

Artificial Intelligence

Tacit Knowledge Is Your Next Competitive Moat

Teresa Tung and Philippe Roussiere



Image Credit | ti_to_tito

Success in the age of agentic AI requires building systems that capture and amplify your company's unique expertise

Most companies see AI as a data and technology challenge. Collect more information. Train larger models. Scale automation across the enterprise. While these efforts are essential, companies making the most progress with agentic AI recognize

RELATED ARTICLES

Dorothy Leonard and Sylvia Sensiper, “[The Role of Tacit Knowledge in Group Innovation](#),” *California Management Review* 40, no. 3 (Spring 1998): 112–32.

Herman Vantrappen, “[Designing a Fluid Organization of Humans and AI Agents](#),” *California Management Review Insights*, October 9, 2025.

RELATED TOPICS

[Agentic AI](#)

[Digital Transformation](#)

[Competitive Advantage](#)

[Teams & Collaboration](#)

something deeper: The real differentiator is not the data or even the models, but the “tacit knowledge” embedded in the judgment of their people.

Tacit knowledge includes the reasoning patterns, informal heuristics, situational awareness and nuanced interpretive skills that experts develop over years of experience. It is the understanding that rarely appears in manuals or dashboards but determines how teams make sense of ambiguity and how they distinguish meaningful signals from noise.

Today, as agentic and generative AI become more integral to business, the need to capture, structure and apply tacit knowledge is accelerating. That Baby Boomers are retiring in large numbers—and taking their tacit knowledge with them—is adding further urgency to act.

What this means is that your next competitive moat will come from designing new systems that capture tacit knowledge and make it explicit. Below, we draw on our research and client work to identify five essential actions that can guide companies in putting this hard-earned expertise to work: mapping where tacit knowledge lives; codifying it without flattening its richness; embedding it into semantic structures that AI agents can reason with; helping people and AI work together effectively; and ensuring leaders shape how AI is used across the business. Let's take each action in turn.

Map Tacit Knowledge

The first step in bringing tacit knowledge into AI systems involves understanding where expertise actually resides and how it flows through the company. Because tacit knowledge tends to be passed informally between colleagues, it's difficult to capture with traditional knowledge-management tools. Yet rather than ask teams to document everything they know, some companies are beginning with a simpler step: mapping where critical know-how resides. This involves identifying who holds key expertise, how it flows informally and where the organization is most vulnerable to knowledge loss.

A leading automaker, for example, recently confronted a problem that had been building quietly for years: The company's most experienced engineers carried critical knowledge that wasn't written down anywhere. Much of their judgment was developed through decades of diagnosing ambiguous issues and resolving tradeoffs. However, as retirements mounted, leaders realized that this tacit knowledge—central to both product quality and development speed—was becoming harder to access and even harder to pass on.

To stem that loss, the company built a set of AI agents designed to capture how its master engineers reasoned. These systems pulled from design archives, test results and production logs. But they also absorbed the informal materials that reflected how

problems were actually solved, such as through issue-resolution histories and the discussions that unfolded in collaborative engineering reviews. The agents could summarize prior debates and frame options in ways that mirrored how veteran engineers approached similar decisions. Instead of codifying a checklist, the company created a living model of its institutional judgment.

Thanks to its efforts to map tacit knowledge, less experienced engineers now navigate early assignments with a clearer sense of how the organization interprets complex signals and why past teams made the choices they did. Likewise, cross-functional groups reach alignment faster because the institutional memory that once depended on knowing whom to ask is now available on demand. And with expertise no longer concentrated in a handful of senior staff, the automaker is able to move more quickly while upholding the craftsmanship that distinguishes its vehicles.

Similar approaches can map hidden expertise in other industries. Social-network analysis, for example, could reveal where tacit knowledge circulates—and where it stalls. At a multinational mining firm, say, maintenance workflows might hinge on a few field engineers on whom colleagues around the world rely informally, despite their know-how being absent from formal documentation.

Mapping these “invisible” experts would show not only who holds critical judgment, but also how often teams lean on them and where decision-making slows when they are unavailable. With that visibility, the mining firm could redesign its AI-enabled maintenance tools to reflect how work gets done, thereby ensuring that the tacit knowledge of its most trusted people becomes a shared, scalable resource.

Codify Tacit Knowledge Without Flattening It

Once companies identify where tacit knowledge lives, the next step is to translate it into formats that AI systems can learn from—all without losing, or “flattening,” its richness or nuance. This requires codifying the logic and domain context on which human experts

rely. AI systems can then reason within that framework through close collaboration between experts and AI engineers, while supported by tools like causal models and structured frameworks.

The key here is to preserve the decision-making context. At a chemicals manufacturer, for instance, it's not enough for AI to know that a certain chemical triggers a compliance rule. AI must also understand when and why a human might override that rule due to situational nuance.

The need to preserve context is evident in other settings, too. One global professional-services firm recognized that important judgment was scattered across teams and embedded in years of project files. This made it harder for consultants to understand how earlier decisions were made; it also often forced them to sift through materials that weren't organized for reuse.

To overcome this challenge, the firm's leaders began building a shared knowledge base to make this organizational expertise easier to reach. Governance specialists began by setting rules for what could be included, such as requiring a minimum number of client examples to protect anonymity. Senior reviewers then examined past proposals and deliverables to determine which were dependable enough to use again. Next, other teams worked on capturing the concepts that shaped the firm's work and used AI-assisted tools to place those relationships into a "knowledge graph"—a structured picture of how ideas and information connect in the business. Still other teams clarified where different kinds of data should sit.

However, making tacit knowledge more widely accessible wasn't enough. So the firm also invested in tools to help its consultants use knowledge in real time. For example, it introduced systems that applied governance rules automatically and created straightforward ways for teams to publish and draw from reusable data products. Meantime, the knowledge graph expanded as more of the firm's work was represented. Separately, the firm is building AI agents to interpret its practical rules and apply them to things like writing proposals, initiating project requirements, assessing current capabilities and recommending work to be done and the order in which to do it. Although

these efforts come from different teams and projects, together they show how the firm is developing a knowledge system that keeps context intact without reducing expert judgment to oversimplified rules.

Build the Semantic Layer

The ability to operationalize tacit knowledge ultimately depends on constructing a semantic layer—a structure that enables AI agents to interpret the meaning of data in different business contexts. Building a semantic layer involves identifying the conceptual entities that experts rely on when reasoning, as well as articulating the relationships that connect these entities and encoding the contextual “modifiers” that cause meanings to shift across regions, product lines or customer segments.

In platforms like SAP, ServiceNow and Atlassian, the semantic layer determines how workflows are orchestrated and how agents collaborate across systems. Nevertheless, if a company embeds its tacit knowledge only in isolated tools or one-off agent deployments, it risks missing the broader opportunity. A more durable strategy is to encode that expertise directly into the semantic layer for an organization’s AI systems.

From banking to high tech to consumer goods, many companies are building enterprise knowledge graphs—a common way to construct a semantic layer—that embed human expertise into the connective tissue of their platforms. These graphs inform agent reasoning across domains, allowing for more contextualized use of data, as well as higher reuse of assets and better coordination across the company. For organizations that have already built a semantic layer, new AI use cases are far faster to develop, too. Indeed, such companies were able to roll out applications in weeks, not months, because more than half of their semantic structure could be reused.

A recent effort at a cosmetics giant shows what a semantic layer built on a knowledge graph looks like in practice. Previously, regulatory compliance for the company depended on the judgment of a small circle of specialists, who interpreted complex, region-specific rules and navigated legacy, rule-based software. While the company maintained extensive structured data, the most crucial decisions still relied on tacit knowledge: understanding

ingredient families, edge-case exceptions, synonyms across databases like COSING and PCPC and the contextual nuances that regulators use when rules refer to categories of substances, rather than individual ingredients.

To scale this expertise, the company built a knowledge graph that unified structured data with years of expert insights. The resulting semantic layer enabled agentic AI systems to interpret intricate regulatory conditions—such as cumulative concentration thresholds across ingredient families and context-dependent rules tied to product type, target audience or application area. And whenever the system ran into a case it couldn't interpret with confidence, it sent it to a human expert for review.

The payoff was substantial. Regulatory evaluations scaled from hundreds to over 40,000 per month. For validated rules, the system delivered 100% accuracy and repeatability with assessment latency measured in seconds. Meanwhile, workload for regulatory experts fell by roughly 80%, as simple cases were automated end-to-end through AI-generated and AI-validated rule interpretation. Because human-in-the-loop review was embedded from the start, each new edge case reinforced the graph and models, allowing the system to continually absorb emerging norms without losing expert oversight.

By integrating explicit data with the tacit knowledge long held by specialists into a semantic layer, the cosmetics giant created a system that could reason with the same contextual understanding as its experts—something structured data alone could never achieve. A similar opportunity also exists in business intelligence (BI), which has long frustrated users with reports that surface data but rarely the context needed to act on it.

At a consumer-goods company, for instance, a knowledge graph can connect sales, marketing and supply-chain data in ways that conventional BI tools rarely manage to do. By encoding that understanding, the company makes its analysts' tacit knowledge reusable across teams. AI systems can then surface insights that mirror how the business actually reasons—allowing teams to explore data conversationally and focus on insights that drive better results.

Help People and AI Work Well Together

Once constructed, the semantic layer becomes the backbone for a new mode of work—one where human and machine reasoning are intertwined and mutually reinforcing. AI agents no longer operate as rule followers; they become interpreters, capable of applying nuanced judgment, understanding context and explaining their reasoning in ways that align with how the organization makes decisions. People, in turn, gain tools that surface relevant precedents, clarify rationale and highlight interdependencies they might otherwise miss.

In practice, this allows companies to transform foundational workflows. In analytics, for example, the semantic layer replaces the static dashboards that typically overwhelm users with undifferentiated data. Analysts can instead engage in conversational exploration, asking natural questions that are interpreted through the organization's expert logic. The system does not merely retrieve data; it explains why certain metrics matter, how they relate to broader patterns and which signals carry the greatest weight.

The semantic layer also accelerates decision-making across functions. Teams that once interpreted data differently—because they used different conceptual frames—can now draw from a shared understanding of how the business operates. New employees, rather than spending years absorbing unspoken norms, learn directly from a dynamic reasoning model that reflects institutional practice.

Yet this new human-machine collaboration will invariably stumble if left on auto pilot: Companies must therefore be deliberate about helping people and AI work well together. Executives, for example, need to clearly articulate how AI supports expert decision-making, not replace it. Tools should be built with human-in-the-loop mechanisms from the start. And incentives should evolve to reward not just human performance or AI accuracy, but the quality of their interaction.

This moment is also critical because many organizations face a wave of retirements among experienced workers—the very people whose tacit knowledge AI systems most need to learn from. Rather than seeing this as a loss, companies can use AI to capture and transfer expertise while creating new kinds of roles that blend human judgment with machine

insight. Prior Accenture **research** on “fusion” skills highlights these emerging capabilities—such as knowledge curation, model interpretation and human-AI orchestration—which allow employees to move up the value chain and guide how AI learns and applies its outputs.

Developing these skills starts with rethinking how people and systems collaborate. Imagine a global financial services firm rethinking its underwriting process. Instead of replacing analysts, it invites senior staff to help train the AI by reviewing edge cases and contributing counterexamples. Over time, the system evolves into more than a decision-support tool; it becomes a reasoning partner. Because collaboration is embedded into performance metrics, adoption rises, loan accuracy improves and analysts see the tool as a way to amplify, rather than replace, their judgment.

Ensure Business Leaders Shape How AI Is Applied

As tacit knowledge becomes a core input for strategic advantage, business leaders can no longer merely delegate AI strategy to technology teams or treat knowledge retention as a back-office function. Instead, they must become active champions of knowledge-rich AI ecosystems.

This imperative includes setting a vision that aligns AI investment with human values. Business leaders need to ensure that their companies are building ethical and trustworthy systems, as well as efficient ones. Executives should also model behaviors that prioritize human-AI partnership (such as involving domain experts in system design), instead of framing AI as a substitute for human roles. As part of these efforts, departments should be encouraged to share insights, rather than hoard them. Our **research**, for example, shows that companies with strong executive sponsorship of human-AI collaboration are twice as likely to achieve significant financial returns from their AI investments, compared to companies without such support.

Building this kind of culture also means raising the organization's data literacy. Leaders should help teams become more comfortable interpreting and questioning data, not just using AI tools. As roles expand, people in every function—from marketing to finance to operations—need the confidence to bring their expertise to data-driven decisions.

Embedding these behaviors at scale also requires a different approach to incentives. If teams are rewarded solely on cost savings or speed, they will prioritize automation over judgment. But if leaders reward knowledge sharing, cross-functional collaboration and the quality of decisions made by human-AI teams, the entire company begins to shift toward long-term learning and resilience.

Tacit Knowledge Is Your Next Competitive Moat

Companies that fail to operationalize tacit knowledge will face escalating costs: brittle AI systems, inconsistent decisions, vanishing expertise and stranded investments. Worse, they'll fall into a false sense of AI maturity—one that ignores the knowledge gap between scalable automation and actual business judgment.

On the other hand, companies that embed human insight into the design of their AI stack will gain a compound advantage: They'll make better decisions, adapt faster to uncertainty and build AI systems that truly reflect the way their best people think.

In a world where data is abundant and models are open, tacit knowledge is your next competitive moat. Don't wait until it retires.



Teresa Tung [Follow](#) Teresa Tung is the global lead for data capability at Accenture.



Philippe Roussiere [Follow](#)

Philippe Roussiere is the global lead for innovation and AI at Accenture Research.