

Cutting through the noise: what differentiates leading AI tools for patent drafting and prosecution

At a glance

The landscape of AI tools for IP work has never moved faster. As we discussed in our previous article [From demos to impact](#), the number of AI tools targeting patent attorneys and IP professionals has expanded rapidly in the last few years. For many practitioners, keeping up has become a job in itself, and cutting through the noise is more difficult than ever before. Having spoken with patent attorney clients and colleagues who have moved past evaluation and are running these tools in live workflows, we have started to see consistent patterns emerge. Six dimensions stand out, separating the tools people adopt from the ones they set aside.

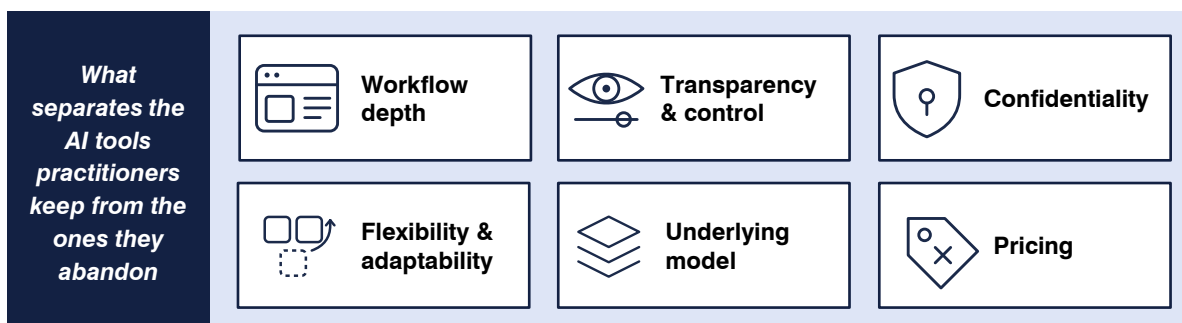
June 2026

AUTHOR



Edvin Brogeland
Senior Associate

Key dimensions differentiating leading AI tools for patent drafting and prosecution



Where do the tools diverge?

1. Workflow depth

On the drafting side, weaker tools generate text. The stronger ones organise a full drafting environment: bundling disclosures, prior art, figures, and metadata into a single managed workspace, rather than leaving attorneys to navigate file directories. They flag disclosure gaps before drafting begins, generate and number figures in compliance with patent office requirements, handle specialised inputs like chemical structures and sequence listings, and allow the user to apply firm- or client-specific style templates.

On the prosecution side, the leading tools follow a consistent pattern. The attorney submits an application number. The tool then automatically parses the prosecution history, generates a structured summary of the examiner's reasoning, and maps claims against rejections. Some platforms extend this further, automatically generating client notification letters and instructions to foreign agents from customisable templates. Tools that require attorneys to manually upload file history, or that produce only a flat summary without automating the workflow fall short.

2. Flexibility and adaptability

The tools that are weakest on this dimension have a fixed structure and require the attorney to work within it. Some have strict workflows that can look neat at first glance but oftentimes lack the flexibility that is in the nature of patent attorney work. The strongest tools are modular: prompts, outputs, workflows, and style are all customisable at the level of individual attorneys and client relationships. This matters particularly for external counsel drafting for multiple clients, each with specific requirements on format and style.

Flexibility also has a direct bearing on adoption. A rigid tool asks attorneys to change how they work, whereas a flexible tool can be shaped around how they already work. For firms with sceptical or experienced practitioners, this difference can be what determines whether a tool gets used at all.

3. Transparency and control

Some tools black-box too much of the process. They take inputs and return outputs, with limited visibility into what happened in between and limited ability to steer or correct mid-task. The better tools are built around an iterative, attorney-directed workflow: documents can be added mid-session, focus can be adjusted, and output can be interrogated step by step. They function as a sparring partner rather than an autonomous generator, and they keep the attorney in a position to direct, challenge, and correct.

This is particularly consequential in opposition and prosecution work, where the real time cost is not drafting responses but getting up to speed on a complex file. The strongest tools can ingest an opposition bundle of twenty to thirty documents and compress hours of reading into a structured, searchable knowledge base, allowing the attorney to use the tool to help build their own understanding. Tools that cannot handle that volume, or that don't allow the attorney to freely interrogate the underlying material, offer a fraction of that value.

4. The underlying model


A well-designed interface on top of a weaker large language model (LLM) will still produce weaker results. All commercially available drafting and prosecution AI tools are built on top of an LLM developed by a handful of major providers and there are meaningful differences between them in reasoning quality, handling of ambiguity, tone, and reliability. Attorneys report these differences most clearly on tasks involving nuanced claim language, where the stronger models reason more precisely and the weaker ones produce output that requires significant correction, if it is to any help at all. As the field develops, the ranking of models may shift, but the underlying model will remain one of the most important variables in how a tool performs, and it is key to understand which model a tool is built on and how well it manages complex tasks.

5. Confidentiality

Most vendors claim that their models are not trained on user input, but that is only part of the picture. The tools that go further stand out. Two additional aspects differentiate the leading tools. First, jurisdictional control: the stronger tools allow clients to elect where data is stored geographically, which matters for both legal and geopolitical reasons, as does whether the underlying model infrastructure is subject to export control regulations that could restrict use from certain jurisdictions. Second, model-level monitoring: model providers are subject to logging and monitoring obligations under the EU AI Act, and the extent to which user inputs fall within that monitoring depends on the specific arrangements a tool has secured with its underlying provider. The best tools have addressed this directly and have secured concrete arrangements to limit exposure of user inputs at the model level.

6. Pricing

Pricing is usually simple in structure but differs in magnitude. While some tools use pricing models such as a fixed price per patent application, most price on a per-seat, per-year basis with module-based add-ons. What varies considerably is the absolute cost level. The spread between the most and least expensive tools is wide enough to be a material procurement decision, particularly for smaller firms or cost-conscious in-house teams. The level of discount that vendors offer also varies considerably, depending on factors such as whether the customer is an early adopter, the length of the commitment, the scale of the implementation, and, of course, the buying power of the firm.



What to take away

There is no need to start from scratch. The market is already testing AI drafting and prosecution tools in live use, and there is a great deal to learn from peers who have gone before you. The six dimensions won't capture everything, but they offer a strong basis for cutting through the noise and judging which tools are worth your time.

That said, this remains a fast-moving field. Some tools are already delivering more value than others, but switching costs are still low and few firms are deeply locked in. As the tools continue to develop rapidly, the sensible posture is to adopt what works today without committing so heavily that you cannot move when the field shifts, because it will.



www.konsert.com