SCG Legal AI Strategy Exchange 2.0 MEETING #6: Prompting & Building [2025 Dec 2]

AI Summary

Note: Below are an (edited by SCG) **executive summary** and a **bulleted summary** of the meeting transcript created by ChatGPT. Listening to the recording – available here – is recommended to ensure full understanding of the presentation.

Executive Summary

The sixth and final meeting of **SCG Legal's AI Strategy Exchange 2.0** focused on *prompting and building* with generative AI. Led by **Nikki Shaver** of Legaltech Hub, the session introduced participants to foundational and advanced concepts in prompt engineering. The discussion emphasized that prompting is an evolving discipline influenced by ongoing research and platform changes.

Key themes included:

- Generative AI works best through **iterative prompt chaining** rather than single prompts, mirroring how lawyers work through complex tasks step-by-step.
- Effective prompts require **clarity, specificity, context, and structure** to reduce ambiguity and minimize user-induced errors.
- Legal professionals must remain vigilant about **hallucinations**, incorrect citations, and biases, while treating all AI outputs as **first drafts** requiring verification.
- The session highlighted the importance of **role-based prompting**, **communicating uncertainty**, and using structured instructions to guide models more reliably.
- Participants were reminded that prompting involves not only the user's instructions but also **backend system prompts**, which can dramatically affect model behavior (e.g., defaults for diversity in image generation).

Overall, the meeting aimed to equip legal practitioners with practical frameworks to interact more effectively with AI tools and prepare for more advanced capabilities in upcoming program iterations.

Bulleted Summary

Foundations of Prompt Engineering

- Prompt engineering is **ever evolving**, with new research continuing to refine best practices.
- Prompting applies across all platforms (ChatGPT, Claude, Perplexity, legal-specific tools).

Types of Prompts

- Distinction between:
  - **User prompts** (what the operator types)
  - **System prompts** (backend instructions that influence behavior)
- Example: Diversity defaults in AI-generated images changed when OpenAI modified hidden system prompts.

Tip 1: Use Prompt Chains

- Complex tasks should be broken into **iterative steps**, not a single question.
- Approach mirrors legal reasoning: extract facts → identify issues → analyze → produce output.
- Enables refinement, correction, and better overall quality.

Golden Rules: Clarity, Specificity, Context (and Structure)

- Avoid ambiguity; unclear instructions (e.g., “make it good”) are ineffective.
- Specificity prevents user-generated errors.
    - Example: “Georgia” must be clarified (state vs. Vermont town vs. country).
- Deep context improves model accuracy, consistency, and ability to handle complexity.
    - Example: Providing detailed trust structure and assets materially improves output quality.

Hallucinations & Verification

- Models respond confidently even when wrong; users must verify facts, citations, and assumptions.
- Legal AI tools increasingly embed verification workflows.
- Even small citation errors may be overlooked without careful review.

Structured Prompt Architecture

- Recommended structure:
    1. **Instructions**
    2. **Context**
    3. **Question**
    4. **Output placeholder** (e.g., “Answer here”)
- Can instruct the model: “Do not begin yet—confirm you understand.”

Tip: Ask AI to Admit Uncertainty

- Explicitly tell it to say “I don’t know” when unsure.
- Example: Model correctly refused to draft an invalid Georgia trust-company agreement when prompted with uncertainty guidelines.

Calibrating Confidence

- Users can request confidence levels or percentages for each statement.

Role-Based Prompting

- Assigning a role (e.g., “act as an estate planning attorney” or “law firm BD writer”) improves tone and relevance.

Format & Output Control

- AI can output in multiple formats—bullets, tables, even graphs.
- Ensure instructions specify format for each stage of a multi-step workflow.

General Best Practices

- Treat AI output as a **first draft**, never final.
- Always check for:
    - Factual errors
    - Citation accuracy
    - Biases
    - Logical inconsistencies