

Integrating Artificial Intelligence Courses into Master's-Level Management Education at Engineering-Oriented Universities: Pathways and Practice

Xinna Zhao^{1,2*}, Chenyang Bao^{1,2}

¹*Beijing Institute of Petrochemical Technology, Beijing 102617, China*

²*Beijing Research Base for Modern Industrial New Area Development, Beijing 102617, China*

**corresponding author*

Keywords: Artificial intelligence; management education; master's programs; engineering-oriented universities; curriculum integration; responsible AI

Abstract: The expanding use of artificial intelligence (AI) in organizational decision-making creates new requirements for postgraduate management education. For engineering-oriented universities, the challenge is to connect technological expertise with managerial reasoning through a coherent educational framework. This paper examines how AI courses can be embedded across master's-level management programs, with particular attention to the Chinese higher education context. Using a design-oriented conceptual synthesis of research on AI literacy, constructive alignment, management, and responsible AI, the paper proposes a framework organized around five competencies: managerial problem formulation, data and model reasoning, AI-supported inquiry, organizational implementation, and responsible judgment. These competencies are developed through entry-level diagnosis, foundational courses, discipline-specific modules, integrative projects, and thesis supervision. A constructed supply chain management scenario illustrates how the framework could translate into teaching practice. The paper also proposes an evaluation strategy that distinguishes AI-assisted task performance from independent understanding and delayed transfer. Its principal contribution is a program-level approach that connects learning outcomes, authentic management tasks, individual evidence, and institutional coordination. The framework has not been empirically validated; its educational effectiveness and implementation costs require subsequent investigation. Nevertheless, it provides a practical foundation for developing AI-informed managerial judgment while preserving disciplinary depth and academic independence.

1. Introduction

Artificial intelligence is changing how organizations analyze information, coordinate activities, and support decisions. These developments raise an educational question: what should management

graduates understand and be able to do when AI becomes part of their professional environment? Workplace research suggests that the benefits of AI depend on users and tasks. For example, Brynjolfsson et al. [1] identified heterogeneous productivity effects in a customer-support setting, with less experienced workers benefiting more substantially than their highly experienced counterparts. Such findings support attention to human capability development, although they cannot establish the effectiveness of any particular university curriculum.

The issue is particularly relevant to engineering-oriented universities, understood here as institutions whose disciplinary strengths and institutional identities are substantially grounded in engineering and related technologies. Their management schools may have opportunities to collaborate with computing departments, engineering laboratories, and industrial partners. These opportunities can support educational activities involving production systems, logistics, technological innovation, and organizational transformation. However, access to such resources varies across institutions and does not automatically produce effective interdisciplinary education.

China's policy environment provides a clear rationale for curriculum reform. The 2026 "Artificial Intelligence + Education" Action Plan calls for AI-related foundational education, interdisciplinary courses, and the revision of existing disciplinary training programs. Translating these policy directions into management master's education requires decisions about learning outcomes, curriculum sequencing, faculty responsibilities, and assessment [2].

Introducing an isolated AI elective may provide initial exposure, but it leaves open how students will apply that knowledge in later coursework, research, and professional practice. Conversely, inserting AI tools into multiple courses without coordination may create repetition, inconsistent expectations, and fragmented learning. The central challenge is therefore program coherence.

This paper addresses three questions:

1. Which AI-related competencies should management master's programs develop?
2. How can these competencies be embedded across the educational process?
3. How can implementation and evaluation distinguish meaningful learning from improved tool-assisted output?

The paper develops a conceptual framework and an illustrative practice design. Its intended contribution is a systematic connection between AI education and the broader organization of postgraduate management training.

2. Literature and Theoretical Foundations

2.1. AI literacy as a multidimensional educational objective

AI literacy extends beyond familiarity with applications. Long and Magerko [3] describe competencies involving the critical evaluation of AI, communication and collaboration with AI, and its practical use. Ng et al. [4] similarly organize AI literacy around understanding, application, evaluation and creation, and ethical issues. Together, these perspectives provide a basis for learning objectives that combine knowledge, action, and judgment.

UNESCO's student competency framework connects a human-centered mindset, AI ethics, techniques and applications, and system design across progressive levels of understanding, application, and creation. Although developed for school curricula, its attention to progression and responsibility offers principles that can be adapted to postgraduate education. Such adaptation requires substantially greater attention to disciplinary reasoning and independent inquiry [5].

For management students, AI literacy should support the ability to translate organizational problems into appropriate analytical questions and interpret results within institutional constraints. Knowing how to operate a model is insufficient if the student cannot explain whether its output addresses the decision at hand.

2.2. Constructive alignment and program coherence

Constructive alignment connects intended learning outcomes with teaching activities and assessment. Its central implication is that students should practice the capabilities they are expected to demonstrate, and assessment should provide evidence of those capabilities [6].

This paper extends that principle across successive stages of a degree program. A competency introduced in a foundational course should reappear in disciplinary applications, integrated projects, and independent work. For example, understanding model uncertainty should develop from interpreting a simple prediction interval to explaining how uncertainty changes a procurement decision.

The proposed mechanism is longitudinal: explicit competencies guide authentic tasks; tasks generate individual evidence; feedback identifies weaknesses; and later tasks require students to apply their learning under greater complexity.

2.3. AI-informed managerial judgment

Management education must address the relationship between technical systems and organizational action. Raisch and Krakowski [7] argue that automation and augmentation are interdependent, creating tensions that organizations must manage over time. This perspective suggests that students should examine how AI redistributes work, authority, and responsibility.

Accordingly, this paper defines **AI-informed managerial judgment** as the ability to formulate organizational problems, evaluate AI-supported evidence, compare feasible actions, and accept responsibility for recommendations. This definition provides the organizing purpose of the proposed curriculum. Technical sophistication matters to the extent that it strengthens justified decisions and independent inquiry.

3. Research Approach and Scope

This study adopts a design-oriented conceptual approach. It synthesizes selected academic publications and official frameworks concerning AI literacy, educational alignment, organizational applications, assessment, and responsible AI. Sources were identified through targeted searches and assessed for relevance to curriculum design and the availability of verifiable publication information.

The approach is a focused narrative synthesis rather than a systematic review. It does not claim exhaustive coverage, estimate the prevalence of existing practices, or establish causal relationships. Policy documents provide contextual direction; conceptual publications support design principles; and empirical studies inform specific cautions about AI use.

The analysis proceeds through three steps. First, relevant educational and managerial concepts are translated into proposed graduate competencies. Second, these competencies are mapped onto curriculum stages, teaching activities, and assessment evidence. Third, a constructed teaching scenario examines how the framework could operate in practice.

No interviews, surveys, institutional interventions, or student performance measurements were conducted for this paper. The credit allocation, project sequence, and evaluation procedures are proposals requiring local adaptation and future testing.

The framework primarily addresses management master's programs situated within engineering-oriented universities. It recognizes differences between academic and professional degrees while retaining a shared expectation of methodological transparency, responsible conduct, and independent reasoning.

4. A Competency Framework for AI Integration

The proposed framework includes five connected dimensions, as summarized in Table 1. These dimensions should recur throughout the program, with increasing analytical independence and organizational complexity.

Table 1. Proposed competencies and corresponding evidence

Competency	Expected graduate capability	Illustrative assessment evidence
Managerial problem formulation	Define a decision, identify stakeholders, and determine whether AI is appropriate	Problem statement comparing AI and alternative approaches
Data and model reasoning	Evaluate data quality, model assumptions, uncertainty, and the limits of inference	Data audit, baseline comparison, and critical interpretation
AI-supported inquiry	Use AI tools transparently while verifying outputs and preserving traceability	Reproducible analysis and documented verification
Organizational implementation	Connect analytical results with workflows, resources, incentives, and accountability	Implementation proposal and stakeholder analysis
Responsible judgment	Identify potential harms and justify oversight, escalation, or non-use	Risk assessment and individual oral defense

Managerial problem formulation should precede tool selection. Students need to distinguish the organizational objective from the available dataset and from the model's prediction target. A model predicting customer dissatisfaction, for example, does not itself determine which service intervention should be adopted.

Data and model reasoning requires a clear distinction between prediction and causal explanation. A variable may improve predictive performance without representing a factor that managers can change to improve outcomes. Students should also recognize data leakage, inappropriate evaluation procedures, and uncertainty about generalization.

AI-supported inquiry includes coding assistance, information extraction, drafting, and other legitimate uses, but requires verification proportionate to the task. Organizational implementation introduces the practical conditions under which analytical recommendations can be acted upon.

Responsible judgment connects these capabilities. Students should be able to recommend a simpler method, further investigation, or non-deployment when evidence is inadequate. The educational objective is therefore justified choice, including the ability to recognize the limits of AI.

5. Pathways for Embedding AI Across the Program

5.1. Entry diagnosis and differentiated preparation

Students may enter management master's programs with different levels of quantitative preparation, programming experience, and professional knowledge. An initial diagnostic activity should identify these differences through practical tasks, such as interpreting a dataset, evaluating a graphical claim, and explaining the limitations of an AI-generated recommendation.

The diagnosis should guide support rather than function as an additional admissions barrier. Optional bridging modules can cover statistics, data handling, research methods, and basic programming. Students with stronger preparation may undertake extension activities.

Accessible interfaces and prepared notebooks can help students begin meaningful analysis, but should be accompanied by explanations of data transformations, assumptions, and evaluation choices. All students should reach a common minimum standard in interpretation and verification,

even when their implementation methods differ.

5.2. A coherent foundation within existing credit structures

A foundational sequence should establish shared concepts before students encounter specialized applications. An illustrative allocation is presented in Table 2.

Table 2. Illustrative eight-credit curriculum allocation

Component	Suggested credits	Main focus
AI Foundations for Management	2	Core concepts, capabilities, limitations, and problem formulation
Data Reasoning and Responsible Analytics	2	Data quality, evaluation, uncertainty, and responsible use
AI Applications in Management	2	Discipline-specific decision problems and analytical methods
Integrated Management Project	2	Problem investigation, implementation planning, and individual defense

This allocation is a design example, not a proposed national requirement. Institutions should adapt it to local credit conventions and degree regulations. Relevant existing methods courses, electives, or project modules may be redesigned to accommodate the content. Embedded activities in other courses should not be counted again as additional credits.

The foundation should cover predictive and generative AI while preserving distinctions between their purposes and evaluation requirements. Students should understand why an apparently fluent explanation may be unsupported and why a highly accurate predictive model may still be unsuitable for a particular decision.

5.3. Horizontal integration into management disciplines

Disciplinary integration should begin with management questions. In operations management, students might assess demand forecasts and their implications for inventory. In marketing, they could evaluate customer segmentation while examining sampling limitations. In organizational behavior, they might analyze how AI-supported monitoring changes discretion and employee trust.

Each module should connect three elements: a disciplinary concept, an analytical task, and a managerial decision. For example, a forecasting exercise should require students to explain how forecast errors affect service levels and inventory costs.

A shared curriculum map can identify where each competency is introduced, practiced, and assessed. This makes repetition purposeful: students revisit common principles in different settings while encountering progressively more demanding applications.

Coordination is essential. Course leaders should agree on terminology, prerequisite knowledge, and minimum expectations for verification. Without such coordination, students may receive conflicting instructions about acceptable AI use or repeatedly complete introductory exercises without advancing.

5.4. Vertical integration through projects and thesis supervision

In an illustrative three-year program, foundational learning could occur during the first semester, disciplinary applications during the second, and an integrated project during the third. Later

semesters would connect these capabilities with thesis development. Shorter programs can compress the sequence while preserving its progression.

Projects should require students to move from structured exercises toward ambiguous problems. Early assignments may specify the dataset and decision. Later work should require students to define the problem, evaluate available evidence, and justify their methodological choices.

Academic master's students should emphasize theoretically grounded questions, research design, reproducibility, and warranted inference. Professional master's students should emphasize organizational diagnosis, feasible action, implementation, and evaluation. Both pathways require analytical rigor; their distinction concerns the primary purpose of the work.

Thesis supervision should include explicit discussion of whether AI is methodologically appropriate. Students should not be required to use AI when another approach better serves their research question. Where AI is used, relevant tools, analytical decisions, verification procedures, and limitations should be documented according to institutional requirements.

5.5. Faculty collaboration and resource coordination

Implementation requires clear responsibility across departments. Management faculty should guide disciplinary interpretation, computing specialists should support technical validity, and educational specialists can assist with assessment design. External partners may contribute organizational context and feedback.

A program coordinator should maintain the competency map and organize periodic review of teaching materials and assessment standards. Joint teaching and supervision need recognition in workload allocation; otherwise, interdisciplinary collaboration may depend on unsustainable individual effort.

Institutions with limited resources can begin with one management specialization and a small set of reusable projects. Public datasets, synthetic teaching data, and campus-based problems can provide alternatives when enterprise data are unavailable.

Program budgets should account for technical support, faculty development, software access, and assessment time. Choosing a complex platform without providing support may create recurring burdens that weaken teaching continuity.

5.6. Responsible use and transparent assessment

UNESCO's guidance emphasizes human agency, privacy, equity, and the educational validation of generative AI. These principles support institutionally coordinated decisions about tool access and learning design [8].

Each assignment should explain which uses of AI are permitted and what evidence students must provide. The AI Assessment Scale offers a framework for relating permitted use to intended learning outcomes [9].

For the proposed curriculum, assessment should combine analytical products with selected process evidence and individual explanation. Students might submit a decision memo, a reproducible analysis, and a brief account of consequential AI assistance. Oral questioning can then examine their reasoning.

Documentation should remain proportionate. Requiring every interaction log would create unnecessary workload and could expose sensitive information. Recording important methodological choices, verification steps, and revisions provides a more focused basis for assessment.

6. Illustrative Practice: AI-Supported Supplier Delay Management

The following is a constructed teaching scenario, not a report of an implemented intervention. It illustrates an eight-week project for management students at an engineering-oriented university.

6.1. Problem definition and data preparation

Students investigate how a manufacturer could prioritize purchase orders for early review when supplier delays threaten production. The decision concerns the allocation of limited managerial attention, rather than prediction alone.

During the first two weeks, teams define the decision point, stakeholders, available actions, and evaluation criteria. They compare the proposed approach with an existing rule, such as reviewing orders from suppliers with recent delays.

During weeks three and four, students examine an authorized, de-identified dataset or clearly labeled synthetic data. They assess missing values, inconsistent records, and label quality. Predictors must be available when the decision would actually be made. Actual delivery information, for example, cannot be used to predict delay at order placement.

The dataset is divided chronologically so that later observations are reserved for evaluation. Students also consider whether the task concerns established suppliers or suppliers absent from the training data, since these imply different generalization challenges.

6.2. Analysis and managerial interpretation

During weeks five and six, teams compare a transparent statistical baseline with a more flexible model where appropriate. Evaluation includes predictive errors and their managerial implications.

Students specify assumptions about the costs of unnecessary review and missed delays, together with limits on staff capacity. They then examine how alternative decision thresholds change the number and type of orders selected for attention.

A crucial distinction concerns intervention effectiveness. Predicting a delay does not demonstrate that contacting the supplier will prevent it. Consequently, students should present cost estimates as scenario-dependent unless supported by evidence about the intervention itself. Historical predictive performance cannot establish realized organizational savings.

Generative AI may assist with code explanation or summarize synthetic supplier communications. Students must verify consequential outputs against the underlying records and identify unsupported statements.

6.3. Implementation and individual accountability

During the final two weeks, teams prepare a proposal describing who reviews alerts, how exceptions are handled, when predictions are overridden, and how performance would be monitored. They should also consider whether the system disadvantages smaller suppliers because of incomplete records or limited historical data.

The NIST AI Risk Management Framework [10] provides a useful reference for organizing discussion around governance, contextual understanding, measurement, and risk management. These functions are interconnected rather than a simple sequential checklist.

Deliverables include a problem statement, data audit, reproducible analysis, managerial recommendation, and limitations statement. Each student completes an individual defense and subsequently addresses an unfamiliar scenario with changed constraints.

The project is designed to connect technical analysis with organizational judgment. Whether it

achieves this purpose remains an empirical question.

7. Evaluating Educational Effectiveness

7.1. Separating performance from learning

Evaluation should distinguish what students accomplish with AI from what they understand independently. Evidence from high school mathematics shows why this distinction matters: Bastani et al. [11] found that improved performance during AI-supported practice could coexist with poorer subsequent independent performance under some conditions. The setting differs from postgraduate management education, so the findings support a design caution rather than a direct prediction.

Three outcomes should therefore receive priority: the quality of managerial problem formulation, the interpretation of analytical evidence, and transfer to an unfamiliar problem. AI-assisted product quality can be assessed alongside these outcomes.

7.2. A prospective evaluation design

An initial pilot could assess students before instruction, after the integrated project, and during a delayed transfer task. Parallel tasks should differ in surface details while requiring comparable reasoning.

A common rubric should evaluate whether students define a decision clearly, identify relevant assumptions, interpret uncertainty, and justify a recommendation. Multiple assessors should moderate a sample of responses, with identifying information and instructional condition concealed where feasible.

A stronger evaluation could compare integrated instruction with a standalone AI course while keeping instructional time and technical coverage comparable. Random assignment across multiple classes would strengthen causal inference where feasible. Otherwise, comparison groups should be examined for differences in prior preparation, professional experience, instructor effects, and previous AI exposure.

Pre-post improvement alone would not demonstrate that curriculum integration caused the change. Results should include uncertainty estimates, missing-data information, and variation across students. Small cohorts may support conclusions about feasibility while remaining insufficient for precise estimates of educational impact.

7.3. Implementation and continuous improvement

Learning measures should be supplemented by evidence about faculty workload, student workload, resource access, recurring costs, and project continuity. Interviews and reflective accounts can help explain why particular activities succeed or fail.

Research participation should be handled through appropriate institutional procedures and separated from grading decisions. Student feedback should inform revision without becoming a substitute for performance evidence.

Annual review can use anonymized competency profiles to identify weaknesses. If students construct models successfully but cannot explain implementation constraints, the appropriate response may involve stronger management discussion and feedback rather than more technical content.

8. Discussion and Research Propositions

The framework positions AI integration as a problem of educational coordination. Its contribution lies in connecting competencies, disciplinary applications, assessment evidence, and institutional responsibilities across the degree program.

Three propositions follow:

Proposition 1: Repeated alignment of AI competencies with management tasks across program stages will be associated with stronger transfer than equivalent instructional time concentrated in a standalone introductory course.

Proposition 2: The relationship between discipline-specific AI instruction and managerial judgment will depend partly on students' ability to connect analytical evidence with organizational assumptions and constraints.

Proposition 3: Assessment combining AI-assisted products, process evidence, and individual defense will predict independent performance more effectively than assessment based on products alone.

These propositions require testing. The third, in particular, concerns assessment validity and should be evaluated against independent or delayed performance rather than faculty preference.

Engineering-oriented universities may offer favorable conditions for implementation, but disciplinary proximity alone is insufficient. Shared standards, teaching incentives, and appropriately governed data access are necessary enabling conditions. Smaller institutions may require a narrower initial scope.

The framework also has limitations. Its narrative source selection may omit relevant approaches, and its proposals have not been tested with students or faculty. Credit allocations and project requirements may require substantial adjustment across institutions. Changes in AI capabilities may alter suitable teaching activities, although the emphasis on problem formulation, evidence, and responsibility should remain relevant.

Future research should compare implementation across institutional types and degree pathways. It should also investigate whether benefits are distributed equitably among students with different prior preparation.

9. Conclusion

Embedding AI courses into management master's education requires coordination across the full educational process. Foundational instruction should connect with disciplinary coursework, integrated projects, thesis supervision, and program evaluation.

This paper proposes five competencies and a longitudinal pathway for developing AI-informed managerial judgment. The illustrative supplier-delay project shows how students could connect model evaluation with organizational decisions while preserving individual accountability.

The framework offers a basis for curriculum development rather than evidence of demonstrated effectiveness. Subsequent research should examine independent understanding, transfer, implementation costs, and differences across learners. The central educational objective is to prepare graduates who can evaluate AI-supported evidence, justify its use, and take responsibility for the decisions that follow.

Funding

Graduate Education Teaching Reform Research and Practice Project of Beijing Institute of Petrochemical Technology (Project Title: Research on the Path and Practice of Embedding Artificial Intelligence Courses into the Cultivation System for Management Master's Students in

Engineering Universities; Project No.: YJ26-105)

References

- [1] Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- [2] Ministry of Education of the People's Republic of China, National Development and Reform Commission, Ministry of Industry and Information Technology, Ministry of Science and Technology, & National Data Administration. (2026, April 2). “Artificial intelligence + education” action plan [Policy document, in Chinese]. Official document
- [3] Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). Association for Computing Machinery. <https://doi.org/10.1145/3313831.3376727>
- [4] Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, 2, 100041. <https://doi.org/10.1016/j.caeai.2021.100041>
- [5] Miao, F., Shiohira, K., & Lao, N. (2024). AI competency framework for students. UNESCO. Publication
- [6] Biggs, J. (1996). Enhancing teaching through constructive alignment. *Higher Education*, 32(3), 347–364. <https://doi.org/10.1007/BF00138871>
- [7] Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210. <https://doi.org/10.5465/amr.2018.0072>
- [8] Miao, F., & Holmes, W. (2023). Guidance for generative AI in education and research. UNESCO. Publication
- [9] Perkins, M., Furze, L., Roe, J., & MacVaugh, J. (2024). The Artificial Intelligence Assessment Scale (AIAS): A framework for ethical integration of generative AI in educational assessment. *Journal of University Teaching and Learning Practice*, 21(6). <https://doi.org/10.53761/q3azde36>
- [10] Tabassi, E. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
- [11] Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakci, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122. <https://doi.org/10.1073/pnas.2422633122>