BLUESHIFT IP
Software Patent Experts

HOW AI IS AUTOMATING THE INVENTIVE PROCESS AND ITS IMPACT ON THE PATENT SYSTEM

BY ROBERT PLOTKIN
SOFTWARE PATENT ATTORNEY

www.blueshiftip.com rplotkin@blueshiftip.com www.linkedin.com/in/robertplotkin

Table of Contents

Introduction 3

How AI Is Being Used to Automate the Inventive Process 4

Impact of AI Automation on Human Inventors and the Inventive Process 8

Impact of AI Automation on the Patent System 17

Conclusion: The Future of AI Automation in the Inventive Process and the Patent System 26

BLUESHIFT IP®
Software Patent Experts

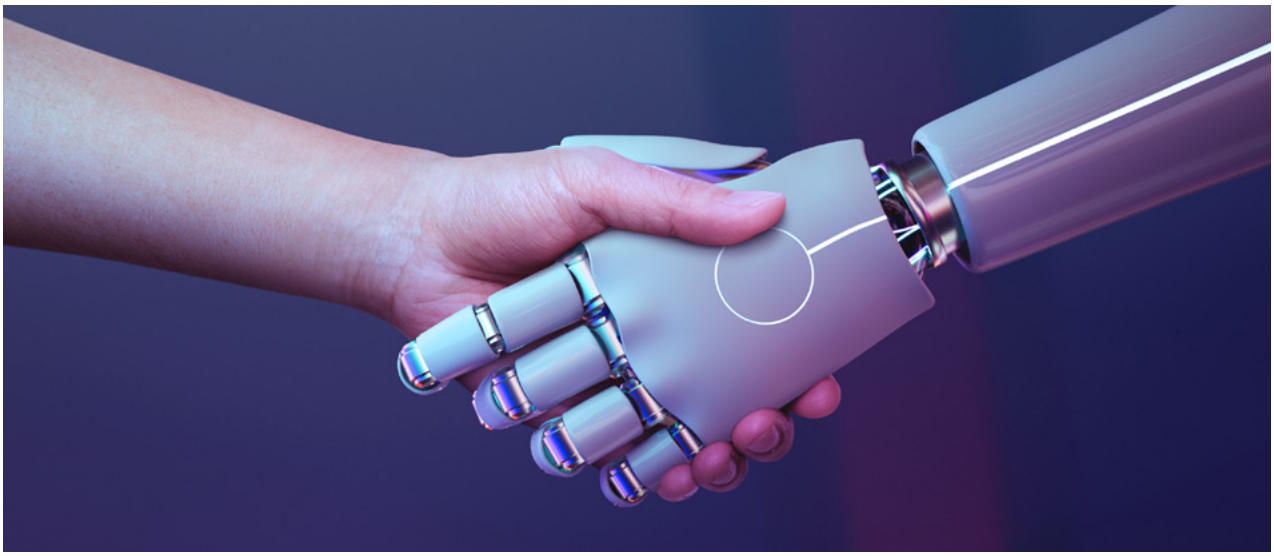

Introduction

Artificial Intelligence (AI) technology is advancing at an incredibly rapid pace. Although recent developments have raised awareness about the ability of large language models to answer questions, search through vast amounts of information, and even write software code, few are prepared for the implications of AI technology that is automating the very process of inventing itself. Such AI-automated inventing is moving us in the direction of making it possible to create new products and services by describing the problems they should solve rather than by specifying how to solve those problems. Such automation of inventing by AI holds the promise of radically increasing the speed and decreasing the cost of innovation, and in turn democratizing innovation and freeing human inventors from tedious tasks so that they can spend their time instead on identifying new problems to solve and working cooperatively with AI to solve those problems.

AI's automation of inventing also raises questions about the role of—and the need for—human ingenuity.

Every previous generation of automation technology has triggered fears that it would make humans obsolete, and the automation of tasks that previously required human creativity elicit this fear in connection with what we consider to be a uniquely human trait.Automated inventing also calls into question whether the patent system is still needed as an incentive to innovation and how to update and administer a patent system that originally was designed to handle industrial-age inventions.

In this white paper, we explore:

- How AI is being used to automate the inventive process;
- The impact of AI-automated inventing on inventors, innovative organizations, and the inventive process; and
- The implications of AI-automated inventing on the patent system.

At every step of the way, we address both the problems and risks of AI as an invention-automating tool and provide practical recommendations for maximizing the benefits and minimizing the risks of AI-automated inventing.

How AI Is Being Used to Automate the Inventive Process

Until recently, inventing a new product or process required some combination of human ingenuity and labor, even if those two didn't always come in equal measures—as Thomas Edison famously said, "Genius is 1% inspiration, 99% perspiration." In fact, the 99% in that equation is something that we not only associate with as a necessary part of the inventive process, but that we take pride in—that manual human effort (sometimes measured in years and thousands of hours) is not only an unavoidable requirement for inventing something new and useful, but also is something that gives us value as humans.

Yet now we are finding that AI and other computer technologies are increasingly automating the perspiration side of the equation, as we'll now see.

Examples of AI Automating the Inventive Process

Before diving into how AI is automating invention, it may help to see some real-world examples of inventions that were generated with significant assistance from AI and the technologies behind them:

Drug Discovery: AI algorithms can analyze large amounts of data to identify potential drug candidates more quickly than humans. This can help to speed up the drug discovery process and reduce the cost of developing new drugs. Stéphane Bancel, the CEO of Moderna, said "AI was critical to the development of our COVID-19 vaccine. We used AI to speed up the process of identifying the most effective vaccine sequence, allowing us to develop a vaccine in record time."

DABUS (Device for the Autonomous Bootstrapping of Unified Sentience): Stephen Thaler, an AI researcher and inventor, has created over 200 inventions using DABUS, an AI system he developed. These inventions range from a device that can sense and report on car emissions to a system that can generate new musical compositions.

Controller: Dr. John Koza used "genetic programming" (a method of problem solving inspired by the principles of natural selection) to create a new controller, a kind of device found in everything from thermostats to automobile cruise control systems. The proof of the controller's novelty is in the patent that the U.S. Patent Office granted over 15 years ago to both the controller itself and to the computer-automated method that Dr. Koza used to invent the controller.

Examples of AI Technologies that Are Being Used to Automate the Inventive Process

It is easy—and dangerous—to use the term "AI" broadly and loosely to refer to just about any computer-implemented technique for automating a task. In reality, a wide range of technologies (not all of which would reasonably qualify as "AI") are being used to automate inventing in a variety of ways. In this section, we will point out the most prominent invention automation technologies as of the time of writing this white paper, some of which overlap with each other:

Machine Learning: Machine learning is a type of artificial intelligence that allows computers to learn from data without being explicitly programmed. It works by using algorithms to identify patterns and make predictions based on the data it has been given . For example, machine learning can be used to teach a computer to distinguish between pictures of dogs and cats. By showing the computer many pictures of dogs and cats, it can learn to identify certain features that are common to each type of animal, like the shape of their ears or the color of their fur. Once it has learned these features, it can use them to classify new pictures as either dogs or cats. Machine learning is also used in many real-world applications, such as image recognition, natural language processing, predictive analytics, and recommendation systems.

Natural Language Processing (NLP): This means that machines can communicate with people in a more natural and intuitive way, like a conversation between two humans. NLP is used in many applications, such as virtual assistants like Siri or Alexa, chatbots that help you with customer service, and even social media platforms that suggest posts or ads based on your interests. Two of the most well-known uses of NLP are automatic speech recognition and language translation. Recently, a new type of NLP—generative language models—has become popular and widely known due to the release of ChatGPT. These models can generate human-like text that is often indistinguishable from something written by a person.that are common to each

Artificial Neural Networks (ANNs): These are algorithms designed to recognize patterns in data in a manner analogous to the human brain. ANNs are made up of artificial neurons, which act as the basic processing unit of the network. Each neuron takes in data, processes it, and then passes the data along to the next neurons in the network. The neurons are connected to each other in a way that allows them to communicate with each other and pass information back and forth. "Deep learning" uses ANNs to learn and make predictions from data. Perhaps the most widely known use of deep learning is ChatGPT, which uses language models in the GPT family to generate natural-language text in response to natural-language prompts.

Genetic and Evolutionary Algorithms: This type of AI works just like nature does to solve problems. These algorithms rely on an artificial version of natural selection, which uses a large group of potential solutions to eventually find the most suitable one for a given problem. The best solutions are then used to create a new generation of solutions, and the process is repeated until the best solution is found. An example of how this is used in the real world is in optimizing transportation routes. By creating a large group of potential routes and evaluating them based on factors like distance and traffic, these algorithms can find the most efficient route for a delivery truck.

Common Features of AI Invention Automation

"Inventing" is a special type of problem solving that involves creating a new product or process to solve a particular problem. Although in some cases an inventor invents something accidentally—what we call "serendipitous invention," as in the famous cases of the invention of penicillin, the microwave oven, and 3M Post-It notes—in the vast majority of cases, an inventor intentionally sets out to tackle a particular problem, and the result (if successful) is a new product or process that solves that problem.

Traditional inventing has involved significant amounts of manual human effort, as reflected in Thomas Edison's famous assertion that genius is "1% inspiration, 99% perspiration." He lived up to this saying in his relentless pursuit to create a successful lightbulb. In his experiments, he tested thousands of materials—including carbonized paper, bamboo, and even cotton thread—as potential filaments for the lightbulb over the course of 14 months. He finally found success in a carbonized bamboo filament, which he used for the first commercially successful lightbulb. In Edison's own words, "I have not failed. I've just found 10,000 ways that won't work."

If Edison had had access to today's AI tools back in his day, he could have eliminated a significant amount of perspiration by using some of the AI technologies listed above to automate various aspects of the inventive process. Although these technologies are distinct from each other, they share some features and benefits in common. For example, AI is being used in the inventive process to:

1. Analyze large amounts of data to identify patterns. Although one problem that plagued AI for many years was that it could not identify patterns that even human children could recognize—such as letters in the alphabet—AI is increasingly able to find patterns in large volumes of data that humans could never find in multiple lifetimes.

2. Generate candidate designs. Today's AI can go beyond merely recognizing patterns (such as recognizing a human face in an image) to generating data in accordance with known patterns. This step of data generation can be used to generate a hypothesis—also called a "candidate solution" or "candidate design" in this paper—which is an essential part of both the scientific method and the process of inventing. In other words, AI can identify existing patterns and create new ones based on the data, which can be used to generate potential solutions or ideas for scientific research and invention.

3. Evaluate candidate designs. Once a candidate invention design—such as a new molecule, airplane wing shape, or electrical circuit—has been generated, it needs to be evaluated to determine whether it solves the problem that the inventors are trying to solve. In the old days, that required physically building and testing a prototype. Now, AI can often evaluate the candidate design automatically, such as by simulating its behavior, without ever having to create it in the real world. Human feedback may be received to supplement or replace the AI's evaluation.

4. Rank, filter, and modify candidate designs. Once candidate designs have been evaluated, they can be ranked according to those scores, then poorly scoring candidates can be removed from consideration and remaining candidates can be modified. Even the process itself can be updated, such as by updating the training of a neural network.

5. Iterate over the above steps. Some or all of the steps above can be repeated, often incredibly quickly and an incredibly large number of times.

The end result is some number of candidate designs—potential inventions—that can be reviewed by human inventors. You can imagine how such a tool would have benefited Edison in his pursuit of a lightbulb filament. Edison could have conveyed his inspiration to an AI tool, which could then have tried out various materials in simulation and, in a matter of minutes, produced the best result, or at least a small number of candidates for real-world testing.

The steps above are the same steps that human inventors often perform when inventing. In fact, these steps can be seen as the kind of hypothesis generation, followed by hypothesis testing, that is fundamental to the scientific method.

What can be different about the process when it is performed using AI is that the volume of data and speed with which the steps can be performed can result in invention designs that would never have been possible for human inventors to produce in any practical amount of time. Combine this with the rapidly decreasing cost of AI and you have a recipe for an explosion in inventive capacity.

Another benefit of AI as an inventive tool is that human inventors can have biases and blind spots that cause them to overlook solutions that AI, in its childlike naïvete, does not rule out. Humans, while having an enormous capacity for creativity, can also be limited by their own experiences, beliefs, and perspectives, which can lead to cognitive biases and blind spots. For example, a human inventor may be biased towards a particular solution because it aligns with their personal beliefs or professional training, even if it is not the most efficient or effective solution. A mechanical engineer, for example, may have been taught in school that steel is the best material for a certain purpose and overlook other possibilities (such as a modern lightweight ceramic) as a result when designing a new machine. On the other hand, AI algorithms are designed to explore all possible solutions without any inherent biases (although note the biases that can result from biases in the data that are used to train AI). This means that AI often can consider and evaluate solutions that a human may have overlooked or discounted due to their biases or limited perspective.

Although I mentioned various types of AI technology at a high level, I intentionally did not describe their technical features in great detail. One reason for that is that specific technologies come and go. Although Deep Learning and generative language models are hot right now, no one knows what the AI flavor of the day will be in just a few years. When I wrote my book, *The Genie in the Machine*, back in 2009, genetic and evolutionary algorithms were dominant as invention accelerators, and artificial neural networks (ANNs) were widely denigrated.

That situation has reversed now. Yet the basic ways in which AI is automating inventing have remained the same, regardless of which specific technologies are the driving force at the moment. For that reason, I've chosen to focus more on the what than on the how of AI invention automation.

To put it another way, I'm focusing on the tasks that are being automated by AI and the impacts of that automation on inventing and the patent system, rather than on the technical details of that automation, so that we don't miss the forest for the trees.

Impact of AI Automation on Human Inventors and the Inventive Process

Addressing the Fear of Being Replaced by AI

> I have created a machine in the image of a man that never tires or makes a mistake. Now we have no further use for living workers."
>
> ***Rotwang,***
> in Fritz Lang's Metropolis

Myth and science fiction are replete with cautionary tales about powerful technology overtaking its creators, ranging from the myth of Prometheus to Frankenstein to I, Robot and The Terminator. Such stories reflect the widespread fear that AI will replace human inventors and make them obsolete. Such fear has long gripped humans when automation has taken hold, as was the case in the Industrial Revolution when the invention of the power loom led to the Luddite movement, and has found its way into our literature and myths, as in the legend of John Henry.

Although such fears are understandable, and although AI may very well eliminate the need for human inventors to perform certain tasks, the actual effects of AI on human inventors are much more complex than simple replacement of humans by AI. In fact, there are many reasons to believe that AI automation will increase the demand for human ingenuity.

The fear that AI automation will make human inventors obsolete might stem from the assumption that the demand for "inventive services" is fixed. However, the well-known “Jevons

Effect" illustrates that demand for a resource often increases when the cost of that resource drops. For example, when gasoline prices drop, people often respond by driving more, which causes them to spend more in a given time period on gasoline than they did before prices dropped.

We frequently see similar effects in the technology world. Yet even the smartest of us fall for the "fixed demand" fallacy, as Bill Gates famously did when he said that "640K [of Random Access Memory] ought to be enough for anybody" back in 1981. The typical computer today has about 50,000 times more memory than that. What Gates failed to take into account was that, as computers provided more useful features to consumers, demand for additional memory-intensive features increased in a feedback loop that shows no signs of stopping.

Similarly, when the cost of a service drops, demand for that service often increases. This phenomenon is particularly evident in the telecommunications industry, where the availability of low-cost email and text messaging services has resulted in people communicating with each other more frequently than ever before.

One less obvious consequence of increased access to a service is referred to as "uncovering of latent demand," meaning that not only do existing customers seek more of that service, but entirely new markets open up for that service. As an example, in my own field of legal services, many people don't seek the assistance of lawyers for wills, real estate transactions, and domestic relations matters because the cost is too high—not because they don't have a need for help. The introduction of low-cost alternatives, such as Quicken Family Lawyer and LegalZoom, has made help available to such people at a price they can afford and has thereby uncovered the latent demand for such services.

For all of these reasons, we should expect that automating aspects of the inventive process will increase demand for "inventive services." There are so many problems that need to be solved that I find it hard to believe that total demand for inventive problem solving by humans will shrink as the result of computer-automated inventing.

At a more philosophical level, the fear-based belief that automating certain inventive tasks will reduce the total demand for human ingenuity was addressed well by Isaac Asimov: "I believe that scientific knowledge has fractal properties, that no matter how much we learn, whatever is left, however small it may seem, is just as infinitely complex as the whole was to start with. That, I think, is the secret of the Universe." I think this is just as true with inventing.

I don't mean to dismiss fears of human obsolescence with any of this. Technology often does eliminate human jobs. Revolutionary times are trying times, even for those who stand to benefit from them, particularly when success is not guaranteed. Technological obsolescence, automation, and outsourcing have all put pressures on engineers in recent years to upgrade their skills or join the ranks of the unemployed.

In just my own lifetime I've witnessed the elimination of the neighborhood newspaper stand, video store clerk, and travel agent, among many others. I've seen my own field of the law radically transformed by technology. At the beginning of my career, I recall a senior attorney telling me of a lawyer he knew who practiced law through the 1960s who refused to have a telephone at his desk, who only spoke to clients in pre-scheduled meetings in his office, and who only advised clients after due consideration by sending them a typed letter by postal mail. That kind of law practice is no longer economically viable.

Yet in many cases it is tasks, not people, that are automated by technology. In response, what we humans need to do is learn to adapt our skills and work processes to focus on where we can bring the most value to the inventive process as a whole. That is where I will turn next.

As AI Automation Moves Up the Waterfall, So Too Can (and Should) Human Inventors

One helpful way to break down tasks in the inventive process is by using the classic "Waterfall Model" of problem solving that is used in engineering schools. According to this model, when you are faced with a problem to solve, you follow a sequence of steps that start at the high-level problem to be solved and that proceed through increasingly low-level design tasks, resulting at the end with a solution to the initial problem. Although the specific steps or "tiers" of the waterfall are defined in many different ways, one useful sequence for our purposes is:

Problem definition: define the problem to be solved (e.g., "to build a human catch-and-release mouse trap")

Requirements analysis: define the requirements that any solution to the problem from the previous step must satisfy (e.g., "catch at least 90% of mice humanely and make it easy to relocate the mouse")

Functional design: design functional modules that solve the problem and satisfy the requirements from the previous steps (e.g., a "luring module" to lure the mouse, a "capturing module" for catching the mouse without killing it, a "transport module" for transporting the mouse and mousetrap, and a "release module" for releasing the mouse safely)

Structural design: design structures that perform the functions, satisfy the requirements, and solve the problem from the previous steps (e.g., a holder for cheese, a pressure-sensitive platform that releases a trapping door in response to the weight of the mouse, a carrying handle, and a quick-release latch)

For our purposes it isn't important that, in practice, inventors don't always carry out all of the steps above. Instead, the reason for introducing the waterfall model is to illustrate how AI can automate some aspects of the inventive process and not others, and how humans can identify the skills that they need to adopt in order to leverage AI, rather than be replaced by it.

In the simplest terms, it would benefit human inventors to identify **which step(s) in the waterfall AI is able to perform automatically and to focus on developing skill at performing the other steps.** For example, as AI automates various steps in the inventive process, humans will need to:

Focus on higher-level tasks that AI cannot (yet) automate: In general, what we have been witnessing in recent decades—and in recent years in the case of AI—is that computer automation has been climbing up the waterfall by automating increasingly high-level (abstract) steps of the inventive process. In such instances, there is less need and opportunity for humans to perform the kind of detailed design work at the "Structural Design" level that has been the exclusive province of human inventors for most of history. Much as individuals in a profession, such as law, tend to start out by performing low-level detailed work in that profession and then gradually move on to higher-level work and supervising people who perform low- and high-level work, so too will it behoove humans to gain skill in performing tasks at higher levels of the waterfall, such as problem definition and requirements analysis.

Decompose and assign tasks: When solving any real-world problem, AI typically cannot be used to solve the entire problem. However, that doesn't mean that the problem can't benefit from the use of AI. Instead, the best inventive process often involves and requires a human to break down the problem into sub-problems, to choose the best approach for solving each sub-problem, and then to delegate each sub-problem accordingly. This may involve assigning some sub-problems to humans and some sub-problems to AI. Even assigning sub-problems to AI may require choosing which form of AI is best to use.

Invent new forms of AI: AI still can't invent new kinds of AI (although it increasingly can improve itself). For the foreseeable future, humans will be needed to develop new AI algorithms and to select the right kind(s) of AI to tackle particular problems.

Configure AI for optimal results: All AI must be configured for use, such as in the form of "meta-parameters." Much skill is required for this. For example, in the case of ChatGPT, initial testing of the GPT-3 language model produced excellent results when ChatGPT was used to complete text (e.g., "The cat chased the..."), but performed poorly when it was given instructions (e.g., "Write a song about a cat chasing a mouse"). In response to this, the human developers of ChatGPT engaged in significant retraining of the GPT-3 language model, specifically on the task of responding to instructions, including using human reviewers to evaluate ChatGPT's attempts to follow instructions. This resulted in a significant improvement that would not have been possible without significant human intervention.

Evaluate AI outputs: If you've used ChatGPT, you know that its outputs can be deceptive, in part because it uses language that sounds completely confident even when the text it generates is flatly false. For this and other reasons, its outputs should always be evaluated by humans before being used. The same is true of the outputs of AI generally, especially for health- and safety-critical applications. AI also may produce outputs that would be unethical to deploy, in which case humans are needed to make ethical choices about whether and how to deploy such outputs.

Refine AI outputs: In almost all cases, even when AI does an excellent job at generating output—such as a set of candidate mousetrap designs—that output is not perfect. Humans will still need to engage in low-level design tasks to refine and finalize the outputs of AI.

Synthesize outputs of AI and humans: Whenever a complex task is decomposed into smaller tasks, it is always necessary to synthesize the results of those smaller tasks to produce an overall solution to the problem at hand. Humans will still be required to perform this critical synthesis task.

As the above examples illustrate, there will be no lack of need for humans in the face of AI automation for the foreseeable future. Humans will just need to adopt new skills and be flexible in response to developments in AI.

AI and Humans Will Work as a Collaborative Inventive Team, with Each Doing What It Does Best

We've now seen that not only has AI not automated all tasks in the inventive process, but that AI has actually created new tasks that humans are required to perform in the inventive process. Although it's an open question whether this will result in a net gain or loss in the amount of human labor that is required to satisfy the demand for inventive services, it should put to rest the twin claims that: (1) the demand for inventive services is necessarily fixed; and (2) AI necessarily reduces that demand.

AI and humans, however, generally have complementary skills. As a result, inventing will be the most effective when humans and AI follow a division of labor in which each does what it does best. For some common examples (which are not true in all cases):

Who Does it Better?	
Humans	**AI**
Problem selection and definition	Identifying patterns in large data sets
Requirements analysis	Generating candidate solutions
Aesthetic judgments	Applying clear rules to large sets of data
Ethical judgments	Searching for data
Problem selection and definition	Identifying patterns in large data sets

Based on the above table, humans tend to be better at:

Problem selection and definition: Humans know what humans need. AI doesn't know or decide which problems to tackle, and even if it could, it probably wouldn't select the problems that best serve humanity. Furthermore, AI needs some definition of the problem to be solved as an input before it can attempt to generate solutions to that problem. As a result, humans are needed to select and define the problems that AI is to solve.

Requirements analysis: Similarly, AI usually cannot take a mere description of a problem as an input and generate a useful solution for that problem as an output, especially for most challenging real-world problems. One reason for this is that the trial-and-error methods that most AI uses run into the problem of "combinatorial explosion"—the number of possible solutions to explore is simply too vast, even for the most powerful AI. As a result, humans are needed to define the problem in more specific terms in the form of the requirements that a successful solution needs to satisfy.

Aesthetic judgments: Despite the recent impressive ability of text-to-image generators such as DALL-E, such AI generates images without evaluating the aesthetic qualities of such images. Humans are still much better at making aesthetic judgments, such as judging how pleasing the scent of a perfume is or how well the visual design of a new smartphone will mesh with other products from the same manufacturer.

Ethical judgments: AI is notorious for generating ethically repugnant output, such as racist rants in response to questions. One reason for this is that most AI gains its skill from large volumes of training data, and if the training data (e.g., text on the Internet) has ethical shortcomings, so too will the output of AI that is trained on that data. ChatGPT, for example, has built-in filters that attempt to suppress offensive output because the core of the software has no such filter. Outputs of AI should always be screened by humans for consistency with our values.

Choosing when to use humans and when to use AI: AI typically cannot decide when a particular problem is better solved by humans or by AI (although there have been some interesting recent efforts at this). Humans are much better at this "orchestration" function of deciding which components of a problem are best solved by humans and which are better-suited for AI.

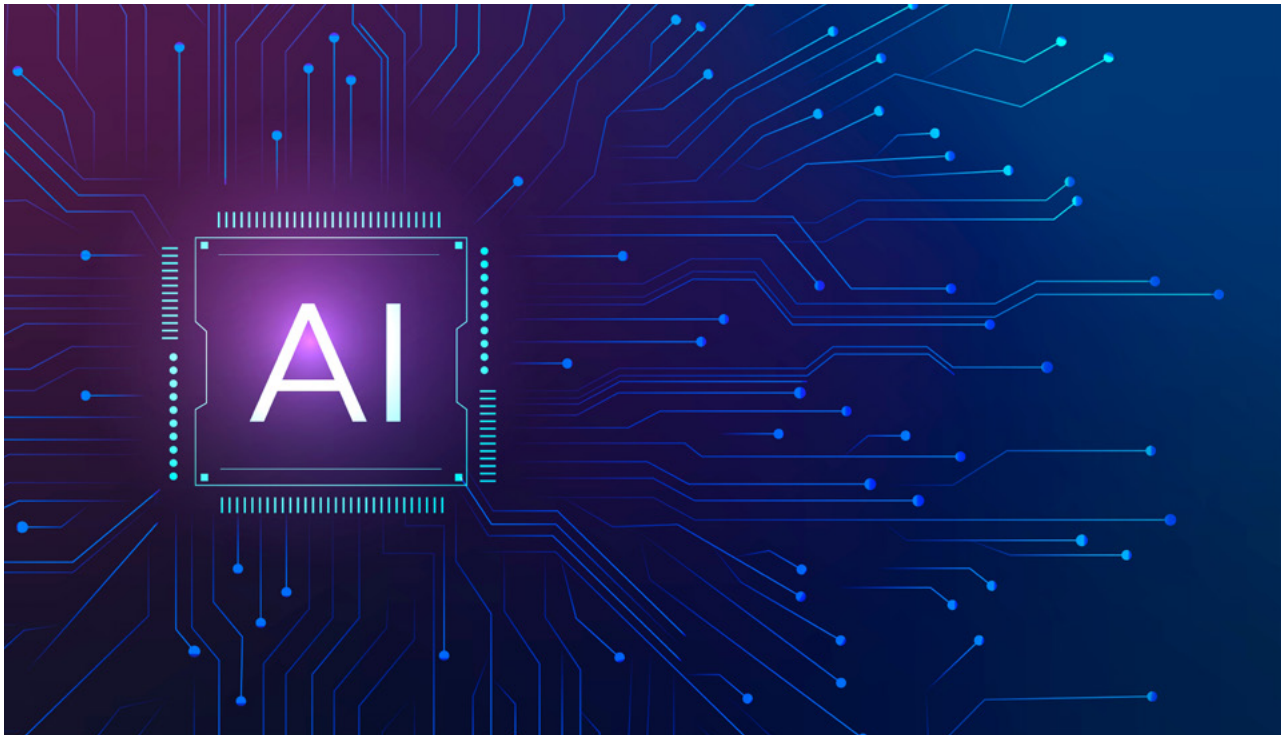

Similarly, computers (AI) tend to be better at:

Identifying patterns in large data sets: This is what most AI is best known for today, especially in the case of Deep Learning and other neural network-based technology. AI increasingly excels at finding and recognizing patterns, which is why it is so widely used for speech recognition, image recognition (e.g., for recognizing tumors in radiology), and web searching.

Generating candidate solutions: AI is getting better and better at generating what I've been calling "candidate solutions" to problems, e.g., text that answers a question, an image that is described by text, or a design for a circuit that satisfies specified requirements. This is particularly true in the case of recent developments in generative language models and other "generative AI."

Applying clear rules to large sets of data: Although expert systems and other rule-based AI gets less attention, you already benefit from it in your tax filing, appointment scheduling, and project management software. Even simple rules can be extremely powerful and time-saving when they are applied to massive volumes of data.

Searching for data: Although we tend to take Internet search engines for granted, everyone marveled the first time they used Google, and for good reason. An incredible amount of sophistication sits behind the software and hardware that delivers results to you nearly instantly when you ask who starred in your favorite movie 20 years ago.

Doing all of the above extremely quickly: Any one of the above tasks might not be that impressive or useful if it were only applied to the amount of text on this page. Yet computers become indispensable—and can produce results that no human or team of humans could attain—when applied to exabytes of data to produce results in seconds. And computers do this reliably and without getting tired or complaining.

Writing Wishes for Computerized Genies. In 2009 I wrote *The Genie in the Machine: How Computer-Automated Inventing is Revolutionizing Law and Business*. I used the metaphor of genie and wisher as an analogy to the roles of AI and humans in the process of inventing. More specifically, I argued that, as AI increasingly automated low-level aspects of the inventive process, humans would be in the position of writing "wishes" for AI to grant in the form of designs for new inventions. At that time, the wishes still took highly technical forms that only experts could write, such as "fitness functions" in the case of genetic and evolutionary algorithms. Now, with the rise of ChatGPT and large language models more generally, we're seeing AI that can take natural language wishes as input and produce natural language wishes as output. Although anyone can write a prompt for ChatGPT, people are already:

- Using the term "prompt engineering" to refer to the process of writing prompts that are better than others for producing desired results;
- Recognizing that writing such optimal prompts requires time and skill that not everyone has; and
- Selling such prompts for people who don't have the skill and/or time to write them (e.g., at PromptBase).

I predicted all of this when I dedicated an entire chapter of *The Genie in the Machine* to "Learning How to Wish," in which I described the need for human inventors to develop the valuable skill of wish-writing and the economic value that wishes themselves would have. Of course, we've already been living through this for quite some time—software began to have independent economic value as a product as far back as the 1960s. What is changing now is that it is increasingly becoming possible for valuable wishes to be written in natural language form.

The Human-Computer Inventive System

> By relieving the brain of all unnecessary work, a good [mathematical] notation sets it free to concentrate on more advanced problems and in effect increases the mental power of the [human] race.
>
> ***Alfred North Whitehead***

Rather than viewing humans and AI as competitors—or, worse yet, AI as a replacement for humans—I prefer to see humans and AI in two ways:

AI as a tool for augmenting human inventive skill. This places AI in the long historical context of conceptual and physical tools that humans have used to make them better and more efficient at inventing, including mathematics, written language, chisels, hammers, rulers, scales, and previous generations of computer technology. A human who skillfully uses AI has the ability to invent more successfully and efficiently than a human without AI.

Humans and AI as an inventive team. Some prefer the more impersonal term "system" to "team," but in either case the idea is that, instead of AI replacing humans, AI joins the team that consists of a combination of humans and AI working together to solve a problem. Such a team—like a human-only team—works best when each member does what it does best and when communications within the team are clear.

Some benefits of such invention augmentation (or human-computer invention collaboration) for humans are:

Humans can focus on the problems to solve rather than on the details of solving them. Recalling Edison, AI will enable humans to focus on inspiration, rather than perspiration, by leaving the perspiring to the AI. Many people prefer to focus on higher-level problem solving and find such work to be more fulfilling than grinding through endless experiments. Those who prefer to perform detailed design, or who are more skilled at that kind of work, can still find opportunities to engage in it (such as by applying that skill to the candidate solutions generated by AI).

Better results can be obtained more efficiently than with human effort alone. When a human-AI team is well configured, it can produce results that no human or machine can produce alone. Producing such results to solve previously unsolved problems can be extremely satisfying and empowering.

Invention will be democratized. Just witness the explosion of creativity that the free public release of ChatGPT has unleashed. Although the honeymoon of completely free access will surely come to an end, it won't be possible to put the genie back into the bottle. More generally, as AI increasingly makes it possible to write "wishes" in natural language, the ability to invent will become available to an increasingly wide range of people who previously lacked the scientific and technical skills to be inventors. This "democratization of invention" can be seen as a kind of "uncovering of latent supply"—as the cost of providing inventive services goes down, an increasing number of "inventive service providers" will have the ability to join the inventive market.

In summary, despite the legitimate fears that AI-automated inventing can raise for human inventors, and the very real effort that such inventors will need to engage in to update their skills to be effective inventors in the age of AI, AI-automated inventing holds the promise of enabling human inventors to exchange control,in the form of designing the details of each new invention),for power,–in the form of the ability to create inventions that solve increasingly complex problems more effectively and efficiently.

Impact of AI Automation on the Patent System

The use of AI in the inventive process is impacting the patent system in ways that will only compound in the future. Patents grant legal rights to their owners to exclude competitors from making, using, and selling patented inventions. Patents are granted to inventors (and the companies who employ them) who are the first to create or improve a specific invention. The use of AI in the inventive process raises questions about who should be credited as the inventor and what legal standards should be applied to determine whether a patent should be granted. Some even question whether the patent system is necessary or serves its original purposes in the face of AI-automated invention.

In this section I address specific challenges posed by AI-automated inventing for the patent system and recommend ways to address each of those challenges.

Is AI Making Patents Obsolete?

Before addressing specific legal issues raised by AI for patents, I'll speak to the elephant in the room—has AI made patents obsolete? Do we need a system that is designed to provide human inventors with an incentive to create new inventions and disclose those inventions to the public anymore when new inventions are just the push of a button away?

The previous sections should make clear that the need for human inventors is here to stay and that automation of aspects of the inventive process may even increase the demand for human inventors while shifting the nature of the tasks that humans perform in the inventive process. For that reason alone, the fundamental basis for the patent system still exists and will continue to exist for the foreseeable future.

Furthermore, the fundamental justification for the patent system is that human inventors require two separate incentives:

- to create inventions rather than not to invent; and
- to disclose their inventions to the public rather than to create inventions and keep those inventions as trade secrets to avoid competition.
-

As long as AI cannot invent on its own or choose to disclose its creations freely to the public, the need for the patent system to incentivize human inventors to invent and to disclose their inventions to the public will remain as strong as ever.

Can AI Be an Inventor on a Patent?

The question whether the law permits AI technology to be named as an inventor on a patent application is one of the first things many people think of when they hear about AI-automated inventing and has garnered attention in recent years because of Dr. Stephen Thaler's attempt to have his DABUS software named as the sole inventor on patent applications that he has filed around the world for inventions (namely, a beverage container and a light-emitting device) that he claims were generated solely by DABUS without any human input. So far he has not succeeded, except in South Africa, although that is a misleading "success" because South Africa merely registers patent applications and does not examine them substantively for compliance with the legal requirements that an invention must satisfy to be granted a patent. In short, therefore, no patent office or court anywhere in the world has yet held that AI technology can be named as an inventor on a patent application.

Although the question of AI inventorship captures our imagination by conjuring up images of humanoid robots with their own identities and personalities, whether AI can be named as an inventor under existing patent laws has little, if any, practical consequences for patent owners or the public, and detracts from the more significant, challenging, and pressing issues that I describe below.

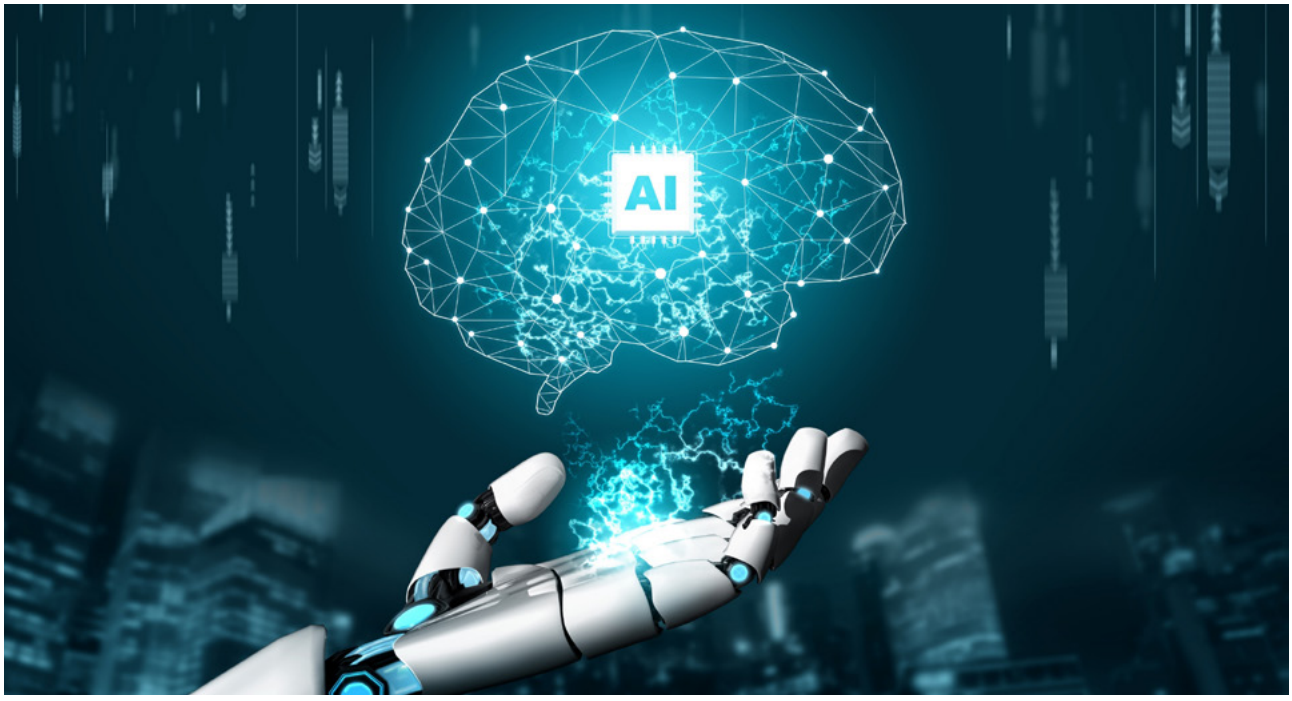

Is Patent Law Suitable for Protecting AI Innovations? ("Patent-Eligibility" and "Patentable Subject Matter")

Although the requirement that an invention be "patent eligible" (also called "patentable subject matter") under U.S. law has proven extremely difficult to define, the basic concept is easy to understand. I think of it as the requirement that a product or process be the "kind of thing" that the patent system was designed to protect—namely, inventions—as opposed to the kinds of things that other kinds of intellectual property were designed to protect—namely aesthetic works (e.g., novels, artwork, and songs) in the case of copyright, and logos, product names, and other "brand identifiers" in the case of trademark law. People generally agree on the easy cases, namely that a song is copyrightable subject matter but not patentable subject matter, and that a steam engine is patentable subject matter but not copyrightable subject matter. People also generally agree that mathematical formulae, laws of nature, and theories are not patent eligible (although so-called "practical applications" of all of those can be patent eligible).

The challenge arises when attempting to use words to define what is and isn't patentable subject matter, especially in borderline cases. In particular, it has always proven challenging to define which types of processes are and are not patent-eligible, and this has always been particularly challenging in the context of software. Although a process for manufacturing a spark plug is very likely patent eligible, a process performed by software to sort data may or may not be.

Some words that have been associated with patentable subject matter in the U.S. are "useful" and "practical," whereas Europe and other jurisdictions use terms such as "industrial" and "technical." As straightforward as these terms may sound at first blush, they have proven to be incredibly difficult to define and apply clearly and objectively when they are applied to specific inventions in practice.

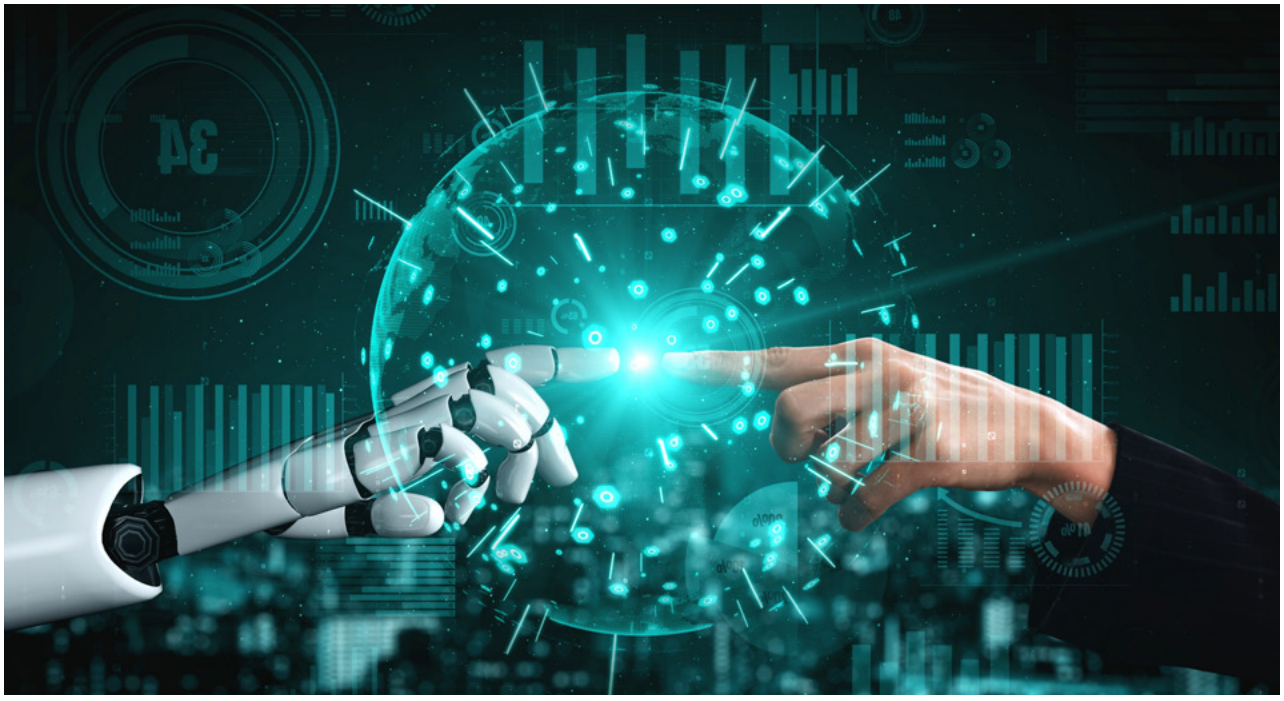

One interesting way in which AI may be bringing patent eligibility full circle is that AI can be used to assist in the process of creating designs for physical products ranging from automobile bumpers to suction cups to cancer drugs. Such physical products are patent-eligible regardless of the process by which they were invented. In such cases, AI raises fewer issues for patent eligibility than traditional software.

Even AI-generated software and other digital creations of AI (such as data structures and trained models) may satisfy the requirements for patent eligibility more easily than previous forms of software to the extent that the process of creating such entities involved testing them sufficiently to demonstrate that they work to perform a useful function. Submitting evidence of such tests may be sufficient to satisfy concerns about the abstractness of software whose high-level functions are described in patent applications, but which have never been implemented in working software.

Even AI-generated software and other digital creations of AI (such as data structures and trained models) may satisfy the requirements for patent eligibility more easily than previous forms of software to the extent that the process of creating such entities involved testing them sufficiently to demonstrate that they work to perform a useful function. Submitting evidence of such tests may be sufficient to satisfy concerns about the abstractness of software whose high-level functions are described in patent applications, but which have never been implemented in working software.

New AI algorithms themselves, such as new kinds of neural networks and new methods of training models may, however, face many of the same difficulties satisfying the requirements for patent eligibility as other kinds of software has faced generally for decades. Such algorithms may, for example, be viewed as mathematical methods, especially if they are described merely as operating on abstract data without any practical, real-world use. To obtain patent protection for such algorithms, it may be helpful or necessary to provide examples in which they have been used to solve some problem in the physical world—as moving a robot arm to grasp an object more accurately—or to demonstrate that they enable computers to perform some fundamental operation on data (e.g., arithmetic, sorting, or decision-making) more effectively and/or efficiently.

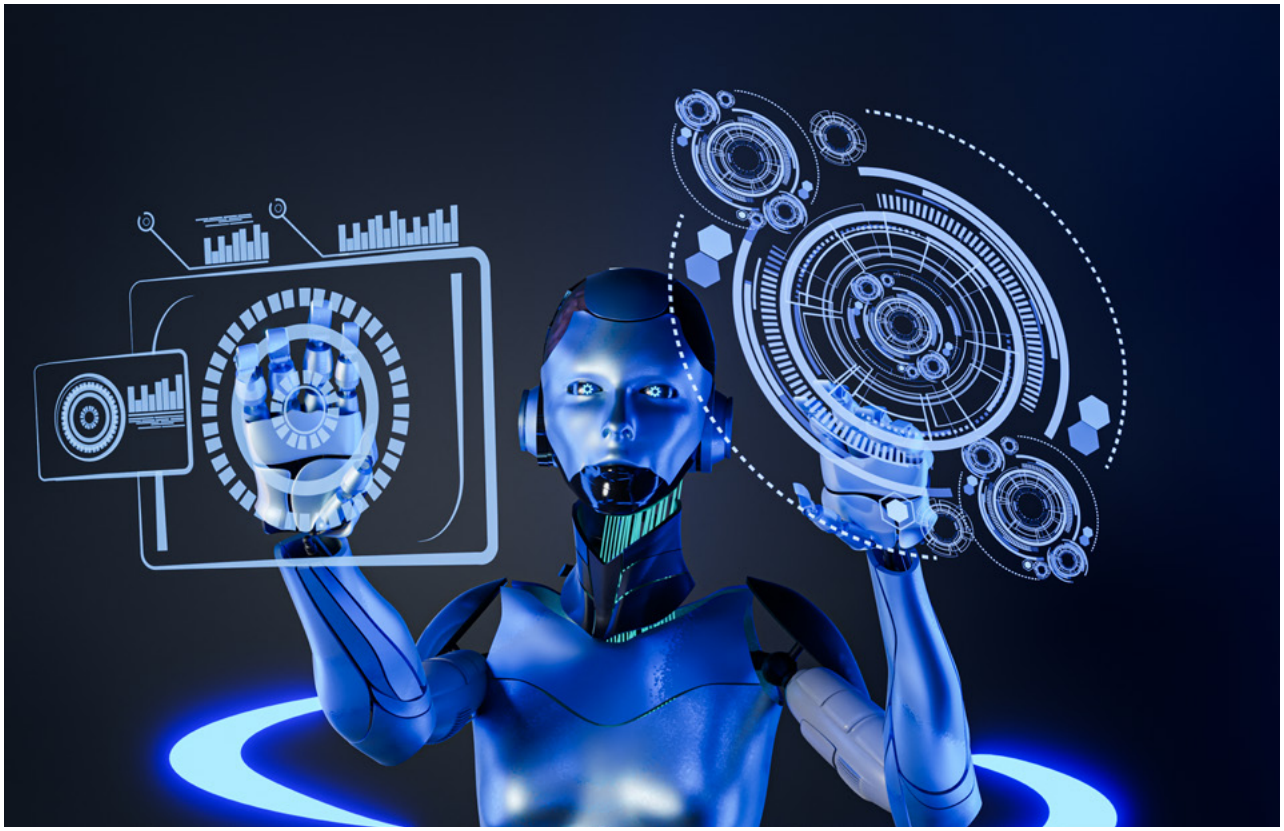

AI Raises the Bar for Patentability ("Nonobviousness" and "Inventive Step")

Patents are intended to provide protection only for inventions that are a true advance over existing “state of the art” technology. As a result, trivial changes to previous inventions—such as changing the color of a mousetrap in a way that doesn't affect its function—do not merit patent protection. As in the case of terms such as "useful" and "technical" in the case of patent eligibility, it is very difficult to define what constitutes a "non-trivial" change that warrants patent protection. Yet this is what patent law's "nonobviousness" requirement (referred to as the "inventive step" requirement in Europe and elsewhere) attempts to do.

One purpose of the nonobviousness requirement is to weed out inventions that would likely be created and made available to the public even without patent protection.

Consider the following example: If any ordinary automotive engineer, working away at the daily grind, would have found it obvious to increase the length of a car frame by an inch, then we don’t need to grant a patent to you if you happen to be the first person to make such a modification. Even if the unavailability of patent protection for your car frame causes you to decide to keep it a secret, another automotive engineer will likely come up with the same idea tomorrow and incorporate it into an automobile made available to the public. Therefore we, the public, don’t need to grant you a patent on the inch-longer car frame in order to ensure that it will become available to us.

> There are no hard problems, only problems that are hard to a certain level of intelligence. Move the smallest bit upwards [in level of intelligence], and some problems will suddenly move from ‘impossible’ to ‘obvious.’ Move a substantial degree upwards, and all of them will become obvious."
>
> ***Eliezer S. Yudkowsky***

Because the nonobviousness requirement has such a high risk of being applied in a subjective, “I know obviousness when I see it” kind of way, the law attempts to interpret nonobviousness objectively as follows. When a patent examiner or judge applies patent law's nonobviousness requirement to a particular invention, it is necessary first to identify the "level of skill" of a "person having ordinary skill in the art" (PHOSITA) to which the particular invention pertains. ("Art" is the archaic term that patent lawyers use to refer to a scientific or technical field.)

The PHOSITA's knowledge, experience, and skill level serves as the baseline for the determination of nonobviousness. It is important to note, however, that the PHOSITA is not necessarily a real person, but rather an imaginary person who is proficient in the relevant art.

In order to identify the level of skill of the PHOSITA, courts will look to evidence in the form of expert testimony, printed publications, patents, and other sources. For example, courts will ask what level of education (e.g., bachelor's degree or Ph.D.) PHOSITA has in a particular field. If the level of education is higher, then any particular proposed invention in that field will have to satisfy a higher bar for nonobviousness. This determination is typically conducted on a case-by-case basis and is necessary for a court to determine whether the particular invention is nonobvious.

We've already seen that AI is boosting the **effective level of skill** of ordinary inventors in a variety of fields. Alternatively, you could see the **inventive system** consisting of human inventor + AI as having a higher level of inventive skill than the human inventor alone. For example, supply every automotive engineer with state-of-the-art artificial invention technology and train them in how to use that technology, and you have effectively boosted the level of ordinary inventive skill in the field (art) of automotive engineering. If instead all of those automotive engineers had gone back to school and obtained a Ph.D. in automotive engineering, the law would recognize that the level of skill of the ordinary automotive engineer had just increased a notch for purposes of applying the nonobviousness requirement. This would raise the bar for nonobviousness in the field of automotive engineering, thereby making it more difficult to obtain patents in the field.

By analogy, if typical inventors in a particular art become conversant in the use of AI as a tool in the inventive process and use AI regularly to help them invent, should the law recognize that this has a similar boosting effect on the level of skill of PHOSITA as if those inventors had all earned advanced degrees, i.e., to **raise the bar** for nonobviousness in that field? I think so.

Let's stay on this point for a moment because it is significant. Perhaps the most important impact of AI on the patent system is that, by increasing the effective level of inventive skill of human inventors in a particular art, AI should cause the law to raise the nonobviousness bar that inventions in that art need to satisfy in order to merit patent protection.

One particular non-intuitive consequence of this is that such a raised bar for nonobviousness—just like a raised bar resulting from an increase in average education level for inventors in an art—should apply to **all** inventions in that art for which patent protection is sought, **not only inventions that were created using AI.** For example, if AI comes into widespread use by automotive inventors, and a particular automotive inventor invents a new spark plug the old-fashioned way without the use of any AI, then whether that manually invented spark plug is nonobvious should be determined based on whether **a typical automotive engineer would have invented that spark plug by making ordinary use of AI, even though the actual inventor did not use AI to invent that actual spark plug.**

Interpreting nonobviousness in this way is consistent with the purpose of the nonobviousness requirement: to ensure that patents are not granted on inventions that would be made available to the public even in the absence of patent protection. If any automotive engineer having ordinary skill would find it obvious to invent a new spark plug by, for example, providing a known set of requirements as input to widely used AI, then there is no reason to grant a patent on such a spark plug, regardless of how it was invented. By putting the spark plug design within the reach of anyone with ordinary skill, artificial invention technology makes granting a patent to someone who happens to be the first person to use AI to design that spark plug not only unnecessary, but positively harmful to the public interest.

Although it may be challenging to demonstrate what the current level of skill is in a particular art in light of the widespread use of AI in that art, it is already challenging to prove what the level of skill of humans in a particular art is. There may be some new challenges involved in proving what could have been invented using a particular AI technology at a particular time. Doing so might require using an archived version of AI technology from a certain date, providing it with certain inputs, and seeing whether it can produce the claimed invention as an output.

Although the impact of AI on nonobviousness might be significant, the nonobviousness standard is already defined by the law in a way that is designed to enable it to adapt to changes in human skill and technology over time. No changes to the patent statute or new court decisions are needed. Patent offices and courts already calibrate the level of skill that is used in making nonobviousness determinations to particular people and technologies. In this regard, the nonobviousness standard is already well equipped to handle the significant leaps forward in effective inventive skill that AI will produce, as long as actors in the legal system are attuned to this and apply the ways in which AI actually is used in particular arts to the nonobviousness standard in particular cases.

How Can and Should AI Innovations be Described in Patents? ("Written Description" and "Enablement")

The word patent is used to refer to two different things:

A set of legal rights that empowers the patent's owner to exclude others from making, using, or selling the patented invention without the patent owner's permission.

The document that confers those legal rights.

Two key reasons for requiring inventors to describe their inventions in a document that is published for the world to see are to:

Put the public on notice of what has been patented: The publication of patents and patent applications promotes transparency in the patent system. By publishing patent applications, the public can see what has been filed and what is pending. This gives the public a fair opportunity not to infringe those patents by not making, using, or selling products and services that are described in published patent documents. It also would be unfair to find someone liable for infringing a patent if that person had no way of knowing that an invention had been patented.

Encourages innovation and research: Patent publication encourages innovation and research by allowing inventors to share their ideas and inventions with the world. This can lead to new ideas and inventions, which can be patented and further developed.

The patent system could have been designed without requiring such documents to be written and published. For example, a patent system could just require inventors to submit a working model of their invention to the patent office, where that model could serve as proof of what had been invented and also be used to compare against future products to determine whether those products infringe (violate) the patent, without requiring inventors to submit any written descriptions of their inventions. (In fact, until the 1880s, the US Patent Office required inventors to submit working models of their inventions as part of the patent application process. However, this requirement was eventually replaced by the requirement to submit detailed drawings and written descriptions of the invention.)

The primary requirement that the description of an invention in a patent application must satisfy is that it must "enable" a "person having ordinary skill in the art to which the invention pertains" (recall PHOSITA from our discussion of nonobviousness) to make and use the invention without having to engage in "undue experimentation." In other words, the patent application must serve as a recipe that PHOSITA can follow to construct and use your invention. Although your description does not need to include every last detail about how to make and use your invention, you can't leave out critical information so that PHOSITA would need to expend an unreasonable amount of effort to make a functioning version of your invention based on your explanation.

AI-enabled inventing poses some challenges for writing patents that satisfy the enablement requirement. I'll touch on these for several categories of AI-related inventions:

New types of AI itself: New hardware and software that implement new kinds of AI technology, such as new AI algorithms, can be described in the same way as traditional hardware and software in the vast majority of cases. As a result, such technologies don't generally pose new problems for patent law's enablement requirement—except if those technologies are themselves invented using AI. (See below.)

Inventions invented using AI: I'll break these down into two sub-categories:

Inventions that are easily understandable by humans: In some cases, AI helps to invent an invention—such as a more aerodynamic airplane wing—that humans can understand and describe in the same way as if that invention had been invented manually. In such cases, a patent for the invention can describe it the old-fashioned way in a patent.

Inventions that are not understandable by humans: In some cases, AI generates output that "works"—it performs a useful function, such as distinguishing images of dogs from images of cats—but it is difficult or impossible for a human to understand how it performs that function. This lack of transparency into the inner workings of some AI systems is referred to as "the inscrutability of AI." In the case of inscrutable AI, it may be necessary to describe the data and process that were used to generate the AI,e.g., to train a neural network, rather than to describe the end product directly. In cases in which the training data and/or process themselves are inscrutable, this can pose additional challenges.

Methods for training AI models: At least as of this writing, most methods for training AI models are invented by humans. For example, humans set the hyper-parameters of AI models and decide which training data to feed those models in order to train them. Humans also typically create the algorithms that train the AI model based on the training data. As a result, I would put new training methods into the same category as "new types of AI itself," above. If, however, a new method for training an AI model is invented using the assistance of AI, then the considerations described above in "Intentions Invented Using AI" will apply.

Trained AI models: Because of the inscrutability of AI described above, often it is not possible for a human to understand or provide an understandable description of a trained AI model, such as a trained neural network, of any non-trivial size. The simplest explanation of such a model may be the model itself, which is not particularly useful for a human to use as a recipe for recreating the invention. As a result, patenting trained AI models may require describing the data and process that were used to train the model.

As we have now seen, although AI-enabled inventing poses some challenges for satisfying the enablement requirement in patents, these challenges appear to be manageable for the foreseeable future, so long as careful attention is paid to the requirements for enablement in each case.

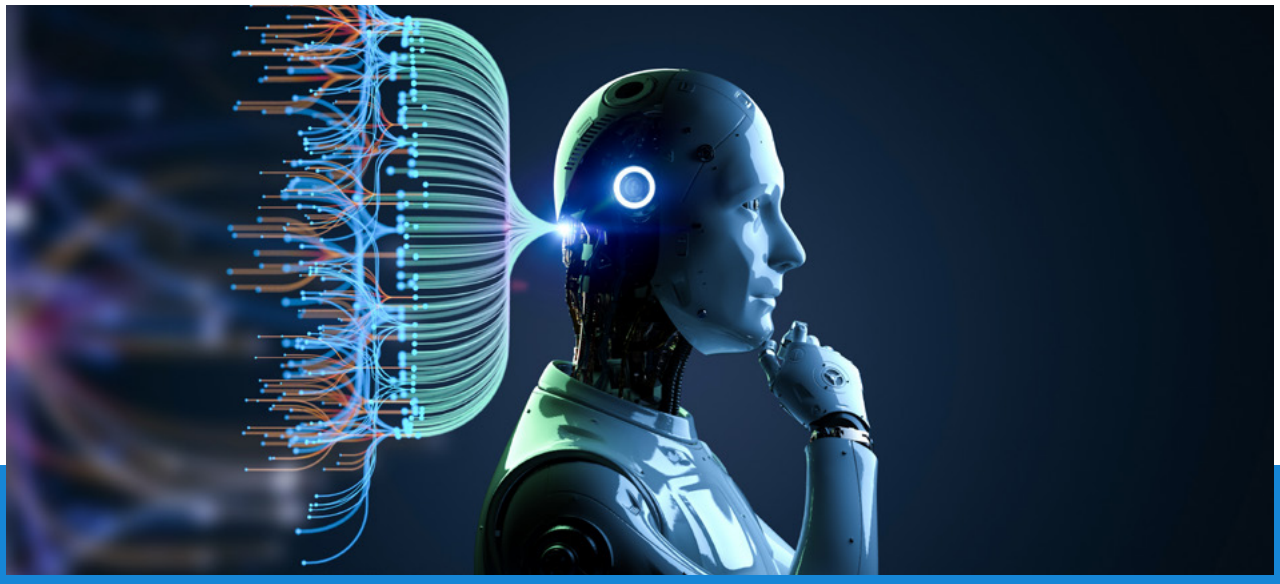

Conclusion: The Future of AI Automation in the Inventive Process and the Patent System

AI technology is revolutionizing the way that we approach the inventive process. By automating tasks that previously required human ingenuity, we are now able to invent new products and services at a faster rate than ever before. However, this also means that human inventors must learn new skills to work effectively in cooperation with AI in order to stay competitive.

As AI becomes increasingly integrated into the invention process, it is essential for the patent system to adapt to the challenges created by AI-automated inventing. Without proper recalibration and oversight, there is a risk that the patent system will be unable to adequately promote and protect innovations in AI and innovations generated using AI.

As AI continues to revolutionize the inventive process, it's important for inventors and innovative organizations to stay ahead of the curve. To ensure that your AI technology is fully protected and can be leveraged to its fullest potential, it may be worth exploring partnering with patent experts who have experience in the field. This is where Blueshift IP can help. To learn more about Blueshift IP and access additional resources about AI and patents, including our blog, podcasts, and webinars, visit www.blueshiftip.com.

About the Author:

Robert Plotkin is a patent attorney at the boutique patent law firm of Blueshift IP who has been specializing in software and AI patents for over 25 years. He is an MIT Computer Science graduate and the author of *The Genie in the Machine: How Computer-Automated Inventing is Revolutionizing Law and Business* (Stanford University Press, 2009)--the first, and still the only, book to explore how artificial intelligence is automating the inventive process and the implications of automated inventing for innovation and the patent system. He is a frequent writer and speaker on the impact of AI-automated inventing on patents, and he uses the legal strategies outlined in his book to obtain strong, broad, and defensible patent protection for his clients in the U.S. and worldwide.

www.blueshiftip.com

rplotkin@blueshiftip.com

www.linkedin.com/in/robertplotkin

Check out our free blog, podcast, and webinars on software patents and the future of AI

www.blueshiftip.com/blog/

www.blueshiftip.com/podcasts/

www.blueshiftip.com/webinars/