PROJECT MANAGEMENT · AI

Managing AI Projects with PMI-CPMAI

A Practitioner's Guide to Delivering AI Projects with PMI-CPMAI

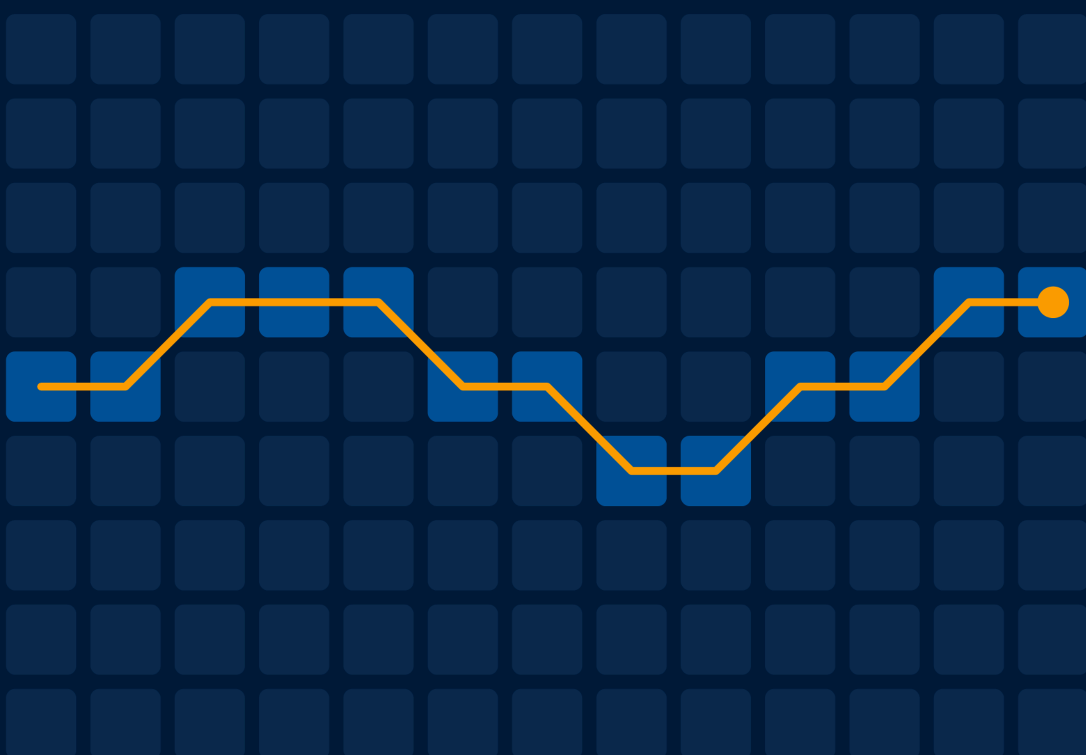

LAMJED CHIHAOUI

PMP · PMI-PMOCP · PMI-CPMAI

Managing AI Projects with PMI-CPMAI

A Practitioner's Guide to Delivering AI Projects with PMI-CPMAI

Lamjed CHIHAOUI

Copyright

This book is an original work. It explains, extends, and contextualizes the publicly known concepts and terminology of the CPMAI methodology and the PMI-CPMAI certification, and it does not reproduce the text, figures, examples, or exercises of any PMI course or publication. All examples, case studies, diagrams, tables, and templates in this book are original.

First edition, August 2026.

Preface

Most AI projects do not fail because the technology does not work. They fail because the business problem was never clear, the data was not ready, the right people were not engaged, or no one planned for life after deployment. These are project management problems, and they are solvable with the right method.

This book exists to give project managers that method. It teaches how to deliver AI initiatives using the CPMAI life cycle and the discipline of the PMI-CPMAI framework, and it prepares those who want it for the PMI-CPMAI certification. It is written from the vantage point of the person accountable for delivery, not the data scientist building the model. You do not need to become a data scientist to lead AI projects well. You need to manage business value, data, risk, stakeholders, and trust, and to know enough about how AI works to make good decisions.

This book has a second origin. Candidates preparing for the certification kept asking two questions: in what order should I learn this, and how does it map to the Examination Content Outline? An outline is a specification of what an exam assesses, not a study path. This book supplies the path, and Appendix A supplies the map, task by task.

The book is for project managers, PMO and program professionals, product managers, business analysts, digital transformation leaders, consultants, and technology managers who are taking on AI initiatives. Whether you are running your first AI pilot or your tenth, the aim is the same: to help you start from the business need, treat data as the foundation, iterate, keep the work trustworthy, and deliver real value.

A single fictional case, RetentionRadar, runs through the book so that each idea lands in a concrete setting. Every example, diagram, and template is original, and the focus throughout is practical. The goal is not to help you memorize a framework, but to help you use one.

Introduction

This book is organized into four parts that move from understanding to execution to responsibility, with an optional track for certification:

- Part 1, Foundations, sets the ground. It explains what makes AI projects different from traditional ones and introduces the CPMAI methodology at a high level, including its six phases and its iterative, data-centric character.
- Part 2, The CPMAI Lifecycle, is the heart of the book. It works through the six phases in order, one chapter each, from Business Understanding to Operationalization, carrying RetentionRadar through every phase and showing the inputs, outputs, deliverables, and decision gates of each.

- Part 3, Responsible AI, treats trustworthy AI as a pillar rather than an afterthought. It covers governance, privacy and security, fairness and bias, transparency and explainability, human oversight, and responsible generative AI, and it shows how these run through every phase of delivery.
- Part 4, the Certification Track, is optional. Readers who only want to manage AI projects can stop after Part 3. Those preparing for the PMI-CPMAI exam will find two chapters. The first describes the exam itself, its five domains and their weights, what it does and does not test, and what makes it difficult for each starting profile. The second covers how to read a question, how to prepare, how candidates lose this exam, and how to read the result.

Five appendices follow:

- an ECO Coverage Matrix that maps every exam task to the chapter that teaches it,
- a Deliverables Toolkit of ready-to-adapt templates,
- an AI Project Manager Toolkit of working aids,
- a glossary,
- a short resources guide.

To look something up quickly, start with the glossary in Appendix D for a term, or the ECO Coverage Matrix in Appendix A to find the chapter that teaches a given exam task.

A few conventions are worth noting:

- The book uses American English and refers to figures and tables by number.
- Each chapter in the main parts shares a common shape, with a PM Perspective box for each major concept, occasional PM Reality and Common Mistake notes, a consolidated look at RetentionRadar, and a closing sequence of a chapter summary and key takeaways.
- Study and exam material is gathered in Part 4 and in the ECO Coverage Matrix rather than repeated in every chapter, so the main chapters read as a clean practitioner guide.
- Official PMI-CPMAI terminology, such as the names of the six phases and the seven AI patterns, is used throughout, while a small number of industry terms outside the official vocabulary are flagged where they appear.

How you read the book is up to your goal:

- To build delivery capability, read Parts 1 to 3 in order and use the appendices as working tools.
- To prepare for the certification, read the whole book and lean on Part 4 and the ECO Coverage Matrix to focus your study.

Either way, the through line is the same: clear business framing, data that is fit for purpose, disciplined iteration, and trust built in from the start.

Contents

1. Introduction to AI Project Management

Learning objectives

- Explain what AI is, what it is good at, and how AI, machine learning, deep learning, and generative AI relate.
- Use the DIKUW pyramid to see where AI creates value and where human judgment stays in charge.
- Explain why most AI projects fail, and why traditional project management falls short for AI.
- Distinguish an AI project from the simple use of an AI tool, and tell the two sides of the coin apart.
- Adopt the mindset of an AI project manager and the common language that holds an AI team together.

Introduction

Most AI projects do not fail because the technology is weak. They fail because they are managed as if they were ordinary software projects. That single observation is the reason this book exists, and it is the reason a project manager, not only a data scientist, sits at the center of AI success.

This first chapter sets up the vocabulary and the mindset you will use for the rest of the book. It is deliberately high level. It explains what AI is and is not, where it creates value, why AI projects behave differently from traditional ones, and what the job of an AI project manager actually involves. It does not yet teach the methodology. PMI-CPMAI, the framework at the heart of the book, is introduced in the next chapter. Here, the goal is to make sure that when the methodology arrives, you already know why it is needed.

1.1 What AI is, and what it is good at

Why this matters. You cannot scope, budget, or de-risk a project around a capability you cannot describe. Before committing money and people, you need a plain answer to what AI can do well and where it struggles.

What to do. Judge a candidate use case against what AI genuinely does well, not against the technology's reputation. At its simplest, AI is the practice of building machines that perform tasks normally associated with human cognition. It is good at a few specific things:

- making predictions and recommendations more accurate and consistent,
- scaling repetitive or error-prone work,
- removing drudgery through automation so people are freed for judgment,
- providing continuous monitoring where round-the-clock attention matters.

Underneath these capabilities is a simple loop often summarized as the three Ps:

- perception, sensing and interpreting data,
- prediction, anticipating what may happen,
- planning, recommending or selecting the next action.

One property separates AI from traditional software and shapes every AI project. Conventional software follows fixed rules and returns the same output for the same input. AI systems are probabilistic. They produce results with a degree of likelihood rather than certainty, and the same model can be wrong on a case it has never seen before. This is not a defect, it is the nature of learning from data, and it explains three things you will meet again and again:

- first, AI work is measured with performance metrics rather than a simple pass or fail. Accuracy, precision, and recall are common examples for classification problems, but the relevant metric depends on the type of problem, for example F1-score for imbalanced classes, ROC-AUC for ranking quality, or MAE and RMSE for numeric prediction,
- second, reaching a good model takes iteration rather than a single build,
- third, because the output is never guaranteed, human oversight stays essential.

How to lead it. Qualify a candidate project by the capability it relies on:

- A fraud detection system earns its place because it monitors continuously and flags anomalies at a scale no team could match.
- A predictive maintenance system earns its place because it anticipates a failure before it happens.

- An e-commerce recommendation engine earns its place because it personalizes at scale.

In each case the value comes from a capability AI does well, expressed as a probability rather than a guarantee.

> **PM Perspective.** A business ambition is not yet a problem AI can solve, and turning one into the other is the project manager's first job. The risk is committing to a deterministic promise, that it will always be right, when the system is probabilistic by nature. Decisions affected: whether the use case fits AI at all, and how success will be measured. Stakeholders: sponsor, business owner, and the data and AI team. Because the output is probabilistic, the success criteria must be expressed as measurable performance and acceptable error, not as perfection.

1.2 The AI family, nested not interchangeable

Why this matters. Treating AI, machine learning, deep learning, and generative AI as the same thing leads teams to overspend, reaching for the most powerful and most expensive option when a simpler one would do.

What to do. Place the problem at the right level of a nested stack rather than reaching for the most advanced tool. The four terms nest inside one another rather than competing:

- AI is the broadest idea, any technique that makes a machine behave intelligently,
- machine learning is a subset of AI in which systems learn patterns from data rather than following hand-coded rules,
- deep learning is a subset of machine learning that uses many-layered neural networks, and it is what made modern progress on unstructured data such as images, audio, and text possible,
- modern generative AI relies on deep learning architectures and large foundation models to produce new content such as text and images.

Figure 1.1 shows the nesting.

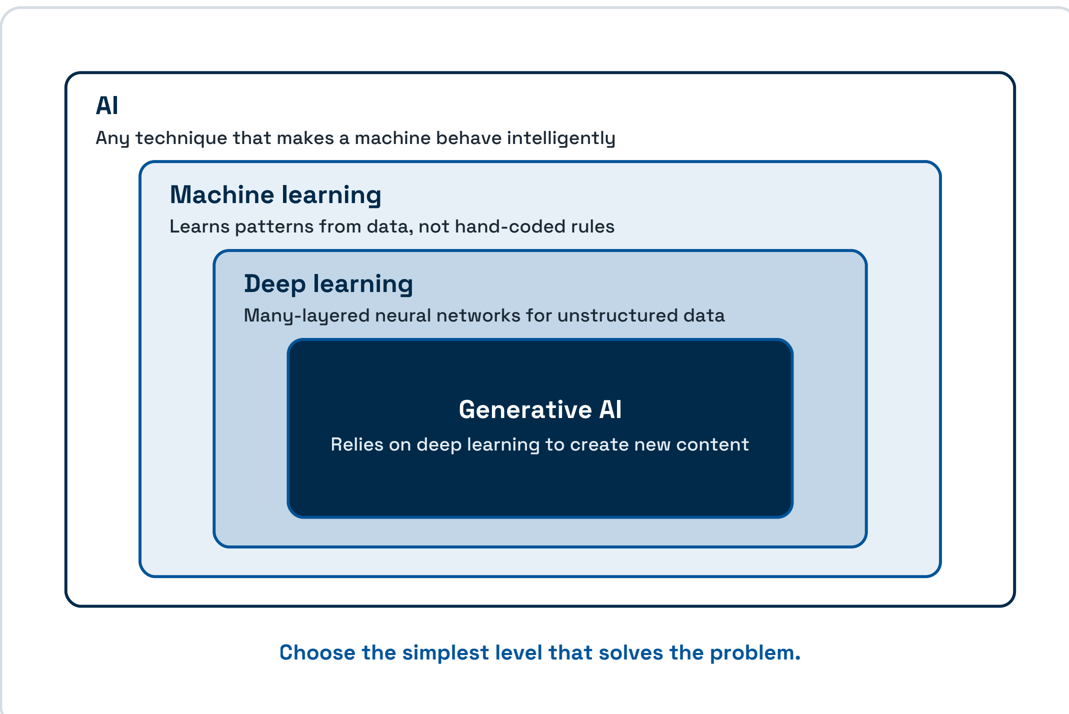

Figure 1.1. AI contains machine learning, which contains deep learning, which modern generative AI relies on.

How to lead it. Choose the simplest level of the stack that solves the problem:

- An autonomous driving system needs deep learning because it interprets raw sensor and image data.
- A legal assistant that drafts and summarizes documents needs generative AI because it produces language.
- A churn prediction on tabular customer data needs only a machine learning classifier, and reaching for generative AI there would add cost and risk without adding value.

Common Mistake. Reaching for generative AI or deep learning because they are impressive, when a simple machine learning model would meet the business need at a fraction of the cost and with far easier explainability. The sophistication of the tool is not the measure of a good solution. The fit to the problem is.

PM Perspective. Overbuilding is a budget decision disguised as a technical one, and the project manager governs it. The risk is scope and budget inflation driven by technology fashion rather than need. Decisions affected: which level of the stack to use, and therefore cost, timeline, skills, and explainability. Stakeholders: data and AI team, finance, and the sponsor. The cheapest option that meets the requirement is usually the right one.

1.3 From data to value, the DIKUW pyramid

Why this matters. A project manager needs a shared map of what AI can and cannot do, so the team and the sponsor calibrate their expectations to the same reality.

What to do. Use the DIKUW pyramid to locate the expected deliverable before promising it. The pyramid describes how value rises as you move from raw data toward wisdom:

- at the base is data, raw facts that on their own do little,
- organize it and you get information, the who, what, when, and where, produced with basic analytics and reporting, where AI is not required,
- above that is knowledge, understanding the how, which is the level machine learning serves best, detecting patterns and forecasting outcomes,
- higher still is understanding, which calls for deeper reasoning and common sense, something AI is still developing,
- at the top is wisdom, the seasoned judgment no data set confers, which people still supply better than machines.

Figure 1.2 shows the pyramid.

A clarification matters here. The DIKUW pyramid is a conceptual model used in the CPMAI material to show where AI adds value. It is not one of the six phases of the methodology, but the term appears in the official CPMAI glossary, so it is worth knowing.

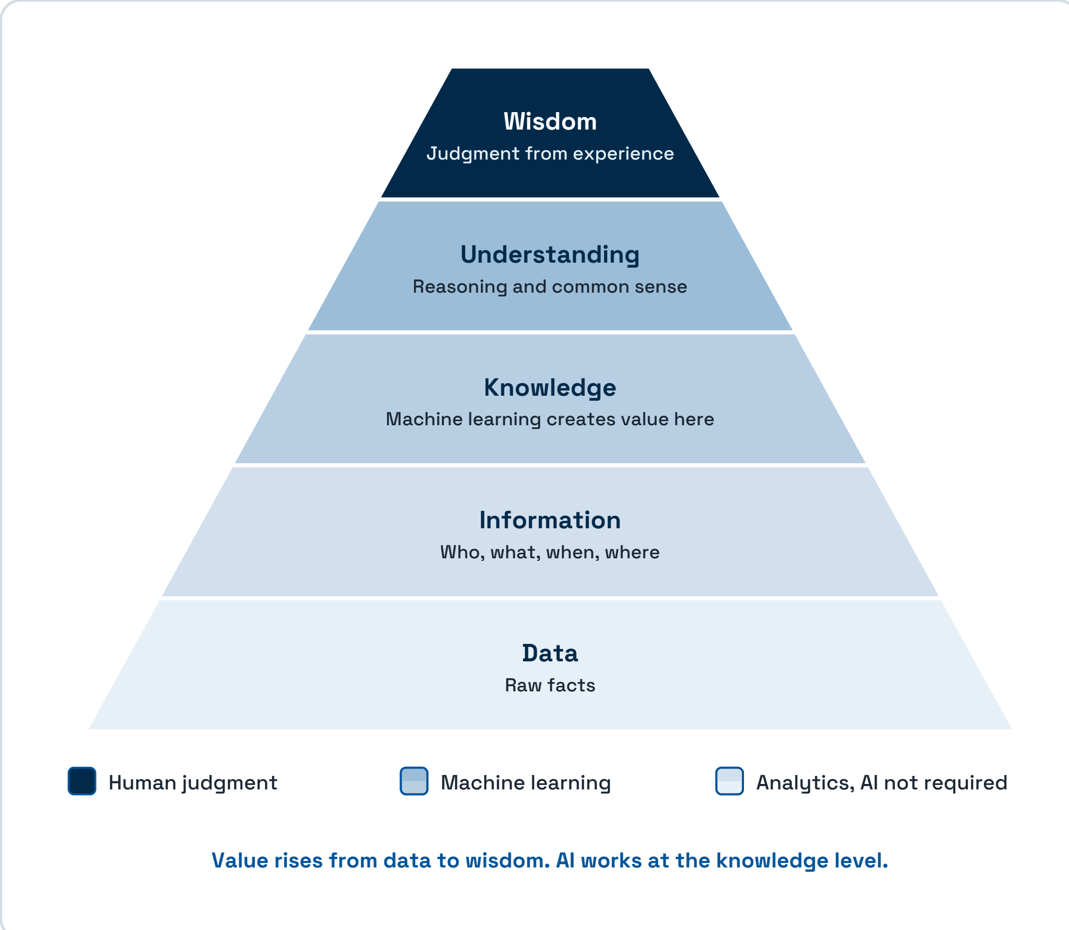

Figure 1.2. The DIKUW pyramid, where machine learning creates value and where humans stay in charge.

How to lead it. Place the expected deliverable on the pyramid before promising it. In a healthcare diagnosis support tool, the model operates at the knowledge level, surfacing patterns and likelihoods, while the understanding and the final clinical decision stay with the physician. Naming the level early prevents the common trap of promising machine wisdom when the system can only provide machine knowledge.

> **PM Perspective.** Name the level of the pyramid before the sponsor names it for you. The risk is a sponsor who expects judgment and wisdom from a system that can only deliver predictive knowledge. Decisions affected: what the system will and will not decide, and where a human stays in the loop. Stakeholders: sponsor, business owner, and end users. Calibrated expectations at the start prevent disappointment at the end.

1.4 Why AI projects fail, and why traditional management falls short

Why this matters. This is the reason the rest of the book exists. If you understand why AI projects fail, you understand what a good AI project manager must do differently. Industry studies consistently report that a large majority of AI initiatives fail to deliver their expected business value, with estimates often around 80 percent, and many never move beyond a prototype. The important insight is that these failures are usually not caused by the AI technology itself. They are caused by how projects are planned, managed, and aligned to a real business need. Three culprits come up again and again:

- a lack of clear business alignment,
- objectives that drift apart between the business and the technical team,
- inadequate data preparation.

Figure 1.3 makes the point that the technology is rarely the villain.

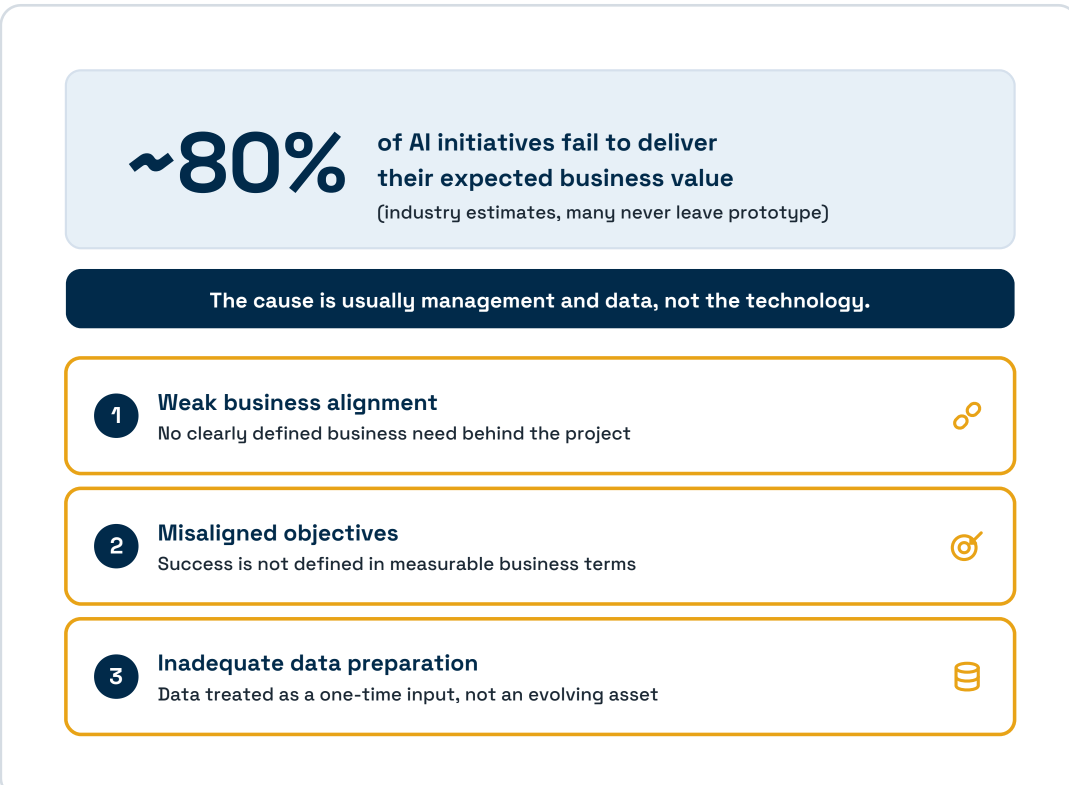

Figure 1.3. Most AI projects fail, and rarely because of the technology.

What to do. Manage the two things that actually decide the outcome, the business framing and the data, and abandon the assumptions of a build-and-ship project. In AI projects, data is not just an input, it is a project asset that evolves throughout the project life cycle. It is collected, cleaned, labeled, validated, and refreshed, and it keeps changing after the model is live. Treating data as a one-time input is one of the surest ways to join the failure statistic.

Traditional approaches rest on assumptions that do not hold for AI. They assume a stable set of requirements, but AI projects require iteration on the data itself, collecting, cleaning, labeling, and retraining, embedded at every phase rather than handled once. They also assume a deployment milestone after which the product is complete, whereas in AI deployment marks the transition to continuous monitoring, evaluation, and improvement that watches for model drift and performance decay. Table 1.1 contrasts the two.

Table 1.1. Traditional project management versus an AI-specific approach.

Assumption in traditional projects	What happens in AI projects	Implication for the PM
Requirements are stable once agreed	Data and models are refined across iterations	Plan for iteration, not a single pass
Data is an input gathered up front	Data is an evolving project asset	Budget continuous data work and validation
Delivery ends at deployment	Deployment marks the transition to continuous monitoring	Plan for drift detection and retraining
Output is deterministic and testable	Output is probabilistic and measured	Define success as performance and acceptable error

How to lead it. Plan an AI initiative from the outset as an iterative, data-driven loop rather than a build-and-ship project. Three moves carry that plan:

- anchor the project to a quantifiable business goal,
- give the data first-class attention and a real budget,
- treat deployment as the beginning of an operating capability, not the end of the work.

> **PM Reality.** The sponsor believes that buying the right AI tool will solve the problem, and that the project is mostly a technical build. The reality the team discovers is that the hardest and most decisive work is the business framing and the data, and that the chosen tool is rarely what makes or breaks the outcome. The project manager who sets that expectation early protects the project from the most common cause of failure.

> **PM Perspective.** The shift from a ship-once mindset to an iterative, data-centric one belongs to the project manager, not to the data team. The risk is running an AI project on traditional assumptions, which is how most of them fail. Decisions affected: the project plan, the budget for ongoing data and monitoring work, and the definition of done. Stakeholders: sponsor, finance, data and AI team, and operations.

1.5 The AI project manager, role, mindset, and common language

Why this matters. If the biggest success factor is how a project is managed, then the AI project manager is central, and the role is distinct from running a traditional project.

What to do. Keep two things apart that are easily confused, and hold the line on what an AI project actually is. There are two different things people mean by AI and project management:

- One is using AI tools to help manage projects, for example asking a model to draft a risk register.
- The other is managing the projects that build and deploy AI systems.

This book is about the second. Building and deploying an AI system, with its data work, training, evaluation, monitoring, and governance, is an AI project. Prompting a generative tool to produce a document, useful as it is, is not an AI project. Table 1.2 draws the line, because mistaking tool use for an AI project is a frequent source of confused scope and budget.

Table 1.2. What is an AI project versus what is not.

It is an AI project when	It is not an AI project when
You build, train, and deploy a model on your data	You type prompts into an existing tool for a one-off output
You run an ongoing loop of monitoring and retraining	You use an AI feature inside software someone else operates
You own the data, the model, and its governance	You consume a finished AI service with no model to manage

How to lead it. Set the mindset and the shared language that hold a cross-functional team together. Managing AI projects rewards a particular posture, and this book treats it as its reference, set once here and only recalled later. It rests on four attitudes shown in Figure 1.4:

- curiosity, to keep asking what the data and the model are really doing,

- adaptability, to adjust as evidence arrives,
- continuous learning, to keep pace with a fast-moving field,
- experimentation, to test small and learn quickly rather than betting everything on one design.

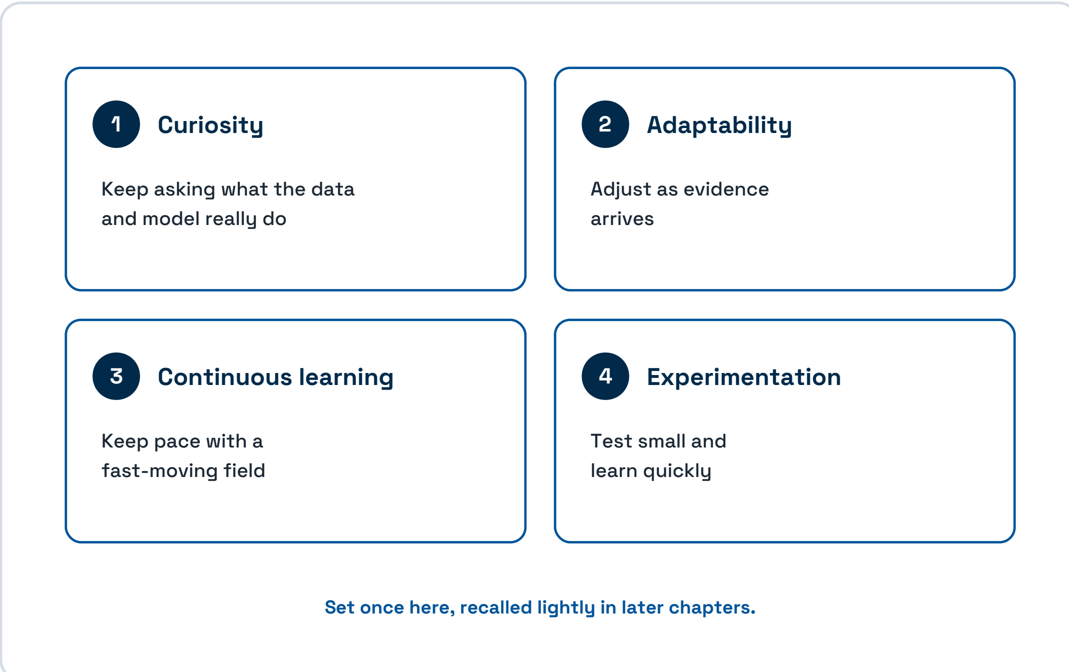

Figure 1.4. The AI PM mindset, four attitudes that AI projects reward.

CPMAI gives business, data, technology, and project teams a shared vocabulary, and that shared language is one of its quiet but decisive benefits. It goes further than words. CPMAI also provides a common structure for governance, collaboration, and decision making across the business, data, AI, and project teams, so that the groups work from the same process rather than improvising their own. Without it, teams can stall for months because each group means something different by the same word. A few messages run through the whole book and are worth fixing now:

- an AI project is not the same as the use of an AI tool,
- the data is the heart of an AI system,
- the framework is stable while the technologies evolve,
- AI is best used to augment human judgment rather than replace it,
- a common language keeps a cross-functional team aligned,
- the goal is to solve the problem rather than to chase the newest tool.

PM Perspective. A shared vocabulary is infrastructure, not etiquette, and the project manager maintains it. The risk is a team that fragments because business, data, and technology talk past each other, or that drifts toward the newest tool instead of the business goal. Decisions affected: how the team communicates, how scope is held to the problem, and how human oversight is preserved. Stakeholders: every group on the project, from the sponsor to the end users.

RetentionRadar in practice

Throughout the book a single running case makes the methodology concrete. RetentionRadar is an illustrative, clearly fictional churn-prediction project for a subscription streaming company, invented for teaching and carried through every phase so you can watch one project move from a vague ambition to a governed model in production.

> **RetentionRadar Example.** RetentionRadar is a small, focused intelligent system. It perceives subscriber behavior, predicts who is likely to cancel, and supports a human in planning the next action, the three Ps in miniature. On the DIKUW pyramid it lives at the knowledge level, turning organized subscriber information into predictive knowledge about churn, while the judgment about how to treat each customer stays with the retention team. It is a machine learning classification problem on tabular data, so it needs neither deep learning nor generative AI for its core prediction. And it is a reminder of the failure lesson: the hardest parts will be framing the business goal and getting the data right, not building the model. The case stays supervised, scored in batches, with a human in the loop and no offer ever applied automatically. After deployment, the model's performance will be monitored continuously to detect any degradation, or model drift, and to decide whether retraining is needed, a thread the later phases pick up in full.

Failure patterns

Even at this foundational level, the ways AI projects go wrong are recognizable. Each has a short fix.

- Building without business alignment. The team starts from a technology rather than a defined business need. Fix: anchor the project to a quantifiable business goal before any model work.
- Underestimating data preparation. The plan treats data as a quick input rather than an evolving asset. Fix: budget continuous data collection, cleaning, labeling, and validation.
- Planning a one-time build. The project ends at deployment with no monitoring. Fix: plan an iterative loop with drift detection and retraining from the start.
- Chasing tools. The team adopts the newest model because it is new. Fix: choose the simplest option that meets the need, and keep the focus on the problem.

PM challenge questions

- Is this a real AI project, or is it the one-off use of an AI tool.
- Does the use case fit something AI is genuinely good at, and can success be expressed as measurable performance.
- What is the simplest level of the stack that would meet the business need.
- Have we planned for data as an evolving asset and for monitoring after deployment, not just a single build.
- Who keeps a human in the loop, given that the system's output is probabilistic.

Chapter at a glance

Figure 1.5 summarizes the chapter on one page, from what AI is good at through the failure lesson to the AI PM mindset.

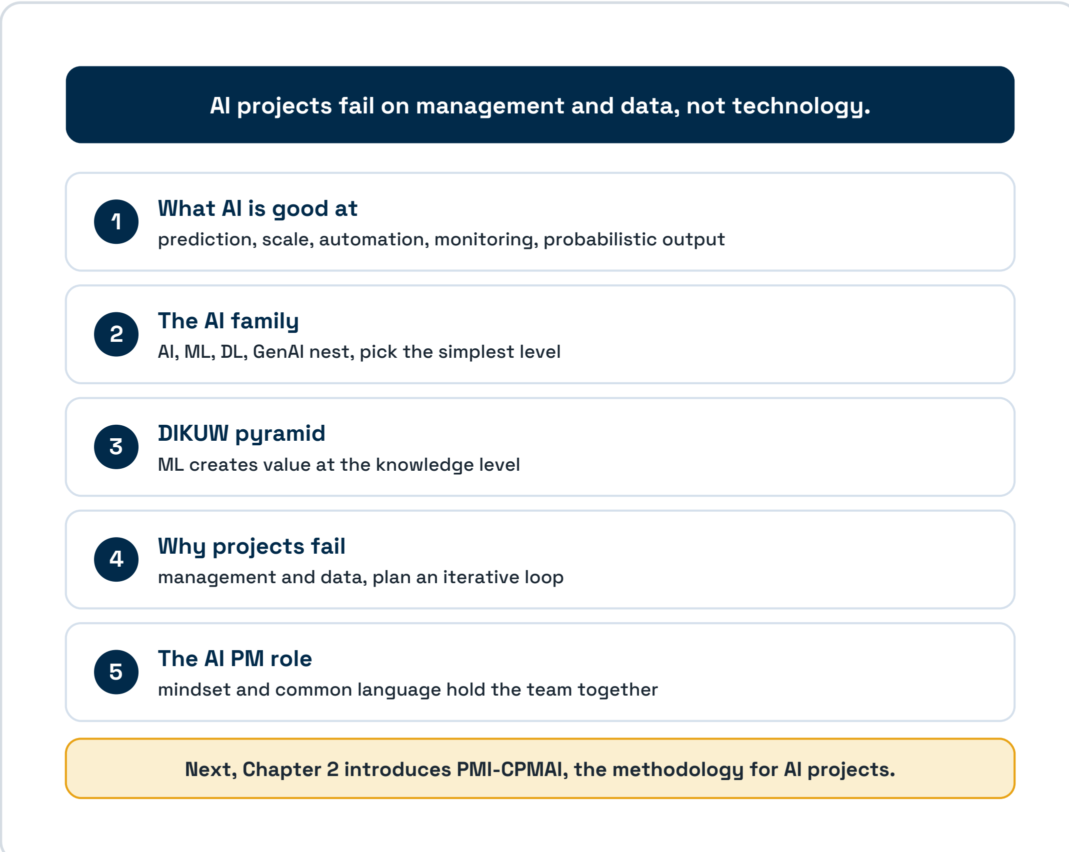

Figure 1.5. Chapter 1 at a glance.

Key takeaways

- AI builds intelligent machines that are good at prediction, scale, automation, and continuous monitoring, and their output is probabilistic, not deterministic.
- AI, machine learning, deep learning, and generative AI nest inside one another, so choose the simplest level that solves the problem.
- The DIKUW pyramid shows that machine learning creates value at the knowledge level, while understanding and wisdom stay human.
- Most AI initiatives fail because of business alignment, data quality, governance, and project management challenges rather than the AI technology itself, and traditional project management falls short because AI is iterative and data-centric and does not end at deployment.

- Running an AI project is different from using an AI tool, and the AI PM mindset of curiosity, adaptability, continuous learning, and experimentation, together with a common language, is what holds the work together.

ECO mapping

This is an orientation chapter and claims no discrete ECO task. It establishes the foundations that the later, task-bearing chapters build on. How this chapter maps to the Examination Content Outline is recorded in Appendix A, the ECO Coverage Matrix.

Looking ahead

If AI projects are fundamentally different, they require a different project management approach. The next chapter introduces PMI-CPMAI, the methodology designed specifically for managing AI projects.

3. Business Understanding

Learning objectives

By the end of this chapter, you will be able to:

- Decide whether a problem is a good fit for AI, and recognize when a non-AI approach such as automation, a rule, a formula, or a person serves the business better.
- Frame a business problem with the nine-step Business Understanding framework, all the way to a defensible decision.
- Map a business problem to the right one of the seven patterns of AI, and always compare it against its non-AI alternative.
- Run the AI go or no-go assessment and define return on investment, key performance indicators, and acceptable performance thresholds across business, model, and technology levels.
- Assess risks, plan adoption, draft the solution, identify resources, and produce the Phase I deliverables and decision gate.

Introduction

Phase I, Business Understanding, is where an AI project earns the right to exist. The discipline is simple to state and easy to skip: align the project with a clearly defined business need before anything technical happens.

AI should not be pursued because it is a trend. It should be applied as a targeted solution to a real business problem. When you articulate the core problem first, you can see where AI actually adds value, translate that need into a viable opportunity, and choose an approach that reaches a measurable business outcome.

One message runs through this entire chapter, and it is worth saying plainly at the start. Not every business problem is an AI problem. Phase I is the place to find out, before money is spent.

Figure 3.1 shows where Phase I sits in the life cycle and what it hands to Phase II.

Three questions kept coming back from candidates preparing for the PMI-CPMAI certification:

- **In what order should I learn this?**
- **How does what I am studying map to the Examination Content Outline?**
- **What do I actually need to master to pass?**

This book answers those three questions. It claims no new method. CPMAI is the six-phase methodology created by Cognilytica and now owned by the Project Management Institute, and the official course remains the authoritative source and a prerequisite for sitting the exam. What this book adds is an order, a map, and a working example.

An order. One continuous path through the six phases, a chapter each, closing on the deliverables and the decision gate that end it. A single case study, RetentionRadar, runs from the first business question to production and operational governance.

A map. Appendix A links all 37 Examination Content Outline tasks to the chapter that teaches each one, so a weak area sends you to a chapter rather than to a search.

A focus. Two chapters describe the exam itself, its five domains and their published weights, what it does and does not assess, and how to read a question under time.

- Six CPMAI phases, six chapters, with their deliverables and decision gates
- Six chapters on responsible AI, from governance to generative AI
- 15 reusable deliverable templates for real projects
- 82 diagrams and 78 tables

Written for candidates preparing the PMI-CPMAI certification, and for project managers who must deliver AI initiatives without becoming AI engineers.

LAMJED CHIHAOUI

PMP · PMI-PMOCP · PMI-CPMAI

is a project manager with more than ten years of experience across banking, retail, e-commerce, and government. He holds an engineering degree in computer science and trains project managers and technology professionals on AI adoption and responsible AI practice.