

## **AI and Workforce Strategy: Preparing Humanity for What We Know and Do Not Yet Know<sup>1</sup>**

Domenico (Mimmo) Parisi<sup>2</sup> and Jonathan Barlow<sup>3</sup>

### **Introduction**

As artificial intelligence (AI) grabs global attention, the primary focus is on what we can do and how it will shape what we may become. At the same time, however, AI is creating a vast space of unknowns, making uncertainty itself a central focus of the discussion on AI and we should prepare humanity for the future. There are so many questions that require an answer and many others we have not even considered. For example, we do not know what kind of people and what kind of societies we should aspire to be in an AI-centered world. We also do not know what ecosystem of rules, norms, incentives, and institutions will ensure AI serves the common good instead of placing power and opportunity in the hands of a few. How do we prepare for the unexpected as we might be confronted with emergent behaviors, unintended consequences, or sudden disruptions?

These are only some of the questions. Once we begin to address them and reflect on their implications, these initial questions quickly multiply. What does it mean to preserve human dignity, agency, and creativity when AI can increasingly replicate tasks once considered uniquely human? How do we build trust in systems that can be powerful yet unexplainable? Who is accountable when AI systems cause harm, and what does justice look like in an AI-centric age? How do we protect privacy and freedom while strengthening security? How do we prevent new forms of poverty and inequality (an “AI divide”)? And perhaps most urgently, how should we redesign education and workforce development so that people are not merely trained to use AI, but empowered to thrive alongside it?

Despite all these questions, today much of the conversation around AI is centered on expanding its capabilities based on what we know and already understand. Far too little attention is being paid to what remains unknown. The gap between known and unknown makes it impossible to fully foresee how AI will transform our future. What we do know is that at one end of this spectrum AI will become deeply woven into every aspect of human life, ranging from education and discovery to security and economic growth. At the other end of the spectrum there is the unknown. If we fail to address the gap between known and unknown, the consequences could be profound and even catastrophic.

---

<sup>1</sup> Adapted from a presentation originally delivered 3/10/2026 at the conference *Fare per imparare: Etica e formazione nell'era dell'intelligenza artificiale*, Università Pontificia Antonianum.

<sup>2</sup> Domenico (Mimmo) Parisi is a professor and the Executive Director of the Data Science Academic Institute at Mississippi State University.

<sup>3</sup> Jonathan Barlow is a professor and the Associate Director of the Data Science Academic Institute at Mississippi State University.

How, then, should we think about addressing or reducing this gap? In a recent article, Hollis Robbins confronts this very question as she reflects on the future of AI in higher education in the United States.<sup>4</sup> She argues that “no one has the omniscience to predict what is coming for the sector.” Rather than assuming that the future can be forecasted with certainty, Robbins suggests that a more useful approach is to organize our thinking not only around what we know about AI, but also around what we do not yet know. She uses Rumsfeld’s “known-unknown” matrix, a useful way to structure the analysis of the gap between what we know and what we do not know about AI by distinguishing four categories: (1) known knowns, (2) known unknowns, (3) unknown knowns, and (4) unknown unknowns. We may adapt the matrix like this:

	Know	Don't Know
Aware	<p><b>Known Knowns</b> Things we're aware that we know.</p>	<p><b>Known Unknowns</b> Things we're aware that we don't know.</p>
Unaware	<p><b>Unknown Knowns</b> Things we're unaware that we know.</p>	<p><b>Unknown Unknowns</b> Things we're unaware that we don't know.</p>

Because there are many things we are unaware that we do not know (unknown unknowns) the most effective way to close the gap is to address the first three categories. By doing so, an organization will be better positioned to respond to the unexpected and build greater resilience in the face of uncertainty. How effectively we address this gap will distinguish winners from losers. The first step is to examine closely what each category contributes to the overall mix. Next, we will develop a strategy that strengthens resilience to unknowns. We can use this general approach to shape an AI-based workforce strategy.

**Known-Knowns: Things we are aware that we know**

AI cannot be understood simply as a piece of digital technology. It is rather best described as a system built on four interlocking components: data, learning algorithms, computational power, and agentic coordination. Together, these components determine what AI systems can learn from

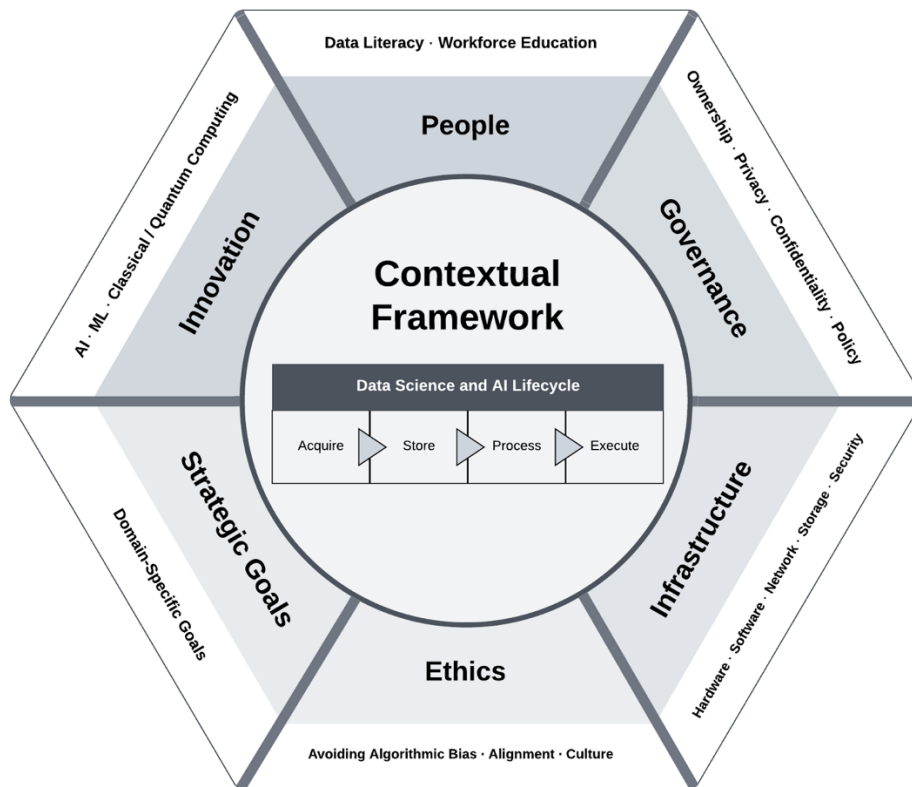
---

<sup>4</sup> Hollis Robbins, “The Rumsfeld Matrix,” *Anecdotal Value*, February 6, 2026.

data, how they improve through training, how they operate at scale, and how they ultimately act in the world.

AI begins with data; this is its “digital experience.” Data is the foundation that makes AI possible. It is the digital representation of the physical world that enables mathematical models to learn patterns, represent aspects of reality, and automate decision-making processes. What an AI system can do is determined by how well data captures and represents the world in digital form, not the other way around.

Data is the backbone of AI, and it defines both its capabilities and its limits. This digital experience is captured through the data business cycle, which begins with data acquisition and ends with its use. The cycle unfolds within a well-defined context that includes ethics and legal considerations, goals and purpose, enabling technologies and innovation, and governance (see the diagram below).



Training is the next pillar of AI. It is the process of exposing learning algorithms to digital experiences (the physical or conceptual world represented through data objects) so that we can derive usable representations of that world. The goal is not simply to learn more about the physical world, but to learn how to enhance human capabilities effectively, thereby reducing the need for human intervention in the execution of tasks and decisions throughout a process.

A third pillar is the provision of appropriate computational scale at several stages. First, computational capacity is needed for pre-training, where models are developed from vast amounts of data. Next, computational capacity is required for post-training and fine-tuning, so that AI systems can acquire specific capabilities such as reasoning, mathematical ability, and task performance, while also incorporating ethical safeguards and behavioral constraints. Finally, computational capacity is needed at the deployment stage to enable AI systems to be used effectively. This requires applications of traditional computing, high performance computing, and, perhaps soon, quantum computing techniques.

The fourth pillar is agentic coordination, in which AI functions not merely as a system that responds to prompts, but as one that carries out work and supports decision-making. Multiple AI systems can be coordinated to achieve organizational objectives. We might quickly reach a point at which AI systems within an organization outnumber its human workforce and become part of a broader system of decision-making and execution.

Taken together, these four components form what is commonly known as modern AI. They also shape the questions we must ask about capability, control, accountability, and the kinds of societies we are building as these systems become more powerful and more autonomous. It helps to identify the things we must know to reduce the gaps with the unknown.

### **Known-Unknowns: Things we are aware that we do not know**

The known-unknowns category is about those unsolved challenges faced by AI of which we're already aware. We know we will face several problems with the use of AI systems or tools, but we still do not know how to fully address them. For example, across both the private and public sectors, organizations face a similar challenge of not knowing how much of what can be done in theory can be done in practice to reap the full AI benefits. In the private sector, the extent of this gap tends to be connected to strategic and cultural challenges such as unclear priorities, fragmented ownership, and a tendency to treat AI as a product to purchase rather than a capability to build. In the public sector, these same challenges are compounded by additional constraints such as legal requirements, political oversight, and democratic expectations around transparency, fairness, and accountability.

We also know that data itself remains a major challenge, and we still do not know how to address it fully. Many organizations continue to rely on siloed, inconsistent, and legacy data systems that were never designed to work together. Yet AI depends on data that is clean, well structured, well documented, and accessible. Without a strong data foundation, even the most advanced AI models will perform poorly.

Another typical challenge is that we know how AI performs in a pilot, but not how it will perform when scaled. Many institutions have experimented with AI through small pilot projects, but far fewer have succeeded in integrating it as a durable and valuable part of their operations. Too often, pilot efforts are driven by external vendors who are not deeply embedded in the

organization's workflows, incentives, and mission. For AI systems to succeed at full scale, they must grow from within, align with real operational needs, and demonstrate tangible value at each stage of expansion.

Securing talents is also a problem that remains unresolved. Organizations across sectors struggle to attract and retain people with the skills required to design, deploy, and govern AI systems. The challenge for governments and universities is that they compete with private-sector salaries and recruiting efficiency. The challenge for the private sector is that investment in internal training and reskilling has not always grown in proportion to the need. Talent gaps can lead to "shadow AI," where employees adopt unapproved tools in ways that may violate privacy, confidentiality, or governance standards.

These problems are exacerbated by legacy systems. As in earlier waves of digital transformation, many organizations depend on older platforms that are essential to daily operations but difficult to modernize or integrate in ways that support AI systems. These systems fail not because the models are poorly designed, but because the organization's infrastructure cannot support them.

Organizations also face the challenge of anxiety and cultural resistance. Many people fear that AI systems will replace them. People in an organization are often afraid that AI will not only change their jobs, but also reduce their importance, authority, and expertise. To be sure, AI is often viewed as a competitor rather than a collaborator. This undermines effective adoption of AI systems because their successful implementation depends on trust, engagement, and repeated use. When people see AI as a threat to their role or expertise, they are less likely to embrace it as tool to improve their performance and productivity. In such situations, people in an organization might use it reluctantly, apply it inconsistently, or resist integrating it into operational and management structure of an organization. As a result, AI never becomes fully part of the organization's operations, and its benefits remain untapped.

When seen as competitors, AI systems can also encourage avoidance as people often respond to perceived threats by distancing themselves from accepting AI as part of the operational and management functions of an organization. This is manifested as people avoiding learning how AI works, hesitating to participate in AI-related initiatives, or declining to use AI tools even when those tools improve performance. To avoid conflicts, senior management becomes reluctant to sponsor AI projects, especially when they expect resistance or fear disruption. This avoidance prevents the organization from developing the literacy and basic experience needed to use AI effectively.

Perceived competition with AI also reduces the willingness to experiment with AI responsibly. AI use requires careful testing, feedback, and adaptation over time. Organizations need people to try new tools, identify weaknesses, refine processes, and learn where AI adds value and where it does not. But when fear and mistrust are strong, people become less open to experiment with AI-powered tools. Instead of testing AI in thoughtful, controlled ways, they might reject AI tools outright or dismiss them before their capabilities and limits are properly understood.

Together, these dynamics create a cycle of resistance. Fear weakens adoption, weak adoption reinforces skepticism, and skepticism makes further use less likely. In such context, an organization struggles not only to implement AI, but also to learn how to govern and improve its use over time.

Finally, organizations struggle with measurement and expectations. Many AI initiatives are abandoned because leaders demand measurable ROI on an unrealistic timeline. Productivity gains from AI are widely anticipated, but they often arrive less like a traditional technology upgrade and more like an increase in team capacity, requiring redesign of workflows, new roles, and sustained learning before the benefits become visible at scale. Without patience and clear metrics, institutions can mistake early complexity for failure and walk away just before value begins to compound.

### **Unknown-Knowns: Things we are not aware that we know**

Unknown-knowns are the things we already know (or could know) but that are not recognized, surfaced, or taken seriously because they are buried in institutions, incentives, or tacit expertise. For example, bias and ideology are two of the most important unknown-knowns. They often operate invisibly. They, however, shape decisions and outcomes even when people believe they are being neutral or purely technical.

Bias is not only an issue to be addressed within data. It can be embedded in problem definitions, labeling practices, proxy variables, evaluation metrics, and deployment contexts. Organizations might “know” that certain groups are historically underrepresented, that measurement is imperfect, or that outcomes are shaped by structural inequities. Yet these realities are frequently treated as irrelevant rather than essential to the correct use of AI systems. The result is that bias is something we could have anticipated and addressed, but unfortunately, often it remains unacknowledged until harms surface.

Ideology operates in a very similar way. It shapes which problems are seen as worth solving, which trade-offs are considered acceptable, whose values are prioritized, and what counts as success. Even choices that appear purely technical such as what to optimize, what to automate, what to measure, and what to ignore carry assumptions about efficiency, productivity, risk, fairness, and authority. When ideology remains implicit, it is often justified as common sense. As a result, organizations may fail to recognize that their own views are embedded in a particular worldview, and that this worldview can significantly shape how AI systems influence operational processes and outcomes.

Bias and ideology call for greater self-awareness by making assumptions explicit, inviting diverse perspectives into AI systems design and governance, and building mechanisms to detect and correct biases and ideological views before they become problematic for the appropriate use of AI tools. Beyond bias and ideology, there are other strong AI-relevant “unknown knowns” that cannot be ignored:

- **Tacit domain knowledge that never makes it into the development of AI systems.** This is knowledge that is not written down but it is essential for the operation of an organization.
- **Hidden assumptions embedded in workflows and policies.** Organizations run on unspoken rules (“how things really get done”). AI projects often automate the formal process while ignoring the informal one, creating failure despite “good” models.
- **Incentive structures that quietly shape outcomes.** Key performance indicators, bonus structures, vendor contracts, procurement rules, and publication incentives can drive behavior and shape choices in predictable ways but are rarely treated as part of the “technical” system.
- **Known operational risks documented somewhere but not acted on.** Prior incident reports, audit findings, safety reviews, and compliance reports are often not integrated into AI-systems.
- **Data origin and quality issues everyone suspects.** Organizations often know their data is incomplete, inconsistently labeled, or modified but those realities remain unofficial. As a result, AI systems fail to reap their full benefits.

The key to addressing unknown knowns, both at the individual and organizational scale, is self-awareness and the courage to take nothing for granted.

### **Unknown-Unknowns: Things we are unaware that we do not know**


Very little may be said about things we do not know that we do not know. The unknown-unknowns category encompasses all the things we cannot anticipate about the use of AI systems. Given the deep uncertainty surrounding AI, we must build on what we already know and understand to create resilience, so that we are better prepared to respond and adapt when the unexpected inevitably arises.

### **AI and the Future of Workforce**

Now that we have examined what we know and what we do not know about the use of AI systems and their implications, we can shift our attention to the workforce. Clearly, an AI-driven economy will require a workforce with strong foundations in mathematics, statistics, computer science, engineering, and data. In response to this growing demand, the last decade has seen the establishment and rapid growth of data science programs across many institutions. We have also seen increasing interest in ensuring that the workforce has a working understanding of ethics and legal considerations, especially those related to responsible use, accountability, and privacy.

In the current environment, many predict large-scale displacement of workers. A recent study shows that early-career workers (ages 22-25) in the most AI-exposed occupations have

experienced a 13 percent relative decline in employment.<sup>5</sup> The study also shows that employment declines are concentrated in occupations where AI is more likely to automate, rather than augment, human labor. This is well represented in the graph reproduced below.<sup>6</sup> Here the fast drops are among young software engineers.



While this early-career effect is important to consider, our focus should be on a more useful and urgent question: what should the workforce look like in an AI-centric society? Given the growing uncertainty surrounding AI, it is important to shift the conversation away from fear of job loss and toward a workforce strategy that captures new economic opportunities by redefining and reimagining roles, skills, institutions, and interventions. We should view the use of AI systems as an opportunity to advance human progress rather than as a threat to it. Workforce development should focus less on simple job preservation and more on building the governance, talent pipelines, and institutional capacity needed to support a resilient workforce. In the end, the goal is to create the conditions for continuous learning and rapid adaptation as AI generates both new opportunities and new challenges. The Rumsfeld “known–unknown” matrix can help develop a workforce strategy that accounts for the opportunities, challenges, and risks associated with the transition to an AI-centric society.

<sup>5</sup> Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen, *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence*, working paper (Stanford Digital Economy Lab, November 13, 2025).

<sup>6</sup> Source of graph: Brynjolfsson et al, 26.

As we move forward, it is useful to consider several key issues through the lens of the Rumsfeld “known–unknown” matrix:

### **Workforce Known-Knowns**

Workforce systems must include strategies that account for low-, middle-, and high-skill jobs (see diagram below).



Much of the attention has focused on high-skill work, as entry-level positions are becoming less common, especially in occupations where tasks can be automated. We must prepare the high-skill workforce to use AI effectively to augment productivity. This emerging trend challenges how universities operate and how they prepare students. The question, then, is: What should the new value proposition of higher education be, and how do we distinguish winners from losers?

### **Workforce Known-Unknowns**

While we know that the workforce system includes low-, middle-, and high-skill jobs, we do not know the optimal mix of low-, middle-, and high-skill workers needed for AI to become an engine of organizational success. There are also latent occupations for which we have neither established titles nor a clear understanding of the education and training required to enter them. When skill gaps are not yet visible, individuals and institutions cannot plan effectively to close them or ensure the development of a high-performance workforce. This raises an important question: how should we redesign the workforce system so that it is nimble enough to pivot and meet the latent demands of the market?

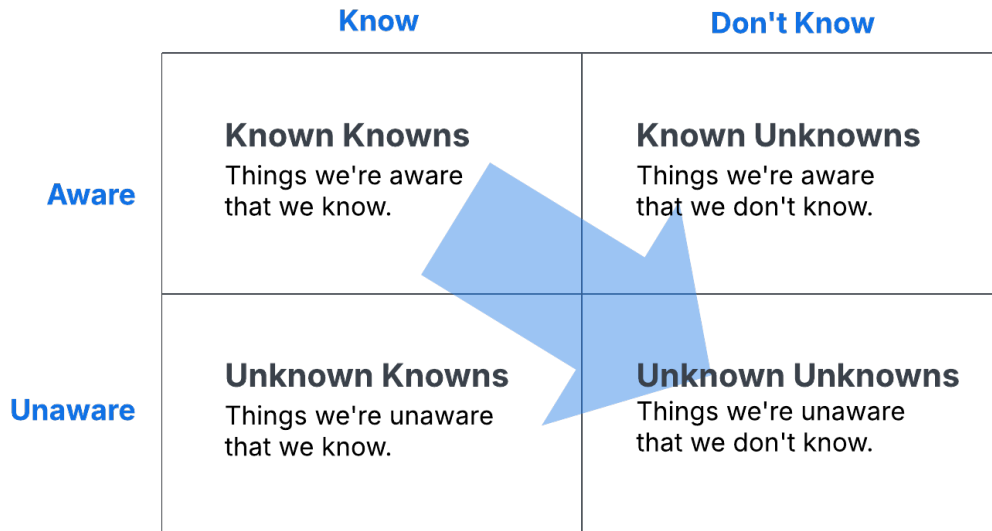
### **Workforce Unknown-Knowns**

We know that much of the workforce remains AI-illiterate. People at all skill levels (low, middle, and high) must develop sufficient AI literacy to ensure the appropriate and responsible use of AI

systems and tools. Middle and senior management must be capable of identifying and addressing biases and ideologies embedded within AI systems and data. In addition, they must be able to recognize the tacit knowledge within an organization that is essential to making AI effective. We must avoid stigmatizing AI use among the very people best positioned to uncover and apply that tacit knowledge for the benefit of the organization. The central question, then, is this: How do we ensure AI literacy, and how do we build pipelines for a new kind of management?

### Workforce Unknown-Unknowns

This category inherently defies the ability to create a list. It raises, however, two important questions: how do we build a culture of curiosity and discovery that compels us to confront the questions we have never thought to ask? And how do we prepare to face the unknown? Our ability to respond effectively to the unexpected will depend on our commitment to closing the gap between what we know and what we do not know.



Since there are things that we do not know that we do not know, the best way to close this gap is to focus on what we know, what we understand, and what we can address (see diagram above). Failure to do so will widen the gap and undermine our ability to charter the unknown. This focus will result in articulating a framework capable of carrying us successfully into the unknown.

### Creating a Workforce Framework for an AI-Centric Ecosystem

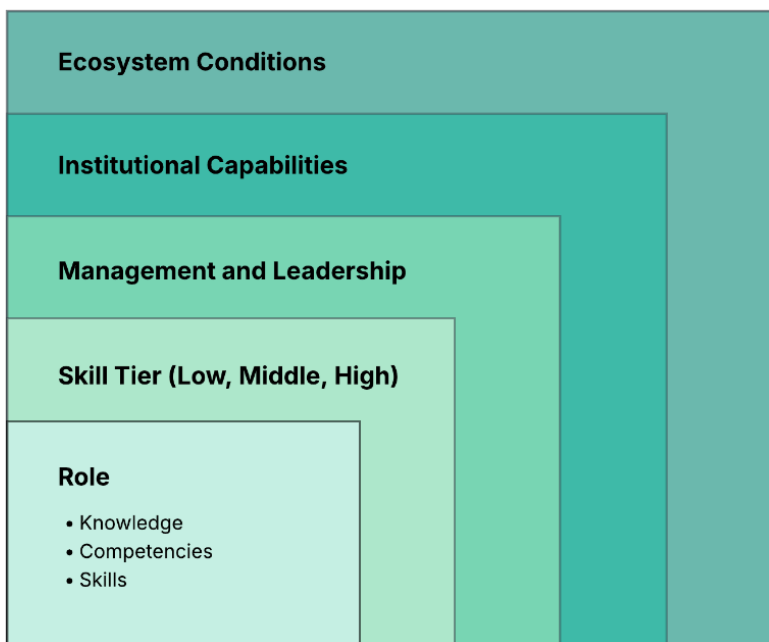
Preparing the workforce for AI cannot be limited to creating technical specialists. It entails establishing a framework that allows AI systems to be built, governed, adopted, scaled, and improved responsibly. This means the framework should be able to support the development of a workforce to support three components at once:

- Capability: the ability to design, deploy, and use AI effectively.
- Governance: the ability to guide AI responsibly, fairly, and lawfully.
- Resilience: the ability to adapt when AI creates new roles, new risks, and new unknowns.

Under this framework, the purpose of workforce development goes beyond simply training people to use AI tools. We must create an ecosystem of human talent, institutions, incentives, and governance structures that allows AI to strengthen productivity, human judgment, innovation, and the common good.

### Schematizing the Framework

An AI workforce framework (see the figure below) should be designed to identify the knowledge, competencies, and skills required to build, govern, adopt, scale, and improve AI systems responsibly. Roles consist of discrete combinations of knowledge, competencies, and skills applied to the focuses of an organization’s mission. At the knowledge level, the framework should include AI literacy, data literacy, quantitative reasoning, organizational understanding, and ethical and legal awareness. At the competency level, it should develop the ability to collaborate with AI, evaluate outputs critically, govern risks, redesign workflows, and adapt to changing conditions. At the skill level, it should prepare workers to use AI responsibly, integrate it into operations, and create value from it. This preparation may occur at many levels of formal training and education (skill tiers), including low, middle, and high-skill.



The framework must also cultivate a new kind of management leadership capable of aligning AI with organizational purposes, surfacing tacit knowledge, confronting bias and ideology, and

building trust in responsible experimentation. This gives the institution the capabilities required to benefit from AI. The goal is to create an ecosystem in which institutions, incentives, governance structures, and talent pipelines work together to ensure that AI advances productivity, resilience, innovation, and the common good.

### Roles

An ideal AI workforce ecosystem requires more than data scientists and engineers. It requires a portfolio of roles including:

1. Builders - the people who design and develop AI systems.
  - A. AI engineers
  - B. Data scientists
  - C. ML researchers
  - D. Data engineers
  - E. Infrastructure specialists
2. Translators - the people who connect technical systems to organizational needs.
  - A. AI product managers
  - B. Domain translators
  - C. Process redesign specialists
  - D. Business analysts
  - E. AI implementation leads
3. Governors - the people who ensure responsible and lawful use.
  - A. AI governance officers
  - B. Risk and compliance specialists
  - C. Ethics and fairness reviewers
  - D. Privacy and legal experts
  - E. Audit and assurance teams
4. Operators - the people who use AI within workflows.
  - A. Frontline workers using AI tools
  - B. Analysts using AI to support decisions
  - C. Managers overseeing AI-supported processes
  - D. Knowledge workers augmenting productivity
5. Integrators - the people who make AI operate at scale.
  - A. Enterprise architects
  - B. Change management leaders
  - C. IT modernization teams
  - D. Training and capability development teams
6. Discoverers - the leaders and innovators who surface unknowns.
  - A. Strategy teams
  - B. Research and innovation units
  - C. Foresight and scenario planners

## D. Organizational learning leaders

### Knowledge

Across roles, and in occupation-specific contexts, the workforce must possess foundational AI knowledge to function in an AI-centric society. Components include:

1. AI and data knowledge
  - A. What AI is and is not
  - B. The four pillars of AI: data, learning algorithms, computational scale, and agentic coordination
  - C. How reality can be represented as data (datafication)
  - D. How models learn from data
  - E. The difference between automation, augmentation, and agency
  - F. Data quality, origin, labeling, structure, and lifecycle management
2. Quantitative and technical knowledge
  - A. Basic mathematics, statistics, and probability
  - B. Computational thinking
  - C. Data analysis and interpretation
  - D. Model limitations, uncertainty, and error
  - E. System integration and digital infrastructure
3. Organizational and operational knowledge
  - A. How workflows operate in practice
  - B. Process design and redesign
  - C. Institutional goals, incentives, and mission alignment
  - D. Tacit knowledge and domain expertise
  - E. Change management and operational scaling
4. Ethical, legal, and civic knowledge
  - A. Bias, fairness, and ideology in AI
  - B. Privacy, accountability, transparency, and explainability
  - C. Human dignity, agency, and creativity
  - D. Legal and regulatory constraints
  - E. Public trust and responsible innovation
5. Strategic knowledge
  - A. When should AI be used and when it should not
  - B. How AI creates value
  - C. How AI changes labor demand and role design
  - D. How to evaluate long-term capacity rather than short-term hype
  - E. How AI interacts with social inequality, power, and opportunity

## Competencies

Likewise, across roles, knowledge must be applied to an organization's goals. Competency refers to the ability to apply knowledge. Competencies include:

1. AI literacy competency - The ability to understand AI systems well enough to use them appropriately, question them intelligently, and recognize their limits.
2. Data competency - The ability to identify what data is needed, assess its quality, understand its limitations, and use it responsibly.
3. Human-AI collaboration competency - The ability to work with AI as a collaborator to augment judgment and productivity rather than replacing human responsibility.
4. Critical evaluation competency - The ability to examine outputs, detect errors, assess reliability, and distinguish signal from noise.
5. Ethical and governance competency - The ability to identify bias, hidden assumptions, evaluate harms, and make decisions consistent with rules, values, and accountability.
6. Systems competency - The ability to understand how models, people, workflows, incentives, and institutions interact.
7. Adaptation competency - The ability to learn continuously, reskill quickly, and pivot as technologies and job demands change.
8. Discovery competency - The ability to ask better questions, identify latent opportunities, uncover hidden risks, and explore unknowns.

## Skills

Based on role, the workforce must possess skills relating to AI. Skills are the practical capabilities the workforce must develop.

1. Universal workforce skills - These are skills everyone should have
  - A. Basic AI literacy
  - B. Data interpretation
  - C. Prompting and tool use
  - D. Verification and judgment
  - E. Digital communication
  - F. Privacy and security awareness
  - G. Responsible use of AI tools
  - H. Continuous learning habits
2. Intermediate workforce skills - These are especially important for supervisors, analysts, specialists, and operational leaders.
  - A. Workflow redesign
  - B. Human-in-the-loop decision making
  - C. AI-assisted problem solving
  - D. Data documentation and stewardship
  - E. Risk recognition

- F. Scenario planning
  - G. Cross-functional collaboration
  - H. Measurement of outcomes and value
3. Advanced workforce skills - These are needed by technical experts, designers, and senior leaders.
- A. Model development and evaluation
  - B. Fine-tuning and system optimization
  - C. AI architecture and infrastructure design
  - D. Multi-agent orchestration
  - E. AI governance design
  - F. Bias and fairness auditing
  - G. Advanced cybersecurity and privacy engineering
  - H. Strategic workforce planning under uncertainty

#### Knowledge, Competencies, and Skills by Skill Tier (Level of Preparation Required)

AI skills should be built across low-, middle-, and high-skill roles, though at different depths. A complete AI workforce framework should address low-, middle-, and high-skill work separately but within one integrated system.

1. Low-skill tier
  - A. Focus: access, literacy, augmentation, and inclusion
  - B. Knowledge
    - i. Basic AI awareness
    - ii. Digital fluency
    - iii. Understanding of task automation and augmentation
    - iv. Privacy and safe use
  - C. Competencies
    - i. Ability to use AI tools responsibly
    - ii. Ability to interpret outputs
    - iii. Ability to recognize when escalation to a human is needed
  - D. Skills
    - i. Basic prompting
    - ii. AI-assisted task completion
    - iii. Digital workflow participation
    - iv. Error spotting
    - v. Simple data entry and quality awareness
2. Middle-skill tier
  - A. Focus: operational adaptation, workflow redesign, and supervisory capability
  - B. Knowledge
    - i. AI in business processes
    - ii. Data quality and workflow logic

- iii. Human-AI coordination
    - iv. Risk and performance measurement
  - C. Competencies
    - i. Ability to integrate AI into team workflows
    - ii. Ability to supervise AI-assisted work
    - iii. Ability to identify process improvement opportunities
    - iv. Skills
  - D. AI-enabled analysis
    - i. Process redesign
    - ii. Quality assurance
    - iii. Exception handling
    - iv. Team training and adoption support
- 3. High-skill tier
  - A. Focus: strategy, advanced design, governance, and innovation
  - B. Knowledge
    - i. Advanced AI concepts
    - ii. Statistics and computation
    - iii. Governance and accountability
    - iv. Organizational strategy under uncertainty
  - C. Competencies
    - i. Ability to design and govern AI systems
    - ii. Ability to align AI with mission and value
    - iii. Ability to anticipate risks and latent workforce needs
  - D. Skills
    - i. Model evaluation
    - ii. System integration
    - iii. Governance design
    - iv. Strategic planning
    - v. Cross-disciplinary leadership

### Management and Leadership

An effective framework to address the AI workforce places special emphasis on a new kind of management. This is critical. Leaders in an AI-centric ecosystem must possess a distinct mix of knowledge, competencies, and skills.

1. What leaders must know:
  - A. AI capabilities and limitations
  - B. Organizational data realities
  - C. Bias, ideology, and tacit knowledge
  - D. Incentives, power, and institutional behavior
  - E. The difference between “buying” AI and building capability

2. What leaders must be able to do:
  - A. Align AI with mission and operational priorities
  - B. Surface hidden assumptions
  - C. Build trust and reduce fear
  - D. Sponsor responsible experimentation
  - E. Redesign workflows, roles, and accountability structures
  - F. Decide when human judgment must remain central
3. What leaders must avoid:
  - A. Treating AI as a plug-in product
  - B. Demanding unrealistic ROI too quickly
  - C. Ignoring tacit knowledge
  - D. Outsourcing strategic judgment to vendors
  - E. Mistaking pilot success for operational readiness

### Institutional Capabilities

Effective leadership helps to build institutions that have six key capabilities required to benefit from AI:

1. AI literacy capability: the ability of the whole workforce to understand and use AI responsibly.
2. Technical capability: the ability to build, deploy, maintain, and improve AI systems.
3. Governance capability: the ability to ensure lawful, ethical, transparent, and accountable AI use.
4. Operational capability: the ability to integrate AI into workflows and scale it beyond pilots.
5. Adaptive capability: the ability to reskill, redesign roles, and respond to emerging labor demands.
6. Discovery capability: the ability to uncover hidden assumptions, latent opportunities, and unknown risks.

Together, these capabilities create the workforce conditions under which AI can generate sustained value within an organization.

### Ecosystem Conditions

Even a skilled workforce, from entry-level through leadership, in an effective institution will fail if the surrounding ecosystem is weak. Building an ideal workforce ecosystem requires attention within a society to:

1. Governance
  - A. Clear policies for responsible AI use
  - B. Defined accountability
  - C. Risk escalation channels

- D. Audit and review mechanisms
- E. Ethical and legal oversight
- 2. Data foundation
  - A. Clean, accessible, interoperable data
  - B. Data stewardship roles
  - C. Strong metadata and documentation
  - D. Clear ownership and quality standards
- 3. Institutional alignment
  - A. AI efforts tied to mission and operational needs
  - B. Incentives aligned with adoption and learning
  - C. Clear decision rights
  - D. Long-term commitment rather than pilot culture
- 4. Learning infrastructure
  - A. Continuous upskilling and reskilling systems
  - B. Stackable credentials and modular training
  - C. Applied learning opportunities
  - D. Career mobility pathways
- 5. Cultural conditions
  - A. Trust in responsible experimentation
  - B. Reduced stigma around AI use
  - C. Openness to redesigning work
  - D. Curiosity, reflection, and disciplined discovery
- 6. Modern infrastructure
  - A. Digital systems capable of integration
  - B. Scalable computing resources
  - C. Secure deployment environments
  - D. Tools that support operational adoption

### **Linking the Workforce Framework to the Known-Unknown Matrix**

An effective workforce framework in the age of AI should be able to address each quadrant of the matrix.

- 1. Workforce known-knowns
  - a. We know the workforce needs:
    - i. AI literacy
    - ii. data literacy
    - iii. quantitative reasoning
    - iv. governance awareness
    - v. role redesign capacity
    - vi. technical and nontechnical collaboration

- b. Framework response: build universal literacy, technical depth, and governance capability.
- 2. Workforce known-unknowns
  - a. We are aware that we do not yet know:
    - i. the optimal skill mix
    - ii. the full range of emerging occupations
    - iii. how fast job categories will evolve
    - iv. how much augmentation versus automation will occur
  - b. Framework response: build modular credentials, flexible pathways, rapid reskilling systems, and labor market sensing capacity.
- 3. Workforce unknown-knowns
  - a. We are often unaware of our:
    - i. tacit knowledge
    - ii. institutional bias
    - iii. hidden assumptions
    - iv. incentive distortions
    - v. informal workflows
  - b. Framework response: create mechanisms to surface tacit knowledge, involve frontline workers, and embed reflection and review into AI design and deployment.
- 4. Workforce unknown-unknowns
  - a. We cannot fully anticipate:
    - i. emergent behaviors
    - ii. disruptive role shifts
    - iii. new social harms
    - iv. unexpected dependencies
  - b. Framework response: build resilience through experimentation, scenario planning, adaptive governance, and cultures of curiosity and discovery.

### **Outcome: An Effective AI Ecosystem**

An ideal ecosystem for AI to reap its full benefits is one in which all workers have a baseline of AI literacy; managers can translate AI into operational value; leaders can govern AI responsibly; technical experts can build and scale systems effectively; institutions can surface tacit knowledge and hidden bias; education and training systems can pivot as occupations evolve; organizations reward experimentation, learning, and adaptation; and governance, trust, and human dignity remain central.

In such an ecosystem, AI is not treated merely as a technology to purchase. It becomes a capability embedded in people, workflows, institutions, and governance systems.