses too much information, the PM can lose situational awareness. The solution proposed here is event-based transparency: every autonomous action is logged and visible, while real-time notifications are prioritized by consequence, uncertainty, novelty, policy conflict and required intervention latency.

4.3 Human–AI teaming architecture

The HG-APM model treats the PM and AI as interdependent contributors to a shared project goal. The team architecture is therefore organized around complementary capabilities. AI contributes continuous attention, data synthesis, forecasting, consistency, workflow coordination, and procedural execution. The human PM contributes contextual judgment, ethical reasoning, leadership, stakeholder legitimacy, negotiation, political and organizational interpretation, and accountability within the organization's governance structure.

Capability	AI-dominant contribution	Human-dominant contribution
Reasoning	Data synthesis, scenario modelling, trade-off analysis, anomaly detection	Context, values, ethics, legitimacy, strategic judgment
Memory	Persistent project knowledge, provenance, retrieval and traceability	Institutional meaning, tacit knowledge, experience

Capability	AI-dominant contribution	Human-dominant
------------	--------------------------	----------------

		contribution
Attention	Continuous monitoring, triage, alerts and early-warning signals	Novelty detection, contextual prioritization and exception judgment
Coordination	Workflow orchestration, reminders, dependencies and routine communication	Role negotiation, conflict resolution, leadership and stakeholder alignment
Governance	Rule enforcement, logging, evidence and escalation	Decision rights, authorization, policy interpretation and accountability
Adaptation	Pattern learning, replanning and response simulation	Creative restructuring, strategic change and organizational judgment

Table 5. Complementary AI-dominant and human-dominant contributions by capability.

4.4 Project-governance decision boundary

A major implication of the model is that autonomy should be assigned to decision classes, not to an AI system in the abstract. The same AI agent may have high autonomy for one activity and almost no autonomy for another. This prevents the common mistake of treating 'the AI' as having a single global autonomy level.




Figure 4. Project-governance decision boundaries by decision domain.

Decision domain	Default AI authority	Required human involvement
Routine monitoring and reminders	High	Notification / audit visibility
Task follow-up and administrative coordination	High within predefined rules	Periodic review
Schedule sequencing with no material external impact	Moderate to high	Review if thresholds are crossed
Resource allocation across workstreams	Moderate	PM or resource owner approval when contention exists
Overtime or worker-welfare implications	Low	Human judgment and policy review
Budget changes / material financial commitments	Very low	Human financial authority
Scope changes	Very low	PM plus authorized change authority
Contractual commitments / changes	Very low	PM plus legal/contract authority
Safety-critical decisions	Restricted	Human competent authority
Major stakeholder or reputational decisions	Restricted	Human leadership / authorized stakeholder
Project termination or strategic redirection	Human-only	Sponsor / governance authority

Table 6. Default AI authority and required human involvement by decision domain.

4.5 Propositions for empirical testing

P1 : Consequence-sensitive authorization and decision quality. Controlling for task complexity and project context, project decisions made under consequence-sensitive autonomy with explicit authorization gates are expected to exhibit higher decision quality than comparable decisions made under undifferentiated autonomous execution. Decision quality should be assessed using criteria such as constraint satisfaction, trade-off quality, downstream consequence management, and decision appropriateness rather than speed alone.

P2: Data quality as a boundary condition. The positive relationship between permitted agentic autonomy and project decision quality is expected to weaken as project data become more incomplete, stale, inconsistent, or poorly governed. Thus, data quality is specified as a moderator of the autonomy–decision-quality relationship rather than as a direct substitute for authorization.

P3: Closed loop adaptation. Compared with a static recommendation architecture that produces recommendations without integrated post-action monitoring and replanning, an agentic PM operating through the full Sense–Understand–Diagnose–Predict–Plan–Evaluate–Authorize–Act–Monitor–Learn–Re-plan cycle is expected to achieve greater adaptation to changing project conditions, operationalized through measures such as recovery from disturbances, replanning quality and time to effective response.

P4: Human–AI complementarity and performance. The positive relationship between agentic PM deployment and project performance is expected to be stronger when human and AI responsibilities are complementary rather than duplicative, particularly under high task complexity and uncertainty. Complementarity should be measured through role differentiation, capability fit and coordination quality rather than inferred merely from the presence of human oversight.

P5: Authorization transparency and calibrated trust. Greater transparency of decision rights, authorization status, action provenance and intervention capability are expected to improve calibrated trust in the agent compared with opaque autonomous execution. Calibrated trust should be assessed by alignment between users' reliance and the agent's demonstrated capability, rather than by trust intensity alone.

P6: Automation, escalation, and human workload. Routine autonomous monitoring and coordination are expected to reduce administrative workload when escalation remains proportionate to decision consequence; as escalation frequency and low-value alerts increase beyond the level that users can effectively process, cognitive workload is expected to increase and the workload benefit of autonomy to diminish. This proposition therefore predicts a conditional, potentially non-linear workload relationship rather than assuming that more autonomy always reduces workload.

P7: Multi-objective reasoning and project performance. Agentic PM performance is expected to decline when the agent optimizes a single project objective without explicitly representing competing time, cost, scope, quality, risk, resource and stakeholder constraints. The proposition is falsifiable by comparing single-objective and multi-objective decision configurations while holding the underlying project problem constant.

The propositions should be tested at the level of the governed project decision wherever possible. This preserves the model's unit of analysis and allows researchers to vary autonomy, data quality, consequence, task interdependence, authorization structure and feedback conditions independently. It also prevents the propositions from being reduced to a simple comparison of "AI" versus "human" project management.

P8: State contingent authorization. The relationship between agentic capability and autonomous execution is expected to depend on a dynamic authorization boundary that recalibrates permitted autonomy as project consequence, uncertainty, reversibility, novelty and data quality change. Autonomy is therefore theorized as an action level governance state rather than a fixed property of the AI system. Evidence against systematic changes in permitted autonomy or decision outcomes when these conditions change would challenge the proposition. These propositions are derived from the literature synthesis and are proposed for empirical testing; the cited studies provide theoretical and empirical premises but do not constitute direct evidence that HG-APM produces these outcomes.

Construct definitions and discriminant logic. The propositions require construct boundaries that prevent conceptual overlap between governance variables, human–AI team variables, and project outcomes. This section therefore defines the principal constructs at the level at which they are intended to be observed and specifies relationships that should not be treated as interchangeable. The purpose is not to prescribe a single measurement instrument, but to establish construct validity requirements for future empirical testing.

Agentic autonomy is the degree to which an AI agent can initiate, sequence or execute project actions without a prior human decision.

Authorization, by contrast, is the legitimate permission for a particular project action to proceed

under the governance policy in force. A distinct and prior, question from whether the agent is technically capable of acting.

Decision quality concerns how good a specific project decision is relative to the objectives and constraints that applied at the time it was made.

Project performance sits one level up: it is an outcome construct capturing aggregate results schedule, cost, scope, quality rather than any single decision.

Adaptation is the project system's ability to detect changed conditions, revise its response, and re-plan accordingly.

Human–AI complementarity describes the extent to which human and AI capabilities are differentiated and jointly exploited rather than duplicated.

Calibrated trust is the alignment between how much a human relies on AI outputs and how reliable those outputs actually are.

Workload, here, means the cognitive and administrative burden experienced by the human project participants, not a proxy for how much the AI is doing.

Multi-objective reasoning is the explicit representation and evaluation of competing project objectives, rather than optimization against a single metric.

These distinctions establish a discriminant logic for the propositions: autonomy and authorization are governance conditions; decision quality, adaptation and calibrated trust are decision or human–AI interaction outcomes; project performance is a higher-level project outcome; complementarity is a team-design condition; workload is a human outcome; and multi-objective reasoning is a reasoning/design condition. Future empirical studies should test these constructs separately rather than using a single composite “agentic PM effectiveness” measure.

Construct	Role in HG-APM	Should not be equated with	Primary proposition(s)
Agentic autonomy	Degree of self-directed execution within a context	Authorization; capability; project performance	P1, P2, P8

Construct	Role in HG-APM	Should not be equated with	Primary proposition(s)
Authorization	Legitimate permission for a specific action	Autonomy; human involvement; accountability	P1, P5, P8
Decision quality	Quality of an individual project decision against objectives and constraints	Speed; prediction accuracy; project performance	P1, P2
Adaptation	Ability to respond and re-plan after changing conditions	Autonomy; project performance	P3
Human–AI complementarity	Fit and differentiation of human and AI capabilities	Human oversight; simple co-presence	P4
Calibrated trust	Alignment of reliance with demonstrated capability	Trust intensity; satisfaction; accuracy	P5
Human workload	Administrative and cognitive burden on human actors	Project complexity; autonomy	P6
Multi-objective reasoning	Explicit handling of competing objectives and constraints	Complexity; performance outcome	P7
Project performance	Aggregate project outcomes over an appropriate horizon	Decision quality; adaptation	P4, P7

Table 7. Construct definitions, discriminant logic and related propositions.

5. Discussion

5.1 What the model adds to existing research

HG-APM should be understood as a focused architectural synthesis, not as the origin of dynamic autonomy, human approval, accountability, or agent governance. Kalluri (2026) addresses implementation and maturity; HAIG addresses authority, autonomy and accountability; Cohen and Sadeh (2026) provide a PM-specific adaptive governance workflow; and LATTICE provides technical runtime governance. HG-APM instead integrates these established mechanisms around the governed project decision.

The model integrates project intelligence with governance, governance with human–AI teaming, and action with learning. AI reasoning is bounded by an autonomy envelope; human and AI responsibilities are differentiated; and execution feeds monitoring, learning and replanning.

The central mechanism is a dynamic authorization boundary between project-state intelligence and autonomous execution. Changing project conditions alter the consequence and evidence profile of a decision, which affects permissible autonomy; authorization then determines whether the

intervention may cross into execution. Execution produces new evidence that can change the next authorization boundary.

5.2 Why authorization is separated from evaluation

Separating evaluations from authorization is theoretically important. If the AI evaluates its own recommendation and immediately executes it, the architecture risks collapsing assessment and authority into a single automated loop. In the proposed model, evaluation produces an evidence-based decision package, while authorization applies organizational decision rights. This makes it possible to preserve human authority without requiring humans to reproduce every computational step performed by the AI.

5.3 Why project dependencies matter

Project decisions are relational. Moving one activity affects downstream activities; reallocating an engineer affects another workstream; increasing overtime affects cost and fatigue; notifying a contractor can create an external commitment; changing sequence can affect testing and quality assurance. Consequently, agentic project management must reason over a dependency graph rather than isolated tasks. The Manager Agent literature supports this direction through graph-based workflow representations, but the proposed model extends the logic by requiring dependency consequences to be evaluated before authorization.

5.4 Accountability should be distributed explicitly

A simplistic statement that 'the PM is accountable for the AI' is insufficient. Autonomous project management creates distributed accountability involving the AI developer, system owner, organization, PM, functional managers, approvers and other governance actors. The model therefore calls for action level traceability and an explicit accountability map. Each autonomous action should be attributable to a system identity, a governing policy, a data state, an authorization condition and a responsible organizational owner.

5.5 The central design tensions

The model exposes a fundamental tension. Too little autonomy makes the AI operationally weak and turns the PM into a bottleneck. Too much autonomy makes human oversight nominal and increases the consequences of model error, data failure, or inappropriate action. The objective is therefore not maximum autonomy. It is appropriate autonomy: enough authority to create meaningful operational value while preserving legitimate human control over consequential decisions.

Operationalizing the HG-APM architecture in project environments

Agentic PM should therefore be implemented as a decision system rather than simply as an AI feature. Each autonomous capability should be associated with a decision class, authority boundary, monitoring requirement, escalation path and audit trail.

Autonomy should be evaluated against consequence, uncertainty, reversibility, novelty and data quality, rather than a single global system label or confidence score.

The architecture therefore supports a decision rights register identifying which actor may recommend, authorize, execute, override and audit each autonomous capability.

Third, the architecture requires explicit treatment of disagreement. An agentic PM should not be designed merely to agree with the human PM, because such a system would reproduce human

errors and provide little independent decision value. The agent should be able to identify contradictions, present evidence, explain the consequences of alternative courses of action, and recommend escalation when a decision exceeds its authority. However, disagreement must remain procedurally bounded. The agent may challenge, recommend, and escalate; the human decisionmaker retains authority where the decision class requires human authorization. This creates a constructive form of dissent without converting the agent into an unaccountable decisionmaker.

5.6 Integrative comparison with existing frameworks

HG-APM is an integrative decision control architecture rather than a replacement for existing capability, teaming, or governance frameworks. Its proposed contribution is the PM specific integration of project-state reasoning, decision rights, evaluation, authorization, execution, and feedback.

Approach	Primary unit / level	Governance or autonomy focus	PM-specificity	Position relative to HG-APM
Kalluri (2026)	Organization adoption	Maturity, implementation pathways, governance, autonomy levels	Project management	Complements HG-APM; organizational implementation rather than action-level decision control.
Masters et al. (2025)	Manager agent / workflow	Goal decomposition, allocation, coordination, adaptation	General managerial workflows	Provides orchestration capabilities; HG-APM adds explicit PM decision rights and authorization semantics.
South et al. (2025)	Agent delegation / infrastructure	Authenticated, auditable authority	General AI agents	Provides authority/delegation foundation; HG-APM

Approach	Primary unit / level	Governance or autonomy focus	PM-specificity	Position relative to HG-APM
				translates it into PM decision classes.
Engin HAIG (2026)	Human–AI governance relationship	Decision authority, process autonomy, accountability continua	Domain agnostic	Overlaps on dynamic authority/autonomy; HG-APM specializes in the dimensions to PM decisions and project dependencies.
Cohen & Sadeh (2026)	PM governance workflow	Graduated autonomy, decision rights, audit, escalation, adaptive balance	Project management	Closest conceptual overlap; HG-APM differentiates through governed-decision unit, Evaluate → Authorize → Act, and dependency-aware reasoning.
Calboreanu (2026) LATTICE	Agent execution architecture	Policy enforcement, approval gates, separation of concerns	Domain-agnostic	Technical governance complement; HG-APM specifies PM decision semantics rather than runtime enforcement.
HG-APM	Governed project decision	Evaluate, authorize, execute, monitor and re-plan under dynamic conditions	Project management	Integrative conceptual architecture for PM-specific, dependency-aware decision control; requires empirical validation.

Table 8. HG-APM positioned against recent agent-governance and comparison approaches.

5.6.1 Relationship to AI-enabled project management

HG-APM extends this literature by treating those capabilities as components of a persistent managerial decision process rather than as isolated applications. Forecasting contributes to Predict; optimization contributes to Plan and Evaluate; risk analytics contributes to Diagnose and Predict; natural-language synthesis contributes to Understand; and event detection contributes to Sense and Monitor. The proposed contribution is therefore not a new claim that AI can perform these individual functions. It is their organization within a governed cycle in which recommendations are assessed against project objectives and constraints before authority to act is established.

5.6.2 Relationship to the Manager Agent

The Manager Agent paradigm is the closest technical analogue to HG-APM. Masters et al. (2025) describe an autonomous manager capable of decomposing goals, allocating work to human and AI workers, modifying workflows, monitoring progress and adapting coordination. Their work establishes the importance of an autonomous managerial layer for orchestrating heterogeneous workers and also identifies challenges involving compositional reasoning, multi-objective optimization, coordination and governance.

HG-APM does not claim that these capabilities are absent from the Manager Agent concept. Instead, it adds a project-governance decision layer around them. First, diagnose is separated from simple observation or workflow restructuring so that causal project-state reasoning is explicit. Second, Evaluate and Authorize are positioned between planning and consequential execution. Third, the decision boundary is expanded to project-specific consequences such as budget, scope, resources, contracts, safety, stakeholder relationships, and strategic commitments.

5.6.3 Relationship to HOGF-AI

HOGF-AI provides a governance-oriented foundation for controlling autonomous agents through mechanisms such as an autonomy envelope, risk gates, intervention, accountability, monitoring and auditability (Montgomery et al., 2025). Recent governance research reinforces this direction by treating delegation, autonomy and control as distinct organizational choices and by documenting the leadership dilemmas that arise when agent autonomy is extended in regulated project-team settings (Srivastava, 2026; Pathak and Panteli, 2026).

HG-APM extends this governance logic into project management decision classes. Rather than treating governance as a general control layer around an AI system, the proposed model asks what the agent may do for project decisions and under what conditions. Routine monitoring may therefore be pre-authorized, while a material budget commitment, contractual change, safety-critical intervention or project termination may require a different authorized human decisionmaker.

5.6.4 Relationship to recent agent-governance architectures

Engin (2026) treats decision authority, process autonomy and accountability as continuous governance dimensions with thresholds. This overlaps with HG-APM's context sensitive autonomy logic; HG-APM therefore claims only the PM-specific instantiation of these dimensions within project decision classes, dependencies and decision rights.

Cohen and Sadeh (2026) are closer in domain: their Agentive Cognitive Workflow combines context profiling, graduated autonomy gates, decision rights, audit trails, escalation and adaptive reconfiguration for agentive PM. The distinction is analytical rather than wholesale model novelty: HG-APM makes the

governed project decision the unit of analysis and centers the semantic separation of Evaluate from Authorize and Act, together with dependency-aware project-state reasoning.

LATTICE provides an engineering-level comparison. Calboreanu (2026) separates planning, execution and governance and uses approval gates to prevent unauthorized actions. HG-APM operates at a different abstraction level: it specifies PM decision semantics and authority rather than runtime enforcement.

5.6.5 Relationship to human–AI complementarity

Recent human–AI teaming research emphasizes that effective teams depend on complementary capabilities, clear roles, coordination and appropriate autonomy rather than simple substitution of humans by machines. Tian et al. (2026) identify human-centered AI, generative AI risk analysis, agile execution and human robot integration as major themes in project-management HAT research, while Zamecnik et al. (2026) show that trust research needs stronger alignment between agent architecture and the characteristics being evaluated. Gonzalez et al. (2026) further frame complementarity around reasoning, memory, attention, role partitioning, shared mental models, trust calibration, and continuous evaluation.

HG-APM operationalizes this complementarity by linking role allocation to decision authority. AI-dominant responsibilities include continuous sensing, data synthesis, scenario analysis, routine coordination, and evidence logging. Human-dominant responsibilities include authorization, policy interpretation, stakeholder-sensitive judgment, and high-consequence strategic decisions. Joint responsibilities arise where the agent provides alternatives or evidence, and the human evaluates them in context.

5.7 Theoretical implications

The model extends project management AI research from function specific applications toward a continuous decision architecture.

It connects human–AI complementarity theory to project governance by locating role allocation and decision authority inside the project decision process rather than treating governance as a separate compliance layer. It translates general oversight and delegation requirements into project-specific decision domains such as scope, budget, schedule, resources, contracts, safety and stakeholders, while acknowledging that the underlying governance mechanisms are established outside PM research. It provides a basis for empirically studying how established concepts of dynamic authority and process autonomy behave when instantiated in project-management decision classes and dependency structures.

5.8 Practical implications

Organizations considering agentic project management should not begin by asking which project-management tasks can be automated. They should first map decision rights. Each candidate AI action should be classified by consequence, reversibility, uncertainty, novelty, data quality and stakeholder impact. The organization should then define which actions are pre-authorized, which require approval, and which are prohibited.

Second, organizations need an auditable project-state layer. An agent cannot safely manage a project from disconnected spreadsheets, stale dashboards and undocumented conversations. Data provenance, update frequency, conflict resolution, and access rights should be explicit. Third, organizations should train PMs in AI literacy, governance, data interpretation and human–AI teaming. The human PM's role is not reduced; it becomes more dependent on judgment, leadership

and governance capability.

Fourth, implementation should be staged. Organizations can begin with high-autonomy routine actions, then expand authority only when monitoring; data quality, intervention mechanisms and evidence of reliable performance justify it. The objective is not to progress toward full autonomy automatically, but to increase autonomy where evidence shows that it is safe, useful, and governable.

6. Limitations and Future Research

This paper is conceptual and does not empirically validate the HG-APM architecture. The propositions require testing in real project environments. Future research should use controlled experiments, simulation, field studies, and longitudinal deployments. MA-Gym and similar environments provide useful starting points for testing task decomposition, multi-objective coordination and workflow adaptation, but project-management-specific benchmarks are needed. Fifth, future research should examine disagreement. A useful agentic PM should be able to disagree with a human recommendation when evidence indicates a serious risk, while still respecting the human authority. This creates an important design problem: how can an agent challenge a decision without becoming an unauthorized decisionmaker? Finally, cross-industry research is needed. Existing AI-in-PM evidence is concentrated in construction and IT, but project governance requirements vary across healthcare, infrastructure, energy, public-sector, creative and other environments.

6.1 Boundary conditions for safe agentic PM

Agentic autonomy should be reduced when project data are insufficient or contradictory; objectives or stakeholder priorities are unclear; decisions have high-impact or irreversible consequences; legal, contractual or safety exposure is material; situations are novel or outside the validated operating envelope; authority is contested; or human intervention cannot occur within the required latency. These are boundary conditions for delegation, not universal thresholds. The proposed architecture is deliberately conditional. Agentic autonomy should be reduced when the project environment violates one or more of the following conditions: (1) insufficient or contradictory data; (2) unclear objectives or changing stakeholder priorities; (3) high-impact decisions with irreversible consequences; (4) significant legal or contractual exposure; (5) safety critical activity; (6) novel situations outside the agent's validated operating envelope; (7) contested stakeholder authority; or (8) evidence that the human reviewer cannot intervene within the required latency.

6.2 Measurement framework for future validation

Future validation should measure four outcome families: project performance (schedule, cost, scope, quality, risk and stakeholder outcomes); decision performance (decision quality, trade-off quality, intervention latency, exception detection and replanning); human outcomes (workload, cognitive load, calibrated trust, perceived control and strategic capacity); and governance outcomes (authorization compliance, audit completeness, unauthorized-action rate, escalation quality, intervention success and accountability traceability).

A credible empirical evaluation of HG-APM should measure at least four outcome families. The first is project performance: schedule adherence, cost variance, scope stability, quality, risk exposure, resource utilization and stakeholder outcomes. The second is decision performance:

decision accuracy, trade-off quality, intervention latency, exception detection and quality of replanning. The third is human outcomes: PM workload, cognitive load, trust calibration, perceived control, strategic capacity, and willingness to use the system. The fourth is governance performance: authorization compliance, audit completeness, unauthorized-action rate, escalation appropriateness, intervention success and accountability traceability.

6.3 Proposed research agenda

A staged research programme is proposed: validate the decision taxonomy with project and governance experts; compare human-only, AI-assisted and governed-agentic conditions in controlled scenarios; vary autonomy and data-quality conditions; evaluate the architecture longitudinally in real projects; and develop benchmarks incorporating dependencies, competing workstreams, stakeholder disagreement, resource contention and changing priorities. Dynamic authority is a particularly important test: future systems should assess whether permitted autonomy changes appropriately as consequence, uncertainty, reversibility, novelty and data quality change. The model supports a staged research programme. The first stage should validate the decision taxonomy with experienced project managers and governance specialists. The second should test the model through controlled scenarios in which the same project problem is handled by a human-only team, an AI-assisted PM and a governed agentic PM. The third should vary autonomy and data-quality conditions to identify thresholds at which human intervention improves or reduces outcomes. The fourth should evaluate the architecture in real project environments through longitudinal field studies. The fifth should develop standardized benchmarks that include dependencies, competing workstreams, stakeholder disagreement, changing priorities, resource contention, and contractual constraints.

7. Conclusion

Agentic AI creates an opportunity to move project management from episodic AI assistance toward continuous, context-aware and governed machine participation. Yet autonomous execution should not be treated as equivalent to autonomous project management. AI-enabled PM research establishes analytical capabilities; Manager Agent research establishes emerging orchestration mechanisms and exposes technical limitations; governance research establishes authority and control; and human–AI teaming research establishes the importance of differentiated roles and calibrated interaction. The Human-Governed Agentic Project Management model integrates these insights through a closed loop: Sense → Understand → Diagnose → Predict → Plan → Evaluate → Authorize → Act → Monitor → Learn → Re-plan. The model's defining principle is that autonomy is not a permanent property granted to an AI system. It is a governed capability assigned to particular actions under particular conditions. Routine, reversible and low-consequence actions can receive greater autonomy; financial, contractual, scope, safety, stakeholder, reputational and strategic decisions require progressively stronger human authority.

The model therefore does not envision an AI replacing the project manager. It envisions a human–AI managerial team in which AI provides persistent attention, analysis, forecasting, coordination and bounded execution, while human and organizational authorities retain responsibility for strategic, ethical, relational and governance-sensitive decisions. The central research question is empirical: whether this architecture can improve project decision quality and delivery outcomes while preserving meaningful human control.

The theoretical position advanced here is deliberately narrower than a claim of fully autonomous project management. Current evidence does not justify that claim. The contribution is a project-

management-specific decision-control architecture in which the governed project decision is the unit of analysis, and autonomy is treated as a state-dependent authorization condition. The architecture integrates persistent project-state reasoning, dependency-aware multi-objective evaluation, explicit decision rights, and post-action feedback into one closed loop. Its value is therefore not that any individual component is unprecedented, but that the model specifies a falsifiable relationship among project state, consequence profile, permissible autonomy, authorization, and subsequent replanning.

References

- Adamantiadou, D. S., and Tsironis, L. (2025). Leveraging artificial intelligence in project management: A systematic review of applications, challenges, and future directions. *Computers*, 14(2), 66. <https://doi.org/10.3390/computers14020066>
- Amershi, S., Weld, D., Vorvoreanu, M., Fournery, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., and Horvitz, E. (2019). Guidelines for human-AI interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13. <https://doi.org/10.1145/3290605.3300233>
- Anghel, I., and Cioara, T. (2026). Generative AI for IT project management: A systematic review and future research agenda. *Systems*, 14(6), 722. <https://doi.org/10.3390/systems14060722>
- Armanios, D. E., and Tucci, C. L. (2025). Building organizational and project readiness for AI: A dialectical approach. *International Journal of Project Management*. <https://doi.org/10.1016/j.ijproman.2025.102792>
- Baccarini, D. (1996). The concept of project complexity—A review. *International Journal of Project Management*, 14(4), 201–204. [https://doi.org/10.1016/0263-7863\(95\)00093-3](https://doi.org/10.1016/0263-7863(95)00093-3)
- Bovens, M. (2007). Analysing and assessing accountability: A conceptual framework. *European Law Journal*, 13(4), 447–468. <https://doi.org/10.1111/j.1468-0386.2007.00378.x>
- Calboreanu, E. C. (2026). LATTICE: A governance-first architecture for authorized autonomous AI operations. *Frontiers in Artificial Intelligence*, 9. <https://doi.org/10.3389/frai.2026.1800407>
- Cohen, Y., and Sadeh, A. (2026). Balancing human and agentic project management: Autonomy, accountability, and transparency. In *Proceedings/Book of the 20th International Workshop on Project Management and Scheduling (PMS 2026)*, 92–95.
- Dacre, N. (2026). AI readiness is not enough: Towards purpose-led AI governance in projects. *International Journal of Project Management*, 44(4), 102850. <https://doi.org/10.1016/j.ijproman.2026.102850>
- Engin, Z. (2026). Human-AI governance (HAIG): A trust-utility approach. *Journal of Responsible Technology*, 26, 100167. <https://doi.org/10.1016/j.jrt.2026.100167>
- Franklin, S., and Graesser, A. (1996). Is it an agent, or just a program? A taxonomy for autonomous agents. In J. P. Müller, M. J. Wooldridge, and N. R. Jennings (Eds.), *Intelligent Agents III: Agent Theories, Architectures, and Languages* (pp. 21–35). Springer. <https://doi.org/10.1007/BFb0026906>
- Gonzalez, C., Donahue, K., Goldstein, D. G., Heidari, H., Jalali, M. S., Schelble, B., Singh, A., and Woolley, A. W. (2026). Toward a science of human–AI teaming for decision making: A complementarity framework. *PNAS Nexus*, 5(3), pgag030. <https://doi.org/10.1093/pnasnexus/pgag030>
- Kalluri, R. (2026). Implementing agentic AI in project management: An evidence-based framework and maturity model from systematic analysis. *International Journal of Managing Projects in Business*. <https://doi.org/10.1108/IJMPB-10-2025-0414>
- Kroll, J. A., Huey, J., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., and Yu, H. (2017). Accountable algorithms. *University of Pennsylvania Law Review*, 165(3), 633–705.
- Lee, J. D., and See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392
- Masters, C., Vellanki, A., Shangguan, J., Kultys, B., Gilmore, J., Moore, A., and Albrecht, S. V. (2025). Orchestrating human-AI teams: The Manager Agent as a unifying research

- challenge. arXiv:2510.02557. <https://arxiv.org/abs/2510.02557>
- Mathieu, J. E., Heffner, T. S., Goodwin, G. F., Salas, E., and Cannon-Bowers, J. A. (2000). The influence of shared mental models on team process and performance. *Journal of Applied Psychology*, 85(2), 273–283. <https://doi.org/10.1037/0021-9010.85.2.273>
- Montgomery, A. K., Gallagher, H. E., and Mercer, D. L. (2025). Human oversight for autonomous AI agents: A governance framework for healthcare, finance, energy, and public infrastructure. *International Journal of Engineering and Computational Applications*, 1(1), 18–32. <https://doi.org/10.54660/IJECA.2025.1.1.18-32>
- Müller, R., Zhai, L., and Wang, A. (2017). Governance and governmentality in projects: Profiles and relationships with success. *International Journal of Project Management*, 35(3), 378–392. <https://doi.org/10.1016/j.ijproman.2017.01.007>
- Müller, R., and Lecoeuvre, L. (2014). Operationalizing governance categories of projects. *International Journal of Project Management*, 32(8), 1346–1357. <https://doi.org/10.1016/j.ijproman.2014.04.005>
- Müller, R. (2011). Project governance. In P. W. G. Morris, P. Pinto, and J. K. Söderlund (Eds.), *The Oxford Handbook of Project Management* (pp. 297–320). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199563142.003.0013>
- Nenni, M. E., De Felice, F., De Luca, C., and Forcina, A. (2024). How artificial intelligence will transform project management in the age of digitization: A systematic literature review. *Management Review Quarterly*, 75, 1669–1716. <https://doi.org/10.1007/s11301-024-00418-z>
- O'Neill, T. A., McNeese, N. J., Barron, A., and Schelble, B. (2022). Human–autonomy teaming: A review and analysis of the empirical literature. *Human Factors*, 64(5), 904–938. <https://doi.org/10.1177/0018720820960865>
- O'Neill, T. A., Flathmann, C., McNeese, N. J., and Salas, E. (2023). Human-autonomy teaming: Need for a guiding team-based framework? *Computers in Human Behavior*, 146, 107762. <https://doi.org/10.1016/j.chb.2023.107762>
- Parasuraman, R., Sheridan, T. B., and Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans*, 30(3), 286–297. <https://doi.org/10.1109/3468.844354>
- Pathak, R., and Panteli, N. (2026). Governing the agents: Leadership dilemmas in deploying agentic AI in a cyber defence project team. *Information and Organization*. <https://doi.org/10.1177/20438869261467768>
- Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., and Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33–44. <https://doi.org/10.1145/3351095.3372873>
- Salimimoghadam, S., Ghanbaripour, A., Tumpa, R. J., Rahimi, A. K., Golmoradi, M., Rashidian, S., and Skitmore, M. (2025). The rise of artificial intelligence in project management: A systematic literature review of current opportunities, enablers, and barriers. *Buildings*, 15(7), 1130. <https://doi.org/10.3390/buildings15071130>
- Shenhar, A. J., and Dvir, D. (2007). Project management research—The challenge and opportunity. *Project Management Journal*, 38(2), 93–99. <https://doi.org/10.1177/875697280703800210>

- Shenhar, A. J., Dvir, D., Levy, O., and Maltz, A. C. (2001). Project success: A multidimensional strategic concept. *Long Range Planning*, 34(6), 699–725. [https://doi.org/10.1016/S0024-6301\(01\)00097-8](https://doi.org/10.1016/S0024-6301(01)00097-8)
- Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe and trustworthy. *International Journal of Human-Computer Interaction*, 36(6), 495–504. <https://doi.org/10.1080/10447318.2020.1741118>
- South, T., Marro, S., Hardjono, T., Mahari, R., Whitney, C. D., Chan, A., and Pentland, A. (2025). Position: AI agents need authenticated delegation. *Proceedings of the 42nd International Conference on Machine Learning*, 267, 82211–82231. <https://proceedings.mlr.press/v267/south25a.html>
- Srivastava, A. (2026). Governing AI-enabled decision making: Delegation, autonomy, and control at the operations–marketing interface. *Production and Operations Management*. Advance online publication. <https://doi.org/10.1177/10591478261473004>
- Tabassi, E. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
- Tatasciore, M., and Loft, S. (2025). Calibrating reliance on automated advice: Transparency and trust calibration feedback. *International Journal of Human-Computer Interaction*, 41(23), 14723–14733. <https://doi.org/10.1080/10447318.2025.2487861>
- Tian, Y., Antwi-Afari, M. F., Seo, J. O., Mi, H. Y., and Fadaie, S. (2026). A systematic literature review of human-autonomy teams in project management. *International Journal of Building Pathology and Adaptation*, 44(3), 789–808. <https://doi.org/10.1108/IJBPA-12-2024-0271>
- Ward, S., and Chapman, C. (2003). Transforming project risk management into project uncertainty management. *International Journal of Project Management*, 21(2), 97–105. [https://doi.org/10.1016/S0263-7863\(01\)00080-1](https://doi.org/10.1016/S0263-7863(01)00080-1)
- Wooldridge, M. J., and Jennings, N. R. (1995). Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2), 115–152. <https://doi.org/10.1017/S0269888900008122>
- Zamecnik, A., Febriantoro, W., Cukurova, M., Joksimović, S., Grossmann, G., and Siemens, G. (2026). Linking AI agent architectures to trust antecedents in human-autonomous teams: A scoping review. *Artificial Intelligence Review*, 59, 178. <https://doi.org/10.1007/s10462-026-11572-z>