

# Diagnosing AI Projects

*The Issue Space Field Guide for Managers, Consultants, and AI Teams*

Dr. Shengxing Yang

## Copyright

First edition. Copyright © 2026 Shengxing Yang. All rights reserved.

## Disclaimer

This book offers practical guidance for diagnosing AI projects. It is written for educational and managerial purposes and does not provide legal, regulatory, financial, investment, tax, employment, cybersecurity, or other professional advice. The cases and examples are illustrative and simplified. They are not a substitute for technical testing, compliance review, domain-specific assessment, or professional judgment in your own organization. Readers remain responsible for decisions about tools, vendors, governance, staffing, and risk.

## How To Use This Book

If you are new to the framework, read the Introduction and Chapters 1 and 2 first.

If you want to apply the method quickly, go straight to Chapter 3.

If you want to see how the method works in practice, read the two cases after Chapter 3 and use the companion toolkit.

If you are reading the Amazon edition, note that the toolkit is not included in every retail version. The full direct edition, available from the author's website, includes the toolkit. If you already have that edition, use the toolkit after Chapter 3.

## Contents

- [Reader Promise](#)
- [Introduction: Why AI Projects Feel Harder Than They Should](#)
- [Chapter 1: The AI Context](#)
- [Chapter 2: What Issue Space Is](#)
- [Chapter 3: How To Use Issue Space](#)
- [Chapter 4: Case 1: "We Need More AI Talent"](#)

- [Chapter 5: Case 2: AI, Platforms, and the Loss of Discretion](#)
- [Chapter 6: From Diagnosis To Action](#)

## Reader Promise

Many AI projects are diagnosed too fast. Teams see a visible problem, give it a simple name, and rush toward a fix: better data, more talent, another pilot, a new tool. Sometimes that works. Often it does not, because the visible problem is only one part of a larger pattern. This guide introduces Issue Space, a practical method for breaking messy AI challenges into clearer parts so action becomes more precise and useful.

You can use it in a steering meeting, a client workshop, or a pilot reset when an AI initiative has started to drift. It is written for managers who need a better agenda, consultants who need a sharper early diagnosis, and AI teams who need a shared language with nontechnical colleagues. Used well, it also builds AI literacy. It helps readers ask better questions about judgment, workflow, capability, and accountability before they buy, build, or scale anything.

## Introduction

### Why AI Projects Feel Harder Than They Should

Imagine a familiar meeting.

- The head of innovation says the company needs to move faster with AI.
- The operations lead says the latest pilot never made it into the workflow.
- Someone from IT says the data is still fragmented.
- A business manager says the model output looks impressive in demos but does not help real decisions.
- Legal is worried about accountability.
- HR says the company does not have enough AI talent.

Everyone is partly right. Nobody is fully wrong. Yet the discussion keeps circling.

Often the problem is not effort. It is diagnosis.

Organizations tend to name the first visible problem and treat it as the whole problem. If a model underperforms, the diagnosis becomes a data problem. If adoption stalls, it becomes a change-management problem. If leaders feel behind, it becomes a talent problem. These labels are convenient. They also hide a lot.

AI projects are rarely clean, single-cause situations. They are usually compound. A complaint about talent may also involve weak business ownership. A complaint about trust may also involve poor explanation, missing domain knowledge, and unrealistic expectations. A complaint about infrastructure may reveal that the organization never agreed on what problem the system was supposed to solve.

Instead of asking, "Which AI tool should we buy?" or "Which capability should we build first?" this guide asks a more basic question:

### What issue are we actually dealing with?

That question sounds modest, but it changes the order of thinking.

You can use that question in a steering meeting, a client workshop, or a project reset when a pilot has stalled.

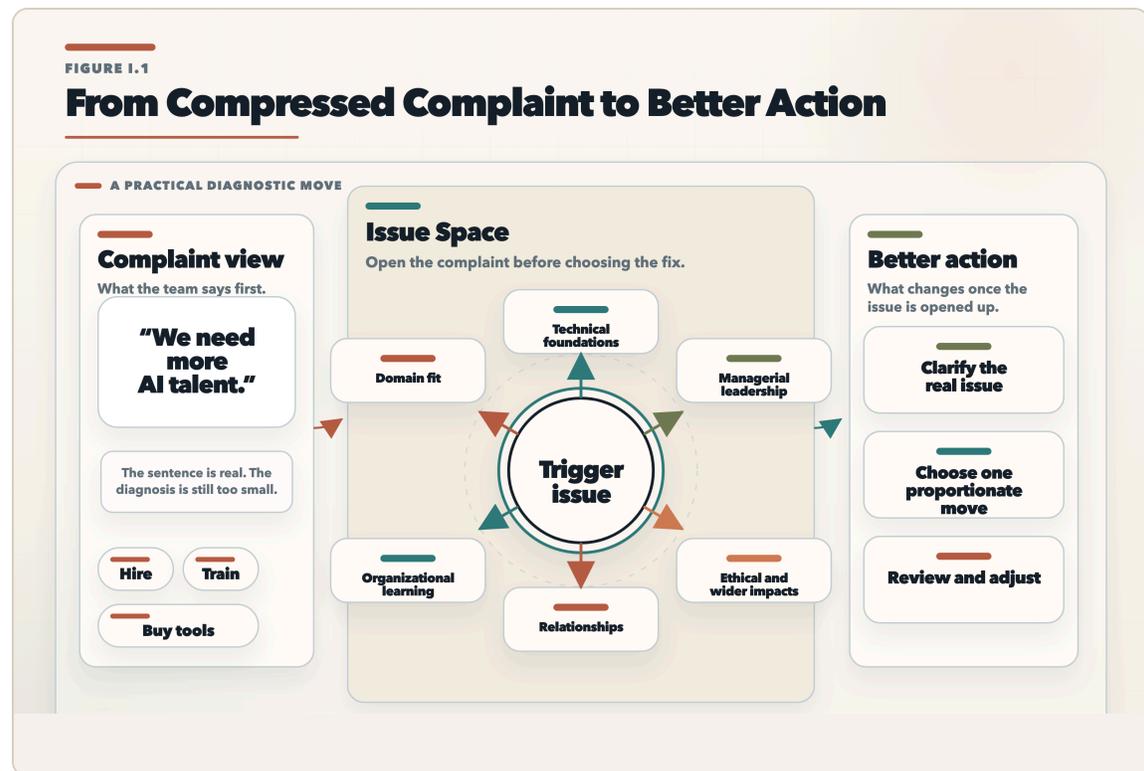


Figure I.1. Issue Space creates a diagnostic layer between a complaint and a costly response.

Most AI conversations are solution-first. Teams jump from signal to answer:

- A board hears that competitors are adopting AI, so it asks for a roadmap.
- A business unit sees a new model demo, so it asks for a pilot.
- A firm struggles to hire machine-learning engineers, so it assumes the core obstacle is technical talent.

Action starts fast, but understanding stays shallow.

Issue Space slows that first jump. It slows premature certainty.

The method in this book grew out of doctoral research on AI transformation and competence design. One pattern kept returning. Organizations were not simply struggling with technology. They were struggling to interpret and organize the issues created by technology.

The same initiative could be described as:

- a workflow problem,
- a leadership problem,

- a data problem,
- a trust problem,
- a skills problem,
- a governance problem.

That does not mean one view is right and the others are useless. It means the issue is spread across several dimensions.

Issue Space maps the structure of an AI problem before you commit to a response. It helps you decompose a messy concern into six dimensions:

1. domain-specific fit,
2. technical and technological foundations,
3. managerial leadership,
4. organizational intelligence,
5. relationships and networking,
6. ethical and wider impacts.

These six dimensions are not a maturity model or a scorecard. They are a diagnostic map. Their job is to help you see where the problem clusters, where dimensions interact, and where your current explanation is too thin.

You do not need to hold all six dimensions in your head at once. Their value becomes clearer as the book moves from concept to use.

Take a familiar sentence: "Our company lacks AI talent." That may be true. It may also hide several different diagnoses:

- the firm lacks machine-learning engineers;
- managers cannot frame use cases clearly enough;
- domain experts are not in position to judge the output;
- the organization lacks people who can work across legal, technical, and business boundaries.

A single sentence can hide several different problems.

When organizations do not unpack that sentence, they often spend money in the wrong place:

- they hire more technical people when the deeper gap is strategic ownership;
- they buy a new platform when the real problem is low-quality process knowledge;
- they demand stronger controls when the project mainly needs better user feedback and better cross-functional translation.

AI creates frustration because it looks like a technology shift, but behaves like an organizational stress test.

AI puts pressure on many things at once:

- strategy,
- data,
- workflows,
- interpretation,
- coordination,
- skills,
- trust,
- accountability.

It forces people with different vocabularies to work together. It moves fast enough to keep plans unstable. It often produces outputs that seem useful before the organization has built the discipline to use them well. It can improve performance in one area while creating hidden risks in another.

This guide does not present AI transformation as a single road with a universal sequence. It treats AI work as a set of recurring diagnostic challenges.

This matters for at least three groups of readers:

- For managers, Issue Space offers a way to reduce vague discussion and turn a broad concern into a more precise agenda.
- For consultants, it offers a structure for early diagnosis that is richer than a tool comparison and more practical than abstract strategy language.
- For AI practitioners and AI transformation teams, it offers a way to explain why a technical answer alone often fails inside a real organization.

Across all three groups, the guide also works as a practical AI literacy tool. It helps people ask better questions before a project becomes expensive, political, or stuck.

The method is simple enough to use in early project conversations, but serious enough to challenge weak assumptions. It gives people enough structure to think clearly without forcing them into false precision.

This short book is organized in a practical order.

Chapter 1 explains why AI projects create confusion even when the technology works.

Chapter 2 introduces Issue Space itself and explains the six dimensions in plain language.

Later chapters show how to use the method and how it works in concrete cases.

If you remember one line from this introduction, let it be this:

**Most AI projects do not suffer from a lack of action. They suffer from a shallow diagnosis.**

## Chapter 1

## **The AI Context: Why Complexity Keeps Returning**

People often talk about AI as if the main question were whether it works. In organizations, that is never enough. A system can perform well in testing and still create confusion, delay, weak adoption, or poor value in use.

AI is more disruptive than a standard software purchase because it makes organizational life more unstable, more interdependent, and harder to interpret. Four features matter especially.

The term AI itself does not have one final, settled meaning. In this guide, I use it broadly to describe systems that shift part of sensing, prediction, generation, recommendation, or decision support into computational intelligence. The practical question is not whether a system counts as "real AI." It is how an organization decides what should remain with human judgment, what can be delegated to computational systems, and how the two should work together.

If you are a manager, think of them as four reasons a good demo can still become a bad project.

### **1. AI changes fast**

Most organizations are used to technologies that change in steps. AI changes in waves. Models improve quickly, vendors reposition fast, and plans built on one generation of tools can look outdated before implementation settles.

This creates drift. Leaders feel pressure to act before the opportunity closes. Technical teams keep rebuilding. Governance, training, and workflow redesign move more slowly than the tools. Projects start with one expectation and get judged by another.

### **2. AI depends heavily on data and context**

AI systems learn from patterns in data, but value depends on more than data volume. It depends on relevance, workflow fit, and local knowledge. Historical data may be incomplete or biased. Important business logic may live in tacit knowledge. The organization may have data without a shared view of what success looks like in the real setting.

This helps explain why promising pilots often disappoint in daily work. A model can succeed in a controlled setting and fail in a shifting one. Users reject outputs that do not fit the task, even when the system looks strong on paper. AI is rarely plug-and-play. It needs local understanding.

In newer generative systems, this also means context engineering: shaping prompts, instructions, retrieved material, examples, workflow state, and guardrails so the model gets the right task context at the right moment. Many failures that look like "bad AI" are really failures of context design inside a live workflow.

### **3. AI often hides its own logic**

Many AI systems produce useful outputs without offering an explanation every stakeholder can understand. Even when a model is not a total black box, the reasoning can still be hard to communicate. Different people need different explanations: technical validity, operational clarity, accountability, reassurance. When explanation breaks down, managers defer too much, users resist, and governance focuses on the wrong controls. Opacity is therefore a management problem as much as a technical one. It shapes trust, ownership, oversight, and decision quality.

#### **4. AI touches several parts of the organization at once**

A conventional software project may stay mostly within one function. AI rarely does. It pulls together data, infrastructure, operations, business logic, workflow design, risk, legal review, user behavior, vendor choices, and often public concerns.

Different groups enter the conversation with different goals.

- Technical teams may care about feasibility, latency, data quality, and model performance.
- Business leaders may care about value, speed, adoption, and competitive pressure.
- Legal and compliance teams may care about documentation, accountability, and exposure.
- Frontline users may care about whether the tool helps them do their job without raising risk or reducing autonomy.

Each group sees part of the picture. Few see the whole issue at once.

#### **AI creates compound problems**

Put these four features together and AI issues become compound. A company may believe it has a talent shortage when the real problem is a mix of shifting tools, unclear ownership, weak workflow design, unrealistic expectations, and poor coordination. Another may call its challenge a governance problem when governance is only the visible expression of poor domain fit and weak explanation.

AI conversations feel slippery because people are often describing different parts of the same issue.

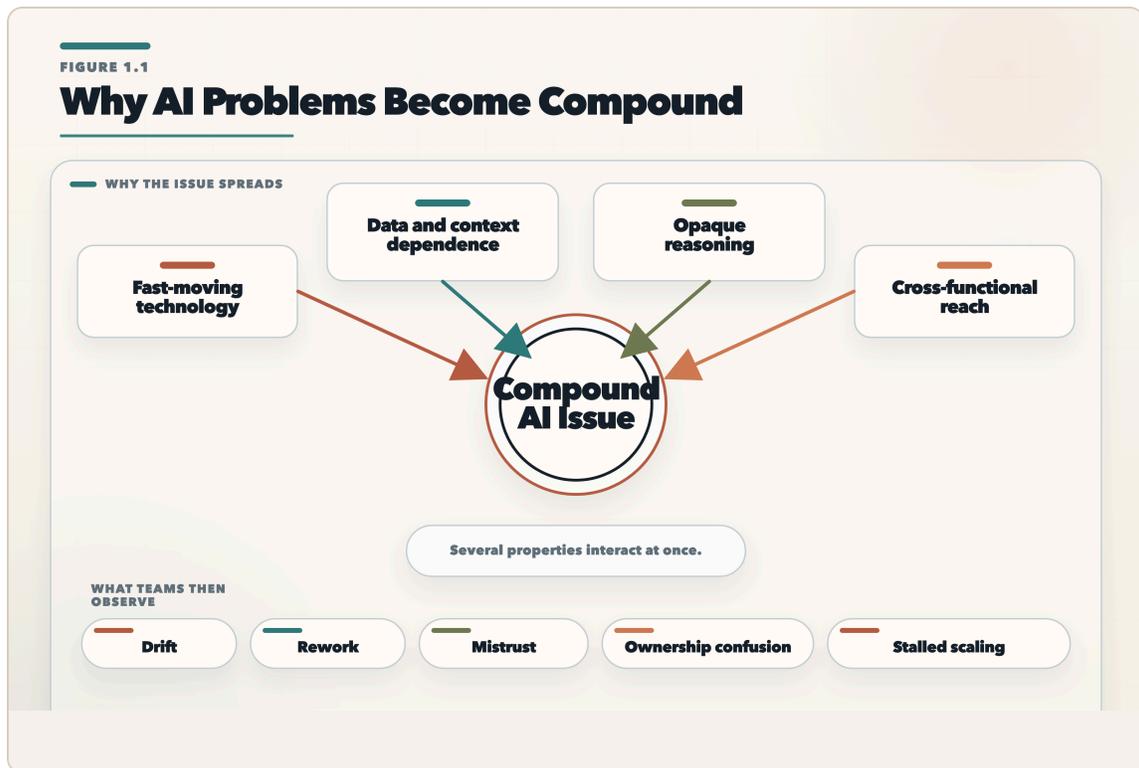


Figure 1.1. Four recurring pressures turn AI work into compound organizational problems.

## The productivity paradox is still with us

AI is associated with real productivity gains in some settings. Many organizations still struggle to convert AI activity into durable value. That gap should not surprise us. Business impact depends on complements: routines, decision rights, training, incentives, and redesign of actual tasks. Without them, AI stays trapped in pilots or demos.

## AI is also changing the role of judgment

Many discussions about AI focus on efficiency. That is too narrow. AI also changes who interprets, who decides, and who takes responsibility. It can augment judgment, narrow it, or quietly redistribute it.

When an organization automates part of a decision, accuracy is only one question. It also needs to decide what human role remains, who can challenge the output, and what happens when the system is wrong.

## Why simple explanations fail

Because AI moves across strategy, infrastructure, workflow, learning, coordination, and accountability, simple explanations fail quickly. Teams may have symptoms, opinions, and even preferred solutions, but no good way to represent the issue in front of them.

That is the gap Issue Space is meant to fill. Diagnosis is not delay. It is how organizations move faster without wasting effort on the wrong fix.

# Chapter 2

## What Issue Space Is

Issue Space is a practical way to diagnose AI problems before jumping into solutions.

It starts from a simple observation: organizations often respond to AI challenges with labels that are too small for the real problem. "Talent," "data," or "governance" may each be partly right and still flatten a more complex pattern. Issue Space is designed to recover that pattern.

At its core, Issue Space is a structured map of the main dimensions along which AI problems tend to appear inside organizations.

It does three jobs at once:

1. it helps people decompose a messy issue,
2. it helps teams compare interpretations across functions,
3. it helps connect diagnosis to action.

The key shift is simple:

**Issue Space puts issue representation before solution choice.**

Suppose a company says, "We need more AI talent."

A solution-first response might be immediate:

- hire more data scientists,
- bring in an external vendor,
- run AI training,
- buy a better platform.

Any of those could help. None should be automatic.

An Issue Space response would pause and ask:

- What kind of talent do we mean?
- For which use case?
- In which workflow?
- Who cannot currently make good decisions?
- Where is coordination breaking down?
- What risks are being ignored?
- Which parts of the issue are technical, and which are managerial or organizational?

The complaint does not disappear. It becomes more useful.

**Issue Space is a map, not a checklist**

A checklist asks whether each box has been completed.

A map helps you understand where you are, where the pressure points are, and how different parts of the terrain connect.

If people misuse Issue Space as a checklist, they will ask superficial questions:

- Did we cover ethics?
- Did we talk about governance?
- Did we mention data?

The real value lies in seeing emphasis, interaction, and blind spots. Some issues cluster strongly in one or two dimensions. Others spread across four or five. Some dimensions are primary. Others are secondary but still important because they shape the consequences of action.

Issue Space therefore helps you ask better questions, not merely more questions.

### **Why issues are a good starting point**

Many organizations prefer to start with capabilities. That is understandable. Capabilities sound strategic. The problem is that capability talk can drift away from the actual difficulty on the ground.

A firm may build AI literacy programs, create a governance committee, hire technical specialists, and still fail to move because it has not clarified what issues it keeps facing and how those issues relate to one another. Issue Space starts somewhere more concrete: the problems people are actually encountering.

These problems matter because they show where the organization is struggling to align technology, work, people, and judgment. The path to better capability often runs through better diagnosis.

### **The six dimensions of Issue Space**

The framework used in this book organizes AI issues into six dimensions.

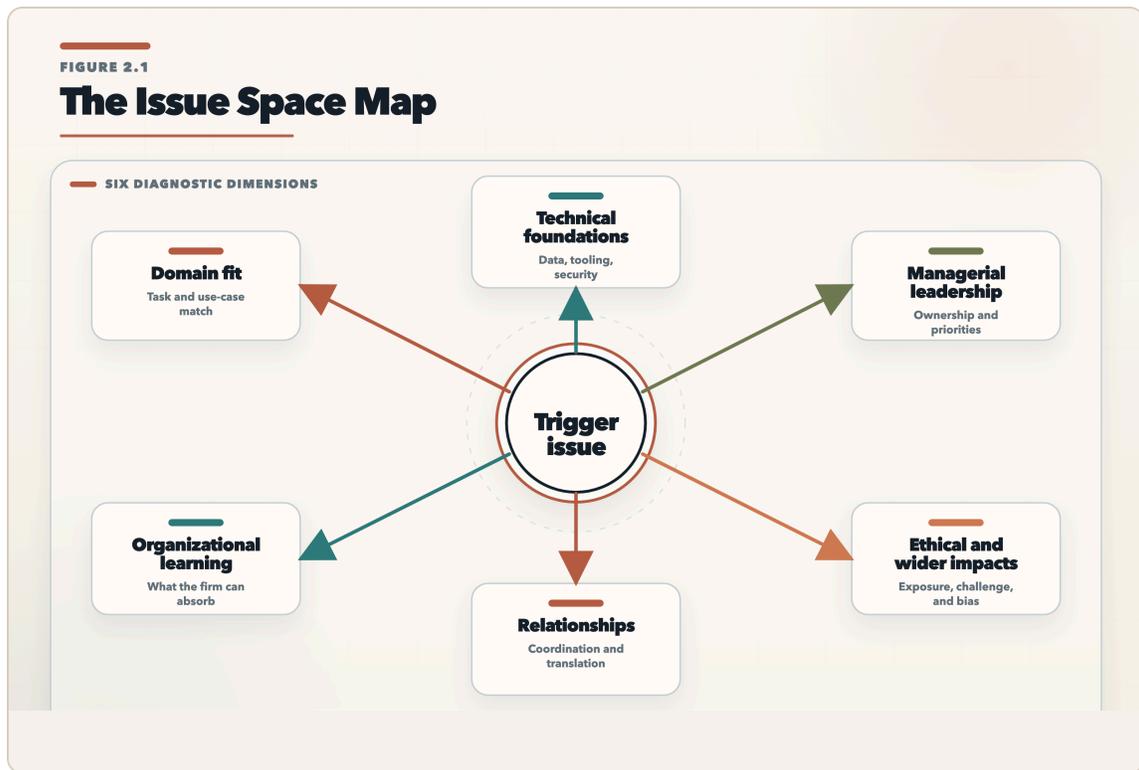


Figure 2.1. Issue Space maps one trigger issue across six interacting dimensions.

### 1. Domain-specific fit

This dimension asks whether the AI approach fits the real logic of the domain. AI is always applied somewhere, and each setting has its own workflows, standards, constraints, tacit knowledge, and error costs.

Ask whether the system matches the real task, whether domain experts stay involved, and whether you understand what a meaningful error looks like in this context. When this dimension is weak, organizations build technically interesting tools that fit poorly into real work.

### 2. Technical and technological foundations

This dimension covers data quality, infrastructure, integration, security, tooling, and the discipline needed to build and maintain the system well.

Ask whether the system fits current workflows, whether the data is fit for purpose, whether it can be monitored over time, and whether reliability or security risks are being handled seriously. Weak foundations can block the whole effort, but technical issues are not always the whole story.

### 3. Managerial leadership

This dimension asks whether the project is being led and framed well. AI projects need more than executive enthusiasm. They need clear ownership, strategic alignment, realistic ambition, and the ability to make trade-offs across functions.

Ask who owns the problem, whether the effort is linked to a real strategic goal, whether commitments on time and scope are realistic, and whether leaders can

resolve cross-functional conflict. Weak leadership usually shows up as drift, confused priorities, or symbolic activity.

#### **4. Organizational intelligence**

This dimension is about how the organization learns and adapts. AI rarely delivers value through technology alone. It usually needs complementary changes in routines, incentives, training, workflow design, and decision processes.

Ask whether the firm knows how ready it really is, whether it learns from mistakes, whether it is building the complements needed for value, and whether work can be redesigned as AI changes the task. When this dimension is weak, organizations confuse activity with progress.

#### **5. Relationships and networking**

This dimension covers coordination across people, teams, and external partners. AI creates translation problems because managers, users, compliance staff, vendors, and technical teams often work with different assumptions and vocabularies.

Ask whether different groups can explain the issue to one another in usable language, whether expectations are aligned, whether the organization is too dependent on external partners, and whether users have a real feedback channel. When this dimension is weak, the project becomes a collection of disconnected views.

#### **6. Ethical and wider impacts**

This dimension asks what broader consequences the organization must take seriously: fairness, accountability, explainability, human oversight, regulatory exposure, privacy, labor effects, and the changing role of human judgment.

Ask who is exposed if the system is wrong, whether someone can challenge the output, whether sensitive data and bias are being handled seriously, and whether the system could reduce useful human judgment even if efficiency improves. When this dimension is ignored, organizations often discover the real cost later.

### **The dimensions interact**

One of the biggest mistakes in AI work is to treat these dimensions as separate silos. They are analytically distinct, but in practice they interact all the time.

- Poor leadership can make technical problems worse because nobody settles priorities.
- Weak domain fit can make ethical review shallow because the real context of harm has not been understood.
- Poor relationships across functions can make trust collapse even when the model is strong.
- Weak organizational learning can trap a company in endless pilots even with good executive support.

## **What Issue Space is not**

Issue Space is not:

- a promise that every AI problem can be solved with a framework,
- a replacement for domain expertise,
- a substitute for technical evaluation,
- a governance system by itself,
- a score that tells you how mature your company is.

It is a way to structure attention. In organizations, attention and shared language are scarce. A method that improves them often changes action more than another abstract strategy statement.

## **From issue diagnosis to action**

Once an issue has been mapped, teams can make better choices. They can see which dimension is primary and which is secondary. They can decide whether the immediate need is better data, tighter ownership, stronger domain input, clearer governance, or a redesign of workflow.

This is also how Issue Space connects to competence building. It does not tell you that one competence matters forever. It helps you see which configuration matters now.

## **A simple example**

Return to the sentence: "We need more AI talent." Issue Space does not reject it. It unpacks it.

In one company, it may reveal a genuine shortage of technical specialists. In another, it may reveal that managers cannot define use cases well enough for technical teams to execute. In another, it may reveal that the firm has specialists but lacks domain experts who can evaluate the output. In another, it may reveal that people use "talent" as a polite way to describe poor leadership, unclear ownership, or weak training.

The same complaint can point to very different action paths. Issue Space gives teams a disciplined way to find those paths before they spend heavily on the wrong fix.

## **The practical value of Issue Space**

If this framework works well, three things improve first. Conversations become more precise. Hidden disagreements become visible. Action becomes more proportional.

That is the practical promise of Issue Space. It does not remove complexity from AI transformation. It helps you face that complexity with better structure, language, and judgment.

Issue Space works because it is strong enough to sharpen attention and simple enough to survive in real organizational conversations. For that reason, the direct-sales version includes a practical toolkit. Readers often do not need another hundred pages of argument. They need a way to use the argument on Monday morning.

## **A final note**

AI will keep changing. New models will appear. Platform logic will spread. Governance expectations will tighten. More work will be pushed toward agentic systems, and more organizations will discover that buying capability is not the same as building judgment.

For that reason, the central question of this book will remain useful:

### **What issue are we actually dealing with?**

Ask that question early enough, and the quality of the whole project can change. Ask it too late, and the organization may spend months or years fixing symptoms that were never the real problem.

Issue Space does not eliminate that risk. It gives you a better way to see it.

## **Note On Cases And References**

The two cases in this book are composite cases. They combine recurring patterns from research, public reports, and official materials. They are meant to clarify issue patterns, not to disguise one company or institution as another.

The selected references at the end of the case chapters show the public materials that informed the cases and give readers a starting point if they want to go further.

## **How To Use This Book With A Team**

If you want to use this guide in a live project, start with one real complaint sentence. Do not begin with a full transformation plan.

Use Chapter 3 to run a short diagnostic session. Keep the first pass small. The goal is not to solve everything in one meeting. The goal is to replace a compressed diagnosis with a clearer one.

Good first questions are simple:

- What issue are we actually naming?
- Which dimension is carrying the most weight?
- Where is the current explanation too thin?
- What should happen next?

## **About The Author**

Dr. Shengxing Yang researches and writes on AI transformation, Issue Space, competence, platform dynamics, and institutional change. This book grows out of doctoral work and related writing on how organizations diagnose AI problems before they commit to solutions; Web: [aigovmap.com](http://aigovmap.com); email: [yangshengxing88@163.com](mailto:yangshengxing88@163.com)