

**INTERNATIONAL JOURNAL OF LAW
MANAGEMENT & HUMANITIES**
[ISSN 2581-5369]

Volume 9 | Issue 5

2026

© 2026 *International Journal of Law Management & Humanities*

Follow this and additional works at: <https://www.ijlmh.com/>

Under the aegis of VidhiAagaz – Inking Your Brain (<https://www.vidhiaagaz.com/>)

This article is brought to you for free and open access by the International Journal of Law Management & Humanities at VidhiAagaz. It has been accepted for inclusion in the International Journal of Law Management & Humanities after due review.

In case of **any suggestions or complaints**, kindly contact support@vidhiaagaz.com.

To submit your Manuscript for Publication in the **International Journal of Law Management & Humanities**, kindly email your Manuscript to submission@ijlmh.com.

Agentic AI in the Project Management Office: A Task- and Risk-Based Framework for AI Autonomy

MADHURI PARASHURAM KERUR*

ABSTRACT

Artificial intelligence is increasingly used in project management offices (PMOs) to support reporting, scheduling, risk identification, forecasting, and stakeholder communication. In most organisations, however, AI remains a support tool: it analyses information and provides recommendations, while people make the final decisions. Agentic AI may change this arrangement. Rather than only responding to instructions, an agentic system can interpret situations, determine responses, and act with limited human involvement. For PMOs, the question is therefore not simply whether to adopt AI, but how much authority it should be given. This paper develops a Task- and Risk-Based Agentic PMO Autonomy Framework to assess AI autonomy for individual PMO activities. The framework considers three factors: task complexity, the risk of autonomous execution, and the degree of human judgement required. These factors are used to recommend appropriate levels of AI autonomy and human oversight. Secondary data from PM Solutions' State of the PMO 2025 report identify areas where AI is currently involved in PMO work. Ten representative PMO activities, mapped against the PMBOK® process framework, are assessed using a researcher-developed rubric informed by the literature. An exploratory Spearman rank-correlation analysis examines whether greater complexity, risk, and human judgement requirements are associated with lower recommended AI autonomy. Across the ten activities, the results show the expected negative relationship. The study argues that the future PMO is unlikely to be fully human-run or fully autonomous. Instead, it is likely to operate as a hybrid environment in which people, digital systems, and AI agents share responsibility according to each activity's requirements. The practical contribution is to position AI autonomy as an activity-level governance decision rather than a single organisation-wide technology choice. The central question for PMO leaders is therefore not only what AI can do, but what it should be authorised to do.

Keywords: *Agentic AI, Project Management Office (PMO), AI Autonomy, Project Management, AI Governance, Human Oversight, Agentic Project Management, Agentic Business Process Management*

* Author is a Student at PES University, Bengaluru, Karnataka, India.

I. INTRODUCTION

A. Background

Project management offices (PMOs) play an important role in bringing governance, standardisation, coordination, and performance monitoring to an organisation's projects. At the same time, artificial intelligence is gradually changing how this work is carried out. Recent studies have highlighted AI applications in areas such as scheduling, risk management, forecasting, reporting, communication, and decision support (Taboada et al., 2023; Nenni et al., 2024; Khalil et al., 2025; Su & Ayob, 2025). At the PMO level, Kohli (2025) suggests that these capabilities could help PMOs move beyond routine administrative work towards a more analytical and strategic role.

Most of this activity, however, is still AI-assisted rather than autonomous. AI analyses information, identifies patterns, and offers recommendations, while a person retains the authority to decide what happens next. Agentic AI represents a further step. Rather than simply responding to a request, an agentic system can interpret a situation, work towards a defined objective, choose between possible actions, and carry out actions within agreed boundaries. When something falls outside those boundaries, it can escalate the matter to a human.

This distinction matters particularly for the PMO because it changes the question being asked. The issue is no longer simply whether AI can perform a particular activity; for a growing number of PMO activities, it increasingly can. The more important question is how much authority AI should be given to perform that activity on its own. PMO work is varied enough that the answer is unlikely to be the same for preparing a routine status report as for making a resource-allocation decision.

B. From AI-assisted to agentic PMO

The move towards an agentic PMO is likely to be gradual rather than sudden. In an AI-assisted PMO, a person identifies a task or asks a question, AI analyses the available information and provides a recommendation, and the person then decides what action to take.

In a more agentic PMO, the relationship changes. Humans and AI agents work towards a shared objective, but the AI agent is able to take a more active role. It can monitor relevant information, analyse a situation, decide on an action within pre-agreed limits, carry out that action, and escalate the matter when those limits are reached or when human judgement is needed.

AI-assisted PMO: Human → AI Analysis → Recommendation → Human Decision

Agentic PMO: Human + AI Agent → Sense → Analyse → Decide → Act → Escalate

Recent research has started to explore this shift. Assalaarachchi et al. (2026) propose an agentic approach to software project management in which AI agents can operate at different levels of autonomy depending on the complexity and risk of the task. The idea is not for AI to replace the human project manager, but to work alongside them. Research on agentic business process management makes a similar argument, emphasising that autonomous agents should operate within clearly defined and governed boundaries rather than being given unrestricted independence (Dumas et al., 2026; Calvanese et al., 2026).

Both streams are relevant to the PMO because PMO activities are not all alike. They differ in complexity, consequence, and the level of judgement they require. Preparing a routine status report is very different from allocating scarce resources across competing projects, escalating a critical issue, or supporting a portfolio-level prioritisation decision.

Because these activities are so different, this study argues that AI autonomy should be decided activity by activity, rather than by applying the same level of autonomy across the entire PMO.

C. Problem statement

The growing capabilities of AI create an important distinction between what AI can do and what AI should be allowed to do. Existing research has already explored AI applications in project management across areas such as analytics, forecasting, risk management, scheduling, reporting, and decision support (Taboada et al., 2023; Nenni et al., 2024; Khalil et al., 2025). More recent research has also started to examine agentic project management and autonomous project-related activities (Assalaarachchi et al., 2026). At the same time, research on agentic business process management has highlighted the importance of setting clear boundaries around autonomous action (Dumas et al., 2026; Calvanese et al., 2026).

What is still missing is a PMO-specific way of deciding how much authority AI should receive across the different activities that a PMO performs. Technical capability alone does not answer that question. An activity may be technically possible for AI to perform, but that does not automatically mean AI should be allowed to perform it independently.

The nature of the activity also matters. Its complexity, the possible consequences of an AI error, and the level of human judgement required can all influence how much authority should reasonably be delegated to AI.

This study addresses this gap through a Task- and Risk-Based Agentic PMO Autonomy Framework. The framework is built around three related but distinct ideas: AI capability, AI use, and AI authority. AI capability refers to what AI can technically do. AI use refers to where

an organisation currently applies AI. AI authority refers to what AI is actually permitted to decide or execute independently.

The question this paper ultimately addresses, then, is not simply where AI can be adopted within the PMO, but where it should be given the authority to act.

D. Research questions

1. How might the role and functioning of the PMO change as agentic AI takes on a more active role across different PMO activities?
2. What levels of AI autonomy, governance, and human oversight are appropriate for different PMO activities?
3. How might increasing AI autonomy change the roles of PMO directors and project managers?

E. Research objectives

1. To develop a Task- and Risk-Based Agentic PMO Autonomy Framework for determining appropriate levels of AI autonomy across PMO activities.
2. To assess representative PMO activities on task complexity, autonomy risk, human judgement, and recommended AI autonomy, using secondary data and a researcher-developed scoring approach.
3. To examine what increasing AI autonomy could mean for PMO governance, human oversight, PMO directors, and project managers.

F. Significance of the study

The main conceptual contribution of this study is the distinction between AI capability, AI use, and AI authority. An AI system may be technically capable of performing an activity, and an organisation may already use AI to support that activity, but this does not automatically mean that AI should be given the authority to make decisions or take action independently.

From a practical perspective, the proposed framework gives PMO leaders a structured way to think about where AI might operate with greater autonomy, where human approval should still be required, and where important decisions should remain human-led.

This distinction is becoming increasingly important as organisations move from AI tools that mainly analyse and recommend towards systems that can act within live project workflows. PMO leaders will therefore need to make deliberate decisions about which activities are suitable for greater autonomy, where AI should recommend rather than decide, which actions should require human approval, and when human intervention should always remain necessary.

The study therefore treats AI autonomy not as a single technology decision that applies across the entire organisation, but as a series of governance decisions to be made for individual PMO activities.

G. Scope and limitations

This study focuses on a selected set of PMO activities and their potential movement towards greater AI autonomy rather than examining the entire project life cycle. The broader range of PMO activities is informed by the PMO literature and mapped to relevant PMBOK® processes (Project Management Institute, 2017), from which ten representative activities are selected for detailed assessment.

The study relies mainly on secondary data, including the *State of the PMO 2025* report (PM Solutions, 2025). The scores assigned to the selected activities and the resulting autonomy recommendations are researcher-developed rather than independently validated. They are intended to demonstrate how the proposed framework can be applied in practice, rather than to serve as a fully validated measurement instrument.

The exploratory statistical analysis should also be interpreted with this limitation in mind. It is used to examine whether the framework produces results that are consistent with its underlying logic, rather than to prove a causal relationship or establish a universal level of AI autonomy that would apply across every organisation.

For this reason, the study should be viewed as an exploratory framework-building exercise. It provides a structured starting point for thinking about how AI authority could be distributed across different PMO activities. Future research can build on this work by testing the framework with PMO professionals and applying it across different organisations, industries, and project environments.

II. LITERATURE REVIEW

A. AI and the changing role of the PMO

The PMO has traditionally supported governance, standardisation, reporting, coordination, performance monitoring, and decision-making. Research on AI in project management identifies applications across risk management, scheduling, forecasting, reporting, communication, data analysis, and decision support (Taboada et al., 2023; Nenni et al., 2024; Khalil et al., 2025; Su & Ayob, 2025). At the PMO level specifically, Kohli (2025) links these capabilities to a shift from administrative activity towards more strategic and value-oriented

work, through improved forecasting, risk detection, resource management, and analytical support.

Despite these developments, most applications remain AI-assisted rather than autonomous: AI provides analysis or recommendations, while humans retain responsibility for decisions and actions. That leaves an open question for a PMO thinking ahead: what happens once AI can not only recommend an action, but also execute it?

B. From generative AI to agentic AI

Generative AI mostly produces content, analysis, summaries, or recommendations in response to human instructions. Agentic AI goes further, letting a system pursue a defined objective through a sequence of actions: it can interpret information, assess a situation, decide on a response, act, and escalate when required. For the PMO, this is the difference between a flow that runs from a human, to AI analysis, to a recommendation on which the human then decides, and one in which a human and an AI agent share an objective and the agent senses, analyses, decides, acts, and escalates within agreed limits. The change that matters here is not only what AI can produce, but what it is allowed to do; and greater autonomy does not so much remove humans from the loop as move their role towards defining objectives, boundaries, approvals, and exceptions.

C. Emerging research on agentic project management

Research on agentic AI in project management is still young. Assalaarachchi et al. (2026) propose an Agentic Project Manager built as a multi-agent system for software project management, working somewhat like a junior project manager alongside a human team, with autonomy levels that vary by task. Their work also raises questions of trust, accountability, ethics, and human supervision, and of how the project manager's own role might change as a result, including a shift towards more strategic leadership and the coaching of people and AI agents alike.

The key implication for this study is that AI autonomy should be task-dependent rather than uniform, a principle that Assalaarachchi et al. (2026) establish at the level of the individual project manager. The PMO, however, carries a broader organisational remit: governance, reporting, portfolio monitoring, resource coordination, risk management, and stakeholder communication across multiple projects at once. This study extends the task-level autonomy idea from the project-manager role to the level of the PMO activity itself.

D. Agentic business process management

Agentic Business Process Management (Agentic BPM) offers a broader lens on autonomous AI within organisational processes. Dumas et al. (2026) describe a shift from conventional automation towards agents that can perceive, reason, and act within defined process boundaries. Calvanese et al. (2026) frame this more formally around four capabilities: “framed autonomy” (autonomy that is constrained by explicit process frames and aligned with organisational goals), explainability, conversational actionability, and self-modification, that is, the capacity of agents to adapt their own behaviour over time.

This is directly relevant to a PMO, since much of what a PMO does can itself be understood as a set of connected processes, among them reporting, risk monitoring, issue management, resource planning, portfolio monitoring, and stakeholder communication. What the Agentic BPM literature does not settle, though, is how much autonomy should be assigned to any one of these activities specifically. That is the gap this study tries to close, by applying the same governance logic (bounded rather than open-ended autonomy) directly to the PMO context.

E. AI capability, AI use, and AI authority

A central distinction running through this study is between three related but non-equivalent ideas: AI capability, or what the technology can technically do; AI use, or what an organisation currently puts AI to work doing; and AI authority, or what AI is actually permitted to decide or execute. An AI system might well be capable of reallocating project resources on its own, while an organisation deliberately restricts it to recommending an allocation and requires a person to approve the final call. AI capability, in short, does not automatically translate into AI authority, and heavy AI adoption in a PMO does not by itself mean that the PMO has become highly autonomous.

F. Human judgement and AI governance

As autonomy increases, so does the importance of governance. PMO activities differ enormously in structure, uncertainty, consequence, and the contextual judgement they demand: producing a routine status report is a very different undertaking from allocating scarce resources across competing projects. An agentic PMO therefore needs clear rules covering what AI can do independently, what requires human approval, when AI must escalate, who remains accountable, and when a human override is required.

This is consistent with the wider Agentic BPM literature, which treats bounded, governed autonomy as a design requirement rather than an afterthought (Dumas et al., 2026; Calvanese

et al., 2026). Human oversight, in other words, should be designed into an agentic PMO from the outset, not bolted on once autonomy has already been granted.

G. Positioning the study

The literature reviewed so far provides a strong foundation, but each stream addresses a somewhat different piece of the puzzle. Table 1 sets out how this study sits relative to that existing work.

Research stream	Main focus	Gap this study addresses
AI in project management	Analytics, forecasting, risk, scheduling, decision support	Mainly AI assistance, not autonomous action
AI-enabled PMO research	AI support and PMO transformation	Limited focus on AI authority
Agentic project management	AI agents and task-level autonomy	Largely project-manager / software-project context
Agentic BPM	Autonomous agents within business processes	Limited PMO-specific application
This study	AI autonomy across PMO activities	Links complexity, risk, judgement, and authority at the activity level

Table 1: Positioning of this study against existing research streams

The contribution here is not a claim that agentic AI or AI autonomy is a new idea. It lies in bringing these research streams together in the PMO context, with a specific focus on how AI authority should be distributed across different PMO activities.

H. Research gap

Taken together, the literature shows AI steadily entering project and PMO work, with emerging agentic research pushing the conversation from AI assistance towards AI-supported action (Taboada et al., 2023; Nenni et al., 2024; Khalil et al., 2025; Kohli, 2025; Assalaarachchi et al., 2026), while Agentic BPM research keeps returning to the importance of boundaries, governance, and controlled autonomy (Dumas et al., 2026; Calvanese et al., 2026). What remains underdeveloped is a PMO-specific way of deciding how much autonomy is appropriate for a given activity: not simply where AI can be used in the PMO, but where it should be given the authority to act. That question sets up the methodology and framework developed in the next part of this paper.

III. RESEARCH METHODOLOGY

A. Research approach

This study follows an exploratory, framework-development approach. The purpose is not to measure how widely organisations have adopted AI or to suggest that PMOs are already

operating autonomously. Instead, the study focuses on a more specific question: how can an organisation decide how much authority an AI agent should be given for different PMO activities?

To explore this question, the study combines secondary evidence on the current use of AI in PMO work with a structured assessment of selected PMO activities. Each activity is examined in terms of its task complexity, the risk involved if AI were allowed to act independently, and the level of human judgement required. These factors are then considered together to recommend an appropriate level of AI autonomy and the corresponding degree of human oversight.

The methodology is therefore intended as a decision framework rather than a predictive model. It does not attempt to identify one universal level of AI autonomy that would be appropriate for every organisation, since PMOs differ in their structures, industries, governance arrangements, and risk environments. Instead, the framework provides a transparent and repeatable way of thinking through how AI authority could be distributed across different types of PMO activities.

The overall research process is as follows:

PMO Activities → PMBOK® Context → Current AI Involvement → Complexity → AI Autonomy Risk → Human Judgement → Recommended AI Autonomy → Human Oversight and Governance

B. Data source

The study draws primarily on secondary data from PM Solutions' *State of the PMO 2025* report (PM Solutions, 2025). The report provides useful evidence on the areas in which AI is already being used or explored within PMO-related work.

Importantly, the report is not used to measure AI autonomy directly. Its role in this study is to provide a picture of the current level and pattern of AI involvement across different PMO activities. This distinction is important because using AI for an activity does not necessarily mean that AI has the authority to make decisions or act independently.

For example, an organisation may use AI to identify potential risks or prepare project reports while still requiring a project manager or PMO leader to review the information and decide what action should be taken. In other words, AI involvement and AI authority are not the same thing.

The secondary data therefore provide the current context for the analysis, while the activity-level framework developed in this part of the paper is used to consider how much authority AI might appropriately receive in the future.

Where the *State of the PMO 2025* report does not provide an exact AI-use measure for one of the selected activities, the closest related measure is used as a proxy. These proxy measures are interpreted carefully and are not treated as direct evidence of AI autonomy.

C. Identification of PMO activities

A PMO can perform a wide range of activities, including governance, reporting, planning, monitoring, risk management, resource coordination, portfolio management, and stakeholder communication. Attempting to assess every possible PMO activity individually would make the study unnecessarily broad and difficult to apply.

For this reason, the study focuses on ten representative PMO activities covering the major areas of PMO work.

The activities were selected to provide variation in terms of:

- task structure and complexity;
- potential consequences of autonomous execution;
- current AI involvement;
- dependence on human judgement; and
- relevance to core PMO responsibilities.

The selected activities are:

1. Status and performance reporting
2. Schedule monitoring
3. Risk identification
4. Risk reporting and escalation
5. Issue monitoring and escalation
6. Portfolio monitoring and analysis
7. Resource planning
8. Resource allocation
9. Stakeholder communication
10. Data analysis and decision support

These activities were also considered in relation to relevant PMBOK® process areas (Project Management Institute, 2017). The PMBOK® mapping does not suggest that a PMO directly performs every project management process. Instead, it provides a recognised project-management structure for understanding where these activities sit within the broader project environment and where PMOs commonly provide governance, coordination, monitoring, support, or control.

The ten activities should therefore be understood as representative analytical cases, rather than a complete list of every activity performed by every PMO.

D. Framework dimensions

Each selected PMO activity is assessed on three main dimensions:

1. Task Complexity (C)
2. AI Autonomy Risk (R)
3. Human Judgement Requirement (H)

These three dimensions were chosen because they capture different aspects of the autonomy decision. An activity may be relatively simple from a technical perspective but still carry serious consequences if AI performs it incorrectly. Similarly, an activity may involve relatively little direct risk but still depend heavily on contextual understanding, stakeholder relationships, negotiation, or strategic judgement.

Looking at these dimensions together therefore provides a more balanced basis for deciding how much authority should be given to AI.

i. Task complexity

Task complexity refers to the extent to which an activity involves multiple variables, uncertainty, interdependencies, changing conditions, or situations that cannot easily be handled through standard rules.

Lower-complexity activities are generally more structured, repetitive, and predictable. They often involve relatively clear rules, predefined data, or standard procedures. Higher-complexity activities, in contrast, may involve competing priorities, multiple projects, changing organisational conditions, or situations where the correct response depends heavily on context.

Task complexity was assessed on a five-point ordinal scale, as set out in Table 2.

Score	Interpretation
1	Very low complexity

Score	Interpretation
2	Low complexity
3	Moderate complexity
4	High complexity
5	Very high complexity

Table 2: Task complexity scale

A higher complexity score does not mean that AI cannot support the activity. AI may still provide useful analysis, identify patterns, or recommend possible actions. The score simply suggests that greater caution may be needed before giving AI independent decision-making or execution authority.

ii. AI autonomy risk

AI autonomy risk refers to the potential risk created if an AI system performs an activity incorrectly, inappropriately, or without sufficient human intervention.

Risk is assessed using the established probability $\times$ impact approach:

$$\text{Risk (R)} = \text{Probability (P)} \times \text{Impact (I)}$$

Probability represents the likelihood that autonomous execution could result in an error, inappropriate action, or undesirable outcome. Impact represents the potential consequence if such an error occurs.

Both probability and impact are scored on a five-point scale (Table 3).

Score	Probability	Impact
1	Very low	Negligible
2	Low	Minor
3	Moderate	Moderate
4	High	Major
5	Very high	Severe

Table 3: Probability and impact scales

The resulting risk score ranges from 1 to 25.

For example, an activity with a probability score of 3 and an impact score of 4 receives a risk score of:

$$\mathbf{R = 3 \times 4 = 12}$$

This approach does not claim to estimate the exact statistical probability of an AI failure. The ordinal probability scale is used as a structured expert-assessment method for comparing the relative likelihood and consequences of autonomous errors across activities.

iii. Human judgement requirement

The third dimension considers the extent to which an activity depends on forms of judgement that may be difficult to standardise fully.

These can include:

- contextual understanding;
- interpersonal judgement;
- stakeholder relationships;
- ethical considerations;
- political or organisational factors;
- strategic trade-offs; and
- decisions made under ambiguity.

Activities requiring less human judgement are generally more structured and can be carried out using clearly defined rules, thresholds, or available data. Activities requiring greater judgement may involve competing priorities, negotiation, stakeholder sensitivity, ethical considerations, or a broader understanding of the organisation.

Human judgement is assessed on a five-point ordinal scale (Table 4).

Score	Interpretation
1	Very low requirement for human judgement
2	Low requirement
3	Moderate requirement
4	High requirement
5	Very high requirement

Table 4: Human judgement requirement scale

The inclusion of human judgement is particularly important to this framework because technical capability alone does not determine whether AI should receive authority. An AI system may be technically capable of producing a recommendation, but an activity involving sensitive stakeholder relationships or major strategic trade-offs may still require a human to make the final decision.

E. Scoring procedure

The scoring process follows four stages.

i. Stage 1: Define the PMO activity

Each activity is clearly defined to avoid assessing broad PMO functions as a single unit. For example, resource planning and resource allocation are treated separately because they involve different levels of authority and potential consequences.

ii. Stage 2: Score the activity dimensions

Each activity is assigned scores for:

- task complexity (C);
- probability of an undesirable outcome from autonomous execution (P);
- impact of such an outcome (I); and
- human judgement requirement (H).

The scoring is based on the definitions established in the framework and informed by the nature of the activity and the relevant literature.

iii. Stage 3: Calculate AI autonomy risk

Probability and impact are combined using $R = P \times I$. This produces an activity-level risk score.

iv. Stage 4: Recommend AI autonomy

Complexity, risk, and human judgement are then considered together to determine the recommended level of AI autonomy. Importantly, the autonomy recommendation is not based solely on current AI use. An activity may already have substantial AI involvement but still be unsuitable for high levels of autonomous authority because of its consequences or its dependence on human judgement.

F. Recommended AI autonomy levels

The framework uses five levels of recommended AI autonomy (Table 5).

Level	AI role	Description
Level 1	Human-led	Humans perform and control the activity directly.
Level 2	AI-assisted	AI provides information, analysis, or support, while humans make decisions and take action.
Level 3	AI-recommended	AI identifies options or recommends actions, but human approval is required before implementation.

Level	AI role	Description
Level 4	Human-approved autonomy	AI independently performs structured activities under approved conditions, within predefined rules, thresholds, and authority limits.
Level 5	Bounded autonomy	AI manages an activity with substantial independence but remains subject to defined authority boundaries, monitoring, escalation, and human override.

Table 5: Recommended AI autonomy levels

Level 5 does not represent unrestricted AI independence. Even at the highest autonomy level, AI operates within defined organisational boundaries. For example, an AI agent may be allowed to identify a schedule deviation, notify relevant stakeholders, and trigger an escalation when a predefined threshold is reached, yet still be prohibited from changing the approved project baseline or committing additional resources without human approval.

The purpose of the scale is therefore not to decide whether AI is simply “used” or “not used”. It distinguishes between different degrees of authority.

G. Research hypotheses

Based on the logic of the proposed framework, three directional hypotheses were developed (Table 6).

Hypothesis	Statement	Expected direction
H1	Higher task complexity is associated with lower recommended AI autonomy.	Negative
H2	Higher AI autonomy risk is associated with lower recommended AI autonomy.	Negative
H3	Greater human-judgement requirements are associated with lower recommended AI autonomy.	Negative

Table 6: Research hypotheses

These hypotheses are treated as analytical hypotheses rather than conventional hypotheses tested on independently collected respondent data. Their purpose is to examine whether the activity-level recommendations are directionally consistent with the logic of the proposed framework.

H. Exploratory Spearman rank-correlation analysis

An exploratory Spearman rank-correlation analysis is conducted across the ten selected activities to examine whether higher complexity, risk, and human judgement requirements are associated with lower recommended AI autonomy.

Spearman's rho is appropriate for this analysis because the assessment dimensions are ordinal or rank-based and the sample is small. Each of the ten PMO activities is treated as one observation. The analysis therefore uses ten activity-level observations rather than individual survey respondents.

The analysis examines the relationships between:

- Task Complexity and Recommended AI Autonomy;
- AI Autonomy Risk and Recommended AI Autonomy; and
- Human Judgement and Recommended AI Autonomy.

A negative correlation would indicate that activities with higher scores on these dimensions tend to receive lower autonomy recommendations.

It is important to be clear about what this analysis can and cannot demonstrate. Because the complexity, risk, human-judgement, and autonomy scores were developed by the researcher as part of the same framework, the analysis cannot establish causality or provide independent empirical validation of the framework. Instead, the Spearman analysis is used as an exploratory internal-consistency check. It examines whether the resulting activity-level recommendations behave in a direction consistent with the theoretical logic of the framework.

The hypotheses are therefore interpreted as consistent or inconsistent with the observed pattern, rather than as conclusively proven or disproven.

I. Human oversight and governance

AI autonomy cannot be separated from governance. Giving an AI agent permission to act changes not only how an activity is performed, but also who or what is authorised to make decisions and take action.

Human oversight is therefore incorporated directly into the framework rather than treated as something added after autonomy has been granted. The governance requirements considered include:

- clearly defined authority boundaries;
- permitted and prohibited AI actions;

- controlled access to data and systems;
- approval requirements;
- escalation thresholds;
- monitoring and audit trails;
- exception-handling procedures;
- human override mechanisms; and
- clear accountability for AI-supported decisions and actions.

The required level of oversight should increase with the authority granted. A structured activity may allow an AI agent to execute predefined actions under approved conditions. A higher-risk or judgement-intensive activity may remain at the AI-assisted or AI-recommended level.

The central governance principle of the framework is therefore that greater AI autonomy should require clearer governance, not less. Delegating an activity to an AI agent does not transfer organisational accountability to that agent. The organisation and its designated human decision-makers remain responsible for defining the agent's authority, monitoring its actions, and intervening when required.

J. Bringing the framework together

The methodology brings the different stages of the study together into a single decision process:

PMO Activity → PMBOK® Context → Current AI Involvement → Task Complexity → Risk ($P \times I$) → Human Judgement → Recommended AI Autonomy → Authority Boundary → Human Oversight → Governance and Review

The framework is designed to move the discussion beyond whether AI is technically capable of performing a PMO activity. Instead, it focuses on whether AI should be authorised to perform that activity, to what extent, and under what conditions. This distinction between AI capability and AI authority forms the central methodological principle of the study.

K. Methodological limitations

Several limitations should be considered when interpreting the findings.

First, the study does not use primary data from PMO professionals. The recommended autonomy levels therefore represent a conceptual, researcher-developed assessment rather than observed organisational practice.

Second, the *State of the PMO 2025* report does not provide an exact AI-use measure for every selected activity. Where necessary, related measures are used as proxies and are interpreted cautiously rather than as direct evidence.

Third, the complexity, probability, impact, and human-judgement scores are researcher-developed. Although explicit criteria are used to improve transparency and consistency, another researcher could reasonably assign somewhat different scores to certain activities.

Fourth, the Spearman analysis is based on only ten activity-level observations. The results should therefore not be interpreted as population-level evidence, causal relationships, or independent empirical validation of the framework. They are intended only as an exploratory assessment of whether the results are consistent with the framework's theoretical logic.

Finally, agentic AI capabilities, organisational practices, and governance expectations are evolving rapidly. The recommended autonomy levels should therefore be treated as indicative and reviewable rather than as fixed or permanent classifications. As AI reliability improves, governance controls mature, and organisations gain experience with agentic systems, the appropriate autonomy level for a particular PMO activity may also change.

IV. DATA ANALYSIS AND INTERPRETATION

A. Overview

This part of the paper presents the analysis in three stages. First, the *State of the PMO 2025* data are used to understand where AI is currently being used in PMO work. Second, the ten selected PMO activities are assessed in terms of task complexity, risk, and the level of human judgement involved. Finally, the resulting scores are examined using an exploratory Spearman rank-correlation analysis.

The purpose is not to suggest that PMOs have already become autonomous. Instead, the analysis looks at where AI is already supporting PMO work, which activities may be more suitable for greater autonomy in the future, and where human control is likely to remain important. In this sense, the analysis moves from the current use of AI towards a possible future distribution of work between people, AI agents, and existing digital systems.

B. Current AI involvement in PMO work

The *State of the PMO 2025* report shows that AI is already being used across several areas of PMO work. Around 37% of organisations with PMOs reported using AI-supported practices, with adoption higher among high-performing PMOs (PM Solutions, 2025).

On the report's scale of 1 (no extent) to 5 (very great extent), the highest level of AI involvement was reported in facilitating communications (3.4). This was followed by automating repetitive tasks (2.9) and providing data analysis to support decision-making (2.8). Strategic thinking received a rating of 2.7, while automating reporting scored 2.6. Identifying potential risks and optimising project scheduling both scored 2.5. Among the activities relevant to this study, portfolio analysis (2.3) and resource allocation (2.1) showed the lowest levels of AI involvement; resource allocation received the lowest rating of any activity in the survey (PM Solutions, 2025).

Overall, the pattern suggests that AI is currently being used most in information-heavy and communication-related work. However, these figures show AI involvement, not AI autonomy. A higher rating does not mean that AI is independently making decisions or taking action. Most of these activities are still likely to involve AI supporting human decision-making rather than replacing it.

C. Activity-level assessment

Each of the ten selected PMO activities was assessed on task complexity (C), the probability (P) and impact (I) of autonomous execution risk, and the level of human judgement required (H). These assessments were then used to determine a recommended level of AI autonomy. The results are presented in Table 7.

PMO activity	Current AI involvement†	C	P	I	Risk	H	Recommended autonomy
Status & performance reporting	2.6	2	2	2	4	2	Level 4
Schedule monitoring	2.5	2	2	3	6	2	Level 4
Risk identification	2.5	3	3	3	9	3	Level 3
Risk reporting & escalation	n/a	3	3	4	12	4	Level 3
Issue monitoring & escalation	n/a	3	3	4	12	4	Level 3
Portfolio monitoring & analysis	2.3	4	3	4	12	4	Level 3
Resource planning	n/a	3	3	4	12	3	Level 3

PMO activity	Current AI involvement†	C	P	I	Risk	H	Recommended autonomy
Resource allocation	2.1	4	3	5	15	4	Level 2
Stakeholder communication	3.4	3	2	4	8	4	Level 3
Data analysis & decision support	2.8	3	2	4	8	3	Level 3

Table 7: Activity-level assessment scores and recommended autonomy

Note. C = Task Complexity; P = Probability; I = Impact; H = Human Judgement; Risk = $P \times I$.

† Current AI involvement is the mean rating (1 = no extent, 5 = very great extent) for the closest corresponding activity in the *State of the PMO 2025* survey (PM Solutions, 2025); n/a = no corresponding survey item.

The scores cluster around Levels 2 to 4, with no selected activity reaching Level 1 or Level 5. That is a feature of this particular sample of ten activities, not a claim that every PMO activity everywhere must sit within that band.

D. Reading the individual activities

i. Status and performance reporting (Level 4)

Status and performance reporting shares the highest autonomy recommendation in the sample with schedule monitoring. Much of the information involved (milestones, costs, progress measures, risks, and deliverables) is already structured and predefined. An AI agent could therefore collect information, consolidate it, compare planned and actual performance, prepare reports, and distribute them according to predefined rules.

However, preparing a report is different from deciding what management should do in response to the information in that report. For this reason, the activity is suited to a high level of autonomy, but not to unrestricted decision-making. Human oversight remains important when the information requires interpretation or action.

ii. Schedule monitoring (Level 4)

Schedule monitoring is also relatively structured. Project schedules contain defined milestones, dependencies, dates, and baselines, making continuous AI-supported monitoring technically possible. An AI agent could identify delays, calculate deviations, flag affected dependencies, and escalate predefined exceptions.

The situation changes when the task moves from monitoring the schedule to changing it. AI may be able to identify that a delay has occurred, but changing an approved project baseline involves a different level of authority and should require human approval. This distinction is important throughout the study: autonomy depends on the specific activity being performed, rather than being applied automatically to an entire area of PMO responsibility.

iii. Risk identification (Level 3)

AI can be particularly useful in identifying potential risks. It can scan project data, historical patterns, dependencies, and performance indicators to identify risks that may not be immediately visible through manual review.

However, identifying a potential risk is only the first step. Someone still needs to judge whether the risk is genuinely significant, what the appropriate response should be, and whether it requires management attention. For this reason, AI is well suited to identifying, analysing, and prioritising potential risks, while the final response remains a human responsibility.

iv. Risk reporting and escalation (Level 3)

An AI agent could detect when a risk threshold has been crossed and prepare the information needed for escalation. However, the same risk can require different responses depending on the project, its importance, stakeholder expectations, contractual commitments, and previous management decisions.

These contextual factors are difficult to capture fully in predefined rules. AI can therefore support the escalation process by identifying the issue, preparing the relevant information, and recommending action. The decision about whether and how to escalate should, however, remain with a human.

v. Issue monitoring and escalation (Level 3)

A similar pattern applies to issue monitoring. AI can continuously monitor issue logs, identify overdue items, detect recurring problems, and flag issues when predefined thresholds are crossed.

The difficulty lies in deciding what the issue actually means in its wider context. An overdue technical issue may not require senior management attention if a recovery plan is already in place. At the same time, a seemingly small issue may be urgent if it affects an important project deliverable.

AI can identify and monitor the issue, but the decision to escalate often requires judgement that goes beyond the available data.

vi. Portfolio monitoring and analysis (Level 3)

Portfolio-level decisions involve a broader organisational perspective than decisions relating to a single project. AI can compare projects, identify trends, highlight underperformance, and identify potential conflicts involving resources or investment.

However, the strongest-performing project is not always the most strategically important one. Similarly, a project that appears to be underperforming may still be important to the organisation's long-term strategy.

AI can provide valuable analysis, but decisions about project priorities, continuation, or strategic importance require people who understand the wider organisational context.

vii. Resource planning (Level 3)

AI can support resource planning by analysing project requirements, workforce capacity, available skills, employee availability, and historical patterns. It can also identify possible shortages and suggest staffing options.

However, resource planning is not based entirely on quantitative information. Employee development, team relationships, individual circumstances, and management preferences may also influence the final decision. AI can therefore provide useful analysis and recommendations, while human decision-makers retain responsibility for the final choice.

viii. Resource allocation (Level 2)

Resource allocation provides the clearest example in the study of why AI capability should not automatically result in AI authority. This activity received the highest impact score (5), along with high complexity and human judgement scores (4 each), resulting in the highest overall risk score of 15.

Allocating people or teams to one project can affect several other projects at the same time. Decisions often involve competing priorities, limited specialist skills, deadlines, organisational objectives, and sometimes internal politics. These factors make resource allocation more than a simple matching exercise.

AI can analyse capacity and generate possible allocation options, including the likely trade-offs involved. However, because of the potential consequences of these decisions, the final allocation should remain under human control.

ix. Stakeholder communication (Level 3)

Stakeholder communication has the highest current AI involvement rating in the dataset (3.4). AI can already support many communication activities, including drafting routine updates, summarising discussions, personalising messages, and distributing scheduled communications. However, communication becomes more complex when it involves conflict, sensitive information, negotiation, senior stakeholders, or reputational risk. In these situations, the wording and timing of communication can have important consequences. AI can support the preparation of communication, but sensitive or high-consequence interactions should remain under human control.

x. Data analysis and decision support (Level 3)

AI can process large volumes of project information and identify patterns that may otherwise be difficult for managers to detect. This makes it particularly useful for analysing project performance and supporting decision-making.

However, analysis alone does not determine what action should be taken. Knowing that a project is underperforming does not automatically reveal whether the best response is to add resources, change the scope, revise the schedule, or take no immediate action. AI can analyse information and recommend possible responses, but the final decision requires human judgement.

E. The overall pattern

The overall results show a clear pattern. Status and performance reporting and schedule monitoring were both placed at Level 4 because they are relatively structured and carry lower consequences than the other activities.

Seven activities were placed at Level 3: risk identification, risk reporting and escalation, issue monitoring and escalation, portfolio monitoring and analysis, resource planning, stakeholder communication, and data analysis and decision support. These activities can benefit significantly from AI analysis and recommendations, but they still require human judgement before important action is taken.

Resource allocation was the only activity placed at Level 2. Its higher complexity, potential impact, and greater need for human judgement make stronger human involvement more appropriate.

The large number of Level 3 activities is itself meaningful. In the near term, the most realistic role for agentic AI in the PMO may not be completely independent execution. Instead, AI may act as a continuous analytical and recommendation layer, monitoring information, identifying

patterns, preparing possible actions, and highlighting exceptions, while people remain responsible for the final decision.

The findings also show that current AI use and recommended AI autonomy are not the same thing. Stakeholder communication, for example, has the highest level of current AI involvement but is still placed at Level 3 because many communication situations require contextual and interpersonal judgement.

Resource allocation shows the opposite pattern. Its lower current AI involvement does not necessarily mean that AI cannot support the activity; rather, the framework recommends caution because of the consequences associated with autonomous decisions.

F. Exploratory Spearman analysis and hypothesis results

Spearman's rank correlation was calculated across the ten activities using the finalised scores, with the results reported in Table 8.

Relationship	Spearman's ρ	p-value	Direction
Complexity → Recommended Autonomy	-0.885	0.001	Strong negative
Risk → Recommended Autonomy	-0.837	0.002	Strong negative
Human Judgement → Recommended Autonomy	-0.748	0.013	Strong negative

Table 8: Exploratory Spearman rank-correlation results ($N = 10$)

All three relationships are negative and statistically significant at the 5% level, which is consistent with the direction proposed in H1, H2, and H3.

The strongest relationship is between task complexity and recommended autonomy ($\rho = -0.885$), followed by risk ($\rho = -0.837$) and human judgement ($\rho = -0.748$). In practical terms, activities that score higher in complexity, risk, or human judgement tend to receive lower recommendations for AI autonomy.

These findings should, however, be interpreted with caution. The analysis rests on ten researcher-coded activities, not on independent respondent data such as the 134 valid survey responses behind the *State of the PMO 2025* report, and the recommended autonomy levels are conceptually derived from the same characteristics being examined. The Spearman analysis therefore does not provide independent empirical proof of the hypotheses.

Its value is instead exploratory. It shows that the pattern produced by the framework is internally consistent with the logic on which the framework is based.

G. Agentic PMO autonomy portfolio

The activity-level results can be translated into a portfolio view, shown in Table 9.

Autonomy level	PMO activities	Typical AI role
Level 4: Human-approved autonomy	Status & performance reporting; schedule monitoring	Execute defined activities under approved conditions
Level 3: AI-recommended	Risk, issues, portfolio, resource planning, communication, analysis	Analyse, identify, recommend, and prepare actions
Level 2: AI-assisted	Resource allocation	Generate options and decision support
Level 1: Human-led	Not represented in this sample	Human performs and controls the activity directly
Level 5: Bounded autonomy	Not assigned in this sample	A future possibility, once boundaries are strong enough

Table 9: Agentic PMO autonomy portfolio

The fact that none of the activities reached Level 5 should not necessarily be seen as a weakness of the framework. Level 5 involves allowing an AI agent to act independently within clearly defined rules, thresholds, permissions, and escalation boundaries.

The results suggest that the ten selected PMO activities still involve enough organisational context, uncertainty, or potential consequence to make this level of independence difficult to justify at present.

This does not mean that Level 5 will always remain unsuitable. Highly structured, repeatable, and reversible activities may eventually move towards bounded autonomy when decision rules are clear, data are reliable, consequences are limited, actions can be reversed, and continuous monitoring and human override are available.

Level 5 can therefore be seen as a future governance possibility rather than the inevitable final stage of AI adoption.

H. From AI-enhanced PMO to agentic PMO

Taken together, the findings suggest that moving towards an agentic PMO is not simply about increasing the amount of AI used. The more important change is how work, responsibility, and decision-making authority are shared between people, AI agents, and existing digital systems.

Traditional PMO: People → Systems → Reports → Decisions

AI-enhanced PMO: People + AI → Analysis → Recommendations → Human Decisions

Agentic PMO: People + AI Agents + Systems → Sense → Analyse → Decide within Boundaries → Act → Escalate

This transition is unlikely to happen at the same pace across every PMO activity. More structured work, such as reporting and monitoring, is likely to move towards greater autonomy first. Activities involving strategic decisions, resource trade-offs, sensitive stakeholder interactions, or significant consequences are likely to require stronger human involvement for longer.

I. Implications for PMO directors and project managers

For PMO directors, the role may gradually move away from directly managing every process and towards governing a broader environment involving people, AI agents, digital systems, and organisational controls.

This could mean deciding which activities can be delegated to AI, which decisions require approval, what an AI agent is allowed to do independently, what situations should trigger escalation, how AI actions are monitored, and who remains accountable when something goes wrong.

In this sense, the PMO may increasingly become responsible not only for governing projects, but also for governing how AI participates in project work.

For project managers, the findings do not suggest simple replacement or displacement. Instead, AI may reduce some of the time spent on repetitive reporting, monitoring, and information processing. This could allow project managers to focus more on areas where human judgement remains particularly important, such as strategic decision-making, stakeholder relationships, negotiation, conflict resolution, exception handling, ethical judgement, and dealing with ambiguity.

The role of the project manager may therefore become less focused on collecting and consolidating information and more focused on interpreting information, making decisions, and supervising AI-supported work.

J. Key findings

Four key findings emerge from this analysis.

First, AI is already becoming part of PMO work, although it is still mainly used as a support capability. Current involvement is strongest in communication, repetitive tasks, data analysis, reporting, risk identification, and scheduling.

Second, AI autonomy should not be treated as a single decision for the entire PMO. The ten activities examined in this study received different autonomy recommendations. Structured

activities such as reporting and schedule monitoring were more suitable for greater autonomy, while resource allocation required stronger human control.

Third, task complexity, autonomy risk, and human judgement were all negatively associated with recommended AI autonomy. The exploratory Spearman analysis showed a consistent pattern: as these factors increase, the recommended level of AI autonomy tends to decrease.

Finally, the future PMO is more likely to become a shared human–AI working environment than either a fully automated or a fully human function. AI agents may take on more responsibility for monitoring, analysing, identifying, and preparing actions, while people continue to provide judgement, accountability, strategic direction, and control over important decisions.

The question for the future PMO is not simply whether AI can perform an activity, but how much authority it should have to perform it.

V. CONCLUSION

A. Summary of the study

This study explored how the role of the PMO may change as AI moves beyond supporting tasks and begins taking a more active role in carrying them out. The main aim was to provide PMO leaders with a practical way to decide how much independence AI should have for different types of activity.

The proposed Task- and Risk-Based Agentic PMO Autonomy Framework approaches this question by considering three factors together: task complexity, the risk involved in allowing AI to act independently, and the level of human judgement required. Rather than assuming that AI autonomy should be applied in the same way across the entire PMO, the framework looks at each activity individually.

When the framework was applied to ten representative PMO activities, a clear pattern emerged. More structured and routine activities, such as reporting and schedule monitoring, were found to be more suitable for higher levels of AI autonomy. Seven activities, in areas such as risk, issue management, portfolio decisions, resource planning, communication, and analysis, were better suited to a recommend-and-decide arrangement, in which AI supports the process but humans retain final decision-making authority. Resource allocation required the strongest level of human control because of the potential consequences of getting the decision wrong.

The exploratory Spearman correlation analysis supported this overall pattern. Higher levels of task complexity, autonomy risk, and human judgement were generally associated with lower

recommended levels of AI autonomy. Given the small sample and the exploratory nature of the analysis, these results should not be seen as proving a causal relationship. Instead, they provide an additional check that the framework produces results consistent with its underlying logic.

Taken as a whole, the main contribution of the study is not a claim that every PMO activity has one universally correct autonomy score. Instead, it offers a structured and repeatable way of asking a more important question: how much authority should AI be given for a particular activity? The framework also keeps AI capability, AI use, and AI authority separate, recognising that the fact that AI can perform an activity does not necessarily mean it should be authorised to carry it out independently. Governance, therefore, needs to be considered from the beginning rather than added only after autonomy has already been given.

B. Theoretical and managerial implications

From a theoretical perspective, this study extends the discussion of agentic AI beyond individual project tasks and applies it to the wider range of activities performed by the PMO. It also highlights an important distinction between AI capability, AI use, and AI authority. A system may be technically capable of performing an activity, but that alone does not mean it should be allowed to make decisions or take action independently.

In this sense, the study builds on recent work on task-specific autonomy in agentic project management, particularly Assalaarachchi et al. (2026), and applies the same underlying idea at the level of the PMO activity. The study does not argue that AI should either replace humans or remain permanently limited to a supporting role. Instead, it suggests that the appropriate level of autonomy depends on the nature of the activity itself.

For organisations, the main implication is straightforward: AI autonomy should be assigned activity by activity, rather than granted to the PMO as one uniform level of authority. PMO leaders need to be clear about what AI can do independently, what requires human approval, when AI should escalate an issue, which decisions should remain human-led, and how AI actions will be monitored over time.

In practice, this could involve a step-by-step approach:

1. Identify and map the main activities performed by the PMO.
2. Assess each activity in terms of complexity, autonomy risk, and the level of human judgement required.
3. Assign an appropriate level of AI autonomy.
4. Define clear boundaries around what AI is authorised to do.

5. Establish approval, monitoring, and escalation mechanisms.
6. Review and adjust autonomy levels as the organisation gains more experience with AI.

The move towards agentic AI is also likely to reshape PMO roles rather than simply remove them. PMO directors may spend less time managing every process directly and more time governing a human–AI environment. Their role may increasingly involve setting authority boundaries, approval rules, escalation thresholds, and accountability mechanisms.

Project managers may also see their work change. If AI takes on more routine reporting, monitoring, and analytical activities, project managers may have more time to focus on leadership, stakeholder relationships, problem-solving, and decisions that require experience and contextual judgement.

Both roles are therefore likely to require new capabilities alongside traditional project management skills. These may include AI oversight, the ability to question and evaluate AI outputs, data interpretation, and the ability to work effectively alongside AI agents rather than simply treating them as tools to which work is delegated.

C. Limitations and future research

The findings of this study should be interpreted with several limitations in mind.

First, the study relies mainly on secondary data and does not capture the direct experiences of PMO professionals who are currently working with agentic AI. As a result, the framework reflects a structured assessment based on the literature and available evidence rather than direct practitioner validation.

Second, the scores used for complexity, autonomy risk, human judgement, and recommended autonomy were developed by the researcher. The scoring approach provides a transparent and consistent basis for comparing activities, but it is not an objective or independently validated measurement scale. Another researcher or PMO professional could reasonably assess some activities differently.

Third, the analysis covers only ten representative PMO activities. The Spearman correlation analysis is exploratory and is based on a small sample. Its purpose is therefore to check whether the results are broadly consistent with the framework's logic, rather than to provide independent empirical validation or establish a causal relationship.

These limitations also point to several opportunities for future research. The framework could be tested directly with PMO directors, project managers, and AI governance professionals

through surveys, interviews, or expert panels. Future studies could also examine whether appropriate levels of AI autonomy differ across industries, organisational cultures, and types of PMO.

Other factors could also be added to the framework. These might include data quality, regulatory requirements, explainability, accountability, and how easily a decision can be reversed if AI makes an error. Finally, longitudinal research following organisations that actually introduce AI agents into PMO work would provide valuable insight into how autonomy works in practice once it moves beyond a framework and into real organisational decision-making.

D. Final proposition

AI autonomy in the PMO should be decided at the activity level: structured and lower-risk activities can support greater autonomy, while activities involving higher complexity, risk, or human judgement require stronger human oversight.

On this basis, the future of the PMO is unlikely to be a choice between humans and AI. A more realistic future is one in which people and AI agents work alongside each other, with responsibilities divided according to the nature of the activity and the level of authority involved.

The challenge for the agentic PMO, therefore, is not simply to identify everything AI is capable of doing. The more important challenge is deciding what AI should actually be authorised to do.

How well a PMO answers that question, activity by activity, may ultimately matter more than how advanced or capable its AI tools become.

VI. REFERENCES

- Assalaarachchi, L. I., Masood, Z., Hoda, R., & Grundy, J. (2026). Toward agentic software project management: A vision and roadmap. In *Proceedings of the 2026 International Workshop on Agentic Engineering (AGENT '26)* (pp. 157–164). Association for Computing Machinery. <https://doi.org/10.1145/3786167.3788409>
- Calvanese, D., Casciani, A., De Giacomo, G., Dumas, M., Fournier, F., Kampik, T., La Malfa, E., Limonad, L., Marrella, A., Metzger, A., Montali, M., Amyot, D., Fettke, P., Polyvyanyy, A., Rinderle-Ma, S., Sardiña, S., Tax, N., & Weber, B. (2026). Agentic business process management: A research manifesto. *Information Systems*, 140, Article 102738. <https://doi.org/10.1016/j.is.2026.102738>

- Dumas, M., Milani, F., & Chapela-Campa, D. (2026). Agentic business process management systems. In I. van de Weerd, B. Estrada Torres, & H. van der Aa (Eds.), *Business process management workshops: BPM 2025 international workshops, Seville, Spain, revised selected papers* (pp. 3–14). Springer. https://doi.org/10.1007/978-3-032-13426-4_1
- Khalil, M., Bravo, A., Vieira, D., & de Carvalho, M. M. (2025). Mapping the AI landscape in project management context: A systematic literature review. *Systems*, 13(10), Article 913. <https://doi.org/10.3390/systems13100913>
- Kohli, R. (2025). The evolution of the PMO: Comparing traditional vs. AI-enhanced project management offices. *PM World Journal*, 14(12), 1–23. <https://pmworldjournal.com/article/the-evolution-of-the-pmo>
- Nenni, M. E., De Felice, F., De Luca, C., & Forcina, A. (2024). How artificial intelligence will transform project management in the age of digitization: A systematic literature review. *Management Review Quarterly*, 75(2), 1669–1716. <https://doi.org/10.1007/s11301-024-00418-z>
- PM Solutions. (2025). *The state of the PMO 2025 research report and data*. Project Management Solutions. https://www.pmsolutions.com/uploads/files/uploads/files/State_of_the_PMO_2025_Research_Report_and_Data.pdf
- Project Management Institute. (2017). *A guide to the project management body of knowledge (PMBOK® guide)* (6th ed.).
- Su, X., & Ayob, A. H. (2025). Artificial intelligence in project success: A systematic literature review. *Information*, 16(8), Article 682. <https://doi.org/10.3390/info16080682>
- Taboada, I., Daneshpajouh, A., Toledo, N., & de Vass, T. (2023). Artificial intelligence enabled project management: A systematic literature review. *Applied Sciences*, 13(8), Article 5014. <https://doi.org/10.3390/app13085014>
