

BOARD ROOM ESSAY

EXECUTIVE BRIEF NO. 002

You Can't Lead What You Haven't Defined

Your executive team is nodding at the same AI strategy slides — while holding three different definitions of AI. That gap is where strategy fragments and capital gets wasted.

AI LEADERSHIP

6 MIN READ • AUGUST 9, 2026



You Can't Lead What You Haven't Defined

Ask your executive team to define AI and you will get as many answers as people in the room. Shared language is not a soft aspiration. It is the operational requirement for everything that follows.

BY TOM LAWRY

Here is a simple test. At your next executive team meeting, give everyone five minutes and ask them to write down their definition of artificial intelligence in 100 words or less.

Then compare the answers.

You will see the problem immediately. Some will describe a chatbot. Others will describe machine learning models. Some will write about automation. The Chief Medical Officer's definition will not match the Chief Financial Officer's. The Chief Information Officer's may not match either of them.

And yet all three just sat through a board presentation on your organization's AI strategy — nodding at the same slides, believing they agreed on the same things.

It is hard to lead something you have not clearly defined. And in the age of AI, vague definitions do not stay vague. They become missed opportunities and expensive mistakes.

What you need to know — and what you don't

Let's be clear about the goal. It is not technical mastery. It is strategic and operational fluency — the foundation every health leader needs to make sound decisions, ask the right questions, and hold their organizations accountable.

Your role is to set direction, allocate capital, manage risk, and hold teams accountable. That demands informed judgment — not coding skills.

The leaders who navigate this era well will not be the ones who can explain transformer architectures. They will be the ones who can separate signal from noise — who know when a vendor claim is credible and when it is marketing, and who can tell the difference between a strategic investment and an expensive experiment.

“Foundational knowledge enables better questions. Better questions enable better decisions.”

Creating a shared language

AI conversations do not stay in one lane. They cut across clinical, financial, operational, compliance, and technical domains simultaneously. Each group brings different expertise, incentives, and assumptions — and each uses the term “AI” to mean something slightly different.

When definitions vary, strategy fragments. When terminology is unclear, governance weakens. When leaders believe they are aligned but use the same words to mean different things, resources flow to different

visions — and no one realizes it until the work is already underway.

Shared language is not a soft aspiration. It is a hard operational requirement — the prerequisite for everything that follows: strategy, governance, investment, workforce development, and vendor accountability.

AI is not just another IT upgrade. It is healthcare's new decision infrastructure.

A working definition

AI is not a single technology. It is a family of capabilities — different tools, built on different principles, designed to do different things. So anchor the conversation with a single, non-technical definition:

Artificial Intelligence is the ability of computer systems to perform tasks that normally require human intelligence — recognizing patterns, learning from data, making predictions, and supporting decisions.¹

That definition is intentionally broad. The AI that flags an abnormality in a chest X-ray and the AI that drafts a clinical note from an ambient conversation are both “AI” — but they work on different principles and serve completely different purposes.

The word alone tells you almost nothing. What matters is what kind of AI, applied to what kind of problem, with what kind of data.

With that foundation, here is an executive map of the landscape — not a glossary, but three clusters of capability, each solving a different category of problems.

Cluster 1: The pattern-recognition core

These are the tools that find meaning in data. They learn from examples rather than following programmed rules — and they power most practical AI in healthcare today.

- **Machine Learning** is the foundation. These systems identify patterns in large datasets and improve as they process more data. In healthcare, machine learning predicts which patients are likely to be readmitted, forecasts emergency department demand, and flags deteriorating patients before a crisis occurs.
- **Deep Learning** is a more sophisticated form that excels at complex data — images, audio, genomic sequences. It is what allows AI to analyze pathology slides, interpret MRI results, and detect subtle patterns trained human eyes might miss.
- **Computer Vision** applies deep learning to visual data. It detects early signs of cancer in a mammogram and flags retinal abnormalities before symptoms appear. It does not “see” the way humans do — but in specific, well-defined tasks, it can match or exceed clinical accuracy.

Cluster 2: The language layer

These tools let AI work with human language — reading, understanding, generating, conversing. This cluster is behind most of what has captured public attention since 2022.

- **Natural Language Processing** enables computers to read, interpret, and generate human language. It powers tools that listen to clinical conversations and generate documentation automatically. NLP has been in clinical use for years; what changed is its sophistication and accessibility.
- **Large Language Models** are the engine behind today's most visible AI tools. Trained on vast amounts of text, they let anyone interact with computers in plain language — no programming, no specialized interfaces. That is what makes them so accessible. And so powerful.

- **Generative AI** is the capability that reset public expectations. Rather than simply analyzing data, it creates — drafting clinical notes, generating patient education materials, synthesizing research. It was built on decades of machine learning research. What changed was scale, accessibility, and capability.

Cluster 3: The automation layer

These tools take action, not just analysis. Where the pattern-recognition core finds insights and the language layer communicates them, the automation layer executes.


- **Robotic Process Automation** handles repetitive, rules-based digital tasks —

scheduling, claims processing, prior authorization workflows. Purists debate whether RPA counts as “true” AI. For leaders, the debate misses the point. What matters is operational value: RPA reduces administrative burden at scale without requiring clinical workflow redesign.

- **Agentic AI** is the emerging frontier. Rather than responding to a single prompt, agentic systems autonomously plan and execute multi-step sequences — checking a patient's data, identifying a concern, alerting the care team, scheduling the follow-up, logging the action. No human directing each step. Still developing, advancing quickly, and destined to change how your organization thinks about automation and oversight.

The AI Toolkit: Three Clusters of Capability

AI capabilities fall into three complementary clusters — from finding patterns, to understanding language, to taking action.



The AI toolkit: an executive map. Every proposal that reaches the board should declare its cluster — and the problem it solves.

Why AI still needs you

None of these capabilities exist in isolation. Every AI system is shaped by three things: the data it was trained on, the infrastructure supporting it, and the humans who design, deploy, and oversee it.

Computers process data at a scale and speed no clinician can match. But humans bring what AI cannot replicate: abstract thinking, lived experience, judgment in genuinely ambiguous situations, and empathy for another person.

In healthcare, those capacities are not peripheral. They are central. The physician who reviews an AI-flagged scan brings context the model was never given — the patient's history, their preferences, the clinical nuance that never appears in structured data.

This is not a limitation of AI. It is the design of a system that works.

The message for the board room

Think of AI as a toolkit — some tools mature, some evolving, some not yet invented. They are most valuable applied to work that is high-volume, data-intensive, and well-defined: predicting risk, processing images,

automating documentation. They are least reliable when the task requires nuanced judgment, contextual awareness, or genuine human connection.

Prediction: Within three years, the health systems pulling ahead on AI will not be the ones with the biggest technology budgets. They will be the ones whose leadership teams share a common vocabulary — where the CMO, CFO, and CIO mean the same thing when they say “AI,” and every investment decision starts from that shared ground.

Recommendation: Run the five-minute definition test with your executive team this month. If the answers diverge — and they will — make building a shared AI vocabulary a governance priority, not a training afterthought. Adopt one working definition. Use the three-cluster map. Require every AI proposal that reaches your desk to state which cluster it draws from and what problem it solves.

Understanding what AI is — and what it is not — is not an academic exercise. It is the prerequisite for every strategic decision that follows.

You can't lead what you haven't defined. So define it.

SOURCE

1. Davenport, T. and Ronanki, R. (2018) 'Artificial intelligence for the real world,' Harvard Business Review, vol. 96, no. 1, pp. 108-116.

Tom Lawry is managing director of Second Century Tech, author of AI in Health, and the former national director of AI for health and life sciences at Microsoft.

www.tomlawry.com

© 2026 Tom Lawry. All rights reserved.

This essay reflects the author's views and is for general information only; it is not legal or clinical advice.