



AI And Machine Learning

4 Steps to Transform the "Middle Office" with AI

by H. James Wilson, Chetna Sehgal, Michael Zimmerman, Ali Arsanjani, Blaise Abderholden and Jimmy Priestas

August 20, 2026



Yaroslav Kushta/Getty Images

Summary. Companies are spending heavily on AI, but many are still struggling to turn that investment into measurable business results. One reason is that they're overlooking the middle office, where work such as contract review, compliance, risk management,... [more](#)

Recent evidence on AI adoption has revealed a sobering paradox. Economists who linked large-scale adoption surveys to administrative payroll records found that although most employers had rolled out AI tools in areas where AI should have the highest probability of impact, the effect on earnings and

hours two years on was statistically indistinguishable from zero. Despite capable technology and workers reporting productivity gains, the promised benefits failed to materialize on the bottom line.


Why? The researchers found that the time the tools saved was eaten by the work of running them. Employees spent their “freed up” time responding to the model’s output and manually wiring it into existing systems and decisions. The gains never surfaced as earnings or reclaimed hours because the models were not truly integrated. The companies lacked the integration that turns them into results.

Based on our proprietary cross-industry analysis and structured experiments in live operations over the past year, we (Accenture and Google) have identified four steps leading companies take to close the gap between adoption and benefit and transform the middle office. This gap between AI adoption and benefit is widest and most costly in the middle office, the largely overlooked domain between so-called front-office customer-facing functions and back-office routine administration. This is where companies handle non-routine, exception-heavy work such as contract

reviews, risk management, and compliance. Decisions often depend heavily on human judgment to resolve ambiguities as business conditions change. And because, according to our analysis, the middle office accounts for more than four-in-10 working-hour tasks across 18 industries, it represents the enterprise's largest single layer of work and one of its biggest untapped sources of AI value.

Yet most AI roadmaps jump from the back office, where automation is mature and the gains are thinning, straight to ambitious front-office bets like dynamic pricing or AI-assisted dealmaking, whose payoff always seems a year away. The middle office, holding the plurality of the work and the nearest large returns, gets passed over.

That's a mistake. The key to closing the adoption-to-benefit gap is improving intelligence integration: the work of turning a capable model into a profitable one by wiring it into how decisions actually get made. The results can be dramatic. For instance, in a live experiment we conducted, improving intelligence integration in a typical middle-office workflow lifted overall success from 42 percent to 80 percent. Resolution of the hardest cases rose from 36 percent to 99 percent, while freeing experts for higher-value work. To help your organization achieve stronger process performance and benefit from humans taking on higher-value work, we recommend the following four steps.



1. Uncover processes that are automated in name only. Such processes run smoothly through easy cases while routing every hard one to your experts. Ask yourself this question: When my AI meets an exception governed by an unwritten rule, does it resolve the case, or does it escalate to a person? If it escalates, you sit at the bottom of the middle-office stage, however much back-office work you have automated.

Through such questioning of processes, you can build your own exception heat map. Score processes by exception density and the cost of escalation, not simply by transaction volume. Target processes where exceptions cluster and each one lands on a senior desk. In our research, we found that nearly a quarter of middle-office decisions rely on reasoning that no policy document fully captures. Those decisions stall AI, and they hold the recoverable value.

2. Pick one process and build a system that learns at runtime.

Resist the pull of the headline front-office project. Instead, first use the criteria in step 1 to choose a single process with these red flags: high exception volume, expensive escalation, and decision logic that lives in people's heads. Then build for learning during operation, not only before it.

The mechanism matters. When our pilot system met an invoice it could not resolve, it captured the surrounding conditions and asked a senior approver a targeted question rather than failing silently or routing the case upward. The expert's answer was distilled into a reusable rule and applied to every similar case afterward. We call this "learning from an example of one." A single correction, rather than the thousands of labeled examples conventional machine learning demands, was enough to generalize an entire class of exceptions.

Consider the effect on one real decision. In one case we studied, two nearly identical invoices arrived for the same purchase order. One was legitimate, citing a valid executive authorization for a parent company to bill on a subsidiary's contract. The other was a near-miss: same amounts, same line items, but its authorization named the wrong entity. Even an experienced reviewer working at volume can miss such discrepancies, and conventional automation, keying on matching fields, often does. After one expert explanation, the system learned to tell the two apart and carried that distinction across thousands of future cases.

The same principle shows up in one of the most credible middle-office results we saw in our analysis. A large Asia-based bank we worked with replaced its rules-based money-laundering monitoring, which generated enormous volumes of false alarms, with a system that learns from its own data. The system now identifies two to four times more genuine financial crime while cutting alert volume by about 60 percent, freeing investigators to work the cases that actually matter. Experts stop being reviewers who inspect output case by case and become teachers who set the unwritten rules once, making good on the value of learning at runtime.

3. Employ metrics that prove the return is real. You cannot know if you have succeeded if you simply measure the raw automation rate. Define success as the exception-resolution rate and the reduction in escalation, and establish the baseline before you deploy so the improvement is provable. Also track the senior staff hours the system gives back, and whether they actually move to higher-value work. Counting only speed measures half of what intelligence integration delivers. And guard against speed that hides a worse decision.

Consider accounts payable, a middle-office process we studied closely in our pilot tests and through industry analysis. A survey

of more than 200 accounts payable executives found that a typical B2B organization takes 17.4 days and spends \$12.88 to approve a single invoice, a bottleneck that can stall vendor relationships and fog up visibility into real-time cash flow. Yet the survey found that best-in-class AI-enabled organizations process the same invoice in 3.1 days at \$2.78, an 82 percent time reduction and a 78 percent cost improvement. Invoice exception rates also fell from 22 percent to 9 percent. Speed and accuracy improved together rather than trading off.

To make improved capability durable, require provenance: every rule should be traceable to the expert exchange that produced it. And test each new rule against a known good case and a known bad case before it goes live. For instance, the bank mentioned above built a money-laundering detection system to produce explainable, auditable risk scores that investigators and regulators could inspect, the kind of transparency that lets a high-stakes process earn trust.

4. Promote your experts into tasks AI can't do. Improved intelligence integration hands hours back to your most expensive people. Decide deliberately where those hours go, or they quietly dissipate into busy work, forfeiting one of the biggest payoffs. That potential recoverable value is significant.

Consider the experience of a major mortgage lender we've been studying. The company takes in over 1.5 million documents a month. Until recently, underwriters spent much of their time sorting and reviewing them. But now, the company's AI platform automatically identifies nearly 70 percent of those documents and extracts the data inside, saving underwriters more than 5,000 hours in a single month, with another 4,000 saved on manual data entry.

But the even bigger payoff came from shifting the attention of the underwriters to much higher-value work. Rather than classifying pay stubs and keying figures from bank statements, they now concentrate on the files that genuinely require a decision: reconciling a self-employed borrower's uneven income, weighing a thin or unusual credit history, resolving the cases a checklist cannot. Formerly, document sorting had consumed their days; but dealing with the difficult files is where their expertise earns its keep.

One European energy company, for example, gave its in-house legal and compliance team an AI assistant over the department's contracts and precedent, built into the tools the lawyers already used, and 96 percent of the department adopted it. Where the lawyers once spent hours combing agreements and first-drafting routine documents, they now hand that work to the assistant and put their time into negotiation strategy and the risk calls that only a lawyer should make.

A recent MIT study found that when automation takes over the routine parts of a job, the work that remains concentrates around harder-to-replace expertise, and its value tends to climb. As software absorbed the rote calculation in bookkeeping, the profession grew leaner but considerably better paid, with bookkeepers' real wages rising nearly 40 percent over recent decades. Freed from exception reviews, your experts move toward the complex, relationship-driven work that AI is least able to do today, and that expertise becomes more valuable, not less. Aim that redeployed expertise at your next, more ambitious project.

...

Leaders frustrated by the lack of returns on their AI investments are looking for answers in the wrong places. The middle office is

far more than the dull step before the exciting one. It holds the largest reachable returns today, and it is where you build the exception-handling capability your front-office ambitions quietly assume you already have.

The path to AI value runs through the middle office. Ignore the messy, exception-heavy workflows of the middle office and your next transformative project inherits the gap rather than the capability. In contrast, focusing on this overlooked territory can result in massive, immediate returns. A winning implementation is the one where the returns actually land, whatever the ambition behind it. And you arrive at your next, more ambitious project standing on capability you built rather than borrowed, with your experts in the lead and their expertise scaling up instead of evaporating.

Stop waiting for the future to arrive in the front office. Build it today in the middle.



H. James Wilson is the global managing director of technology research and thought leadership at Accenture Research. He is the coauthor, with Paul R. Daugherty, of *Human + Machine: Reimagining Work in the Age of AI*, New and Expanded Edition (HBR Press, 2024).

X @hjameswilson

Chetna Sehgal is senior managing director at Accenture.

Michael Zimmerman is in the Google Cloud Office of the CTO and has served as a product executive and CEO in both Fortune-100 companies and start-ups.

Ali Arsanjani is the Director of Applied AI Engineering at Google Cloud, and Head of the Global AI Center of Excellence.

Blaise Abderholden is global Google Cloud AI Lead for partnerships.

Jimmy Priestas is a global AI and Data lead for Digital Core at Accenture.



Read more on **AI and machine learning** or related topics **Generative AI, Digital transformation, IT management, Innovation, Disruptive innovation, Technology and analytics, Algorithms, Analytics and data science, Data management, Performance indicators and Automation**