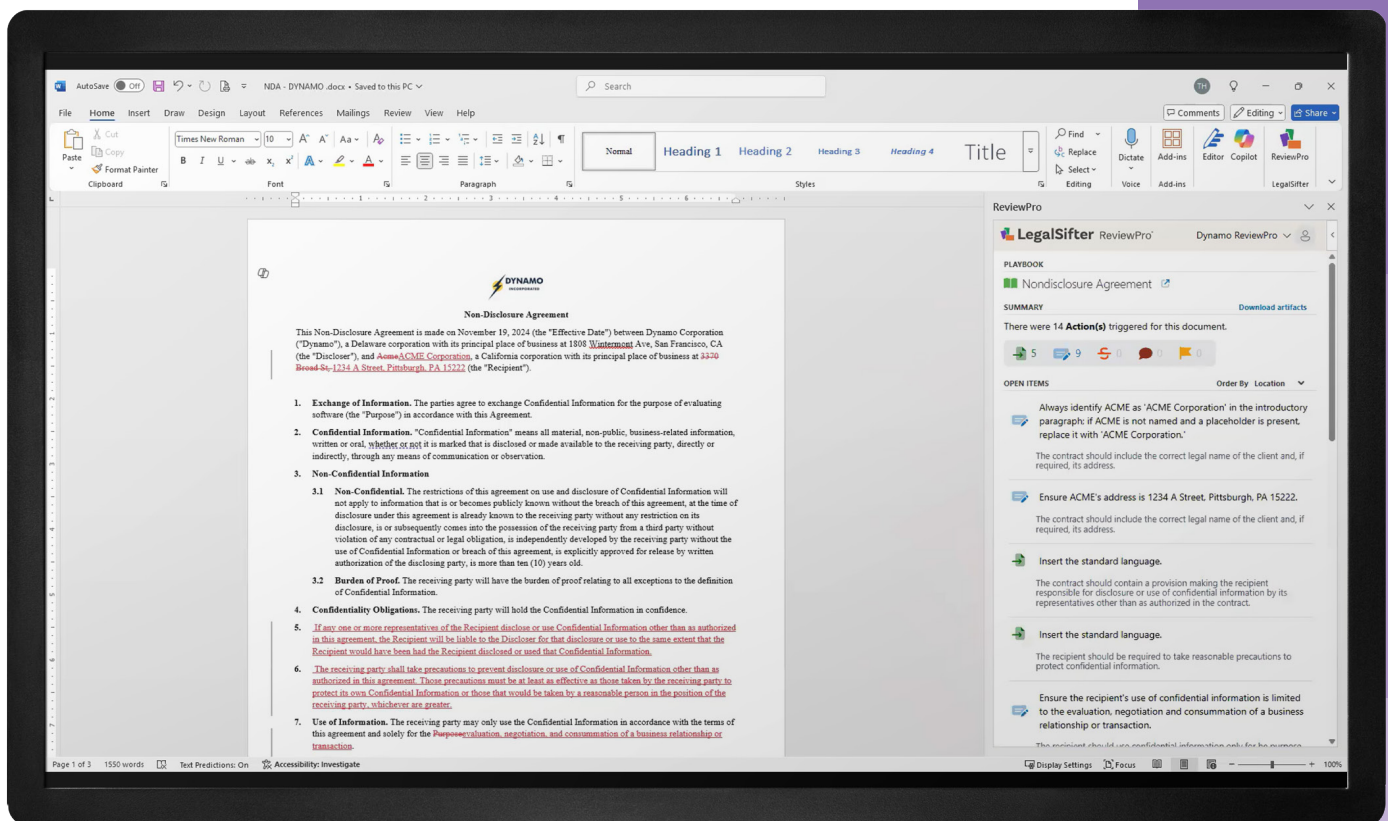


A New Model for AI-Powered Contract Redlining

First-Pass Automation with Human Oversight,
Built for Real-World Legal Workflows



Built on a Decade of Contract Intelligence

The rise of generative AI has made it easier than ever to suggest edits to a contract, but hard to know which edits to trust. Many tools offer flashy demos or conversational interfaces, but lack the rigor, control, and legal reliability that professionals demand.

LegalSifter ReviewPro™ is the result of more than 10 years spent building and honing contract-specific AI. Long before the rise of generative AI, we were applying natural language processing (NLP) to the world of contracts, training models to understand risk, compliance, and negotiation context with legal-grade precision.

At the heart of ReviewPro is our proprietary AI Sifter Engine®, a system of more than 2,200 high-accuracy Sifters, each designed to detect nuanced legal concepts across a wide range of agreements. These Sifters, trained on thousands of real-world

contracts and continually refined by legal experts, provide a deep, battle-tested foundation of domain intelligence.

Now, we've paired that foundation with the power of large language models, within a controlled, auditable framework. The result is ReviewPro: a platform that combines the precision of expert-coded rules with the fluency and flexibility of generative AI. It acts as a true first-pass reviewer, automating the routine while keeping professionals in charge.

This guide outlines:

- The architecture behind ReviewPro.
- The principles that guide our AI design.
- Real-world results from enterprise testing.

Our Philosophy: First-Pass Review, Done Right

WE BELIEVE AI SHOULD REDUCE WORKLOAD, NOT REPLACE JUDGMENT.

ReviewPro is built to handle the repetitive, rules-based aspects of contract redlining, so professionals across legal, procurement, and business operations can focus on the work that requires expertise and strategy. Whether you're ensuring compliance, negotiating risk, or aligning with internal standards, ReviewPro accelerates review without sacrificing oversight.

Unlike generic AI tools that generate open-ended suggestions, ReviewPro performs a structured first-pass review based on curated rules and expert-authored playbooks. It inserts proposed redlines directly into the document, making every change transparent, traceable, and ready for human approval.

This approach delivers:

- **Faster reviews**, by eliminating the manual work of checking for boilerplate, clause presence, and policy alignment
- **Fewer errors**, through consistent application of rules and language
- **Higher-value work**, by freeing teams to focus on judgment, negotiation, and decision-making

All while keeping humans firmly in control of final outcomes.



Our Architecture: 3 Control Layers, Engineered for Trust

Our system architecture reflects a deliberate set of choices, built to meet the quality, control, and auditability requirements of enterprise legal teams. We've designed each layer to minimize risk and maximize consistency.

1. Intent Layer: Structured Playbooks

We don't rely on ad hoc chat prompts or user memory. Instead, ReviewPro uses structured playbooks that define exactly what the AI should look for and how it should behave across each type of contract. This eliminates the guesswork and inconsistency that can come from chat-driven approaches, while also making the system more auditable and repeatable.

→ System control is centralized and deliberate, ensuring that review behavior is predictable, scalable, and consistent across users and documents.

2. Reasoning Layer: Controlled LLM Use

Rather than letting the LLM reason freely or improvise language edits, ReviewPro places tight constraints around the LLM's behavior, ensuring it behaves like a precision tool, not a free-form assistant. Behind the scenes, we've adapted over a decade of contract-specific AI work—including 2,200+ high-accuracy classifiers—to steer the system behavior. This foundation enables predictable, auditable behavior aligned with enterprise review needs.

→ The LLM is steered with precision, reducing hallucinations and off-policy edits while ensuring explainable, accurate results.

3. Execution Layer: Automatic Edits

ReviewPro applies automatic edits to the document, without requiring manual approval for each change. This "0-Click" model turns the AI into a true first-pass reviewer, letting human reviewers focus on higher-level reviews while retaining complete control to reject or modify the AI edits.

→ This reduces manual effort and accelerates review, without giving up control.

Why LegalSifter's Approach Is Different

WE'RE NOT A CHATBOT. WE'RE NOT A GENERIC AI WRAPPER. WE'RE A CONTRACT INTELLIGENCE COMPANY WITH REAL-WORLD EXPERIENCE AND DEEP PRODUCT DISCIPLINE.

Proven Contract-Specific AI Expertise

ReviewPro is built on over a decade of experience developing contract-specific AI. Under the hood, it uses AI Sifter Engine®, our proprietary AI system that includes more than 2,200 AI Sifters, each trained to identify legal concepts with precision, whether present or missing. These models have been trained and tested across thousands of real-world contracts, consistently achieving 95%+ accuracy in production, and they are refined through ongoing feedback from attorneys and legal professionals.

This isn't a proof-of-concept—it's contract AI that's been tested, tuned, and trusted across thousands of real-world negotiations. Clients are already using it to reduce review time, improve consistency, and stay aligned with strategy.

Highly Controlled AI Framework

Building on this foundation, ReviewPro uses a tightly engineered approach to guide the behavior of large language models. We don't rely on open-ended reasoning or improvisational chat. Instead, we control the model with structured prompts, curated examples, and rule-based logic, that build upon and extend our existing AI.

This control framework ensures that the system behaves like a precision tool, not a conversational assistant. Every edit is traceable to a specific instruction. Every rule is executed deliberately. The result is a redlining engine that behaves predictably, avoids hallucinations, and produces edits that are auditable, defensible, and negotiation-ready.

Playbooks Built by Legal Experts

Our redlining engine is powered by expert-built playbooks that encode decades of legal and negotiation experience. These playbooks are created and maintained by expert attorneys, capturing best practices, preferred language, and additional guidance, all tailored to specific contract types. Each playbook reflects not only legal know-how but real-world judgment about risk and negotiation dynamics. This depth of expertise enables ReviewPro to suggest edits that are not only technically correct, but commercially sound.

A Rich Library—and Full Playbook Control

Unlike generic AI tools, ReviewPro comes ready with a growing library of 50+ playbooks that cover common contract types and use cases. These playbooks work out of the box to accelerate adoption and deliver value immediately.

Even better, clients aren't locked into one-size-fits-all content. With ReviewPro, legal teams can create, modify, and govern their own playbooks through an intuitive interface. This flexibility ensures that your AI redlines reflect your policies, your voice, and your standards.



What “Good” Looks Like: Our 5 Quality Standards

WE DON'T JUST BUILD FOR AUTOMATION. WE BUILD FOR QUALITY.

We've designed ReviewPro to address 5 quality attributes that define high-performing AI in legal review. These are grounded in the real needs of legal and procurement teams and based on the industry framework introduced in our Evaluating AI Redline Systems guide.

1. **Steerable.** You control what the AI looks for and how it behaves. ReviewPro supports structured playbooks, adjustable redline strategies, and tailored settings by contract type, so your intent stays in charge.
2. **Thorough.** No skipped steps. No silent failures. The system executes every rule in your playbook and surfaces every instance that matters, across even the most complex contracts.
3. **Accurate.** We don't guess. ReviewPro detects with 95%+ accuracy whether provisions are present, missing, or require edits, and applies edits with precision.
4. **Explainable.** Every edit is tied to a clear reason, whether a playbook rule, fallback position, or user instruction. ReviewPro shows its work, making it easy for you to audit, defend, and refine your playbook.
5. **Fluent.** Language matters. Our edits read like they belong—preserving tone, structure, formatting, and legal polish. You get automation without sacrificing professionalism.

These 5 benchmarks shape how we train, test, and deliver the ReviewPro experience. It's not just about what the AI can do, it's about how reliably and professionally it does it.

AI redlining is no longer hypothetical. It's here—and it works—when built with the right architecture, data, and quality standards.

LegalSifter ReviewPro is built for serious review work:

- **Structured Playbooks** ensure consistent, policy-aligned behavior
- **Controlled Reasoning** delivers explainable, legal-grade precision
- **Automatic Edits** speed up review times without sacrificing oversight

LegalSifter ReviewPro can automate the routine and reclaim your time for what matters.

Ready to see it in action?

Request a Demo

Talk with Our Team

Start a Free Trial

legalsifter.com | (724) 221-7438