PROGRAM MATERIALS
Program #3676
April 29, 2026

The Practical Lawyer's Guide to AI Use and Privacy

Celesq® AttorneysEd Center
www.celesq.com

5301 North Federal Highway, Suite 150, Boca Raton, FL 33487
Phone 561-241-1919

The Practical Lawyer's Guide to AI Use and Privacy

Erich Dylus

Varia Law · CamoText

varia.law · camotext.ai · camovoice.com

Introduction

Varia Law

Law, consulting & programming boutique.

CamoText & CamoVoice

Fully offline anonymization, redaction, and speech-to-text software.

Experience

Prior firm: international aviation financings and secured transactions
Programming: smart contracts, custom software, and AI implementation advisory.

My Perspectives

- Price and solution flexibility via general LLMs outpaces alternatives
- Decision paralysis leads to insecure shadow AI use
- Privacy is imperative, one-way, highly subjective, and concerns differ meaningfully from other cloud services
- Policies must be simple to be followed

Today's Agenda

1. **GenAI as a Collaborative Tool**

   Framing AI effectively for legal work
   Use cases and prompting strategies

2. **Privacy Fundamentals**

   Logging, training, prompt injection fundamentals, and how to protect

3. **Implementation**

   Policies, tools, vendor evaluation, and practical bottom-line guidance

GenAI as a Genius Intern

① Brilliant and Creative

② Lacks *Full* Context and Experience

③ Overly Eager, Requires Supervision

GenAI as a Genius Intern

IDEAL FOR:

- Issue-spotting & analysis
- Subjective first drafting
- Second opinions & research
- Brainstorming and strategy critique
- Mock negotiations
- Regulatory guidance summaries
- Non-billable writing and marketing copy

AVOID:

- Final decisions without human review
- Client PII or privileged information
- Unverified AI output in filings or advice
- Mission-critical entrustment
- Unsupervised work or access

Effective Prompting for Lawyers

If you receive a clearly wrong or unexpected output, start over with a better initial prompt rather than trying to patch it.

1 Structure Your Prompt

Role → **Task** → **Instructions** → **Context**. LLMs process from left to right; context follows directives (or prompt twice).

Provide example and ask for prompt (reverse prompting).

2 Require Clarification

Use “Plan mode” instruct AI to ask you questions before proceeding—surfaces assumptions and avoids latent issues.

3 Demand Citations

Require verifiable sources for all claims. Check citations manually. You are responsible.

4 Save Useful Prompt and Skill Templates

Build a library with global rules, specific prompts, and workflow “Skills” automatically applied
These are portable to other services!

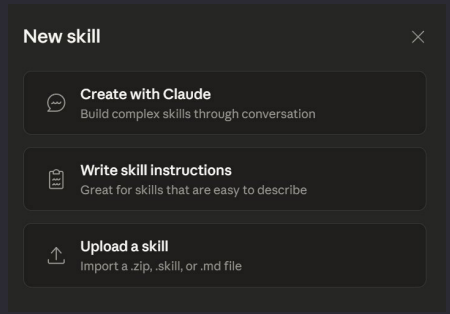

Three Primary Privacy Concerns

1

Prompt Material

PII and privileged data should never enter AI systems. Treat AI providers as you would an adverse counterparty: assume they can see everything you submit (even if they don’t train, they log).

2

Trained and Logged Data

Prompts, metadata, IP addresses, and usage data are retained for 30 days to 2+ years for "optimization" and "safety"—often without meaningful notice. User data often trains models in free tiers, often quietly used by staff.

3

Prompt Injection

Innocent looking file or link (whether saved or accessed by user, or an agent) can contain malicious instructions or code

AI Hosts Log Your Data, Inputs, and Outputs

What providers like OpenAI collect (from published privacy policies):

Prompts & Uploaded Files

Prompt, context, file, image, and audio submission is logged for performance and "safety review“ regardless of training

Log Data

Your IP address, account data, browser type, date and time of each request, and how you interact with the service.

Usage Data

Content types viewed, features used, actions taken, time zone, country, and device details.

Device Information

Device name, operating system, device identifiers, and browser version.

"[Your] chat is scheduled for permanent deletion from OpenAI systems within 30 days, unless we are legally required to retain it." — OpenAI Privacy Policy

Taking Privacy Into Your Own Hands

ON-DEVICE / LOCAL

Redact & Anonymize Before Prompting

Remove PII, privileged identifiers, and metadata locally before submitting to any cloud AI, even if a private instance.

ENTERPRISE

Private Enterprise Services

Use team/enterprise subscriptions with zero-data-retention options and data processing agreements. Understand exactly what is logged, for how long, and whether it trains the model.

Read the policy and settings closely as to notice/consent, and whether privacy is opt-in or out.

MONITORING

Opt In to Controls & Stay Alert

Enable data deletion wherever offered. Use VPNs and incognito/temporary modes.

Monitor privacy policies proactively—providers can update terms without meaningful notice, and court orders can override stated policies.

Disclosure can be Forced

Accordingly, OpenAI is **NOW DIRECTED to preserve and segregate all output log data that would otherwise be deleted on a going forward basis until further order of the Court** (in essence, the output log data that OpenAI has been destroying), whether such data might be deleted at a user's request or because of "numerous privacy laws and regulations" that might require OpenAI to do so.

SO ORDERED.

s/ Ona T. Wang

Dated: May 13, 2025
New York, New York

Ona T. Wang
United States Magistrate Judge

A Simple Data Classification Framework

RED — Do Not Use

Do not include in any AI prompts. Redact or anonymize, strip metadata.
Examples: Client PII (names, emails, addresses, phone), matter-specific confidential facts, privileged communications, personnel records, source code with competitive risk, internal financials, credentials.

YELLOW — Anonymize First

Case-by-case. Anonymize to protect sensitive details before prompting.
Examples: Anonymized engagement or usage data, de-identified survey results, aggregated client metrics, general competitive intelligence.

GREEN — No additional exposure

Generally acceptable for AI prompts, subject to firm policy.
Examples: Your role and general company description, publicly available legal standards, general drafting assistance unrelated to a specific matter, publicly available research.

What to Ask Legal AI Vendors

What differentiates your product from using privacy tools and general LLMs?

Privacy Policy

Is prior consent required for changes? What notice is provided? What data is used for training? What third-party dependencies exist—and what are their policies?

Data Retention

What is logged and for how long? Is client data isolated? Can specific conversation data be deleted on request (e.g., for terminated clients)?

Security

Is data end-to-end encrypted? Where are servers located? What happens to your data in a breach? Is there a BAA or DPA available?

Models & Dependencies

Which underlying models are used? Off-the-shelf or fine-tuned? What is the training set? What third-party APIs does the product call?

Agentic Controls

For AI agents: what permissions are granted? How are actions logged? Who reviews agent decisions? What human-in-the-loop controls exist?

Implementation Checklist

✓ **Start with Low-Risk Tasks**

Issue-spotting, analysis, strategy critique, and subjective first drafting. Build confidence before handling sensitive work.

✓ **Invest Time in Prompting**

Build prompt and skill templates, require citations, maintain a library.

✓ **Review All AI Output**

Check citations manually. Validate sources. Maintain professional skepticism. You are ultimately responsible.

✓ **Protect Privacy First**

Redact PII and privileged data locally before any AI interaction. Assume all cloud AI conversations are logged and potentially reachable by third parties.

✓ **Check Your Vendors**

Confirm logging, training, deletion, and security configurations, and revisit as policies change. Get contractual protections, but don’t rely on them.

Start with high-upside, low-downside tasks.

Invest time in effective prompts and source verification.

Control what data leaves your device.

Thank you!

Erich Dylus

varia.law · camotext.ai · camovoice.com · linkedin: erich-dylus