INNOVATION COUNCIL

IP Protection for Artificial Intelligence

by Philip R. Wadsworth

Senior Advisor to Innovation Council and former Chief Patent Counsel and Legal Counsel for Global IP Policy at Qualcomm

and Jennifer L. Brant

Director of Innovation Council

Introduction

Artificial Intelligence is rapidly evolving to provide accurate information and solutions to problems. AI solutions can be extremely useful in many fields of endeavor.

There are various components and aspects of such systems. Currently, there is some uncertainty - especially with respect to patent protection - as to how intellectual property (IP) rights can be used to protect those components along with the outputs of AI systems.

To understand the applicability of IP rights to these components and outputs, it is necessary to first have a basic understanding of AI and its constituent technologies. This context is set out in Section 1. Then, Section 2 links the components and outputs to different types of IP rights. Finally, Section 3 identifies areas for possible future study and applies the analysis to another frontier technology, quantum computing.

What is AI?

As defined by IBM:

> *"In its simplest form, artificial intelligence (AI) is a field, which combines computer science and robust datasets, to enable problem-solving. It also encompasses sub-fields of machine learning and deep learning which are frequently mentioned in conjunction with artificial intelligence. Most of today's AI systems are referred to as "weak" or "narrow" AI, as they are only capable of performing specific tasks. For the future, there are ongoing efforts to create "strong" or "general" AI systems that will approach the intelligence of human beings."*

AI solutions are comprised of algorithms that are used to create expert systems that can make predictions or classifications based on input data. There are several steps involved in creating an AI system:

- First, a computer must be appropriately configured to solve a desired problem.
- Algorithms must be created that are likely to be able to solve that problem. These will run on the computer.
- Datasets must be collected and structured in a way to "train" the algorithms to accurately solve the problem in accordance with known outcomes.
- There are typically many iterations of the training process, with necessary modifications to the data and/or algorithms to achieve the desired solution.
- Once the training is satisfactorily completed, unstructured data may be used to achieve accurate problem-solving.

Examples of how AI systems are applied include speech recognition, facial recognition, manufacturing robots, self-driving cars, smart assistants, proactive healthcare management, disease mapping, automated financial investing, virtual travel booking agents, and social media monitoring.

Components of AI Systems

1. Hardware

Computers

One of the main components of an AI system is a traditional computer itself. Computers come in various sizes and computing capacities, such as main frame computers, desktop computers, personal computers and even mobile phones. These types of computers are comprised of one or more processing units (e.g. microprocessors) and software algorithms (e.g. specific software applications), which are configured to perform the desired function(s) of the system.

Neural Network

A more specific and sophisticated type of computer system that may be used in AI systems is referred to as a neural network, or in some instances an artificial neural network. A neural network is a system of hardware and software patterned after the operation of neurons in the human brain. It usually involves a large number of processors operating in parallel and arranged in tiers. The first tier receives the raw input information - analogous to optic nerves in human visual processing. Each successive tier receives the output from the tier preceding it, rather than the raw input - in the same way that

neurons further from the optic nerve receive signals from those closer to it. The last tier produces the output of the system.

Each processing node has the ability to leverage its own knowledge base, which includes information it has seen and any rules it was originally programmed with (i.e. algorithms) or that it developed for itself. The tiers are highly interconnected such that the specific information from each node can be shared with other relevant nodes in the network. There may be one or multiple nodes in the so-called output layer, from which the answer it produces can be read.

Neural networks are highly adaptive, and they can modify themselves as they learn from initial training, and as subsequent runs provide more information to enable the systems to provide solutions to problems that are likely to be accurate. The learning model is centered on weighting the input streams, since each node weighs the importance of input data from each of its predecessors. Inputs that contribute to getting right answers are weighted higher and, using this approach, the system learns and improves its performance over time.

2. Training Tools

Algorithms

Machine learning refers to the use and development of computer systems that can learn and adapt without following explicit instructions, by using algorithms and statistical models to analyze and draw inferences from patterns in data. UC Berkeley breaks the machine learning algorithm into three main parts:

1. *Decision Process:* In general, machine learning algorithms are used to make a prediction or classification. Based on some input data, which can be labelled or unlabelled, the algorithm will produce an estimate about a pattern in the data;
2. *Error Function:* An error function serves to evaluate the prediction of the model. If there are known examples, an error function can make a comparison to assess the accuracy of the model; and
3. *Model Optimization Process:* If the model can do a better job of evaluating the data points in the training set, then weights can be adjusted to reduce the discrepancy between the known example and the model estimate. The algorithm will repeat this evaluate-and-optimize process, updating weights autonomously until a threshold of accuracy has been met.

There are different types of machine learning algorithms, including those described below.

Deep learning is a sub-field of machine learning and typically uses artificial neural networks to process the data sets. However, deep learning automates some of the process, eliminating some of the manual human intervention required and enabling the use of larger data sets. "Deep" machine learning can leverage labelled datasets, also known as supervised learning, to inform its algorithm. At the same time, it doesn't necessarily require a labelled dataset. It can process unstructured data in its raw form (e.g. text, images) and automatically determine the hierarchy of features that distinguish different categories of data from one another. Unlike machine learning, it doesn't require human intervention to process data, which allows for scalable machine learning in more interesting ways.

Computer vision is a field of artificial intelligence (AI) that enables computers and systems to derive meaningful information from digital images, videos, and other visual inputs — and to take actions or make recommendations based on that information. If AI enables computers to think, computer vision enables them to see, observe, and understand. Computer vision trains machines to perform these functions using cameras, data, and algorithms. For example, because a system trained to inspect products or watch a production asset can analyze thousands of products or processes a minute, noticing imperceptible defects or issues, it can quickly surpass human capabilities, with their

reliance on retinas, optic nerves, and a visual cortex.

Datasets

Data annotation is the technique by which data is labelled so as to make objects recognizable by machines. Data labelling refers to adding more info and metadata to various data types (text, audio, image and video) in order to train machine learning systems. Annotated, or labelled, data is the basic requirement to train machine learning models.

Structured data - which is typically categorized as quantitative data - is highly organized and easily decipherable. Examples of structured data include dates, names, addresses, and credit card numbers. Normally, a quality assurance phase for the data is completed before proceeding to the training steps.

Training Processes

Training processes are critical to ensure the successful development of an accurate AI system. Training entails teaching the system to properly interpret data and learn from experience, to ultimately perform tasks with accuracy.

First and foremost, AI training starts with data. While the actual size of the dataset needed is dependent on the project, all machine learning projects require high-quality, well-annotated, or structured data in order to be successful.

In the initial training step, an AI model is given a set of training data and asked to make decisions based on that information. This is an iterative process to allow tweaking of the data and other components to achieve the desired outcome. Once the AI system has completed the initial training, it can proceed to the next step of validation. In this step, assumptions are validated based on how well the AI will perform using a new set of data. As with the training step, it is necessary to evaluate the results to confirm the AI system is behaving as expected, and to account for any new variables that were not previously considered. In the final training step, a real-world test is conducted using an unstructured dataset that does not include any tags or targets. If the AI system can make accurate decisions based on this unstructured information, it is ready to go live. If not, it's back to the training step and the process is repeated until the outcome is accurate and the AI system is performing as expected. Simply put, unstructured data is comprised of data that is usually not as easily searchable as structured data; it may include, for example, formats like audio, video, and social media postings.

Upon the successful results, based on the input of the unstructured data, a viable and innovative AI system has been developed. Then the AI system may be used to generate outputs that represent solutions to the problems it is capable of handling. There are many ways those outputs may be represented to best convey the solution to the problem, and many ways the outputs can be further used according to the needs of users.

How IP Protection Interacts with AI

Patents

To provide a brief refresher, patent protection is limited to a particular country or region, with each office having its own definition of patentable inventions. For example, in the US, Section 101 of the Patent Act provides that an inventor can patent any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof. Many countries and regions also require that an invention, to be patentable, must also have a technical character. Other requirements of patentability are that inventions must be novel and non-obvious to one skilled in the art. Below is a general discussion of potential patent protection for the various components of an AI system.

Computer Architecture

As mentioned earlier, there may be various hardware configurations for the processors used in the AI system. To the extent that a new computer architecture or an improvement to an existing architecture has been developed, patent protection may be considered. For example, the basic configuration for a neural network is well established, but there may be improvements to or modifications of that basic configuration that may be patentable. It should be noted that in many cases where a general computer is used as part of the AI system, the software algorithms used to program the system to perform its AI functions are likely to include patentable content.

Algorithms

For purposes of this section, algorithms are those programming steps, in the form of software code, that direct the computer to perform its AI functions. Some algorithms simply represent a particular mathematical formula to calculate a number, whereas others are much more complex. They can allow the AI system to be trained and to provide solutions to problems that are much more sophisticated than just calculating a mathematical formula. The latter type of algorithms may be considered for patent protection in many jurisdictions. For example, in the case of neural networks, even though the basic architecture may be used, a particular node (i.e. processor) may have new or improved algorithms which may be patentable. However, it should be noted that there is much uncertainty regarding what constitutes patentable subject matter for algorithms/software and the law is continuing to evolve. Thus, working with an experienced patent attorney or agent, in the jurisdiction where patent protection is being considered, to obtain advice regarding software implemented inventions is highly advisable.

Data Structures

As mentioned in the background section, annotated and structured data is necessary for training an AI system to be able to successfully analyze unstructured data to provide an accurate solution to a problem or output. Structured data, to the extent it establishes relationships among the dataset and/or includes algorithms that enable the AI system to learn to accurately draw conclusions, may be eligible for patent protection. The process used to create the structured data may also be eligible for patent protection.

Training Processes

The training process for an AI system can be both tedious and iterative, in order to assure a trustworthy system that can accurately process unstructured input data. This may lead to observations about, and modifications to, the process which substantially improve it. The improved training process is likely eligible for patent protection.

Protecting AI Solutions and Outputs

A well-trained AI system will be capable of providing solutions to problems or other outputs that may be inventive in their own right. This phenomenon has raised a significant inventorship issue. Many countries require inventors to be natural persons, which means that inventions created by an AI system are unavailable for patent protection. The United States and Europe take this approach, but this is not the case in all countries' IP offices. There is ongoing litigation, plus legislative activities, regarding this issue. These should be monitored to stay abreast of any changes to present policy and law, across jurisdictions. It's worth noting that further refinement by a person of the rendering of the output from the AI system, for example a novel graphical user interface, may be eligible for patent protection.

Commercially Available or Open Source AI systems

Over time, many AI systems have become generally available for use by consumers. They are typically designed to perform a specific purpose. However, it is possible for someone to use one of these systems as a starting point then determine that the system can be used, with or without modification, to solve a different problem. In this case, patent protection may be sought for a new use of the original system.

Copyright

Copyright protects original works of authorship including literary, dramatic, musical, and artistic works, such as computer software. For purposes of the section, the copyright law provides the exclusive rights to the owner, among others, to copy and distribute their original works and derivative works thereof. It should be noted that the copyright does not protect the functional aspects of computer software. Computer software is created in different forms to ultimately provide instructions executable by a computer. First source code is created which represents a human understandable form of software. The source code is then compiled and assembled into the binary object code form, which is then converted into executable code which is understandable by the computer. All of these forms of computer

Algorithms

Since the algorithms used in AI systems are generally created in the forms mentioned above, are eligible for protection by copyright.

Datasets

As noted above, structured data sets typically require the creation of labelling schemes that create the relationships among the individual pieces of data. To the extent that these schemes represent original content, the dataset may be eligible for copyright protection. Similarly, unstructured data could be protectable by copyright if arranged in an original manner. However, if the data consists of factual information in no particular form, it is unlikely to be protectable by copyright.

Trade Secrets

Protection of trade secrets is recognized in some form by most countries in the world. World Trade Organization members are bound by the TRIPS Agreement to provide protection against unauthorized use of trade secrets. In the United States under the Uniform Trade Secrets Act (“UTSA”), a trade secret is defined as information that derives independent economic value because it is not generally known or readily ascertainable, and it is the subject of efforts to maintain secrecy. One advantage of trade secret protection is that, unlike patent and copyright protection, there is no time-limited term of protection. As long as the trade secret information is maintained in confidence, trade secret protection may be relied upon.

There are various considerations to assess before relying on trade secrets. First, within an organization, a program must be developed to identify trade secrets resulting from R&D efforts. An important part of the program is to understand if the information being considered for trade secrets protection might be disseminated in a product that could be reverse engineered to detect the trade secret. If so, patent protection may be a better strategic approach. All employees must be required to sign confidentiality agreements and trade secret information should only be disseminated to those employees with a “need to know”. Finally, if products or software are being distributed to outside parties, appropriate agreements with confidentiality provisions must be developed, and these must be signed by the recipients of the information.

Algorithms

The algorithms used in an AI system, in addition to copyright protection, can potentially be protected by way of trade secrets. However, the precautions mentioned above should be taken to prevent disclosure of the trade secret information.

Datasets

Since substantial effort goes into creating the structured data to optimize the training of an AI system, the structured dataset is a good candidate for trade secret protection. Depending on the sources of the unstructured datasets and the uniqueness of their arrangement by the AI developer, they may or may not be conducive to trade secret protection. For example, data gathered from public sources is unlikely to be a good

candidate for such protection.

Computer Configurations

If the AI system includes the development of a unique computer architecture or configuration, trade secret protection for that configuration may be considered. Again, it is important to consider the precautions mentioned above.

Problem Solution and Outputs

The results that come from using the AI system may be considered for trade secret protection. If such results are only going to be used internally for competitive advantage, then reliance on trade secret protection may be a strategic approach. One example of this is the use of AI systems to support R&D efforts. The AI systems can help to prevent an organization from going down dead-end paths, saving a lot of time and money. Such information is invaluable for competitive advantage and is an excellent candidate for trade secret protection. If, however, the information is likely to be shared externally, then the precautions mentioned earlier should be considered before making a decision as to whether to rely, or not, on trade secret protection.

Trademarks

A trademark is any word, name, symbol, or design, or any combination thereof that is used in commerce to identify and distinguish the goods of one manufacturer or seller from those of another, and to indicate the source of the goods. Service marks are similarly used to identify the services of an entity. Like patents, governments issue registrations for the protection of trademarks. Like patents and copyright, such IP rights are national in scope.

Product or Solution Name

Developers of AI products that are offered commercially to other parties may seek trademark protection for their products. An example of this is Ipiphany. There are also companies that offer services using their AI systems. A good example of this is IBM's service mark "Let's Create".

Concluding Remarks

The use of AI in new technologies and services is rapidly growing. For example the International Data Corporation (IDC) reports: "Worldwide revenues for the AI market, including software, hardware, and services, are forecast to grow 16.4 per cent year over year in 2021 to $327.5 billion … by 2024, the market is expected to break the $500 billion mark with a five-year compound annual growth rate of 17.5 per cent and total revenues reaching an impressive $554.3 billion".

This projected growth portends high rates of innovation and, likely, increasing demands on IP protection systems. Due to the nature of AI solutions, present-day IP systems may be challenged or even find themselves ill-equipped to provide protection for certain aspects of AI systems. For instance, as mentioned earlier, when the AI system is an innovator/inventor/author, patent and copyright protection may not be available because protection is generally only available to natural persons. Even the question as to how software should benefit from IP protection is still unsettled, and the views of IP offices around the world differ on this point. Decisions by courts and legislative initiatives can be expected to settle some of these issues over the coming years. Their activities should be closely followed to stay abreast of important changes to the IP laws and scope of protection for AI-related technologies and solutions.

Looking to the future, the application of AI to other fields of technology will undoubtedly continue to grow. One somewhat nascent area is the use of quantum computing to deliver more powerful AI solutions.

Quantum computing is a rapidly emerging technology that harnesses the laws of quantum mechanics to solve problems that are too complex for classical computers. Quantum computing systems typically rely on the use of super conducting computers as their foundation.[1] Like IP protection for the components and solutions associated with the more conventional AI systems mentioned above, such protection will likely be available also for the components and solutions associated with quantum computing systems that provide AI solutions. Again, activities by the courts and legislative bodies should be closely monitored to stay abreast of these likely rapidly developing changes.

[1] Quantum mechanics is a fundamental theory in physics that provides a description of the physical properties of nature at the scale of atoms and subatomic particles. More specifically, it is the branch of physics that deals with the mathematical description of the motion and interaction of subatomic particles, incorporating the concepts of quantization of energy, wave-particle duality, the uncertainty principle, and the correspondence principle. Examples of problems that could be solved with quantum computing systems include improving the nitrogen-fixation process for creating ammonia-based fertilizer, creating a room-temperature superconductor, removing carbon dioxide to improve the climate, creating solid-state batteries, developing new drugs, and weather forecasting.