

AI Is Turning M&A into a High-Impact Learning Machine

By [Jens Kengelbach](#), [Daniel Friedman](#), [Anant Shivraj](#), Yiran Wang, Ouassima El Bouri, and [Dominik Degen](#)

ARTICLE JUNE 10, 2026 12 MIN READ

AI is transforming every phase of the M&A deal life cycle, but its most important impact is more fundamental. It is shifting dealmaking from an episodic, experience-driven process into a continuous, data-driven learning system in which each transaction informs the next. The competitive edge is no longer just access to information but rather the ability to rapidly capture it, synthesize it, and learn from it systematically.

For M&A teams, this addresses a longstanding constraint. Small teams are expected to run complex, multistage processes—pipeline development, due diligence, negotiation, and integration handoff—typically with fragmented data, heavy advisor dependence, and legacy workflows. Much of the insight generated in one deal rarely carries forward, particularly as professionals move on to other roles and institutional knowledge dissipates. AI changes that equation by structuring information across the life cycle, compressing timelines, surfacing risks earlier, allowing teams to focus on the most critical questions, and, crucially, enabling insight to compound across transactions.

The tools to do this are already available. What remains rare is a coherent approach that connects them into an integrated, self-improving system that links how opportunities are identified with how decisions are made and refined over time.

Here we examine how this shift is already taking hold across the upstream deal life cycle—from building the pipeline to signing the deal—and what it takes to turn it into a compounding advantage. A companion article will address the implications for post-merger integration.

Where AI Already Reshapes Dealmaking

Corporate acquirers, private equity firms, and advisory practices are deploying a range of AI capabilities. Some are mature and widely adopted; others are at an earlier stage but increasingly proven. Together, they define the current state of AI-enabled M&A.

Building and Managing the Deal Pipeline

Teams are using AI to identify, track, and prioritize potential opportunities.

Target Identification and Sourcing. Relationship-driven outreach remains the foundation of proactive deal origination, but AI is reshaping how opportunities are surfaced and prioritized. Leading teams have moved from periodic screening of databases to always-on monitoring, with AI-generated alerts and automated briefing packs replacing annual or quarterly manual sweeps.

Natural-language processing models scan news, filings, patent activity, hiring patterns, and web traffic to identify companies that match strategic criteria. Predictive signals—such as holding periods, competitive pressures, or employee attrition—help in estimating a company’s likelihood to sell. Vector-based company mapping is expanding the universe of potential targets, surfacing nonobvious candidates that traditional filters miss. The result is broader coverage along with a continually refreshed pipeline informed by live market signals.

Relationship Management and Institutionalized Memory. As the top of the funnel expands, the next challenge is retaining what the organization learns over time. In most M&A teams, pipeline data sits in spreadsheets, relationship context lives in individual inboxes, and institutional knowledge leaves with the people who hold it.

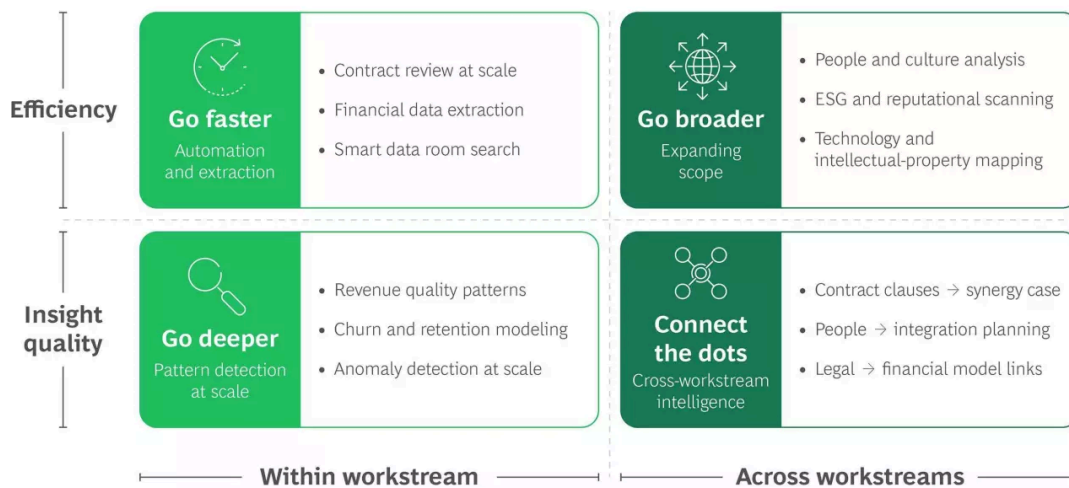
AI-enabled tools are beginning to change this. Relationship intelligence platforms can map networks automatically from emails and calendars, and pipeline scoring models rank opportunities by fit and readiness. Increasingly, companies are building institutionalized memory, capturing past interactions, decisions, and rationales in a form that is searchable and reusable. Adoption remains the constraint: these systems only work if dealmakers contribute to them. But where they do, they lay the foundation for a compounding advantage, with each interaction and decision strengthening the next. Every cycle of sourcing, screening, and outreach feeds signals back into the model, so the pipeline gets smarter with every deal pursued—whether won or lost.

Empowering Deal Execution

AI is shifting due diligence and decision making from fragmented analysis to a more connected, insight-driven system.

Due Diligence. This is where AI adoption is most advanced and where its impact extends beyond efficiency to fundamentally improving how insight is generated and connected. The use of contract review platforms that ingest entire data rooms and extract deal-critical provisions is already standard practice. But the real shift is broader: AI is enabling diligence that is faster, deeper, broader, and more systematically integrated. (See the exhibit.)

AI-Powered Due Diligence Unlocks Value Across Four Dimensions



Source: BCG analysis.
Note: ESG = environmental, social, and governance.

First, it accelerates the basics. AI handles high-volume, repetitive tasks—extracting financial data, normalizing formats, and structuring information—compressing work that once took days into hours.

Even more value comes from increasing analytical depth. By processing entire datasets simultaneously, AI can surface patterns that are difficult to detect within deal timelines. These range from revenue quality issues and churn signals to working-capital anomalies and inconsistencies across financial statements.

AI is also broadening the scope of diligence. Areas that were often deprioritized because of time and resource constraints—such as ESG exposure, reputational risk, people and culture dynamics, and technology and intellectual-property mapping—can now be analyzed without proportionally increasing cost or timeline. The result is a more complete view of the target, rather than one shaped by triage.

Most significant, AI is beginning to connect insights across workstreams. Consider a contractual clause that affects revenue durability, an attrition trend that alters integration risk, or a customer concentration issue that reshapes the synergy case. Traditional processes, where insights sit in separate teams and meetings, often miss these critical linkages or reveal them too late, when a final report is circulated. AI can surface such connections systematically, turning diligence from a set of parallel analyses into a more integrated view of value and risk.

This shift is also pulling risk identification earlier in the process. Rather than waiting for formal diligence, teams can use AI to interrogate available data during preliminary assessments, focusing attention before significant time and resources are committed. The same approach can be applied to the upside: AI-enabled synergy assessment can quickly identify, quantify, benchmark, and prioritize the most material value creation levers before full diligence and inform the post-deal value creation roadmap.

Early adopters emphasize a related discipline: value comes not from applying AI everywhere but from applying it to the right questions. When guided well and supported by high-quality, structured data, AI produces sharper, more actionable insights. (See “AI-Powered Customer Due Diligence.”)

— AI-Powered Customer Analytics in Due Diligence

In traditional due diligence, the acquirer’s view of a target’s customers is largely shaped by what the target chooses to present, such as curated customer satisfaction scores, selective reference calls, and management narratives. While teams may receive churn data and customer lists, pressure testing that story is time consuming and often incomplete. The result is that critical judgments about revenue quality are frequently based on a partial, target-mediated picture.

AI changes this by making customer data analyzable at scale. Where access is available, it can process large volumes of unstructured data—such as reviews, support tickets, and call transcripts—alongside behavioral data like usage, renewal rates, and expansion patterns. Instead of relying on sampled inputs, teams can build a comprehensive, forward-looking view of customer health, identifying early signals of churn, shifts in sentiment, and differences across segments.

The greatest value lies in connecting these insights to the deal model. Revenue concentration in a small number of accounts represents one risk profile if those customers are satisfied and expanding but a very different one if underlying signals point to deterioration. By linking customer sentiment and behavior directly to revenue quality, AI enables a more integrated view of value and risk across commercial and financial diligence.

As with other applications, the depth of insight depends on data access. When detailed customer data is available, the analysis can be highly granular; when it is not, publicly available sources still provide directional signals. In both cases, this shifts how acquirers approach diligence—placing greater emphasis on accessing the underlying data needed to independently assess the durability of revenue.

Valuation, Deal Structuring, and Negotiation. The next step, translating better insights into better decisions, is where many corporate acquirers have historically been disadvantaged. Unlike private equity firms, they typically execute fewer deals, limiting the accumulation of pattern recognition in valuation and negotiation.

AI closes that gap by codifying and scaling pattern recognition. Large language models (LLMs) can parse teasers and confidential information memoranda into structured datasets in minutes. Automated peer group selection combined with AI-enabled financial normalization improves comparability across accounting standards, reporting conventions, and different business mixes. Scenario modeling tools move beyond base, upside, and downside cases to generate probability-weighted business plans and valuation ranges—shifting the analyst role from assembling inputs to exercising judgment on outputs.

In negotiation, adoption is at an earlier stage but is accelerating. Deal term benchmarking—mining databases of completed transactions to show where proposed terms fall relative to market—is among the highest-value applications already available. LLM-powered tools draft and review share purchase agreements and generate first-pass markups that highlight deviations from standard positions, focusing legal effort on genuinely negotiable issues. Earn-out structures can be modeled across a range of scenarios rather than shaped by intuition alone. And increasingly, counterparty intelligence—analyzing a seller’s or an advisor’s prior deal behavior—is informing negotiation strategy with insights that would be impractical to compile manually.

Taken together, these capabilities shift valuation and negotiation from experience-driven processes to more data-informed ones. While they do not replace judgment, they change its foundation by grounding decisions in a broader and more systematically derived set of signals. When captured systematically, those signals become the raw material for sharper screening and valuation of the next deal.

Enabling the Organization

Organizational readiness, not technology availability, is often the limiting factor in capturing these benefits. Three dimensions matter most.

Talent and Ways of Working. AI literacy is quickly becoming a core competency for M&A professionals, not as a standalone skill but as part of how work gets done day to day. Leading teams are starting pragmatically: using AI in routine tasks such as conducting research, summarizing content, and modeling first drafts before embedding it into formal deal processes. Fluency builds through use, not training decks.

At the same time, the role of judgment becomes more important. AI can surface insights and accelerate analysis, but it cannot interpret nuance in a negotiation, make a final investment decision, or persuade a skeptical board. It also cannot assume liability for its outputs. The risk is automation bias: overrelying on outputs without sufficient scrutiny. In an AI-enabled process, accountability still rests firmly with the deal team. M&A will remain a human-in-the-loop discipline, where competitive advantage comes from combining machine-scale analysis with experienced judgment.

Data Quality, Security, and Confidentiality. The effectiveness of AI in M&A depends directly on the quality and accessibility of underlying data, much of which is highly sensitive. Financial models, board materials, and target information are often subject to nondisclosure agreements and regulatory constraints, requiring highly controlled use of AI tools. Enterprise-grade platforms with strong data isolation are improving rapidly, but organizations still need clear guardrails.

M&A teams must work with IT, legal, and compliance to decide what data can be used, in which tools, and under what conditions. The risk is not using AI, but using it informally, without structure or oversight. Teams that establish clear protocols early are better positioned to scale usage safely and consistently.

Build-versus-Buy Decisions. General-purpose AI tools now require little upfront investment. At the same time, specialized M&A platforms—relationship management systems with embedded AI, contract analysis tools, and due diligence workflow platforms—have matured significantly and are often more practical to purchase than to build.

Proprietary development is most beneficial when a company has a genuine data or expertise advantage. A serial acquirer can use its accumulated deal data to train models that steadily improve screening and valuation. Companies that systematically capture structured data from every transaction create proprietary datasets that off-the-shelf tools cannot replicate.

Over time, this advantage compounds. Better data enables better AI outputs; better outputs support better deals; and better deals generate more data. This is where the shift to a learning system becomes tangible. The companies that invest early in capturing and structuring that data will build an advantage that is difficult for others to replicate.

Toward a Self-Improving M&A System

As today's applications improve individual steps in the deal process, companies must prepare for the next phase: connecting and orchestrating them. Leading organizations will turn AI from a set of tools into a system that actively manages workflows, surfaces insights, and continuously learns from outcomes.

Agentic Deal Management. The most immediate shift is from AI that supports analysis to AI that coordinates the process itself. Today, a significant share of dealmakers' time is spent on orchestration rather than judgment: tracking diligence progress, chasing inputs from internal stakeholders and advisors, and assembling materials for decision forums. Emerging agentic AI systems can absorb much of this overhead: managing diligence workflows end to end, routing insights across workstreams, and assembling outputs such as board materials from structured inputs. The effect is not to replace deal teams but to refocus them on the decisions that require experience and judgment.

Agentic Due Diligence. Nowhere is this more visible than in due diligence. The next step beyond AI-assisted review is AI-driven inquiry. These systems can pursue a line of questioning across datasets, linking a contractual risk to revenue exposure, testing whether it is reflected in the financial model, and surfacing the implications for valuation. Human oversight remains critical, but it shifts from every step to key decision points. The result is diligence that is both faster and more exhaustive, because more lines of inquiry can be pursued concurrently.

Signing-to-Close Monitoring. The period between signing and closing is among the least automated stages of the deal life cycle, yet it is where material adverse changes (MACs), regulatory developments, and interim covenant breaches can derail a transaction. AI can replace periodic manual checks with continuous monitoring that evaluates news, regulatory filings, market data, and operational metrics against deal-specific MAC definitions and closing conditions in real time. When a relevant development is detected, the deal team receives an early warning rather than an after-the-fact surprise. This capability is particularly valuable in cross-border transactions, where extended regulatory approval timelines increase exposure to changes.

Integration Handoff. One of the most persistent challenges in M&A is the gap between due diligence findings and integration planning. Due diligence generates large volumes of information across multiple workstreams, but the integration team often inherits only a subset—frequently unstructured and disconnected from the operational decisions required on day one.

AI can bridge this gap by converting due diligence outputs into structured inputs for integration plans or early drafts of plans, with ready-to-execute plans a possibility in the future. It can translate contractual obligations into sequenced task lists with deadlines, turn identified people

risks into targeted retention actions, and link technology compatibility assessments to draft target architectures, potential migration timelines, and initial cost estimates. It can also flag findings that should trigger specific integration steps.

The objective is for the integration team to begin its work with a structured, actionable plan grounded in diligence, rather than relying on a generic template. Done well, this can materially improve the historically uneven track record of post-merger integration and accelerate planning.

Closing the Feedback Loop. The most transformative long-term capability is closing the loop between deal outcomes and upstream decision making. How did realized synergies compare with pre-deal estimates? Which due diligence findings predicted post-close issues, and which proved immaterial? Where did the valuation model diverge from reality and why?

Today, most corporate acquirers capture this information inconsistently, if at all. Teams move on to the next transaction, and the institutional learning that should follow each deal dissipates. AI enables a different model: systematically tracking outcomes against assumptions and feeding those insights back into screening models, assessment frameworks, and valuation approaches.

Over time, this creates a true learning system that improves with every transaction. It is also where the advantages of proprietary data become most tangible. Organizations that close this loop early will compound their advantage with each deal they complete.

The M&A organizations benefiting most from AI will not necessarily be those that adopt the most sophisticated tools first. Rather, they will be those building the systems that enable value to compound. Organizational discipline, high-quality data, and closed feedback loops will determine whether AI remains a set of tools or becomes a source of sustained advantage.

The practical starting point is straightforward: focus on the biggest constraint in the current deal process, apply AI to improve it, and build from there. But the ambition should be broader. Each application should contribute to a system that captures what was learned, connects it to future decisions, and improves with every transaction.

The technology will continue to evolve rapidly. The differentiator will be how deliberately firms integrate it—turning individual use cases into an integrated, self-improving system that makes every deal smarter than the last.

Authors



Jens Kengelbach

Managing Director & Senior Partner; Global Leader of Mergers & Acquisitions
Munich



Daniel Friedman

Managing Director & Senior Partner; Global Leader of Transactions & Integrations
Los Angeles



Anant Shivraj

Managing Director & Partner
Singapore



Yiran Wang

Knowledge Analyst
Munich



Ouassima El Bouri

Senior Analyst - BCG Vantage
Casablanca



Dominik Degen

Senior Director, Transactions & Integrations, BCG Vantage
Munich



ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.

© Boston Consulting Group 2026. All rights reserved.

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow Boston Consulting Group on [Facebook](#) and [X \(formerly Twitter\)](#).