

AI in Patent Drafting and Prosecution: Our Approach

Stevens Center is committed to efficiency and cost-effectiveness in everything we do. We take the question of AI in patent work seriously — and this document reflects our current approach, informed by our experience, our outside patent counsel, and the broader tech transfer community.

We Are Actively Evaluating The Role of AI in Patent Work

Can AI be used to draft patents faster and more cheaply? We hear this question regularly from inventors, faculty, and members of the USC community. It is a fair question, given how useful AI is for many tasks. Stevens Center is actively looking for ways to increase efficiency and reduce costs across everything we do, and patent prosecution is one of our most significant line items. AI is a natural place to look.

We drew on our own experience integrating AI tools into our workflows, our decades of experience in patent prosecution and licensing, and most importantly, direct conversations with our patent law firms about how they are using AI today and where they see it heading. We also stay actively connected to the broader tech transfer community through AUTM, where peer offices share what is and is not working in real portfolios. This document reflects that collective input.

In practice: AI is already part of how our law firms work, and Stevens Center has already seen cost reductions as a result. But the picture is more nuanced than it might appear from the outside. This document explains what AI can and cannot do in patent practice today, why the limitations matter for licensing outcomes specifically, and how we intend to continue evolving our approach as the technology and legal landscape develop. Our commitment is to move in step with the legal industry so that everything we do has a sound legal basis and serves the long-term interests of our inventors and institution.

What AI Can and Cannot Do in Patent Work

Because AI can generate text that looks like a patent application, it is reasonable to ask whether it can substantially reduce attorney time and cost. Our firms are asking the same question, and while AI adds value in some areas, others continue to require experienced legal and technical judgment.

Patent drafting is not primarily a writing exercise. It is a legal, technical, and commercial strategy function. The value of a patent does not arise from how well it reads — it arises from how well it is conceived, structured, and positioned relative to prior art, competitors, and future enforcement and licensing opportunities. AI can assist with targeted drafting tasks; while strategic decisions require experienced attorney judgment and inventor feedback.

Equally important: the licensability of a patent depends heavily on demonstrating and substantiating that the technology works and outperforms the state of the art. Data generation with targeted experiments — not drafting cost — is one of the primary constraints on licensing success. Filing more applications does not change that equation.

Where Attorney Judgment Adds Value

1. Strategic Judgment Drives Value

The most important decisions in any patent application, what to claim, how broadly to claim it, and how to position the invention, require judgment informed by the prior art landscape, competitor positioning, and licensing potential. These are legal and commercial strategy decisions grounded in:

- Experience with prosecution dynamics and how to interact with examiners so the broadest claims are allowed without creating perilous prosecution history
- Knowledge of what claim types are attractive to licensees in specific industries
- Understanding of the kinds of claim structures that are most durable in enforcement and diligence proceedings (e.g. overcoming subject matter eligibility requirements, such as 101)
- Understanding relevant case law (Mayo, Alice, etc)

2. Claim Language Precision

Small wording choices can have significant consequences. They determine whether a patent is valid, whether competitors can design around it, and whether the asset holds licensing value. Currently, AI-generated claims tend toward one of common outcomes: too narrow and easily circumvented, or too broad without adequate specification support and therefore may create challenges during diligence or litigation.

3. Areas That Require Experienced Review

Certain areas require experienced review to identify and manage:

- Characterizing inventive features as prior art in the Background section, which courts and examiners treat as admissions that can be novelty destroying
- Overclaiming beyond what is enabled by the specification
- Creating internal inconsistencies that can undermine patentability

These issues are invisible until they surface during prosecution, diligence, or litigation, precisely when they are most costly to address.

4. Prosecution Strategy Is Built Into Every Application

An experienced attorney drafts with examination in mind. That means layered dependent claims, fallback claim positions, and specification support carefully constructed to allow future amendments without conceding scope. AI systems optimize for immediate outputs. They may suggest amendments that improve allowance likelihood while narrowing claims to the point of not being commercially valuable, or put another way, short-term wins that become long-term liabilities.

5. Privilege, Confidentiality, and Institutional Risk

Operating without attorney oversight creates institutional risks that extend well beyond any single application:

- Inputs to AI tools may unintentionally create prior art that bars your own later-filed patent
- Attorney-client privilege and work product protections may be compromised depending on how AI tools are used
- Duty of candor and inventorship obligations are the attorney's responsibility to manage, they cannot be delegated to a language model

6. AI Still Requires Expert Supervision

Some AI can introduce risks related to things like hallucination, overconfidence, and agreeableness. This can reinforce user assumptions rather than challenge them. AI can generate technically plausible content that was never actually disclosed by the inventors.

AI tools may not be well suited for the technical elements that many patents depend on most such as mechanical drawings, chemical structures, genetic sequences, and similar specialized content. These require precise, domain-specific expertise that current AI cannot reliably provide at this point. What makes this particularly risky is that errors in these areas are often subtle and difficult to spot without deep technical knowledge. Unlike a poorly worded claim, a flawed chemical structure or an inaccurate figure may pass unnoticed through review only to become a challenging issue during prosecution or litigation.

Law firms consistently find that the time required to review, validate, and correct AI-generated drafts offsets the perceived cost savings. What appears to reduce expenses at filing can create hidden liabilities downstream such as failed licensing, increased litigation exposure, and reputational damage that far exceeds what was saved.

Risk at a Glance

Risk Factor	With Attorney Oversight	AI-Only (No Attorney)
Claim Quality	Strategic breadth, enforcement-ready language, prosecution history carefully managed	Narrowly drafted or overbroad; legally fragile; easy to design around
Prior Art Risk	Attorney identifies and mitigates disclosure risks in Background sections	Inventive features inadvertently characterized as known art; validity undermined
Prosecution Strategy	Layered fallback claims, examiner anticipation, amendment-ready specification	Optimizes for immediate allowance; sacrifices commercial scope
Privilege & Confidentiality	Attorney-client privilege preserved; duty of candor managed	Privilege may be lost; AI inputs may create unintentional prior art
Hallucination Risk	Human review catches errors before filing	Unsupported embodiments filed; hidden defects emerge in diligence or litigation
Licensing Potential	Patent built as a licensable asset aligned with deal objectives and supporting data	Technical document with reduced licensability; fails diligence scrutiny
True Cost	Higher upfront; dramatically lower downstream risk	Lower upfront; hidden liabilities may cost money to fix that exceed savings

The Licensing Imperative

The purpose of a technology transfer office's patent program is not to file patents for the sake of growing a large docket. It is to generate revenue through licensing. These are different objectives that require different standards.

Weak patents that are narrow, internally inconsistent, legally fragile, lacking data support are significantly less likely to be licensed. AI-only workflows can increase quantity but may reduce quality, which will not improve licensing outcomes. Filing high volume, AI-driven applications does not support a viable licensing strategy. It is an administrative burden that consumes resources and creates portfolio risk.

There is also a timing dimension that AI-driven volume ignores. Filing too early starts a clock. The more productive strategy is to delay filing strategically, until there is enough of a data package and a compelling technology story to attract licensees. In a university TTO setting, that is usually when there is an accepted manuscript to a journal. AI-based volume filing accelerates the wrong part of the timeline (the earliest stage, when there is not enough data), creating mismatches between patent timelines, technology maturity, and commercial readiness.

Each premature filing also increases the risk of generating prior art that narrows future claim scope in future filings and locks in immature embodiments thereby, degrading the portfolio over time, not strengthening it.

Ultimately, successful licensing requires an eager and energetic inventor, competent tech transfer staff, strong proof-of-concept data, and a solid IP position. AI is useful as a tool in this process, but it cannot address the real bottleneck of identifying the right corporate partner, engaging in fruitful discussions, establishing trust and negotiating a win-win agreement for a long-term relationship.

How Stevens Center Is Using AI Today

Stevens Center is not waiting on the sidelines. We are actively using AI tools in our day-to-day operations, and we expect that usage to grow as the technology matures and as better purpose-built tools become available. As of 2026, our team uses two primary platforms in support of patent and licensing workflows.

ChatGPT is used across a range of internal tasks including drafting and editing communications, summarizing technical disclosures, preparing marketing pitches for technologies, identifying competitive technologies and potential licensees, and supporting research tasks that inform licensing strategy.

Tradespace is an AI-based tech transfer platform we use to support prior art searches, generate technology marketing materials, and identify potential industry partners and company lists for outreach. It is purpose-built for the tech transfer context, which makes it particularly well-suited to the early-stage commercialization work that defines much of what we do.

We are also actively evaluating additional tools as they emerge. FirstIgnite, an AI-powered platform focused on research partnership and commercialization matching and which also offers prior art reviews, is among the tools we are currently assessing. Our approach is to evaluate new platforms rigorously, adopt those that demonstrably improve outcomes, and retire or replace tools as better options become available. The specific tools we use will evolve; the standard we hold them to will not.

Through our participation in AUTM and ongoing dialogue with peer institutions, we also benefit from a broader view of what is working across the tech transfer community. When other offices identify tools or workflows that are producing real results, we learn from that experience and factor it into our own decisions.

The Right Role for AI in Our Practice

Used effectively, AI is a meaningful efficiency tool. The following represent current, high value uses that are currently being used by skilled attorneys.

- Drafting efficiency: generating initial structural drafts that attorneys shape, refine, and strategically position
- Office action responses: accelerating preparation of routine responses under attorney review
- Prior art searches: surfacing relevant literature for attorney analysis
- Claim charting and mechanical comparison tasks

- Targeted, scoped questions, rather than open-ended drafting tasks
- Finally attorneys are building their own tools and agents and using them in practice today to address the above

What AI should not be permitted to do unsupervised:

- Determine claim scope or prosecution strategy
- Generate figures
- Identify and resolve legal traps in specifications
- Manage privilege, confidentiality, or inventorship obligations
- Make the judgment calls that determine whether a patent has licensing potential

The bottleneck in technology transfer is not drafting or filing — it is finding a licensee and convincing them to take a chance on early-stage technology. Our goal is to use every available tool, including AI, to get there faster and more efficiently.

Fewer, better-conceived patents built on sound data, drafted with prosecution strategy in mind, and aligned with licensing objectives will always outperform a larger docket of weaker assets. This is not an argument against efficiency, it is an argument for investing in the right efficiencies. AI adoption that improves drafting speed, surfaces prior art faster, and reduces routine prosecution costs is exactly what we are pursuing. What we are being deliberate about is ensuring that efficiency gains do not come at the cost of portfolio quality or legal integrity.

The goal is not more patents. The goal is more licenses. AI is one of the tools we are committed to using in that pursuit, thoughtfully, in step with the legal industry, and in ways that hold up when it matters most.