LexisNexis®

AI in IP

Learn from IP thought leaders.
Understand the impact.
Prepare for what's next.



Foreword



In the late 2010s, artificial intelligence was still a relatively marginal topic at intellectual property conferences and was typically confined to areas such as search. By 2023 and 2024, that had changed dramatically. From the major IP gatherings in the United States to the Intellectual Property Information Fair & Conference in Tokyo and Europe's IP Service World in Munich, AI had become one of the dominant themes across the global IP community. At times, it seemed as though almost every second presentation had AI in the title.

That intensity reflected both excitement and uncertainty. Today, the conversation is beginning to mature. The question today is where AI can create meaningful and sustainable value, and where we can rely on it to do so.

We first saw AI entering patent workflows primarily through search: helping professionals navigate ever-growing volumes of patent information more efficiently. From there, its role expanded into drafting, prosecution, and other process-oriented applications. General-purpose AI has added another dimension, giving IP professionals new ways to summarize information, explore unfamiliar topics, generate ideas, and streamline everyday work.

While these developments matter, I believe the more consequential opportunity is still emerging. It is the use of AI not only to make existing IP processes faster, but to help turn intellectual property information into better strategic decisions.

Patent data can provide a unique perspective on technologies, competitors, innovation trajectories, and emerging areas of opportunity. AI has the potential to make those insights more accessible, to allow us to ask more complex questions, and to deliver IP answers at business speed.

At the same time, expectations of AI are becoming more demanding. Fluency is rapidly becoming commonplace, while trustworthiness is not. For consequential IP and business decisions, an answer must be grounded in authoritative information, based on reliable data, and capable of being understood and challenged. In patent intelligence in particular, the quality of the underlying data remains fundamental. No amount of AI sophistication can compensate for an unreliable analytical foundation.

A comparison for 15 multi-national, well-known companies showed that some commercial patent databases miss up to 71.7% of the relevant patents in a portfolio. Imagine a cross-licensing scenario or risk assessment, and that gap in understanding can translate directly into an incomplete assessment of value, exposure, and negotiating position. No executive wants to explain to the board how that happened.

This is also why AI should not be considered an objective in itself. A patent landscape is not the outcome. An analysis is not the decision. The value is created when technology, trusted data, and professional judgment come together to answer a question the business genuinely needs to address: where to invest, which technologies to develop, whom to watch, where opportunities are emerging, and where risks may be accumulating.

For us at LexisNexis Intellectual Property Solutions, this is the direction that matters most: not AI for its own sake, but AI combined with trusted data and IP expertise to help professionals and business leaders make better-informed decisions. This helps to elevate the role of IP further, from an administrative function to a true business partner and advisor. The next question is therefore not simply what AI can do for IP, but what organizations can do differently when AI, trusted data, and expertise come together.

I invite you to explore the perspectives in this report, challenge some of your own assumptions about AI and IP, and consider what this next phase of AI adoption could mean for your organization.

Marco Richter
Managing Director & Chief Commercial Officer
LexisNexis® Intellectual Property Solutions



Contents

Foreword 1

Editor's Introduction 4

By Holger Zimmer, Global Director Strategic Communication & Partnerships for LexisNexis® Intellectual Property Solutions

AI's impact on the patent system 5

By Bowman Heiden, Co-Director of the Center for Intellectual Property (CIP) at the University of Gothenburg, and Executive Director of the Tusher Strategic Initiative for Technology Leadership, UC Berkeley.

How AI is changing the work of IP professionals 11

By Ruud Peters, Co-Director at the Center for Intellectual Property (CIP), CEO at Peters IP Consultancy B.V., former Chief IP Officer and Executive Vice President at Koninklijke Philips N.V.

Using general-purpose AI across the IP value stream 17

By Dr. Stephan Wolke, CEO of thyssenkrupp Intellectual Property GmbH, and Stephan Ising, Head of IP Strategy / Research of thyssenkrupp Intellectual Property GmbH

Democratizing IP intelligence: How AI broadens participation and strengthens decision-making 29

By Sakae Nakamura, Executive Advisor to LexisNexis® Intellectual Property Solutions /Senior Advisor, IP Landscape Association, Japan

From inventor culture to invention impact: How AI is changing the challenge for IP management 37

By Tim Alexander Oelmann, VP, Head of Competence Center Business IPR Management at Deutsche Telekom AG, and Daniel Uhl, Patent Information Specialist at Deutsche Telekom AG

From embedded AI to agent-ready patent intelligence 42

By William Mansfield, Principal Product Manager, Strategy, Discovery & Market Intelligence (SDMI) for LexisNexis® Intellectual Property Solutions

Closing Summary 56

About the contributors 57

About LexisNexis® Intellectual Property Solutions 63

Editor's Introduction

This report brings together contributors whose work spans corporate IP leadership, legal practice, patent intelligence, innovation strategy, and IP policy, including work with policymakers, patent offices, and other institutions involved in the development and application of intellectual property frameworks. Their chapters examine AI in IP from several professional and practical perspectives, from the patent system and the changing role of IP professionals to patent intelligence, invention management, and high-stakes decision-making.

The report begins with the wider context. **Bowman Heiden** considers how AI is affecting the patent system, while **Ruud Peters** examines its implications for the IP profession, including the roles and skills that will matter as AI becomes part of everyday work.

The following chapters examine how these questions arise in practice. **Stephan Wolke** and **Stephan Ising** address the use of general-purpose AI in patent workflows. **Sakae Nakamura** considers how AI can expand access to patent intelligence and support technology landscaping. **Tim Alexander Oelmann** and **Daniel Uhl** explore the implications for inventors and invention disclosures within a corporate environment. **William Mansfield** examines the analytical foundations required to use patent data in flexible and meaningful ways.

These contributions do not prescribe a single model for adoption. The tasks, decisions, and organizational contexts discussed in the report differ substantially. A workflow intended to help a professional explore information has different requirements from an analysis used to compare competitors, support an investment decision, or assess a legal position. The appropriate approach depends on the purpose of the work, the underlying information, and the consequences of getting the answer wrong.

The report is also deliberately limited in scope. It does not seek to assess AI technologies as a licensing market or to provide a general account of AI development. Its subject is the application of AI within IP: how it changes work, how it changes access to information, and how IP teams may adapt their practices and responsibilities.

Read from beginning to end, the report follows a progression from the wider environment in which IP operates to the practical choices organizations face. It shows how questions about systems, people, workflows, information, and decision-making are connected. The report is intended to help organizations consider where AI can contribute to their IP work and what needs to change around it.

Holger Zimmer
Editor of this report & Global Director Strategic Communication & Partnerships
for LexisNexis® Intellectual Property Solutions



Bowman Heiden

Co-Director of the Center for Intellectual Property (CIP) at the University of Gothenburg, and Executive Director of the Tusher Strategic Initiative for Technology Leadership, UC Berkeley.

AI's impact on the patent system

By Bowman Heiden, Co-Director of the Center for Intellectual Property (CIP) at the University of Gothenburg, and Executive Director of the Tusher Strategic Initiative for Technology Leadership, UC Berkeley.

Artificial intelligence is frequently discussed as a new generation of productivity technology for intellectual property professionals. Patent searching can be accelerated. Drafting can be supported. Large portfolios can be analyzed more efficiently. Administrative processes can be automated. These developments are significant, but they describe only the first-order effects of AI.

The more fundamental question is what happens when AI becomes embedded not only in the administration of the patent system, but in every facet of business, including the process of invention itself. At that point, AI no longer merely changes how efficiently the patent system operates. It begins to affect the foundations on which the system has historically been built.

Precisely what the resulting system will look like remains uncertain. But the relevant question is no longer whether AI will affect patents. It is how the patent system will adapt when invention, patent creation, examination, and enforcement increasingly take place in an AI-enabled environment.

AI changes the economics of invention

The traditional patent system assumes that invention is a comparatively scarce activity. Inventors and R&D organizations invest expertise, time, and capital to develop technical solutions, and the patent system provides a temporary exclusive right in return for disclosure of qualifying inventions.

AI can alter this equation. Generative and computational AI systems can increasingly support researchers in identifying possible solutions, exploring technical alternatives, analyzing prior approaches, and generating candidate designs. The cost of exploring a large solution space can therefore decline substantially. This does not mean that human creativity disappears. Nor does it mean that every AI-generated suggestion becomes an invention. But it does mean that the economics of producing potentially patentable technical ideas may change.

If considerably more technical alternatives can be generated and assessed with the same amount of human effort, the patent system could face a paradox: invention becomes easier to produce, while demonstrating that an invention deserves patent protection becomes harder. The consequence could be that AI increases the rate at which technical solutions can be generated while simultaneously raising the threshold those solutions must clear to qualify for patent protection.

The benchmark for inventive step may move

AI complicates the benchmark for patentability. If engineers routinely use AI systems capable of reviewing large bodies of technical literature, identifying combinations across different technical fields and proposing possible solutions, the capabilities available to the skilled person may change as "ordinary" becomes "extraordinary".

This creates a fundamental question for patent law:

> *What constitutes an inventive step when sophisticated AI assistance becomes an ordinary tool of technical development?*

The answer will not necessarily require a formal rewriting of patent law. Patent systems have historically adapted their interpretation of concepts such as obviousness and inventive step as technologies and professional practices have evolved. Nevertheless, AI could gradually raise the effective bar for patentability. Solutions that once required significant ingenuity may eventually become relatively straightforward outputs of an AI-assisted development process.

It will be less about whether AI contributed to an invention, and more about whether the resulting invention represents a meaningful technical contribution in an environment where AI assistance is now normal.

More invention does not necessarily mean more patents

A common assumption is that easier invention generation will lead directly to higher patent filing volumes. In the short term, that is certainly possible. Lower drafting costs and easier generation of technical alternatives may reduce barriers to filing.

But the longer-term effect could be very different. If patent offices respond to increasingly AI-assisted invention by applying more demanding standards of inventive step, the number of patents that ultimately qualify for protection could stabilize or even decline.

This illustrates why the effects of AI on the patent system cannot be understood simply by counting applications. Filing activity, grant rates, claim scope, prosecution behavior, and portfolio value may develop in different directions.

For IP leaders, this makes patent quality increasingly important. In a world in which creating and filing patent applications becomes easier, the strategic value of individual patents may depend even more on their technical significance, legal robustness, and relevance to commercially important technologies.

Bowman Heiden
Co-Director of the Center for Intellectual Property (CIP) at the University of Gothenburg, and Executive Director of the Tusher Strategic Initiative for Technology Leadership, UC Berkeley.

AI will operate on both sides of the system

AI will not be used exclusively by applicants. Patent offices are also beginning to incorporate AI into searching, classification, examination, and administrative processes. Applicants may use AI to identify prior art, generate claims, develop prosecution arguments, and assess potential patentability. Examiners can increasingly use similar technologies to identify relevant documents, compare technical concepts, and evaluate applications.

The result is not simply faster processing. It creates a technological escalation on both sides of the patent process.

More capable applicant tools encourage more capable examination tools, which in turn require applicants to improve the quality and sophistication of their filings. Over time, this interaction could change both the speed and the standard of patent prosecution. AI could therefore make the patent system simultaneously more efficient and more demanding.

The information asymmetry around patents will decline

Another structural effect concerns access to patent knowledge. Historically, meaningful use of patent information required considerable expertise. Patent databases are complex. Ownership information can be difficult to interpret. Patent families, legal status, classifications, and claims require specialist knowledge. Sophisticated patent intelligence consequently remained concentrated in patent departments and specialist analytics teams.

AI can reduce this barrier considerably. Natural-language interfaces increasingly allow users without deep patent-search expertise to interact with highly complex patent information. Business leaders, R&D teams, corporate strategists, and other functions may now be able to ask questions of patent data that previously required specialist analysts to formulate and answer.

This has consequences beyond productivity. Patent information can move from being a specialist resource used predominantly within IP departments toward becoming a more broadly accessible component of corporate intelligence.

That matters because patents are more than legal rights. They constitute one of the world's largest structured repositories of technical information. Making that information significantly easier to interrogate could increase its relevance for technology strategy, competitive intelligence, R&D planning, and investment decisions.

Easier access makes reliable data more important, not less

The democratization of patent intelligence also introduces a critical dependency. AI can make patent information easier to interrogate, but it cannot eliminate deficiencies in the underlying data. Incorrect ownership information, incomplete legal-status information, inconsistently consolidated patent families, or poorly harmonized entities can generate answers that appear authoritative while being fundamentally misleading.

The apparent simplicity of asking a natural-language question can therefore obscure the complexity involved in producing a defensible answer.

This distinction will become increasingly important as AI interfaces become easier to use. The quality of AI-supported patent intelligence depends not only on the capability of the model, but also on the quality, completeness, currency, and structure of the patent data to which it has access. As the interface becomes simpler, the infrastructure behind it becomes more important.

Bowman Heiden
Co-Director of the Center for Intellectual Property (CIP) at the University of Gothenburg, and Executive Director of the Tusher Strategic Initiative for Technology Leadership, UC Berkeley.

AI may change the balance between quantity and quality

The same dynamic complicates the relationship between quantity and quality. If AI lowers the marginal cost of generating patent applications, patent offices and businesses alike will have to become more capable of distinguishing between volume and value. For companies, this raises strategic questions. Should lower drafting costs simply result in more filings? Or should AI instead allow IP departments to devote greater resources to identifying the inventions that matter most, understanding their competitive significance and constructing stronger portfolios around them? Over time, AI will likely be used to support these higher-value activities as well.

For patent offices, similar questions arise around examination capacity, and patent quality. AI therefore has the potential to increase both sides of the equation: the volume of potentially patentable output and the analytical capacity available to evaluate it. The long-term impact may consequently be less about producing more patents and more about raising expectations regarding which patents justify protection and investment. In theory, AI could allow patent offices to grant patents with such high quality that validity is no longer questioned.

Taken together, these developments show that AI does not affect only one stage of the patent system. It enters the system as patentable subject matter, as a contributor to invention and prior-art generation, as a tool for patent professionals, and as an examination capability for patent offices. The figure below summarizes these four direct points of impact.

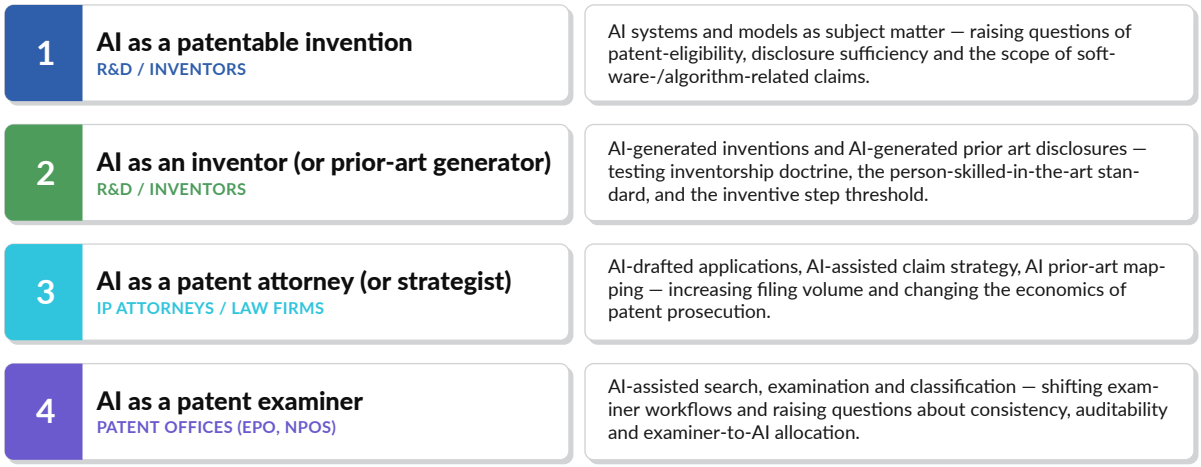


Figure 1: Four ways AI can enter the patent process

These points of impact will increasingly interact. The patent system must therefore adapt as an interconnected whole rather than treating each use of AI as an isolated development.

Bowman Heiden
Co-Director of the Center for Intellectual Property (CIP) at the University of Gothenburg, and Executive Director of the Tusher Strategic Initiative for Technology Leadership, UC Berkeley.

The patent system will have to adapt without undermining its purpose

Technological change has repeatedly required patent systems to adapt. AI is different in scale because it can influence several foundational stages simultaneously: how inventions are generated, how patent applications are produced, how prior art is discovered, how applications are examined, and how patent information is interpreted.

The objective should not be to preserve every existing process simply because it is familiar. Nor should AI-driven change be treated as evidence that the underlying rationale of the patent system has disappeared. The core policy objective remains to maintain an effective mechanism for incentivizing and disclosing meaningful technological innovation while providing sufficient legal certainty for inventors, companies, and markets.

As with any significant institutional change, the challenge will be avoiding overcorrection. Policy responses should be based on evidence of how AI is actually changing inventive activity and patent-system behavior rather than assumptions about what AI might ultimately be capable of doing. This will require careful observation of filing behavior, examination outcomes, inventive-step standards, patent quality, and the economic use of granted rights.

From an AI-enabled patent system to an AI-shaped patent system

Most organizations are currently experiencing AI primarily as a new tool. That is understandable. The immediate manifestations are better searches, faster analysis, assisted drafting, and automation.

But these are likely to be transitional effects. The deeper transformation occurs when AI simultaneously changes invention, patentability, filing behavior, examination, access to patent knowledge, and the economics of patent analysis. At that point, the patent system is no longer merely AI-enabled. It becomes increasingly AI-shaped.

The architecture of the patent system may remain recognizable – inventions will still be assessed, applications examined, rights granted, and patents enforced. But the volume of information, the capabilities available to participants, and the meaning of concepts such as routine technical skill may evolve substantially. That is why the impact of AI on patents should not be framed only as an efficiency discussion.

How far this transformation proceeds will depend on two separate dimensions: how capable AI systems become and how deeply those capabilities alter the patent system. More advanced AI does not automatically result in institutional transformation; it may simply automate existing processes. Conversely, even current Generative AI could make certain established activities unnecessary. The framework below separates these two dimensions.

Bowman Heiden
Co-Director of the Center for Intellectual Property (CIP) at the University of Gothenburg, and Executive Director of the Tusher Strategic Initiative for Technology Leadership, UC Berkeley.

	Automation	Innovation	Elimination
Hype AI progress stalls; overclaiming fades	Isolated gains in search, drafting, and administration	Experiments rarely scale	Established workflows remain
GenAI+ Current systems improve and spread widely	Faster analysis, drafting, and examination with human review	Natural-language intelligence and AI-supported IP decisions	Routine administration and first-pass review decline
AGI Near-human general capability arrives	End-to-end invention, filing and examination	Autonomous agents connect R&D, IP strategy, and patent offices	Routine search, drafting, and prosecution largely disappear
Superintelligence AI exceeds human capability across domains	Machine-scale invention and assessment	Patent standards and innovation incentives are redesigned	Legacy filing and examination models may become obsolete

Table 1: Four Possible Futures for AI in Intellectual Property. Adapted from Susskind, R. (2025). How to think about AI: A guide for the perplexed. Oxford University Press.

The framework is not intended as a forecast. It provides a way to test different assumptions about AI development against three possible forms of impact: automating existing work, redesigning how outcomes are delivered, and eliminating processes that no longer serve a necessary function.

The more consequential question is whether a system designed around human-scale invention and human-scale information processing can adapt to an environment in which machines dramatically expand both.

The answer will emerge gradually through patent offices, courts, applicants, and markets. But the direction of travel is already visible. AI is moving from being a tool used within the patent system to becoming one of the forces reshaping the system itself.

Bowman Heiden
Co-Director of the Center for Intellectual Property (CIP) at the University of Gothenburg, and Executive Director of the Tusher Strategic Initiative for Technology Leadership, UC Berkeley.



Ruud Peters

Co-Director at the Center for Intellectual Property (CIP), CEO at Peters IP Consultancy B.V., former Chief IP Officer and Executive Vice President at Koninklijke Philips N.V.

How AI is changing the work of IP professionals

By Ruud Peters, Co-Director of the Center for Intellectual Property (CIP), and CEO at Peters IP Consultancy B.V.; former Chief IP Officer and Executive Vice President at Koninklijke Philips N.V.

For many years, the work of intellectual property professionals has been defined by specialist expertise and experience.

Patent attorneys learned to draft and prosecute patent applications, conduct freedom-to-operate studies, and handle oppositions. Search specialists developed sophisticated methods for identifying prior art. Patent information professionals learned to navigate complex databases and classification systems. Licensing specialists developed expertise in building business cases for transactions, conducting negotiations, and drafting agreements.

Artificial intelligence will not make this expertise irrelevant overnight. But it will fundamentally change where and how professionals can create value. The important question is therefore not simply which activities AI can automate. It is what IP professionals should focus on, given that many activities that historically required substantial specialist effort are increasingly handled by AI.

> *I expect that essentially every activity we currently recognize within the IP profession will be affected by AI.*

Search already makes extensive use of AI technologies. Patent drafting and prosecution are increasingly handled by AI. We see the same trend in handling oppositions, detecting infringements, claim-charting, developing business cases for IP transactions, and drafting licensing and other IP agreements.

It therefore makes less sense to focus on a single part of the IP workflow that AI will transform. The more useful assumption is that the entire workflow will be transformed.

From performing tasks to applying judgment

Much of the discussion around AI in professional services focuses on efficiency. How much faster can a patent application be drafted? How much time can be saved in a search? How many documents can be reviewed automatically? These are legitimate questions, particularly in the short term. But they describe only the operational impact of AI. The more important development is that AI can gradually separate professional value from the execution of individual tasks.

Consider, for example, patent drafting. Historically, the ability to draft a 'strong' patent application was a substantial part of the value a patent attorney provided. Drafting required technical understanding, patent legal knowledge, and extensive experience. AI is increasingly taking over significant parts of that process.

That does not eliminate the need for an IP professional, but it will change the nature and level of their contributions. The professional increasingly needs to determine the broader portfolio strategy to support the business, whether a patent filing for the invention at hand aligns with that strategy, and, if so, how and where to protect it.

The same principle applies to other patent activities. If AI can produce a search result, professional value shifts to interpreting the results and the impact on the patent portfolio. If AI can generate a claim chart, professional value moves toward understanding its legal and commercial significance. If AI can analyze portfolios, professional value increasingly lies in determining what the results mean for your portfolio and business strategies. AI therefore shifts emphasis from producing information to interpreting information and making decisions.

Technical IP expertise will no longer be enough

This development reinforces a change that was already long overdue and partly underway before generative AI emerged.

> *The modern IP professional needs to understand much more than IP law and patent procedures.*

To contribute meaningfully to a business, an IP professional needs to understand the business itself and the environment in which it operates: its competitors, the technological direction in the relevant fields, R&D investments, and developments in the markets where it intends to compete.

IP professionals need to understand how IP can support those objectives, which IP value model is appropriate in a particular situation, and what value IP brings to the business. To do this, they need sufficient understanding of business financials to discuss the contribution of IP in the language business leaders use. AI makes these capabilities more important, not less. As AI becomes better at performing specialist analytical tasks, expertise that is difficult to automate becomes relatively more valuable. Context, judgment, strategic understanding, communication, and the ability to connect IP decisions with business consequences therefore move toward the center of the profession.

This creates a challenge for the way IP careers have traditionally developed. In many organizations, IP professionals begin by developing deep expertise in the narrow area of patent drafting and prosecution for many years before gradually expanding into IP management or strategy. In the future, organizations may need to create broader exposure much earlier.

Ruud Peters
Co-Director at the Center for Intellectual Property (CIP), CEO at Peters IP Consultancy B.V., former Chief IP Officer and Executive Vice President at Koninklijke Philips N.V.

Indeed, it may sometimes be easier to teach IP to someone who already understands business than to teach business thinking to someone whose professional development has remained entirely within the IP function.

That is a provocative proposition, but it points to an important change in the profession. IP expertise remains valuable, but without business understanding it is no longer sufficient.

AI will make IP knowledge accessible far beyond the IP department

It is not only AI that drives the changing role of IP professionals. AI is also changing who has access to IP knowledge. Historically, it was difficult for non-specialists to use patent information. Even relatively straightforward questions required knowledge of patent families, legal status, classification systems, ownership structures, claims, and specialized search syntax.

This created information asymmetry: IP professionals and patent information specialists could access and interpret information that business executives, R&D managers, or corporate strategists often could not use independently. AI is reducing that asymmetry. One broader effect of AI is the democratization of knowledge. Information that previously required specialist expertise to retrieve and interpret is increasingly becoming accessible through natural-language interaction.

The same is happening with IP. Business and R&D people will increasingly be able to interact directly with patent data, ask questions, and request analyses without involvement of IP specialists. The AI tools will support them in learning the traditional mechanics of patent searching or analytics. For IP professionals, this should not be viewed primarily as a threat. It is an opportunity to extend the influence of IP throughout the organization. If technology leaders, R&D managers, and business executives can more easily access patent intelligence, IP information will become part of business decisions and it also will change the relationship between the IP function and its internal stakeholders.

The IP department will increasingly cease to be the gatekeeper through which every patent-related question must pass. Its value will instead come from ensuring that IP information is correctly interpreted and connected to the business strategy.

Easier access means expertise shifts

There is a risk in assuming that a simpler interface means a simpler underlying problem. It does not. A business executive may soon be able to ask an AI system, in ordinary language, which competitors have the strongest patent positions in an emerging technology. The interaction itself may take only seconds.

But the quality of the answer still depends on multiple complex factors. Which patents should belong to the relevant technology set? How should corporate ownership be consolidated? Which patent families are still active? How should different jurisdictions be treated? Which indicators should be used to compare portfolios? Is the analysis measuring patent volume, quality, competitive impact, or something else?

These questions were relevant in the pre-AI era and remain relevant in the AI era. Business executives will be sufficiently versed in AI to know they cannot simply rely on what AI tools deliver. They know that they have to ask the right questions to get the right answers. IP professionals need to help them understand the IP field well enough to ask better questions. In my view, the wrong approach would be for IP departments to try to prevent or dissuade business and R&D managers from using IP information generated by AI tools. Such efforts will be seen as defensive and will have the opposite effect: keeping IP people out of these matters altogether.

Ruud Peters
Co-Director at the Center for Intellectual Property (CIP), CEO at Peters IP Consultancy B.V., former Chief IP Officer and Executive Vice President at Koninklijke Philips N.V.

One factor will remain the same and that is the quality of the underlying data. Reliable IP intelligence requires complete, accurate, and current information about matters such as inventorship, assignments, ownership, and legal status. The underlying metadata must be trustworthy. This becomes particularly important with generative AI because an answer can sound fluent and confident even when the underlying information is incomplete or incorrect.

The professional skill therefore moves from knowing how to manually retrieve every piece of information toward helping others understand when an answer can be trusted.

Patent intelligence specialists move closer to decision-making

This transformation is particularly significant for professionals working in patent intelligence and patent analytics. Traditionally, much of their value has come from their ability to formulate complex searches, construct technology landscapes, clean patent datasets, interpret ownership structures, and generate analytical outputs.

Most of these activities will increasingly be done by AI, but that does not make the patent intelligence function necessarily less relevant. If the technical barriers to analysis fall, more sophisticated analysis and reporting can be done around technology development, competitive positioning, partnerships, M&A, R&D investments, portfolio development, and licensing using patent and other IP data.

> *The role of the analyst therefore evolves from providing patent landscapes and other patent statistics to delivering knowledge derived from interpreting the data.*

That is a fundamentally different position. The most valuable patent intelligence professionals will not simply be those who know how to operate the analytical tools. They will be those who understand what the underlying business is and can translate that question into an analytically robust IP perspective.

AI should support a more proactive approach to IP

The development also creates an opportunity to address a long-standing weakness in many IP organizations. Too many companies still manage their IP portfolios reactively. They select the best inventions from the incoming stream of submissions to file as patents. Portfolios built this way are mostly scattered across the various R&D projects. Decisions about where to invest in R&D are generally taken without involvement or input from the IP department.

A more strategic approach works in the opposite direction. The organization first considers where the business intends to compete in the future. It identifies the technologies and market positions that will matter and determines where IP could create strategically important control points. The IP department can then develop the portfolio deliberately around those future focal areas.

AI can make this form of active portfolio management considerably more practical. Large amounts of patent, technology, and competitive information can be analyzed more quickly. Emerging fields can be monitored more continuously. Competitor movements can be identified earlier. Portfolio gaps can be examined more systematically.

These capabilities should not merely automate existing processes. They should help IP organizations move from reacting to inventions to actively shaping portfolios that support the business strategy.

Ruud Peters
Co-Director at the Center for Intellectual Property (CIP), CEO at Peters IP Consultancy B.V., former Chief IP Officer and Executive Vice President at Koninklijke Philips N.V.

The value of IP must still be explained in business terms

Technology alone will not solve another long-standing challenge for IP departments: communicating their contribution to senior management.

IP professionals frequently speak a language that other functions do not understand. The solution is not to educate every board member in the details of patent law. IP professionals need to become better at translating IP into the language of business. That includes financial language. IP departments should be able to explain what they contribute to the business, how IP supports commercial objectives, and what return the organization receives from its investment in IP.

AI may make some of the supporting analysis easier, but it does not remove the responsibility to make the IP business case more transparent. In fact, it may raise expectations. If information becomes easier to obtain, management will reasonably expect the IP function to provide clearer answers about what the IP function contributes to the business and what it will do to improve that contribution.
The future IP professional therefore needs to be comfortable moving between business, technology, strategy, finance, and IP.

What happens to traditional IP roles?

For young professionals considering a career in IP, this raises an uncomfortable but necessary question.

Would I recommend becoming an IP professional today? Yes—but not necessarily an IP professional as we generally know them today.

I would be cautious about building a career exclusively around becoming an expert in activities such as patent drafting and prosecution. Those activities will continue to exist for a while, but AI will materially change how they are performed with decreasing involvement from IP professionals.

That does not mean the profession becomes less interesting. Quite the opposite.

> *New roles will emerge around the role of IP in business strategy, including providing IP intelligence to develop the business strategy, translating that strategy into building an IP portfolio, capturing value from that portfolio, and reporting and forecasting current and future value captured to higher-level management.*

AI will accelerate and support all these activities in ways that were not possible before. IP professionals should embrace these developments and use them instead of clinging to current practices, which are doomed to fail.

The IP profession will evolve, and those new roles will be even more fascinating than the traditional ones, as they will elevate the role of the IP professional to a higher level, directly involved in the business.

Ruud Peters
Co-Director at the Center for Intellectual Property (CIP), CEO at Peters IP Consultancy B.V., former Chief IP Officer and Executive Vice President at Koninklijke Philips N.V.

The IP professional becomes an orchestrator

A useful way to think about this transition may therefore be to stop asking whether AI will replace IP professionals. That question is too simplistic. AI will replace many activities, especially repetitive tasks, and it will accelerate and transform others. It will raise questions about the contribution the IP function makes to the business and the company overall, which will require the IP department to focus more on how to use and measure the value created for businesses through IP.

But organizations will still need people who understand intellectual property well enough to know what questions should be asked, whether the resulting answers can be trusted and how those answers should influence business decisions.

The IP professional increasingly becomes an orchestrator at the interface between several domains, including business, technology (R&D), and IP (broadly, not just IPRs). That requires a broader professional profile than many traditional IP roles demanded. The IP professional will not disappear, but its role has to change from being a specialist in the narrow field of IPR to a specialist who can combine that expertise with business understanding, financial literacy, strategic thinking, and the ability to use AI critically.

The opportunity is to raise the level of IP practice

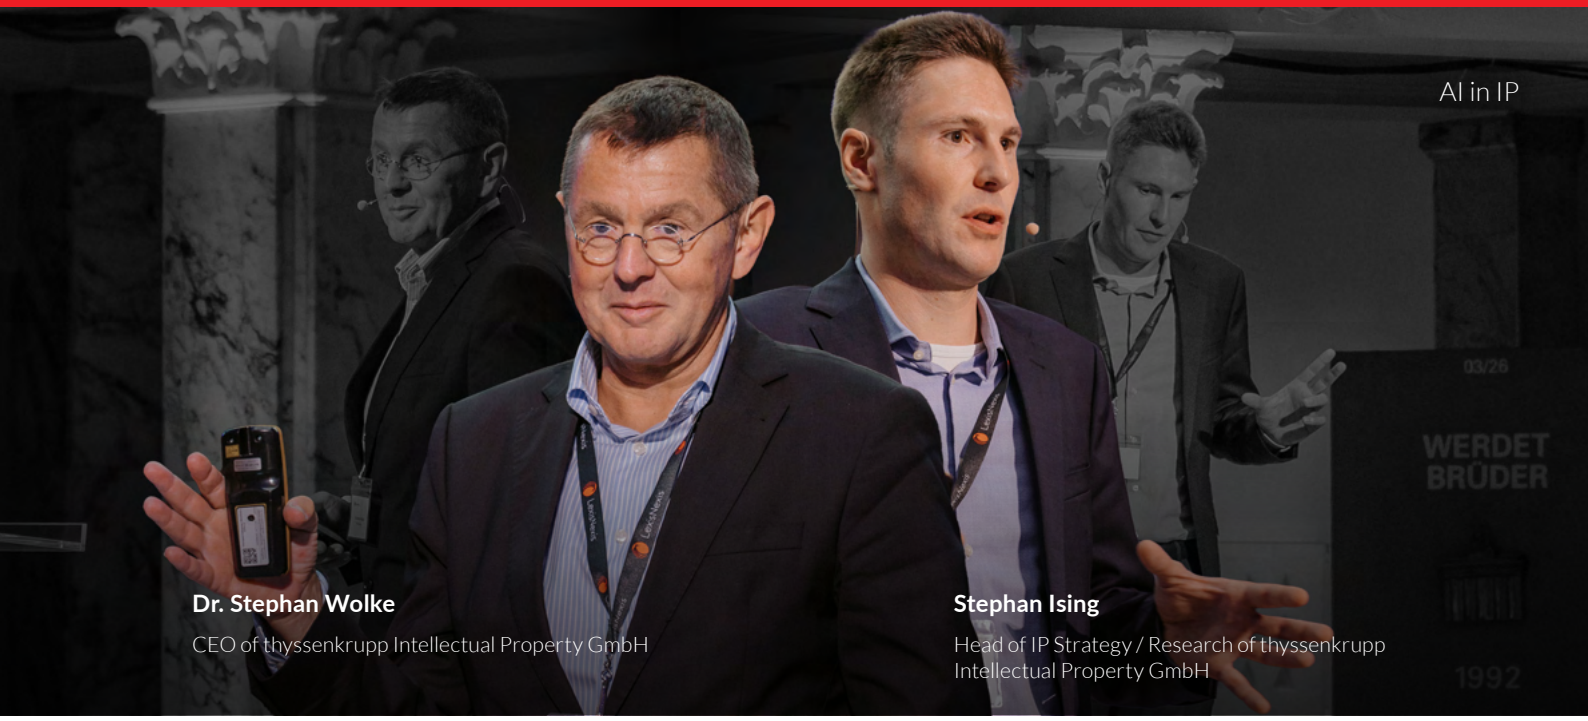
Throughout my career, one of my objectives has been to help companies use intellectual property more effectively for the benefit of their businesses. That requires continuously raising the level at which IP is practiced.

AI offers another opportunity to do exactly that. If it is used only to reduce the time required for drafting, searching, or administration, the profession will capture only a fraction of its potential value. Many traditional workflows will need to be redesigned when AI tools are increasingly used to perform activities.

The larger opportunity is to use AI to move IP professionals closer to the strategic decisions of the organization: understanding where the company is going, identifying how IP can support that direction, creating portfolios deliberately around future opportunities, and communicating the resulting value in terms that the business understands.

The central purpose of the IP function should become clearer rather than less important: **to use intellectual property as an instrument for creating business value.**

Ruud Peters
Co-Director at the Center for Intellectual Property (CIP), CEO at Peters IP Consultancy B.V., former Chief IP Officer and Executive Vice President at Koninklijke Philips N.V.



Dr. Stephan Wolke
CEO of thyssenkrupp Intellectual Property GmbH

Stephan Ising
Head of IP Strategy / Research of thyssenkrupp Intellectual Property GmbH

Using general-purpose AI across the IP value stream

By Dr. Stephan Wolke, CEO of thyssenkrupp Intellectual Property GmbH, and Stephan Ising, Head of IP Strategy / Research of thyssenkrupp Intellectual Property GmbH

Practical experience from a corporate IP department

AI affects an IP department in two fundamentally different ways. The first is **IP for AI**: applying the established IP mandate to AI-enabled solutions developed in the business—identifying and protecting valuable innovations, respecting third-party rights and, where necessary, enforcing the resulting IP. The second is **AI for IP**: using AI itself to support IP work, from invention disclosure and patent drafting through prosecution, search, and monitoring to portfolio analysis and IP strategy.

At thyssenkrupp, we view both perspectives along the same basic IP logic: **protect our own innovations – respect third-party rights – detect infringements of our IP**. AI therefore changes both what we protect and how we perform the work of protecting it.

This contribution focuses primarily on AI for IP. Our starting point has been deliberately pragmatic: **we are not an AI company; we do not develop large language models (LLMs); we use them**. Rather than building a bespoke patent-specific LLM, we have mainly worked with available general-purpose models and other AI technologies and adapted them to concrete IP tasks. While our starting point is general-purpose AI, some examples also show where analytical or specialized methods are better suited to the task.

The relevant question was therefore not simply what AI could do, but where it would make sense to use AI within our actual IP processes.

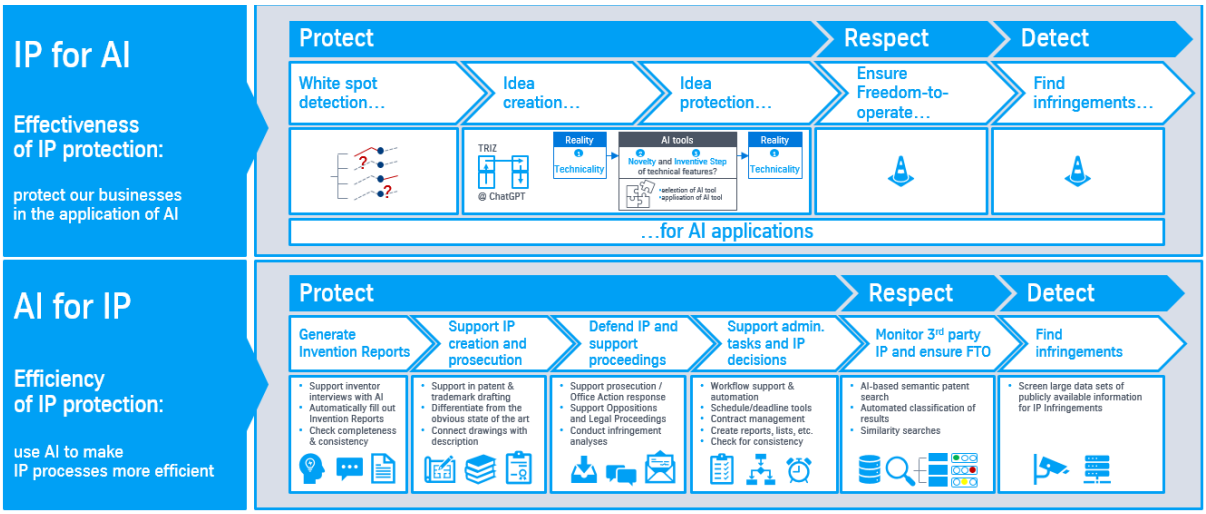


Figure 1: AI for IP and IP for AI across the Protect–Respect–Detect value stream

AI for IP: start with the process

Our approach did not begin by compiling a list of ChatGPT capabilities. We first looked at the work already being performed in the IP department. Using Protect–Respect–Detect as the overall framework, we reviewed our processes and asked three pragmatic questions: How much manual effort does an activity require? Where could AI sensibly support it? How much effort would be required to change the process?

Manual effort was estimated roughly in working hours, while implementation effort was assessed qualitatively. The objective was not to calculate a precise business case for every possible use case, but to identify activities where significant manual effort coincided with a realistic AI opportunity and manageable implementation effort. Search, analysis, and monitoring emerged as obvious early candidates because they are frequent, data-intensive activities with substantial preparatory work.

We soon found, however, that labels such as **patent drafting**, **search**, **monitoring**, or **prosecution** were themselves too broad to be useful units for AI implementation. A patent application, for example, results from gathering and structuring information, understanding the invention, considering prior art, defining a protection concept, drafting different parts of the application, dealing with drawings, and checking consistency and formal requirements. These activities require different information, different forms of analysis, and different combinations of human and AI work.

This exposes a recurring misconception: treating professional work as an information-in, finished-product-out problem. If an invention disclosure is supplied to a model with the expectation of receiving a perfect patent application, disappointing output may reflect the definition of the task as much as the capability of the model.

The more useful question is therefore not whether AI can draft a patent or respond to an office action, but which individual tasks within those processes can be supported, what information each task requires, and where professional judgement must remain decisive.

We consequently started to assess AI opportunities in two dimensions: **where the task sits in the IP process** and **what kind of operation it requires** – for example, extracting, summarizing, comparing, classifying, validating, generating, or discovering relationships. Different operations may require very different technical methods.

Dr. Stephan Wolke, *CEO of thyssenkrupp Intellectual Property GmbH*
Stephan Ising, *Head of IP Strategy / Research of thyssenkrupp Intellectual Property GmbH*

Our working logic became:
process → task → type of operation → appropriate method

rather than:
AI tool → search for a use case

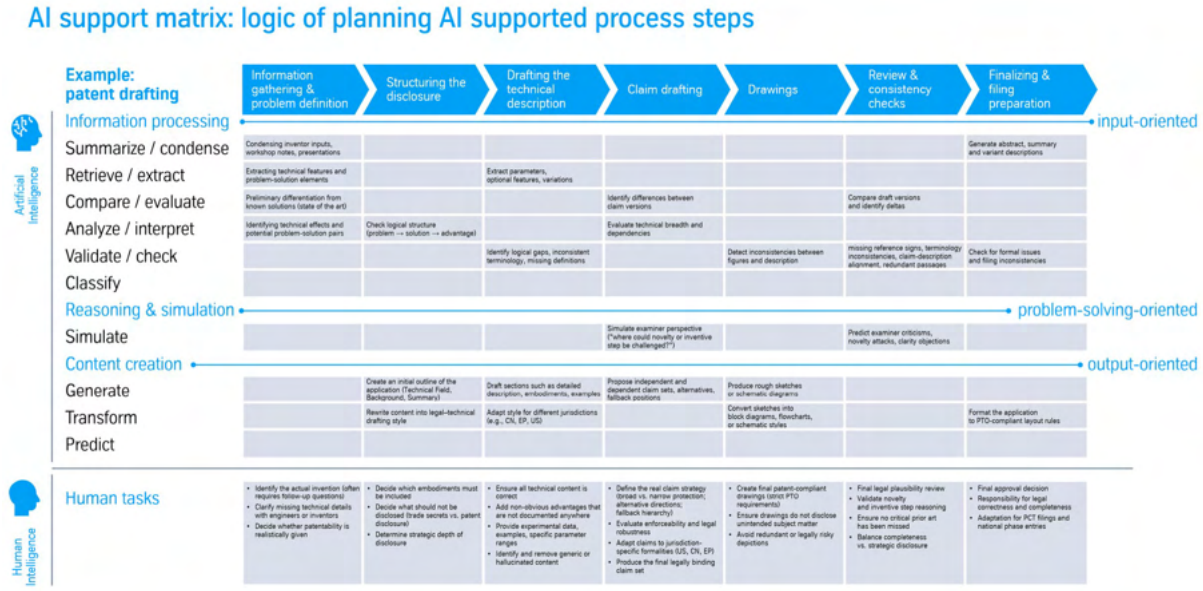


Figure 2: Task-based planning of AI support across the patent drafting process

From experimentation to repeatable methods

Our first use of generative AI in IP work was unsurprisingly simple: we started with chat. That remains appropriate for many activities. Summarizing a document, preparing questions for an inventor meeting, explaining an unfamiliar technology, or challenging an initial idea does not necessarily justify a dedicated application. For such tasks, flexibility is often more valuable than standardization.

The situation changes when the same task occurs repeatedly. We then began replacing spontaneous instructions with more structured prompts defining context, objective, analytical steps, and expected output. The value of such prompting should not be mystified: a good prompt does not make the model magically more intelligent; it makes part of the professional **working method explicit**.

AI can also help develop that instruction. We have used meta-prompting to turn informal task descriptions into more structured methods by providing examples, describing typical mistakes and asking the model to identify ambiguities or missing criteria. Recurring tasks are then tested against known or **gold-standard cases**. Once an approach proves useful, it can be shared rather than remaining in an individual user's chat history.

For some applications, the instruction alone is not sufficient. The model also needs structured domain knowledge. In our classification work, we use the working term “knowledge container” for structured domain knowledge supplied to the model as context, such as definitions, characteristic features, terminology, exclusions, examples, and distinctions from neighboring categories.

The distinction is useful: **the prompt describes what the model should do; the knowledge container provides domain-specific information it should use when doing it.**

Dr. Stephan Wolke, *CEO of thyssenkrupp Intellectual Property GmbH*
Stephan Ising, *Head of IP Strategy / Research of thyssenkrupp Intellectual Property GmbH*

Human review then becomes more than a final quality check. If reviewers repeatedly observe the same error, the underlying prompt, category definition, or knowledge container can be revised. The individual result is corrected, but the method improves as well.

A further transition occurs when a task is sufficiently stable but manual interaction becomes inefficient. A classification prompt may work well for twenty patents and still be impractical for several thousand. At that point, surrounding software can read the patent data, add the validated instruction and context, call the model through an API, and write the structured result back into the dataset.

Our practical development path has therefore often looked like:

explore interactively → structure the method → test → standardize → automate where volume and stability justify it

This is not a maturity model. An occasional task may remain best served by chat, while another may require an entirely different technical architecture.

The organizational side developed in parallel. We did not begin with elaborate AI governance, but with two clear principles: colleagues were allowed to experiment, and professional responsibility remained unchanged. Experimentation took place within boundaries regarding confidentiality, approved tools, and sensitive information. Ownership also stayed close to the underlying IP work, with colleagues responsible for AI activities in areas such as invention disclosure, search, drafting, office actions, and IT infrastructure.

Adoption followed usefulness rather than formal usage targets. Faster preparation, more focused monitoring, and less repetitive work proved more persuasive than generic AI demonstrations.

Practical use cases across the IP value stream

Supporting inventors

The invention disclosure process shows how the same general-purpose model can perform very different roles.

Before a formal invention report exists, AI can act as a challenger. An inventor may have a promising idea but may not yet have articulated the underlying problem, technical effect, critical parameters or alternative embodiments clearly enough. A conversational model can ask targeted questions and help make the concept more explicit. The objective is not to let AI decide whether the idea is patentable, but to identify aspects relevant to subsequent technical and patentability assessment.

AI can also make established innovation methodologies easier to apply. TRIZ is one example. The methodology itself is not new; what AI changes is the effort required to use it. In one example involving a hammer crusher, applying TRIZ principles with a general-purpose model produced directions such as a modular rotor shaft, variable geometries and integrated additional functions. These were treated as starting points for engineers, not as finished AI-generated inventions.

A second pattern uses AI as an interviewer. Instead of asking an inventor to complete an empty disclosure form, the model can conduct a structured conversation, follow up where information is missing, and organize the answers into the required format. A third pattern applies the same logic to technically thin disclosures – the familiar Bierdeckel-Erfindung, or "beer-mat invention" – by asking for missing embodiments, effects, alternatives or boundary conditions. Finally, AI can turn transcripts of inventor–attorney discussions into structured technical memos for subsequent drafting.

***Dr. Stephan Wolke**, CEO of thyssenkrupp Intellectual Property GmbH*
***Stephan Ising**, Head of IP Strategy / Research of thyssenkrupp Intellectual Property GmbH*

The AI therefore acts, depending on the stage, as **challenger**, **interviewer**, **structuring assistant**, **completeness checker**, **or summarizer**. It does not replace the inventor's technical contribution; it changes the interface through which that contribution is captured.

Patent drafting and prosecution

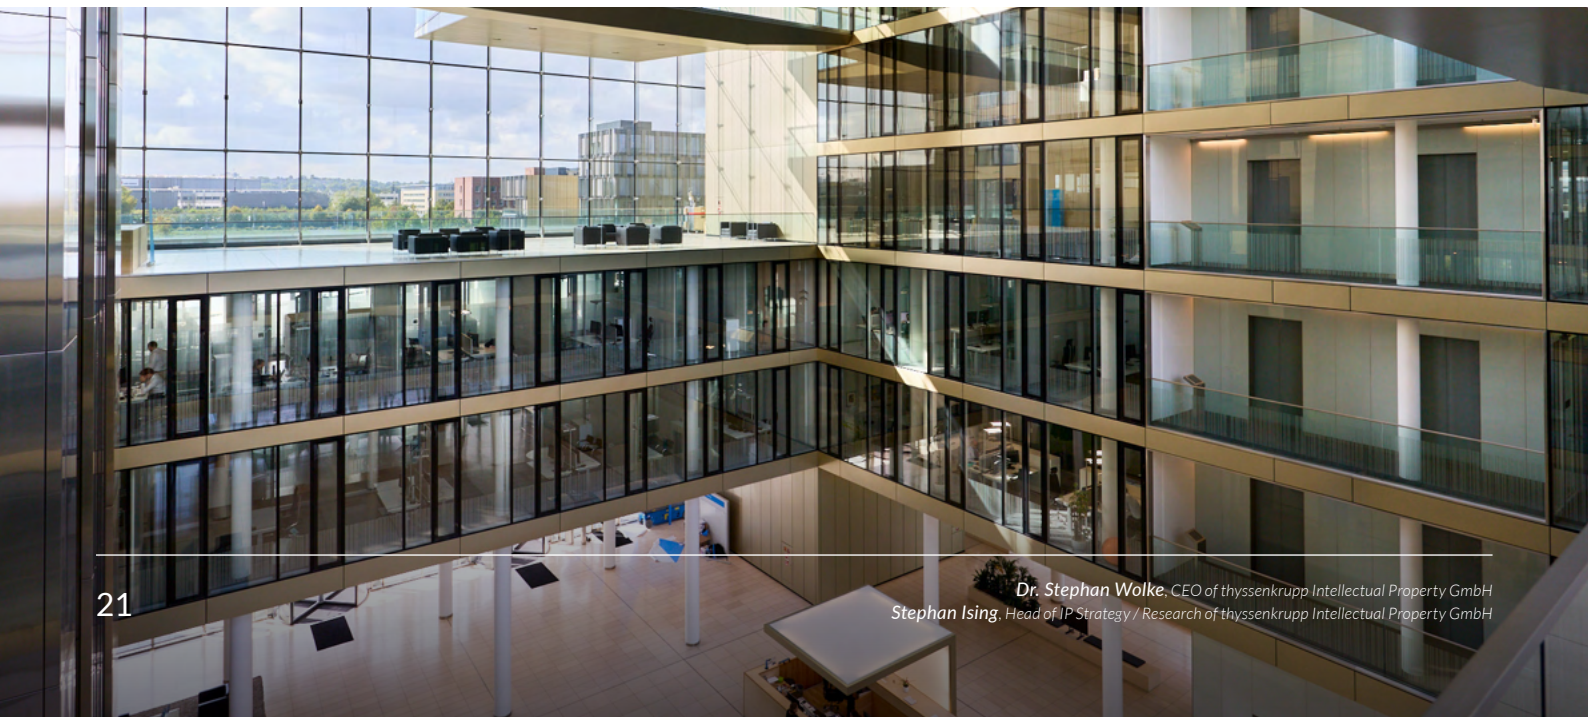
Patent drafting illustrates especially clearly why task decomposition matters. The intuitive expectation of invention disclosure in patent application out proved to be the wrong abstraction.

In our approach, the claim set remains a central responsibility of the patent attorney because it involves strategic decisions about scope, fallback positions, enforceability, and the overall protection concept. Once the claims and relevant technical information exist, however, AI can support subsequent steps: preparing parts of the background, assisting with the summary, describing embodiments from drawings and technical information, and checking terminology, grammar, reference signs, or internal consistency.

The relevant information differs from step to step. A model asked to describe an embodiment needs different input from one asked to check reference signs. Drafting therefore becomes an iterative collaboration between human judgement and AI rather than a one-shot generation exercise.

The same principle applies to prosecution. Rather than asking a model directly to respond to an office action or comparable examiner communication, we separate the work into analytical stages. AI first structures the communication, including cited documents, affected claims, and examiner objections. It then analyses the cited documents individually and can subsequently decompose relevant claims into technical features and map them against the prior art. The resulting feature matrix helps expose possible distinguishing features or amendment positions.

Only then does the workflow move towards potential response options. The model may prepare alternative amendment directions or the basis for an argumentative response, but we deliberately prefer options over one apparently definitive recommendation. The patent attorney assesses their legal and strategic implications and decides how to proceed.



***Dr. Stephan Wolke**, CEO of thyssenkrupp Intellectual Property GmbH*
***Stephan Ising**, Head of IP Strategy / Research of thyssenkrupp Intellectual Property GmbH*

The useful unit of AI support is therefore not **“respond to the office action”**, but **prepare the analytical basis and available options for the attorney who decides how to respond**.

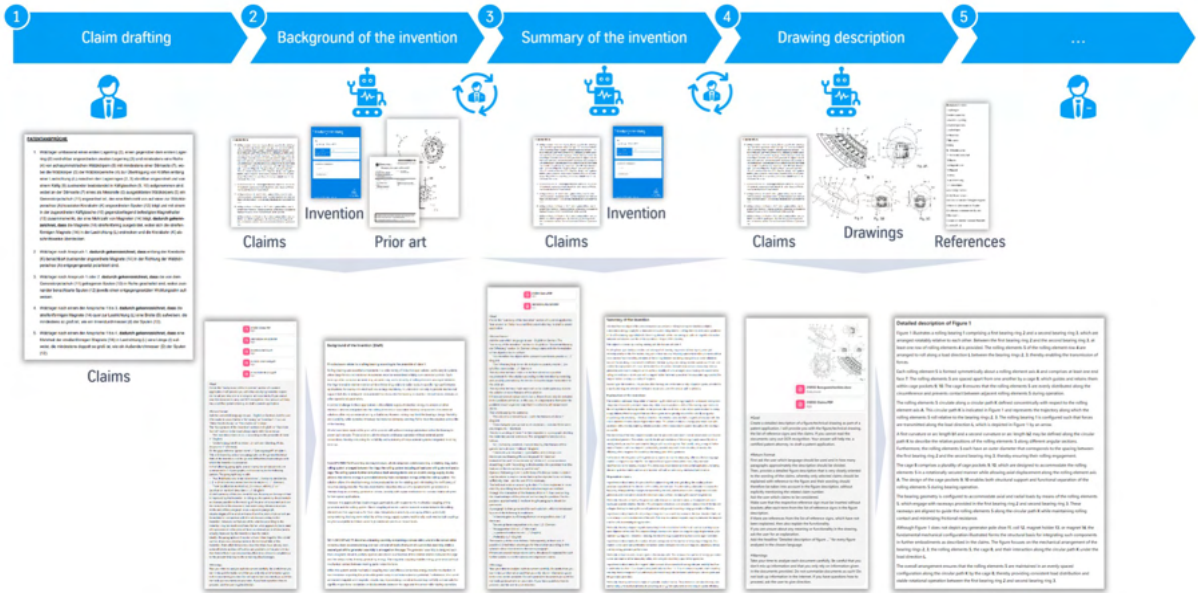


Figure 3: AI-supported patent drafting as a sequence of task-specific human–AI interactions

Search and IP research

Patent research also begins before the first database query. An internal request may consist of an email, presentation, drawing, or short technical description, while the terminology used by engineers may differ substantially from patent language.

AI can first support what we call subject simplification: extracting the technical core from a broader request and converting it into concepts that can actually be searched. It can then support search-strategy development through synonyms, alternative terminology, IPC/CPC suggestions, or Boolean queries. Commercial semantic search functions provide useful additional retrieval routes and cross-checks, but we do not treat them as replacements for professional search strategy.

Once documents have been found, the problem changes from where to search to why a document matters. AI can summarize disclosures, extract features, identify relevant passages, and prepare comparisons. We increasingly focus on evidence extraction rather than generated conclusions. “D1 appears relevant” provides little value; identifying which features are disclosed in which passages creates something a professional can verify.

Relevant patents can also be converted into structured data containing technical categories, problem statements, solution principles, applicants, legal status, or project-specific relevance. The result can then support classification, monitoring, competitor analysis, dashboards, or strategy instead of remaining merely a collection of PDFs.

The workflow therefore becomes:

understand request → simplify subject → develop search strategy → retrieve → extract evidence → structure → assess → report

The legal boundaries remain clear: semantic or technical similarity does not determine novelty, a relevance classification is not an infringement finding, and a search result indicating that no blocking patent was identified is not a freedom-to-operate (FTO) opinion.

Classification and monitoring

Patent classification is one of the areas where AI becomes particularly valuable at scale. For hundreds or thousands of publications, manual classification quickly becomes a substantial analytical workload.

Our first lesson was that the difficult part is often not the model but the taxonomy. Categories that appear obvious to an experienced analyst may contain implicit assumptions, overlapping boundaries, or inconsistent terminology. Before AI can classify consistently, much of this knowledge has to become explicit.

Knowledge containers can capture definitions, characteristic features, exclusions, and distinctions. An existing patent portfolio or other curated dataset can help generate an initial version, which experts then review. The process therefore looks less like **patent → AI → label** and more like:

domain knowledge → explicit classification logic → patent → AI-supported classification → review

Once the method is stable enough, a surrounding application can process the dataset through an API. We normally request not only the category but also a concise technical justification, because the objective is a r**eviewable classification**, not simply an AI label.

Human review feeds back into the method. Repeated false positives may indicate that a category is too broad; repeated omissions may expose missing technical criteria. The cycle becomes:

classify → review → identify recurring error → refine method → classify again

Once sufficient reliable labelled data exists, a different architecture becomes possible. In one approach, patent texts were converted into vector representations using the patent-specific embedding model PatentSBERTa. These representations were then used to train an XGBoost classifier. Such a specialized model can be evaluated with held-out test data and methods such as confusion matrices, which show not only overall performance but different types of error.

This is not a simple maturity progression from LLM to classifier. General-purpose models are attractive while definitions still evolve or rich contextual instructions are needed. A specialized classifier becomes interesting once the task is stable and sufficient labelled data exists, particularly where reproducibility, benchmarking, processing cost, or local deployment matter.

Classification then becomes an **enabling layer for patent intelligence**. Consistently structured data can support landscapes, competitor analysis, technical review prioritization, trends, and management dashboards.

Patent monitoring demonstrates the process effect particularly clearly. In a conventional workflow, an IP professional may screen new publications before routing relevant documents to technical experts. AI-supported classification can perform much of this initial structuring in advance, allowing experts to spend less time deciding what deserves attention and more time assessing what the relevant patents mean.

This became an example in which inserting AI into the old workflow was less valuable than redesigning the workflow around the new capability. We describe such a process as AI-ready: AI performs much of the initial structuring while human expertise is concentrated where substantive judgement creates the most value.

***Dr. Stephan Wolke**, CEO of thyssenkrupp Intellectual Property GmbH*
***Stephan Ising**, Head of IP Strategy / Research of thyssenkrupp Intellectual Property GmbH*

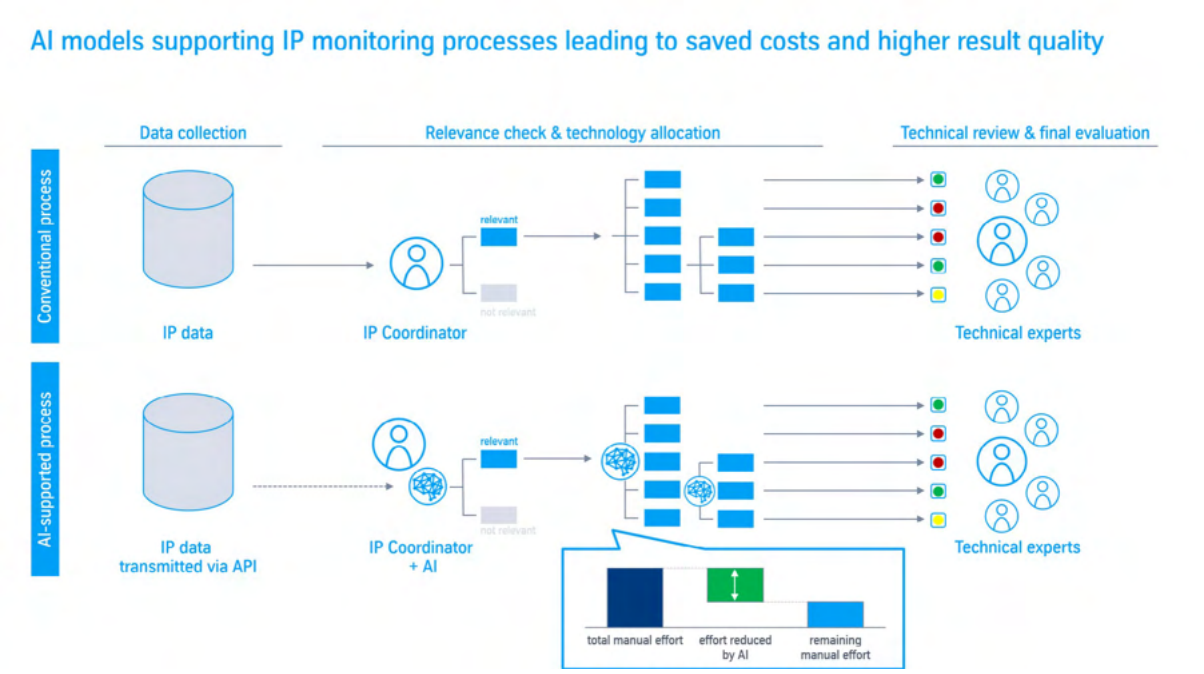


Figure 4: Conventional and AI-supported patent monitoring workflows

Patent translation as a multi-step workflow

Patent translation provides a compact example of moving from one AI task to a quality-controlled workflow. Machine translation itself is not new; the more interesting question is what happens around it.

In our implementation, long documents are segmented, translated section by section, and reassembled. Further steps then check substantive consistency, technical terminology, formal requirements and the final output:

document intake → segmentation → translation → reassembly → substantive review → terminology review → formal review → final quality control

Each stage addresses a different failure mode. A translation may be fluent but technically wrong, correct sentence by sentence but terminologically inconsistent, or substantively correct while failing formal requirements.

The durable asset is therefore not the particular orchestration platform, but the workflow logic. The fact that a task can be technically automated does not mean that it should be delegated professionally. An internal working translation and a document used in a legally significant process may be technically processed in similar ways while requiring very different levels of verification.

From patent analytics to strategic insight

Not every AI application starts with a prompt. Embeddings provide a different analytical approach by representing text as vectors and allowing semantic relationships across large datasets to be compared mathematically.

One straightforward use is similarity. A patent, claim, or technical problem description can be compared with a larger dataset to identify neighboring documents even where terminology differs. Similarity should not be overinterpreted, however: a semantically close patent is not automatically novelty-destroying prior art. It identifies a candidate for review.

Embeddings also support clustering, which differs fundamentally from classification. Classification starts with a predefined taxonomy and asks where a document belongs; clustering begins with the data and asks which documents appear to form groups. Such clusters can reveal structures that were not obvious beforehand, but still require expert interpretation.

This suggests a useful sequence:

explore → understand → codify → classify

An embedding analysis may reveal recurring groups. Experts interpret them, convert meaningful distinctions into a taxonomy or knowledge container, and that structure can subsequently be applied consistently through classification.

For strategic work, whole-document similarity may still be too coarse. Two patents can appear similar because they concern the same component while solving different technical problems. Conversely, patents from different fields may rely on the same solution principle.

We have therefore explored separating the problem addressed from the solution approach and representing both independently. This allows us to ask whether competitors solve the same problem in the same way, use alternative approaches, transfer one solution principle across different applications, or address relevant problems absent from our own portfolio.

The similarity measure itself is not the strategic insight. It exposes patterns from which better questions can be derived.

AI-supported patent analytics can then contribute to four broader strategic functions:

understand → diagnose → innovate → decide

To **understand** means turning large datasets into situational awareness regarding technologies, competitors, portfolio concentrations, and emerging themes. To **diagnose** means asking what these structures imply: Which problems recur? Which solution paradigms dominate? Where are competitors converging or diverging?

To **innovate** means using that understanding to expand the problem and solution space. Patent intelligence can be combined with functional decomposition, failure modes, physical constraints, customer needs, cross-domain analogies, or methodologies such as TRIZ. AI can support this exploration, but engineering judgement remains essential.

This is particularly important in white-space analysis. An empty area in a patent map is not automatically an opportunity. Patent documents represent **documented technical space**, not the full range of physical possibilities, customer needs, manufacturing constraints, economics, or future regulation. A visible gap may represent an opportunity, an unattractive technical direction, incomplete data, or simply a visualization artefact.

A more useful progression is:

describe topology → analyze patterns → expand problem space → test engineering feasibility → assess strategic and IP value

or more compactly:

patent gap → technical hypothesis → feasible solution space → strategically valuable IP opportunity

A gap in a patent landscape is not yet a strategic white space.

Finally, AI can support the **decision** stage by reducing complexity before human judgement is required. Structured patent data can help identify competitor developments deserving attention,

***Dr. Stephan Wolke**, CEO of thyssenkrupp Intellectual Property GmbH*
***Stephan Ising**, Head of IP Strategy / Research of thyssenkrupp Intellectual Property GmbH*

patent families requiring expert review, expired, lapsed, or otherwise non-blocking third-party positions, portfolio gaps, or technologies warranting deeper R&D analysis.

Dashboards are useful here, but their value depends primarily on the quality of the underlying classification and data structure. A strategically relevant patent does not automatically need to be maintained; a potentially non-blocking position is not an FTO opinion; and an apparent portfolio gap does not automatically justify additional filings.

> *AI can help us understand the documented problem space. Strategy begins when we connect that space to engineering reality and business value.*

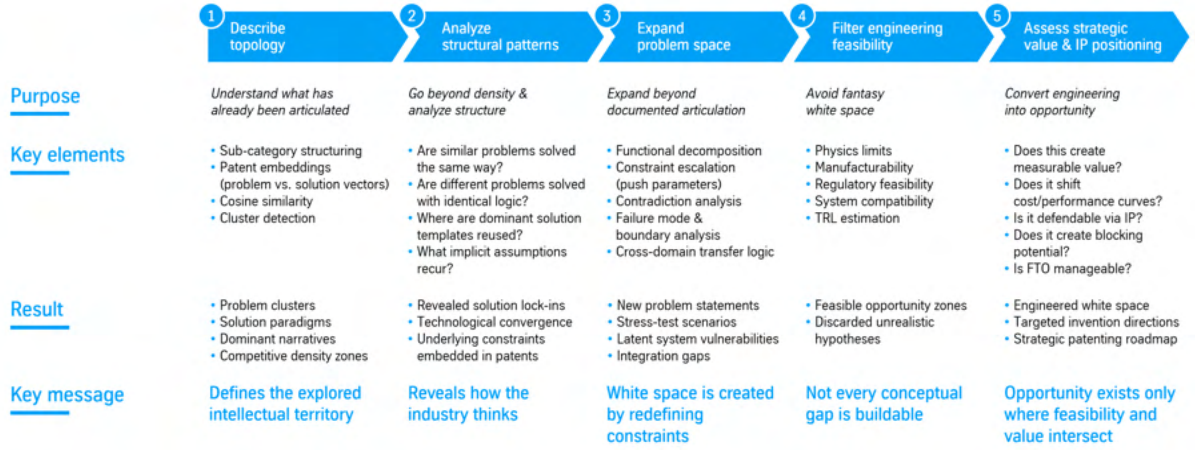


Figure 5: From patent mapping to engineered opportunity

Choosing the right approach and building capability

The range of possible AI applications is broad, but AI is not the right method for every IP problem. We have encountered use cases where early search or classification approaches produced too many irrelevant results to provide sufficient value. In such cases, it was not always possible to determine whether the main cause was the model, the data, the instruction, or the task definition.

Strategic decisions therefore require information beyond patent data. A general-purpose model may be useful for flexible interpretation or generation; structured prompts and knowledge containers for recurring tasks requiring domain context; embeddings for semantic exploration; specialized classifiers for stable high-volume classification; deterministic software for rule-based checks; and conventional databases where structured numerical information dominates.

Steel-alloy research provides a useful counterexample to the assumption that semantic methods are always the better choice. The decisive information may lie in numerical composition ranges. Two documents can be linguistically similar while defining materially different compositions, or linguistically different while containing important numerical overlap. In our earlier experiments, semantic approaches produced too many false results for this type of task, while structured numerical databases remained more useful.

Dr. Stephan Wolke, *CEO of thyssenkrupp Intellectual Property GmbH*
Stephan Ising, *Head of IP Strategy / Research of thyssenkrupp Intellectual Property GmbH*

Not every patent problem is a language problem, and not every language problem is best solved by a generative model.

Validation should likewise fit the method. LLM workflows can be tested against known cases, gold-standard examples, and underlying evidence. Specialized classifiers with labelled test data can be assessed using confusion matrices and similar methods. Which error matters most depends on the process; in some monitoring applications, missing a relevant patent may be more problematic than sending an additional irrelevant document to an expert.

A model does not need to be perfect to be useful, but its failures need to be understood well enough for the surrounding process to manage them.

Reviewability is therefore more valuable than apparent reasoning. General-purpose models can generate convincing explanations, but those explanations are not proof. We ask for concise technical justifications because they help a reviewer understand and challenge the output; wherever possible, the result should point back to a verifiable passage, claim feature or cited document.

Errors should also improve the surrounding method. A recurring problem may indicate that a category is poorly defined, a knowledge container is incomplete, or another technical method is more appropriate. In many practical applications, the underlying general-purpose model does not learn from an individual correction. What improves is the **prompt**, **knowledge container**, **examples**, **workflow**, **data**, **or specialized model built around it**.

The model often does not learn from our experience; the surrounding method does.

Building capability therefore includes knowing when not to use AI. Some tasks are too infrequent to justify automation, some are better solved deterministically, some require specialized databases and some matters may be too sensitive for a particular AI environment.

Several lessons have remained stable throughout our work. We start with an actual IP problem rather than the newest model and experiment before investing in automation. Once a method proves useful, we try to make it transferable through prompts, knowledge containers, validation cases, or workflow definitions.

At the same time, responsibility and confidentiality boundaries remain unchanged. Adoption should follow practical usefulness, while evaluation should focus primarily on observable process effects such as workload, throughput, consistency, and the quality of decision input.

Most importantly, AI should not merely accelerate an existing process. It can change the sensible allocation of work. Our monitoring example is instructive: the significant improvement was not simply faster classification, but shifting expert time away from pre-filtering towards substantive assessment.

The goal is not to automate the old process as completely as possible. It is to design a better process around the combined strengths of AI and human expertise.

***Dr. Stephan Wolke**, CEO of thyssenkrupp Intellectual Property GmbH*
***Stephan Ising**, Head of IP Strategy / Research of thyssenkrupp Intellectual Property GmbH*

Conclusion

General-purpose AI has made sophisticated analytical and language-processing capabilities accessible to IP departments without requiring them to develop their own large language models. The opportunities extend across almost the entire IP value stream, but our experience suggests that the decisive question is not how much AI an IP department uses. It is how well the underlying work is understood and structured.

The most useful applications emerged when professional activities were decomposed into meaningful tasks and the appropriate method was selected for each. Extraction, comparison, classification, generation, validation and discovery require different information, technologies, and levels of professional review.

The development we have experienced is therefore not simply a progression from better prompts to better models and more automation. It is an increasingly explicit division of work between IP processes, AI methods, and **professional judgement**.

Models and tools will change. The more durable assets are an understanding of the process, explicit professional methods, domain knowledge, validated examples, and the ability to redesign workflows where technology enables a better allocation of work.

For us, the transition from **AI as an experiment to AI as a capability** has therefore been less about finding the perfect AI system than about understanding our own work well enough to combine AI, conventional tools and professional expertise appropriately.



***Dr. Stephan Wolke**, CEO of thyssenkrupp Intellectual Property GmbH*
***Stephan Ising**, Head of IP Strategy / Research of thyssenkrupp Intellectual Property GmbH*



Sakae Nakamura

Executive Advisor to LexisNexis® Intellectual Property Solutions / Senior Advisor, IP Landscape Association, Japan

Democratizing IP intelligence: How AI broadens participation and strengthens decision-making

By Sakae Nakamura, Executive Advisor to LexisNexis® Intellectual Property Solutions; Senior Advisor, IP Landscape Association, Japan

For many years, the ability to work effectively with patent information has largely remained the domain of specialists. Patent attorneys, search professionals, and IP intelligence teams developed the expertise required to formulate complex searches, interpret classifications, understand patent families, and analyze competitive portfolios. Business leaders, R&D teams, and corporate strategists typically relied on these specialists when they wanted to understand what patent information could tell them.

Artificial intelligence is beginning to change this relationship. Its significance is not simply that it makes existing activities faster. AI lowers the barrier to engaging with patent information in the first place. Search terminology can increasingly be generated through natural-language interaction. Large collections of patents can be summarized rapidly. Technologies, competitors, and emerging areas can be explored without every user first mastering the mechanics of traditional patent databases.

I believe this creates the potential for the **democratization of IP intelligence**. Here, it is important to be precise about what democratization means. It does not mean that every engineer, business manager, or strategist becomes a patent analyst. Nor does it mean that specialist expertise becomes unnecessary. Democratization changes who can participate meaningfully in the **early stages** of IP intelligence. More people can formulate questions, explore technologies, and test initial hypotheses for themselves.

It does not change who is responsible for the robustness of the analysis. The IP professional remains central to validating the analytical foundation, identifying omissions, and interpreting what the evidence really means. If anything, broader access makes this responsibility more important.

In practice, the bottleneck begins to move from **access to information** toward **confidence in information**. This distinction is fundamental to understanding the role AI can play in IP landscape analysis. IP landscape analysis itself should not be understood simply as patent analysis. Its purpose is to integrate patent and technology information with market developments, competitive activity, societal change, regulation, and policy in order to support strategic decisions. In other words, the objective is not simply to understand a patent dataset, but to construct a strategic view of where a technology, market, or competitive environment may be heading.

The changes brought by AI therefore affect more than searching. They affect the entire process through which an organization moves from a question to a decision. Based on my experience, I would like to propose a newly structured five-step framework for thinking about this process.

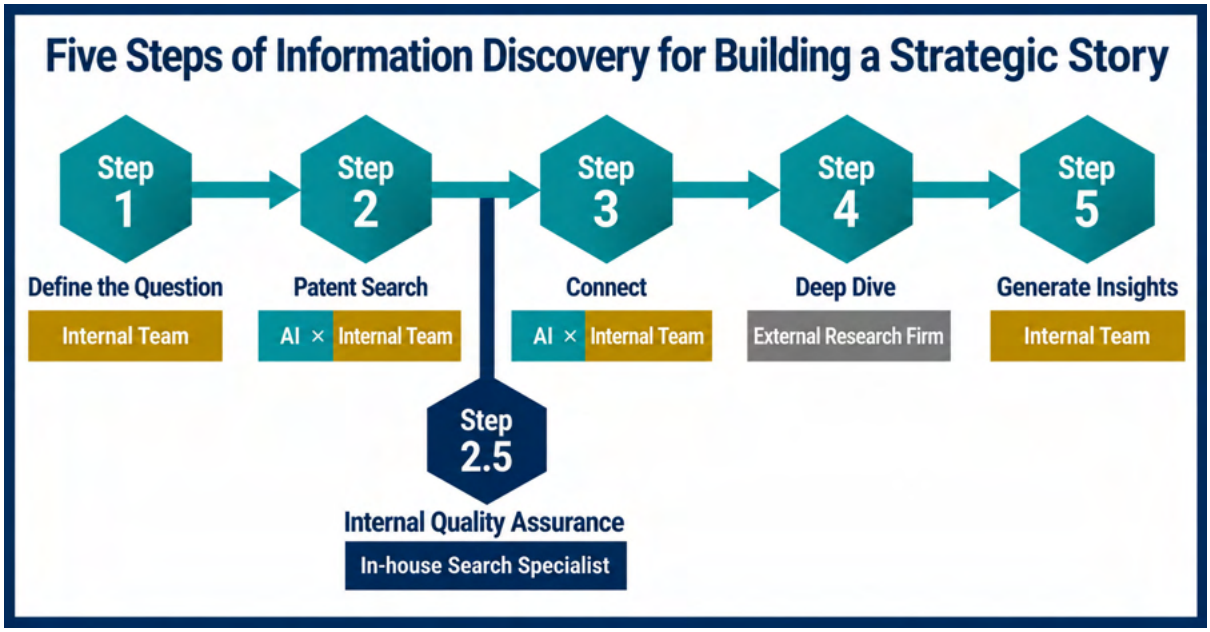


Figure 1: The five-step IP intelligence process, from strategic questions to actionable insights.

Step 1: Define the strategic question

Every strong analysis starts with the question. This sounds obvious, but it is also where many analytical projects begin to go wrong. The problem is often not that the search strategy is poor. It is that the question being investigated is too narrow, too technical, or insufficiently connected to a decision the organization actually needs to make.

There is an important difference between asking:

> *"What patents exist in this technology?" and asking "What do we need to understand about this technology before deciding whether to invest?"*

The first question produces information. The second creates the basis for a decision. This is why the starting point should be the business problem rather than the search query. What are we trying to achieve? Which customer or market need are we trying to address? Where are we looking for growth? Which capability may become strategically important? What would we need to know before management could make a decision?

The practical method is therefore to translate an initial business concern into hypotheses that can be investigated. The objective is not simply to determine "what to search," but to establish **what must be clarified in order to make a business decision**. This becomes even more important in the age of generative AI. When retrieving and processing information becomes faster, the ability to formulate the right question becomes more valuable.

AI can suggest terminology. It can broaden a technical concept. It can help structure hypotheses. But it cannot determine which strategic problem matters most to a particular company. That requires understanding the company's strategy, technologies, customers, competitive position, and ambitions. The quality of the question therefore becomes one of the most important determinants of the quality of the analysis.

Step 2: Explore broadly with AI

Once the strategic question has been defined, the process moves into exploration. This is where AI can substantially change the speed and breadth of IP intelligence work.

In traditional patent searching, considerable effort may be required simply to become familiar with the terminology of an unfamiliar technology. Different companies may describe similar concepts in different ways. Technical language changes over time. New terminology emerges. Relevant inventions may span several classifications.

AI allows the exploratory stage to become much more iterative. A user can begin with a concept, explore synonyms, identify adjacent terminology, test different technical interpretations, and progressively refine the field under investigation. AI can also help summarize large numbers of patents, identify possible technology structures, and surface competitors or emerging players that deserve closer examination.

The important point is that AI is not necessarily producing the final answer. It is creating an **exploration map**.

AI is used to move from unfamiliar terminology toward competitive structures and technology trends. It is used to draw a map of where to look, rather than delivering the definitive strategic conclusion. This distinction matters. AI makes it possible to explore more possibilities, more quickly. It allows engineers, business teams, and strategists to participate earlier in the process rather than waiting until an IP specialist has completed a full analysis.

That is one of the most tangible forms of democratization. But exploration must remain exploration. The ability to produce a result quickly should not be confused with the ability to demonstrate that the result is analytically sound. That leads to an additional stage that I regard as essential.

Step 2.5: Quality assurance — the bridge from exploration to intelligence

Between exploration and interpretation sits a step that is easily underestimated: **quality assurance**.

AI can be extremely effective at breadth. But breadth does not automatically produce precision. An AI-assisted search may capture terminology that is too broad.

Sakae Nakamura
Executive Advisor to LexisNexis® Intellectual Property Solutions / Senior Advisor, IP Landscape Association, Japan

Adjacent technologies may enter the result set. Important classifications may be missing. Apparently relevant documents may prove to address a different technical context.

Therefore, a specialist validation stage needs to be placed directly after AI-assisted exploration. Search professionals review the search strategy and resulting patent population, correct classifications, remove noise, identify gaps, and refine the dataset until it represents the technology being investigated with sufficient confidence.

I sometimes describe the relationship in simple terms:

> *AI provides the wide-angle lens. The specialist provides the microscope.*

AI expands the field of view. The expert tests whether what is being seen is relevant. This is where the changing role of IP professionals becomes particularly visible.

If business and R&D users can increasingly conduct initial exploration themselves, the IP specialist no longer needs to act as the gatekeeper through which every question must pass. Instead, specialist expertise becomes concentrated where it has the greatest value: assuring quality, recognizing blind spots, challenging assumptions, and determining whether an analytical result is robust enough to support further interpretation.

Democratization therefore does not remove expertise. It changes where expertise enters the process. The bottleneck shifts from **"Can I access this information?"** to **"Can I trust this analysis?"** That is a much more consequential responsibility.

Step 3: Connect patent information with its wider context

Validated patent information is still not strategy. A patent tells us something about a technical development. A portfolio may reveal where an organization has invested inventive activity. A trend may indicate that interest in a technology is increasing.

But none of this tells us automatically why it matters to a particular business. The next stage is therefore to create connections. Patent information needs to be considered alongside market developments, regulation, policy, demographic trends, customer needs, investment behavior, and the company's own capabilities. This is the point at which individual pieces of technical information begin to form a strategic narrative, moving from the **"point" of patent information to the "lines" of market, policy and societal context**. AI can help identify possible connections, while people determine which connections are meaningful for the company.

That division of roles is important. AI may indicate that a particular technology coincides with demographic change, new regulatory requirements, or increasing investment. But significance depends on context.

- Does the company already have relevant capabilities?
- Does the technology fit its channels and customers?
- Would investment reinforce an existing competitive advantage?
- Does the regulatory environment create an opportunity or a barrier?

A useful strategic connection is therefore not merely a correlation between datasets. It is a connection that changes how the company understands its options. This is also why IP intelligence needs to become increasingly cross-functional. The IP team brings patent knowledge. R&D brings technical understanding. Business teams bring market and customer context. Strategy brings corporate priorities. But no single function alone possesses the complete picture. The value emerges through their interaction.

Step 4: Investigate what is still uncertain

The next stage is deliberate investigation. By this point, the analysis should have generated hypotheses. Some will appear attractive. Others may challenge existing assumptions. But before a hypothesis becomes the basis for investment or strategic action, the organization needs to determine what is still missing.

Patent data and AI cannot answer every question. A regulatory issue may require specialist legal interpretation. A market assumption may require direct validation. A technically attractive opportunity may be economically unrealistic. A competitor's apparent patent position may not correspond to its commercial capabilities. This stage is therefore about challenging the emerging story.

- What evidence could disprove our hypothesis?
- What assumptions remain uncertain?
- Which information cannot be obtained from patent data?
- What additional expertise is required?

In the five-step methodology, deeper external investigation is used precisely where AI and internal analysis cannot provide sufficient confidence, for example around regulation, competitor behavior, technical realities, or international experience.

This step prevents IP intelligence from becoming an exercise in confirming what the organization already wants to believe. Good intelligence should not merely make a strategy look convincing. It should make the strategy more robust by exposing weaknesses before resources are committed.

Step 5: Turn intelligence into a decision

The final step is not another analysis. It is a decision. At this stage, the role of AI changes again. AI can help illuminate possible directions. IP specialists can validate the analytical evidence. R&D, business teams, and external experts may contribute additional knowledge.

But strategic decision-making and accountability must ultimately sit with the leadership responsible for the outcome. This deserves emphasis because the distinction is sometimes blurred in discussions about increasingly capable AI systems. The purpose of intelligence is to support leadership, not replace it. Specialists ensure analytical integrity. External advisors may contribute expertise. AI can broaden the available options. But choices about what the company will pursue, where it will invest, which risks it is willing to accept and which opportunities it will decline belong to those responsible for the business consequences.

The methodology structures this final output around three components:

- Direction—where should we go?
- Options—how could we proceed?
- Risks—what could prevent success?

Sakae Nakamura
Executive Advisor to LexisNexis® Intellectual Property Solutions /
Senior Advisor, IP Landscape Association, Japan

Only when these elements are brought together does information become a strategic story that management can translate into concrete action. This is the real endpoint of IP intelligence. Not the report. Not the landscape. Not the visualization. The decision.

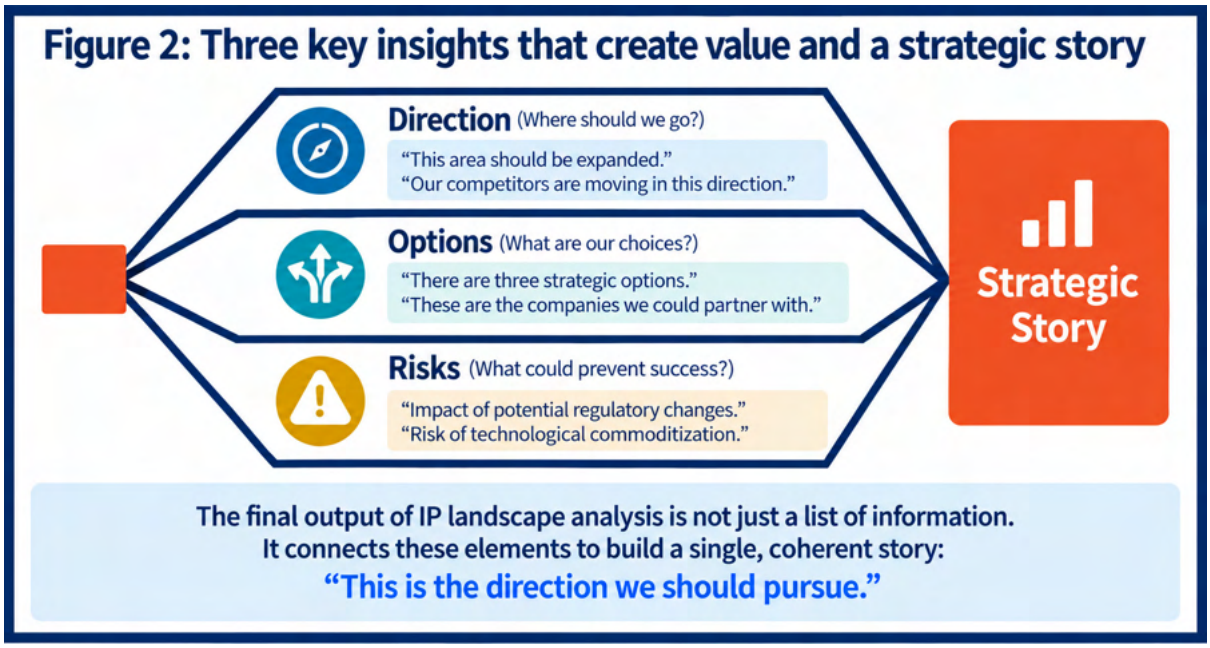


Figure 2: Three key insights that create value and a strategic story: direction, options and risks.

From an IP service function to an intelligence hub

This process also has important implications for the way IP organizations operate. In many companies, IP intelligence still resembles a service model: A business unit submits a request. The IP team conducts research. Several days or weeks later, a report is returned.

The underlying organizational structure remains largely siloed: management makes decisions, business develops strategy, R&D develops technology, and IP provides specialist analysis.

The presentation describes the limitations of this model clearly. Business may not understand competitor technology; R&D may operate without sufficient market context; IP may investigate a question without fully understanding management's intent; and management may make investment decisions without a sufficiently detailed view of the technology landscape.

As AI makes patent information easier to use across functions, a different model becomes possible. IP intelligence can become a shared information layer across R&D, business development, strategy, and management.

The IP function then changes from **holding the information** to the function that **enables the organization to use it well**.

It becomes an intelligence hub. Its role is to help colleagues formulate better questions, provide access to appropriate data and tools, validate analytical methods, challenge weak conclusions, and translate patent information into business implications. This is a much more strategic role than operating as a research service desk. It also means that IP professionals become involved earlier. Instead of receiving the question after the business has already determined what it wants to investigate, they can participate in framing the problem itself. Instead of delivering a report at the end of the process, they can participate continuously as hypotheses are developed, tested, and revised.

Ideally, this turns into a movement from a silo structure toward a hub structure in which patent information becomes a common language connecting different functions. This may ultimately be one of the most important consequences of AI for IP intelligence.

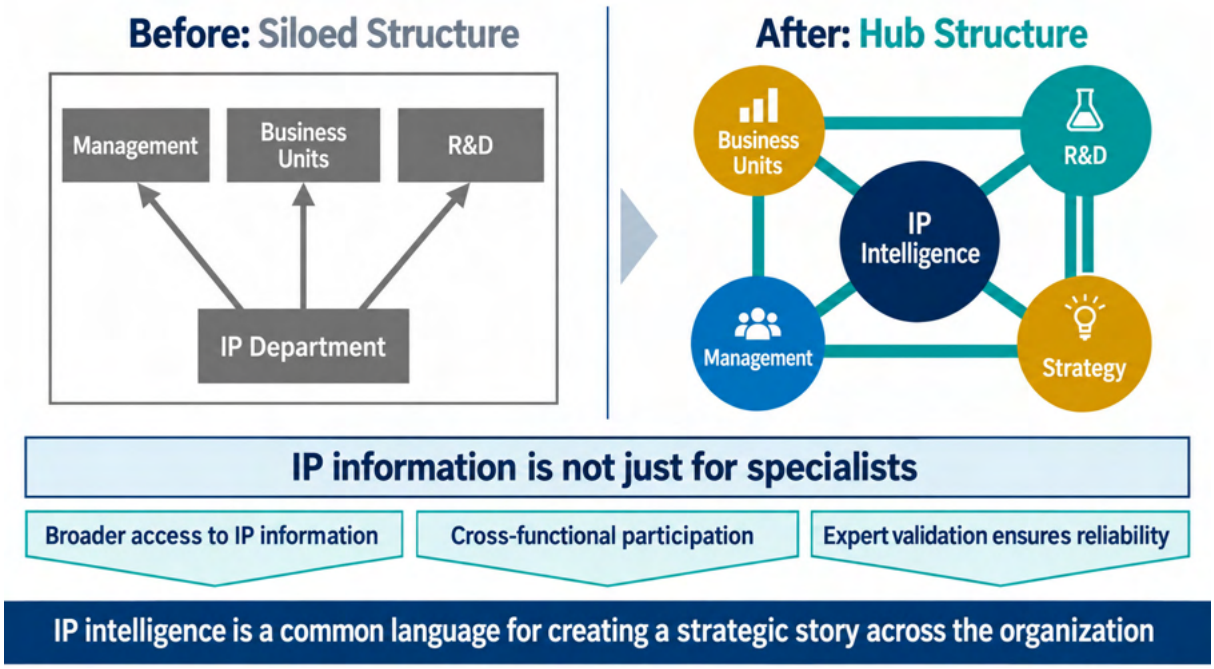


Figure 3: From a siloed IP function to an intelligence hub, enabling broader participation and stronger decision-making.

Democratization is about participation, not the removal of expertise

AI makes patent information easier for more people to use than ever before. But easier use is not the objective in itself. The objective is to improve the organization's ability to make decisions. Democratization therefore should not be interpreted as giving everyone equal responsibility for patent analysis. It means enabling people beyond the IP function—including R&D, business, corporate strategy, and management—to participate meaningfully in asking questions, exploring possibilities and contributing their own expertise.

At the same time, specialist IP expertise becomes more—not less—important at the points where analytical robustness matters. The result is a different division of work:

- AI expands exploration
- IP professionals assure analytical integrity
- Business and R&D contribute context
- Leadership owns the strategic decision

Sakae Nakamura
Executive Advisor to LexisNexis® Intellectual Property Solutions / Senior Advisor, IP Landscape Association, Japan

This is not simply a more efficient version of the traditional IP landscape process. It is a different operating model. And as AI makes answers increasingly easy to generate, another scarce resource becomes more valuable: the ability to formulate the questions that are worth answering.

This leads to a final point:

> *When AI can return information rapidly, the quality of research and IP work increasingly depends on the quality of the questions being asked.*

That, in my view, is where the real opportunity lies. The future of AI-enabled IP intelligence is not a world in which everyone becomes a patent expert. It is one in which people beyond the IP function—including R&D, business, corporate strategy, and management—can participate more directly in understanding technology and the competitive environment, while IP professionals ensure that the resulting intelligence remains reliable, relevant, and connected to decision-making.

The goal, therefore, is not simply to make patent information easier to use. It is to **enable IP, R&D, business, corporate strategy, and management to bring their distinct perspectives into the process—asking better questions, exploring possibilities, and combining their expertise to support better-informed decisions**. That is what the democratization of IP intelligence is intended to achieve.



Sakae Nakamura
Executive Advisor to LexisNexis® Intellectual Property Solutions /
Senior Advisor, IP Landscape Association, Japan



Tim Alexander Oelmann

VP, Head of Competence Center Business IPR Management at Deutsche Telekom AG

Daniel Uhl

Patent Information Specialist at Deutsche Telekom AG

From inventor culture to invention impact: How AI is changing the challenge for IP management

By Tim Alexander Oelmann, VP, Head of Competence Center Business IPR Management at Deutsche Telekom AG, and Daniel Uhl, Patent Information Specialist at Deutsche Telekom AG

Invention starts with people

At Deutsche Telekom, invention and R&D are often one aspect of broader roles and responsibilities. Employees contribute to technological development while balancing other tasks in their day-to-day work. Developing an idea further and preparing an invention disclosure therefore requires dedicated time and attention.

At the same time, invention itself is becoming more demanding. Increasing technological complexity makes it more challenging to understand the relevant state of the art, develop new technical solutions, and describe increasingly complex inventions. Development cycles are getting shorter and competition more dynamic. In simple terms, inventors face more effort in less time.

Our experience shows that having ideas is only part of the equation. Inventors also need the time, space, and mindset to develop those ideas and make them visible. The challenge is therefore not simply to generate ideas, but to create the conditions in which people are willing and able to take them further.

For us, the first lever is therefore not technology. It is creating an environment that gives people the motivation, space, and support that invention requires.

Inventor Culture activates

This is where our Inventor Culture comes in. Its purpose is to put inventors at the center and create an environment that strengthens their intrinsic motivation.

Inventors want to solve problems, create something new, and see their ideas make a difference. Our role is to support this through knowledge, recognition, visibility, connection, and dedicated

opportunities to invent. The aim is to move invention away from being perceived primarily as an administrative task and towards a contribution people can identify with and be proud of.

This approach is grounded in research on motivation, creativity, and employee engagement. Scientific findings underline the importance of autonomy, competence, and belonging for intrinsic motivation, while research on team creativity highlights the value of exchanging and integrating different perspectives. Recognition, positive learning experiences, and identifying with the role of an inventor can further strengthen this motivation. These insights inform how we develop our Inventor Culture.

Different formats address these needs. The **Patent Garage** provides knowledge and practical support around patents and invention disclosures. The **Inventor Hall of Fame** creates visibility and recognition. **Ideation Labs** offer dedicated time and space to explore ideas. The **Inventathon** brings people from different backgrounds together, while the **Inventor Community** connects people and expertise beyond individual events.

Co-creation is particularly important. At the Inventathon, deliberately diverse teams combine technical expertise and practical experience that might otherwise never meet. This exchange can lead to solutions less likely to emerge within homogeneous teams or established organizational boundaries.

The environment matters too. Taking people out of their normal routines and creating space for exchange can make invention feel less like another corporate task. Identification, community, and intrinsic motivation can create a dynamic that financial incentives alone cannot sustain. Inventor Culture is therefore not an add-on to the patent process. It is how we activate the people behind it.

IPR Intelligence gives direction

Motivation alone is not enough. What inventors work on matters just as much. For inventors, knowing that they are working on the "right" topics matters too: where could their ideas be relevant, and where can they contribute to Deutsche Telekom's technological priorities?

This is where IPR Intelligence, Data-to-Decision, and strategic direction become important. Patent data and analytics provide insights into technological developments, competitors, and the state of the art. They provide a stronger basis for evaluating technological whitespaces and invention potential, while our strategic perspective helps focus on topics relevant to Deutsche Telekom.

Data does not make that evaluation on its own. Determining whether a technological field offers relevant invention potential requires technology know-how and patent know-how to come together.

Inventors contribute their understanding of technologies, practical challenges, and potential solutions. Business IPR Management adds the patent perspective and strategic context. Together, we can evaluate relevant fields and provide a better-informed starting point for co-creation. In formats, such as the Inventathon, this helps bring inventors together around topics that are both technologically promising and relevant to Deutsche Telekom.

We therefore do not only motivate and connect people. We also help focus their creativity on relevant challenges.

AI comes in as an additional booster

Alongside this development, AI tools are becoming part of everyday work. They are not currently a centrally designed element of our Inventor Culture; inventors are adopting them themselves and using them where they find them useful.

***Tim Alexander Oelmann**, VP, Head of Competence Center Business IPR Management at Deutsche Telekom AG*
***Daniel Uhl**, Patent Information Specialist at Deutsche Telekom AG*

This includes research and information processing, structuring technical content, refining formulations, or preparing an invention disclosure. To provide a clear framework, our IP Legal team has developed guidance on how AI can be used in the context of inventions, including what is permitted and where limitations apply.

AI can reduce some of the effort surrounding invention and help an already activated inventor base work more efficiently. At the same time, it creates opportunities to automate and streamline parts of the IP process and support a growing Inventor Community more effectively.

This is where the combination becomes interesting. Inventor Culture creates motivation, connection, and space to invent. AI provides tools that can make parts of the work easier. In that sense, one plus one can become three.

The reverse is just as important. Without motivated people who recognize problems, bring technical expertise, and want to develop new solutions, there is little for AI to support. Zero plus one remains zero.

The invention itself still depends on people. AI supports the work around it.

The effect becomes visible

Over the past three years, we have systematically developed our Inventor Culture, while inventors have increasingly incorporated AI tools into their daily work. AI is also creating opportunities across the broader IP organization, including IPR Management, IP Legal, and Standardization, to improve analysis, processes, and inventor support.

> *Together, these developments have contributed to a significant increase in invention activity. In the first half of 2026, invention disclosures more than doubled compared with the same period in the previous year.*

We also see a qualitative development in our day-to-day assessment: **submissions are increasingly detailed and clearly structured**, while inventors **respond faster** and provide **richer technical information** when clarification is needed.



***Tim Alexander Oelmann**, VP, Head of Competence Center Business IPR Management at Deutsche Telekom AG*
***Daniel Uhl**, Patent Information Specialist at Deutsche Telekom AG*

This does not mean that AI creates better inventions. Their technical substance, innovative character, and relevance still depend on the people behind them. AI can, however, make the environment around invention more effective, both for inventors and for the functions supporting them.

This can create a reinforcing dynamic: a strong Inventor Culture encourages invention, AI can make it easier to develop and communicate ideas, and greater efficiency across the broader IP organization can strengthen support for a growing inventor base.

The result is more inventive activity and better information entering the patent process. But as this success scales, it raises the next challenge.

What happens when success scales?

More invention disclosures also mean more work for the IP organization. Each disclosure has to be assessed, promising inventions require technical and legal evaluation, filing decisions need to be made, applications prepared, and resulting assets managed within the portfolio.

> *The productivity gains inventors experience through AI do not currently reduce this workload at the same rate.*

If Inventor Culture and new tools succeed in making more inventions visible, IP management has to absorb the additional output. This can mean more patent work, more portfolio decisions, and potentially higher costs. Increasing budgets and scaling existing processes may be part of the response, but additional measures will be needed to manage this growth effectively.

From more inventions to a stronger patent portfolio

The focus now shifts from generating more inventions to making the right decisions about what follows from them.

At Deutsche Telekom, we deliberately connect innovation management and IPR management. Our objective is not to generate ideas for their own sake, but to develop promising inventions into patents and ultimately strategic assets.

As invention activity grows, quality and prioritization become increasingly important. The relevant question is which inventions can contribute most effectively to the strategic objectives of the portfolio.

For Deutsche Telekom, those objectives include protecting the business, supporting innovation leadership, creating monetization opportunities, and strengthening deterrence.

A larger pipeline can therefore become an opportunity to build a stronger portfolio, provided the IP organization can evaluate, prioritize, and develop it effectively.

The next challenge for IP management

Inventor Culture sets the dynamic in motion. IPR Intelligence helps give creativity direction. AI adds tools that can make parts of inventors' work faster and easier. If that combination continues to increase inventive output, IP organizations need to scale not only invention, but also their ability to evaluate, prioritize, and manage what it produces.

This brings us to the central question:

> *How do we turn significantly more inventive output into a stronger, higher-quality patent portfolio?*

***Tim Alexander Oelmann**, VP, Head of Competence Center Business IPR Management at Deutsche Telekom AG*
***Daniel Uhl**, Patent Information Specialist at Deutsche Telekom AG*

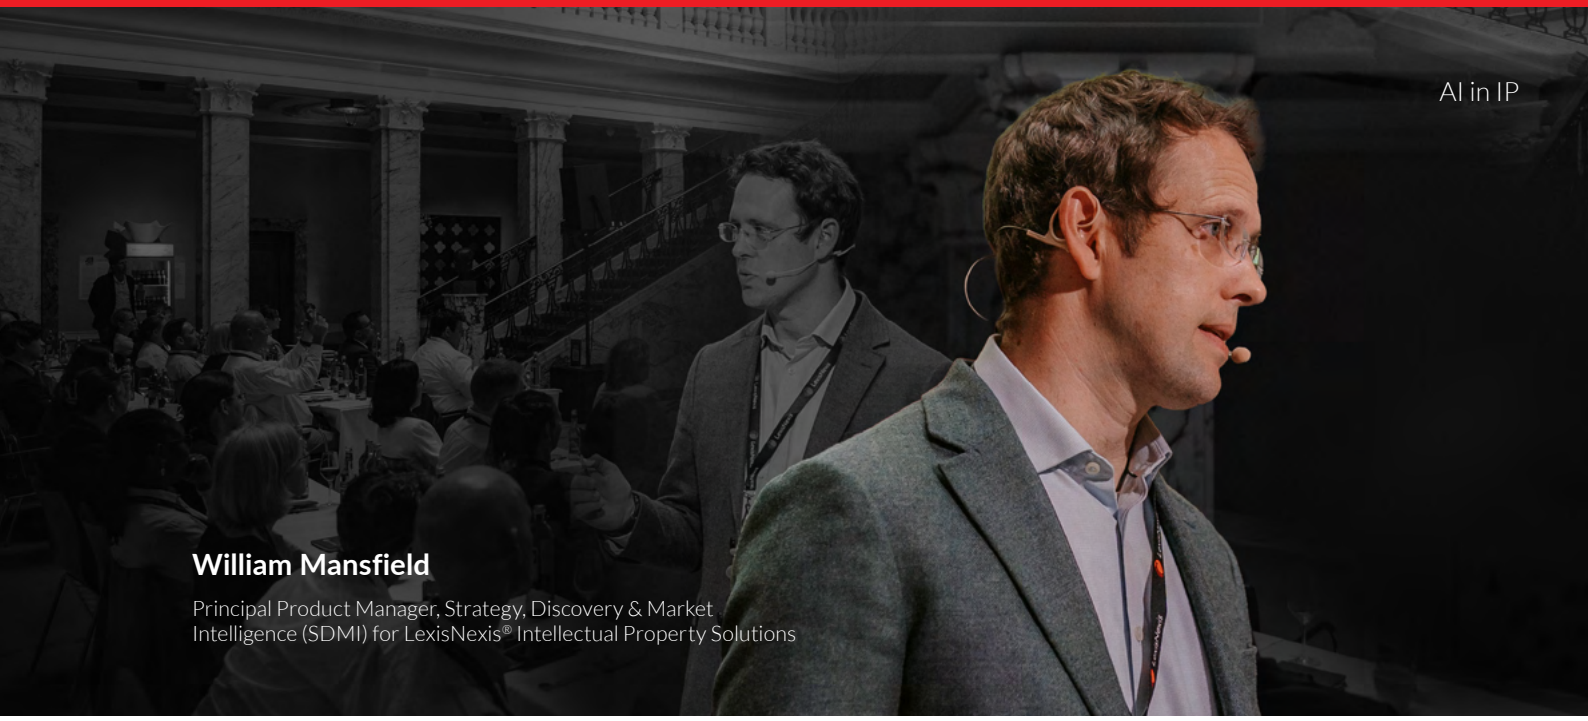
There is no single answer. Do portfolio and prioritization processes need to be adapted? Do budgets need to grow? Do we need to become more selective? Each option has implications for how IP organizations operate and where they focus their resources.

AI may eventually become part of the answer on this side of the process as well. Can AI also help IP management evaluate and steer the additional output more intelligently in the future? And looking beyond individual AI applications, an even more fundamental question remains: Could targeted and consistent process automation be the holy grail?

These are the questions that follow when increasing invention activity is no longer the only challenge. The next task is to determine how IP management itself needs to evolve to turn that potential into a stronger patent portfolio.



***Tim Alexander Oelmann**, VP, Head of Competence Center Business IPR Management at Deutsche Telekom AG*
***Daniel Uhl**, Patent Information Specialist at Deutsche Telekom AG*



William Mansfield

Principal Product Manager, Strategy, Discovery & Market Intelligence (SDMI) for LexisNexis® Intellectual Property Solutions

From embedded AI to agent-ready patent intelligence

By William Mansfield, Principal Product Manager, Strategy, Discovery & Market Intelligence (SDMI) for LexisNexis® Intellectual Property Solutions

Discussions about AI in intellectual property often focus on individual features: an AI-assisted search, an automated summary, a drafting copilot or a chatbot for patent data. Yet the more consequential question is not simply **whether a tool uses AI**, but where **AI sits within the technology stack—and what surrounds it**.

Organizations are beginning to implement these layers through three principal models:

1. AI-supported or AI-enhanced software-as-a-service tools.
2. Direct access to patent datasets combined with an enterprise AI layer.
3. MCP-based access to patent data and analytical capabilities.

These models should not be treated as a simple maturity ladder. An MCP-based architecture is not automatically more advanced than a purpose-built SaaS application, just as direct access to data is not inherently superior to a finished analytical product. Each model solves a different problem and assigns responsibility for data quality, analytical logic, user experience, and governance differently.

AI-supported and AI-enhanced SaaS tools

In the first model, AI is embedded within a purpose-built IP or patent application. The software provider typically supplies most of the stack: the patent data, data harmonization, search and analytical functionality, domain-specific metrics, user interface, AI models, and security framework.

The user does not need to assemble these elements independently. Instead, AI augments an established workflow—for example by allowing users to formulate questions in natural language, helping structure an analysis, identifying relevant portfolios or technologies, summarizing results, or producing visual outputs.

LexisNexis® PatentSight+™ with Protégé™ is one example of this approach. Protégé allows users to begin with a plain-language business question and then guides them towards structured patent analyses grounded in PatentSight+ data, portfolio metrics, and business context. The application is designed to show how an analysis has been constructed and to produce outputs that can be reviewed and communicated to non-specialist stakeholders.

The principal advantage: a governed analytical environment

The primary strength of the SaaS model is not simply convenience. It is the combination of AI with a defined analytical environment.

Many apparently simple patent questions depend on methodological choices that must be settled before any analysis can be trusted: how a patent family is delimited, which legal entities constitute a corporate group, whether a family counts as pending, in force or inactive, and at which point in time a portfolio is being observed. Section 6.5 examines these choices in detail.

The relevant point for this model is where the choices are made. A purpose-built platform can embed them in the data and in the workflow, so that they are applied consistently and identically for every user, rather than asking the language model or the individual analyst to reconstruct them each time a question is asked. The methodology becomes a property of the product rather than of the prompt.

This has a practical consequence for AI interaction. Because the definitions are fixed and documented, the model's task is reduced to interpreting intent, selecting an appropriate analysis, and explaining the result. It is not required to decide what an owner is.

Where this model is strongest

AI-enhanced SaaS is particularly well suited to:

- Recurring competitive-intelligence and benchmarking analyses
- Portfolio reviews and technology landscapes
- Trend identification and strategic reporting
- Analyses that need to be shared with R&D, corporate strategy, or senior management
- IP teams that require rapid time to insight without constructing their own data and AI infrastructure

Because the application defines the data model and analytical ways, results are generally easier to reproduce and compare across users. The model can also broaden access to patent intelligence: business stakeholders may be able to start with a strategic question without first learning a patent-query language or the details of the underlying database schema.

The trade-off

The same structure that creates consistency also places boundaries around customization.

A SaaS application can only support the data, metrics, workflows, and integrations available within its product environment. Organizations may find it more difficult to incorporate highly proprietary taxonomies, internal product data, confidential R&D information, financial models, or organization-specific scoring methods.

William Mansfield
Principal Product Manager, Strategy, Discovery & Market Intelligence (SDMI)
for LexisNexis® Intellectual Property Solutions

Dataset access combined with an enterprise AI layer

The second model shifts the point of control. Instead of accessing patent intelligence primarily through a finished application, the organization receives structured patent data within its own data and analytics environment.

Modern enterprise data platforms make this practical. Open sharing protocols can deliver live, governed datasets to a customer's environment without repeated copying or bespoke extract pipelines, and federation capabilities allow queries to reach data that has not been migrated at all. Commercial data catalogues and marketplaces have made third-party datasets a routine component of the enterprise data estate rather than an exception.

A harmonized patent dataset delivered in this way typically includes patent-family records, corporate ownership structures, legal-status information, technology classifications, and patent-level or portfolio-level indicators. The organization can then combine these elements with its own data and analytical models. The PatentSight+ Dataset, available through cloud data platforms and data-sharing mechanisms e.g. Databricks, is one example of a patent dataset offered in this form.

The customer then creates the AI layer. This might include:

- A natural-language interface that translates questions into SQL
- A patent-analytics copilot for analysts and business users
- Retrieval-augmented generation over patent texts and internal documents
- Agents that execute approved analytical functions
- Custom classifiers for technologies, products or strategic themes
- Automated monitoring, reporting and alerting workflows
- Models that combine patent indicators with financial, product, market, or R&D data

The enterprise platform supplies much of the surrounding machinery. Governance catalogues now extend beyond tables to models, notebooks, dashboards, and agents, providing centralized access control, lineage tracking, and audit records of who ran which analysis, against which data, and when. Natural-language query interfaces can inherit those permissions automatically. Agent frameworks increasingly include evaluation tooling, automated judges, quality metrics, and human-feedback loops, so that agent behavior can be measured rather than assumed. Low-latency serving layers and application hosting allow the resulting agents to be deployed close to the data instead of shipping data to an external application.

These capabilities provide a genuine enterprise foundation. What they do not provide is patent-domain meaning.

William Mansfield
Principal Product Manager, Strategy, Discovery & Market Intelligence (SDMI)
for LexisNexis® Intellectual Property Solutions

The principal advantage: analytical differentiation

Dataset access becomes valuable when the organization needs to answer questions that cannot be addressed adequately within a standardized application.

This is not because the necessary data is unavailable in tools such as PatentSight+. Customers value PatentSight+ precisely for its flexibility and the many ways in which metrics, filters, views, and visualizations can be tailored to their needs. Like any SaaS product, however, it is designed primarily around common, repeatable use cases and must make choices about which analyses to provide as standard functionality for a broad user base.

Some questions require a more tailored analytical approach: a different calculation, a non-standard segmentation of the data, a bespoke comparison over time, or a visualization designed for a particular management or strategic decision. These requirements may be highly valuable to one organization without being sufficiently general to become standard product functionality. Direct access to the underlying harmonized data allows analysts to work with the available data more flexibly and to develop those organization-specific analyses.

A company could, for example:

- Link patent strength to product revenue or R&D expenditure
- Compare patent positioning with internal technology road maps
- Identify acquisition targets using patent, financial, and market indicators
- Build a proprietary technology taxonomy reflecting its business structure
- Combine patent portfolios with supply-chain, standards, or litigation data
- Develop organization-specific portfolio optimization or renewal models
- Create longitudinal company-level datasets for econometric or investment analysis

The ability to join harmonized ultimate-owner identifiers with corporate and financial identifiers is particularly important for these analyses. Entity resolution is often the hidden bottleneck in patent-to-business analytics: even sophisticated models will produce unreliable conclusions when company names and corporate structures are not aligned across datasets.

Dataset access can also support a more explicit treatment of time. Where ownership history and portfolio snapshots are represented in the data model, analysts can reconstruct which company controlled a portfolio at a given point in time and build company-year panels rather than projecting today's structure backwards across a decade of filings.

The AI layer requires a semantic layer

Giving a language model access to structured patent tables is not equivalent to creating a patent-analytics solution.

Enterprise platforms increasingly claim a semantic understanding of the customer's own data estate, derived from metadata, column descriptions, business definitions, dashboards, query patterns, and user feedback. This is a real capability, and it addresses a real problem: platform documentation is fond of the example of a column named sales_NE, where only business context reveals whether the reference is to Nebraska, New England, or the Netherlands.

William Mansfield
Principal Product Manager, Strategy, Discovery & Market Intelligence (SDMI)
for LexisNexis® Intellectual Property Solutions

Patent data poses the same class of problem at a far higher level of domain specificity, and it is not solvable by observing internal usage. No amount of metadata inference will establish that several hundred distinct applicant names belong to one corporate group, or that a particular sequence of office actions means a family has lapsed in one jurisdiction while remaining in force in another. That knowledge has to be supplied by the dataset or modeled deliberately by the customer.

Concretely, the AI layer needs to understand:

- The difference between publications, applications, grants, and families
- The appropriate join paths between patents, owners, technologies, and legal events
- The definitions and limitations of each metric
- Whether a question refers to current ownership or historical ownership.
- Whether a trend is based on filing date, publication date, priority date, or reporting date
- Which data may be disclosed to which users
- Which calculations require deterministic functions rather than generative reasoning

For this reason, the most robust architecture does not allow the language model to invent analytical logic each time it receives a question. Instead, the model should operate through a governed semantic layer containing approved definitions, queries, metrics, and functions.

For numerical patent analytics, deterministic SQL or validated analytical services should perform the calculation.

The language model can then interpret the user's intent, select the appropriate function and explain the result.

> *The model should orchestrate and communicate the analysis—not improvise the underlying patent methodology.*

The trade-off

Dataset access provides much greater flexibility, but it transfers significant responsibility to the customer.

The organization must build and maintain:

- Data pipelines and update processes
- A patent-domain semantic model
- Access controls and data entitlements
- AI instructions and tool definitions
- Evaluation datasets and quality tests
- Monitoring, lineage, and audit processes.
- User interfaces and workflow integrations
- Processes for reviewing AI-generated conclusions and recommendations

It is worth noting how platform literature tends to divide this responsibility. Vendor material is typically explicit that high-quality, well-governed data is foundational to model quality, and equally explicit that governance, security, and compliance remain the customer's obligation. The domain semantics in between, what an owner is, when a right is enforceable, which portfolio existed in 2018, are generally assumed to arrive with the data.

The principal risk is therefore not only hallucination. It is the production of

William Mansfield
Principal Product Manager, Strategy, Discovery & Market Intelligence (SDMI) for LexisNexis® Intellectual Property Solutions

outputs that are technically valid but analytically inappropriate—for example, a syntactically correct query that uses the wrong date concept, ownership definition. or portfolio metric. Such results are difficult to detect precisely because nothing has failed.

Dataset access is consequently most appropriate where the organization has both a differentiated analytical need and the technical and domain resources required to govern it.

MCP-based access to patent data and analytics

The third model is based on the Model Context Protocol, or MCP.

MCP is an open protocol that provides a standardized way for AI applications to connect to external data sources and executable tools. Its architecture distinguishes between an AI host, clients managed by that host, and specialized servers that expose capabilities. MCP servers can make resources, tools, and reusable prompts available to an AI application.

The problem MCP addresses is integration cost. Platform vendors describe the situation it replaces in consistent terms: every system an agent needs to reach, a Customer Relationship Management, a messaging platform, a code repository, a document store, has historically required a custom connector for every agent framework. MCP offers one standard way to expose data and tools instead. This framing is worth keeping in mind, because it locates MCP precisely: it is a connectivity standard.

MCP is not itself a patent database, an AI model, or an analytical methodology. It is an interface through which an AI system can discover and use approved external capabilities.

MCP is an access and orchestration layer—not a substitute for patent-data quality or analytical expertise.

In a patent context, an MCP server could be implemented by a patent-data provider or by an enterprise using licensed patent data. It might expose three categories of capability.

Resources

Resources provide data or contextual information that the AI application can read, subject to the user's permissions. Patent-related resources could include:

- Patent-family records
- Corporate ownership hierarchies
- Technology taxonomies
- Saved patent landscapes
- Portfolio definitions
- Analytical methodology documentation
- Previously approved reports

Tools

Tools are executable functions that an AI model can call. A patent-data MCP server would not necessarily expose a separate tool for every analytical use case. Instead, it might offer a small set of flexible, composable functions, such as:

search()
patent_details()
aggregate()

search() could perform full-text or structured patent searches. patent_details() could retrieve the underlying patent records and source text needed to inspect or substantiate a result. aggregate() could apply a defined set of filters, groupings and measures to create portfolio, competitor, or technology analyses in different combinations.

William Mansfield
Principal Product Manager, Strategy, Discovery & Market Intelligence (SDMI)
for LexisNexis® Intellectual Property Solutions

This approach retains analytical flexibility without giving the model unrestricted access to the underlying database. A generic execute_sql() tool may allow an agent to construct inappropriate queries or apply definitions inconsistently. By contrast, an aggregate() tool can embed the relevant business logic while still allowing the user or agent to define the analytical question.

That business logic matters. For example, a patent may have multiple owners or multiple IPC classifications. When aggregating a portfolio, those relationships must be handled carefully to avoid double counting. Similarly, ownership definitions, patent-family rules, reporting dates, and portfolio metrics need to be applied consistently. Tools that return defined, traceable responses reduce the risk that an AI model reaches a plausible looking but methodologically incorrect result.

This places the analytical logic on the server side, where it can be tested, versioned, and audited. The names of the tools listed above are themselves instructive: each one encapsulates a domain decision that would otherwise have to be improvised.

Prompts and workflows

An MCP server can also make reusable prompts or workflow templates discoverable to an AI application. A patent-intelligence workflow might guide an agent through the steps required to:

- Define a technology field
- Validate the relevant competitors
- Establish a reporting period
- Compare portfolio size and quality
- Identify areas of strength and weakness
- Retrieve the underlying evidence
- Generate an executive summary with stated assumptions

The official MCP specification defines resources, tools, and prompts as distinct server capabilities, allowing AI hosts to combine specialized services while retaining control over the context supplied to each server.

The principal advantage: reusable patent capabilities

MCP can make the same governed patent-data capabilities available across multiple AI environments.

For example, a company might want patent intelligence to be accessible through:

- An IP department assistant
- An R&D copilot
- A corporate-strategy agent
- An M&A due-diligence workflow
- A Microsoft Teams or enterprise-search interface
- A custom application used by innovation managers

Without a common access layer, each application may require a separate integration, separate query logic, and separate security model. An MCP server can centralize the approved patent tools while allowing different AI hosts to invoke them.

In principle, an enterprise agent could also combine a patent-data MCP server with separate financial, market, product, or R&D services. The agent might call one tool to identify a company's patent strength, another to retrieve financial performance, and a third to compare the result with the internal technology road map. MCP's composable host-client-server architecture is intended to support this form of specialized integration.

The trade-off

MCP standardizes connectivity, but it does not determine:

- Whether the underlying patent data is accurate
- Whether ownership has been harmonized correctly
- Whether legal-status information has been interpreted appropriately
- Whether a metric is suitable for the business question
- Whether the AI should be permitted to execute a particular tool
- Whether the resulting recommendation is legally or strategically sound

MCP implementations therefore require careful authorization, access control, logging, and user-consent mechanisms. The protocol's security guidance emphasizes user control, data privacy, and caution around tools that can access data or execute actions. Its authorization framework provides mechanisms for restricted HTTP-based servers, but organizations still need to implement permissions and policies appropriate to their own data and risks.

For patent workflows, this should include at least:

- Role-based access to confidential portfolios and analyses
- Enforcement of contractual data entitlements
- Separation of public patent data from privileged internal information
- Tool-level permissions rather than unrestricted system access
- Logging of the data, tools, and parameters used in an analysis
- Human review before consequential legal or portfolio decisions
- Evidence links allowing users to inspect the underlying patent records

MCP is most valuable where an organization already has reliable patent-data or analytical services and wants to make those capabilities safely reusable across multiple agents and applications.

William Mansfield
Principal Product Manager, Strategy, Discovery & Market Intelligence (SDMI)
for LexisNexis® Intellectual Property Solutions

Comparing the three integration models

Decision dimension	AI-enhanced SaaS	Dataset access with an AI layer	MCP-based on patent data
Primary deliverable	Finished patent application and user experience	Governed patent-data foundation	Standardized interface to patent resources and tools
Where the domain logic resides	Primarily in the provider's application	In the customer's semantic and analytical layer	In the MCP server and the services behind its tools
Time to value	Generally fastest	Longer implementation cycle	Depends on whether governed data services already exist
Customization	Moderate and bounded by the application	Very high	High at the workflow level, but bounded by exposed tools
Combination with internal data	Possible through supported integrations or exports	A core strength	Possible through composition with other approved services
Technical responsibility	Primarily the provider, with customer usage governance	Primarily the customer	Shared across server provider, AI host and customer
Analytical consistency	High when users remain within supported workflows	Depends on semantic-layer and model design	Depends on the quality and specificity of server-side tools
Reusability across AI assistants	Usually limited to the application	Requires custom interfaces	A principal design objective
Best suited to	Repeatable patent analyses and rapid decision support	Proprietary, cross-domain, and highly customized analytics	Agentic workflows requiring governed access to reusable patent capabilities
Principal risk	Overreliance on a prescribed analytical framework	Technically correct but semantically inappropriate analysis	Over-permissioned agents or poorly designed tools
Patent-analytics quality depends on	The application's data, metrics, and workflow design	The underlying data, semanti model, and engineering disciplinec	The server-side methodology and the data services behind it

Table 1: Comparison of integration models. Source: Author's synthesis.

The final row of the table is the one that does not vary. Whichever model an organization selects, the quality of the result is determined by the patent-data layer beneath it. That layer is the subject of the next section.

Data quality as the prerequisite for the AI layer

The three models distribute responsibility differently, but none of them creates patent-domain meaning. A finished application can apply it consistently, an enterprise semantic layer can encode it, and an MCP tool can enforce it—but in each case the meaning must already exist in the data.

Four elements matter most.

Patent families: defining the unit of analysis

Family definitions often sound simple: a set of documents covering the same invention. In practice, they are not always straightforward. Different patent offices have different rules, which may require filings in their jurisdictions to be materially different. An applicant's understanding of an invention may also change over the application timeline or during examination. This means wording can change for reasons beyond translation, and an invention may be split into multiple filings.

The 'simple family' applies a clear rule: it groups documents with exactly the same set of priority applications. But continuations, continuations-in-part and divisionals can muddy those relationships. Unexamined provisional applications may also connect filings that later develop into potentially very different inventions. At global scale, the task is further complicated by patent-office errors, such as a 7 being entered as a 1 in a priority number, and inconsistent patent-number formats, which can create incorrect or missing links.

For portfolio comparison, the patent family is generally the most meaningful unit of analysis. But, as the examples above show, "family" is not a single concept. Family definitions differ in how broadly they group related filings, and a portfolio comparison is only valid if the same definition has been applied consistently to every company in scope. This decision belongs in the data layer, not in the query.

Ultimate ownership: from applicant names to economic control

Patent offices record applicants as they were named at the time of filing. They do not record corporate structure, and they do not update when structure changes.

The result is that a single corporate group may appear in patent data under hundreds of distinct applicant names: legal entities in different jurisdictions, subsidiaries, joint ventures, historical company names, transliterations, spelling variants, and entities acquired decades ago. Ownership harmonization consolidates these into an ultimate-owner structure that reflects economic control, incorporating reassignment events, and corporate transactions.

Without this layer, a competitor's portfolio is scattered across dozens or hundreds of nominal applicants. The resulting analysis is not merely imprecise; it is systematically biased. Groups with complex legal structures, extensive acquisition histories, or many national subsidiaries appear smaller than they are, while companies that file under a single entity appear disproportionately strong. Any comparison, ranking or trend built on unharmonized applicant names inherits that distortion, and no language model placed on top of the data can detect it—the data looks entirely consistent.

Ultimate ownership is also the join key between patent data and the rest of the business. Linking patent positions to revenue, R&D spend, market data, or acquisition candidates requires that the same corporate entity be identifiable in both datasets. Entity resolution, not model capability, is usually what determines whether patent-to-business analytics succeeds.

Legal status: from office events to analytically usable states

Patent offices publish very large volumes of procedural events. In their raw form these events are not analytically usable: they are jurisdiction-specific, inconsistently coded, and meaningful only in sequence.

There is an additional data-quality challenge. Many patent offices provide incomplete, delayed, or otherwise poor updates of legal status.

William Mansfield
Principal Product Manager, Strategy, Discovery & Market Intelligence (SDMI)
for LexisNexis® Intellectual Property Solutions

Establishing the most likely current status therefore requires more than collecting the published events; it requires interpreting their sequence and significance, while accounting for gaps and inconsistencies in the underlying office data.

Interpreted legal-status data translates these events into family-level states—pending, in force, inactive—that can be used to filter and compare portfolios. This is a different exercise from asking a general-purpose language model to infer status from an isolated document or event, which will produce a plausible answer with no reliable basis.

Two distinctions deserve emphasis. First, legal status is a property of a right in a jurisdiction, and a family may simultaneously be granted in some jurisdictions, pending in others, and lapsed elsewhere; a family-level state is an analytical summary of that pattern, not a substitute for it. Second, an analytical legal-status categorization is not a legal opinion on the enforceability of a particular right. It is suitable for portfolio analysis and unsuitable as the sole basis for a freedom-to-operate or enforcement decision. An AI system that presents the former as the latter has made a category error, not a factual one.

The reporting date: observing a portfolio at a point in time

Most temporal confusion in patent analytics comes from a single unexamined assumption: that a portfolio can be described by grouping today's patents by their filing or publication year.

It cannot. That approach silently applies today's portfolio composition, today's ownership, and today's metric values to the past. Patents that have since lapsed disappear from the historical record, patents acquired last year are attributed to a company that did not own them, and indicator values calculated from today's citation data are projected backwards.

More fundamentally, this approach answers a different question. A historical portfolio snapshot shows not only which inventions a company created, but which rights it chose to maintain by continuing to invest in them and pay the relevant fees. It can also reveal rights the company chose not to maintain or allowed to lapse. Such decisions may indicate a change in strategic focus, although the reason for an individual lapse cannot be inferred from status data alone. Invention activity, continued investment in a portfolio and the decision to discontinue particular rights are all relevant, but they should not be treated as the same thing.

The reporting-date concept addresses this by defining the point in time at which a portfolio is evaluated: which families were active at that date, and which information was available for calculating metrics then. A series of reporting dates produces genuine historical snapshots, making it possible to observe how portfolio size, technological positioning, and patent impact actually developed, rather than how today's portfolio happens to be distributed across past years.

The practical consequence is that filing year, publication year, priority year, reporting date, and ownership date answer different questions. An AI system must not choose among them implicitly, and a user must be able to see which was used.

Reporting date and historical ownership are separate concepts

These two ideas are frequently conflated, and the distinction matters whenever corporate transactions are involved.

A historical reporting date reproduces the active portfolio and the analytical information available at a past point in time. The treatment of ownership is a separate question: whether the portfolio is attributed to the entity that controlled it at that date, or to the group that owns it today, depends on the data model and on the purpose of the analysis.

Both views are legitimate, and they answer different questions. Current ownership applied historically shows how a group's present portfolio was built, including through acquisition. Historical ownership shows what each company actually controlled at the time, which is what is required for panel analysis, for assessing the effect of a divestiture, or for evaluating innovation performance before and after a merger. Reporting one while describing the other is a common and consequential error, particularly in due diligence, where the question is usually what was acquired, not what is now held.

Comparability: metrics, not counts

The final prerequisite concerns the metrics themselves. Patent counts are not comparable across jurisdictions and technologies: filing strategies, procedural norms, family sizes, and citation practices vary too widely. A portfolio that looks large in one classification area may be strategically marginal, and a small portfolio may be highly influential.

Analytical comparability therefore requires indicators designed for cross-jurisdictional and cross-technology comparison, with a documented methodology, a stated reference basis and a transparent relationship to the underlying records. Whether such an indicator is appropriate for a given business question is itself a domain judgement—one that a model can explain but should not make unilaterally.

Why this layer cannot be delegated to the model

Each element above shares a characteristic that makes it unsuitable for generative reasoning. The correct answer is not inferable from the immediate context, it is not verifiable by the model, and an incorrect answer is indistinguishable from a correct one in the output.

A model asked to consolidate applicant names will produce a plausible grouping. A model asked to interpret a legal event will produce a plausible status. A model asked for a portfolio in 2018 will produce a plausible number. None of these will be flagged as uncertain, because from the model's perspective nothing went wrong. This is why the failure mode described earlier, technically valid, analytically inappropriate, is more dangerous than outright hallucination, which at least tends to be visible.

The corollary is a division of labor that holds across all three models:

> *Resolve ownership, legal status, and time in the data layer. Execute calculations deterministically. Let the model interpret intent, select the appropriate analysis and explain the result.*

Where each model places this layer differs. In the SaaS model it is supplied and maintained by the provider and applied uniformly through the application. In the dataset model it arrives with the data but must be preserved by the customer's semantic layer, which can just as easily discard it through an incorrect join or an unconsidered date filter. In the MCP model it must be enforced inside the exposed tools, since an agent will only ever see what

William Mansfield
Principal Product Manager, Strategy, Discovery & Market Intelligence (SDMI)
for LexisNexis® Intellectual Property Solutions

a tool returns. In none of the three does it emerge from the AI layer itself.

One business question, three implementations

Consider a business stakeholder asking:

> *How has our strongest competitor's position in battery recycling changed since 2018, and what might that mean for our R&D strategy?*

Note how much of the question is already domain-dependent. "Our strongest competitor" presumes resolved corporate ownership. "Position" presumes a comparable metric. "Since 2018" presumes a reporting-date logic and a decision about whether portfolios are attributed to their historical or current owners. None of this is supplied by the phrasing.

In an **AI-enhanced SaaS environment**, the user could ask the question in natural language. The application could help define the technology area, identify the relevant corporate portfolios, apply consistent metrics and reporting dates, and generate a structured comparison with visualizations. The domain decisions are made by the product and are visible to the user.

In a **dataset-plus-AI environment**, the organization could combine patent trends with internal R&D projects, product plans, technical taxonomies, investment levels, and market forecasts. The answer could be tailored closely to the company's strategy, but the organization would need to design and validate the analysis—including the ownership, status and date logic the SaaS application would have applied by default.

In an **MCP-based environment**, a corporate-strategy agent could call a governed patent tool to retrieve the competitor trend, an internal R&D tool to identify overlapping projects, and a financial or market tool to assess commercial relevance. The AI host would synthesize the outputs, while each specialized server would remain responsible for its own data and calculations.

The business question is the same. What changes is the allocation of responsibility between the application provider, the enterprise data team, and the AI agent.

The models are complementary, not mutually exclusive

A mature IP technology architecture may use all three models.

A purpose-built SaaS platform can remain the primary workspace for IP professionals performing recurring patent analyses. A patent dataset in the enterprise data platform can support proprietary models and combinations with confidential business information. Selected patent capabilities can then be exposed through MCP to approved AI assistants used by R&D, strategy, or corporate-development teams.

In such an architecture:

- SaaS provides the governed analytical experience
- Dataset access provides the configurable analytical foundation
- MCP provides the reusable agent-access layer
- The harmonized patent-data layer underpins all three

MCP may sit above either a customer-controlled dataset or provider-hosted analytical services. It does not necessarily replace the SaaS or dataset model. Equally, dataset access should not be selected merely because it appears more technically sophisticated. Rebuilding a standardized analysis internally may create cost and risk without producing meaningful differentiation.

William Mansfield
Principal Product Manager, Strategy, Discovery & Market Intelligence (SDMI) for LexisNexis® Intellectual Property Solutions

The appropriate model depends on three questions:

1. **How standardized is the analytical problem?** Repeatable problems are often best addressed in a purpose-built application.
2. **How much proprietary data or methodology is required?** Analyses that depend on internal data, custom taxonomies, or company-specific models may justify dataset access.
3. **How widely must the capability be reused?** Where multiple assistants and workflows need access to the same patent intelligence, an MCP-based service layer may add significant value.

A fourth question applies to all three: who is responsible for the ownership, status, and temporal logic on which the analysis depends, and can that responsibility be demonstrated?

From AI functionality to trustworthy patent-intelligence infrastructure

The evolution of AI in patent tools is moving the discussion beyond isolated features. The strategic issue is increasingly how patent data, domain logic, and AI capabilities can be assembled into an infrastructure that is usable, explainable, and governable.

Across all three models, several principles remain constant.

Data quality precedes model quality. As section 6.5 sets out, ownership harmonization, patent-family construction, legal-status interpretation and consistent classifications determine what the AI is able to know. No layer above the data can recover meaning that was never established within it.

Temporal context must be explicit. Filing year, publication year, priority year, reporting date, and ownership date answer different questions. An AI system should not select among them implicitly.

Calculations should be deterministic. Language models can interpret intent and explain results, but portfolio metrics, legal-status filters, and trend calculations should be executed through validated analytical functions.

Outputs require provenance. Users should be able to identify the datasets, assumptions, tools, metrics, and patent records supporting a conclusion.

Recommendations require accountability. AI can surface patterns, generate hypotheses, and propose actions, but decisions on filing, prosecution, licensing, enforcement, acquisition, or portfolio abandonment remain the responsibility of qualified professionals.

The central choice is therefore not between "buying AI," "building AI," or "connecting AI." It is deciding which patent-intelligence capabilities should be packaged in a finished application, which should be customized within the enterprise data environment, and which should be made reusable across agents.

The strongest architecture will not necessarily be the one with the most autonomous AI. It will be the one that keeps data quality, patent-domain semantics, analytical methodology and human accountability aligned as AI becomes more deeply integrated into IP decision-making.

William Mansfield
Principal Product Manager, Strategy, Discovery & Market Intelligence (SDMI)
for LexisNexis® Intellectual Property Solutions

Closing Summary

Shaping the next chapter of AI in IP—together

The contributions in this report make one thing clear: AI is no longer a future consideration for intellectual property. It is already changing how inventions are created and captured, how IP professionals work, and how patent intelligence is accessed and applied to innovation and business decisions. Yet we are still at an early stage of this transformation. As AI capabilities advance, not only will the role of the IP professional evolve, but established concepts within the patent system, including inventorship, inventive step, and the role of human contribution, will also continue to be tested. The opportunity lies not simply in using AI to do existing work faster, but in using it to strengthen the contribution IP makes to innovation and business strategy. **Greater emphasis will shift toward judgment, validation, strategic thinking, and the ability to connect IP with inventors, R&D, and the wider business**. At the same time, natural-language interaction will broaden access to patent intelligence, bringing information that was once largely the domain of specialists to a much wider group of users.

But greater accessibility and increasingly powerful AI do not make expertise or trusted information less important. **They make them more important**. As AI-generated insights move closer to decisions with legal, financial, or strategic consequences, understanding the quality of the underlying data, the reliability of the analysis, and where human judgment is required becomes critical.

There is unlikely to be a single model for AI in IP. General-purpose AI, AI embedded within specialist applications, purpose-built agents, and MCP-based approaches connecting AI to trusted data sources can all have a role, depending on the task, the organization, and the consequences of the decision being supported.

The conversation should not end here

The questions explored in this report do not have final answers. Technology is evolving rapidly, new use cases are emerging, and expectations around accuracy, transparency, governance, and professional responsibility will continue to develop.

At LexisNexis, we believe that developing AI tools that genuinely support IP professionals requires continuous interaction with the IP community. Needs and expectations differ across industries, regions, and organizations of different sizes and maturity levels. That breadth of experience is essential to creating AI that works in real-world IP environments.

That is why we value conversations about where AI is delivering value, where it falls short, and what IP professionals need it to do next. Through our AI Insider program, we invite IP professionals to engage with us more closely, exchange perspectives on emerging use cases, and provide feedback that can help shape what comes next.

If the ideas in this report resonate with you or if you disagree with them, we would like to hear from you.

The future of AI in IP should not simply be something that happens to the profession. **It should evolve through continuous dialogue with the IP community, guided by real-world needs and the ambition to create meaningful value for the profession and the organizations it serves**.

About the contributors

The perspectives presented in this report are shaped by the individual experience and expertise of each author. Their contributions examine AI and IP from different angles, spanning the patent system itself, the changing role of IP professionals, practical corporate experience with AI, the democratization of patent intelligence, and the different ways in which AI can be integrated with patent software and data.

The following profiles provide further background on the contributors and the professional perspectives they bring to the discussion around AI and intellectual property.

The views expressed by the contributors are their own and do not necessarily reflect the views, positions, or policies of their affiliated organizations. The contributors are listed in order of appearance in the report.

Bowman Heiden

Co-Director of the Center for Intellectual Property (CIP) at the University of Gothenburg, and Executive Director of the Tusher Strategic Initiative for Technology Leadership, UC Berkeley



Dr. Bowman Heiden is the Co-Director of the Center for Intellectual Property (CIP) at the University of Gothenburg, which is a joint platform between academia and industry focused on the transformation of knowledge into wealth and welfare. He is also the Executive Director of the Tusher Strategic Initiative for Technology Leadership at UC Berkeley and co-chair of the Technology, Innovation, and Intellectual Property program at the Classical Liberal Institute at the NYU School of Law. Dr. Heiden was recently a member of the European Commission Expert Group on Standard Essential Patents.

Dr. Heiden is the co-founder of the Berkeley Deep Tech Innovation Lab, the Dynamic Competition Initiative, ICM Global, and Increasing Diversity in Innovation. Over the past fifteen years, Dr. Heiden has also managed over 150 innovation projects with industry, university research institutes, healthcare providers, and start-up ventures.

Dr. Heiden holds degrees in engineering, technology management, and economics, and his research is at the interdisciplinary interface of economics, law, and innovation, in particular, intellectual property, innovation economics, and competition policy in knowledge-intensive sectors. Before turning his focus to the fields of innovation strategy and policy, Dr. Heiden played professional basketball in a number of European countries.



Ruud Peters

Co-Director at the Center for Intellectual Property (CIP), CEO at Peters IP Consultancy B.V., former Chief IP Officer and Executive Vice President at Koninklijke Philips N.V.

Ruud Peters is an IP business leader and entrepreneur with both strategic and operational capabilities, who created and managed one of the world's leading and most successful IP businesses as Chief Intellectual Property Officer of Philips for a period of 15 years until the end of 2013. In 2014, he established his own IP Consultancy company providing consultancy in IP strategy and organizational matters to enterprises and institutes around the globe. In addition, he is acting as board member/advisory board member/advisor for Iprova Sarl, Lausanne, Switzerland; School of Management, Zhejiang University, Hangzhou, China; Institute for Intellectual Property Management, Zhejiang University, Hangzhou, China; Center for Intellectual Property (CIP), Gothenburg, Sweden; AI Patents, North Carolina, USA and Arbell Energy, Tel Aviv, Israel. In 2017 he was appointed as Guest-Professor at the School of Management, Zhejiang University, Hangzhou, China.

He was inducted into the IP Hall of Fame in 2010 and he received an Outstanding Achievement Award for his lifetime contributions to the IP field from Managing Intellectual Property magazine in 2014. He frequently speaks at major international IP conferences/webinars and also writes articles regularly in leading IP and business magazines. Ruud has a background in physics (Technical University Delft, The Netherlands).



Stephan Wolke

CEO of thyssenkrupp Intellectual Property GmbH

Dr. Stephan Wolke studied physics, philosophy and macroeconomics at Bonn University and graduated with a Doctorate in physics. He spent the first five years of his professional career at McKinsey & Company. After founding (and having sold) two businesses of his own, Stephan served in several functions at Bayer AG as well as Danaher Corp. He joined thyssenkrupp in 2011 and is responsible for Intellectual Property at the corporate level as well as head of management board of thyssenkrupp Intellectual Property GmbH.

Stephan Ising

Head of IP Strategy / IP Research of thyssenkrupp Intellectual Property GmbH



Stephan Ising is the Head of IP Strategy & IP Research at thyssenkrupp Intellectual Property GmbH, a role he has held since July 2020. With nearly a decade of experience at thyssenkrupp, Stephan has developed a deep expertise in intellectual property strategy and research, helping drive innovation and strategic decision-making within the company. Prior to his current role, Stephan served as an Expert in IP Strategy & Research from 2015 to 2020. He began his journey at thyssenkrupp as a Werkstudent in 2014.

Stephan holds a Master of Science in Industrial Engineering from the University of Duisburg-Essen and a Master 2 in IP Law and Management from the Center for International Intellectual Property Studies (CEIPI). His academic background, combined with his professional experience, positions him as a key figure in shaping the future of intellectual property strategy at thyssenkrupp.



Sakae Nakamura

Executive Advisor to LexisNexis® Intellectual Property Solutions / Senior Advisor, IP Landscape Association, Japan

Sakae Nakamura is an Executive Advisor to LexisNexis Intellectual Property Solutions and a recognized pioneer in IP landscaping and strategic patent intelligence. She spent nearly four decades at Asahi Kasei Corporation, where she held senior leadership positions in intellectual property and helped establish and lead the company's IP intelligence activities. Nakamura has played a prominent role in advancing IP landscaping in Japan, including serving as Representative Secretary of the IP Landscape Association. Her work has focused on integrating patent and technology intelligence into business strategy, management decision-making and innovation planning.



Tim Alexander Oelmann

VP, Head of Competence Center Business IPR Management at Deutsche Telekom AG

Tim Alexander Oelmann heads the Competence Center for Business IPR Management at Deutsche Telekom AG, shaping how intellectual property creates strategic and measurable value across a global technology group.

His mission is to move IP beyond protection — turning patents, technology, and innovation into business impact. He connects IP strategy with corporate strategy, R&D, innovation processes, business models, and technology value chains to build sustainable competitive advantage. A strong believer in the power of people and ideas, Tim drives initiatives that strengthen the connection between inventors, Research & Development, and the business.

Through programs such as the Inventathon, Ideation Lab, and Patent Garage, he promotes a culture where great ideas are identified faster, developed further, and translated into valuable intellectual property and business opportunities.

His areas of responsibility include Business Patent Management, strategic patent portfolio management, inventor recognition and reward, IP commercialization, governance, and the strategic positioning of IP in emerging technology domains.



Daniel Uhl

Patent Information Specialist at Deutsche Telekom AG

Daniel Uhl is a Patent Information Specialist at Deutsche Telekom AG. He holds a Diplom in Geography from Freie Universität Berlin (2006) and an MBA from Steinbeis Universität Berlin (2009). With extensive experience at Deutsche Telekom, Daniel specialized in Customer Insights and Market Intelligence from 2008 to 2011, and since 2012 has served as a Patent Intelligence Specialist, recently expanding his role to include Ideation Specialist responsibilities in 2024. Additionally, from 2023 to 2026, Daniel served as Chair of the Patent Development Group (PDG) Task Force on AI Sorting, where he demonstrated leadership and expertise at the intersection of technology innovation and intellectual property.



William Mansfield

Principal Product Manager, Strategy, Discovery & Market Intelligence (SDMI) for LexisNexis® Intellectual Property Solutions

William Mansfield is part of Product Strategy and Discovery for LexisNexis Intellectual Property Solutions, including PatentSight+ and Protégé. He directs the discovery process that turns customer and expert input into working prototypes and product specifications, working closely with engineering, strategy, and customer-facing teams to bring new capabilities from concept through to release. His work spans data engineering, competitive intelligence, and AI integration. He previously oversaw LexisNexis Intellectual Property Solutions' global consulting work and Customer Success teams, and led initiatives on patent data quality. He has an Economics background and has previously worked on developing Market and R&D strategy in the solar industry.

About LexisNexis® Intellectual Property Solutions

LexisNexis® Intellectual Property Solutions brings clarity to innovation for businesses worldwide. We enable innovators to accomplish more by helping them make informed decisions, be more productive, comply with regulations, and ultimately achieve a competitive advantage for their business. Our broad suite of workflow and analytics solutions, including LexisNexis® PatentSight+™ with Protégé™, LexisNexis® Classification, LexisNexis® IPlytics™, LexisNexis PatentOptimizer®, LexisNexis PatentAdvisor®, and LexisNexis TotalPatent One®, LexisNexis® IP DataDirect, enables companies to be more efficient and effective at bringing meaningful innovations to our world.

By combining trusted, enriched data, industry-proven metrics, and AI-powered capabilities such as Protégé, we help organizations turn complex intellectual property data into clear, actionable insights that support confident strategic decision-making.

We are proud to directly support and serve these innovators in their endeavors to better humankind.

LexisNexis® Intellectual Property Solutions
Bringing clarity to innovation

To learn more, visit us at:
LexisNexisIP.com

LexisNexis, LexisNexis PatentAdvisor, LexisNexis PatentOptimizer, LexisNexis TotalPatent One and the Knowledge Burst logo are registered trademarks registered in the U.S. or other countries and LexisNexis Protégé and Protégé are trademarks of RELX Inc. PatentSight is a registered trademark and PatentSight+ and IPlytics are trademarks of LexisNexis Intellectual Property Solutions GmbH. Other products or services may be trademarks or registered trademarks of their respective companies. Copyright © 2026 LexisNexis.

LexisNexis®