



CHALMERS
UNIVERSITY OF TECHNOLOGY



Intellectual Property Management in the Age of Artificial Intelligence:

Exploring the Value and Organizational Implications of AI in Patent Management

Master's thesis in Management and Economics of Innovation

KLARA SPETZ
AGNES MAGNUSSON

DEPARTMENT OF TECHNOLOGY MANAGEMENT AND ECONOMICS
DIVISION OF ENTREPRENEURSHIP AND STRATEGY

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

DEGREE PROJECT REPORT 2026

Intellectual Property Management in the Age of Artificial Intelligence:

Exploring the Value and Organizational Implications of AI in Patent Management

KLARA SPETZ
AGNES MAGNUSSON



CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Technology Management and Economics
Division of Entrepreneurship and Strategy
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026

Intellectual Property Management in the Age of Artificial Intelligence:
Exploring the Value and Organizational Implications of AI in Patent Management
KLARA SPETZ
AGNES MAGNUSSON

© KLARA SPETZ, 2026.
© AGNES MAGNUSSON, 2026.

Supervisor: Rabih Darwish, Department of Technology Management and Economics
Examiner: Marcus Holgersson, Department of Technology Management and Economics

Degree project report 2026
Department of Technology Management and Economics
Division of Entrepreneurship and Strategy
Chalmers University of Technology
SE-412 96 Gothenburg
Sweden
Telephone +46 31 772 1000

Gothenburg, Sweden 2026

Intellectual Property Management in the Age of Artificial Intelligence:
Exploring the Value and Organizational Implications of AI in Patent Management

KLARA SPETZ
AGNES MAGNUSSON

Department of Technology Management and Economics
Chalmers University of Technology

Abstract

Artificial intelligence (AI) is rapidly changing how organizations handle knowledge-intensive and analytical tasks. Within patent management, increasing amounts of patent data, growing technological complexity, and rising competitive pressure have resulted in a growing interest in how AI can support intellectual property (IP) work. At the same time, the integration of AI into patent management raises new organizational challenges. While previous research has demonstrated promising technical applications of AI in patent-related activities, there is still limited understanding of how AI is currently used, perceived, and implemented in practice within corporate patent management functions. Building on a qualitative research approach, this study is based on 32 interviews with professionals from corporate IP departments, IP firms, and AI tool providers to examine current applications and the potential of AI in patent management. The findings illustrated that AI is perceived to hold significant potential across several patent management activities, particularly in search optimization, drafting support, portfolio analysis, and analytical work involving large amounts of data. However, despite strong interest and optimism surrounding AI, implementation remains relatively limited. Several organizational barriers were identified, including confidentiality concerns, lack of trust in AI, fear of losing professional judgment and expertise, psychological inertia, unrealistic expectations regarding AI performance and introduction of bottlenecks. The study concludes that the challenge of integrating AI into patent management is not primarily about whether AI has potential, but about organizations' ability to adapt in ways that allow this potential to be realized without compromising the requirements and professional standards under which patent management operates. The study thereby contributes to the emerging literature on AI in patent management by providing empirical insights into how organizations currently approach AI adoption and the organizational conditions shaping implementation in practice.

Keywords: Artificial Intelligence, Patent Management, Intellectual Property, AI adoption, Patent Processes

Acknowledgments

We would like to thank our supervisor at Chalmers University of Technology, Rabih Darwish, for his high engagement, valuable advice and encouragement during the project. We would also like to thank our examiner, Marcus Holgersson, for contributing with his expertise.

Furthermore, we would like to express our sincere gratitude to our supervisors at the collaborating company for their engagement, valuable knowledge and insights. Their guidance and support helped us navigate challenges along the way and their openness made us feel welcome and included as a part of the team. We would also like to thank the IP team and interview respondents for their collaboration and contributions, which have been essential for this thesis. A special thanks is extended to a person at the collaborating company for his mentorship during this project. Purely out of generosity and interest, he not only provided us with invaluable contacts, but also helped us understand the complexities and mechanisms of IP work that are rarely articulated explicitly.

Moreover, we would like to thank all external respondents who contributed to this thesis. Without their time, openness and willingness to share their experiences, this study would not have been possible. Beyond this, our discussions were not only invaluable to this study but also enjoyable and inspiring.

Klara Spetz & Agnes Magnusson, Gothenburg, May 2026

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis:

AI	Artificial Intelligence
ML	Machine Learning
NLP	Natural Language Processing
FTO	Freedom To Operate
IDF	Invention Disclosure Form
IP	Intellectual Property

Contents

List of Acronyms	ix
1 Introduction	1
1.1 Background	2
1.2 Problem statement	3
1.3 Purpose and research questions	4
1.3.1 Purpose	4
1.3.2 Research Questions	5
1.4 Limitations	5
2 Theoretical framework	7
2.1 Intellectual property	7
2.2 Patent management	8
2.3 Diffusion of innovations	9
2.4 Artificial intelligence	11
2.4.1 Definition and foundation of artificial intelligence	11
2.4.2 AI concepts and methods	12
2.4.3 AI in organizations	13
2.4.4 Organizational adoption of AI	15
2.5 Defining capabilities of AI	17
2.6 Patent management processes and the potential of AI	20
2.7 Theoretical mapping of AI capabilities in relation to patent manage- ment	28
3 Methods	31
3.1 Research approach	31
3.2 Research design	32
3.3 Data collection	33
3.4 Data analysis	35
3.5 Ethical considerations	36
4 Empirical findings	39
4.1 Identified patent management processes	39
4.2 AI supported patent management tools	43
4.3 How is the potential of AI realized in corporate patent management processes today?	46
4.3.1 Patent filing	46

4.3.2	Portfolio management	50
4.3.3	Intelligence	52
4.3.4	Enforcement & exploitation	54
4.3.5	Risk assessment	57
4.4	What are the organizational barriers of adoption in relation to integrating AI into corporate patent management processes?	59
4.4.1	Confidentiality and security requirements as a barrier of adoption	59
4.4.2	The search for the perfect tool as a barrier of adoption	60
4.4.3	Lack of trust as a barrier of adoption	62
4.4.4	Introduction of bottlenecks as a barrier of adoption	64
4.4.5	Fear of losing competence and professional judgment as a barrier of adoption	66
4.4.6	Inertia and social stigma as a barrier of adoption	68
5	Discussion	71
5.1	The gap between the potential of AI and practical implementation . .	71
5.2	Understanding the barriers of AI adoption in patent management . .	75
5.3	Practical implications	83
6	Conclusion	87
6.1	Limitations and future research	88
	References	91

1

Introduction

There is no doubt that Artificial Intelligence (AI) is transforming our society and reshaping how organizations operate. What still remains unclear is to what extent and how to capture the value of this promising opportunity (Sætra, 2023). As with all technological changes we must ask ourselves if this transformation can generate value, will cause us harm or is an unavoidable necessity for competitive survival. One domain where these broader questions have recently emerged is within the landscape of intellectual property (IP). Intellectual property functions now face both complex challenges and promising opportunities because of AI-driven transformation.

As businesses strive to prevent others from commercializing their intellectual creations or inventions, efficient IP management becomes critical. The shift towards a pro-patent era, beginning in the 1980s, has strengthened the role of IP protection in business strategy (Granstrand & Holgersson, 2015). At the same time, the volume of patent documents continues to grow. According to World Intellectual Property Organization (WIPO) the number of patent applications is currently increasing, reaching 3.7 million during 2024, an increase of 4.9%, representing the fastest growth rate since 2018 (WIPO, 2025). As organizations are required to manage these growing amounts of patent data, the need for more efficient IP management processes becomes increasingly important. In this context, recent advancements in AI have introduced possibilities for supporting patent-related activities. The potential of these technologies has generated considerable interest among IP professionals. A recent study by Questel (2025) found that 77% of respondents expressed enthusiasm toward AI, while 58% reported that their regular use of AI tools had resulted in a positive impact on their work.

In alignment with these evolving opportunities, the purpose of this study is to explore how AI could be applied in patent management in corporate organizations. Rather than analyzing individual tools or isolated activities, the study aims to understand the potential and current applications of AI in patent management and its processes. To address this, the study combines insights from multiple organizations with a more detailed investigation of one multinational industrial firm in the automotive sector, in this study referred to as the collaborating company. This design makes it possible to capture both industry level challenges and process level insights.

1.1 Background

It has become increasingly common for companies to consider their patent portfolio as a competitive advantage and as a source of income. However, the patent process is costly, and a firm must therefore weigh the potential benefits against the costs of a specific patent strategy (Al-Aali & Teece, 2013). Because of this, firms need to effectively manage their patents, for example through planning and evaluation, in order to realize value from them (Soranzo et al., 2017).

Patent management comprises several interrelated processes, and these should be adjusted to the firms' specific needs and strategy (Soranzo et al., 2017). Some core processes include patent generation, securing and analyzing freedom to operate (FTO), portfolio management, exploitation and enforcement and intelligence (Agostini et al., 2022). Each of these processes involves a wide variety of activities necessary for effective patent management. For example, activities such as freedom to operate analysis, novelty analysis and state of the art analysis require extensive searching of patent databases and non-patent literature, as well as interpretation of complex patent texts to judge similarity and infringement risks (Soranzo et al., 2017; Setchi et al., 2021). As the number of patents worldwide increase, these activities become even more data-intensive.

Another analysis intensive activity in patent management is portfolio management, which involves continuously assessing and monitoring the patent portfolio in terms of its value for the firm (Agostini et al., 2022) a process that can become increasingly challenging when the number of patents grows. Furthermore, patent drafting includes the requirement of producing high quality patent texts and claims (WIPO, 2023), which requires both legal and technical knowledge.

Given that patent management comprises data intensive and analytically demanding processes, the rapid advancement within AI becomes particularly relevant. These developments raise the question of whether AI could offer meaningful benefits in corporate patent management.

Literature shows that AI demonstrates high potential in patent management. For example, AI has shown potential for assisting patent examiners in patent prior art search, which can offer both potential time and cost savings (Setchi et al., 2021). Further literature has shown successful use of large language models (LLMs) for producing high quality patent summarizations (Trappey et al., 2025). In addition, deep learning models such as convolutional neural networks have demonstrated potential in optimizing the activity of patent classification which has traditionally been a manual task (Li et al., 2018). Deep learning can further be used for technology forecasting, shown by Zhou et al. (2020) where a tested model was able to predict emerging technologies one year before they appeared. Deep learning models have also been tested for automatic patent value evaluation, which showed improved accuracy compared to previously used methods tested on manufacturing IoT patents (Trappey et al., 2021). Furthermore, (Hain et al., 2022) developed a method based on natural language processing and k-nearest neighbor techniques for

assessing technical similarity between patents, which showed promising potential for effective patent comparison. Such models, or similar, could potentially improve prior art search, state of the art analysis, freedom to operate analysis as well as landscaping.

In addition, the increasing number of AI based tools available for patent management operations indicates growing commercial interest in AI applications in this domain. For example, DeepIP offers products in both patent drafting, patent prosecution, FTO analysis, landscaping and portfolio intelligence (DeepIP, n.d.). IPRally offers services such as patent search, patent review and portfolio analysis (IPRally, n.d.), and soon to be launched Patentext that offers patent drafting services (Patentext, n.d)

However, with rapid advancement of AI, several risks should be taken into consideration (WIPO, 2024). For example, generative AI can produce incorrect or unreliable information. Furthermore, the regulatory landscape around AI is evolving, which can pose challenges for business to adapt to changing requirements. Moreover, several risks related to the handling of confidential information should be considered, which are especially important in patent management as it involves handling unpublished information about inventions where leakage can lead to significant consequences. For example, some AI tools might use user input to further train their models which raises the question whether this information could become available to the public. In addition, third party suppliers of AI tools monitor user prompts for inappropriate use which gives them direct access to the information in the prompts (WIPO, 2024).

Previous studies show that AI tools can deliver promising results in patent-related activities. However, most studies discuss specific models developed for specific tasks and tested on specific datasets. In addition, the increasing number of AI tools on the market indicates that there is promising potential in using AI in various tasks within patent management. However, there is still a research gap regarding whether the potential is actually realized within corporate patent management.

1.2 Problem statement

In the modern business environment, intellectual property serves a purpose that extends far beyond protecting innovations, with a company's IPRs acting as a key driver of competitive advantage (Holgersson & van Santen, 2018). Additionally, the accelerating pace of innovation, combined with the surge in patent applications underscores the need for more efficient patent management processes (WIPO, 2025). AI has demonstrated significant potential to improve efficiency and enable automation across a wide range of organizational functions (Singla et al., 2025), which explains the growing interest among IP management functions in adopting AI into their workflows. However, enthusiasm and potential do not necessarily translate into value creation.

To realize the value promised by AI, IP functions must not only address its potential benefits but also the associated risks (WIPO, 2024). The amount of failed AI investments grew in 2025 and financial returns are not guaranteed (Pettersson, Björkdahl, & Holgersson, 2025). Consequently, firms face multiple strategic decisions, including whether to adopt AI, which AI solutions to implement, if the IP function is ready, and which patent management tasks are most suitable for AI. Answering these questions is further complicated by the novelty of this subject. At the same time, organizations are pressured to act quickly to avoid losing their competitive strength as a result of failing to keep pace with the effectiveness of other IP functions.

Patent management provides favorable conditions for the application of AI. Patent documents are available in large quantities, are well structured, follow strict guidelines and are publicly accessible. Together, these characteristics provide high quality training data, enabling a strong foundation for developing and applying AI models. Although previous studies highlight that AI is increasingly used in patent management (Lin & Chou, 2025), how organizations apply, or are planning to apply these technologies, and the value they bring remains unclear.

Existing literature identifies several promising use cases, such as prior art search, patent summarizations, patent classification, technology forecasting and patent evaluation (Hain et al., 2022; Li et al., 2018; Setchi et al., 2021; Trappey et al., 2021; Trappey et al., 2025; Zhou et al., 2020). However, there is limited insight into how AI is integrated into organizational patent management processes and what barriers might influence their ability to adopt.

This reveals a gap in literature, while enthusiasm, perceived potential and early adoption are evident, there is still limited understanding of how AI can be implemented in patent management processes within established organizations, and which adoption barriers may constrain that implementation.

Organizations need a better understanding of the current landscape of AI in patent management in order to evaluate if it is a good fit for their organization and if there truly are opportunities for value creation. This thesis will provide a starting point for narrowing this gap in research and provide useful information about the current applications and perceived potential of AI in patent management and the organizational barriers associated with its implementation.

1.3 Purpose and research questions

1.3.1 Purpose

The purpose of this study is to broaden the understanding of AI in patent management by examining its potential within patent-related processes, its current implementation in practice, and the organizational barriers that influence its adoption.

1.3.2 Research Questions

This thesis aims to answer the following research questions:

RQ1: What are the potential applications of AI in corporate patent management processes?

RQ2: How is the potential of AI realized in corporate patent management processes today?

RQ3: What are the organizational barriers of adoption in relation to integrating AI into corporate patent management processes?

1.4 Limitations

This thesis focuses exclusively on patents in order to maintain a clear and manageable analytical scope. Other forms of IPRs, such as trademarks, design rights and copyrights are excluded.

Furthermore, the study is focused on the IP management function within the collaborating company, which is a multinational original equipment manufacturer within the automotive industry (OEM). However, the study is not limited to the automotive industry as companies with IP departments in other industries are also included in the analysis to enable comparisons with the collaborating company and provide broader perspectives on AI use in patent management.

Patent rights are territorial, meaning that each country has its own standards and patent system for protection of inventions, and the invention is only protected within the country where the patent was granted. Consequently, regulatory frameworks and standards differ between countries. Since the collaborating company of this thesis has its main operations in Sweden, the focus is limited to Swedish patent standards.

The scope of this thesis is limited to the application of AI tools in relation to patent management activities conducted in IP departments. Accordingly, issues concerning AI-generated inventions and the legal ownership of inventions created by AI are excluded.

2

Theoretical framework

2.1 Intellectual property

An intellectual property (IP) is an intangible resource, in other words, a non-physical asset (Granstrand & Holgersson, 2015). Examples include technical inventions, designs and trademarks. World Intellectual Property organization (WIPO) defines intellectual property as: “ [...] creations of the mind, such as inventions; literary and artistic works; designs; and symbols, names and images used in commerce.” (WIPO, n.d.-a).

Intellectual property rights (IPRs) are legal rights that, under certain conditions, exclude others from commercializing the IP. IPRs are typically temporary, transferable and licensable. Examples of IPRs include patent rights, trademarks, copyrights, design rights and trade secret rights. Each are bound to different legal frameworks, which can also differ between countries and change over time (Granstrand & Holgersson, 2015).

Patent rights can be granted for technical inventions. To be granted a patent, the invention must at least be novel, non-obvious and industrially applicable. This means that at least one feature of the invention must not have been known to the public before patent filing, that the invention is not obvious to a person having average skill in the relevant technical area and that the invention provides a practical benefit (WIPO, 2023). A patent is a legal exclusive right that gives the inventor the right to decide if and how others are allowed to use the invention. In other words, the patent owner can stop others from making, importing, selling, distributing and using the invention in the region where the patent was granted within the time frame the patent is valid by taking legal action for patent infringement in court (WIPO, n.d.-b). Patents are territorial, meaning that they only apply in the country within which they were granted and are only valid during a limited time frame, which is usually 20 years (WIPO, 2023). It is possible to seek patent protection in one or multiple countries (WIPO, n.d.-b).

IPRs are important for several reasons. They can for example create incentives for individuals and organizations to invest in innovation and R&D. IPRs also result in disclosure of new knowledge, thereby enabling the innovations to benefit society in general. There are also moral justifications for IPRs, such as the idea that cre-

ators should have rights to the result of their own work (Holgersson & Granstrand, 2015). Regarding patents, research show that firms' motives behind patent applications vary widely (Holgersson & Grandstand, 2017). The primary drivers include protecting product technologies, securing freedom to operate (FTO) which means the ability to commercialize the product without infringing IPRs of other actors, blocking competitors from technology areas and creating retaliatory power against competitors.

2.2 Patent management

Patent management refers to the coordination of patent related activities and processes within a company. Organized patent management is becoming increasingly important as innovation processes are becoming more complex and patents are increasingly seen as important assets for competitiveness (Agostini et al., 2022-a). Agostini et al. (2022-a) revisited and refined the patent management framework initially presented by Agostini et al. (2019), using data collected from European patent managers. The analysis resulted in a framework consisting of five core processes and two supporting dimensions of patent management.

The first core process is patent generation, consisting of state-of-the-art analysis that involves assessing existing knowledge relevant to the field of the potential patent application, in other words generating an overview of existing knowledge (Agostini et al., 2022-a). Geographical scope is also a part of this process, where decisions about in which countries patent applications should be filed for is analyzed (Agostini et al., 2022-a). The next process is freedom to operate (FTO) which focuses on assessing the risk of infringement and securing FTO for the product or business (Agostini et al., 2022-a; Soranzo et al., 2017).

Furthermore, the process of portfolio management involves continuously reviewing the portfolio in terms of its value for the firm and taking action based on this (Agostini et al., 2022-a). The exploitation and enforcement process includes external exploitation questions, such as licensing and selling, proactive licensing, signaling through for example disclosure of information and activities related to detecting and responding to infringement (Agostini et al., 2022-a). The intelligence process, which spans across all these activities, involves monitoring and analyzing patents from other firms. Through this review, the company can gain an understanding of the current landscape in terms of competitiveness and technology (Agostini et al., 2022-a). Furthermore, two supporting dimensions are also presented. These include the alignment of the patent strategy and the overall strategy of the firm as well as organizational features such as patent culture, management involvement and patent committee (Agostini et al., 2022-a) This framework serves as a baseline for understanding which activities are included in patent management work conducted in IP departments in this thesis.

2.3 Diffusion of innovations

The theory of diffusion of innovations is most commonly connected to Everett Rogers and the book *Diffusion of Innovations* (1983). By Rogers, an innovation is defined as an idea, object or practice that is perceived as new by an individual and diffusion refers to how these innovations are spread among individuals in a social system. Rogers argues that the newness of an innovation naturally creates uncertainty among individuals regarding the possible consequences the innovation may bring. This uncertainty acts as a motivation to search for more information about the innovation in order to reduce this uncertainty. Not until this uncertainty is reduced will an individual make a decision about whether to adopt or reject the innovation. Therefore, the most central part of the diffusion of innovations theory presented by Rogers (1983) is the gathering of information in order to reduce the uncertainty associated with the innovation.

Rogers (1983) presents four main elements of the diffusion of innovations. This is the innovation itself, communication channels, time and the social system. Communication channels refers to the ways in which information about the innovation is spread, for example through media or interpersonal communication. Time is another central element, which involves both the innovation decision process of individuals, individuals' innovativeness and the rate of adoption, in other words how quickly an innovation is spread and adopted within a social system. The social system refers to the group of individuals, organizations or other entities within which the diffusion takes place. Norms and other structures within this system can also influence how innovations are diffused and adopted.

While these elements simultaneously influence the diffusion process, it has been shown that a large part of the explanation to the rate of adoption of an innovation lies in the characteristics of the innovation itself (Rogers, 1983). However, Rogers argues that it is not the characteristics of the innovation in objective terms that affect adoption, but rather how these characteristics are perceived by potential adopters. To describe this, Rogers presents five perceived attributes of innovations, including relative advantage, compatibility, trialability, complexity and observability.

Relative advantage refers to the extent to which individuals perceive an innovation as superior to existing alternatives. Generally, innovations that are associated with a greater perceived relative advantage tend to diffuse more rapidly within a social system. It can be measured in several dimensions, such as economic profitability and social prestige. The economic dimension relates to the financial benefits of an innovation such as reduced costs, increased profitability or a lower price. While innovations that are associated with economic advantages often diffuse more rapidly, Rogers (1983) emphasizes that economic factors alone can not explain the rate of adoption. In addition to financial motives, the social dimension can also strongly influence adoption, meaning that individuals often adopt innovations primarily motivated by that they are associated with social status. However, the importance of different dimensions may vary depending both on the type of innovation and the characteristics of the adopter.

Furthermore, Rogers (1983) argues that the perceived relative advantage is one of the most important aspects of determining the rate of adoption. Because of this, it is common to try to increase the rate of adoption by influencing this attribute through different types of incentives for adopters. These can be in both monetary and non-monetary terms. By providing such incentives, the perceived relative advantage of an innovation may increase thereby accelerating adoption.

However, relative advantage is not the only attribute of an innovation to consider. The second attribute presented by Rogers (1983) is compatibility which refers to how well the innovation aligns with the potential adopters values, previous experiences and needs. This is important, as innovations that in some way conflict with the norms or values of a social system are likely to diffuse more slowly. Furthermore, individuals often assess a new innovation based on earlier experiences. Therefore, a negative experience with previous innovations may negatively affect the rate of adoption of a new innovation. An additional aspect of compatibility concerns the needs of the users. When individuals perceive that an innovation fulfills a need, the likelihood of adoption increases.

Another key attribute identified by Rogers (1983) is trialability, which refers to the extent to which it is possible to experiment with an innovation before adoption. If an innovation can be tested on a limited basis before full adoption, individuals have the opportunity to learn by doing and gain information about the innovation, thereby reducing uncertainty about the innovation. Such innovations are generally adopted more rapidly.

Additionally, the attribute of complexity is how complex the innovation is perceived by the potential adopters. If an innovation is difficult to understand and requires that the individual develop new knowledge, adoption might be slower compared to innovations that are perceived as easy to use and understand.

Finally, the fifth attribute discussed by Rogers (1983) is observability, which refers to how apparent the results of an innovation are to other individuals. When the benefits or consequences can be observed by others, it can encourage discussion and information sharing. This makes it easier for potential adopters to gain knowledge about the innovation and thereby reduce uncertainty and positively influences the rate of adoption.

To summarize, how potential adopters perceive these attributes can help explain the rate of adoption of an innovation. However, the rate of adoption is influenced by several interacting elements such as, the type of innovation-decision, communication channels and the social system.

2.4 Artificial intelligence

2.4.1 Definition and foundation of artificial intelligence

The term artificial intelligence itself was first introduced by John McCarthy, in a preparatory document for the 1956 Dartmouth conference (Cordeschi, 2007). The workshop was founded by the Rockefeller Foundation with the aim of bringing together scientists from different disciplines to explore the possibility of creating machines capable of simulating human intelligence. After this conference, many years of significant progress in AI research followed. Early programs such as ELIZA, one of the first tools to simulate a conversation and the General Problem Solver, that were capable of automatically solving simple problems such as towers of Hanoi, contributed to an optimistic outlook for AI research. These advancements resulted in substantial investment in AI. However, by the early 1970s this progress stalled, in 1973, the US congress criticized the large spending on AI, and in the same year James Lighthill published a report arguing that the exciting promises about AI were unrealistic (Haenlein & Kaplan, 2019).

The true rise of AI, as we know it today, did not occur until the 2010s. The reason for this was that the early systems like ELIZA, were so-called expert systems that worked well for structured tasks but failed at tasks requiring flexibility, they could not learn from data or understand context, important factors for the modern understanding of AI. Although ideas of learning systems such as neural networks had existed early on, computers were not powerful enough to operate them. In 2015 AlphaGo, developed by Google, demonstrated a modern form of neural networks by the type of deep learning. It received recognition for beating the world champion in the complex board game Go. This achievement overturned the long-standing assumptions about the limits of AI. Today, neural networks and deep learning are the foundation of most AI systems and are delivering on the promises made decades ago (Haenlein & Kaplan, 2019).

Researchers have historically approached AI from different conceptual angles, leading to a lack of a universal definition. Some have defined it by how closely machines can replicate human performance while others use a more formal definition relating to rationality or doing the “right thing” (Russell & Norvig, 2022). Today, AI is often described as “A technology that enables machines to imitate various complex human skills” (Sheikh et al., 2023). However, this definition is limited as it does not clarify what these complex human skills include and therefore fails to provide a precise description of what AI is. Other studies attempt to explain these skills or tasks, for example by defining abilities such as taking action, pursuing objectives or learning through feedback (Sheikh et al., 2023). One example of a definition that relates AI to abilities is the definition by the European commission’s High-level expert group on Artificial Intelligence, which describes AI as “Systems that display intelligent behavior by analyzing their environment and taking actions – with some degree of autonomy – to achieve specific goals” (High-Level Expert Group on Artificial Intelligence, 2018). However, as Sheikh et al. (2023) argue, this definition still has shortcomings, as it relies on imprecise concepts and is too broad. Definitions

like this could in theory fit multiple technologies that most people would not classify as AI.

The difficulty of defining AI with precision relates to the fact that it attempts to replicate human intelligence, something that remains only partially understood. Without a complete understanding of natural intelligence, a precise definition of how to imitate it artificially becomes impossible. The research on both human and artificial intelligence will continue to evolve and insights from each domain will affect the other, driving a deeper understanding of intelligence overall (Sheikh et al., 2023). Defining AI becomes even more challenging because the field is changing rapidly with new technologies constantly emerging (Gbadegeshin et al., 2021).

Due to AI's wide range of meanings, the broader computer science and information technology community often defines it by the techniques that have emerged from it, such as neural networks or data mining. Over time, AI has branched into many sub-fields such as machine learning or robotics and many researchers identify more with these areas than AI as a whole (Wang, 2019).

Given the scope of examining AI in relation to patent management, it is sufficient for the purposes of this thesis to adopt WIPO's description of AI as: "learning systems; that is, machines that can become better at a task typically performed by humans with limited or no human intervention". This definition can be divided into distinct technological categories and covers a broad range of underlying techniques (WIPO, 2019).

2.4.2 AI concepts and methods

The evolution of AI has led to the emergence of several sub-fields and methods that form the foundation of modern AI systems. One of the most widely deployed AI subsets is Machine learning (ML). ML refers to models that improve their performance by learning patterns from data. ML models hold the abilities of classification, decision making and predictions on new data by learning patterns from training data rather than relying on a set of manually defined rules. A central branch of ML is deep learning, which is based on deep artificial neural networks. Unlike traditional ML methods that rely on manually defined algorithms and rules, deep learning uses interconnected networks of mathematical operations that enable models to capture highly complex patterns in data. However, it also requires greater computational resources and larger datasets (Bergman, n.d)

One of the most common deep learning models is large language models (LLMs). LLMs have the ability to generate and understand natural language. The process of transforming human language into formats that computers can work with, as well as generating human-readable language from machine based information is called natural language processing (NLP) (Trappey et.al., 2020). NLP builds on algorithms that understand context, structure, intent and entities of the human language (Taulli, 2023). LLMs are the first AI system that enables humans to communicate naturally with computers, as it can handle large amounts of unstructured

language, while capturing both context and nuance (Stryker, n.d.).

The emergence of LLMs has contributed to the rise of Generative AI. Generative AI is a branch of AI focusing on generating new content in response to a user's prompt (Bergman, n.d.). Commonly known examples of generative AI are ChatGPT and Google Gemini (Stryker, n.d.). The new content that generative AI creates is based on extensive training on large amounts of data (Stryker & Shapicchio, n.d.). The outputs can include, for example, text, video, code and images (Bergman, n.d.). Generative AI tools are based on foundational deep learning models and the specific type of foundational model depends on the type of intended content generation. For example, for text generation LLMs constitute the foundational model (Caballar, n.d.). Because of the creative capabilities of generative AI, there are many use cases. Examples include AI-chatbots, creation of synthetic data, generation of visual content customized for specific audiences and code generation (Caballar, n.d.).

More recently, a new branch of AI called agentic AI has gained significant interest among researchers, organizations, and the public. Agentic AI advances generative AI by not relying on predefined rules or continuous human intervention. Instead, it acts independently with limited supervision to achieve goals and solve problems. While generative AI creates content, agentic AI broadens this by taking actions to achieve specific objectives (IBM, n.d.). Agentic AI can retain memory, perceive context, operate continuously, and improve through feedback (Al-Mahdawi, 2025). It also enables collaboration across different agents (Bandi et al., 2025). These features enable human-like problem-solving (Sakhare et al., 2025). Use cases include performing routine activities, assisting in decision making, managing complex problems, and delivering personalized support by adapting to the surrounding context (Bandi et al., 2025).

When discussing agentic AI, it is important to distinguish between AI agents and agentic systems. AI agents are autonomous software systems capable of perceiving their environment and reasoning about it, resulting in actions directed towards achieving a specific task. In other words, it operates within a specific task scope. On the other hand, agentic AI is a multi-agent system in which multiple specialized agents collaborate to handle more complex tasks (Sakhare et al., 2025).

While these AI systems differ in their primary objectives, functionalities and levels of autonomy, they are not isolated technologies but can instead be combined and integrated to complement each other within broader AI applications (Sakhare et al., 2025).

2.4.3 AI in organizations

AI has evolved from a largely obscure scientific concept to an integrated element of everyday life. Its influence now extends beyond personal use and is reshaping how organizations make decisions and engage with stakeholders (Haenlein & Kaplan, 2019). Many researchers predict that AI is a disruptive technology capable of reshaping both organizations and society (Scheepers, Lui, & Ngai, 2023). The

magnitude of this transformation is illustrated by an analysis by McKinsey & Company that estimates that generative AI alone has the potential to add between 2.6 and 4.4 trillion dollars annually in additional value to the global economy (Chui et al., 2023). However, implementing AI is a time-consuming process that involves substantial financial investment as well as significant risks (Lee, Scheepers, Lui, & Ngai, 2023). Levels of organizational adoption vary and further study by McKinsey & Company estimates that nine out of ten organizations report using some type of AI tool on a regular basis. However, most companies are still in the pilot stage and have not scaled the implementation across the entire organization (Singla et al., 2025).

AI brings different benefits to different organizations, and common advantages include enhanced operational efficiency, improved customer experience and increased profitability (Scheepers, Lui, & Ngai, 2023). However, it is also expected that implementation of AI can lead to undesired effects. For example, while AI can strengthen knowledge management, creation and sharing, it may also lead to a decline in employees' skills or knowledge development as organizations increasingly rely on AI systems instead of human knowledge (Lee, Scheepers, Lui, & Ngai, 2023).

The increasing organizational integration of AI reflects a broader technological evolution of how AI systems have developed and been applied over time. Early adoption of AI in organizations started with rule-based systems designed to enhance operational efficiency, followed by the integration of ML that enabled predictive analytics and decision support. Today, organizations have entered a new era defined by generative AI, which has the capability of not only enhancing existing processes but also reshaping entire business models (Simón, Revilla, & Saenz, 2024).

Generative AI can be implemented in organizations in different ways (Calvino et al., 2025). For example, it has the ability to fully automate certain parts of work, particularly clearly defined and limited tasks. However, it is unlikely that generative will automate entire professions and replace workers, since most professional roles involve many different types of activities that vary in how much human involvement they require. Therefore, generative AI is more likely to support workers in a collaborative way by supporting in different tasks, referred to as AI augmentation. For organizations and individuals, AI augmentation can lead to increased productivity, which in turn allows more time to be devoted to analytical and complex tasks (Calvino et al., 2025).

However, generative AI also introduces specific risks that are important to consider in organizational use. For example, one of the key limitations of generative AI is the phenomenon of hallucinations. Hallucinations occur when the model generates outputs that appear credible but are incorrect or unsupported. This problem stems from the underlying mechanisms of LLMs. They are designed to generate text by estimating the most likely token in a sequence using the patterns from their training data which can result in the generation of fabricated information (Ciubotaru, 2025). In addition, generative AI models have inherent biases. Biases are systematic errors in how the models process information, which can result in unequal treatment of

groups or reinforcement of inaccurate assumptions (Wei et al., 2025). Some generative AI tools also introduce what is referred to as the black box problem, which relates to the difficulty of tracing how the model arrived at its conclusions (Kosinski, 2024).

In addition to generative AI, the introduction of agentic AI into organizational environments presents opportunities. Although still at a relatively early stage of development, agentic AI is predicted to transform a wide range of organizational activities due to its increased autonomy and reasoning capabilities. Agentic AI is believed to have strong potential for innovation-related activities because of its ability to plan, synthesize ideas and execute tasks autonomously, which simplifies experimentation. Multi-agent systems will also enable analysis and processing of large volumes of data in a fraction of the time required by human teams. Additionally, the advanced reasoning capabilities of agentic AI help reduce some of the risks associated with generative AI, including hallucinations (Purdy, 2024).

Despite these promising opportunities, agentic AI also introduces several challenges for organizations. In addition to potentially reshaping workforce structures (Purdy, 2024), current agentic AI systems still struggle with long-term planning and face challenges with accountability when errors occur. Additionally, multi-agent collaboration sometimes leads to communication errors that introduce safety risks or biased decisions. Performance and tool integration is also one of the challenges with agentic AI, where complex and multi-step tasks are sometimes ineffective or delayed, especially when multiple agents interact with external tools (Bandi et al., 2025).

2.4.4 Organizational adoption of AI

It is clear that organizations recognize the potential of AI, however implementing it at scale is highly complex. Successful adoption requires appropriate technical infrastructures, relevant skills, and governance to address ethical and regulatory issues, including concerns about privacy, fairness and accountability (Daoud et al., 2025). Adoption of AI also poses strategic risks; without redesigning processes and decision models, there is a risk that organizations will not fully realize the benefits of their investments (Ransbotham et al., 2025).

The challenge of adopting AI in organizations is more related to organizational barriers than technical ones (Li et al., 2025). Li et al. (2025) identify three focus areas to enable successful adoption in organizations. The first area relates to *people's readiness*. The problems related to readiness are uncertainty, fear of replacement and what is described as the self-image problem. Uncertainty relates to opinions formed by employees because of a lack of knowledge, where some are enthusiastic about AI and dismiss its limits, and others see it as a trend that will pass. This problem is described as one of the explanations as to why many firms abandon AI initiatives even if there is observed potential. If risks such as regulatory risks exist, the uncertainty and lack of knowledge may result in intended users resisting adoption because of perceived risks in combination with a lack of knowledge. Li et al. (2025) suggest handling this issue by integrating AI governance into daily practices and

making it intuitive for employees to use AI responsibly. The fear of replacement problem relates to a reluctance to contribute because of the fear of being replaced by an AI system. In addition to a tendency of fault-finding where AI is held to higher standards than humans and demands of unrealistic levels of accuracy can halt implementation. To overcome this barrier, companies should allow employees to see and gain from the benefits AI creates and involve employees in the process. Examples to achieve this could include training royalties, productivity bonuses, or career guarantees that promise reskilling rather than firing if what an employee does becomes automated.

The self-image problem refers to a fear of losing status in relation to the adoption of AI. This fear shapes how AI is (or isn't) used. Li et al. (2025) note that some engineers use AI secretly to avoid damaging their reputation or being perceived as less competent or lazy. To address this, organizations must reposition AI use as a mark of forward-thinking, creating positive incentives for employees to openly adopt AI. At the same time, it is crucial to ensure that humans remain responsible for final decisions, which reassures employees that their judgment is not replaced, only strengthened.

The second focus area to enable successful adoption is described as *process*, which relates to redesigning workflows. A McKinsey & Company study also highlights workflow redesign as one of the most important success factors for achieving meaningful business impact. Their study showed that high performers succeed because they aim for large transformation, including redesign of workflows, rather than small efficiency gains (Singla et al., 2025). Li et al. (2025) describe that these process changes require transformation of workflows on multiple levels: individual, cross-functional and system-wide.

On the individual level, this involves redesigning workflows so that AI performs the tasks for which it is best suited, enabling humans to focus on higher-value tasks. At the cross-functional level, this involves strengthening connections between departments. Implementing AI creates the opportunity to improve communication across teams and enable faster decision-making, but this also requires changing how information flows within the organization. At the system-wide level, workflows must be redesigned to reflect how the entire system of teams, processes and individuals interacts. AI can improve individual tasks and strengthen connections between departments, but local improvements do not automatically translate into overall organizational gains. Without a system-wide perspective, companies risk shifting bottlenecks to new parts of the workflow rather than improving overall performance. To avoid this, AI must be coordinated across the system by mapping the entire workflow network to understand connections, and AI adoption needs to be synchronized. For example, if one team's speed increases because of automation, all connecting teams need matching capacity to capture the value (Li et al., 2025).

The last focus area is *politics*. To succeed in AI adoption, organizations need to adapt their governance structures and incentive mechanisms. Li et al. (2025) describe three common problems related to politics. The first is resource hoarding,

which happens when departments or individuals withstand sharing information or AI tools because of competitive instincts. This barrier is also confirmed in a study by Deloitte HKU lab, where one third of the executives in their study attributed the failures of AI initiatives to siloed departments where cross-collaboration was blocked (HKU Center for AI, Management and Organization, n.d.).

The second is described as disruption of traditional workplace hierarchies. Traditionally, workplace hierarchy is based on experience and team size. Junior employees with strong AI skills may outperform more experienced colleagues on certain tasks, creating tensions when greater contributions are not matched with recognition. Some organizations have addressed this by incorporating AI skills into competence models. AI also challenges the status associated with managing large teams. Managers whose influence or status depend on team size may be reluctant towards AI. To overcome this, organizations can redefine status and reward managers whose teams excel in AI (Li et al., 2025).

A third political factor is accountability. Since AI can track in detail, it can also pinpoint where mistakes occur. However, assigning responsibility too precisely can disturb organizational harmony. To effectively adopt AI, organizations need to understand when precision strengthens performance and when it triggers internal disputes (Li et al., 2025).

2.5 Defining capabilities of AI

In this study the artificial intelligence capabilities are conceptualized as functional abilities that reflect what AI systems do in practice rather than the underlying technologies. This functional perspective provides a basis for understanding AI's potential in relation to patent management activities. There is no universally established framework for categorizing AI capabilities within organizational processes. Existing frameworks differ in terminology, scope and level of abstraction. The National Institute of Standards and Technology (NIST) proposes the "AI Use Taxonomy", that aims to describe AI systems in terms of activities that contribute to outcomes, traditionally produced by humans and is not bound to underlying AI techniques. This approach enables analyzing AI systems' usability in new application areas (Theofanos et al., 2024).

The AI use taxonomy identifies 16 activities: content creation, content synthesis, decision making, detection, digital assistance, discovery, image analysis, information retrieval, monitoring, performance improvement, personalization, prediction, process automation, recommendation, robotic automation and vehicular automation.

Whereas the AI use taxonomy focuses on describing specific activities performed by AI systems, the European institute of innovation and technology (EIT) defines AI capabilities at a broader level by describing them as general problem domains that AI systems can address (EIT, n.d.). By using this categorization, eight capabilities are identified: computer vision, computer audition, computer linguistics, advanced

robotics and control, forecasting, discovery, planning and creation.

The AI use taxonomy and the capabilities identified by EIT(n.d.) largely overlap but have some structural differences (Theofanos et al., 2024; EIT, n.d.). Forecasting in the AI use taxonomy aligns with prediction as both describe the estimation of future outcomes. Similarly, computer vision and image analysis both relate to the ability to understand visual data. The ability of computer audition defined by EIT (n.d.) is included in content creation in the AI use taxonomy (Theofanos et al., 2024).

Discovery in both descriptions relates to recognizing something new, but the capability identified by EIT (n.d.) explicitly mentions the ability of analyzing data to identify patterns and relationships by using pattern recognition and clustering as underlying abilities. The capabilities of creation and computer linguistics are comparable to content creation and content synthesis in the AI use taxonomy since they all revolve around generating new artifacts such as text or images, and include summarization of text. While the AI use taxonomy does not specify text generation in creation, the ability to generate narratives is included, which is comparable to what is described as the ability to interact in human language, included in the description of computer linguistics defined by EIT (n.d.). Planning, defined by EIT (n.d.), is expressed by combinations of activities rather than a single category in the AI use taxonomy. For example, to conduct planning as defined by EIT (n.d.), activities such as decision-making, prediction, information retrieval, and recommendation, as well as abilities related to automation, could be applied (Theofanos et al., 2024; EIT, n.d.).

Although the capabilities proposed by EIT (n.d.) and the AI-assisted activities described by Theofanos et al. (2024) adopt different approaches, they are conceptually similar. Their combined perspectives serve as a baseline for identifying relevant AI capabilities in this study, allowing an assessment of AI's potential within the patent management process (Theofanos et al., 2024; EIT, n.d.).

To account for overlaps in content creation and content synthesis in the AI use taxonomy and computer linguistics and creation defined by EIT (n.d.), this study defines two distinct capabilities: content creation and summarization. They account for the abilities to generate images, text, and video, as well as to generate narratives, in addition to the ability to summarize. The ability to understand human language, which is included in computer linguistics in the definition by EIT (n.d.) and not specifically mentioned but implied in content creation in the AI use taxonomy, is in this study understood as an underlying ability enabling all the capabilities defined, except for image analysis (Theofanos et al., 2024; EIT, n.d.).

Similarly, to account for differences in descriptions of discovery, the capability is referred to in this study as pattern recognition and discovery to highlight that discovery also includes the identification of relationships, similarities and differences, as suggested by EIT (n.d.). To account for the differences in how the ability of planning is described, this study aligns with the AI use taxonomy and understands

the ability of planning and providing optimal solutions to problems as an ability combining multiple of the capabilities defined. This is also true for the capabilities of information retrieval, monitoring, recommendation and detection which are explicitly mentioned in the AI use taxonomy but are either not mentioned or implicitly included in the broader categories described by EIT(n.d.). To address these capabilities, this study adopts the definitions of the AI use taxonomy and extends the capability of detection to include search and detection, highlighting that the ability of detection builds on the act of searching (Theofanos et al., 2024; EIT, n.d.). In addition, the capability of decision-making, originating from the AI use taxonomy and not explicitly mentioned by EIT (n.d.), is extended to better reflect developments of AI. The definition is expanded beyond selecting a course of action from alternatives to the ability to autonomously execute decisions, reflecting the evolving role of agentic AI (Theofanos et al., 2024; EIT, n.d.).

However, not all of the capabilities listed in the two are included. Robotic automation from the AI use taxonomy and the corresponding advanced robotics and control capability identified by EIT(n.d) are excluded as they relate to the control of physical systems, which is not central to the patent management process. Similarly, vehicular automation is excluded because it describes the automation of physical transportation, which is not relevant to patent management activities. The capabilities of performance improvement, process automation, personalization and digital assistance described in the AI use taxonomy but not mentioned by EIT(n.d) are also excluded from this study. While these categories are relevant for describing AI use in practice, they are not considered as distinct or deep-level capabilities. Instead, they either represent outcomes of AI use, such as improving quality, or are not directly applicable to how patent management activities are conducted (Theofanos et al., 2024; EIT, n.d.).

The capabilities, together with how they are defined for the purposes of this study, are presented in Table 2.1

Table 2.1: AI Capabilities and definitions

Capability	Definition
Content creation	The ability to generate new content such as images, graphs, text, presentations, videos and perspectives.
Summarization	The ability to summarize information from one or multiple sources.
Image analysis	The ability to understand and interpret visual data.
Search & detection	The ability to process large datasets and search based on specific requirements to identify the existence of something.
Pattern recognition & discovery	The ability to process large datasets to identify existing or discover new relationships, clusters, similarities and differences.
Information retrieval	The ability to locate and extract information on defined subjects.
Monitoring	The ability to continuously track the state of an entity over time and deliver a response.
Forecasting	The ability to analyze historical and current data to predict the likelihood of future outcomes.
Recommendation	Suggest options or actions acting as support for decision-making.
Decision making	The ability to generate and assess alternative actions and autonomously or semi-autonomously execute decisions and achieve specified goals.

2.6 Patent management processes and the potential of AI

This section explores the potential applications of AI across the patent management process. The purpose is to both describe the core activities that are typically conducted within patent management and to conceptually map established AI capabilities to these activities. In doing so, the section provides a perspective on how AI could support existing patent management practices based on prior research. The processes investigated are based on the previously described framework by Agostini et al. (2022-a), although the framework is adapted to support the purposes of this study.

The process of patent generation defined by Agostini et al. (2022-a) included state-of-the-art analysis and geographical scope. For the purposes of this study, this process is expanded to include all activities occurring in relation to applying for

patents to allow for a complete analysis of all of the activities in relation to corporate patent work. The supporting dimensions are also excluded because this study focuses on the applications of AI within core processes relating to patent work.

It is important to note that many of the AI applications described are not established in practice and are based on emerging research and early developments. While the sources mentioned provide valuable insights into current advancements and potential use cases, they should be interpreted as emerging potential rather than validated solutions.

Patent generation

The act of filing patents is a complex and text-intensive process. The first step of generating a patent is the preparation of an Invention Disclosure Form (IDF), it is typically performed by the inventor and consists of generating a detailed written explanation of the idea (WIPO, 2016). Akhtar (2025) suggests that NLP presents many opportunities regarding implementing AI into the process of completing IDF's, where AI can assist in categorization of inventions, identification of technical features and provide suggestions on claims. The ability to compare new IDFs with existing patents and ongoing internal research to enable identification of overlaps between projects and avoiding duplication is also highlighted by Akhtar (2025).

According to World Intellectual Property Organization (WIPO, n.d.-c), the patent application itself, following from the IDF, generally consists of a request containing the title of the invention and formal information about the applicant and relevant dates, a detailed description of the invention, claims that distinguish the invention from prior art, and a summary of the invention. All of these activities relate to generating text and require both technical and legal wording. In some cases, the application also contains images, such as technical drawings or diagrams, that illustrate the invention's visual characteristics and complement the written description (WIPO, n.d.-c). AI systems are capable of performing tasks such as manipulating and analyzing text, generating text, images and other content (Stanford Institute for Human-Centered Artificial Intelligence, 2023). These capabilities are, from a theoretical perspective, closely aligned with the activities described above, where the activities of the tasks correspond to known AI capabilities.

One of the fundamental steps in the application process is to distinguish the invention from prior art (WIPO n.d.-c). To do so, a prior art search consisting of both a product search and a patent search, must be completed. The European patent office (n.d.-a) explains that to complete a prior art search, an extensive and systematic review of multiple sources, with the objective of ensuring that all relevant information has been identified is required. Traditionally, this process is performed by searches based on keywords and classifications across patent databases and other public sources. It requires examining large volumes of patent documents including an evaluation of the relevance of the documents in relation to the invention (EPO, n.d.-a). This process is time-consuming and depend on the manual work of the patent attorney, who must construct search queries consisting of word combi-

nations intended to find documents describing similar inventions to the one being investigated. In addition, searches often require multiple iterations of refinement to identify relevant documents. Moreover, a limitation of the searches is that they do not account for synonyms or related themes. As a result, relevant prior art may be overlooked, especially because abstract language is often used strategically to broaden protection scope (Helmets et al., 2019).

AI systems are capable of information retrieval (IR). Jurafsky and Martin (2026) defines IR as the process of identifying the most relevant documents from a collection, based on a set of terms expressed by a user. In addition, IR can be combined with LLMs trained on external sources of knowledge to allow both retrieval of relevant documents and generation of answers based on them. AI also holds the capability of text classification where information from an input is extracted based on a set of features and is classified into a predefined group (Jurafsky and Martin, 2026). Helmets et al. (2019) highlights the connection between these opportunities that AI offers and patent searches and notes that by using NLP and ML techniques there is strong potential in automating searches for prior art and overcoming some of the disadvantages associated with traditional ways of working. Bui (2025) adds to this by explaining that unlike keyword-based approaches AI models have the ability to understand the semantic meaning of text which enables the identification of relevant documents even when different terminology or language is used. Bui (2025) also describes the opportunity to include more sources than patent databases, such as technical reports or academic articles, which results in a more comprehensive search.

After an application is submitted, a formal examination and a technical assessment are conducted by a patent office (Patent- och registreringsverket, n.d.). The examination is done by a patent attorney. During this process, novelty and inventive step is evaluated and the result includes a conclusion on whether there is existing prior art that hinders patentability. If relevant documents are located, the applicant will be notified by an office action and search report. The applicant can respond to the office action and address the objections (Gómez Lagerlöf, 2024).

The ability to generate text has already been identified as a central component of the patent application process. This capability is clearly relevant to the task of generating text to answer office actions. However, beyond text generation, generative AI can also support the development of legal reasoning. When responding to office actions, you interpret the feedback from the examiner, analyze prior art in search reports, and formulate new arguments. Regalia (2024) suggests that AI can highlight new or less apparent legal strategies and arguments. This aligns with the demands of answering office actions, which require the ability to interpret complex information, refine descriptions, and generate new arguments based on additional information.

Freedom-to-operate(FTO)

According to Agostini et al. (2022-a) the act of conducting an FTO analysis relates

to assessing risks of infringement and litigation. In order to ensure that you have freedom to operate WIPO(2020-a) describes that firms conduct freedom to operate searches. The searches are complex and complicated as they require specialized expertise in patents and regulations that relate to court decisions (WIPO, 2020-a). The process of conducting an FTO search can generally be described as a three-stage process where you first gather information about the invention and its planned use. The second stage is to examine public patent databases and other literature to find relevant prior patents. The third stage is assessing the legal status and the claims of the patents identified in the search and deciding which rights might impact the invention (WIPO, 2020-b).

For the first stage, WIPO (2020-b) suggests that the outcome should be a summary describing the invention, its intended purpose, where it will be used, and the expected timing of its use. In this context, AI has potential to support the process, as generative AI is recognized for its ability to produce and summarize textual information. By assisting in structuring and refining descriptions of the invention, AI could contribute to answering these questions and producing the summary more efficiently.

Similar to prior art searches, FTO analysis relies on identifying relevant documents in patent databases and assessing their potential impact in relation to the investigated invention. In this context, as described previously, AI capabilities such as the ability to understand semantic meaning of words and phrases can be applied to improve the effectiveness of FTO searches. Its ability to detect relationships between technologies also enables identification of relevant patents beyond simple keyword search. In addition, its ability to conduct broader searches, including other public sources, allows for a more comprehensive search (Bui, 2025).

In terms of assessment of potential infringement risks, there is potential in applications of AI for predictive analytics. By analyzing market developments and competitor behavior, AI tools can assess potential infringement risks and by using historical data forecast future IP violations which could be beneficial to the third stage of an FTO (Bui, 2025).

Portfolio management

The patent portfolio of an organization is the collection of patents owned by the firm. Portfolio management involves determining the commercial value of patents, handling patents as a strategic assets and evaluation of the portfolio's position in relation to market dynamics and technological developments (EPO, n.d-b). Furthermore, portfolio management is an ongoing process that requires continuous review of the portfolio's value contribution. This includes assessing whether current and future key technologies are protected. Based on these assessments, strategic decisions are made regarding patent disposals (Agostini et al., 2022-a).

Akhtar (2025) describes that AI tools can be used for the purpose of aligning patents against organizational goals by analysis of the existing portfolio, R&D activities and

market position to assess the alignment of new inventions and the IP strategy of the organization. This opportunity is enabled by AI's ability to retrieve information from multiple sources such as company reports, market studies and industry insights. Based on these, AI can provide information or recommendations on whether proceeding with a patent aligns with strategic objectives.

Furthermore, Akthar (2025) emphasizes that AI can enable a more accurate and dynamic valuation of patents. This is because AI can analyze multiple sources of data simultaneously such as market trends, technological relevance and citation data and based on this identify patterns that may be difficult to detect using traditional methods. By including shifting market demands and emerging technologies it is possible to analyze how the value of patents might change over time. In relation to this, due to its advanced analytical capabilities, AI can help in assessing and recommending which patents to maintain or allow to lapse by including factors such as associated costs and the company's business objectives (Akthar, 2025). Together, this can enable organizations to manage patents more dynamically, ensuring that the portfolio is aligned with business priorities and market developments, thereby potentially increasing the economic value generated from the portfolio (Akthar, 2025). Akhtar (2025) further notes that AI systems can automate administrative tasks related to the patent portfolio, such as keeping track of due dates for maintenance fees, calculating fees and prepare payments.

Enforcement & exploitation

Enforcement is described as the act of ensuring that others comply with your intellectual property rights (European Commission IPR Helpdesk, 2014). Agostini et al. (2019) describe enforcement & exploitation as one of the core processes of patent management. This process involves activities that relate to how companies generate value from their patents. It can be in terms of licensing or selling, through which firms obtain compensation, either in monetary or non-monetary form for their patents (Agostini et al., 2022-a). The process also includes proactive licensing which refers to searching for other organizations with a potential interest in applying the firms patented technologies, or signaling effects which refers to voluntary disclosure of information about patents and technologies in order to signal capabilities to external parties (Agostini et al., 2022-b), it also involves monitoring and identifying infringements and taking action if infringement is identified (Agostini et al., 2019).

According to the European Commission IPR Helpdesk (2014), in most IP infringement proceedings accused parties challenge the validity of the IP right and file counterclaims, making it essential to have the ability of analyzing and proving validity of IP rights. Ikoma & Mitamura (2025) suggest that LLMs open up new opportunities to compare claims with cited prior art to establish validity. They explain that LLMs can assist both examiners and applicants in patent evaluations by comparing the content of claims and prior art. This capability is relevant both during the previously mentioned activity of answering office actions, where it can support the refinement of claims, and during the enforcement and exploitation process. AI could assist in evaluating patent validity, both when claiming infringement against

others and when assessing legal risk and determining appropriate responses if the organization is accused of infringing.

Organizations should also actively monitor market activity. This includes regularly reviewing intellectual property databases to track competitors and identify potential infringements (European Commission IPR Helpdesk, 2014). Bui (2025) highlights that AI capabilities in both text and image analysis can assist these activities. Through NLP, AI systems can analyze product descriptions and technical specifications, while computer vision techniques enable the examination of visual elements such as product images. This enables reduction in the manual monitoring of competitor products and patent activity. Thus, AI systems could be used to identify potentially infringing products, which can then be further evaluated by enforcement teams, reducing the time and effort required for continuous search (Akhtar, 2025).

If a potential infringement is identified, the enforcement strategy varies depending on the situation. In some cases, relatively simple measures such as issuing a cease-and-desist letter can be sufficient. In this letter the infringing party is notified that legal action may be initiated if the infringing activities do not stop. In more complex situations organizations must make a strategic decision on how to proceed, typically choosing between seeking a settlement or initiating litigation (European commission IPR Helpdesk, 2014). In this decision-making process, several options are available. Disputes can be addressed through private enforcement actions, public prosecution, or alternative dispute resolution (ADR). ADR provides a way to resolve conflicts without expensive court proceedings and includes methods such as negotiation, mediation, and arbitration. Negotiation involves the parties reaching an agreement, with legal representatives assisting in drafting the terms. Mediation is a process in which an independent third party helps the parties find a mutually acceptable solution. Arbitration involves submitting the dispute to a tribunal, which concludes a binding decision (European commission IPR Helpdesk, 2014).

Regardless of which strategy is chosen, these processes are dependent on legal reasoning and the formulation of arguments. Regalia (2024) suggests that generative AI has large potential in innovative problem-solving. AI can examine previous case outcomes and based on those recommend new strategies or identify less apparent arguments to a problem. This could be beneficial in situations such as negotiations, where there is a need to interpret the opposing party's arguments and evaluate one's own position. In such contexts, AI could support the development of stronger strategies by offering alternative perspectives, identifying potential arguments, and providing inspiration for new ways of approaching the case.

If out of court solutions are insufficient, organizations may initiate formal legal proceedings. Prior to litigation, an assessment of the likelihood of success is required. This involves evaluating several factors, including the basis for bringing a claim, the validity of the owned IP right, potential defenses available to the opposing party and the associated costs (European commission IPR Helpdesk, 2014). There is alignment with the need to assess likelihood and the abilities of AI. For example, Juranek and Otneim (2024) highlight the connection of AI's ability of predictive analytics and

assessment of likelihood of a patent being subject in a court case. They describe that ML techniques such as classification models can be used for prediction on whether a patent will get litigated. In addition, Bui (2025) describes how predictive analytics can be used to evaluate infringement risks. Predictive models can analyze market trends and actions of competitors to predict risks early, which leads to the opportunity to take preventative action.

While it has been described how AI can assist in infringement detection. AI could also support detection of potential licensing opportunities, thereby contributing to proactive licensing activities. By assessing market trends, technological advancements and patent landscapes, AI can extract companies that may be interested in the organization's technology. Furthermore, by analyzing historical data from prior licensing agreements in combination with current market conditions and the specific needs of potential licensees, AI could support in formulating licensing offers and pricing strategies (Akhtar, 2025).

Intelligence

Agostini et al. (2022-a) explains that the process of intelligence involves activities aimed at understanding and mapping competitor and technology landscapes. In this process, analyzes of current trends and future developments are conducted, for example, by monitoring existing and emerging patent grants, publications, and activities aimed at mapping the patent landscape. To support this work, different types of reports and analyses of the external environment are conducted. One example is patent watching, which aims to monitor new patents and applications, to evaluate if these are of interest to the firm. It could be monitoring applications that might be a risk to the patent rights or it could be monitoring of competitor organizations to identify which new patents arrive in relation to technologies of high interest (WIPO, 2015). There are automated alert systems for this, such as the European Patent Office's register alert and email notification services, used for monitoring changes in selected patent applications (EPO, 2024). AI could have potential to enhance the process of monitoring by improving the analysis. It has already been concluded that AI can process large amounts of data and can compare text beyond keyword matching by using semantic search and NLP, by doing so AI can deliver more relevant search results by understanding context and relationships (Google cloud, 2026). Another example is patent maps, which are visual representations of patent data. These maps are often structured around one specific attribute of the dataset, for example classification of documents based on content similarity. Generative AI has high potential in generation of visual material based on structured data, and thereby has potential to assist in the visual representations (Ye et al., 2024).

One of the most common reports produced in landscape analyses is a patent landscape report which provides an overview of patenting activity within a specific technological field and serve as an input to strategic decision-making. These reports are designed to answer key questions regarding which technologies are covered, which organizations hold relevant patents, and in which geographical regions these patents

are held. The process of producing a patent landscape report is both time-consuming and expensive, as it requires the analysis of large volumes of raw data and the interpretation of results (WIPO, 2015).

When examining the details of producing reports such as patent landscapes it becomes evident that there is clear theoretical alignment with the capabilities of AI and the activities performed. A fundamental activity is data preparation, including data cleanup and grouping, which refers to correcting inconsistencies and combining synonymous terms in patent data, as well as grouping related concepts such as technology categories or classification codes into higher level categories that can be analyzed. World Intellectual Property Organization (2015) notes that the task of preparing patent data for analytics is often more time-consuming than the analysis itself. There are many different applications of AI that hold the ability to detect synonyms or similarity of words (Jurafsky & Martin, 2026). Furthermore, Panwar (2024) suggests that artificial intelligence has the potential to reduce the time and resources required for data preparation by automating the detection and correction of inconsistencies through ML and NLP techniques.

Following data preparation, list generation is a common activity. In this activity, elements within a technological field are compared and quantified, for example to identify the most active organizations or inventors. The activity involves both sorting data and counting entities and the lists are often visualized with bar charts or histograms to enable comparison. Traditionally, these tasks are performed using software such as Microsoft Excel (WIPO, 2015). Akhtar (2025) describes that AI can be used to identify key competitors in a technological field and produce forecasts of future market sizes. Apart from supporting in extracting relevant insights for these tasks, generative AI can also assist in generating visualizations such as the mentioned graphs or histograms.

An additional activity is clustering and classification. Clustering relates to the grouping of patent documents based on similarity, and classification relates to assigning documents to predefined categories. These activities are aligned with the capabilities of AI. This is highlighted by the fact that the World Intellectual Property Organization (WIPO) describes that these tasks are often performed using algorithms based on neural networks and support vector machines, in their report from 2015. This indicates that, unlike other areas of the patent management process, the potential of AI is not only theoretical, but already represents an established practice (WIPO, 2015).

In addition, geographic representations are often included, where patent data is mapped to specific locations based on inventors or applicants. This allows for the analysis of the distribution of innovations and supports the identification of potential collaboration opportunities (WIPO, 2015). This process could also be enhanced by integrating generative AI models, both for the analysis and for the search of locations, but also in the generation of visual representations (Ye et al., 2024).

Finally, patent landscape analysis often involves some type of semantic analysis

where key problems and corresponding solutions within patent documents are identified, thereby supporting a deeper understanding of technological developments and challenges within a field (WIPO, 2015). Bui (2025) highlights that there is high potential in the use of AI-driven analytics to provide insights into trends and evaluation of competitor portfolios, suggesting that AI can support the analytical work conducted by both innovators and attorneys in terms of understanding landscapes. He further explains that AI enables identification of areas that are unexplored. Akhtar (2025) adds to this by describing that AI can assist in predictions of future market sizes which would be an asset in understanding future directions.

2.7 Theoretical mapping of AI capabilities in relation to patent management

Table 2.2 illustrate how the previously identified AI capabilities can be mapped in the context of the patent management process described in 2.6. Each capability has been mapped according to the descriptions of capabilities in table 2.1. The mapping highlights that many capabilities are applicable across multiple of the processes rather than isolated activities. The table reflects a theoretical mapping between AI capabilities and the activities of patent management and serves as a foundation for the subsequent analysis, where these theoretically identified applications are examined in relation to how AI is currently applied in practice.

Table 2.2: AI capabilities in patent management

Potential AI application within patent management	Processes	AI Capabilities	Notes
Draft and structure innovation disclosures	Patent filing	Content creation	Content creation enables generation of text for innovation disclosures
Identify technical features and suggest claims	Patent filing	Information retrieval, recommendation	Information retrieval enables identification of key features, recommendation enables suggestions on claims
Identify overlaps between new ideas and existing patents	Patent filing	Search & detection, pattern recognition & discovery	Search & detection enables finding similar patents, pattern recognition & discovery enables assessing similarities and differences of identified patents
Generate text, images and drawings for patent drafts	Patent filing	Content creation	Content creation enables generation of text, images and drawings
Generate invention summaries	Patent filing, FTO	Summarization	
Search patent databases and non-patent literature using semantic matching beyond exact keywords	Patent filing, FTO, Intelligence	Search & detection, information retrieval	Search & detection and information retrieval enable identification of relevant matches across large volumes of data to support in prior art-, FTO- and intelligence searches
Analyze office actions and generate new arguments and texts for responses	Patent filing	Content creation, recommendation	Content creation and recommendation enable interpretations of office actions and generation of arguments
Analyze the alignment of inventions and the IP strategy of the organization	Portfolio management	Pattern recognition & discovery	Pattern recognition & discovery enables discovering similarities and differences between invention descriptions and IP strategy
Automation of administrative tasks	Portfolio management	Monitoring	Monitoring enables AI to keep track of clearly defined repetitive tasks such as handling of patent fees
Accurate and dynamic valuation of patent value and analysis of how it may evolve over time	Portfolio management	Pattern recognition & discovery, information retrieval	Information retrieval and pattern recognition & discovery enable analysis of multiple sources of data simultaneously and identification of hard to detect patterns and connections
Recommendation regarding which patents to keep or allow to lapse	Portfolio management	Pattern recognition & discovery, information retrieval, recommendation	Information retrieval enables retrieval of relevant data, pattern recognition & discovery enables analysis of multiple sources of data simultaneously to identify relevant patterns and connections, recommendation enables suggestions on action based on the insights
Support in validity analysis	Enforcement & exploitation	Pattern recognition & discovery	Pattern recognition & discovery enables identification of similarities and differences between claims and prior art claims

Potential AI application within patent management	Processes	AI Capabilities	Notes
Active searching and monitoring for potential infringements	Enforcement & exploitation	Monitoring, search & detection, image analysis	Monitoring and search & detection enable tracking of competitor activities and search for potentially infringing solutions, image analysis expands the search to detect infringing products in images or videos
Support in negotiations	Enforcement & exploitation	Content creation, recommendation	Content creation and recommendation can enable stronger negotiation strategies by generating alternative arguments and perspectives and recommending the best alternative
Assessing the likelihood of litigation	Enforcement & exploitation	Forecasting	Forecasting enables analysis of historical data to predict the likelihood of litigation for a specific case
Analyze market trends and actions of competitors to predict risk of infringement	Enforcement & exploitation	Forecasting	Forecasting enables analysis of market trends and competitor's actions to predict their likely future behavior, thereby estimating the risk of potential infringement of your IP rights
Finding potential licensing opportunities	Enforcement & exploitation	Search & detection, pattern recognition & discovery, recommendation	Search & detection and pattern recognition & discovery enable finding opportunities with similar or adjacent technologies and strategic fit, uncovering new relationships through similarity patterns and recommendation enables prioritizing the best alternatives
Support in the formulation of new licensing offers and pricing strategies	Enforcement & exploitation	Recommendation, pattern recognition & discovery	Pattern recognition & discovery enables analysis of historical licensing data, recommendation enables suggestion on best alternative based on the analysis
Generation of visualizations, including charts, graphs or geographical mappings of search results	Intelligence	Content creation	Content creation enables the creation of visualizations and mappings from data
Preparing data for analysis by grouping and categorizing data into relevant categories	Intelligence	Pattern recognition & discovery	Pattern recognition & discovery enables identification of similarities and differences across large volumes of data, supporting grouping and categorizing into relevant categories
Predictions of future market sizes	Intelligence	Forecasting	Forecasting enables analysis of historical market sizes over time and relevant drivers and industry indicators to predict the likely market size at a specific point in time
Identification of unexplored areas and opportunities	Intelligence	Information retrieval, pattern recognition & discovery	Information retrieval enables retrieval of data on technologies, market trends and competitor portfolios, pattern recognition & discovery enables identification of opportunity areas, for example where competitors hold few patents

3

Methods

This study was conducted as an interview-based investigation of AI in corporate patent management. Empirically, the work combined an in-depth analysis at a collaborating company with external perspectives from corporate IP departments, IP firms, tool providers, public authority and industry experts. To create a robust basis for comparing AI potential with practice, a conceptual map of patent management processes and activities were developed, grounded primarily in empirical insights from the collaborating company and iteratively refined through input from external respondents. This map provided a common structure for locating in which tasks AI was and could be applied in patent management.

In parallel, the study assessed AI potential from two complementary angles. First, potential applications were identified through existing literature. Second, AI tools identified in interviews were reviewed to contextualize the practical availability of AI-enabled support across the identified process activities. Together, literature-based potential and the tool-availability identification were used to define the theoretical AI potential as indicated by both prior research and the current tool landscape. This theoretical potential was then compared with empirical evidence of perceived potential value and current use. Finally, barriers of adoption were identified.

3.1 Research approach

This study applies a qualitative and abductive research approach. The potential of AI in patent management is a concept that is both context-dependent and is currently evolving. Although research has begun to describe AI use in IP management, there is no comprehensive explanation of how the value of AI is perceived or realized by different actors. This gap in literature supports the adoption of a qualitative, abductive research design.

Bell et al. (2019) defines qualitative research as the collection and analysis of non-numerical data such as written texts, spoken words and visual material. Qualitative research is aimed at developing an understanding of social phenomena from the perspective of participants involved. It emphasizes rich descriptions of people, interactions and behavior.

A qualitative approach was chosen because it allows for examination of how individuals and groups within the intellectual property landscape perceive adoption of AI in patent management. In addition, a qualitative approach allows for subjective meanings and situational differences. This approach enables the study to capture diverse perspectives and account for how context and individual interpretations shape views on AI's potential and implications for adoption (Bell et al., 2019).

In addition, this study followed an abductive logic to allow theoretical understanding to be refined in response to emerging empirical insights. Bell et al. (2019) describes the abductive approach as a model for addressing problems that existing theory cannot fully explain. The data collection was conducted iteratively and rather than applying a fixed set of concepts from the outset, empirical insights from the early interviews were used to refine the focus of subsequent data collection. In addition, the empirical data was used to refine the theoretical understanding of the study. Throughout the study, emerging findings about AI use, perceived potential and adoption barriers were continuously compared with relevant literature, and the theoretical framework was adjusted when new patterns or insights appeared.

3.2 Research design

This study was designed with special emphasis on a single company. The aim of the research is to provide a detailed analysis of the patent management process within organizations. To achieve this purpose, the study was conducted in collaboration with a multinational industrial company operating within the automotive sector. By focusing on a single IP management function and its processes in relation to patent management, the study enables deep analysis of workflows, sub-processes and their interrelations. The detailed analysis was complemented by insights from IP departments across multiple organizations, firms specializing in IP strategy, and individuals with expert knowledge of IP management. Additional perspectives were also obtained from firms developing AI-enabled tools for patent management processes. This design allowed for an analysis that is both detailed enough to represent the realities of organizational practice and broad enough to capture the wider trends and benchmark findings across the broader patent-management landscape.

The collaborating company supported the research by providing access to interview participants, internal documents, and opportunities to observe the patent management processes of the IP department. The collaborating company represents an appropriate context for this study, as it has maintained high investments in R&D during recent years, underscoring its commitment to innovation. In addition, there is strong organizational emphasis on continuous innovation and integration of emerging technologies. Its strategic emphasis on innovation results in a context where patent management is highly relevant. Alongside this, the collaborating company provided access to observations of patent-management work in practice, which enabled a detailed investigation of how activities are carried out in corporate settings.

3.3 Data collection

The main data collection method of this study was qualitative interviews. This approach allowed for a flexible structure in which the researchers could depart from the interview schedule and ask follow-up questions when relevant.

There are two main types of qualitative interviews: unstructured and semi-structured (Bell et al., 2022). Both interview types were applied in this study. Unstructured interviews resemble conversations and the discussion is shaped by what seems relevant and interesting in the moment, although some broad topics can be decided in advance. Semi-structured interviews are guided by an interview guide with pre-defined topics, while still allowing for follow-up questions and flexibility in order of questions (Bell et al., 2022).

In total, 32 interviews were conducted. The data collection started with four unstructured exploratory interviews at the collaborating company to establish an initial understanding of patent management work and current AI-related practices. Building on these insights, interview guides for the semi-structured interviews were developed and tailored to each respondent group.

28 semi-structured interviews were conducted. These included nine respondents from the collaborating company, eight respondents from external companies with IP departments, six respondents from IP firms, two tool providers, one consultant firm, one domain expert and one public authority. Interviews were typically 45-60 minutes long and were conducted in person or via Microsoft Teams depending on the respondent's availability and location. Separate interview guides were used for each respondent group to reflect differences in perspective and scope. In accordance with Bell et al.'s (2022) description, all participants in their respective groups were asked all the listed questions and in approximately the same way.

External companies were included to provide benchmarks for current practices and activities, in relation to insights from the collaborating company. They also contributed to identifying which AI applications were perceived as relevant across organizational contexts. Consultancy firms were included because they support multiple organizations in IP and technology-related change initiatives and could contribute with a broader perspective on current trends, opportunities of AI and adoption patterns. IP firms were included because they conduct patent work across multiple clients and processes, offering a cross-company view of what is already being used in practice and where there exists untapped potential of AI. Tool providers were included to contextualize what AI-enabled functionalities are currently offered and insights about current implementation considerations. An industry expert was included to provide an independent perspective on patent-management practices. Finally, a public authority was included to add the perspective of institutional requirements, standards and constraints that affect AI implementation in patent management.

To select interview participants, a purposive sampling approach was chosen. Pur-

positive sampling is common in qualitative research and involves sampling in relation to the goal of the study so that the analysis can meaningfully answer the research questions (Bell et al., 2022). In addition to this approach, a snowball effect occurred where respondents suggested additional relevant contacts.

Participants from the collaborating company were selected based on their role and expertise. Selections were made with the aim of fair representation of the different roles and processes within the collaborating company’s IP function, ensuring that all major stages of the patent management process were covered. External companies were selected based on the criteria that they had an IP function and were R&D intensive. Within these companies, interviews were conducted with individuals with different roles within the IP function, mainly patent managers and patent attorneys. Consultancy firms were selected based on whether they specialized in IP work and were knowledgeable about emerging technologies. Tool providers were selected based on the requirement of offering AI-enabled tools specifically designed for patent-management work. The domain expert was selected based on demonstrated expertise and academic achievements in IP work. The public authority respondents were selected based on knowledge in patent examination and AI’s influence on the broader patent system.

When respondents consented, interviews were audio recorded, in interviews where recordings were not possible, due to confidentiality concerns related to the topic being IP work, both authors took structured notes during the interview. Recorded interviews were transcribed using Microsoft Word transcription.

To complement interview data, selected internal documents from the collaborating company were reviewed, including strategy documents, training material and role descriptions. According to Bell et al. (2022), documentary data have the advantage of being non-reactive as they are not produced specifically for the purposes of the study. However, the reliability of internal documents should be considered, as they may reflect selective reporting and present the organization in a favorable way. In alignment with this, the documents were used as contextual input and as support for mapping the processes of patent management rather than as a standalone dataset.

Respondent number	Respondent group	Date	Format
1	Collaborating company	2026-02-20	In-person
2	Collaborating company	2026-02-20	In-person
3	Collaborating company	2026-02-20	In-person
4	External IP department	2026-02-23	Microsoft Teams
5	External IP department	2026-02-23	Microsoft Teams
6	External IP department	2026-02-24	In-person
7	External IP department	2026-02-25	In-person
8	External IP department	2026-02-25	Microsoft Teams
9	IP firm	2026-02-25	Microsoft Teams
10	External IP department	2026-02-25	In-person
11	IP firm	2026-03-02	Microsoft Teams
12	Collaborating company	2026-03-03	In-person
13	Tool provider	2026-03-03	Microsoft Teams
14	Area expert	2026-03-05	Microsoft Teams
15	Consultant	2026-03-05	Microsoft Teams
16	Public authority	2026-03-05	Microsoft Teams
17	External IP department	2026-03-06	In-person
18	External IP department	2026-03-06	In-person
19	Collaborating company	2026-03-09	In-person
20	Collaborating company	2026-03-10	In-person
21	IP firm	2026-03-10	Microsoft Teams
22	Collaborating company	2026-03-11	In-person
23	Collaborating company	2026-03-11	Microsoft Teams
24	Collaborating company	2026-03-11	In-person
25	Collaborating company	2026-03-12	In-person
26	Collaborating company	2026-03-13	In-person
27	Tool provider	2026-03-16	Microsoft Teams
28	IP firm	2026-03-18	Microsoft Teams
29	IP firm	2026-03-25	In-person
30	IP firm	2026-03-25	In-person
31	Collaborating company	2026-03-26	In-person
32	Collaborating company	2026-04-10	In-person

Table 3.1: Overview of conducted interviews

3.4 Data analysis

The data analysis was conducted iteratively through three interconnected analytical parts that were developed in parallel and refined as new empirical insights emerged.

Mapping the patent management process

A conceptual map was developed to represent common patent management processes and activities. The mapping was primarily based on the insights from the interviews with the collaborating company and internal documents. The mapping was refined iteratively throughout the data collection period. During the devel-

opment process, the map was shown to internal respondents to validate that key activities were captured. To avoid the process map reflecting only one firm's process, it was also discussed with external respondents and adjusted based on their feedback to better represent a general process typically conducted by organizations with IP departments. The resulting process map provided the baseline for addressing RQ1 and RQ2 by defining the set of activities against which AI applications could be analyzed.

Mapping AI tool availability to process steps

AI-supported patent management tools identified during interviews were reviewed to contextualize tool availability across the mapped process steps. For each tool, descriptions were gathered from the respective vendor websites. Tool features were then mapped to the relevant process steps. This overview was not intended as an extensive market analysis, but as a structured way to understand whether availability of potential applications of AI exists across the processes included in this study, and to complement the theoretical potential of AI. Together, the two inputs were used to answer RQ1.

Thematic synthesis of realized AI potential and adoption barriers

The empirical data was synthesized using thematic analysis to examine how the theoretical potential identified was realized in practice to answer RQ2, and to identify barriers that help explain gaps between indicated potential and realized use to answer RQ3. Thematic analysis is a flexible approach to identifying patterns and structure data into themes. A theme can be understood as a meaningful pattern in the data that is relevant to the research questions (Bell et al., 2022). A common way to identify themes is to look for repetitions and recurring ideas within and across sources (Bell et al., 2022), which in this study were interview transcripts and internal documents. For RQ2, empirical insights about AI use and perceived potential among respondents were organized into themes defined by the activities in the process map, enabling comparison of where AI was used and where implementation remained unrealized. For RQ3 recurring concerns were coded across interviews and structured into themes constituting barriers limiting adoption. Coding and categorization were initially performed separately by both authors and then cross-checked and combined through joint comparison. Microsoft Word was used to categorize findings and citations in relation to themes.

3.5 Ethical considerations

This study followed the ethical guidelines established for qualitative research with human participants. This means that the well-being of the participants was prioritized over the research questions (Mirza et al., 2023). One ethical consideration of this study relates to the principle of relationships with participants and conflict of interest (Mirza et al., 2023). Because part of the study involves conducting interviews within the collaborating company, some participants were individuals the

researchers had previously interacted with. This is due to the thesis work being conducted on-site at the collaborating company, where the researchers participated in department meetings or other workplace activities. To mitigate any risks in relation to prior familiarity, researchers maintained an unbiased stance and upheld a formal relationship with all participants throughout the data collection process.

As part of the data collection, informed consent of every participant was established. Prior to the interviews, participants were informed about the aim of the study, confidentiality and anonymity. Participation was voluntary and respondents could choose whether the interview should be recorded or not. Due to the sensitivity of patent and IP-related information, particular attention was paid to confidentiality, including the use of note-based documentation when recording was not possible. All material in relation to interviews (recordings, transcripts and notes) was stored secure locations in alignment with the established guidelines of Chalmers university of technology (2024). In accordance with the General Data Protection Regulation (GDPR), the personal data collected was gathered only for the purpose of this study. In addition, participants were permitted to examine any data collected from their participation, if requested, and had the option to withdraw from participation at any time (Mirza et al., 2023).

4

Empirical findings

The following section presents the findings from the data collection. First, an overview and explanation of common processes and activities within patent management is presented and described. Second, an overview of AI supported tools identified in the interview material is presented in relation to the patent management processes. Third, the interview material related to the use of AI in the patent management processes is presented and analyzed in relation to theoretical potential and identified tools. Finally, the identified organizational barriers to adoption are presented and analyzed.

4.1 Identified patent management processes

The aim of this section is to provide an overview of common processes and activities associated with patent management in organizations. The visualization serves as a foundation for the subsequent analysis in this thesis, where the use and potential of AI is examined across the different patent-related processes. It is also intended to provide an understanding of how patent work can be organized within organizations and which activities are commonly involved. The identified processes and activities are based on the existing patent management framework and activities presented in chapter 2.2 and chapter 2.6, but the mapping is primarily grounded in empirical data. It should be noted that patent management activities may vary between organizations, therefore it should not be seen as a representation every organization but as a synthesis of commonly occurring processes and activities within the scope of this thesis.

The patent management processes are presented in Figure 4.1. Dark blue ovals represent patent management processes, while lighter blue rectangles indicate activities within each process. Grey rectangles represent sub-activities. Activities outlined with a blue border show activities that are not always performed but occurs occasionally within the process. With exception for the patent filing process, the activities within the processes are not carried out in a specific order and may be carried out in parallel or iteratively. However, within the patent filing process, the order of activities follows the arrows in the picture. Each process and activity are explained below.




Figure 4.1: Conceptual overview of patent management processes

Patent filing

The patent filing process is presented as a series of activities that are carried out in order. The process begins when the inventor completes an Invention Disclosure Form (IDF) and submits it to the company’s IP function. In this study, this step is called *invention identification*. When the IDF has been received, the IP function performs the next step, which is a *preliminary assessment* of the IDF. This assessment typically includes *specifying the core of the invention*, *analyzing detectability*, conducting a *business assessment* and in some cases an *initial state of the art analysis*. Specifying the core of the invention means identifying the key technical idea of the invention, deciding which parts should be protected and included in the patent and separating the key features from supporting ones. Furthermore, analysis of detectability involves assessing how easy or difficult it would be to detect potential infringements, as this can influence whether it is worth filing for a patent or not. In this activity a business assessment is also carried out, with the aim to ensure that the invention are relevant to the overall business strategy. Additionally, an initial state of the art analysis may be carried out. The purpose is to get an overview of the technical area of the invention and what is already known in the field, which in turn can influence how the IP function decides to proceed. Based on the preliminary assessment, it is decided whether to continue the process with the invention or not.

The next step in the patent filing process is the *patentability assessment*, which aims to confirm that the invention meets the requirements for a patent to be granted. In this conceptual overview, the patentability assessment includes a *prior art search*,

analysis of non-obviousness and *confirmation of industrial applicability*. The prior art search involves searching for earlier publications, patents or other disclosures that could potentially challenge the patentability of the invention in terms of novelty. Non-obviousness means assessing whether the invention would have been obvious to a person skilled in the art. For example, it can involve identifying the closest prior art and assess whether a skilled person in the art could easily come up with the invention. If yes, the invention may be considered as obvious and thereby not patentable. Furthermore, it is also confirmed that the invention is industrially applicable. The subsequent activity is *drafting*, which is the task of writing the patent application. This involves describing the invention and defining what should be protected, it may also include images or drawings. Before filing a patent application, the geographical scope is determined, in other words, in which countries patent protection should be applied for. The final step within the patent filing process is *prosecution* which involves handling the communications, also called office actions, from the patent office after an application has been examined. It often involves responding to office actions by presenting arguments against the objections and, when necessary, amending the claims to address those objections.

Intelligence

The intelligence process involves completing analyzes and reports that can act as a basis for better informed decisions related to business, IP strategy and R&D. It can include collecting and synthesizing information on trends and outlooks, actors relative portfolio strength and patent landscapes. The activity of *analyzing trends & outlooks* relates to identification of trends on for example top patent holders, filing trends or new entrants in relation to a technical domain. *Actor analysis* involve investigating whether a certain actor, such as a competitor or a supplier, holds relevant patent rights, it can also include insights on other actors' technical solutions. The activity *analysis of relative strength* involves visualizing the relative strength of patent portfolios between relevant actors. It can for example include comparisons of size, filing intensity or geographical coverage of portfolio. *Patent landscape* includes generating a broader overview of a technical field. It can for example include descriptions of covered technologies, the organizations that hold patent rights or in which regions patents relating to this technological field are held.

Portfolio management

Portfolio management involves the overall management of the firms' patents, including monitoring, evaluating and optimizing the portfolio to maximize value and ensure alignment with both IP and business strategy. In this study, the process is broken down into several smaller activities.

One activity is *implementation, technology & product status assessment*. On one hand, it involves reviewing whether and how patented inventions are implemented within products or services, or when they are planned to be, in order to understand which patents are actively used. On the other hand, it involves technology and product status assessment, which refers to reviewing ongoing R&D projects

or roadmaps in order to assess if relevant coverage for upcoming technologies exist. A further activity is *competitive strength analysis*. This includes assessing the company's portfolio compared to the portfolios of competitors, thus evaluating the relative competitive advantage. This analysis often relies on material from intelligence reports in which portfolio data of relevant competitors have been collected and synthesized. It can also include citation analysis to examine how often patents in the portfolio are cited by later patents. This can indicate the value of the patent, both in terms of its technical influence in the field and its potential blocking power.

Furthermore, the activity *strategic value mapping* refers to the evaluation of patent portfolio in relation to business strategy. It includes ensuring that the portfolio is relevant to the overall business strategy. *Business value evaluation* refers to the activity of evaluating the portfolio in terms of business value. This can for example include analyzing if the portfolio generates exclusivity, generates direct income, has the potential to prevent others from establishing exclusivity, or has the potential to enhance the company's negotiation power or collaboration attractiveness. The activity *portfolio visualization & mapping* involves visualizing or mapping the portfolio based on different criteria, for example how patents relate to each other, to products, projects or techniques. It enables an overview and thereby understanding of the portfolio, making it possible to identify overlaps as well as gaps.

Additionally, the activity *identify areas for patent creation* is also included in the conceptual overview of this study. It refers to the activity of determining areas where the company seeks to file additional patents and expand its portfolio. These can be both areas where the company aims to strengthen its strategic position, as well as areas where sufficient protection is currently missing. In other words, it can relate to both opportunities and gaps in the existing portfolio. The identification can be informed by intelligence analyzes on trends and outlooks. Another activity is *docketing* which includes tasks related to tracking and organizing patent information. For example, tracking deadlines is an important part, which can include due dates on maintenance fees and office action responses. This is the administrative part of the portfolio management. Additionally, *pruning* includes evaluating patents based on economic value and related costs and based on this taking decisions on removing patents that do no longer provide value.

Enforcement & exploitation

In this study, the activities *infringement search*, *infringement analysis*, *monetization*, *legal action & commercial negotiation* and *validity analysis* are included in the enforcement & exploitation process.

The infringement search activity relates to searching for potential infringement products that may violate the company's granted patents. This can for example be done by monitoring or searching online sources, examining competitors' products or observing products at fairs. Depending on the company strategy, it can be carried out continuously or occasionally. Infringement analysis refers to the assessment of the potential infringement by comparing the patent claims with the potentially infring-

ing product or solution. A common method that can be used for this purpose is a claim chart, where the claim is broken down and mapped in a structured way to product features. A further included activity is legal action and commercial negotiation. This activity is typically executed when the company is either accused of infringing third-party patents or believes that a third party is infringing its patents. The action based on this can either be legal action or commercial negotiation, but in either case, it typically includes tasks such as preparing cases and supporting arguments or negotiating and drafting agreements. The monetization activity refers to generation of direct economic income through for example finding licensing opportunities or selling patents.

Another included activity is validity analysis, which involves assessing a patent in terms of its validity, with the aim of identifying potential ground for invalidation. This includes analyzing the patent in terms of novelty, non-obviousness and industrial applicability, and typically whether the invention is clearly described. The basis for the analysis is often a prior art search and an analysis of claims and description. A validity analysis can be performed in different situations. For example, if a third-party accuses the company of infringement, a validity analysis can be performed to evaluate the litigation risk and to develop invalidity arguments for potential negotiations. Likewise, a company can assess its own patent to ensure that it's enforceable before taking legal action against a third party. Additionally, a validity analysis can also be conducted on a competitor's patent during the post-grant period, where third parties can file oppositions against the patent to challenge competitors.

Risk assessment

The risk assessment process aims to detect, analyze and mitigate patent-related risks. In this study, the activity *FTO analysis* is included within this process. FTO analysis refers to the activity of evaluating whether an invention can be commercialized without infringing other third party patents in a specific market. It involves searching patent databases for relevant active patents and analyzing these to assess infringement risk and relevant mitigation actions. An FTO analysis can be carried out as an input to different patent management processes, for example before filing for a patent. It can also be performed independently of patent processes, for example in connection with projects or before product launch. The activity can either be carried out continuously, or when an event or situation has triggered the need for a risk assessment. In the risk assessment process, intelligence activities may also be carried out to provide for example a preliminary overview or identification of potential risks, while an FTO analysis is conducted when a more in depth analysis of risks are needed.

4.2 AI supported patent management tools

This section provides an overview of AI-enabled tools that were mentioned by one or more respondents during the interviews. The purpose is to provide a contextual background for the subsequent interview findings by illustrating the availability of

tools and their functionalities across the different patent management processes. As this overview is solely based on the tools mentioned during interviews, it represents only a fraction of the tools currently available on the market.

To provide a clearer understanding of tool functionality, interview material was complemented with publicly available descriptions of tool functions from respective vendor website. However, tool offerings in this area evolve rapidly and reported functionalities may change over short periods of time. The information presented in the table is based on features stated on the tool websites in April 2026

The overview is presented in Table 4.1, where each row represents a tool. The references to the interviews in which the tool was mentioned are provided in parentheses. The columns represent the patent management processes applied in this study. For each tool and process, Table 4.1 summarizes activities the tools claim to support based on the information provided on their websites.

Based on Table 4.1, it can be observed that even when considering only the tools mentioned in the interviews, tools are available across all patent management processes. Several tools cover more than one process area, suggesting that tool offerings often span multiple parts of the patent management process rather than being limited to a single process or activity. Patent filing demonstrates the highest level of coverage, with 10 out of the 12 identified tools supporting activities in this process. Furthermore, portfolio management, intelligence and enforcement & exploitation also show a high level of coverage with between seven and nine tools providing support. Risk assessment appears to be the least covered process, with only three tools explicitly supporting FTO analysis. However, it should also be noted that many of the tools claim to support patent searches which may also be utilized across several processes even if it is not explicitly mentioned on the websites.

Table 4.1 AI supported patent management tools

Process → Tool ↓	Patent filing	Portfolio management	Intelligence	Enforcement & exploitation	Risk assessment
Qthena - Questel (R8; R11; R15; R20; R28; Questel(n.d.))	Invention disclosure generation; Drafting; Prosecution; Patent search review	Automated portfolio review	Industry intelligence	Claim chart generation	FTO search review
Solve Intelligence (R8; R13; Solve Intelligence(n.d.-a-b-c-d))	Invention harvesting; Drafting; Prosecution			Claim chart generation	
IPRally (R2; R3; R5; R8; R9; R11; R13; R20; R23; R25; R26; R28; R32; IPRally (n.d.))	Patent searches	Portfolio analysis and classification	Patent monitoring; Patent searches	Patent searches	FTO searches
Eureka by Patsnap (R8; R13; Patsnap (n.d.))	Prior art search; Patent drafting; Prosecution; Invention disclosure				
XLscout (R20; R25; xlscout(n.d.-a-b-c-d-e-f-g))	Prior art analysis; Ideation and brainstorming; Patent drafting	Portfolio analysis	Technology and competitor monitoring	Infringement analysis; Validity searches; Identification of licensing opportunities	
Patlytics (R13; Patlytics(n.d.-a-b-c-d-e-f-g-h))	Drafting; Patent invention disclosure; Prosecution; Prior art search	Patent pruning; Portfolio classification; Patent-to-product mapping		Validity analysis; Claim chart generation; Infringement detection	
Lexis Nexis – Patentsight (R12; R19; Lexixnexis(n.d.-a-b-c-d-e))		Portfolio optimization	Tracking competitors; Technology scouting	Infringement risk mitigation; Identification of licensing opportunities	
Iprova (R3; Iprova(n.d.))	Identification, creation, evaluation and documentation of inventions				
Derwent Patent Intelligence - Clarivate (R19; R25; Clarivate(n.d.-a-b))	Prior art search	Portfolio management	Competitive intelligence	Validity search; Monetization	FTO analysis
Stilta (R32; R3; Stilta(n.d.))		Portfolio assessment		Infringement monitoring; Validity analysis	
Lightbringer (R7; R8; R14; R26; Lightbringer(n.d.-a-b-c-d))	Prior art search; Drafting	Portfolio management	Monitoring technical fields and trends		
Orbit Intelligence - Questel (R8; Questel(n.d.))	Prior art search	Portfolio management	Patent landscaping		

4.3 How is the potential of AI realized in corporate patent management processes today?

The following section presents an overview of how AI is used and perceived across different patent management processes and activities, based on respondents descriptions of current practices, experiences and perceived potential of AI. These findings are further compared to the theoretical potential of AI from Table 2.2 and identified tools in Table 4.1.

4.3.1 Patent filing

Patentability assessment

Among the activities within patentability assessment, prior art search stands out in the interviews as one of the most frequently discussed activities in relation to the use of AI. In total, 17 respondents from 14 distinct companies mentioned the use of AI in prior art search. Among these, 12 respondents from 10 different companies state that they have adopted AI tools for prior art search. This makes prior art search the activity where AI is most clearly adopted in practice compared to the other activities in this study.

The level of adoption is also reflected in the number of advantages mentioned by respondents. Several respondents emphasized that traditional prior art search based on keywords and patent classifications is highly time consuming, and that AI enabled tools perform better than traditional methods both in terms of speed and quality of results. One IP manager with extensive experience in the field noted that available tools deliver relevant results significantly faster than traditional searches (R8). Respondents from the collaborating company confirmed this, with one respondent stating that, *“Compared to how we did it traditionally, this way allows us to find relevant matches quicker”* (R26). R6, a patent attorney from an external IP department, and R11, a patent attorney from an IP firm similarly reported substantial time savings. One respondent from an external IP department further explained that when less time is needed for the search, it is possible to spend more time on analytical tasks (R5). R6 further noted that in addition to time savings, the thoroughness of searches has improved. This was further confirmed by a respondent at the collaborating company that noted that, *“I get better search results with AI, and compared to outsourcing to an IP firm I am now in control of the search”* (R23), indicating that AI can make it possible to do searches in-house and thereby enabling further control.

At the same time, respondents also highlighted limitations and risks. R2 referred to the black-box problem, where the user cannot understand how the AI model reached the conclusion, thereby resulting in a skepticism to the result. Another respondent from the collaborating company emphasized the importance of remaining aware of the limitation of these tools and that they may occasionally produce irrelevant results (R26). A respondent from an IP firm further stressed that, *“it is*

important with a human-in-the loop approach, because AI might miss relevant documents" (R29), meaning that searches should not be conducted solely by AI but with human oversight.

Taken together, the findings related to prior art search suggests that the theoretical potential of AI to enhance and broaden searches, primarily enabled by search and detection and information retrieval capabilities, was largely realized across all respondent groups in this activity. The fact that eight out of the 12 identified tools in table 4.1 support prior art search further confirms the applicability of AI in this activity.

Drafting

Apart from prior art search, drafting was also discussed by a significant number of respondents. Most often, it was discussed as an activity where AI could have great potential, partly because drafting is perceived as a cost intensive and time-consuming part of patent-related work. One respondent from the collaborating company stated that drafting had some of the largest monetary potential based on how much money that is currently spent on it (R1). At the same time, compared with prior art search, the interview data indicates that adoption of AI tools in drafting is more uneven and often described as partial or not adopted at all.

Even though there were high expectations regarding the potential of using AI in drafting, a recurring pattern that appeared both among respondent from the collaborating company and in interviews with respondents from external IP departments was the dissatisfaction with drafting tools that had been tested, where multiple respondents, ranging from managers to patent attorneys, stated that tools they had tried did not meet required quality levels (R1; R2; R4; R18; R23). R2 from the collaborating company even expressed that they had realized that they initially had overly optimistic expectations regarding the capabilities of AI in this area. Another respondent, also from the collaborating company expressed that, *"There are available AI tools that writes patent applications, but we have not yet implemented it since there are still quality issues"* (R23).

Importantly, the interviews suggest that quality in drafting is not only about producing text, but also about making strategic judgments related to the written material. One IP manager from an external IP department described a well written patent as, *"A good patent should include a mix of breadth, focus, technology and the legal dimension, and the 'golden area' lies in between them"* (R18). This shows the complicated and strategical work related to patent drafting and in turn the high expectations on the quality of the text, which possibly can help to explain why several respondents were critical of patent drafts produced by AI tools. For example, respondents expressed that AI tools could add details that the drafter did not intend to disclose, which reduces control over what is included and what can be saved for later applications and as well as what strategically remains open for interpretation in litigation (R2; R18).

On the other hand, one respondent noted that the ability to use AI to generate patent applications could be valuable for firms pursuing a high-volume patenting strategy, where the number of filings is prioritized over the quality of the applications (R7). This suggests that the perceived value of AI drafting tools available today may be perceived differently depending on the firms patent portfolio strategy or intended use of the patents. If the strategy is mainly to build a large portfolio quickly, the tools available today may be considered valuable because they can enable the production of many patent drafts in a short amount of time. On the other hand, the same respondent highlighted that when the strategy relies on licensing patents, the quality of the claims and the patent application must be high and therefore AI tools cannot be used currently (R7).

In addition, confidentiality and disclosure risk appeared as key constraints specifically for drafting, as this activity involves information about the invention, which if the information is leaked before the patent is granted can create novelty destroying art and limit the patentability. Several respondents therefore emphasized that public AI tools should not be used in this part of the patent management processes (R6; R15;R16), even though they could perform better than other tools (R6).

Despite these barriers, respondents also described use of AI in drafting, for example as use for language improvement and faster iterations rather than generation of a full patent draft. Moreover, while many respondents were cautious about tool performance a consultant with insight in many firms stated that the use of AI in drafting could already now save up to 20% of time (R15) which was also confirmed by R9 from an IP firm who stated that, “*The most time consuming part of my work is writing patent drafts, but when I can use AI as support, it makes the process much faster*” (R9).

To summarize, the interviews indicate that AI usage in patent drafting was perceived by respondents as having strong future potential for value creation. The capabilities recommendation, information retrieval and content creation show the theoretical potential of AI in this activity, particularly through recommending claims and generating text and drawings. The identification of tools in Table 4.1 suggests that this potential is reflected in the available tools, as five of the identified tools offer solutions for drafting. However, the interviews show that no AI tool specifically for drafting was implemented at the collaborating company, although this may partly be explained by the fact that drafting is a task that can often be outsourced. At the same time, interviews with both a consultant and IP firm respondents indicated that AI used in a supportive way in drafting can already result in time savings. This suggests that the potential of AI are already being realized at IP firms, while a gap remains between the potential and practical application at the collaborating company.

Prosecution

Prosecution is another activity within the patent filing process which was mentioned by several respondents, specifically as an area where they see significant potential

for AI support. One respondent stated that prosecution was among one of the most time-consuming parts of the patent process (R15) and another respondent that there are high expectations for cost savings in this area (R10).

Beyond time and cost savings, the use of AI in prosecution was also described by a consultant as having the potential to enhance quality: *"AI can create value in terms of quality improvements by suggesting new arguments and alternatives ways of reasoning"* (R15). This potential was further discussed by several respondents that expressed that there are available AI tools that offer solutions for the prosecution task, helping with breaking down office actions and based on that suggesting response strategies and revised drafts to enhance the likelihood of the patent being granted (R2; R23; R28).

At the same time, the interview material provides a mixed picture regarding current adoption. Two respondents from the collaborating company expressed that they had tested several tools, but none of them yet met their quality expectations (R2; R23). For example, R2 stated that, *"We have tested tools for prosecution and see clear opportunities for using AI to draft office action responses and improve applications in order to obtain granted patents, but the tools that we have tested are not yet good enough"* (R2).

In contrast to this, two respondents from different IP firms mentioned a more advanced use. One respondent described using AI enabled tools for prosecution in daily work but always with a human in the loop approach (R28) and one mentioned being in the pilot phase of two internal tools that could reduce the time needed for these tasks from a couple of hours to approximately half an hour (R21). In addition, R15, who in the role as a consultant interact with multiple IP departments, reported that some companies using AI enabled tools for prosecution and drafting already claim time savings of around 10-20%, with the potential for further savings.

Taken together, the interview material suggests that prosecution was seen as a promising area for AI support by respondents, mainly because the activity is time-consuming and costly. The perceived potential among respondents is also reflected in the theoretical potential, where the capabilities content creation and recommendation enable AI to analyze office actions and generate new arguments for office action responses. In addition, the identification of tools shows that four of the identified tools offer solutions for prosecution, which suggest that the functionality is available across several tools. However, at the same time, respondents from the collaborating company stated that they had not implemented any AI tool for this activity. In contrast, IP firm respondents described both active use and late-stage testing. While none of the respondents from external IP departments discussed prosecution in the interviews, the insights from a consultant indicate that there are external IP departments that have adopted AI enabled tools for prosecution and have consequently seen time savings. This indicates that IP firms and external IP departments are already realizing aspects of AI's potential in prosecution, while there is a gap between the potential of AI and the actual use at the collaborating company.

Remaining activities within patent filing

Apart from prior art search within patentability assessment, drafting and prosecution, there are additional activities within the patent filing process. However, interview material contains limited detail on AI use in these activities, which makes it difficult to present findings per activity. Since invention identification and preliminary assessment were briefly covered in interviews, they are therefore discussed together in this section, while activities that were not discussed by any respondent are excluded.

Regarding the invention identification step, several respondents from the collaborating company discussed AI use and related challenges. For example, R3 mentioned that they had tried an AI tool specifically intended to support invention disclosure creation, but decided not to adopt it because it was perceived as too expensive and because the activity of invention generation was not seen as a problem area to them. Related to this activity, several respondents at the collaborating company mentioned that AI generated invention disclosure texts had become more common. Although no interviews were conducted with inventors themselves, R22 noted that generative AI can make it easier for engineers to fill out disclosure templates. Relating this to the theoretical applications of AI, this suggests that ability of AI to draft and structure innovation disclosures through the content creation capability is at least partly realized. However, since interviews were not conducted with the innovators writing the IDFs, it is difficult to assess to what extent this is actually used.

The subsequent activity is the preliminary assessment stage, which involves specifying the core of invention. Although several respondents expressed that AI generated IDFs increased the difficulty of identifying core of inventions, multiple respondents emphasized that AI was currently used to support this task. In particular, AI was perceived as helpful in identifying the core of invention and suggesting alternative interpretations (R2; R23; R26). None out of the identified theoretical applications of AI in table 2.2 explicitly state that AI can support in specifying core of an invention. However, since one identified application is the identification of technical features based on the information in an IDF through the information retrieval capability, it seems reasonable that this capability can also make it possible to extract the core of an invention. Thereby, this potential application of AI appears to be realized at the collaborating company.

4.3.2 Portfolio management

Although the Figure 4.1 distinguishes several activities within portfolio management, respondents most often discussed portfolio management in general terms rather than activity by activity. When respondents referred to specific activities, the discussion primarily centered around portfolio visualization and mapping. For this reason, findings for portfolio management are presented as one integrated section.

Portfolio management was brought up in many of the interviews. Most statements were in relation to anticipated use cases and possible potential of AI. A recurring

theme was the perceived potential of AI enabled portfolio visualization and mapping to reduce the manual effort required to create and maintain an overview of large and complex patent portfolios and their relationships to products and projects (R22; R26; R20). As one respondent noted, *"From a portfolio management perspective there is significant potential, especially in large organizations where portfolios are extensive and products are highly complex. Even when work is divided into portfolios and sub-units, it is still complex and can lead to a reliance on one person's long-term experience"* (26). In particular, respondents described AI-enabled mapping of patents to products as a highly valuable capability that they would like to have, for example to support stakeholder communication and to answer recurring internal questions regarding patent coverage for specific products or areas (R18; R22).

Beyond mapping to products, respondents described the potential of AI enabled visualization to make portfolios more interpretable. For example, one respondent highlighted that *"it would be valuable to be able to visualize how patents relate to each other and overlap, preferably through formats such as graphs or tree structures"* (R32). Visualization and mapping were also linked to strategic decision making, including identifying strengths and gaps relative to competitors and supporting portfolio actions such as keeping or abandoning rights (R32; R28).

Importantly, one respondent at the collaborating company described that they had carried out experiments related to mapping and visualization by testing AI based portfolio classification solutions with the goal of improving the internal categorization of the portfolio. This was done by comparing tool generated classification with the organization's internal existing structure. The performance of the solution did not reach the desired level, and the effort was therefore not continued with this tool (R32).

Some respondents also linked portfolio mapping and visualization to competitive strength analysis, for example by comparing the firm's portfolio with competitors portfolios in a technical context to identify strengths and gaps (R28; R32). In addition, R32 described potential for strategic value mapping, such as assessing whether the existing portfolio supports strategic objectives and that products that needs to be protected are being protected (R32).

While these themes were largely expressed as future oriented possible potential of AI, the interviews also indicate some current practical use of AI in portfolio management at the collaborating company. Respondents described the use of AI extensions in existing tools to generate easy to read summaries of patents and claims, which was perceived to create value, particularly through time savings and improved communication with internal stakeholders such as engineers (R22). R22 further explained that, *"Patent claims can be difficult to read for those unfamiliar with them, and AI can therefore help by generating shorter and easier-to-read versions of the text"* (R22). Additionally, R32 from the collaborating company also mentioned occasional use of AI in specific tasks such as citation analysis and clustering of documents.

Overall, portfolio management was in the interviews characterized by high interest

and several articulated potential use cases for AI solutions. This potential is further reflected in the wide range of theoretical applications of AI in portfolio management, for example assessing inventions in terms of alignment with strategy, enabling a more accurate and dynamic valuation of patent and portfolio as well as supporting in pruning related decisions. Most of this potential is enabled through the capabilities of pattern recognition & discovery and recommendation. Furthermore, the theoretical potential suggests that AI can enable automation of administrative task. This broad potential is also reflected in the tool availability, where nine of the 12 identified AI tools offer solutions within portfolio management. However, current use of AI was limited to summarization of patents and occasional specific tasks at the collaborating company, while this activity was not discussed in terms of practical implementation by other respondent groups. This suggests that only a fraction of the potential applications of AI within portfolio management is realized at the collaborating company.

A further interesting observation was that several respondents, both from the collaborating company and external IP departments, expressed strong belief in the potential of AI within portfolio management and stated that it would create value if it was possible to use AI for categorization, classification, mapping and visualization of the portfolio. This potential is also reflected in theory, which suggests that AI can support visualization and classification and clustering of data by leveraging the capabilities of pattern recognition & discovery and content creation. The identification of tools in table 4.1 further reinforces this potential, for example, IPRally offers solutions in portfolio classification, while Patlytics claim to support patent to product mapping and classification. Despite this alignment between perceived potential by respondents, theoretical potential and tool availability, AI was not used for these purposes at the collaborating company.

4.3.3 Intelligence

Although the conceptual overview distinguishes between several activities within intelligence, respondents typically referred to these activities more broadly as intelligence work.

The interviews indicate that the use of AI is perceived as highly relevant in intelligence work, particularly as the process is closely related to monitoring as well as the need to handle large volumes of information. For example, R11 suggested that AI could possibly support competitive monitoring through agents that provide alerts when competitors patents are granted, indicating opportunities for more continuous monitoring of the patent landscape enabled by AI. In relation to this, R1 mentioned that AI could possibly be used to *"get better insights into portfolios faster and identify, understand, and capture emerging trends"* (R1).

The possibilities of AI in intelligence work were further illustrated by respondents descriptions of the extensive efforts currently required in this activity. R22, who occasionally requests intelligence reports described intelligence work as particularly demanding due to the large amount of data that must be processed, where much

is handled manually. In addition, intelligence requests often come from different internal stakeholders with varying expectations regarding content and presentation, making alignment with these requests both demanding and time consuming. This perspective was reinforced by R19, who emphasized that intelligence work extends beyond conducting searches to include filtering of results, assessing the most relevant findings in relation to the context and visualizing and presenting findings in a meaningful way, which are all time-consuming tasks. Both respondents expressed that AI could have the potential to save time or simplify these tasks. R19 further described intelligence work as more high level and general compared to activities such as prior art searches: *"Intelligence is mostly about broader landscape analysis. What matters is that it is correct overall, over time and when comparing large amounts of data, rather than at the detailed patent level"* (R19). As a result, the same level of detail and precision in the results is not required, which according to the respondent reduces the perceived risks associated with using AI in this process.

R19 from the collaborating company further mentioned having tested existing commercial tools which were described as powerful. One advantage was that these tools do not rely on traditional keyword and classification based searches, instead natural language searches is used which reduced the need for the user to be skilled in traditional search methods. Additional perceived benefits included that the tool had the possibility to summarize findings and support analysis by relating results to external data such as competitors. R19 described that these functionalities could enable both time savings and more advanced analysis. Despite these experienced benefits, the respondent mentioned that they had only used these tools a few times and had not fully adopted it in daily work. One mentioned reason was the preference of working with traditional search methods, but also that the respondent wanted to maintain traceability when presenting results: *"One reason I don't use AI based search very often is that when I report my results, particularly in enforcement related contexts, I usually report the actual search string I used. With class-based searching I can show exactly what I did, which is not possible with AI based tools"* (R19). This suggests that experiencing the value of AI tools does not always translate into actual use. Instead, adoption can be hindered by personal preferences for established ways of working and in this case the preference to maintain transparency when communicating results.

Taken together, the interviews shows that respondents saw significant potential of using AI in intelligence work, for example to enable continuous monitoring and handling large volumes of information. This is confirmed by the broad theoretical potential of AI in intelligence. For example, the capabilities of monitoring and search & detection can enable monitoring and searches in large amounts of data, which confirms this specific perceived potential. Additionally, AI could also help in identifying unexplored opportunity areas or predict future market sizes, enabled by the capabilities of forecasting, pattern recognition & discovery and information retrieval. In addition, AI also has the potential to create visualizations which in turn leverage the content creation capability. Further, the potential is reflected in the identification of tools, where seven out of the 12 identified tools in table 4.1 offer solutions within intelligence. This shows an alignment between the respondents

perceived potential, the theoretical potential and the functionalities offered by tools. However, the use of AI in intelligence work at the collaborating company remained limited to only sporadic use. The interview material did not provide indications of AI use for intelligence-related activities among the other respondent groups. This shows that in the intelligence process, there is a gap between theoretical potential of AI, respondents perceived potential and the actual use at the collaborating company.

4.3.4 Enforcement & exploitation

Infringement search and analysis

Respondents described a mixed picture of AI in infringement analysis and search. While several respondents highlighted clear opportunities, others expressed skepticism and emphasized that current tools have not yet reached the required level of quality.

Skepticism was expressed in different forms. One respondent argued that AI has limited capabilities in this domain and will fail to find relevant infringements, which constrains its usefulness in infringement purposes (R14). Others highlighted that the feasibility of AI-supported infringement search is strongly shaped by detectability and the availability of observable evidence. One argued that AI may offer less value when searching for hardware-related infringements than for software (R18). It is often difficult to determine from online sources whether a physical product infringes a patent, whereas potential software may in some cases be easier to identify. Online photos of products are often not detailed enough to identify infringing features, and an assessment may therefore require inspecting the product in person or even opening it up to evaluate the components against the patented features (R18). At the same time, R20 mentioned that hardware functionality is often more visible and likely described in manuals, product documentation or marketing materials. In comparison, many software patents describe functionalities that are deeply embedded into systems which could make them more difficult to detect (R20). This suggests that detectability constraints are case dependent and may limit when AI can be meaningfully applied to infringement search.

However, several respondents expressed strong potential for AI in infringement work. R10 suggested that if a company's strategy relies on claiming infringements to generate revenue, AI could be used to continuously monitor the market for potential infringements. However, this use case was also framed as a risk by R1, *"There is a risk that if the tools get better we might get accused of infringements, and in this case it may be preferable that the technology is not fully developed"* (R1). R2 similarly confirmed the perceived risk that competitors could use AI to map a company's products against their own patent portfolios to identify potential infringements, while also noting that the same approach could be used proactively to assess competitors' products (R2). This indicates that what is perceived as potential for one actor may be perceived as a risk for another. Apart from the expressed concerns of increased strength of competitors, R20 also mentioned that patent trolls could gain significant possibilities from AI search tools to identify infringement cases and

generate higher volume of infringement claims, potentially increasing the number of cases organizations must handle (R20). Although respondents mainly discussed this as a risk, it also indicates that AI-enabled infringement search is perceived as strategically valuable as it could strengthen an actor's ability to identify and pursue infringement more effectively.

Apart from perceptions of opportunities and risks, respondents who had tested AI tools emphasized that limited customization of output and a high-quality requirements remain key reasons for not implementing AI in infringement work. R20 from the collaborating company stated that they were actively looking for infringement detection solutions as it would save large amounts of time and enable more work to be handled internally, but it would only be valuable if the tool delivered acceptable performance. In this context, "acceptable performance" meant that the tool would need to match the quality of an infringement search delivered by external IP firms, but within a shorter amount of time. Since the tools tested did not meet this requirement the work remained mostly outsourced and Microsoft Copilot was described as the best available alternative at the time (R20). Another respondent from the collaborating company mentioned having tested a specialized AI tool aimed at infringement monitoring and invalidation but also assessed Copilot as a better alternative (R25). According to R25, a key limitation of the tool was that results could not be presented in the way the user wanted, partly because predefined prompts could not be changed. This shows that perceived usefulness of a tool is not only dependent on model performance, but also on whether the interface supports the user's preferred way of working and output format.

Overall, infringement search and analysis show strong alignment between perceived potential among respondents, theoretical potential and the relevant AI capabilities. The activity aligns well with capabilities such as search and detection and monitoring, which suggests that AI can support extensive infringement searches and continuous monitoring for potential infringements. A further interesting potential application is the possibility to predict infringements due to the forecasting capability of AI. Regarding availability, four out of the identified tools in table 4.1 offers solutions for infringement detection and analysis. Although respondents confirmed the potential and emphasized the value of using AI for infringement search and active competitor monitoring, no AI solutions for this purpose were implemented at the collaborating company and the study did not identify clear indications of established use in other respondent groups. Overall, the theoretical potential for infringement search and analysis do not yet appear to be realized in practice, although to a limited extent at the collaborating company in regards to the use of general-purpose tools.

Validity analysis

Several respondents expressed positive experiences of using AI in validity analysis, primarily when searching for novelty destroying prior art (R8; R25; R2; R28). However, experiences with specific tools were mixed. R25 from the collaborating company described tool X as useful in invalidity work and outlined a workflow where

tool X can be used to generate a list of documents, after which Copilot can be used to structure a comparison by for example creating a claim chart. In contrast, R20 who also works at the collaborating company expressed dissatisfaction with tool X for validity searches, arguing that it often fail to identify relevant prior art. At the same time, the same respondent noted that search for novelty destroying art in this phase is challenging. Since the target patent has already been examined for novelty by a patent office, truly novelty destroying prior art is expected to be difficult to find. Nevertheless, despite these opinions, R20 continued to the tool to search for prior art occasionally. However, the different experiences of the same tool suggest that assessment of tool performance is subjective and influenced by for example individual working styles, familiarity with the tool and expectations regarding output.

Beyond individual perceptions, the perceived performance of AI tools can also vary depending on the patent-related activity for which they are used, a point further emphasized by R20. In invalidity-oriented searching, identifying one sufficient strong reference is a plus, while not finding such material may simply mean that no clear opposition case can be built. By contrast, in activities such as for example FTO analysis, missing a relevant document may have significant legal, commercial and or economic consequences, thereby leading to lower tolerance for imperfect AI support. This suggests that tolerance for imperfect AI support can be tied to the risk profile of the activity, making validity analysis a context where exploratory use of AI is more acceptable than in activities where errors can cause larger consequences.

Similar to infringement search and analysis, respondents expressed concerns about how their work might be affected when competitors start using AI in validity related activities. One respondent from another organization described that they had received an opposition containing an unusually large number of objections, “*We recently received an opposition with an extremely large number of objections, which made me suspect that it was AI generated. These kinds of oppositions will increase our workload and costs because we have to analyze and respond to these documents*” (R18).

Relatedly, R28, a division manager at an IP firm, described AI-generated opposition as a competitive tactic. A faster, lower cost and lower quality opposition that was intended to deter or disrupt competitors. Together, this suggests that AI in validity related work is associated both with concerns about increased opposition workload and with offensive opportunities to apply pressure through faster and lower quality oppositions.

Overall, the interviews showed that AI tools are already used in validity analysis, mainly to support novelty destroying prior art searches. This was reported by respondents from both the collaborating company, a respondent from an IP firm and one from an external IP department. This indicates that the practical application of AI in validity analysis appears to reflect several of the potential applications, particularly regarding comparison of claims with prior art and support in patent searches, both at the collaborating company as well as by respondents from other respondent groups. This potential is furthermore reflected in the identification of

tools, where six out of the identified tools explicitly stated to support either claim chart generation, validity analysis, or both.

Remaining activities within enforcement & exploitation

The remaining activities in the enforcement & exploitation process, legal action and commercial negotiation and monetization, were not covered in sufficient detail in the interviews to support activity level analysis of current AI use. However, one respondent from a patent law firm described established use of AI in legal enforcement work, including support for drafting legal documents such as agreements but also support in other infringement related contexts (R28). This description can be related to the potential of AI to support negotiation activities. Although negotiations were not explicitly mentioned by the respondent, the ability of AI to generate new perspectives and thereby strengthen arguments aligns with such applications. However, the theoretical potential of AI to assess the likelihood of litigation enabled by the forecasting capability was not reflected in the interview data. Neither of these potentials was explicitly identified among the tools presented in Table 4.1, which can suggest that these applications are currently less explored, at least by tools identified in this study.

In addition, regarding the monetization activity, several theoretical applications were identified, for example identifying licensing opportunities which are primarily enabled by the search & detection as well as the pattern recognition & discovery capability of AI. This potential is further supported by the fact that some of the tools identified claim to support such applications. However, these potential applications were not reflected in the interview data.

4.3.5 Risk assessment

FTO analysis

Respondents who discussed FTO described it as an area with clear potential for AI support, but also as an activity with particularly high requirements for reliability and traceability.

Perceived potential was connected to the knowledge that AI already performs well in prior art search which was seen as a capability that could potentially be transferred to FTO contexts (R2). In addition, several respondents at the collaborating company described FTO activities as a bottleneck in the patent-related workflow due to the amount of manual work required to screen and interpret large volumes of patents. This work is often carried out in collaboration with engineers that are directly involved in the development of the assessed invention, which further increases the time and effort required. R20 from the collaborating company similarly noted that, *"Risk assessments and FTO analyses can take a long time, which can become problematic since it leaves us with less time to perform the actual risk mitigation"* (R20)

Despite clear drivers for more efficient and less manual approaches to FTO work, the interviews highlight trust and traceability as fundamental factors preventing the use of AI tools in this activity. R11, a patent attorney at an IP firm, emphasized that FTO analysis is not only about identifying potential risks, but about being able to demonstrate, to the greatest extent possible, that an invention does not infringe third-party patents. In this context, AI based search tools were described as potentially valuable when they identify a problematic patent, as such findings can directly trigger mitigation actions such as design arounds or licensing negotiations. However, when an AI based search produces no critical results, respondents expressed difficulty in justifying confidence in the outcome if the underlying search process cannot be explained.

R11 characterized this challenge as a black-box problem and noted that the AI tools they had tested for FTO purposes did not provide sufficient insights into how conclusions were reached. In contrast to manual searches, where analysts can explain the reasoning process and take responsibility for the scope and limitations of the search, the lack of transparency in AI generated results was perceived as a critical reason for not adopting AI in this activity. In FTO analysis, the primary concern lies in the risk of missing critical documents, and without the ability to explain the search process, it becomes difficult to justify what may have been overlooked. R11 furthermore stated that they had not yet found any AI tool on the market that offered the structure that they wanted and explained how it reached the conclusions. This concern was further confirmed by R20, who noted that even when FTO analyses are outsourced and conducted externally, organizations still expect to be able to inspect search queries and understand why certain documents were excluded: *"I do not always check what the firm has done or which documents was filtered out, but I want to have the possibility to review their search strings in case questions later arise regarding a problematic patent. In such situations, you want to be able to see whether the IP firm representative has considered that patent in the analysis or not"* (R20). Taken together, this suggests that for AI tools to be valuable in FTO work, they should support documentation, explainability and traceability, thereby enabling users to justify and defend the findings.

At the same time, a respondent from an external IP department said that commercial search tools already save time and also may reduce the risk of missing relevant documents and that leading organizations already perform automated FTO searches before human review (R6). In contrast, R8 mentioned that suppliers of AI tools are beginning to develop AI solutions for FTO work, but that the solutions are not fully mature yet. Although maturity was not defined in detail, this assessment aligns with R11 experiences of testing AI tools that did not meet their requirements.

Taken together, the answers related to FTO analysis indicate that AI is perceived as having significant potential to enable more efficient searches, but also as involving substantial risks. This potential is further confirmed by the theoretical potential of AI, as its advanced search and detection capabilities could enable faster and more comprehensive FTO searches compared to traditional methods. This potential is also reflected in the fact that three out of the identified tools claim to offer solutions

explicitly for FTO analysis. However, the collaborating company gave no indication of the use of AI tools in this process and only one external respondent mentioned the use of AI tools in this area. In contrast, the majority of respondents that mentioned this activity emphasized that FTO requires defensible and traceable results and that the AI supported solutions they had encountered so far do not yet meet these requirements. Overall, this shows that there is a gap between the theoretical potential of AI and the actual practical application in FTO work at the collaborating company, while one respondent from an external IP department appeared to have realized value in terms of time savings.

4.4 What are the organizational barriers of adoption in relation to integrating AI into corporate patent management processes?

This section presents an analysis of empirical findings relating to organizational barriers affecting AI adoption in patent management functions. The analysis focuses on how the actors in this study perceive challenges associated with implementing AI in patent-related activities. The analysis aims to provide an understanding of how organizational structures, capabilities and contextual factors influence adoption of AI in practice.

4.4.1 Confidentiality and security requirements as a barrier of adoption

One of the most recurring concerns about adopting AI in the patent management process relates to a fear of exposing sensitive information. In 17 of 32 interviews, respondents spontaneously raised concerns about data leakage, tool security or loss of confidentiality when discussing risks of AI adoption. One respondent highlighted this by explaining, “*There is very high potential in using search tools, but we do not adopt it because of security issues*” (R17).

Another respondent highlights a similar description in relation to drafting, “*There are many drafting tools but we cannot use them because of security issues*” (R16). This is also highlighted by R28 who is very positive towards AI in general but still describes, “*We constantly evaluate drafting tools but we have not found one who is both good enough and secure enough*” (R28). From an adoption perspective, this implies that individuals who are willing to use AI for efficiency gains may still refrain from doing so when they cannot verify how external companies handle data.

This concern is especially relevant in the context of patent management because confidentiality is commercially critical. This means that the practical implementation of AI is constrained by the strict confidentiality requirements inherent to IP-related work. Information disclosure can create irreversible consequences, which makes risk tolerance very low. As one respondent describes, “*The problem of confidentiality slows down the AI transition, especially because of the requirements of IP*

work” (R10). Another respondent expresses the same underlying logic but from a cultural perspective, “*This industry is like a zero error business*” (R21). Together, this indicates a low risk tolerance that limits adoption.

When tools are not bound to the organization’s controlled environment or when it is unclear where the data is stored, users in this domain act cautiously, even if external tools perform better, “*The public tools work the best for the early stages of the patent management process, but we cannot use them because of the risk of losing novelty if the information is released before filing*” (R6). This restricts the possibilities of adopting the tools that perform the best. At the same time, respondents explain that internally developed tools do not perform as well as the public ones and describes a fear that some people might still use public tools because they have access to more data and perform better, “*Internal tool X does not have all the data, so there is a risk that people still use the public tools*” (R3). This perspective is strengthened by the perspective of one of the IP firms, “*Our customers do not want us to use cloud-based tools, so we are forced to develop our own tools*” (R28).

Collectively, these findings indicate that even when AI tools offer clear advantages, their use is constrained by the high stakes associated with patent protection, limitations in secure alternatives, and a risk-aware culture. This results in a cautious approach to AI adoption, where risk mitigation is prioritized over rapid adoption.

4.4.2 The search for the perfect tool as a barrier of adoption

The interviews indicate that one of the barriers of adoption is the performance of tools. At first glance, this concern appears to stem from the current limitations of AI tools in terms of quality. This perception is not unfounded, as many tools are still in their development phase and have not yet reached their full potential in terms of quality. However, a closer examination indicates that this barrier may also reflect a mindset characterized by waiting for the perfect tool before committing to actual use. The rationale behind this is that there is evidence that suggests that while the tools might not produce perfect quality, there is still high potential in capturing value from them in their current state. Several respondents describe abandoning AI tools after small-scale trials, often because early interactions were perceived as inefficient or unreliable compared to established ways of working.

This tendency is reflected in how respondents explain both their own behavior and that of the broader industry. As one respondent notes: “*People in our domain just look for errors and problems instead of focusing on how we can use AI*” (R28). Similarly, another explains that attempts of using AI are often unsuccessful, “*Every time I try using AI and spend time on the prompt, the resulting answer becomes that the AI can’t do what I asked for, and in the end I have spent as much time as if I had just done it manually*” (R24).

This perception is further reinforced by observations that some individuals dismiss AI use without meaningful interaction. As one respondent describes, “*When I have discussions with colleagues who claim the tools do not work, and I look into their AI*

use afterwards, I can see that they have not even used the tool” (R28). The same respondent also highlights the influence of external narratives: *“People hear a lot in the media and form their own negative opinion of AI without grounding them in facts”* (R28). This perspective is further illustrated by respondents who express skepticism despite limited or no personal experience. For example, one respondent stated that, *“AI tools do not work well enough because of their hallucinations”*, while later clarifying that this view was based on what colleagues had said rather than using the tools (R14).

This dynamic is reinforced when users set high performance thresholds, such as expecting AI tools to replicate the quality of their own work while delivering results at a substantially greater speed. As one participant expressed, *“The level of quality that a tool should deliver if I were to adopt it should correspond to exactly what I can deliver myself, but in a hundredth of the time”*(R20). This expectation of a perfect tool is sometimes expressed more implicitly. One respondent compared testing an AI tool to *“trying on a good but too small jacket, it wasn’t quite right, so why would I spend money on it?”* (R7). While such expectations may appear unrealistic, they are not irrational. From an organizational perspective, adopting a tool that initially reduces efficiency while also adding costs may seem unjustifiable. However, the challenge lies in the fact that the true value of AI often only becomes apparent after long-term use and learning, which is never reached if the tool is abandoned too quickly.

The common factor across these observations is that these respondents have typically engaged in limited testing but have not integrated AI into their daily workflows. This suggests a mechanism in which high initial expectations combined with low tolerance for continuous trial or iteration prevent users from reaching a stage where the value of AI becomes visible. As a result, the benefits remain unexperienced, reinforcing the perception that the tools are not useful.

This interpretation is further supported by contrasting experiences among respondents who have adopted AI into their work processes more extensively. These individuals report substantial productivity gains but emphasize that such outcomes required continuous experimentation and learning, as well as iteration with the tools. One respondent, whose work tasks include the full patent application process explained that the transition did not occur immediately, but developed over time and now describes his productivity as *“50% faster with AI”* (R9). Another respondent highlighted that the search tools are highly effective but stressed the importance of iterative use, *“You have to go back and forth with the tool, you have to refine and ask new questions”* (R5).

Interestingly, these more experienced users tend to frame iterating with the tool as a natural part of using AI rather than a reason to abandon it. As one respondent explains, *“Sometimes the result is really bad, but then you just need to reframe your prompt or need a new angle, you can choose to give up right there. . . . I have had enough successful results that I choose not to give up”* (R25).

Taken together, these contrasting perspectives suggest a self-reinforcing pattern: individuals who discontinue their use of AI after failed attempts never experience the benefits that would justify continued adoption, while those who persist through the initial learning phase are able to unlock significant productivity gains, despite the tools not performing perfectly. From this perspective the search for a perfect tool becomes a barrier in itself, as it limits the experimentation and skill development required for the value to materialize.

An additional factor that reinforces this dynamic is the limited time available for experimentation. Several respondents, including patent attorneys themselves, as well as managers and other IP professionals emphasized that patent attorneys are heavily time constrained because of workload. As one respondent stated, “*Patent attorneys have no time to test, they are too heavily constrained by workload*” (R9), while another noted that, “*Even if we would like to explore the tools, we are too heavily constrained by workload to do so*” (R23). This lack of time further restricts the possibility of moving beyond the initial trial, thereby strengthening the cycle in which limited use leads to limited perceived value, which in turn discourages adoption.

4.4.3 Lack of trust as a barrier of adoption

A recurring concern described as a reason for limited adoption is that generative AI can produce incorrect output (hallucinations). The interview material confirms that practitioners frequently encounter incorrect output. Respondents describe risks such as, “*AI delivers facts that are not true*” (R2). Another respondent describes, “*Wrong input can lead to completely wrong conclusions, particularly in niche areas such as legal questions*” (R1).

This supports the idea that hallucination is an issue that could hinder adoption. However, multiple respondents challenge the assumption that incorrect output is an AI-specific threat. One respondent highlights this by saying: “*AI still makes a lot of mistakes, but humans do too*” (R9). This implies that hallucinations may not represent a fundamentally new risk. A senior respondent with extensive experience in patent operations notes that the patent management process is already equipped to catch errors, regardless of whether they were produced by a human or a machine: “*Hallucinations have always existed within IP management, we are already equipped to catch them*” (R10). Another respondent reinforces this point by describing that they already have established safety mechanisms, “*Patent firms usually ensure that they always have two pairs of eyes on patent work, one person who drafts and one who reviews. Internally, we always have both a patent attorney and an engineer that reviews*” (R18).

In addition, some respondents alter the perception of hallucinations as a serious risk factor. As one respondent explain, “*When AI is used in a workflow where the model does not have an immediate effect on the world, and a manual control step is feasible, hallucinations are not that big of a problem*” (R12).

These perspectives suggest that the core issue may not be hallucinations per se, but rather a lack of trust in AI systems which results in resistance of adoption. The interpretation of lack of trust being the core issue is supported by one respondent who notes, “*We place higher demands on AI than on humans because we trust humans, we have the mindset of ‘we are all human’, which we do not extend to machines*” (R18). Similarly, R27 highlights how trust is established in traditional professional contexts, explaining that when hiring a reputable IP firm, clients can see the individuals involved on the firm’s website, “*You can connect to the individual because you see their face, this helps in building trust*” (R27). The respondent further adds, “*Many of these firms generate clients because of their history or reputation*” (R27). Emphasizing that the sense of security may not primarily stem from the absence of errors, but rather from personal trust in the actor producing the outcome.

This implies that the problem of adopting AI in patent management is not about if risks of factual errors or wrong interpretations are present, these risks already exist, with or without AI. What becomes evident is that to generate value in patent work, you need experienced actors that you trust and that can produce results that reach the desired quality level. From this perspective, the main challenge is building trust in AI systems, which is further complicated by AI’s limited explainability. The respondents highlight this by mentioning “the black-box problem”, that relates to AI’s inability to explain how it arrived to its conclusions. For example, “*It is a black box when you use AI. I want to have control and be able to show what I have done*” (R19). Another respondent strengthens the perspective that trust in an AI system will not be established as long as it lacks the ability of explainability, “*So far I haven’t seen an AI tool that can actually account for exactly what it does ... if there were explanations of how the AI made its decisions, I would start to trust them*” (R11).

This implies that with the same rationale that you would not trust a human who cannot explain how they arrived at their conclusions you will not trust an AI system who lacks that ability. The respondents explain that this issue must be handled by keeping humans with enough understanding of technology and patent law in the driver’s seat, “*We can’t let go of the human-in-the-loop control*” (R6). Or as one respondent who has adopted AI extensively describe, “*I always sanity check, it is always my own analysis. I would never submit something I cannot stand behind*” (R9), and thereby indicates that the responsibility of the output is never delegated to the tools.

This also aligns with established guidelines in the field where accountability remains with individuals. As one respondent with experience in patent examination notes, “*It is still very important that humans have responsibility for what is submitted*” (R11). Consequently, AI integration is not perceived as a fully automated push-button solution, but rather as a support tool under human control.

However, if AI is assumed to be less trustworthy than skilled professionals at its current level of maturity, its potential value becomes uncertain. The empirical findings suggest that for experienced users, AI does offer meaningful benefits. Nevertheless,

realizing this value while maintaining a level of quality comparable to that of experienced professionals depends on what the respondents describe as "competent operators".

Respondents highlight AI as a tool for efficiency improvement, but only when the tool is operated by people who combine skills in using AI and domain expertise, both in terms of technology and patent law, "*If the person using the tool is sufficiently knowledgeable about the problem, both technically and legally, and understands what they are doing, AI is not a problem*" (R26). Another respondent describes the competent operator as someone who is both highly experienced in the use of AI and has extensive knowledge in the patent domain and notes that, "*The combination of a competent operator and an AI tool is very powerful, while the same tools in the hands of someone who is not as competent carry substantially greater risk*" (R31), while describing that competent operators can produce results faster with the same level of quality.

This indicates that because of the lack of trust in AI systems, stemming from their limited explainability, the role of competent operators becomes critical in capturing the efficiency gains offered by AI. To achieve this, existing safeguard mechanisms that currently protect IP departments such as double review procedures should be maintained. At the same time, it is essential that AI systems are used by individuals with appropriate expertise. Competent operators are both skilled in patent work, but also understand the limitations of AI and how to use the tools efficiently.

Taken together, these findings suggest that although hallucinations may present a risk, they do not in isolation constitute a barrier of adoption. The findings indicate that the central challenge lies in establishing trust in AI systems and integrating them into workflows in ways that align with the norms and expectations of the patent domain. This is further complicated by AI's lack of explainability. Establishing trust and responsible use includes recognizing that errors that are made by AI do not fundamentally differ from those made by humans, but rather require similar processes of control. At the same time, successful adoption depends on the presence of competent operators who do not blindly trust AI-generated outputs but instead critically evaluate them and, where necessary, rely on their own analysis or perform parts of the work manually when explainability is required.

It should be noted that this barrier reflects the current conditions described by the respondents. As AI systems continue to develop, both in terms of performance and explainability, and as familiarity with AI increases, reducing the current asymmetry in how human- vs AI- errors are perceived, the lack of trust as a barrier of adoption may be less pronounced over time.

4.4.4 Introduction of bottlenecks as a barrier of adoption

The adoption of AI in one part of a workflow can improve local efficiency while increasing workload downstream. This results in a system-level adoption barrier where AI makes one task faster but slows down the end-to-end process. This problem was

identified in the empirical findings and observations from within the collaborating company. This became especially apparent in the flow between drafting IDFs, which is performed by inventors within the organization, and subsequent steps of the patent process. While AI tools are perceived to simplify the process of writing an IDF, the introduction of AI simultaneously increases the workload in downstream roles.

The empirical findings suggest that AI-generated IDFs often increase text volume and contain a substantial amount of irrelevant information, “*When IDFs are generated with AI you receive very long texts, where 90% of the information is irrelevant*” (R22). This creates additional work because “*it makes it difficult understand what the invention is because it is hidden in 20 pages of irrelevant information*” (R22). This highlights an asymmetry where adoption in one part of the process leads to efficiency gains, while making other stages more complex and time-consuming.

For inventors, AI reduces the time invested, but for patent attorneys and those in portfolio roles, AI-generated content increases workload, as they must analyze more complicated documents in which relevant information is hidden, requiring additional time and effort to identify the core of invention. As one respondent explained, “*AI generated IDFs generate an additional review step for patent attorneys, where we have to connect with the inventor and update the information*” (R23). Another respondent strengthens this perspective by explaining, “*When inventors use generative AI to write arguments, they often become very well-articulated, but this moves the workload to reviewing roles, the problem is that the material look very convincing but it needs a higher level of critical verification*” (R26).

The same respondent also introduces an additional perspective by describing that AI use in drafting of IDFs is not per definition negative, “*Some inventors that receive a negative patentability assessment because of my interpretation of what they wrote, now has the opportunity to use AI to understand my reasoning and submit an answer if I misinterpreted, this opens up for more inventions that would never had passed this initial step previously*” (R26). This indicates that while adoption of AI introduces more work in some steps of the process it also has the potential to enhance communication between departments. In addition, the option of using AI to write invention descriptions may lower the threshold for requesting a patentability analysis. This is beneficial, as it may lead to a greater number of inventions being formally assessed. These dual effects imply that there is a need for clear governance structures and guidelines for AI use in order to gain from the benefits while not introducing bottlenecks.

Taken together, this implies that there is a misalignment with incentives for how and where AI should be implemented across full workflows. There is also a risk that organization-wide initiatives reinforce this challenge. This observation is true in the collaborating company, where there is a strong push for integration of AI into all work processes, “*Writing IDFs with AI is problematic, but we can’t criticize it because the organization want people to use it*” (R3). This perspective suggests the emergence of an unintentional “don’t-question-it” norm, where a strong push for AI adoption makes it difficult to raise concerns, even when a use case is not sufficiently

mature from an end-to-end perspective. This becomes a barrier of adoption because the teams that absorb the increased workload are also accountable for efficiency of their own tasks, and without clear guidelines or mandate to control upstream AI use they may be incentivized to resist or limit AI generated requests, preventing AI from maturing into an integrated end-to-end process solution.

4.4.5 Fear of losing competence and professional judgment as a barrier of adoption

One of the central findings relates to how respondents anticipate how AI will affect knowledge and professional judgment over time. The data reveal two contrasting interpretations of this competence-related risk.

On one hand, several respondents frame this risk as an adoption barrier. They argue that while the implementation of AI is likely to increase efficiency, it may also introduce unintended negative consequences for professional judgment. The opportunity of increasing efficiency will put pressure on individuals to produce results faster. Over time, as users become more comfortable relying on AI, this productivity pressure may reduce the extent to which outputs are critically reviewed. As R10 describes, *“If the pressure for efficiency increases as a result of AI adoption, and tools are available to produce faster, people will stop checking because of time constraints”* (R10). R19 strengthens this further by describing how the pressure to deliver faster may introduce an element of laziness, *“If AI makes it possible to work faster, you may feel pressured to use it to keep up, and the risk is that you become lazy and lose the understanding of what you are working with”* (R19).

This dynamic raises concerns about declining quality of judgment and overreliance of AI systems, particularly in high stakes contexts. As a result, some respondents view this risk as a reason to be cautious about adopting AI. At the same time, several respondents point to the risk that increased reliance of AI could lead to a gradual decline of the knowledge required for patent work, especially regarding new entrants, *“When adoption of a new technology starts, it is operated by highly skilled people, but over time the skill level of people using it goes down... it is unclear how we could get the next generation to be as skilled”* (R14). The same respondent argues that in the future fewer people will have the experience needed to judge correctness, *“In the future people will get wrong answers but not know that it is not right because they do not have the experience”* (R14). This concern is reinforced by R2, who highlights the specific vulnerability for new hires, *“For those entering the profession, there is a risk of losing the fundamentals if they begin relying on AI right away”* (R2).

On the other hand, some respondents portray AI as competence-enhancing, positioning it as an adoption enabler rather than a barrier. For these respondents, the same speed and information accessibility questioned previously is interpreted as a way to strengthen learning and communication. For instance, R1 suggest that the quality of patent applications may actually improve as AI enables faster competence development. Similarly, R32 highlights that AI can support technical understanding,

explaining that when handling cases in unfamiliar domains, AI can assist in processing of information and accelerate understanding. These perspectives indicate that AI will enable more rapid skill acquisition.

In addition to individual learning, some respondents explain AI's role in collaboration across organizational and technical boundaries. One respondent describes AI as a translator, enabling individuals to understand views of other teams better, "*AI can act as a translator between roles*" (R5), explaining that employees often lack time to familiarize themselves with other functions expertise areas. R5 further describes that AI will provide higher levels of transparency by lowering the threshold for understanding colleagues with different expertise. This will reduce misunderstandings and enable better coordination because of the increased access of knowledge in multiple areas (R5).

Several respondents also draw parallels to previous technological transitions to argue that AI is unlikely to erode competence in the long-term. One respondent compares the shift to AI with the transition from hand writing to using Microsoft Word, emphasizing that the underlying craftsmanship remains even if the tool change, "*Within a longer time period the craftsmanship will still remain the same, but the tools will change, just as what happened before with the transition from using a pen to using Word*" (R15). This perspective is further strengthened by R18 who draws parallels to previous changes in ways of working, where concerns about declining competence proved to be overstated. The respondent argues that AI is likely to reconfigure professional roles rather than diminish expertise, "*Previously, practitioners were required to spend time in roles such as translators before advancing to drafting, this step was later removed, but it did not lead to a decline in the competence of those performing drafting*" (R18).

The earlier descriptions of concerns regarding junior professionals are also challenged. One respondent argues that the younger generations are already accustomed to using AI in their daily lives and are therefore likely to understand how to use it correctly and naturally integrate it as a work tool. From this perspective, AI enables junior practitioners to become productive faster, rather than hindering their ability to develop relevant expertise, "*Many talk about the risk of how we will educate our young. Personally, I do not see this as a major risk. Younger employees already use AI in their everyday lives, so for them this isn't a revolution. Rather it should be seen a work tool that can help them become productive in the company much faster*" (R28).

R28 further highlights the problem with viewing AI as competence destroying by highlighting the risk of tension between generations within organizations, where concerns about using AI may lead senior practitioners to restrict AI use among new hires. The respondent describes that many senior employees argue that juniors should avoid using AI during their first years in order to develop foundational skills the same way as previous generations. However, R28 challenges this reasoning suggesting that restrictions will be counterproductive using an analogy: "*When you are teaching someone how to drive, you don't switch off the anti-skid function*

or all of the cars sensors, just so the learner ‘really learns the basics’” (R28).

While there appear to be contradicting perspectives regarding whether AI will enhance or erode competence, there is consensus regarding that introduction of AI will introduce the risk of laziness or less critical examination of outputs over time because of productivity pressure, but some respondents do not view this as a reason to refrain from AI adoption. Instead, the emphasis is placed on managing this risk through appropriate practices. One respondent describes that organizations must find a balanced pace of adoption while continuing to develop internal competence (R8). R15 strengthens this by noting, *“There is definitely a risk that you become lazy and lose competence, but you have to make sure you have proper quality controls and actively choose to review critically, that protects knowledge to some extent”* (R15).

Collectively, these findings suggest diverging perceptions on AI’s long-term implications for competence. Some respondents empathize the risk of deskilling, which may lead to selective adoption, particularly if AI use is restricted for junior professionals, while others view AI as competence enhancing, supporting information accessibility, accelerated learning and cross-boundary collaboration. However, respondents agree on a shared risk mechanism: As AI increases speed and lowers the effort required to generate output, productivity pressure may weaken review discipline and increase overreliance. Consequently, respondents highlight that effective adoption requires that risk-mitigation practices are implemented. This finding is particularly important because it directly connects to the previously highlighted requirement of maintaining a human-in-the-loop at the current maturity level of AI.

4.4.6 Inertia and social stigma as a barrier of adoption

One of the themes that emerged as adoption barriers in the data relates to social and psychological factors that hinder adoption. Even when respondents recognized potential value, for example in time savings or improved quality, frictions such as social stigma or comfort with existing routines were observed.

This pattern is not apparent across all interviews in the collaborating company, rather the findings indicate a mixed level of adoption. Some respondents describe extensive integration into their daily work, one respondent states, *“I use AI every day”* (R25). The respondent further explains many use cases that are not specifically limited to the specific role of this person, for example, for all types of summarizations, for answering emails, to compare documents, to find information, and emphasizes that the associated time savings are substantial (R25). On the other hand, there are respondents that describe inertia where they feel efficient enough or stuck in their ordinary ways of working, in addition to observed perceived lack of urgency. These factors reduce motivation to invest time in learning. One example was observed in R3’s description, *“I feel like a person standing on the side of this transition. I get stuck in doing things my usual way, and I feel like I am just as effective on my own, but I can see the potential in other departments, for example HR”*(R3).

This implies that the perceived value of AI in the individuals' own role is low. When employees feel efficient enough, AI becomes an optional addition rather than a necessity which may lower the willingness to invest in the initial cost of learning. This indicates that there exists a passive resistance where change or adoption is not perceived as negative, but not positive enough to be prioritized. R19 strengthen this perspective by describing, "*I haven't really gotten started with AI because I like my way of doing things*" in addition to, "*I don't use it consistently in my work and I do not clearly see the benefit*" (R19). These descriptions indicate that the value proposition is not sufficiently clear for this user, which may reduce the motivation to adopt.

In addition to inertia and established ways of working some of the respondents of the collaborating company describe a persisting social stigma associated with using AI. For example, R3 noted that, "*sometimes when it is obvious that that someone has used AI, it is still mentioned in a negative way*". This suggests that norms around AI use remain ambiguous: while AI may be accepted in principle, it is not yet consistently viewed as legitimate across all tasks and contexts. Notably, this is despite an observed organizational push for AI adoption and multiple respondents reporting that they feel encouraged by departmental managers to use AI. R25 illustrate how these social dynamics can translate into individual hesitation, describing a feeling of reluctance to disclose AI support in certain tasks due to anticipated negative reactions, such as people interpreting it as "cheating" (R25). R25 connects this to an internalized sense of shame:

"I might not always want to disclose that I have used AI, even if everyone knows, I try to be very conscious about being transparent when I use AI but I still have some trouble with how to formulate it, and I think that is because I still have some internal shame about using AI. I can talk about how great AI is, but I am still ashamed" (R25).

Interestingly, stigma was not only described as shame associated with using AI, but also as discomfort associated with not using it. R19 suggest that transparency about using AI is not the main concern, rather the respondent sometimes feels uneasy about being unable to claim more frequent use of AI because this could signal that they are not "*an AI user*" (R19). This divergence suggests that norms surrounding AI are currently in transition where some individuals may still associate AI use with negative perceptions and others experience the opposite pressure, where not using AI may signal lack of adaptability.

5

Discussion

This section first discusses the gap between the theoretical potential of AI in patent management and the current implementation in practice, as well as the value that some respondents have realized from using AI. Furthermore, the previously identified barriers of adoption are discussed by connecting the empirical findings to existing literature on AI adoption in organizations. Finally, based on the discussions, six recommendations regarding how organizations can support the integration of AI in patent management are presented.

5.1 The gap between the potential of AI and practical implementation

The findings show that AI has significant theoretical potential in patent management and the identification of tools suggests that there are available AI tools across all processes. Respondents furthermore expressed strong belief in the potential of AI in patent management and stated that they had tested AI tools in many of the activities in this study. Despite this, the findings show that there is a gap between the theoretical potential of AI and how it is currently applied in practice.

Prior art search appeared to be the activity with the clearest evidence of established use of AI across all respondent groups, where respondents reported benefits in terms of time savings and improved search quality. Overall, this indicates that the potential of AI is realized to a great extent in this area for all respondent groups. In addition, there were also indications from each of the respondent groups of realized value from AI usage in the activity validity analysis.

Furthermore, the activities of drafting, prosecution, FTO, legal action & commercialization are activities in which the collaborating company did not show clear signs of AI use, while indications of practical implementation were identified among other respondent groups. In drafting and prosecution, IP firm respondents described that supportive use of AI resulted in substantial time savings. Findings also indicated that there are external IP departments that already experience efficiency gains in these activities. In addition, a respondent from an external IP department expressed that the use of AI search tools as a first step before human control in FTO saves time and improves search results. Thereby, for companies conducting patent management

activities that have not yet implemented AI in the above mentioned activities, the findings suggest that the potential is not merely theoretical. Rather, the findings show already active use of AI and resulting efficiency gains, indicating that such value can be realized in practice.

The findings of this study indicate that the value currently realized in patent management activities is primarily derived from the use of AI as a supportive tool rather than from full automation. Examples included language improvement and generating initial ideas for drafts, AI-supported searches before human review and summarization for time savings and easier communication. In other words, AI was only used to enhance and complement professional expertise rather than conducting a full task. This type of use of AI aligns with what Calvino et al. (2025) refers to as AI augmentation, where AI is used in a collaborative way with workers rather than to replace them. According to Calvino et al., such augmentation can improve productivity and free up time for more complex tasks. For example, it is described that generative AI can be used in an augmenting way to support data analysts through automation of initial data processing and initial analysis, which can enable more time for interpreting results and finding meaningful insights (Calvin et al., 2025). In this context AI augmentation has been shown to result in time savings which in line with Calvin et al. can enable professionals to focus more on strategic and analytical tasks. This is particularly relevant for patent management, where IP is increasingly viewed as tool for value creation rather than solely for protection, which calls for a more strategic and offensive approach to IP management (Holgersson et al. 2018). In relation to this development, the findings of this study suggests that even supportive use of AI can be strategically important, as the identified efficiency gains can allow more time to be allocated to analytical and strategic patent management activities. Organizations that do not realize such opportunities risk lagging behind competitors that are increasingly using AI to enable more time for strategic patent management work.

One contrasting finding was that some respondents also described using AI in more automation-oriented ways. This included generating large volumes of patent drafts, as well as automatically generating oppositions to invalidate patents and deter competitors. Nevertheless, these forms of automation were primarily viewed as valuable only in specific contexts where speed and volume are prioritized over quality. This suggests that AI-driven automation may create value under certain conditions, but that the relevance remain limited when higher quality outputs are required. However, the common perception among respondents seems to be that AI is primarily viewed as an enhancing tool rather than a tool for full automation given the current state of the technology.

Furthermore, for activities such as portfolio management, intelligence and infringement search and analysis, the findings provided limited indications of practical application of AI across respondent groups, apart from sporadic and limited use at the collaborating company. However, this does not necessarily mean that AI is not used within these activities, as some respondents may simply not have explicitly mentioned such usage during interviews.

Observations from this study suggests that use of AI within both organizations and processes may be more extensive than reflected in the findings. This may partly be explained by the semi structured nature of the interviews, but also by the fact that the level of AI use appeared to vary significantly between individuals, both within the organization and among individuals working with the same activities. Consequently, the experiences described by individual respondents may not fully capture the overall extent to which AI are used within a specific activity or process. The observed variation in AI adoption may indicate a lack of shared organizational structures or processes related to AI.

Overall, the findings of this study show no evidence that the theoretical potential of AI has been fully realized for any activities or respondent groups. Even in activities where some realized value was evident, the theory still suggests that there exists additional potential that remains unrealized. For example, at the collaborating company, AI summarizations and AI for limited specific tasks was used as support within portfolio management, while all of the remaining theoretical potential applications for which AI could be used for in this activity remains untapped. For example, AI could support in more accurate and dynamic portfolio valuation, support in pruning decisions and automate administrative tasks (Akhtar, 2025).

This pattern was evident across multiple processes. At the same time, the identification of tools in table 4.1 shows that many of the theoretically described potential applications of AI are reflected in existing AI supported tools, suggesting that at least part of this potential should be technically possible to realize in practice. However, the findings show that technical availability does not necessarily translate into actual implementation. A recurring topic during interviews was that several respondents expressed dissatisfaction with the quality and functionality of the tools tested, which was often described as a main reason for not implementing AI within the discussed activity. As this study does not include assessment of tool functionality or performance, it is not possible to determine to what extent the limited realization of theoretical potential of AI reflects limitations of tools. At the same time, the rapid development of AI technologies suggests that opportunities which may not be fully realizable today could be possible to realize in the very near future. Findings also show that the assessment of tool performance is subjective, as respondents described different experiences and perceptions regarding usefulness of the same tools, which confirms that the gap between potential and implementation can not be solely explained by tool performance. In addition to this, the theoretical potential largely aligned with the potential applications expressed by respondents regarding how AI could be used in patent management. Although this does not confirm that all theoretical potential applications of AI can be realized in practice, it indicates that professionals within the patent management field consider these applications possible in the near future.

However, when examining how the collaborating company and other external IP departments included in the study currently use AI, it is important to note that many IP related activities are often partly or sometimes fully outsourced to IP firms. This may explain the limited evidence of AI use among the IP departments

in this study. IP firms appear to have implemented AI more extensively in certain tasks such as drafting and prosecution compared to the collaborating company. This can be explained by the fact that it is likely that the incentive to adopt AI tools within IP departments that outsource parts of its patent process is weaker than for IP firms, since tasks such as drafting may not be core internal responsibilities for all of the investigated IP departments.

In relation to this, it is interesting to consider how AI may reshape the domain, as indications from the interviews suggest that some respondents questioned whether some activities could potentially be transferred in-house with the help of AI. In other words, if AI could be a potential enabler for internalizing activities that are currently outsourced. A comparable example can be observed in the historical transition from mainframes to minicomputers. During the 1960s, it was common for companies to outsource information-processing activities to so-called service bureaus that specialized in developing application software. These service bureaus adopted innovations in both hardware and software and developed programs tailored to business needs, for example payroll systems. Rather than organizations running the programs by themselves, the service bureau sold the activity of information-processing to the companies. By relying on these bureaus, firms could avoid the substantial costs associated with both software and hardware during this period (Steinmueller, 1995). This can be compared to how IP departments today often outsource certain activities within the patent process to improve cost efficiency. Although hardware is not the primary expense in this case, the personnel costs associated with maintaining an IP department large enough to manage the entire application process in-house often still exceed the cost of outsourcing.

For the firms relying on service bureaus, this outsourcing arrangement became less necessary as computing technology became more affordable and accessible. The introduction of minicomputers created entirely new opportunities for companies to purchase and operate their own computers instead of outsourcing information processing to service bureaus, enabling them to conduct tasks such as automating payroll activities themselves (Steinmueller, 1995).

This historical pattern provides a useful lens for interpreting the findings of this study. From this perspective it is reasonable that IP departments with a high degree of outsourcing currently do not experience strong pressure to invest in AI tools for tasks they do not perform internally. For these organizations, the most plausible near-term scenario is that AI adoption in currently outsourced tasks will primarily occur within external IP firms that already perform those activities. Moreover, the perceived value of AI for activities such as drafting may currently appear lower from the perspective of corporate IP departments because internalizing these activities would require substantial organizational changes. Bringing outsourced activities in-house would not only require new workflows and capabilities, but could also initially increase personnel demands, as organizations would need to absorb an additional activity into their internal operations. As a result, the potential benefits of AI adoption may not outweigh the organizational costs associated with such a transition. Nevertheless, the longer-term trajectory may depend on whether AI lowers the bar-

riers to performing these activities internally. If so, patent management could follow a similar path to other domains where technological advancements first introduced adoption among specialized providers but ultimately enabled a reconfiguration of responsibilities toward in-house capabilities.

Taken together, the findings indicate that the current use of AI among the respondents in this study remains more limited than the theoretical potential. Although some respondents reported practical applications of AI and resulting benefits, these applications were only a fraction of the potential applications suggested in theory. This suggests that a substantial part of the possible benefits AI may bring to patent management remains untapped. At the same time, respondents demonstrated a clear understanding of the potential of AI and expressed interest in increased use. Despite this, the implementation of AI still remained limited, which suggests that there are factors that make the transition towards increased use of AI within patent management processes difficult.

5.2 Understanding the barriers of AI adoption in patent management

The following section discusses the identified barriers of adoption in relation to existing literature on AI adoption in organizations. By doing this, the aim is to explain the underlying organizational mechanisms behind these barriers in the specific context of patent management and provide insights into how they can be addressed in practice. To structure the discussion, the findings are analyzed through the key overarching categories of obstacles related to AI adoption in organizations identified by Li et al. (2025): **people’s readiness, process and politics**. In addition, selected aspects of Rogers’ (1983) theory on perceived attributes of innovations that affect diffusion are included to further explain how characteristics of AI influence adoption dynamics related to people’s readiness.

From the **people’s readiness** perspective, Li et al. (2025) identify uncertainty, fear of replacement and self-image concerns as central problems affecting employees’ willingness to adopt AI. The findings of this study indicate that uncertainty appeared repeatedly across several aspects of AI adoption in patent management. More specifically, the findings suggest uncertainty related to understanding how AI currently creates value, uncertainty regarding the trustworthiness of AI-generated outputs, and uncertainty regarding the conditions for responsible use.

The first form of uncertainty identified in the findings relates to understanding how AI currently creates value in practice. Li et al. (2025) describe uncertainty as a condition in which limited knowledge leads to polarized perspectives, where some dismiss AI as a temporary hype while others overestimate its capabilities. This pattern is evident in relation to *The search for the perfect tool as a barrier of adoption*. The findings indicated that rejection of AI could often be attributed to limited interaction or opinions shaped by external narratives rather than direct experience. This connects closely with a lack of knowledge, where in this context,

the lack of understanding of how AI generates value today, and how value is only realized through experimentation and iterative use, results in skepticism that limits adoption.

This can further be related to Rogers' (1983) diffusion of innovation theory. Rogers (1983) explains that the trialability of an invention, meaning the degree to which it can be experimented with on a limited basis influences adoption positively. This may explain why those who had experimented with AI for longer periods of time were also those who had adopted to the greatest extent. This is particularly relevant in relation to the findings highlighting that many IP professionals lacked sufficient time to experiment with AI tools. In this case, the limited trialability is not connected to the characteristic of the innovation itself, but the lack of time to experiment with it restricts users progressing from initial trials. Rogers (1983) further argues that observability of an invention, meaning the degree to which its results and advantages are visible and communicable to others influences adoption. In this context, *The search for the perfect tool as a barrier of adoption* suggested that the value generated by AI within patent management is not immediately observable, but instead emerges gradually through iterative interaction. When users do not engage with the tools over longer periods of time, their practical benefits remain difficult to observe, contributing to skepticism and abandonment after isolated negative experiences.

The second form of uncertainty identified in the findings relates to the trustworthiness of AI-generated outputs. This pattern was evident in relation to *Lack of trust as a barrier of adoption*. While some respondents frame hallucinations as a critical risk and reason to avoid adoption, others argue that such errors are not fundamentally different from those already managed in the patent process. This divergence reflects differences in how the risks of AI are interpreted, where some view them as prohibitive, while others consider them manageable through existing quality controls and keeping a human-in-the-loop. This finding can further be related to what Ransbotham et al. (2025) explain, they argue that organizations should not view autonomous AI systems as fully predictable tools, but rather as organizational actors that require supervision comparable to that applied to human employees. From this perspective, AI-generated outputs should not necessarily be interpreted as evidence that the technology is unusable, but rather as a challenge of governance and oversight.

At the same time, the findings indicated that patent management processes already operate under uncertainty and already contain review-based quality assurance. Consequently, the challenge for patent management functions may not primarily be to establish completely new governance structures for AI, but rather to extend existing governance logic to AI-supported output and create a more uniform understanding that risks such as hallucination can be managed. Such an understanding may reduce the polarized opinions by legitimizing AI-supported work as compatible with existing governance practices, which could strengthen the organizational trust in AI. However, the findings of this study also suggest that uncertainty related to trustworthiness of AI within patent management contexts is not only connected to error management. Several respondents described that trust in AI systems was further

complicated by their limited explainability. This is particularly important in patent management, where professional legitimacy depends not only on producing correct outcomes, but also on being able to justify, trace and defend legal and technical reasoning throughout the patent process. Consequently, the lack of explainability makes it difficult for practitioners to confidently rely on AI, even where review mechanisms to catch errors are present. This also helps in explaining why respondents consistently stressed the importance of maintaining a human-in-the-loop.

The third form of uncertainty identified in the findings relates to uncertainty regarding the conditions for responsible use. Li et al. (2025) explain that uncertainty becomes particularly problematic when organizations lack clarity regarding responsible use. They explain that when risks such as regulatory or compliance-related risks are present, the uncertainty and lack of knowledge can result in resistance of adoption. This dynamic is reflected in the findings through *Confidentiality and security requirements as a barrier of adoption*. In patent management the uncertainty element is intensified due to the strict security requirements associated with intellectual property. The findings suggested that even respondents who were highly positive towards AI refrained from adoption when security guarantees could not be established, indicating that the rejection is not of AI itself but a lack of clear conditions for its use.

This may suggest that patent management exhibits adoption patterns similar to those commonly associated with domains where AI-systems are classified as high-risk, such as healthcare, finance or law (Bartsch et al., 2025). Although the risks associated with patent management differ from those in domains such as healthcare, as they do not directly involve human life, the findings still demonstrate a high level of risk awareness. In patent management, disclosure of confidential information before a patent is granted may reduce the likelihood of obtaining patent protection, while also exposing important innovations to competitors, potentially affecting an organization's operations and long-term competitiveness. This points to similar underlying mechanisms related to risk sensitivity and the need for clear guidance on appropriate AI use as in high-risk domains. Previous research in healthcare has shown that confidentiality concerns are among the most significant barriers to AI adoption. In addition, multiple studies have shown that clear guidelines for how AI should be implemented in healthcare are often absent, and that the lack of regulations of AI use negatively affects healthcare professionals' attitudes toward AI (Abdelwanis et al., 2026). This reflects a pattern similar to that identified in this study, where concerns regarding confidentiality, combined with a lack of clear direction for secure use of AI influence perceptions of the technology, and ultimately limit its adoption.

This can also be understood through Rogers' (1983) concept of compatibility, which describes how innovations diffuse slower when they are perceived as inconsistent with existing values, past experiences and needs. Because protecting confidential information is a central norm within patent management, the idea of using AI, especially public tools, may therefore be perceived as conflicting with professional responsibilities, even when efficiency benefits are recognized. This may help explain

why respondents who were otherwise highly positive toward AI still refrained from adoption, and suggests that successful implementation in IP departments requires clear conditions for responsible use. Employees need to understand how AI can be used without violating confidentiality requirements so that uncertainty regarding acceptable use does not become a barrier to adoption.

Bringing these three forms of uncertainty together, Li et al. (2025) suggest that organizations can reduce uncertainty through clear governance structures and making the use of AI intuitive for employees. In the context of patent management, the findings of this study indicate that existing review and quality assurance mechanisms already provide an important foundation for managing incorrect output. This means that organizations may not necessarily need to establish entirely new governance structures for AI-supported work, but rather extend existing governance logic to AI-supported outputs. At the same time, Li et al.'s (2025) argument about making AI use intuitive for employees remains essential. Employees need practical clarity regarding how AI can create value today and that this value will only materialize with experimentation and iterative use. In addition, they need guidance on how AI can be used responsibly without increasing confidentiality risk.

However, the findings also indicated that governance structures and increased familiarity alone may not fully resolve uncertainty within patent management contexts, where explainability is a requirement in many situations. While ensuring that employees act as competent operators, and clearly defining in which tasks AI can be safely applied may partly reduce this challenge, explainability remains a practical concern. While the limited explainability of tools was shown to limit adoption among respondents, current developments in AI indicate that these limitations will soon be reduced. One such development is Explainable Artificial Intelligence (XAI), which refers to AI systems that are designed to provide users with explanations of how the model reached the decision, thereby increasing transparency. Recent advancements in XAI already show increased applicability across industries (Mohamed et al., 2025). Consequently, limited explainability should not necessarily be interpreted as a reason to refrain from AI adoption, especially considering that the findings of this study highlighted that organizations need to spend time on iterative learning. From this perspective, avoiding AI adoption until explainability challenges are resolved will limit the preparedness for future AI developments.

Beyond uncertainty, Li et al. (2025) further identify fear of replacement as one of the problems in relation to people's readiness. However, direct concerns regarding job displacements were not strongly reflected in this study. The respondents did not frame AI as a direct threat to their job. Instead, they generally described it as a tool that needs to be under human control at its current maturity level. Rather than replacing roles, respondents perceived it as an opportunity to optimize tasks and shift effort towards verification and review activities. This may be partly explained by accountability requirements and professional norms in patent management, where responsibility must ultimately remain with a legally accountable person. Because liability cannot be delegated to an AI system, full automation is perceived as difficult, thereby reducing the perceived feasibility of near-term replacement.

Although fear of replacement was not evident, one mechanism associated with this fear, described by Li et al. (2025) as a fault-finding tendency where AI is held to higher standards than humans, which ultimately leads to resistance was clearly evident in *The search for the perfect tool as a barrier of adoption*. It was observed that some respondents expected AI to perform above human levels. Ransbotham et al. (2025) further highlight this problem and describe that organizations often evaluate AI errors differently from human errors. While mistakes made by employees are seen as learning opportunities, errors generated by AI systems are more likely to be interpreted as defects. This tendency was also reflected in the findings of this study, where respondents described how mistakes are generally accepted as an inherent part of human behavior in patent management but that similar tolerance is not extended to errors produced by machines.

Interestingly, the findings in this study actually extend Li et al.'s (2025) discussion of how fault-finding affects AI adoption. The findings suggest that this fault-finding behavior has implications beyond disproportionate expectations leading to resistance. It highlights that fault-finding also shapes subsequent actions where the result is a self-reinforcing mechanism where individuals, after limited trials or due to underlying negative beliefs, resist adoption and thereby prevent the learning process and experimentation through which AI's value often emerges.

The last problem in relation to people's readiness is the self-image problem which relates to a fear of loss of status that can result in employees hiding their AI use because of fears of appearing less competent, lazy or dishonest (Li et al., 2025). This mechanism is clearly observed in the findings relating to *Inertia and social stigma as a barrier of adoption*. Interestingly, in the collaborating company the self-image problem was two-sided. It was visible not only as reluctance to disclose AI due to fears of negative judgment and shame but also in discomfort associated with disclosing not using AI, as it may signal lack of adaptability. This indicates that employees are operating within unclear and evolving norms regarding what appropriate AI use is. This problem can also be found in Rogers' (1983) description on relative advantage. Rogers (1983) explains that innovations are adopted faster if users perceive clear advantages in relation to existing ways of working, and suggests that one of the most important motivations for adoption is the desire to gain social status.

What is particularly interesting is that neither Li et al. (2025) nor Rogers (1983) fully explain the pattern evident in the findings. In this case, it was unclear whether adopting AI would increase status by signaling adaptability and competence, or decrease status by signaling laziness or incompetence, as both perceptions were evident within the same organization. This ambiguity has important implications. When both using and not using AI carry potential reputational risks, employees are left without a stable reference point for behavior. As a result, adoption becomes inconsistent and depends on individual interests rather than organizational direction. This can lead to fragmented use patterns, where some employees adopt AI extensively while others do not. This was evident at the collaborating company where people who had a personal interest in AI had adopted, while those who did not

were lagging behind. This leads to a setting where shared learning becomes difficult because everyone is at different levels. On top of that, the hesitation to disclose AI use may restrict opportunities of knowledge-sharing.

At the same time, Rogers' (1983) concept of relative advantage helps explain the inertia identified in the findings. Several respondents described feeling sufficiently effective in their current ways of working and therefore did not perceive any clear advantage in adoption. This can also be interpreted through Rogers' (1983) concept of compatibility, where one factor influencing adoption diffusion is the degree to which an innovation is compatible with users' existing needs. From this perspective, employees who already feel comfortable in their current ways of working may not perceive a need for AI integration. However, Rogers (1983) also argues that needs for an innovation are not always recognized prior to gaining sufficient understanding of the innovation itself. As discussed previously, understanding the practical value of AI within patent management often requires interaction with the tools over longer time periods. It is therefore unsurprising that respondents who had limited experience with AI struggled to identify any strong need for adoption.

Li et al. (2025) suggest that the self-image problem can be addressed by actively reshaping how AI use is perceived, positioning it as a factor of professional competence rather than a threat to it. In the context of the collaborating company, this could help reduce the negative social stigma surrounding AI use. However, the findings of this study suggest that this only addresses one side of the problem. While promoting AI use may reduce stigma for some employees, it may simultaneously increase pressure on those who already feel discomfort regarding their limited use. The findings therefore suggest that the collaborating company should first focus on reducing the ambiguity surrounding AI, creating a uniform perception of AI use as something positive, but not a determinant of your competence or adaptability.

In relation to the inertia identified in the findings, Rogers (1983) argues that incentives may increase the perceived relative advantage of an invention and thereby encourage relative change. However, as discussed previously, understanding the practical value of AI within patent management often requires experimentation and use over longer periods of time for the value to materialize, or in this sense to establish a need or relative advantage. Employees may therefore first need opportunities to explore how AI can create value within their own work practices before stronger adoption expectations or incentive structures are introduced.

Having discussed the people's readiness perspective, the attention now switches to **process** factors. Li et al. (2025) emphasizes that successful adoption requires redesigning of workflows rather than simply integrating AI into existing processes and distinguish between three connected levels, individual (nodes), interactions (edges) and the overall system(network).

From the nodes perspective, Li et al. (2025) describe redesigning of how work is executed, for example by allowing it to handle standardized elements while shifting human efforts towards tasks that require expertise. This means that users must

adapt their ways of working to collect the value that AI offers. This is interesting in terms of what became evident through what is described in *The search for the perfect tool as a barrier of adoption* previously discussed in relation to fault-finding behavior. The findings showed that rather than adapting tasks in accordance with what AI actually can do today, AI tools were abandoned when performance did not meet expectations, or from this perspective, did not fit into the current ways of working. This could imply that, in addition to lowering expectations on performance, adaptations of tasks may be required to visualize the value of AI.

Li et al. (2025) suggest that this problem could be handled by forcing process change through strong performance pressure and high targets. However, in this context such an approach may be counterproductive. As became evident in *The fear of losing competence and professional judgment as a barrier of adoption*, efficiency pressure could result in reduced critical review which is of particular importance in domains such as patent work. On the other hand, if tasks are rearranged, for example by letting AI assist in tasks that are not critical or changing the workflow so that AI does groundwork and a professional does the review, introducing somewhat harder performance targets on those tasks could be an option worth investigating.

At the edge level, Li et al. (2025) describe how AI can transform the connections between tasks by improving how information flows and how decisions are coordinated. In this respect, the findings indicate a positive outlook. As described in *Introduction of bottlenecks as a barrier of adoption* it was explained that AI can enhance communication between roles. For example, inventors can use AI to understand feedback from patent attorneys and respond with a more refined input which might speed up the end-to end process. On the other hand, from the network perspective, Li et al. (2025) describe that improvements at one level must be aligned across the others to generate system-wide value and avoid shifting bottlenecks. The findings clearly demonstrated this issue where AI implementation in one part of the organization has already introduced bottlenecks in another. While AI increased efficiency in upstream activities, such as generating IDFs, it simultaneously introduced additional complexity in downstream roles. In line with the argument of Li et al. (2025), this indicates that individual nodes are optimized without aligning the broader network.

Li et al. (2025) suggest that network challenges should be handled by coordinated action across interconnected tasks. The entire network should be mapped to allow synchronization of capacity improvements across tasks and throughout the workflow. In the context of the collaborating company this implies that AI use in upstream processes such as IDF drafting should be evaluated in connection to its downstream impact. Without mechanisms to balance workload distribution, adoption risk introducing inefficiencies and discouraging downstream roles from remaining positive towards AI adoption. In addition, incentives across the full workflow should be aligned, and guidelines for how AI should be used in different stages should be established. This could involve defining quality standards for AI generated content acting as input for a subsequent processes such as specifications defining IDF content and format.

The discussion now moves to what Li et al. (2025) describe as the **political** dimension where the effects of power dynamics and incentives are described. First, Li et al. (2025) describes the problem of resource hoarding where individuals withhold knowledge to preserve their own advantage. This dynamic is not evident in the findings of this study. While limited knowledge sharing was identified as a risk, this was more in relation to reluctance to disclose AI use, which was discussed previously in connection to the self-image problem.

Second, Li et al. (2025) describe the risk of hierarchy disruption, where AI challenges existing structures of expertise and authority. This risk was not directly observed in the collaborating company, but elements of it were identified in insights from external IP firms and experienced professionals in the findings related to *The fear of losing competence and professional judgment as a barrier of adoption*. It became evident that concerns about declining expertise may lead to senior practitioners discouraging AI use among less experienced colleagues. This suggests that such dynamics exist within the broader patent domain.

Li et al. (2025) propose that these challenges can be addressed by changing how competence is valued, for example by recognizing AI-related skills as a part of professional expertise. In the context of patent management this could mean framing AI experience as complementary to technical and legal expertise. This would ensure that adoption does not threaten established professional identities.

Even though the hierarchal dynamics were not observed as an issue in the collaborating company in this study, an approach like this would still be beneficial. The findings indicated there are existing concerns about how increased reliance on AI could reduce critical thinking within the collaborating company. By positioning AI as a tool that supports existing expertise rather than replacing professional judgment, these concerns could be weakened. In addition, such an approach will support the development of what respondents described as “competent operators” in relation to *Lack of trust as a barrier of adoption*, where the findings identified combining domain specific knowledge with AI expertise as a key condition for capturing value of AI in patent management.

Finally, Li et al. (2025) highlights the problem of accountability attribution where introduction of AI could disrupt responsibility attribution and the balance of blame because of the increased traceability of decisions and actions, making it easier to identify who or which department did what. This becomes a problem because it can create tension and may expose individual contributions to an extent that is not preferable. However, this problem was not raised in the findings of this study. This can be explained by the scope of this study where the analysis focuses on a single department and examines how AI is used to support task execution, rather than on how it is used to monitor or track performance across workflows. In addition, findings in *Lack of trust as a barrier of adoption* suggested that accountability structures are likely to remain intact at the current state of AI adoption. In patent work, responsibility structures are already well defined, and the findings emphasize maintaining a human-in-the-loop approach and stressed that responsibility in this

domain remains with humans regardless of the tools involved.

Bringing this together, the discussion highlights that the barriers of AI adoption in patent management are shaped by the characteristics of the domain itself, where confidentiality requirements, professional accountability and a risk-aware culture influence how AI is perceived and integrated into practice. The discussion further shows that these barriers are not isolated, but rather strengthen each other, suggesting that successful adoption depends on how organizations manage these constraints in combination.

5.3 Practical implications

To support organizations in addressing the identified barriers of adoption, the authors propose the following recommendations for integrating AI into patent management processes.

Establish guidelines

The authors recommend that organizations establish guidelines for responsible AI use within patent management processes. These guidelines should specify which AI tools are approved for different purposes, how AI-generated output should be critically reviewed and verified, and what constitutes appropriate AI use.

Establishing guidelines addresses *confidentiality and security requirements as a barrier of adoption* by clarifying which tools are approved for handling patent-related information, thereby reducing the actual confidentiality risks and the uncertainty regarding acceptable AI use.

Furthermore, establishing guidelines addresses *Lack of trust as a barrier of adoption* by establishing mandatory verification of AI-supported work. While hallucinations and limited explainability contribute to reduced trust in AI systems, requiring employees to critically evaluate AI-generated output will reduce the risks associated with these limitations. This will help reduce the perception that AI-supported work introduces fundamentally uncontrollable risks, while supporting employees in acting as competent operators.

Finally, establishing guidelines reduces *Fear of losing competence and professional judgment as a barrier of adoption* by requiring employees to critically evaluate AI-supported work. This will reduce concerns regarding overreliance on AI by clarifying that AI can be used responsibly without abandoning critical examination or human judgment, provided that employees follow established guidelines.

Provide education

The authors recommend that organizations should provide education on: the value that AI can bring today, how AI systems work, how AI can be used responsibly, and training in AI-related skills relevant to patent management work.

Providing education regarding the value AI can generate at its current maturity level addresses *The search for the perfect tool as a barrier of adoption* by reducing unrealistic expectations, while highlighting that AI may still generate value despite not performing perfectly. In addition, providing education on how AI systems function together with training in AI-related skills will help address *Lack of trust as a barrier of adoption* by supporting employees in developing the skills required to act as competent operators. An increased understanding of how AI systems function will also assist in establishing trust. Finally, education on responsible AI use addresses *Confidentiality and security requirements as a barrier of adoption* by helping employees understand how to use AI without violating requirements. This reduces uncertainty and the risk of unsafe practices.

Encourage experimentation

The authors recommend that organizations encourage and enable experimentation with AI tools. This involves allocating time, providing access to AI tools, and offering opportunities for hands-on interaction, such as collaborative forums, workshops, or delegation of tasks with encouragement of completing them with AI.

Experimentation addresses *The search for the perfect tool as a barrier of adoption* by allowing employees to discover the value that AI actually can provide at its current maturity. It will also address *Lack of trust as a barrier of adoption*, as experience is essential for developing the necessary skills to become competent operators. Having the skills of a competent operator will help employees manage limitations such as hallucinations and explainability, which currently contribute to reduced trust in AI systems. In addition, experimentation helps overcome *Inertia and social stigma as a barrier of adoption* by reducing resistance to change by revealing new needs or use cases. Firsthand experience with AI's benefits increases motivation to adopt it.

Although the authors recognize that organizations face limited time and competing priorities, avoiding experimentation with AI will simply postpone the issue while increasing the risk of a widening knowledge gap as the technology continues to develop. Furthermore, organizations that delay experimentation risk falling behind competitors that are already building capabilities and experience. Delayed experimentation will also result in missed opportunities to capture the value that existing AI tools can already provide.

Establish shared organizational norms around AI use

The authors recommend that organizations establish shared norms for AI use. Actions to do this may include identifying AI champions, recognizing AI-related skills as valuable, and organizing collaborative activities to discuss and shape views around appropriate AI-supported work.

By reducing ambiguity regarding how AI use is perceived and by establishing a more stable reference point for behavior, this recommendation addresses *Inertia and social stigma as a barrier of adoption*. By creating a more uniform perception of AI use as something positive, but not a requirement for professional legitimacy, organizations

will reduce both reluctance to disclose AI use, and discomfort associated with limited AI use. Establishing AI-related skills as complementary to professional expertise rather than a replacement for it will also reduce concerns related to *Fear of losing competence and professional judgment as a barrier of adoption*.

In addition, reducing social stigma will create better conditions for knowledge sharing and organizational learning, which will support overcoming both *The search for the perfect tool as a barrier of an adoption* and *Lack of trust as a barrier of adoption*. Observing how colleagues successfully create value through AI will help employees understand how AI can improve their own work, while also increasing exposure to practical examples of trustworthy AI use.

Map and coordinate AI activities across interconnected tasks

The authors recommend that organizations map and coordinate AI activities across interconnected tasks. This involves evaluating how AI adoption in one workflow stage affects subsequent tasks, aligning incentives across roles, and defining standards for the structure, relevance, and quality of AI-generated materials used in later stages.

This recommendation addresses *Introduction of bottlenecks as a barrier of adoption* by reducing the risk that local efficiency gains create additional workload or complexity downstream. Coordinating AI adoption helps maintain positive attitudes by preventing situations where one role benefits while another absorbs extra work.

Develop objective models for evaluating AI tools

The authors recommend that organizations develop objective models for evaluating AI tools. Structured evaluation models will reduce the tendency to evaluate AI tools based on subjective perceptions or isolated negative experiences that may not accurately reflect their practical ability to generate value. This will also reduce the fragmented perceptions of AI across the department.

This recommendation addresses *The search for the perfect tool as a barrier of adoption* by establishing clear criteria for what constitutes a valuable AI tool. This will reduce fault-finding behavior and premature dismissal based on isolated failures or subjective perceptions. Structured evaluation approaches will also make assessing AI tools more manageable by providing systematic criteria for comparing solutions, especially as the number of available tools increases.

Furthermore, this will also address *Lack of trust as a barrier of adoption*, as tools that have been evaluated based on predefined and established criteria can be perceived as more reliable and trustworthy by employees. Additionally, it will also address *Confidentiality and security requirements as a barrier of adoption* by reducing the risk that inappropriate tools are implemented within the organization. It can also reduce employees fear that using the tool may result in disclosure of sensitive information, both because the tool has been formally evaluated and because responsibility is shifted to the functions responsible for choosing and approving the

tool.

To end this discussion, the authors would like to stress the importance of adopting an iterative and adaptive strategy for AI adoption. Continuous refinement ensures organizations remain prepared when more capable and affordable models enter the market and make current concerns about hallucinations, stigma, and finding the perfect tool less relevant.

This implies that the recommendations presented should not be treated as static implementation activities. Instead, organizations should continuously educate employees on new AI capabilities and use cases, update guidelines and verification requirements as AI systems become more capable, continue encouraging experimentation, and regularly evaluate new tools. Ongoing experimentation is critical for keeping pace with rapid AI developments and ensuring that organizations can effectively translate AI potential into tangible value. This becomes particularly relevant in relation to the ongoing development toward increasingly agentic AI systems. As AI systems become more capable of reasoning and differentiating between information sources, today's concerns about hallucinations and trustworthiness will become less pronounced (Purdy, 2024). Organizations that postpone adoption risk missing out on current opportunities and falling behind competitors in competence, experience, and preparedness for future AI advancements.

6

Conclusion

To fulfill the purpose of this thesis, the study examined AI applications in patent management by investigating the potential of AI across patent-related processes, the current realization of this potential, and the organizational barriers influencing adoption.

To answer the questions of what the potential applications of AI in corporate patent management processes are and how this potential is realized in practice this study concludes that AI has significant potential within corporate patent management processes. This potential is not only supported by literature, but also by identified AI tool offerings and by the perceived potential expressed by professionals in the field. Much of the perceived potential was related to how AI could facilitate time consuming activities and enhance analytical work within patent management.

At the same time, the findings show that the current use of AI still only reflects a limited part of this potential. The clearest evidence of established use of AI and realized value was within prior art search, where AI contributed to time savings and improved search quality. The same benefits were also found in the drafting and prosecution activities when AI was used in a supportive way. Beyond these areas, identified use of AI were connected to smaller tasks or sporadic use. However, even limited use appeared to generate value, mainly in terms of time savings. Still, the practical applications of AI identified in this study only reflect a fraction of the theoretical potential. This suggests that much of the value that AI could bring to patent management remains unrealized.

To answer the question on which organizational barriers of adoption that exist in relation to integrating AI into corporate patent management processes this study identified six barriers of adoption, including confidentiality concerns, lack of trust, psychological inertia and shame, fear of losing competence and professional judgment, introduction of bottlenecks, and waiting for perfect AI tools.

This study showed that the barriers of AI adoption in patent management are connected to the characteristics of the patent domain itself. Patent management is a risk-sensitive environment characterized by confidentiality and accountability requirements, low tolerance for mistakes, and professional norms surrounding expertise. Organizations appeared to place substantially higher demands on AI systems

than on human professionals, despite acknowledging that human experts also make mistakes. At the same time, this study showed that the value of AI is often not immediately visible, since effective use requires experimentation, iterative learning, and adjustments of workflows and ways of working.

However, the same organizational conditions that create a need for AI, such as high workload, complexity, and pressure for quality, simultaneously limit the time and flexibility required for organizations to experiment and adapt to the technology. Taken together, this thesis concludes that the challenge of integrating AI into patent management is not primarily about whether AI has potential, but about organizations' ability to adapt in ways that allow this potential to be realized without compromising the requirements and professional standards under which patent management operates.

6.1 Limitations and future research

First, this study was limited to semi-structured interviews, which enabled nuanced insight into AI use across different IP management activities. However, when split across activities and respondent groups, the number of observations remained limited as not all respondent groups discussed each activity. This in turn limits the depth across different parts of the IP management process. Future research could therefore build on these findings through for example structured interviews or survey based methods, in order to collect a larger number of data points for each activity and respondent group. Together, more data per activity and respondent group could enable a better comparison across activities and organizations.

Second, it was evident in the study that respondent expressed concerns about quality and maturity of current AI tools, although this study did not include a technical evaluation of the tools. The findings in this study therefore reflect organizational perceptions rather than technological performance. Future research could build on these findings by assessing the technical quality of the tools in relation to the intended use, to enable assessment of whether the available tools are technically mature to generate practical value and thereby realize the theoretical potential of AI in patent management.

Third, a recurring discussion during interviews was how the work flow within the patent field is going to be affected by the increased use of AI, more specifically the future role of IP firms. This study touched upon some of the possible future scenarios, for example whether AI could enable firms to perform more IP related activities in-house and thereby affect the current outsourcing patterns to IP firms. However, these developments were not empirically investigated within the scope of this study. Future research could therefore study whether increasing AI capabilities may change the balance between different actors in the patent field.

Fourth, as AI technologies are developing at a rapid pace, the findings of this study should be understood within the context of the current state of AI. This affects the

long-term relevance of the findings. Rather than focusing on the ongoing solution of current problems in the patent field, future studies could investigate how organizations can build structures and processes that enable continuous adaptation to emerging technological developments.

Fifth, the AI potential is identified in this study is based on both existing literature and the perceptions of professionals within the patent management field. Given the novelty and rapid development of AI, the available literature remains limited. In addition, respondents' perceptions of potential may be constrained by their limited AI expertise. Consequently, neither source may fully capture the potential of AI. Future research could therefore incorporate perspectives from AI specialists to provide a broader understanding of its potential.

Finally, the study touched upon findings suggesting that certain actors may benefit more from the current state of AI than others, particularly in contexts where volume and speed are prioritized over precision and quality. Examples included the use of AI to generate large volumes of patents and automatically generating oppositions to deter competitors. A future research suggestion could therefore be to study whether certain type of actors, business contexts or IP strategies are more likely to benefit from the current state of AI.

References

- Abdelwanis, M., Simsekler, M. C. E., Gabor, A. F., Sleptchenko, A., & Omar, M. (2026). Artificial intelligence adoption challenges from healthcare providers' perspectives: A comprehensive review and future directions. *Safety Science*, 193, 107028. <https://doi.org/10.1016/j.ssci.2025.107028>
- Agostini L, Nosella A, Holgersson M (2022-b), "Patent management: the prominent role of strategy and organization". *European Journal of Innovation Management*, Vol. 26 No. 4 pp. 1054–1070, <https://doi.org/10.1108/EJIM-09-2021-0452>
- Agostini, L., Nosella, A., & Teshome, M. B. (2019). Towards the development of scales to measure patent management. *World Patent Information*, 58, 101909. <https://doi.org/10.1016/j.wpi.2019.101909>
- Agostini, L., Nosella, A., Teshome, M. B., & Holgersson, M. (2022-a). What is patent management? The purification and validation of an integrated measurement scale. *IEEE Transactions on Engineering Management*. <https://doi.org/10.1109/TEM.2021.3138331>
- Akhtar, S. (2025). The role of artificial intelligence in enhancing patent lifecycle management. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*. https://www.researchgate.net/publication/390207842_The_Role_of_Artificial_Intelligence_in_Enhancing_Patent_Lifecycle_Management
- Al-Aali, A. Y., & Teece, D. J. (2013). Towards the (Strategic) Management of Intellectual Property: Retrospective and Prospective. *California Management Review*, 55(4), 15-30. <https://doi.org/10.1525/cmr.2013.55.4.15>
- Al-Mahdawi, E. (2025). From generative to agentic: How AI pes the future of work. *European Journal of Sports Science Technology*. <https://doi.org/10.58305/ejsst.v15i59.736>
- Bandi, A., Kongari, B., Naguru, R., Pasnoor, S., & Vilipala, S. V. (2025). The Rise of Agentic AI: A Review of Definitions, Frameworks, Architectures, Applications, Evaluation Metrics, and Challenges. *Future Internet*, 17(9), 404. <https://doi.org/10.3390/fi17090404>

- Bartsch, S., Behn, O., Benlian, A., Brownsword, R., Bücken, S., Düwell, M., Formánek, N., Jungtäubl, M., Leyer, M., Richter, A., Schmidt, J.-H., & Will-Zocholl, M. (2025). Governance of High-Risk AI Systems in Healthcare and Credit Scoring. *Bus Inf Syst Eng* 67, 563–581 (2025). <https://doi.org/10.1007/s12599-025-00944-4>
- Bell, E., Bryman, A., & Harley, B. (2019). *Business research methods* (5th ed.). Oxford University Press.
- Bell, E., Bryman, A., & Harley, B. (2022). *Business research methods*. (6th ed.). Oxford University Press.
- Bergmann, D. (n.d). *What is machine learning?* IBM. <https://www.ibm.com/think/topics/machine-learning>
- Bui, L. V. (2025). Advancing patent law with generative AI: Human-in-the-loop systems for AI-assisted drafting, prior art search, and multimodal IP protection. *World Patent Information*. <https://www.sciencedirect.com/science/article/pii/S0172219025000080>
- Caballar, R.D. (n.d). *Generative AI vs predictive AI: what's the difference?*. IBM Think. <https://www.ibm.com/think/topics/generative-ai-vs-predictive-ai-whats-the-difference>
- Calvino, F., J. Reijerink and L. Samek (2025), “The effects of generative AI on productivity, innovation and entrepreneurship”, *OECD Artificial Intelligence Papers*, No. 39, OECD Publishing, Paris, <https://doi.org/10.1787/b21df222-en>
- Chalmers University of Technology. (2024). *Behandling av personuppgifter*. <https://www.chalmers.se/utbildning/dina-studier/kandidat-och-examensarbete/behandling-av-personuppgifter/>
- Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., & Zimmel, R. (2023). *The economic potential of generative AI: The next productivity frontier*. McKinsey Global Institute. <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier#/>
- Ciubotaru, B.-I. (2025). *The hallucination problem in Generative Artificial Intelligence: Accuracy and trust in digital learning*. 35-45. https://www.researchgate.net/publication/390560703_The_hallucination_problem_in_Generative_Artificial_Intelligence_accuracy_and_trust_in_digital_learning
- Clarivate. (n.d.-a). *Derwent Patent Search*. <https://clarivate.com/intellectual-property/derwent/patent-search/>
- Clarivate. (n.d.-b). *Derwent Patent Analytics*. <https://clarivate.com/intellectual-property/derwent/patent-analytics/>

- Cordeschi, R. (2007). AI TURNS FIFTY: REVISITING ITS ORIGINS. *Applied Artificial Intelligence*, 21(4–5), 259–279. <https://doi.org/10.1080/08839510701252304>
- Daoud, M., Ennoui, A., & Sabri, M. A. (2025). *The strategic role of AI in business digital transformation: Opportunities and challenges*. In Proceedings of the 2025 International Conference on Circuit, Systems and Communication (ICCSC).<https://ieeexplore.ieee.org/document/11134890>
- DeepIP. (n.d.). *Patent Intelligence, From Idea to Enforcement*. <https://www.deepip.ai/>
- European IPR Helpdesk. (2014). *Defending and enforcing intellectual property*. European Commission. <https://www.ipoi.gov.ie/en/understanding-ip/ip-infringement/enforcing-your-ip/defenfig-and-enforcing-ip.pdf>
- European Patent Office. (2024). *Monitoring services in the European Patent Register*. <https://www.epo.org/en/searching-for-patents/helpful-resources/patent-knowledge-news/monitoring-services-european-patent>
- European Patent Office. (n.d.-a). *Novelty and prior art*. <https://www.epo.org/en/learning/learning-resources-profile/business-and-ip-managers/inventors-handbook/novelty-and-prior-art>
- European Patent Office. (n.d.-b). *What is a patent portfolio?*<https://www.epo.org/en/service-support/faq/searching-patents/patent-management-and-valuation/what-patent-portfolio>
- Gbadegeshin, S. A., Al Natsheh, A., Ghafel, K., Tikkanen, J., Gray, A., Rimpiläinen, A., Kuoppala, A., Kalermo-Poranen, J., & Hirvonen, N. (2021). What is artificial intelligence (AI): A simple buzzword or a worthwhile inevitability? In *ICERI2021 Proceedings* (pp. 468–479). <https://doi.org/10.21125/iceri.2021.0171>
- Google Cloud. (2026). *What is semantic search?* <https://cloud.google.com/discover/what-is-semantic-search>
- Gómez Lagerlöf, C. (2024) Grundkurs. Patentavdelningen, patent - och registreringsverket.
- Granstrand, O., & Holgersson, M. (2015). Intellectual property. In D. D. (Ed.), *The Wiley-Blackwell encyclopedia of consumption and consumer studies*. Wiley Blackwell. <https://doi.org/10.1002/9781118989463.wbeccs151>
- Haenlein, M., & Kaplan, A. (2019). A brief history of artificial intelligence: On the past, present, and future of artificial intelligence. *California Management Review*, 61(4), 5–14. <https://doi.org/10.1177/0008125619864925>
- Hain, D. S., Jurowetzki, R., Buchmann, T., & Wolf, P. (2022). A text embedding

based approach to measuring patent to patent technological similarity. *Technological Forecasting and Social Change*, 177, Article 121559. <https://doi.org/10.1016/j.techfore.2022.121559>

Helmers, L., Horn, F., Biegler, F., Oppermann, T., & Müller, K.-R. (2019). *Automating the search for a patent's prior art with a full text similarity search*. https://www.researchgate.net/publication/331509595_Automating_the_search_for_a_patent's_prior_art_with_a_full_text_similarity_search

High-Level Expert Group on Artificial Intelligence. (2018). European commission. *A definition of AI: Main capabilities and scientific disciplines*. <https://digital-strategy.ec.europa.eu/en/library/definition-artificial-intelligence-main-capabilities-and-scientific-disciplines>

HKU Centre for AI, Management and Organization. (n.d.). *Deloitte-HKU Lab for Organizational Transformation*. <https://camo.hku.hk/research-labs/research-lab-s-deloitte-hku-lab-for-organizational-transformation/>

Holgersson, M., & Granstrand, O. (2017). Patenting motives, technology strategies, and open innovation. *Management Decision*, 55(6), 1265–1284. <https://doi.org/10.1108/MD-04-2016-0233>

Holgersson, M., & van Santen, S. (2018). The business of intellectual property: A literature review of IP management research. *Stockholm Intellectual Property Law Review*. https://www.researchgate.net/publication/325995884_The_Business_of_Intellectual_Property_A_Literature_Review_of_IP_Management_Research

Holgersson, M., Granstrand, O., & Bogers, M. (2018). The evolution of intellectual property strategy in innovation ecosystems: Uncovering complementary and substitute appropriability regimes. **Long Range Planning*, 51*(2), 303–319. <https://doi.org/10.1016/j.lrp.2017.08.007>

IBM. (n.d.). *What is agentic AI?* <https://www.ibm.com/think/topics/agentic-ai>

Ikoma, M., & Mitamura, T. (2025). *Can AI Examine Novelty of Patents?: Novelty Evaluation Based on the Correspondence between Patent Claim and Prior Art*. https://www.researchgate.net/publication/388883603_Can_AI_Examine_Novelty_of_Patents_Novelty_Evaluation_Based_on_the_Correspondence_between_Patent_Claim_and_Prior_Art

IPRally. (n.d.). *Features*. <https://www.iprally.com/product/features>

Iprova. (n.d.). *Iprova*. <https://www.iprova.com/>

Jurafsky, D., & Martin, J. H. (2026). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition* (3rd ed). Stanford University. <https://web.stanford.edu/~jurafsky/slp3/>

- Juranek, S., & Otneim, H. (2024). *Predicting patent lawsuits with machine learning. International Review of Law and Economics*. <https://www.sciencedirect.com/science/article/pii/S0144818824000486>
- Kosinski, M. (2024, October 29). *What is black box AI? IBM*. <https://www.ibm.com/think/topics/black-box-ai>
- Lee, M. C. M., Scheepers, H., Lui, A. K. H., & Ngai, E. W. T. (2023). *The implementation of artificial intelligence in organizations: A systematic literature review*. <https://www.sciencedirect.com/science/article/abs/pii/S0378720623000642>
- LexisNexis. (n.d.-a). *Competitive Intelligence*. <https://www.lexisnexisip.com/solutions/ip-analytics-and-intelligence/patentsight/competitive-intelligence-benchmarking/>
- LexisNexis. (n.d.-b). *Patent Portfolio Optimization*. <https://www.lexisnexisip.com/solutions/ip-analytics-and-intelligence/patentsight/portfolio-management-optimization/>
- LexisNexis. (n.d.-c). *Technology Scouting*. <https://www.lexisnexisip.com/solutions/ip-analytics-and-intelligence/patentsight/trend-scouting/>
- LexisNexis. (n.d.-d). *Risk Mitigation*. <https://www.lexisnexisip.com/solutions/ip-analytics-and-intelligence/patentsight/risk-mitigation/>
- LexisNexis. (n.d.-e). *Licensing*. <https://www.lexisnexisip.com/solutions/ip-analytics-and-intelligence/patentsight/licensing/>
- Li, J., Zhu, F., & Hua, P. (2025, November 11). Overcoming the organizational barriers to AI adoption. *Harvard Business Review*. <https://hbr.org/2025/11/overcoming-the-organizational-barriers-to-ai-adoption>
- Li, S., Hu, J., Cui, Y., & Hu, J. (2018). DeepPatent: patent classification with convolutional neural networks and word embedding. *Scientometrics* 117, 721–744 (2018). <https://doi.org/10.1007/s11192-018-2905-5>
- Lin, T.-Y., & Chou, L.-C. (2025). A systematic review of artificial intelligence applications and methodological advances in patent analysis. *World Patent Information*, 82, Article 102383. <https://doi.org/10.1016/j.wpi.2025.102383>
- Lightbringer. (n.d.-a). *Get a free novelty search*. https://www.lightbringer.com/free-novelty-search?_gl=1*v8z0hx*_up*MQ..*_ga*NTA1MjEzODI1LjE3NzY2NzIzMtU.*_ga_CJWBWTFZCC*cZ3NzY2NzIzMtUkbzEkZzEkdDE3NzY2NzI1NzAkajYwJGwwJGgyMDg3MjEwNTk5
- Lightbringer, n.d.-b). *File patents with speed and confidence Designed for innovators, not law firms*. https://www.lightbringer.com/our-solutions/services/patent?_gl=1*11p5hj1*_up*MQ..*_ga*NTA1MjEzODI1LjE3NzY2NzIzMtU.*_ga_C

JWBWTFZCC*czE3NzY2NzIzMTUkbzEkZzEkdDE3NzY2NzI0NjgkajYwJGwwJGgyMDg3MjEwNTk5

Lightbringer. (n.d.-c). *Agentic patent services for deep tech founders*. https://www.lightbringer.com/patent-management-service?_gl=1*198xqv4*_up*MQ..*_ga*NTA1MjEzODI1LjE3NzY2NzIzMTU.*_ga_CJWBWTFZCC*czE3NzY2NzIzMTUkbzEkZzEkdDE3NzY2NzIzOTQkajYwJGwwJGgyMDg3MjEwNTk5PatentManagementService

Lightbringer. (n.d.-d). *Get clarity on your next patent move*. https://www.lightbringer.com/our-solutions/company-stage/patents-for-growing-tech-companies?_gl=1*1y79u1e*_up*MQ..*_ga*NTA1MjEzODI1LjE3NzY2NzIzMTU.*_ga_CJWBWTFZCC*czE3NzY2NzIzMTUkbzEkZzEkdDE3NzY2NzI2NDakajU3JGwwJGgyMDg3MjEwNTk5

Mariani, M., & Dwivedi, Y. K. (2024). Generative artificial intelligence in innovation management: A preview of future research developments. *Journal of Business Research*, 175, 114542. <https://doi.org/10.1016/j.jbusres.2024.114542>

Mirza, H., Bellale, F., & Mirza, C. (2023). Ethical considerations in qualitative research: Summary guidelines for novice social science researchers. *Social Studies and Research Journal*. https://www.researchgate.net/publication/370838199_Ethical_Considerations_in_Qualitative_Research_Summary_Guidelines_for_Novice_Social_Science_Researchers

Mohamed, A., Abdelqader, K., & Shaalan, K. (2025). Explainable artificial intelligence: A systematic review of progress and challenges. *Intelligent Systems with Applications*, 28, 200595. <https://doi.org/10.1016/j.iswa.2025.200595>

NVIDIA. (n.d.). *American Express prevents fraud and foils cybercrime with NVIDIA AI solutions*. NVIDIA. <https://www.nvidia.com/en-us/case-studies/american-express-prevents-fraud-and-foils-cybercrime-with-nvidia-ai-solutions/>

Panwar, V. (2024). AI-powered data cleansing: Innovative approaches for ensuring database integrity and accuracy. *International Journal of Computer Trends and Technology*. https://www.researchgate.net/publication/380366163_AI-Powered_Data_Cleansing_Innovative_Approaches_for_Ensuring_Database_Integrity_and_Accuracy

Patent- och registreringsverket. (n.d.). *Examination of the application*. <https://www.prv.se/en/patents/the-advanced-patent-guide/examination-of-the-application/>

Patenttext. (n.d.). *Patenttext*. <https://www.patenttext.com/about/>

Patlytics. (n.d.-a). *Draft Patent Applications in a Fraction of the Time*. <https://www.patlytics.ai/patent-application-drafting>

Patlytics. (n.d.-b). *AI-Powered Patent Invalidity Analysis*. <https://www.patlytics>

ai/invalidity-analysis

Patlytics. (n.d.-c). *Pinpoint Your Most Valuable Patents and Prune Your Least Valuable*. <https://www.patlytics.ai/patent-pruning>

Patlytics. (n.d.-d). *AI-Powered Patent Invention Disclosure*. <https://www.patlytics.ai/ai-powered-patent-invention-disclosure>

Patlytics. (n.d.-e). *AI-Powered Patent Infringement Detection*. <https://www.patlytics.ai/infringement-detection>

Patlytics. (n.d.-f). *AI-Powered Patent Claim Charting*. <https://www.patlytics.ai/claim-charts>

Patlytics. (n.d.-g). *Patent Vault + Classification*. <https://www.patlytics.ai/patent-vault-classification>

Patlytics. (n.d.-h). *AI-Powered Office Action Analysis*. <https://www.patlytics.ai/office-action>

Patsnap. (n.d.). *Eureka by Patsnap*. <https://eureka.patsnap.com/ip-landing>

Pettersson, M. O., Björkdahl, J., & Holgersson, M. (2025). Profiting from AI: Evidence from Ericsson's pursuit to capture value. *California Management Review*. <https://research.chalmers.se/publication/547748>

Purdy, M. (2024) What is agentic AI, and how will it change work? *Harvard Business Review*. <https://hbr.org/2024/12/what-is-agentic-ai-and-how-will-it-change-work>

Questel. (n.d.). *Orbit Intelligence — IP Intelligence Software*. <https://www.questel.com/patent/ip-intelligence-software/orbit-intelligence/>

Questel. (2025). *Pathways to productivity: What our 2025 industry outlook research reveals about AI in IP*. <https://www.questel.com/resourcehub/pathways-to-productivity-what-our-2025-industry-outlook-research-reveals-about-ai-in-ip/>

Ransbotham, S., Kiron, D., Khodabandeh, S., Iyer, S., & Das, A. (2025). The emerging agentic enterprise: How leaders must navigate a new age of AI. *MIT Sloan Management Review and Boston Consulting Group*. <https://sloanreview.mit.edu/projects/the-emerging-agentic-enterprise-how-leaders-must-navigate-a-new-age-of-ai/>

Regalia, J. (2024). *From briefs to bytes: How generative AI is transforming legal writing and practice*. <https://scholars.law.unlv.edu/cgi/viewcontent.cgi?article=2482&context=facpub>

Rogers, E. M. (1983). *Diffusion of innovations* (3rd ed.). Free Press.

Russell, S. J., & Norvig, P. (2022). *Artificial intelligence: A modern approach* (4th

- ed., Global Edition). Pearson Education Limited. https://api.pageplace.de/preview/DT0400.9781292401171_A41586057/preview-9781292401171_A41586057.pdf
- Sakhare, A., Chavan, P., Nandaniwala, Z., Patne, A., & Desai, M. (2025). A step forward to AGI: Integrating agentic AI and generative AI for human like intelligence. *International Journal for Research in Applied Science & Engineering Technology* <https://doi.org/10.22214/ijraset.2025.66831>
- Setchi, R., Spasić, I., Morgan, J., Harrison, C., & Corken, R. (2021). Artificial intelligence for patent prior art searching. *World Patent Information*, 64, Article 102021. <https://doi.org/10.1016/j.wpi.2021.102021>
- Sheikh, H., Prins, C., & Schrijvers, E. (2023). Artificial Intelligence: Definition and Background. In *Mission AI* (pp. 15–41). Springer. https://doi.org/10.1007/978-3-031-21448-6_2
- Simón, C., Revilla, E., & Saenz, M. J. (2024). Integrating AI in organizations for value creation through human–AI teaming: A dynamic-capabilities approach. *Journal of Business Research*, 182, 114783. <https://doi.org/10.1016/j.jbusres.2024.114783>
- Singla, A., Sukharevsky, A., Hall, B., Yee, L., & Chui, M. (2025, November 5). The state of AI in 2025: Agents, innovation, and transformation. McKinsey & Company. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- Solve Intelligence. (n.d.-a). *Invention harvesting*. <https://www.solveintelligence.com/#invention-harvesting>
- Solve Intelligence. (n.d.-b). *Patent application drafting*. <https://www.solveintelligence.com/#patent-application-drafting>
- Solve Intelligence. (n.d.-c) *Patent Prosecution*. <https://www.solveintelligence.com/#patent-prosecution>
- Solve Intelligence. (n.d.-d). *Claim Charts*. <https://www.solveintelligence.com/#claim-charts>
- Soranzo B, Nosella A, Filippini R (2017), "Redesigning patent management process: an Action Research study". *Management Decision*, Vol. 55 No. 6 pp. 1100–1121. <https://doi.org/10.1108/MD-04-2016-0226>
- Stanford Institute for Human-Centered Artificial Intelligence. (2023). *AI index report 2023*. Stanford University <https://hai.stanford.edu/ai-index/2023-ai-index-report>
- Steinmueller, W. E. (1995). *The U.S. software industry: An analysis and interpretive history*. https://www.researchgate.net/publication/266217991_The_US_Soft

ware_Industry_An_Analysis_and_Interpretive_History

Stilta. (n.d.). *Stilta*. <https://www.stilta.com/>

Stryker, C. (n.d). *What are large language models (LLMs)?* IBM. <https://www.ibm.com/think/topics/large-language-models>

Stryker, C., & Scapicchio, M. (n.d.). *What is generative AI?* <https://www.ibm.com/think/topics/generative-ai>

Taulli, T. (2023). Large Language Models. In: Generative AI. Apress, Berkeley, CA. https://doi.org/10.1007/978-1-4842-9367-6_5

The European Institute of Innovation and Technology (EIT). (n.d). *Creation of a taxonomy for the European AI ecosystem*. <https://www.eit.europa.eu/library/creation-taxonomy-european-ai-ecosystem>

Theofanos, M., Choong, Y.-Y., & Jensen, T. (2024). *AI use taxonomy: A human-centered approach* (NIST AI 200-1). <https://www.nist.gov/publications/ai-use-taxonomy-human-centered-approach>

Trappey, A. J. C., Lin, Y. Y. C., & Wu, C.-Y. (2025). Generative AI-based intelligent patent summarization for intellectual property knowledge communication and cooperation. *World Patent Information*, 83, Article 102410. <https://doi.org/10.1016/j.wpi.2025.102410>

Trappey, A. J. C., Trappey, C. V., Govindarajan, U. H., & Sun, J. J. H. (2021). Patent value analysis using deep learning models—The case of IoT technology mining for the manufacturing industry. *IEEE Transactions on Engineering Management*, 1–13. <https://ieeexplore.ieee.org/document/8941278>

Trappey, A. J. C., Trappey, C. V., Wu, J.-L., & Wang, J. W. C. (2020). Intelligent compilation of patent summaries using machine learning and natural language processing techniques. *Advanced Engineering Informatics*, 43, 101027. <https://doi.org/10.1016/j.aei.2019.101027>

Wang, P. (2019). On defining artificial intelligence. *Journal of Artificial General Intelligence*, 10(2), 1–37. https://www.researchgate.net/publication/335279198_On_Defining_Artificial_Intelligence

Wei, X., Kumar, N., & Zhang, H. (2025). Addressing bias in generative AI: Challenges and research opportunities in information management. *Information & Management*, <https://www.sciencedirect.com/science/article/pii/S0378720625000060>

World Intellectual Property Organization. (n.d.-a). *What is Intellectual Property?* <https://www.wipo.int/en/web/about-ip/>

World Intellectual Property Organization. (n.d.-b). *Patents*. <https://www.wipo.int/en/web/patents>

- World Intellectual Property Organization. (n.d.-c). *How to protect inventions through patents*. <https://www.wipo.int/en/web/patents/protection>
- World Intellectual Property Organization. (2015). *Guidelines for preparing patent landscape reports*. <https://www.wipo.int/publications/en/details.jsp?id=3938>
- World Intellectual Property Organization. (2020-a). *Using inventions in the public domain: A guide for inventors and entrepreneurs* <https://www.wipo.int/publications/en/details.jsp?id=4502>
- World Intellectual Property Organization. (2020-b). *Identifying inventions in the public domain: A guide for inventors and entrepreneurs* <https://doi.org/10.34667/tind.41976>
- World Intellectual Property Organization. (2019). *WIPO technology trends 2019: Artificial intelligence*. <https://www.wipo.int/publications/en/details.jsp?id=4386>
- World Intellectual Property Organization. (2023). *WIPO Patent Drafting Manual* (Second edition). <http://dx.doi.org/10.34667/tind.44657>
- World Intellectual Property Organization. (2024). *Generative AI: Navigating intellectual property*. World Intellectual Property Organization. <https://www.wipo.int/publications/en/details.jsp?id=4713>
- World Intellectual Property Organization. (2025). *World Intellectual Property Indicators 2025: Patents highlights*. <https://www.wipo.int/web-publications/world-intellectual-property-indicators-2025-highlights/en/patents-highlights.html>
- XLSCOUT. (n.d.-a). *Accelerate Patent Prior Art Search & Validate Initial Ideas*. <https://xlscout.ai/novelty-checker-llm-ai-prior-art-search-tool/>
- XLSCOUT. (n.d.-b). *ENHANCE CREATIVE IDEATION AND BRAINSTORMING WITH LLMs & GENERATIVE AI*. <https://xlscout.ai/ideacue-ai-ideation-brainstorming/>
- XLSCOUT. (n.d.-c). *ACCELERATE AND ENHANCE THE PATENT DRAFTING PROCESS USING LLMs*. <https://xlscout.ai/drafting-llm/>
- XLSCOUT. (n.d.-d). *Invalidator LLM: A Litigator's New Best Friend in Patent Litigation*. <https://xlscout.ai/invalidator-llm-a-litigators-new-best-friend-in-patent-litigation/>
- XLSCOUT. (n.d.-e). *ELEVATE YOUR PATENT LICENSING & ENFORCEMENT WORKFLOWS WITH CLAIMCHART LLM*. <https://xlscout.ai/claim-chart-llm-ai-patent-infringement-analysis/>
- XLSCOUT. (n.d.-f). *ELEVATE YOUR PATENT LICENSING WORKFLOWS WITH LLMs*. <https://xlscout.ai/patdigger-llm/>

- XLSCOUT. (n.d.-g). *Noise-Free Patent Monitoring with TechScaper LLM*. <https://xlscout.ai/techscaper-llm-patent-monitoring-and-landscape/>
- Ye, Y., Hao, J., Hou, Y., Wang, Z., Xiao, S., Luo, Y., & Zeng, W. (2024). Generative AI for visualization: State of the art and future directions. *Visual Informatics*. <https://www.sciencedirect.com/science/article/pii/S2468502X24000160>
- Yumiko Hamano. (2016). *Invention disclosure: Form, reporting mechanism, receipt and processing*. World Intellectual Property Organization. https://www.wipo.int/meetings/en/doc_details.jsp?doc_id=376757
- Zhou, Y., Dong, F., Liu, Y., Li, Z., Du, J., & Zhang, L. (2020). Forecasting emerging technologies using data augmentation and deep learning. *Scientometrics* 123, 1–29 (2020). <https://doi.org/10.1007/s11192-020-03351-6>

DEPARTMENT OF TECHNOLOGY MANAGEMENT AND ECONOMICS
DIVISION OF ENTREPRENEURSHIP AND STRATEGY
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden



CHALMERS
UNIVERSITY OF TECHNOLOGY