

USING FRAM TO MODEL AND IMPROVE AI-HUMAN INTERACTIONS IN LEGAL CONTRACT CHECKING

David Slater, Cardiff University, UK, dslater@cambrebnsis.org and

John Kunzler, Marsh Specialty UK, Marsh Ltd., John.Kunzler@marsh.com

Abstract

The integration of Artificial Intelligence (AI) into legal contract review processes promises faster, more consistent detection of errors and risks, yet introduces new complexities that traditional workflow models cannot adequately capture. This report applies the Functional Resonance Analysis Method (FRAM) to reframe AI-assisted contract review as a dynamic socio-technical system characterized by interdependent human and AI functions, each subject to performance variability. Drawing on real-world case studies, scenario simulations, and system modeling, the report demonstrates how minor fluctuations in AI reliability, human judgement, and task conditions can interact to produce resonant error patterns. By constructing a detailed FRAM model of contract review workflows, we identify critical pathways where variability amplifies risk and propose resilient system designs that maintain human interpretive authority. The findings underscore that optimizing AI-human collaboration requires not only better tools, but systemic redesign grounded in complexity science principles. FRAM provides a structured, predictive approach to making AI integration in legal practice safer, smarter, and more accountable.

Keywords

- Functional Resonance Analysis Method (FRAM)
- Legal Technology
- Artificial Intelligence in Law
- Human-AI Collaboration
- Contract Review Reliability
- Socio-Technical Systems
- Risk Modeling and Resilience Engineering

Introduction

The review of legal contracts for risks, errors, inconsistencies, and omissions has long been a critical pillar of professional legal services. Traditionally reliant on meticulous human expertise, this painstaking work demands high levels of concentration, deep domain knowledge, and contextual understanding of client needs and regulatory frameworks. Although legal contract review can be conducted by many businesses themselves, law firms generate value from the level of assurance they bring to such tasks by providing high reliability, and heavily insured and guaranteed recourse in the event of an error. Over recent years, however, the rise of Artificial Intelligence (AI) systems, particularly those based on large language models (LLMs), has begun to reshape this landscape. LegalTech companies, internal law firm innovations, and broad technology integrations are increasingly supplementing human effort with machine-based assistance, promising faster, cheaper, and in some cases more consistent contract review.

Yet alongside these promises come new, underappreciated challenges. AI is often treated as a static tool within a human-controlled process, but in reality, it acts as a dynamic function, introducing its own performance variability, interaction patterns, and emergent behaviours to those inherent in the work as done. In addition, and perhaps as a consequence of this lack of

understanding, failure (failing to deliver planned benefits) rates in relation to AI projects reportedly run higher (85%)(1) than the general rate indicated in separate studies (70%)(2)

The State of AI-Assisted Contract Review: Promise and Pitfalls

Across the legal industry, AI adoption is accelerating. Specialized platforms like Kira Systems, Luminance, and LawGeex offer powerful contract analysis capabilities, flagging missing clauses, risky language, and compliance gaps. At the same time, major players like Microsoft are integrating LLMs into everyday drafting tools (e.g., Copilot in Word), while firms like Allen & Overy have launched internal AI tools (e.g., Harvey) based on fine-tuned GPT models.

Empirical case studies highlight both the strengths and the critical limitations of current AI deployments. The LawGeex 2018 study (3), famously showed AI achieving 94% accuracy against human lawyers' 85% when reviewing NDAs. Kira Systems helped uncover hidden change-of-control clauses during a high-stakes M&A deal, averting major legal liabilities. Luminance assisted law firms in identifying GDPR compliance risks across thousands of legacy contracts.

However, the dangers of over-reliance are equally stark. In one cautionary case, a corporate legal department trusted an AI-suggested liability clause without senior review, leading to devastating financial losses after a supplier dispute. Even Allen & Overy's Harvey platform, despite its sophistication, carries a disclaimer that it has occasionally "hallucinated" unnecessary clauses based on generalized best practices that were contextually inappropriate.

The picture is clear: AI can greatly assist with detection and initial screening, but nuanced legal judgement remains an inherently human function. The risks lie not in AI alone, but in the ways human and AI variability interact across the contract review system.

Safe Workflow for Using AI in Contract Review

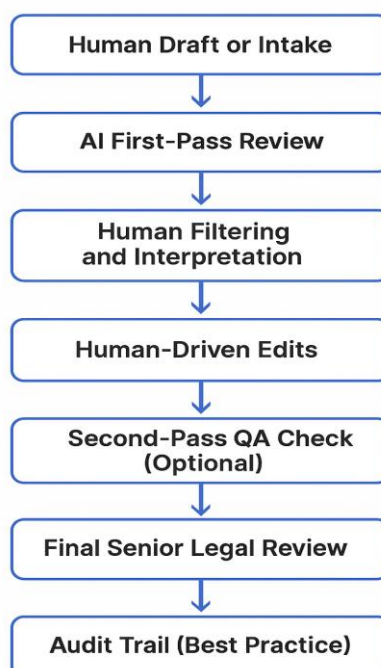


Figure 1 Suggested “safe” workflow

The paper *Identifying AI Hazards and Responsibility Gaps* by Mary L. Cummings (4), highlights that traditional accident models focused on human error are insufficient for analyzing failures involving AI systems, particularly in self-driving vehicles. To address this, Cummings proposes the **Taxonomy for AI Hazard Analysis (TAIHA)**, introducing layers that capture failures in AI oversight, design, maintenance, and testing. Through case studies, the paper shows how poor corporate safety cultures, inadequate testing, model drift, and regulatory indifference are leading to accidents and widening responsibility gaps. The analysis argues that responsibility gaps emerge from deliberate human choices rather than technological inevitability.

Traditional linear workflow models fail to account for the complex ways human and AI activities interrelate. It is here that the Functional Resonance Analysis Method (FRAM) (5), offers a breakthrough perspective: allowing us to model the entire human-AI work system as a living, variable, and potentially resonant network. This report explores how FRAM can rigorously map, predict, and guide the safer integration of AI tools into legal contract checking.

Rethinking Contract Review as a Complex Socio-Technical System

Legal contract review, particularly when AI tools are involved, is not a linear assembly line. It is a dynamic, complex, adaptive system. Every actor—human or machine—introduces variability. These variations interact, sometimes cancelling each other out, sometimes amplifying into major failures.

Most current approaches to integrating AI in contract review assume a linear augmentation model, treating AI as a static tool that reduces human error or workload. However, this perspective ignores the dynamic and emergent nature of socio-technical systems. In reality, AI acts as an operational function within a broader human-AI work system, introducing its own variability, couplings, and dependencies.

Today, most discussions about **AI in contract review** treat AI like a *perfect assistant* or just another “tool” sitting outside the system. But **FRAM shows us** that **any participant (human or AI)** is a **function** inside a bigger **dynamic work system** — with Inputs, Outputs, Preconditions, Controls, Resources, and Timing aspects.

And crucially, **variability** (performance fluctuations) happens in both humans **and** AI.

Thus:

- AI is *not* a perfect “spellchecker” — it’s another actor whose reliability can vary depending on conditions.
- Human reviewers and AI interact in **coupled functions** — sometimes reinforcing each other’s strengths, sometimes amplifying each other’s weaknesses (resonance effects).
- Errors (missed amendments or ineffective drafting issues) arise *not* because AI or humans individually “fail” — but because the **whole system** (humans + AI + process design) resonates in an unexpected way. One function fluctuating significantly is often linked to other functions all showing increased performance variability, sometimes with the result that **none** of them function as expected.

The Functional Resonance Analysis Method (FRAM) offers a structured approach to model such complex systems. Originally developed in resilience engineering to understand how variability in everyday performance can lead to unexpected outcomes, FRAM models systems as a network of interdependent functions characterized by six aspects: Input, Output, Preconditions, Resources, Control, and Time.

Applying FRAM to AI-assisted contract review enables us to:

- Treat AI and human actors equally as dynamic functions.
- Identify how variability propagates across functions.
- Recognize resonance effects, where small variations compound into major failures.
- Model and test different scenarios to predict system behavior under varying conditions.
- Design interventions that improve systemic resilience rather than just optimizing isolated tasks.

Bringing FRAM modeling into the analysis of AI-assisted contract review would let us see the systemic couplings between AI and human work.

- It would help predict when "resonance" (errors compounding) is likely to happen — for example, under conditions like tight deadlines, AI overtrust, lack of senior oversight.
- It would show where to insert stronger controls (e.g., enforce human interpretive filtering between AI output and final editing).

FRAM offers a methodology to model this complexity realistically. Rather than viewing tasks as sequential steps, FRAM sees work as a set of interdependent functions, each defined by:

- **Inputs** (what is needed to begin)
- **Outputs** (what is produced)
- **Preconditions** (what must be in place)
- **Resources** (what is needed to perform the function)
- **Controls** (what regulates the function)
- **Time constraints**

Importantly, each function can vary in how it is performed, depending on local conditions such as fatigue, time pressure, tool performance, and organizational policies.

Applying FRAM to AI-assisted contract review means recognizing that AI models are not just passive tools but active functions with their own Inputs, Outputs, Preconditions, Resources, Controls, and Timing dynamics.

Building the FRAM Model: Functions, Variability, and Couplings

In a typical AI-assisted contract review workflow, the major functions might be modelled as:

- **Human Draft Contract:** Taking client requirements and translating them into a first draft under the guidance of templates and legal standards.
- **AI First-Pass Review:** Running the draft through an AI system trained to flag missing clauses, inconsistencies, and risky language.
- **Human Filter AI Suggestions:** Evaluating AI outputs critically, deciding which suggestions to accept, modify, or reject.
- **Human Revise Contract:** Editing the document based on filtered AI suggestions and human judgement. Draft documentation shared with counterparty and proposed

updated versions (Travelling draft) discussed and agreed by client sometimes passing through several cycles of change until agreed in principle.

- **Human Final Review and Approval:** Senior lawyers conducting a comprehensive review, focusing on business risk, client context, and negotiation sensitivity.

Each function experiences variability. For example, AI performance can fluctuate based on model updates, prompt designs, or data drift. Human performance varies with fatigue, time constraints, expertise, and trust levels. Task characteristics (e.g., complexity, ambiguity) introduce further variability.

Critically, these variabilities can interact. Minor issues unflagged by AI may pass unnoticed if human reviewers are rushed. AI hallucinations may be mistakenly accepted if human reviewers overtrust machine outputs. Cumulative small variations can resonate, compounding into major contract errors.

Quantifying System Risk: Scenario Simulations

By assigning performance probabilities to each function—such as the likelihood of an AI reviewer missing a critical clause, or a human filterer overlooking a problematic suggestion—we can simulate overall system error rates under different scenarios.⁽⁶⁾

Ideal conditions (well-trained humans, calibrated AI, ample time) produce low estimated system error probabilities (~8%).

Under high-stress conditions (tight deadlines, AI model drift, fatigued human reviewers), system error probabilities spike dramatically (~26.1%).

With targeted mitigations (extra senior review, stricter AI filtering guidelines, peer audits), error probabilities can be brought down again (~12%).

These simulations vividly demonstrate that risk emerges from system dynamics, not individual incompetence. It is the coupling and resonance of performance variabilities across human and AI functions that determine overall reliability.

Redesigning Contract Review Workflows: FRAM-Informed Best Practices

Instead of simplistic linear workflows, FRAM suggests a more resilient architecture:

Contract review should be structured as an iterative, critically filtered, and redundantly safeguarded process. AI tools should serve as high-speed anomaly detectors, not autonomous drafters. Human judgement must dominate all substantive decisions.

Workflows must ensure:

- **Structured human interpretation of AI outputs** (not blind acceptance)
- **Manual incorporation of changes** following legal risk assessment and client input
- **Independent final reviews** focused on client-specific risk profiles
- **Audit trails** recording AI contributions, human decisions, and review outcomes

Contextual judgement, negotiation sensitivity, and strategic client advice remain the exclusive domain of experienced human lawyers. No AI, no matter how advanced, can currently replicate this dimension of legal work.

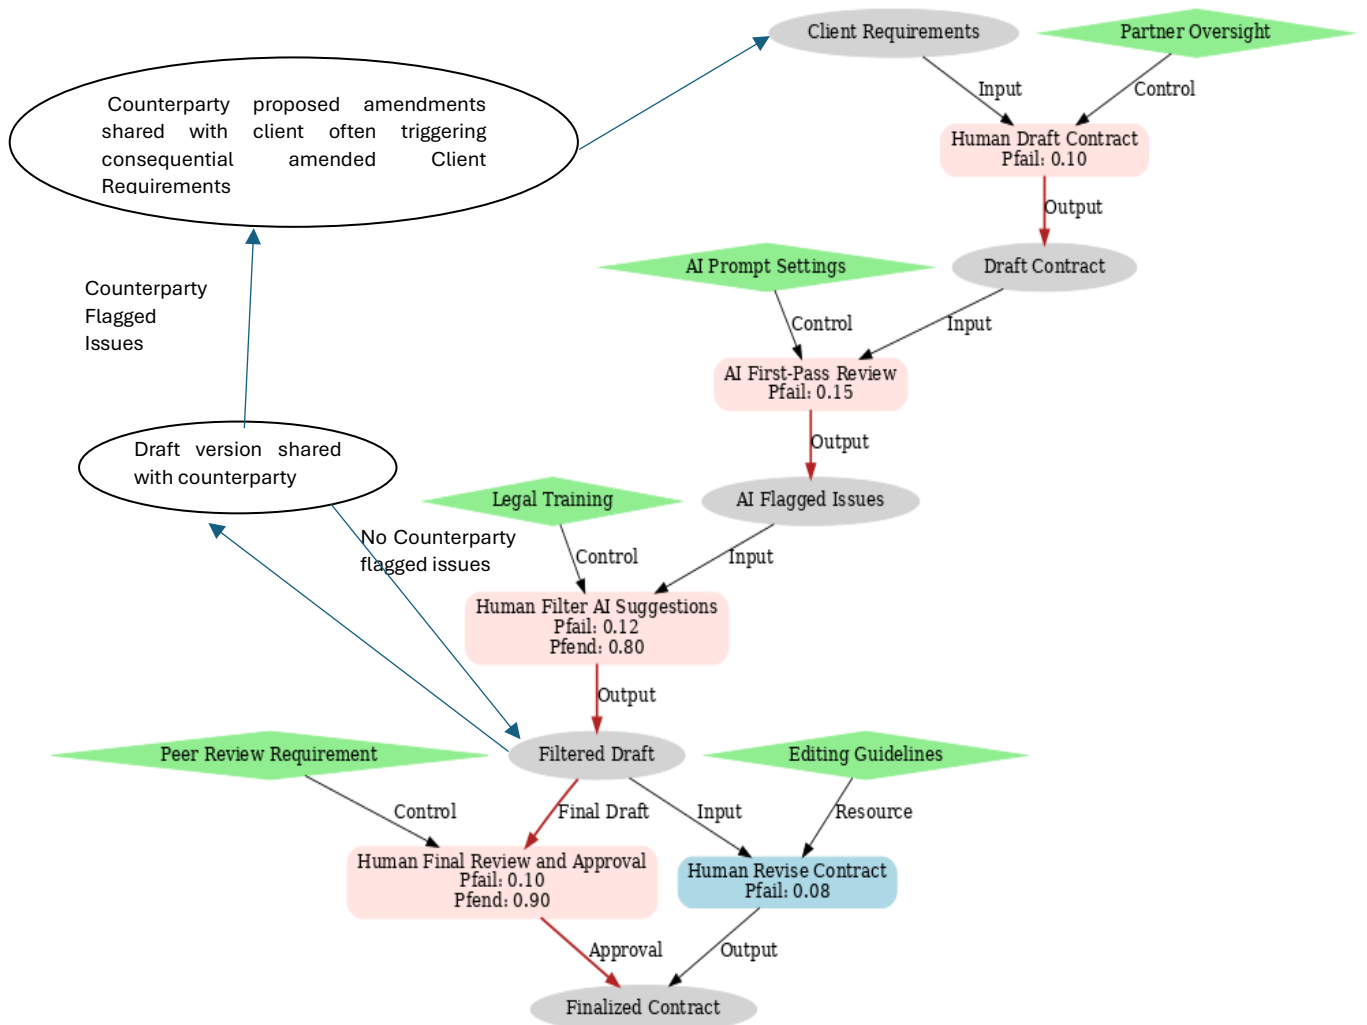


Figure 2 – Notional Functional Model

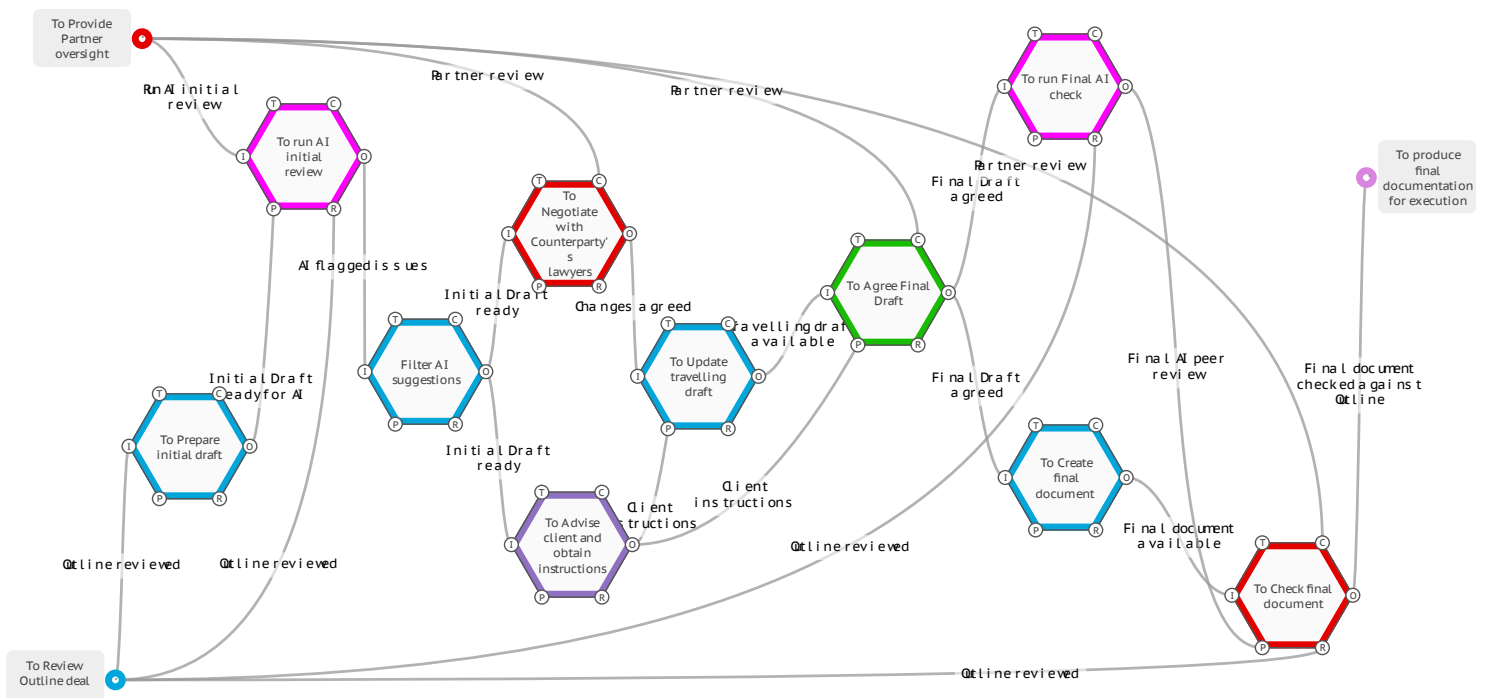


Figure 3 – Equivalent FRAM model

Conclusion: Toward Safer, Smarter Integration of AI in Legal Practice

The future of contract review will be increasingly AI-assisted, but it must also be increasingly system-aware. FRAM modeling shows that legal practices cannot treat AI as a black box adjunct to human workflows. Instead, through the FRAM lens AI must be seen as a dynamic function within a socio-technical system, whose reliability depends on how its outputs are interpreted, filtered, and integrated by human professionals.

By adopting FRAM-informed design principles—explicitly modeling functions, mapping variability, predicting resonance risks, and engineering resilient workflows—the legal industry can move beyond simplistic automation of the current “what is” toward truly intelligent augmentation. The opportunity to reduce the variability beyond current levels is tantalisingly close but requires design informed by these principles. AI can become a powerful ally in enhancing legal services—but only if we design, predict, and manage the human-AI interaction system as a whole.

The opportunity is enormous. So too is the responsibility to get it right.

References

1. Project Management Statistics: Trends and Common Mistakes in 2024, <https://teamstage.io/project-management-statistics>
2. Rayome, A.D, (2019), Why 85% of AI projects fail, <https://www.techrepublic.com/article/why-85-of-ai-projects-fail/>
3. LawGeex, (2018), “Comparing the Performance of Artificial Intelligence to Human Lawyers in the Review of Standard Business Contracts”, <https://images.law.com/contrib/content/uploads/documents/397/5408/lawgeex.pdf>
4. M. L. Cummings, (2025), "Identifying AI Hazards and Responsibility Gaps," in *IEEE Access*, vol. 13, pp. 54338-54349, 2025, doi: 10.1109/ACCESS.2025.3552200.
5. Hollnagel, E., (2012), *FRAM: The Functional Resonance Analysis Method: Modelling Complex Socio-technical Systems*, CRC Press.
6. Kunzler, J., and Slater, D. (2025), Testing the applicability of FRAM for better understanding and improving the performance and design of legal contract preparation, International FRAM Workshop, FRAMily 2025 Delft.