

Dealing with AI errors in legal practice

A governance model for leveraging the benefits while mitigating the risks of the use of AI by lawyers

“A danger known is a danger avoided”

Dr. Adrian Schaub
March 2026

Why This Matters

AI adoption in legal practice is accelerating. So is overconfidence.

This guide addresses the structural risk most lawyers underestimate: not that AI fails — but that it fails **without warning**, even when output looks right.

- **Objective 1:** Understand the structural limits of generative AI
- **Objective 2:** Implement a risk-based governance model for defensible AI use

-> scope is limited to dealing with AI errors (not handling of other issues like confidentiality, Intellectual Property or Data Privacy)

The governance model presented is consistent with the [Formal Opinion 512 of the American Bar Association on Generative AI Tools \(July 29, 2024\)](#)

Do Not Humanize AI

- We trust a junior lawyer more as we observe consistent, improving performance
- We apply the same logic to AI — and reduce our review as outputs impress us
- But: AI does not accumulate reliable competence: each run is a new probabilistic generation

Key insight: Fluency of output ≠ legal correctness

Example: AI drafts a beautifully structured clause analysis — and invents the case law it cites.

Why the Human Trust Model Fails

- AI generates output based on statistical likelihood — not verified legal reasoning
- Outputs are plausible, not validated
- Each run is a new statistical generation: past performance does not reduce future error risk
- Errors are non-systematic and can occur in core and peripheral issues alike
- Confident tone does not signal correctness — AI always looks polished, even when wrong

This is structural — it applies equally to free or premium models, including special legal AI tools.

Predictable Errors

Example: AI comparing a supplier contract against your template

- Missed references and incorporation risks (internal and external)
- Failure to detect consequences of missing clauses (e.g. unlimited liability under Swiss CO)

These errors are generally detectable through targeted checklists — manageable with structured review.

Unpredictable Errors

Hallucinations in legal work go far beyond “making things up”:

- Fabricated content — invented case law, non-existent statutory clauses
- False positives or false negatives — missed or phantom deviations in contract review
- Logical slips — correct analysis followed by an unsupported legal conclusion
- Material omissions — failure to identify a critical gap

Critical consequence: errors are not pattern-based — structured spot checks are not a reliable control mechanism.

Model Drift

AI tools change over time: new model versions, updated safety layers, new training data, changed enterprise configurations

- Same prompt may yield different outputs at different times — without visible notice
- Prior reliability does not guarantee future reliability

A workflow reliable in January may produce different results in June. Re-validate critical prompts periodically.

Governance model: Risk-Based Human Review

Key question: Is an elementary mistake in a material aspect tolerable?

- **Not tolerable** → **Full human verification** (client advice, court filings, high-value contracts)
- **Limited tolerance** → **Validate core aspects** (internal drafts, negotiation prep, early analysis)
- **Low impact** → **Cursory review** (brainstorming, formatting, early idea generation)

-> Make review depth a conscious, deliberate choice

Checklist available at www.adrianschaub.com/ai/

Note: ABA Formal Opinion 512 (2024) confirms the risk-based review standard. Note: prior validated experience with a tool may reduce verification requirements for narrow, repeatable tasks — but does not reduce the risk of non-systematic errors in new or complex matters.

Client Transparency

- Lawyers remain fully accountable — regardless of AI involvement
- **Full human verification is the rule for any document sent to a client or a court (duty of competent representation)** -> the use of AI must not be disclosed.

However, when full human verification was not performed (e.g. due to time and cost constraints imposed by the client) Make the residual risk explicit to the client and obtain acceptance of

- Scope of AI involvement and extent of human validation
- Residual risk and worst-case consequence
- Time and cost of full human verification as alternative

Document acceptance — but prepare documentation consistently with privilege principles and with awareness of potential litigation use

Full checklist available at www.adrianschaub.com/ai/

Note: Framework consistent with ABA Formal Opinion 512 (July 2024) which however further requires transparency when AI output materially influences a significant decision in the representation (regardless of review depth).

Bottom Line

- AI materially increases legal productivity — the risks identified here are no reason to avoid it
- Elementary errors can occur in core areas even in the most sophisticated AI models
- Errors follow no pattern — they cannot be reliably detected by structured spot checks

Responsible AI use = a risk-based approach calibrated to the tolerability of error in material aspects.

“A danger known is a danger avoided.”

Feedback

Don't hesitate to send any questions or feedback to adrian.schaub@gmx.ch

This slide deck, checklists and other AI related materials can be downloaded from www.adrianschaub.com/ai/