

The Legal AI Curve from 2023–25: Why Slow and Steady Is Winning the Race

By Brandi Pack, Director of Innovation, UpLevel Ops

For legal professionals and tech leaders alike, the past 18 months have brought a complete reset in how we think about AI. 2023 may have captured headlines as the year generative AI hit the mainstream, but 2024 and 2025 will likely be remembered as the years legal teams held back, cautious, deliberate, and under pressure to make the right call in the face of fast-moving change.

That caution wasn't a sign of resistance. It was a calculated wait-and-see posture, shaped by the challenges of managing risk, defining governance, and assessing unclear returns. Early excitement gave way to harder questions: What does AI really solve for us? Where does it fit in a legal setting? How do we move forward without compromising client trust or team integrity?

As we enter the second half of 2025, answers are starting to emerge. The AI adoption curve in legal has followed a distinct trajectory: first, the *Prompting Craze*, where experimentation was the name of the game and teams raced to learn what these tools could do. Then came *Contextual Customization*, as legal departments began deploying tools trained on internal policies and knowledge bases, building GPT assistants for intake, training, research, and more. Now, we're seeing the onset of the *Agentic Era*, where AI systems are no longer passive helpers but active participants in workflows, filing documents, retrieving data, and triggering follow-up actions without manual intervention.

What's striking about this progression isn't just the speed, it's the shift in accessibility. The adoption strategies taking hold in 2025 aren't reserved for top-tier law firms or innovation labs. They're modular, flexible, and cost-effective enough to fit organizations of all sizes. Return on investment is no longer a future goal. It is happening now. The barriers of uncertainty, cost, and complexity that held teams back last year are starting to fade as the technology improves and best practices take shape.

In addition, we're beginning to see early gains from agents working in coordination, offering new possibilities for tasks that demand nuance and judgment, such as contract analysis and invoice review. The shift from one-off tools to collaborative AI ecosystems marks a turning point in how legal teams approach complexity and opens the door to deeper, more scalable transformation.

The Prompting Craze

As 2024 began, much of the legal tech conversation was still centered on prompts. The early adopters had spent the back half of 2023 learning how to "talk" to tools like ChatGPT and Gemini, treating prompt engineering as both a new literacy and a competitive edge. Legal teams experimented with tone, phrasing, and structure, hoping to extract better summaries, tighter clauses, or clearer explanations.

Some firms even created internal roles or workshops for "AI whisperers" or those who could reliably coax usable output from large language models with just the right combination of words. For a brief period, crafting prompts felt like unlocking a secret code. Vendors leaned into it, offering templates and prompt libraries for tasks ranging from case law summaries to compliance FAQs.

But this phase didn't last long. By spring, the models themselves had advanced to the point that much of the prompt finesse became unnecessary. Tools like GPT-4o and Gemini 1.5 became more

forgiving, capable of interpreting shorthand, vague phrasing, or incomplete instructions with far more accuracy than their predecessors.

That didn't make prompting irrelevant, far from it. Knowing how to structure a clear request remains valuable, especially for more nuanced work. But the real shift was in where legal teams focused their energy. Instead of spending time refining individual prompts, the emphasis turned toward integrating AI into daily workflows in practical, sustainable ways.

Prompting is still a skill, but it's no longer *the only* skill. What matters more now is how strategically and intentionally AI is woven into the rhythm of your work. Used well, it acts as an extension of your skills and training, not a replacement for them. The goal was never to offload your thinking to a machine. It was to bring on an additional brain to work with, not to turn off the one you have.

Contextual Customization

As the year progressed, legal teams began moving past the experimentation stage. The early focus on "talking to the AI" gave way to more foundational questions about how to actually *work with* it. Prompting alone wasn't enough. What professionals wanted, and what the next generation of tools began to deliver, was relevance. Not a clever answer, but a *correct one* grounded in their firm's policies, templates, and workflows.

This marked a clear transition from convenience to context. Legal operations teams began configuring custom GPTs tailored to their internal knowledge bases and processes. General-purpose chatbots, designed for mass use, can only deliver so much customization. Their strength was broad accessibility, but that came with trade-offs. They lacked the precision and adaptability that legal teams require when working with nuanced policies or bespoke workflows. While some embedded bots were genuinely helpful, others proved too rigid to meet evolving legal needs. In contrast, modular assistants built directly on top of base models offered greater control, faster iteration, and better alignment with specific operational goals.

This shift in direction gave rise to more structured AI strategy efforts. IT, InfoSec, and legal operations began aligning around centralized governance. Teams set policies for what data could be used, where models would be hosted, and how usage would be tracked. These efforts became cross-functional initiatives, built with input from security leads, risk officers, and department heads.

As a result, AI adoption began to look less like experimentation and more like implementation. Organizations that had been cautious in early 2024 were now taking confident steps forward. Not because they had all the answers, but because they had a framework for asking the right questions.

This phase was not about launching flashy, large-scale pilots. But it did give way to focused, well-scoped pilot programs that delivered surprising ROI, especially when targeted to specific pain points like intake automation, knowledge access, and internal training. These smaller efforts demonstrated that meaningful results didn't require massive investment or months of development. AI could be configurable, governable, and accountable from day one. In doing so, legal teams began to see tangible returns, not just in time saved, but in credibility earned across the business.

That cautious scoping has proved wise. A recent study from MIT Sloan Management Review found that most corporate AI pilots fail to deliver business value, not because the technology can't work, but because projects launch without clear objectives, governance, or integration plans.

The Agentic Turn

In early 2025, autonomy entered the picture. The conversation shifted from tools that support decisions to systems that make them. Across the tech landscape, early agentic capabilities began to surface. One example came from OpenAI, which introduced Operator, a feature that allowed users to schedule tasks for later execution. It offered a glimpse of what delegated automation could look like in everyday use.

More recently, I experienced the full impact of this shift while using OpenAI's built-in agent mode for the first time. It was one thing to understand the idea in theory, and quite another to witness it in action. I gave ChatGPT a multi-step task to make a dinner reservation for me. It opened its own web browser and began navigating, clicking, searching, and selecting items on my behalf. It was like watching another person use the internet for me, only faster and without distraction. At one point, it encountered a CAPTCHA and paused, flipping control back to me for a quick human review, then seamlessly picked up where it left off. The whole interaction felt surprisingly smooth and, to be honest, a little exhilarating.

That moment captured the essence of the agentic era. AI is stepping out of the role of assistant and into the role of operator. It is not just offering suggestions or preparing drafts. It is taking action. This capability has accelerated workflows, but it will also force a reassessment of risk, governance, and compliance. As AI takes more autonomous steps, the role of humans has shifted up the chain, from executor to overseer. The definition of 'human in the loop' is shifting from micromanaging actions to managing boundaries.

This is where the MIT study's warning becomes urgent: autonomy without strong guardrails is where most AI pilots collapse.

Adapting to the Future

The rise of agents signals a fundamental shift in how legal professionals manage their workflows. This isn't just about smarter tools but about embedding intelligence directly into the way work gets done.

The future belongs to teams that treat AI not as a bolt-on but as a customized layer built around their unique pain points. Because we're still at the start of the AI curve, the pace of change will only accelerate. The smartest strategy is to stay close to the base models and build simple, modular agents that can evolve as the technology does. That's how you protect ROI and avoid the endless cycle of resetting strategy every time the technology advances.

The MIT Sloan research supports this. Their findings show that most corporate AI pilots collapse under the weight of unclear goals and over-engineering. Legal's more measured approach, where pilots that are scoped, governable, and designed to solve real problems, is proving steadier and more effective.

Learn more about the AI Advantage from **UpLevel Ops** at UpLevelOps.com.