

Cracking the Code on AI Contract Redlining

By Brandi Pack, Director of Innovation, UpLevel Ops,
and Sumi Trombley, Senior Advisor, UpLevel Ops

For years, contract redlining has been the holy grail of legal tech: the automation everyone wants and no one seems to get right. Every few months, a new product promises effortless review, automatic markup, or instant risk flagging. Yet anyone who's actually tried one of these tools knows the truth: AI can read a contract, but it still can't redline like a lawyer.

Why Redlining Remains a Human-First Challenge

The challenge isn't that technology lacks sophistication. It's that redlining lives in the messy middle of law, where context, business judgment, and language precision collide. AI can draft, summarize, and analyze. However, the moment it attempts to *negotiate* within a Word document, the seams start to show. Lists renumber themselves, definitions wander off, and formatting chaos takes over.

In other words: The problem isn't just legal complexity, it's structural fragility. And until technology can handle the microstructure of legal documents with the same care lawyers do, every "automated redline" will still need a cleanup crew.

It's Not the Tech. It's the Framework.

The biggest breakthrough didn't come from a shiny new tool. It came from rethinking the *order of operations*. Instead of asking AI to do everything at once, teams are learning that redlining only becomes reliable when the process itself is designed around how lawyers think and how documents behave.

That means building the right foundation: where substance comes before markup, and structure comes before automation. When that alignment exists, AI finally starts to deliver usable, trustworthy results.

Not perfect, but *practical*: edits that don't break formatting, suggestions that match business tone, and changes that hold up under scrutiny.

Rethinking the Role of AI in Redlining

Too many legal teams treat AI as a shortcut. But anyone who's ever led a major contract negotiation knows that a redline isn't just a list of tracked changes; it's a reflection of strategy and risk tolerance.

The real opportunity lies in teaching AI to analyze like a lawyer before it edits like one. When AI learns to evaluate before it executes (such as identifying clauses worth negotiating, not just sentences worth changing), it becomes a genuine collaborator.

Used this way, AI doesn't replace the lawyer; it amplifies the lawyer's intent. It highlights what matters and gets out of the way everywhere else.

What We've Learned

Through testing, piloting, and (yes) plenty of formatting meltdowns, a few principles have emerged:

1. **Structure first, substance second.** Get the document framework right before asking AI to touch the text.

2. **Think like a lawyer, not a machine.** AI should mirror the logic of risk evaluation, not the mechanics of markup.
3. **Human judgment is non-negotiable.** The goal isn't full automation; it's freeing lawyers to focus on judgment calls that actually matter.

Each of these insights reshapes how legal teams approach AI in real workflows, whether for NDAs, MSAs, or procurement contracts. It's not about teaching technology to be perfect. It's about teaching it to be useful.

What This Means for Legal Teams

The lesson is bigger than redlining. It's about how AI fits into work that's inherently relational, contextual, and high-stakes. Most in-house teams don't need more automation; they need automation they can trust.

And trust doesn't come from software features; it comes from strategy. From understanding where human oversight adds value, where structure drives efficiency, and where AI can quietly carry the load.

Because when it comes to redlining, the magic isn't in the markup. It's in making the technology finally think the way lawyers work.

What This Signals for Legal Tech

AI redlining isn't just a product milestone; it's a cultural one. It signals a shift from tools that execute tasks to tools that reflect judgment. The early wins in redlining automation show that legal tech is entering an era where success isn't defined by speed alone, but by *alignment*: alignment with playbooks, tone, and business priorities.

As Generative AI continues to evolve, the same principle applies across the legal ecosystem: contract review, intake triage, knowledge management, and even outside counsel selection. The most successful implementations won't come from buying more technology; they'll come from building the right *operational spine* to support it.

Legal Operations sits squarely at the center of this transformation. Legal ops leaders are now the translators between innovation and adoption, ensuring that AI tools integrate into workflows, deliver measurable ROI, and preserve the integrity of legal judgment along the way.

The teams that thrive aren't the ones chasing every new feature. They're the ones creating conditions where AI can do its best work, guided by structure, governed by process, and grounded in human insight.

So as legal departments begin their own experiments with AI, the question isn't "*Can this tool redline?*" It's "*Can our operations support AI that actually works?*"

When the answer is yes, redlining becomes more than a technical problem solved; it becomes proof that legal AI has finally entered its useful era.

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