4 KEY TAKEAWAYS

Beyond the Prompt: Four Takeaways on Agentic AI in Patent Prosecution

Kilpatrick's Charles Gray recently spoke at the **31st Annual Patent Law Institute** in Austin, Texas, on the topic of "Beyond the Prompt: Agentic AI in Patent Prosecution." The presentation discussed how many practitioners continue to use AI as a chatbot, hunting for the right prompt. Charles discussed moving past just prompting to agentic systems that can perform multi-step prosecution tasks under attorney supervision.

Charles provides the following key takeaways from his presentation:

1. **A prompt is not a procedure.** A chat answer sounds plausible every time and has nothing to be tested against. A procedure runs the same way each time, has criteria that are true or false, and has a checker that can refuse.

2. **The limit is capacity, not intelligence.** A model works from one "whiteboard," its context window. Overfill it and the answer is written from whatever is left, with nothing failing loudly. A bigger window does not fix this; keeping each board small, with one focused job, does.

3. **Split, specialize, and argue.** A response council pairs an orchestrator that holds the plan with specialists for amending, arguing and estoppel, plus an examiner simulator that attacks the draft. The goal is not agreement but an argument that survives the attack. At volume, such as 300 references, ask about each one separately and then reduce.

4. **The judgment, and the signature, stay with the lawyer.** Under 37 CFR 11.18(b), signing a paper certifies the inquiry behind it. A practical start is one deliverable you produce every month: write what makes it acceptable as true-or-false statements, build a check that can refuse, and put a person in front of anything that leaves the building.

For more information, please contact:
Charles Gray: cgray@ktslaw.com