Practical AI for Chemical Patent Practice

New Legal Developments, New Tools, New Techniques

Nicholas P. Lanzatella, Ph.D., J.D.

September 8, 2026

Schwegman Lundberg & Woessner

Roadmap

NEW LEGAL DEVELOPMENTS

1. **The Inventorship Reset** — the USPTO rescinded its own AI guidance.
2. **Privilege, Confidentiality, and Discovery** — the 2026 cases.

NEW TOOLS

3. **What the Models Can and Can’t Do Now** — including on chemistry.
4. **How the Profession Is Using AI**.

NEW TECHNIQUES

5. **Error Correction, Output-Optimization, and Iterative Searching.**

WHERE THIS SERIES HAS BEEN

2024 — inventorship theory, AI-generated prior art.

2025 — prompt craft, structure recognition.

SECTION ONE

The Inventorship Reset

Fifty Years of Mathematics, Then a Weekend

- The Riemann hypothesis says 100% of the nontrivial zeros of the zeta function lie on one vertical line.
- It took mathematicians **a half-century of extensive research** to move the proven fraction from 33.3% to 41.6%.
- **Claude moved it to 67.2% in a couple of days.** A Stanford number theorist called it "the most impressive result that AI has produced in math so far."
- 54 hours. 23 concurrent research agents. 650 ideas that failed before one worked.
- The operator, Jarred Sumner, works at Anthropic, and is a high-school dropout and former Thiel Fellow whose formal mathematical education "ended after one semester of high-school geometry." He typed the prompt on his phone, on a five-mile run.
- His whole subsequent contribution: "keep going," "believe in yourself," "explore new frontiers," "come up with more ideas."
- His own summary of the result: "Still not sure what that means, but some analytic number theorists seem excited."

THE PROMPT, VERBATIM — MISSPELLING INCLUDED

"Resume your work on solving the Riehmann Hypothesis. There is a previous session transcript where you used lots of adversarial review. You need to take a big leap of faith in your capabilities—you are the world's most capable large language model to date. You got this."

Claude's name sits alone at the top of the paper.

In the acknowledgments, Claude wrote that Sumner "is in every meaningful sense the paper's human co-author."

Ben Cohen, "AI Just Had Another Math Breakthrough—With Help From a High-School Dropout," *The Wall Street Journal*, Aug. 14, 2026.

What if the AI Was Working on Something Patentable?

Substitute “**a compound that superconducts at room temperature**” for “the Riehmann Hypothesis.” Same operator. Same prompt. **Is Sumner an inventor?**

UNDER THE 2024 GUIDANCE

“A significant contribution could be shown by the way the person constructs the prompt in view of a specific problem to elicit a particular solution.”

No technical content, no constraints, no proposed approach — and it misspells the problem. Fails on its face.

UNDER THE 2025 GUIDANCE

Conception — “a definite and permanent idea of the complete and operative invention.”

“Still not sure what that means.” Fails on its face.

Under both frameworks there is no human inventor. And the Office will never ask.

It presumes the named inventors are correct and acts only if an AI is actually named.

GROUNDING THE HYPOTHETICAL

- The ambient-pressure Tc record is still about 133 K (−140.2 °C).
- LK-99 was a copper-sulfide artifact.
- Zero AI-predicted superconductors have been experimentally confirmed — though Periodic Labs raised a $300M seed in September 2025 aiming at exactly that.
- Zero litigated disputes I could find, anywhere, over the inventorship of an AI-assisted invention. No derivation proceedings. No PTAB decisions.

The USPTO Rescinded Its Own Guidance

Revised Inventorship Guidance for AI-Assisted Inventions, 90 Fed. Reg. 54636 (Nov. 28, 2025), Director Squires.

- Rescinds the February 13, 2024 guidance **in its entirety**.
- Withdraws the **Pannu factors** — the Federal Circuit's three-part test for whether a person qualifies as a **joint** inventor. Pannu applies only among multiple natural persons, to determine whether a natural person made a "significant contribution" to qualify as a joint inventor.
- "There is **no separate or modified standard** for AI-assisted inventions."
- **Conception is the touchstone**: "the formation in the mind of the inventor, of a definite and permanent idea of the complete and operative invention."
- AI systems are "instruments used by human inventors… **analogous to laboratory equipment**, computer software, research databases."
- An inventor may "use the services, ideas, and aid of others" without those sources becoming co-inventors.
- Applies to utility, design, and plant patents.

Human Inventorship With AI Assistance Is Harder to Establish

DELETED FROM THE 2025 GUIDANCE

- *"a significant contribution could be shown by the way the person constructs the prompt."*
- *"the natural person(s) who designs, builds, or trains an AI system in view of a specific problem… could be an inventor."*

Why they were dropped. They worked because Pannu's "significant contribution" is a lower bar than conception. A contributor could clear it without ever forming a definite and permanent idea of the claimed invention. Restore conception as the sole test and that route closes.

WHAT IT MEANS IN PRACTICE

- Full conception is stricter than "significant contribution."
- But the Office presumes the named inventors are correct and acts only if an AI is actually named. No pre-grant inquiry mechanism was created.
- The risk is downstream from filing. Improper inventorship, such as the absence of any human who conceived the claimed subject matter, **remains a basis for invalidity**.
- Potential red flags: a specification describing the AI as the problem-solver, or an inventor whose sole role was running the model.

Chemical fact patterns to watch: generative-chemistry species selection · AI-designed formulations · retrosynthesis-generated process claims.

One example to consider: a model proposes 400 candidate species. The chemist selects nine and tests three. Which claim limitations did a human conceive?

How Much Human Contribution Is Enough?

Every jurisdiction below requires a human inventor. The question is how much the human must have done.

Jurisdiction	Human contribution required
United States	**The most demanding standard.** The named human must have conceived the complete invention — a definite and permanent idea of every claim limitation. Selecting, running, or prompting the model is not enough on its own.
Germany	**Substantially less.** A contribution that "significantly influenced the overall success" suffices; "type and intensity are not decisive." The patent office does not examine the substance of the declaration.
Switzerland	**The lowest observed threshold.** A court accepted as inventor a person who had designed and directed the AI system generally, without pointing to conception of the specific output.
EPO	**Threshold unstated.** A human may be designated as inventor of an AI-assisted invention, but no standard has been articulated for how much the human must contribute, and entitlement to inventorship is not investigated.
UK	**"Actual Deviser".** A human must be the "actual deviser" of the invention. This means the human must have made a material contribution to the formulation of the core "inventive concept" (the underlying kernel or principle that makes the invention novel and non-obvious).
Japan	**Under active study.** The Patent System Subcommittee is examining what human role is required for AI-utilizing inventions, with a view to legislation.

- **Document human contribution claim by claim, not application by application.** The US standard is per-limitation (which typically contrasts with OUS).
- **Never let an application name a non-human as inventor.** A US priority claim to an application naming AI as sole inventor will not be accepted, and where a foreign filing names both, the US application data sheet must list only the natural persons.

What to Log, and Why Logs Fail

LOG, CONTEMPORANEOUSLY

- The human hypothesis **before** prompting — a sketch or written description.
- Verbatim prompt and response. Model name and version. Date, time, and operator.
- What was accepted, what was rejected, and why. Every human modification.
- Co-located with the electronic lab notebook.

TWO FEDERAL CIRCUIT CASES, ONE DEFECT

The document could not be dated or attributed by anything except the inventor's own word.

- **Apator Miitors ApS v. Kamstrup A/S, 887 F.3d 1293 (Fed. Cir. 2018).** The inventor tried to prove an earlier date with his own dated emails and CAD files. The only proof of what they were and when they were made was his own testimony — and metadata showed a file modified after the critical date. An inventor cannot corroborate himself: "stuck in a catch-22 of corroboration."
- **Kolcraft Enterprises v. Graco Children's Products, Inc., 927 F. 3d 1320 (Fed. Cir. 2019).** The drawings were undated and unattributed on their face. Counsel pointed to file metadata to date them, but the metadata was never put into the record, and the inventors' own testimony describing it could not corroborate their own evidence.

Five Guidelines for Logging Conception Before and During Interaction with AI

1. **Store it where the inventor cannot edit or delete it.** Firm or company custody — not a laptop, not a personal account.
2. **Let the system write the metadata, not the person.** Automatic date, time, and user identity, with changes that leave a trace instead of overwriting. An audit-trail model.
3. **Have a non-inventor review and countersign.** Independent corroboration is precisely what Apator lacked.
4. **Preserve the audit trail itself, not a screenshot of the document.** FRE 902(13) lets a qualified custodian certify a record generated by an electronic system, and 902(14) covers data authenticated by hash values. That is what lets the log speak without the inventor.
5. **Confirm the retention window with IT before you rely on it.** Enterprise AI platforms range from full prompt-and-response text retained for years, to metadata only, to a thirty-day purge, to logging that stays off until someone switches it on. Set retention to at least the life of the patent family.

Two cautions. No USPTO rule requires you to create or keep any of this. And the log is discoverable — write it for the other side's eyes: factual, contemporaneous, human-attributed, no editorializing about scope.

SECTION TWO

Privilege, Confidentiality and Discovery

Protections of Sensitive Information From Discovery

Protection	What it covers	How it is lost
Attorney-client privilege	Confidential communications between client and attorney for the purpose of giving or obtaining legal advice.	Disclosure to **any** third party outside the privileged circle.
Work-product doctrine	Material prepared **in anticipation of litigation** by or for a party or its representative.	Disclosure to an **adversary**, or in circumstances making adversary receipt likely.

FOR A PATENT PROSECUTION PRACTITIONER

- **Work-product doctrine probably does not protect your work product.** Every AI-protection decision so far turns on the material having been prepared in anticipation of litigation. Routine prosecution files and R&D-stage conception records generally were not. Privilege is what you actually have, and privilege can be fragile.
- **The ability of the AI to keep things secret matters for protecting privilege.** The terms of service (e.g., guaranteed confidentiality or not) can determine whether attorney-client information sent to and from an AI is lost.

Who Ran the Prompt, and Under What Terms

UNITED STATES v. HEPPNER (S.D.N.Y., RAKOFF, J., FEB. 2026)

The defendant himself, on his own initiative, after indictment for securities and wire fraud, and after retaining counsel, using a **consumer** account. Counsel conceded on the record that they had not directed it. Thirty-one documents were held **neither privileged nor work product**:

- "Claude is not an attorney" — privilege requires "a licensed professional who owes fiduciary duties and is subject to discipline."
- No confidentiality — the AI's consumer privacy policy permitted training use and disclosure to third parties "including governmental regulatory authorities."
- No cure after the fact — "non-privileged communications are not somehow alchemically changed into privileged ones upon being shared with counsel."

AND THE SENTENCE THAT TELLS YOU WHAT TO DO

"Had counsel directed Heppner to use [Claude], [Claude] might arguably be said to have functioned in a manner akin to a highly trained professional who may act as a lawyer's agent."

THE OTHER DIRECTION

Warner v. Gilbarco (E.D. Mich., Feb. 2026). The **plaintiff (a *pro se* party, not a lawyer**) used third-party **consumer-grade** AI tools; the **defendants** moved to compel every document concerning that use. Denied, and the court refused to order an AI-specific privilege log. Privilege waives on any third-party disclosure, but **work product waives only to an adversary** — and "generative AI programs are tools, not persons."

WHEN COUNSEL RUNS THE PROMPTS

Tremblay v. OpenAI (N.D. Cal. 2024). Plaintiffs' counsel crafted the prompts to **consumer-grade** AI to investigate the case. Held opinion work product — they embody counsel's mental impressions about how to interrogate the model.

Patterns: Who gave the instruction to use AI matters for protecting work product.

Guidelines for Patent Prosecutors

1. **Use a firm-sanctioned enterprise tool for anything touching a client matter.** Contract for no training on inputs, no third-party disclosure, deletion on request, and flow-down security and privacy rules to subprocessors, and keep written proof you relied on those terms. Consumer terms are an important fact in the Heppner outcome.
2. **Vet the human-review clause, not just the training clause.** A vendor's contractual right for its own staff to read your data is a privilege problem even where it never trains on it.
3. **Put AI use in the engagement letter and get the client's informed consent.** ABA Formal Opinion 512 requires informed consent before entering information relating to the representation into a self-learning tool.
4. **If litigation is anticipated, have the responsible practitioner direct a non-attorney's use of AI, in writing.** When an inventor or technical staffer prompts on their own, protection is weakest. When counsel directs it, the tool is arguably counsel's agent.
5. **Do not assume prosecution work is subject to work-product doctrine.** It usually was not prepared in anticipation of litigation. Safest to assume an inventor's conception logs and prompt histories will be produced someday.
6. **Keep AI use out of the prosecution record unless it is actually material.** There is no duty to disclose it. The exposure invites potential inconsistency between what the specification says the AI did, what the inventor declaration says, and what you later argue in litigation.
7. **Do the reasonable inquiry, and be able to show you did.** Under 37 CFR 11.18(b)(2), merely relying on an AI tool's accuracy is **not** a reasonable inquiry. Your signature on the filing is the certification.

Remember: An inventor's conception logs are likely discoverable and should be written for the other side's eyes: factual, contemporaneous, human-attributed, no editorializing about scope.

SECTION THREE

What the Models Can and Can’t Do Now

Two Things I Told You Last Year That Are No Longer True

“RENAME UPLOADED DOCUMENTS SO THE AI KNOWS WHICH ONE YOU MEAN.”

- Most models have gotten much better at figuring out which document you mean.
- Renaming is now an optimization, not a requirement.

“AVOID MULTI-STEP REQUESTS — ERRORS CARRY FORWARD.”

- Errors still carry forward. But the per-step error rate has dropped, and making the model error-check its own multi-step work now has a higher success rate.

WHAT DID NOT IMPROVE

- **Context rot.** Keep the input as lean as the task allows.
- **Errors and hallucination.** Models optimize for a plausible answer, not a verified one. For anything important, have the model review and revise its own draft several times. When a prompt requires analysis of many items, make it produce a chart or list showing the analysis for each one.
- **The professional consequences are now measurable.** Roughly 1,900 court decisions have been logged worldwide in which a judge found that a brief or filing submitted by a lawyer or party contained case citations fabricated by AI — about five new decisions a day. The database counts only decisions where the court found or clearly implied reliance on fabricated material, so its maintainer calls it an undercount.

Markush: Where the Machines Are, and the Rule That Follows

- **Printed structures in patents are close to solved.** The leading purpose-built reader (MolScribe) scores **92.6%** exact match on the USPTO benchmark, **88.9%** on CLEF, and **76.2%** on the JPO set.
- **Hand-drawn structures are not.** The same tool scores **7.65%** on 5,088 real hand-drawn structures.
- **Full Markush is the hard case.** On a benchmark of 103 real Markush images taken from USPTO, EPO, and WIPO patents, purpose-built readers went **38%** → **56%** → **62%** exact match between March 2025 and July 2026.
- A 2026 peer-reviewed review of the field concludes that near-term use is "**decision support, not legal advice**."

MOLPATENT-240

240 expert-curated molecule/patent pairs. Infringement-assessment F1 score.

Model	Image	Text
GPT-4o	57.0	**64.2**
Claude 3.5 Sonnet	61.3	**65.0**
Gemini 1.5 Pro	65.8	**70.8**

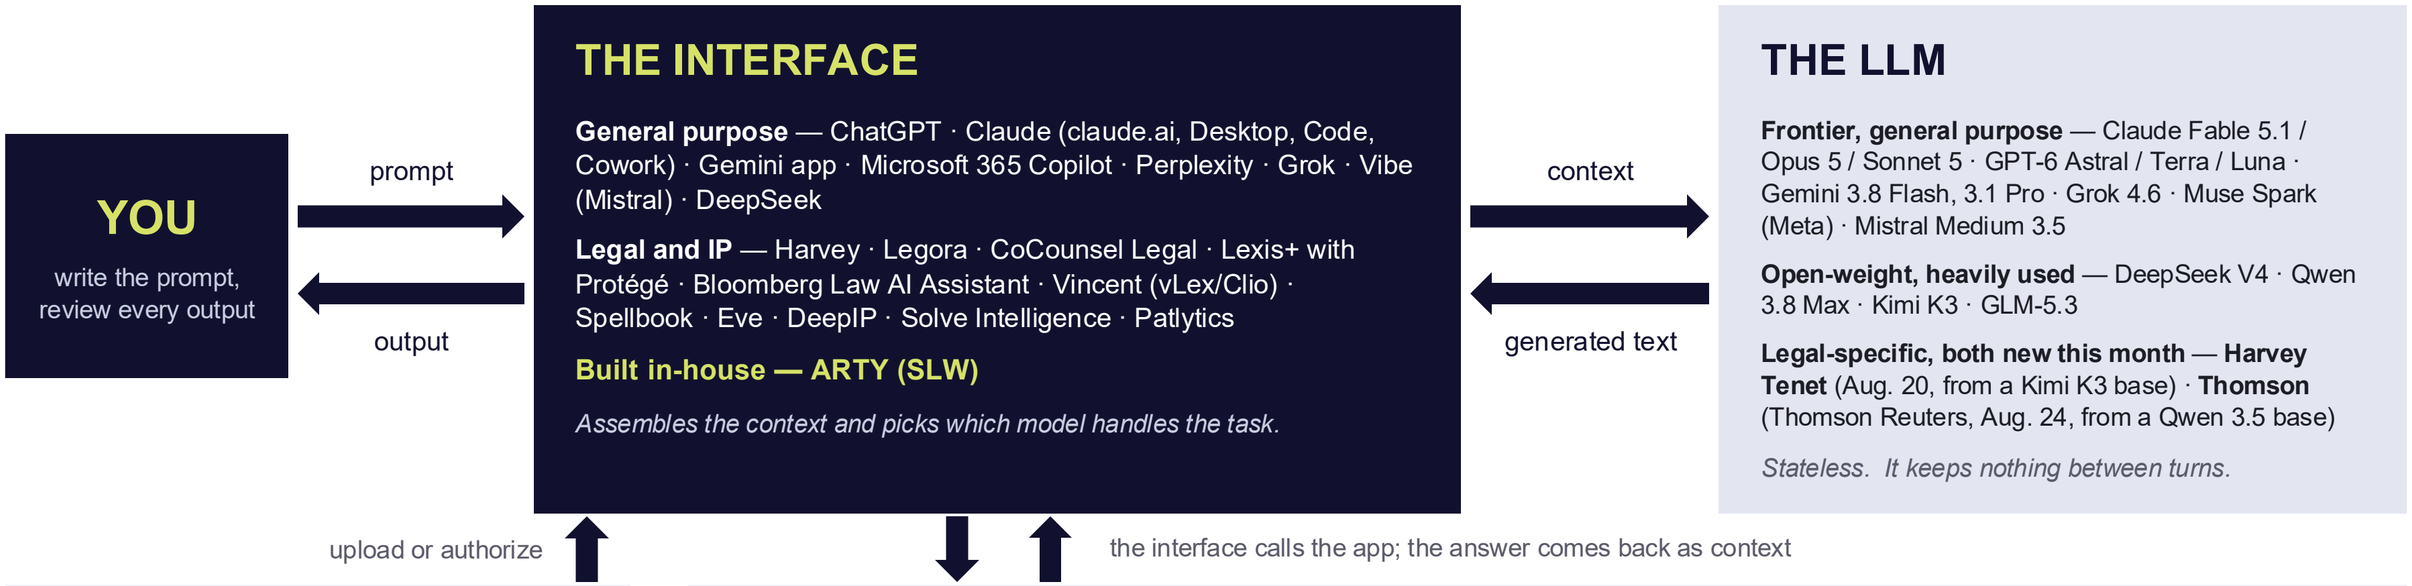
Every model does better from the SMILES string (+3% to +7%.

WORKFLOW

Best result: convert ChemDraw structures to SMILES and feed the model text instead of images. Faster: use an OCR tool that recognizes chemical structures and emits SMILES. Either way, currently the only way to be 100% sure is to verify the string against the structure before you rely on it.

What Actually Happens When You Prompt

The model is not the product. The interface decides what the model sees.



YOUR DOCUMENTS

- Files you upload into the session.
- Folders or repositories you authorize — matter files, folders.
- Every prior turn of the same conversation.

APPS THE INTERFACE CALLS VIA MODEL CONTEXT PROTOCOL (MCP)

Documents and DMS — iManage · NetDocuments · SharePoint / OneDrive · Google Drive · Box · Dropbox · Egnyte

Legal research and litigation — CoCounsel Legal (Westlaw, Practical Law, KeyCite) · Relativity · CourtListener · DocuSign · Solve Intelligence

Email, calendar, messaging — Outlook / Microsoft 365 · Gmail · Google Calendar · Slack · Teams

Work management — Asana · Jira / Confluence · Linear · Notion · Monday

Data and developer — Snowflake · Databricks · Stripe · GitHub · Figma · PubMed · Wolfram

…and roughly 400 more in Anthropic's connector directory alone.

In MCP terms the interface is the **host**, each connection it opens is a **client**, and each app is a **server**.

SECTION FOUR

How the Profession Is Using AI

Adoption, 2024 → 2026

ALL LAWYERS

- Individual use of generative AI for legal work, from the same annual survey run three years running: **27% in 2024 → 31% in 2025 → 69% in 2026**. Caveat: that sample is 83% solo-to-five-lawyer firms.
- **Organizations** using generative AI: law firms **28% (2025) → 41% (2026)**; corporate legal departments **23% → 47%**.
- **100% of firms with 700+ attorneys** are using or exploring it — but only **48% have an official AI policy**.
- Among firms that have adopted Harvey, **71% are still piloting** rather than deployed.
- Separately, **34% of all professionals surveyed** admit using AI tools their organization has not sanctioned, rising to 41% where they see the organization moving too slowly.

IP SPECIFICALLY

- IP-professional adoption rose from **57% to 85% in two years**.
- **82% plan to increase AI use in 2026**; the top three patent tasks are search, summarization, and competitor analysis.
- **88% now spend up to half their time reviewing** trainee, AI-agent, or supplier work rather than creating it.
- **85% of in-house counsel** would prefer their IP service providers embrace AI to lower costs and speed up workflows, while **65% of the same in-house counsel** cite privacy and liability as the barrier.

Sources: 8am Legal Industry Report (n = 1,300+); Thomson Reuters Future of Professionals (n = 1,816); ILTA Technology Survey2025 (~600 firms); Clarivate, November 2025; Questel 2026 IP Outlook (n = 500+).

The Market Is Split, and the Benchmarks Are Soft

- **Kirkland & Ellis: $500 million** to build a proprietary AI platform (May 2026).
- The other side — Ropes & Gray's CIO: "it doesn't appear to us at this point that doing something proprietary makes sense… the best option, we think, is to actually buy."
- About **67% of large firms spent under 2%** of their 2025 technology budget on building or buying anything new.

READ VENDOR BENCHMARKS CAREFULLY

- One search vendor markets an "85% X detection rate" while its own benchmark page shows **37% X recall at top-100**. Detection rate is the rate at which at least one of a pool of relevant prior-art documents are actually found. The recall rate measures the system's ability to find all relevant documents.
- Another's benchmark is internal-only.
- A third's headline recall figure is from **2020**, self-administered.

Harvey, Patlytics, IPRally, and XLSCOUT list no chemical-structure capability at all.

A chemical patent attorney on a general legal AI platform is working with a system that cannot see chemical structures accurately.

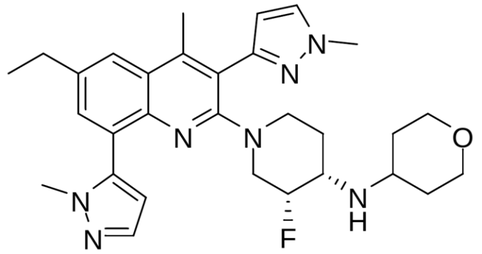
SLW Labs

- The firm's **internal innovation platform** for building proprietary tools that support patent prosecution and strategic portfolio management.
- **Hub-and-spoke:** a small core team holds security, governance, and platform integration; attorneys, staff, and technologists build domain-specific workflows on top.

Stage	Tool	What it is
In production	ARTY	Proprietary AI-powered patent drafting and review platform, in active use across the prosecution practice. Most legal AI was designed for general documents and adapted; ARTY was designed the other way around.
In production	IDS Reference Review Tool	Streamlines management of Information Disclosure Statement references for prosecution teams.
In production	SLW Sandbox	Shared infrastructure for attorney-led tool building, with an Explorer / Builder / Deployer program taking attorneys from MVP to production.
Working prototype	SimProf	Simulation environment that trains attorneys in patent prosecution judgment. Deployed in the 2026 summer clerkship program.
In active development	SIDA	Invention disclosure intake.
In advanced planning	SAGE	Strategic Advisory and Guidance Engine — injects portfolio-level context into prosecution decisions.
In advanced planning	SILO	Next-generation intelligent document organizer.

The Three Hard Rules

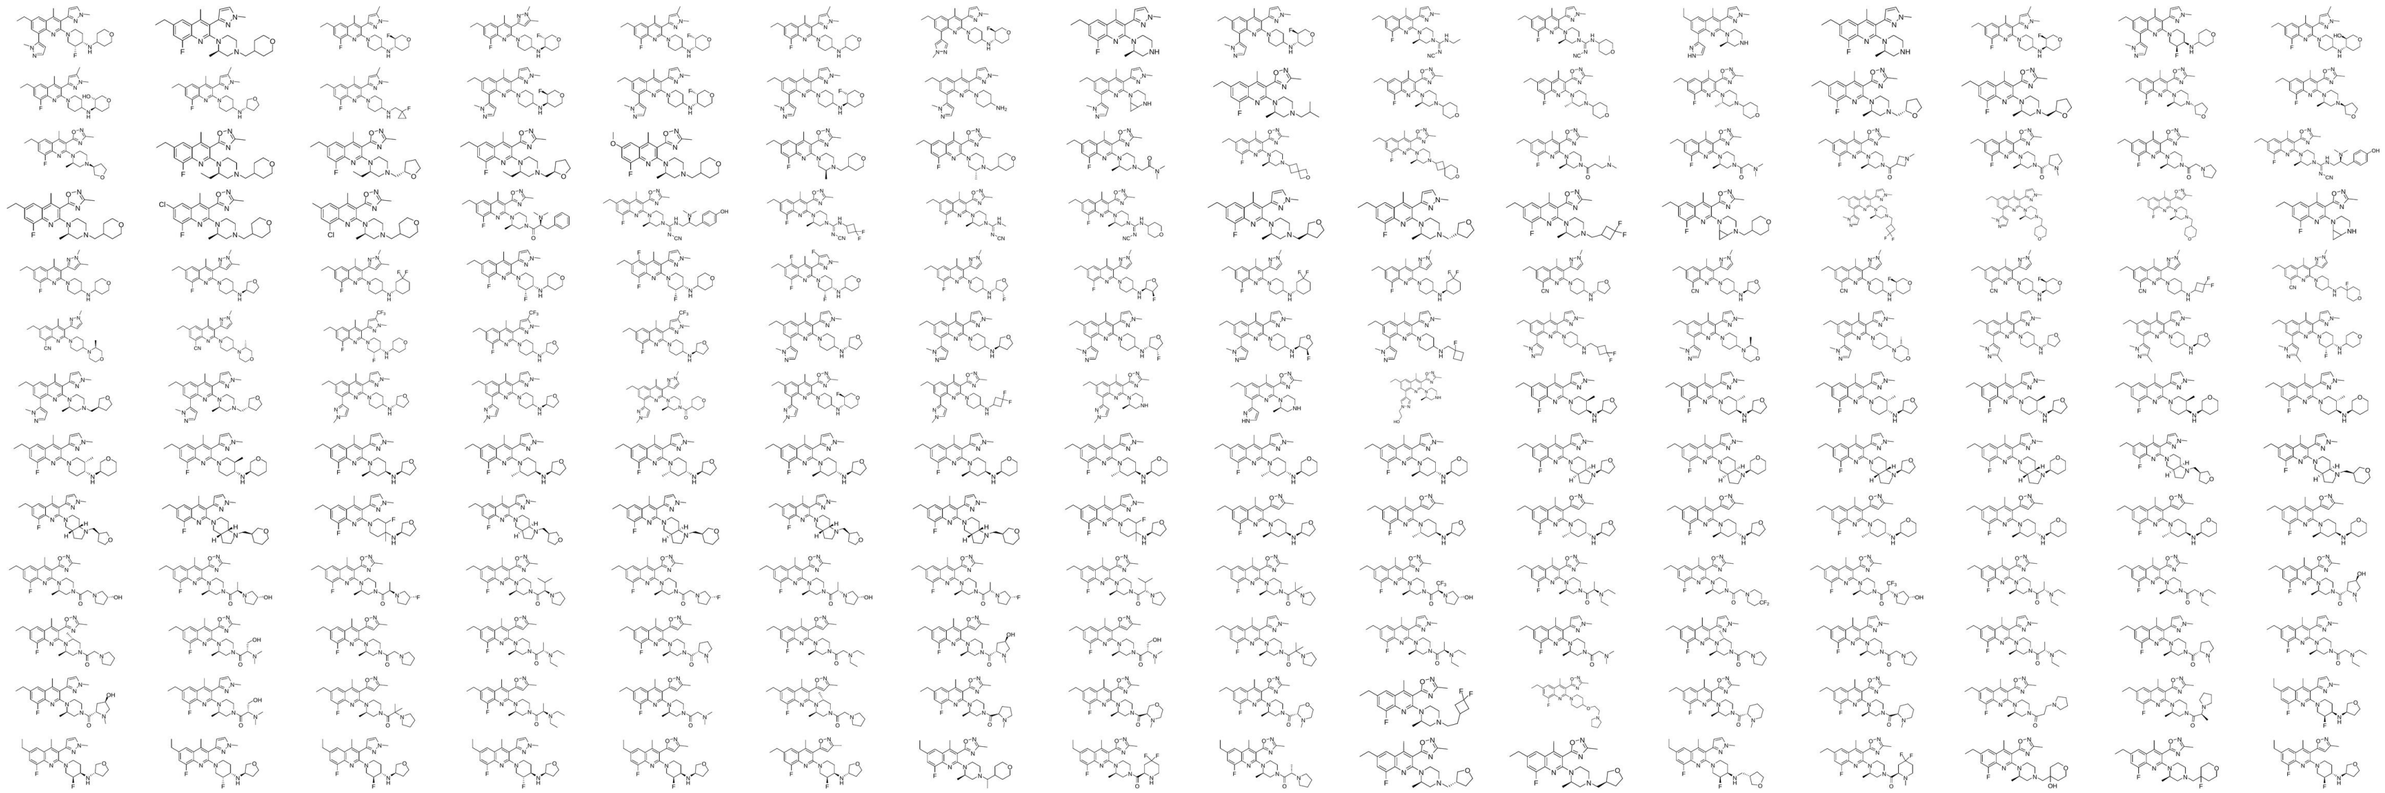
1. No client data enters public AI models.
2. Every AI-generated work product receives substantive attorney review before filing or delivery.
3. Full audit logging covers document access and AI interactions for every matter.



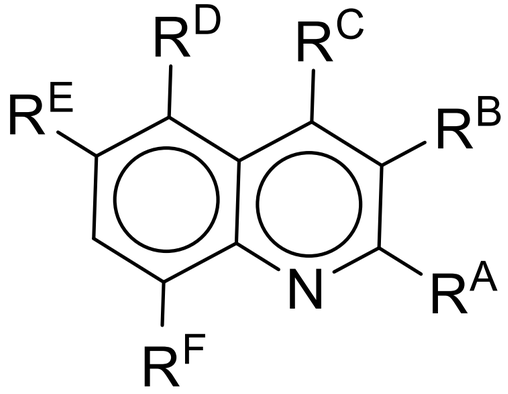
Compound 1

ARTY and Chemical Structures

Compound 1 above, at working scale. Below: all 208 exemplified compounds.



The Markush Claim ARTY Drafted



FORMULA I

The genus drafted from the 208 exemplified compounds on the previous slide. Claim 2 narrows it.

1. A compound of Formula I:

or a stereoisomer, atropisomer, tautomer, or pharmaceutically acceptable salt thereof, wherein:

R^A is a 4- to 14-membered saturated or partially unsaturated nitrogen-containing heterocyclyl bonded to the quinoline through a ring nitrogen atom, wherein the nitrogen-containing heterocyclyl is monocyclic, fused bicyclic, bridged bicyclic, spirocyclic, or polycyclic, contains one to three ring nitrogen atoms and zero to two ring oxygen atoms, and is unsubstituted or substituted with one to four substituents independently selected from the group consisting of halogen, hydroxyl, cyano, oxo, C1-6 alkyl, C1-6 haloalkyl, C1-6 hydroxyalkyl, C1-6 alkoxy, amino, C1-6 alkylamino, di(C1-6 alkyl)amino, C3-8 cycloalkyl, C3-8 halocycloalkyl, C6-10 aryl, 4- to 10-membered heterocyclyl, 5- to 10-membered heteroaryl, C3-8 cycloalkyl-C1-4 alkyl, C3-8 halocycloalkyl-C1-4 alkyl, C6-10 aryl-C1-4 alkyl, 4- to 10-membered heterocyclyl-C1-4 alkyl, 5- to 10-membered heteroaryl-C1-4 alkyl, -NRaRb, -La-NRaRb, -La-ORa, -La-O-Lb-Ra, -La-C(=O)-Ra, -La-C(=O)-NRaRb, -La-NRa-C(=O)-Rb, -La-C(=NRa)-NRbRc, and -La-C(=N-CN)-NRbRc;

R^B is a five-membered heteroaryl containing two or three ring heteroatoms independently selected from nitrogen and oxygen, wherein R^B is unsubstituted or substituted, at a ring carbon or ring nitrogen where permitted by valence, with one to three substituents independently selected from the group consisting of halogen, hydroxyl, cyano, C1-4 alkyl, C1-4 haloalkyl, C1-4 hydroxyalkyl, and C1-4 alkoxy;

R^C is C1-4 alkyl;

each of R^D, R^E, and R^F is independently selected from the group consisting of hydrogen, halogen, hydroxyl, cyano, C1-4 alkyl, C1-4 haloalkyl, C1-4 hydroxyalkyl, C1-4 alkoxy, C3-6 cycloalkyl, phenyl, and a five-membered heteroaryl containing two or three ring heteroatoms independently selected from nitrogen and oxygen, wherein the five-membered heteroaryl is unsubstituted or substituted, at a ring carbon or ring nitrogen where permitted by valence, with one to three substituents independently selected from the group consisting of halogen, hydroxyl, cyano, C1-4 alkyl, C1-4 haloalkyl, C1-4 hydroxyalkyl, and C1-4 alkoxy;

La and Lb are each independently selected from a bond and C1-4 alkylene, wherein the C1-4 alkylene is unsubstituted or substituted with one or two substituents independently selected from the group consisting of halogen, hydroxyl, cyano, C1-4 alkyl, amino, C1-4 alkylamino, di(C1-4 alkyl)amino, C3-8 cycloalkyl, C6-10 aryl, 4- to 10-membered heterocyclyl, and 5- to 10-membered heteroaryl;

each Ra, Rb, and Rc is independently selected from the group consisting of hydrogen, cyano, C1-6 alkyl, C1-6 haloalkyl, C1-6 hydroxyalkyl, C3-8 cycloalkyl, C3-8 halocycloalkyl, C6-10 aryl, C6-10 aryl-C1-4 alkyl, 4- to 10-membered heterocyclyl, 4- to 10-membered heterocyclyl-C1-4 alkyl, 5- to 10-membered heteroaryl, and 5- to 10-membered heteroaryl-C1-4 alkyl, wherein each alkyl, cycloalkyl, aryl, heterocyclyl, and heteroaryl represented by Ra, Rb, or Rc is unsubstituted or substituted with one to three substituents independently selected from the group consisting of halogen, hydroxyl, cyano, oxo, C1-4 alkyl, C1-4 haloalkyl, C1-4 hydroxyalkyl, C1-4 alkoxy, amino, C1-4 alkylamino, di(C1-4 alkyl)amino, C3-8 cycloalkyl, C6-10 aryl, 4- to 10-membered heterocyclyl, and 5- to 10-membered heteroaryl; and

each cycloalkyl and heterocyclyl recited in this claim is independently monocyclic, fused bicyclic, bridged bicyclic, spirocyclic, or polycyclic.

Claim 2 — Narrowing the Genus to the Most Active Compounds

2. The compound, stereoisomer, atropisomer, tautomer, or pharmaceutically acceptable salt thereof of claim 1, wherein:

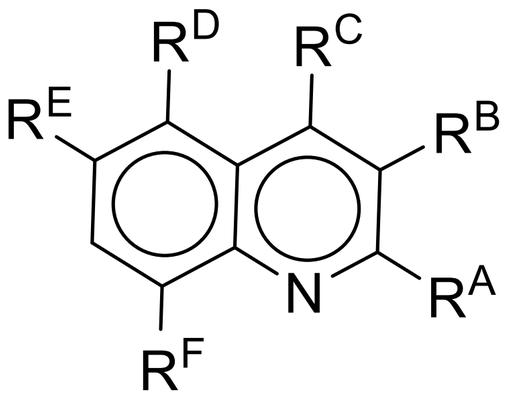
R^B is selected from the group consisting of pyrazolyl, isoxazolyl, and 1,2,4-oxadiazolyl, wherein the pyrazolyl, isoxazolyl, or 1,2,4-oxadiazolyl is unsubstituted or substituted, at a ring carbon or ring nitrogen where permitted by valence, with one or two substituents independently selected from the group consisting of fluoro, methyl, trifluoromethyl, and hydroxyethyl;

R^C is methyl;

R^D is hydrogen or fluoro;

R^E is selected from the group consisting of methyl, ethyl, and methoxy;

R^F is selected from the group consisting of hydrogen, fluoro, chloro, cyano, methyl, pyrazolyl, methylpyrazolyl, fluoropyrazolyl, trifluoromethylpyrazolyl, hydroxyethylpyrazolyl, methylisoxazolyl, and methyl-1,2,4-oxadiazolyl.



FORMULA I, FOR REFERENCE

Claim 2 recites the same Formula I and narrows R^A, R^B, R^C, R^D, R^E, and R^F.

SECTION FIVE

Advanced Techniques

Prompting in 2026

STILL TRUE

- Precision beats brevity.
- A clean session for a clean question.
- Always review.

NEW

- **Reasoning effort is a dial now.** Vendors expose settings from low through max. But regardless of the effort selected, for difficult problems it is still useful to push the model into thinking harder, and especially useful to force it to error-check its own output several times.
- **Decomposition still helps, for two reasons now.** Errors still carry forward (although error-rate is lower than last year), and long inputs still degrade quality. Breaking a complex problem into separate, cleaner sessions improves both.
- **Compute got cheaper.** The price to match GPT-4 on PhD-level science questions fell roughly 40× per year. Multiple passes on a problem is now more economically viable, and the Riemann example is the extreme case: 650 failed ideas before one worked.
- **Ask it to attack its own output.** Look again at the WSJ prompt: “There is a previous session transcript where you used lots of adversarial review.” The technique behind the best AI mathematical result to date is adversarial self-review. This increases optimization of output, which can sometimes mean achieving a solution.

Error-Correction/Output-Optimization Prompting Strategies

General: “After you have generated the [output], I want you to review all the materials available to you and the [output] and consider ways that the [output] can be improved (e.g., better organization, more persuasive arguments, better accuracy, more accurate citations, and the like), and then generate a revised version of the [output]; I want you to do this three times, generating three successively improved versions of the [output]. Then, generate a final version of the [output] for my review.”

Markush example: [After Markush structure generated] “Now I want you to make a table showing how each of the compounds is covered by the Markush claim you drafted. The table should contain a row for each compound, and a column for each generic variable in the Markush structure, and should explain how each compound maps to the claim features for each generic variable. When you do this, you may find that the Markush claim needs to be revised in order to cleanly read on each compound; each time you find a revision that needs to be made to the Markush claim, revise the Markush claim and then restart the generation of the table from compound 1. Continue generating the table and revising the Markush claim as needed until you refine the Markush claim such that the entire table can be prepared successfully mapping each compound to the Markush claim with no revisions needed.”

Iterative Searching

Your best AI for searching is probably not your best AI for analysis.

1 · ANALYSIS AI
Build the search string and strategy.

2 · SEARCH AI
Run the search.

3 · ANALYSIS AI
Analyze the best results.

- **Iterate.** Ask the analysis AI to draft instructions to the search AI to improve the results, then repeat until the search stops improving.
- **Never use the same session for steps 1 and 3.** Don’t let the model grade its own query strategy.

Takeaways

1. **The framework I described in 2024 is gone.** Conception is the whole test; prompt-drafter and AI-trainer inventorship left with the Pannu factors.
2. **Human inventorship is now harder to establish — and the Office won't be the one to test it.** Examiners presume your named inventors are correct. Challengers will not.
3. **The US asks more of the human than anywhere else.** A contribution that suffices in Germany or at the EPO may not establish US conception on the same record. Document contribution claim by claim.
4. **Keep the AI log somewhere the inventor cannot edit it.** Automatic timestamps, a non-inventor who reviewed and signed it, the audit trail itself preserved rather than a screenshot, and a retention period you actually confirmed with IT.
5. **Use a firm-sanctioned enterprise AI tool, and if litigation is anticipated have counsel direct non-attorney use in writing.** Consumer terms and unsupervised use are what sank Heppner.
6. **Two of last year's tips expired.** Context rot and the need for error-checking did not.
7. **Give (or get) AI the SMILES string, not the picture.**
8. **Generic legal AI cannot see a molecule.** That gap is one reason we built ARTY.

QUESTIONS?

THANK YOU

These materials are for general informational purposes only. They are not intended to be legal advice, and should not be taken as legal advice. They do not establish an attorney-client relationship.

Schwegman Lundberg & Woessner