E

WIPO
WORLD
INTELLECTUAL PROPERTY
ORGANIZATION

SCP/30/5
ORIGINAL: ENGLISH
DATE: MAY 28, 2019

Standing Committee on the Law of Patents

Thirtieth Session
Geneva, June 24 to 27, 2019

BACKGROUND DOCUMENT ON PATENTS AND EMERGING TECHNOLOGIES

Document prepared by the Secretariat

Table of Contents

I. INTRODUCTION

1. At the twenty-ninth session of the Standing Committee on the Law of Patents (SCP), held in Geneva from December 3 to 6, 2018, the Committee agreed that the Secretariat would prepare a background document on patents and emerging technologies and submit it to the thirtieth session of the SCP. This document is submitted to the SCP pursuant to that decision.

2. The term "emerging technologies" might have the broad meaning, covering various new technologies, including artificial intelligence (AI) and machine learning, blockchain, synthetic biology and gene editing etc. However, AI and blockchain, for example, are, from the technology point of view, different technologies that may involve different issues in relation to patents. From the discussions at the twenty-ninth session of the Committee that had led to the above decision, many delegations who had taken the floor had referred to AI as an issue to be discussed by the Committee. Accordingly, this document covers background information on patents and AI.

3. The document consists of three parts. The first part of the document provides background information about the AI technology. With the assistance of an AI technology expert,[1] the first part of the document illustrate the basic technological concept of AI, particularly on machine learning technology, which is the core of the current AI development. Such an introductory description of the technology is considered necessary, since implication of a particular technology to the patent system requires at least the basic understanding of the technology itself.

4. The second and third parts of the document describe the intersection between patent systems and AI. They address two distinct issues: the second part looks at the AI technology (or AI-related inventions) as the subject of patent protection, and the third part discusses use of the AI technology as a tool for the authorities and users of the patent systems.

5. As to the term "quality of patents", although no single definition is identifiable, two main concepts arose from the earlier activities of the SCP. They are: (i) the quality of a patent itself; and (ii) the quality of patent procedures before patent offices and beyond (document SCP/27/4 Rev.). From this viewpoint, it could be said that the issues under patent protection of AI-related inventions touch upon the first aspect of patent quality, while the issues about improvement of patent procedures using AI technology relate to the second aspect of patent quality.

6. In addition, the document contains an Annex, which lists conferences organized, and publications made available, by WIPO and its Member States.

---

[1] The Secretariat greatly benefitted from the contribution of Mr. Patrice Lopez (Science-Miner) for the preparation of the first part of the document, "Introduction to AI Technology: Neural Network and Deep Learning". He also assisted the Secretariat in reviewing the accuracy of the document with respect to the description of the AI technology.

II INTRODUCTION TO AI TECHNOLOGY: NEURAL NETWORK AND DEEP LEARNING

7. While there is no single definition of AI, AI systems can be viewed primarily as learning systems. The first part of the document introduces the most important technical concepts around Neural Network (NN) and Deep Learning (DL), which are today the booming technologies in AI.[2] It provides an understanding of how these emerging technologies work in an accessible manner for non-computer specialists in order to assist better understanding about the intersection between AI technology and patents.

A. Machine Learning

8. Historically the first approaches to AI were to *program* a machine. Program here means that a human provides step-by-step instructions to the machine for completing a certain task. In the 80s for example, the dominant AI approach was the *Expert Systems*, using rules written by specialists of their domain to reproduce human expertise. Costly and limited, these approaches led to the so-called second AI winter between years 1987 and 1993.

9. In contrast, the Machine Learning (ML) approaches explore how a machine can learn to solve a task from examples of input and expected output, without being explicitly programmed how to do so in a step-by-step sequence of instructions. This approach is closer to actual biological cognition: a child learns to recognize objects (such as cups) from examples of the same objects (such as various kinds of cups). It is today by far the dominant and most successful approach in AI.

10. Generally speaking, a Machine Learning method takes in an input of observations, and uses them to predict an output. Given a dataset of input and output pairs, the learning method will try to build a mathematical model that minimizes the difference between its predictions and expected outputs. By doing this, it tries to learn the associations/patterns between given inputs and outputs that can be generalized to new inputs not seen before.

11. To illustrate this learning process, let us consider the simplest approach to machine learning, a linear regression. Suppose that we want to learn how to correlate the height of the person with the size of her/his hand. We have a certain number of observations of height and hand size pairs (left table), represented as crosses in the figure below:

height (cm)	**hand size (cm)**
170	19.0
155	17.5
184	20.2
188	19.8
178	19.7
172	19.1
165	18.7
180	20.1
161	18.3
171	19.5
164	17.9
166	19.4

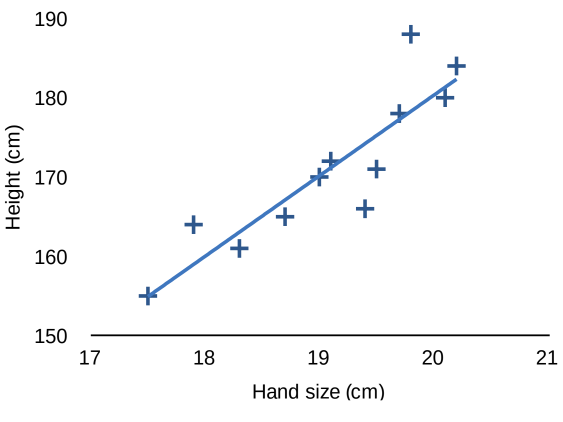

hand size =
0.098 * height + 2.23

if height = 180 cm,
hand size =
0.098 * 180 + 2.23
hand size = 19.9 cm

---

[2] WIPO Technology Trends 2019 − Artificial Intelligence, page 31. Machine learning represents 89% of AI-related patent filing and 40% of all AI-related patents. Within the machine learning technique, deep learning showed the annual growth rate of 175%, and neural networks grew at the rate of 46%, from 2013 to 2016.

12. Linear regression is a technique for finding a straight line between these points with the least possible error. The process for minimizing the error is the training. A mathematical method realizes this training by finding the straight line with highest proximity to the data points. Once this line with minimal error is found, the hand size of a person based on her/his height can be predicted. For instance, if the height of a person is 180 cm, the model will predict that it's hand size is 19.9 cm (see the right box).

13. Such a simple method is of course too restricted for learning more complex problems, involving for instance, more than two numerical variables. In the example above, it appears that age and gender should be added to the person height for more reliable hand size predictions. More sophisticated mathematical models are also used, in particular non-linear models that are not limited to straight lines.

14. Among those more advanced methods, Neural Networks (NN) offers a universal predictor, able to accept any kind of input. NN excel more particularly for solving tasks involving unstructured data as input, such as image or speech. As an advanced type of NN, Deep Learning is today booming as core technique in all AI patent applications.

B. Neural Networks

15. The fundamental building block of a NN is artificial neuron, also called *perceptron* or *node*. It was developed by Frank Rosenblatt in the 1950s and 1960s. A neuron takes *n* inputs, known as *features*, which are numerical representations of the data to be processed (pixels, words, signal, etc.). Each input is multiplied by a weight and sum-up (see the diagram, below). A bias *b* is added to this weighted summed combination. Finally, this value is passed to an activation function *f*.

16. For example, coming back to the example of predicting hand size, if the data on height, age and gender of a person are available, the artificial neuron will be as follow:

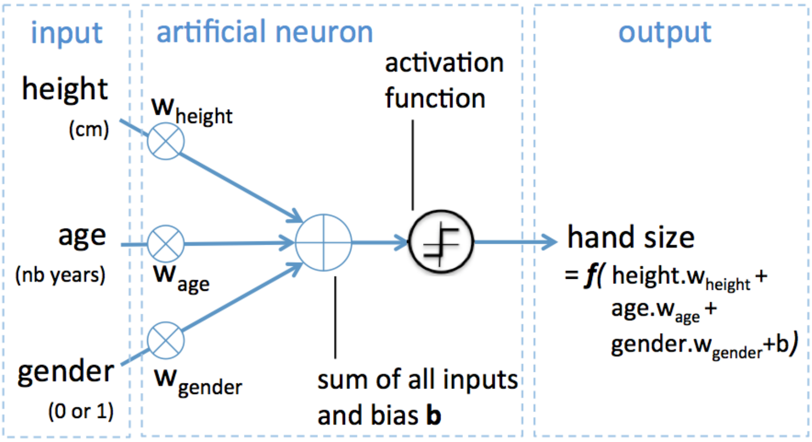

17. The weights capture the strength of the corresponding input features, in other words how much a particular feature influences the final results on its own.

18. The activation function models the "firing rate" of a biological neuron – propagating either a final signal or no signal. It takes the weighted sum of input and performs a certain fixed simple mathematical operation on it. One of the most commonly used activation function today is called ReLU (Rectified Linear Unit).[3]

---

3 The ReLU takes a number as input and returns the maximum of 0 or that number. For example, if the input is "1", the output will be "1", and if the input is "-1", the output will be "0".

19. An artificial neuron is a relatively simple function. It can be programmed in less than 25 lines of codes. A full neural network is then composed of at least three layers: an input layer, one or more hidden layers and an output layer. Input and output layers contain nodes performing no computation. They simply pass the numerical information to hidden layer for the input nodes, or transfer information from the network to the outside world for the output layer. Hidden layers contain artificial neurons as presented above. Nodes from adjacent layers have connections (or edges), shown in the arrows, between them.

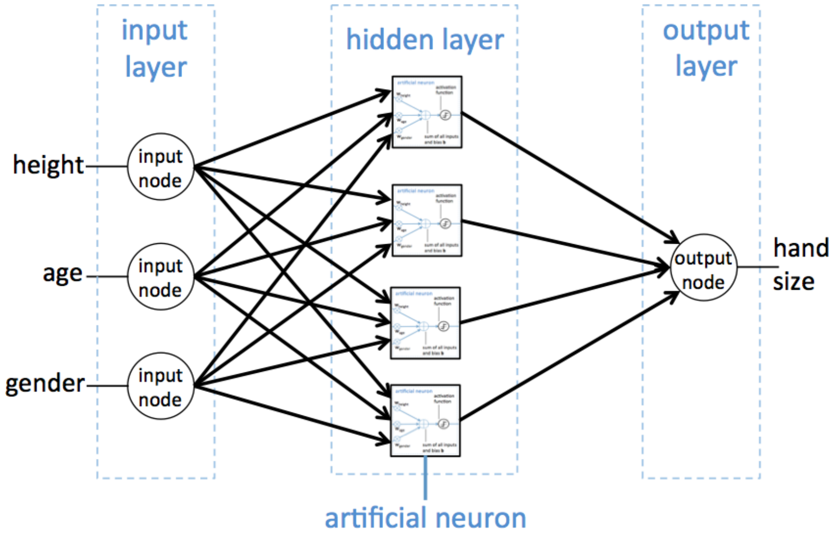

20. The input layer is filled with numerically encoded information, and then propagated forward through the hidden layers. The initial numerical values are modified by the neurons of the hidden layer and then propagated to the output layer corresponding to the final output. The number of output nodes matches the number of answers expected from the NN. For instance in this example, a single value, the hand size, is expected. The flow of data is here always forward through the layers.

21. Training a neural network consists of setting the parameters *weights* and the *bias* of all the neurons of the hidden layers to minimize the error observed on a set of examples, similarly as for the linear regression presented in section A, above. The mechanism for training a NN is basically "learning from mistakes". The training data consists of a number of input/output pairs. When a neural network is presented with an input, it makes a random "guess" as to what the corresponding output might be. It then sees how far its answer was from the actual output, and makes an appropriate adjustment back to its weights and bias. The process continues repeatedly with all input/output pairs until we reach optimum weights and bias.

22. It should be noted that the artificial neurons are only very loosely inspired from the mammalian biological neuronal structure and on a much lower scale. Biological neurons are considerably more complex and diverse than artificial neurons. A large number of factors (synaptic structure and geometry, type of neurotransmitter, etc.) have an effect on the signal propagation. A synapse for instance is composed of more than 2000 different proteins, presenting a large variety of physicochemical properties.[4]

---

4 "The differences between Artificial and Biological Neural Networks", Nagyfi Richárd, Blog entry at Toward Data Science, September 2018. https://towardsdatascience.com/the-differences-between-artificial-and-biological-neural-networks-a8b46db828b7.

C. Deep Learning

23. Although neural networks are known since the fifties, the usual number of hidden layers remained in practice only *one* until the years 2000s. The improvement of computational power has made possible in the last decade to increase (so "deepen") the number of layer of neural networks. For example, considering a cat or dog image classification problem (do we have a cat or dog on an image?), a deep neural network today looks as follow:

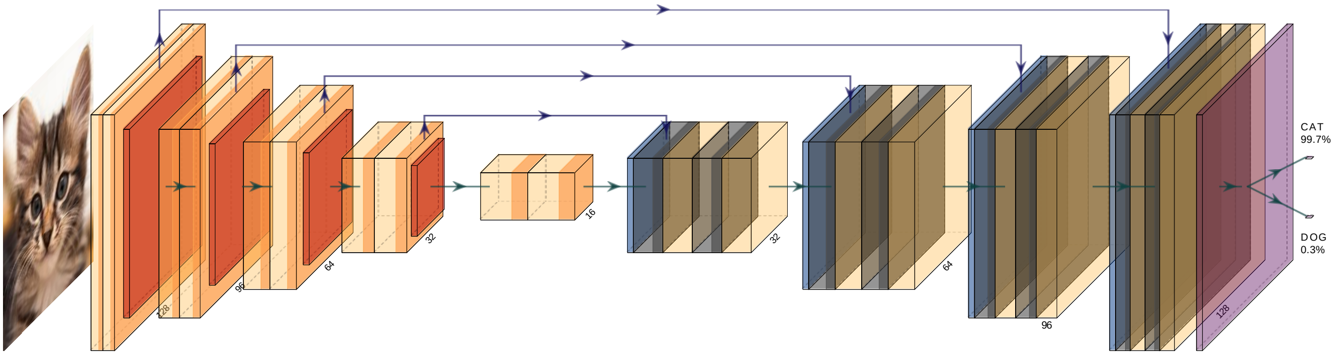

24. In the above example, we see a major shift in term of scale as compared to the simple NN previously described:

(i) The number of input nodes is very high; each input node receives the information of one pixel of an image. For a cat and dog image classification, we can use typically images of size 128*128 pixels, with each pixels defined by three values for Red, Green, Blue levels, i.e., 49,152 input nodes, and consequently 49,152 input features for *each* following neurons.

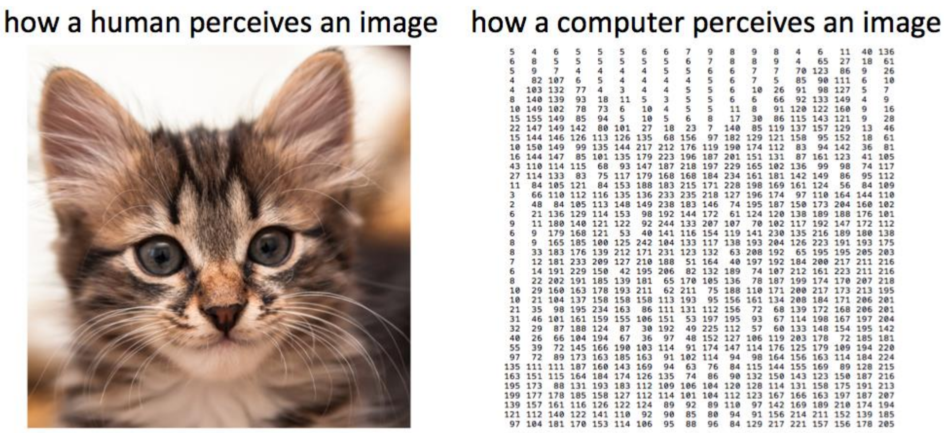

(ii) Multiple layers of neurons are introduced to process successively the input information. More than ten layers are not rare for image processing, each layer possibly containing hundreds of neurons, usually organized differently to provide particular advantages.

(iii) A typical deep neural network like this one can have several ten millions of weights and bias parameters to be set during the training, requiring ten thousand of labeled images.

25. Surprisingly, with an existing Open Source framework, such as Keras,[5] a trained data scientist can implement this Deep NN in less than 100 lines. With an online open dataset of dog and cat images, the network will achieve more than 93% accuracy of classification with commodity hardware: a level that is not far from the human performance (estimated around 95% for such a task).

26. The multiplication of layers introduces the notion of hierarchy for the representations and the process involved in the global prediction task. The first layers usually capture low level patterns in the input data (like lines, colored areas, etc. when processing an image), intermediary layers identify higher-level structures (like prototypical ears or muzzle of cats for the cat and dog classification) and finally the last layers self-specialize for performing the final prediction tasks based on identified structures.

27. Deep neural networks present several key properties, as compared to traditional neural networks, which explain their current success.

Discovery of features representations

28. Traditional machine learning uses features handcrafted by an engineer for solving a problem. For instance, for the prediction of the hand size, the ML engineer will select himself some features based on its own intuition and experiments, for instance the height, gender and age of the person. This step is called *features engineering*. A feature is an aspect of the data to be used by the ML algorithm to predict an output. This step is in general highly time-consuming, and when processing unstructured data (images, text, voice, videos), it is relatively inefficient.

29. For the first time in machine learning, Deep Neural Networks show a practical ability to discover automatically such features from raw data. By deepening the number of layers, neural networks both learn the useful features and how to use them to solve tasks. For example, for predicting the hand size, one would simply feed in a deep neural network the largest possible set of biometric measures, and let the network identifies automatically the ones to exploit for the final selection. Similarly for image classification, raw pixel data are sent to the network, which will identify patterns, like shapes of ears, tongues or teeth that are discriminant to decide if the input is a picture of a dog or a cat.

Data scale and deep learning performance

30. With traditional machine learning techniques, the performance quickly reaches a plateau as the amount of training data increases. It means that adding more training data is useless, after a while, the training algorithm somehow “saturates”. One of the key properties of deep learning is that the performance continuously increases with an increase in the training data. This property explains why the largest networks existing today in machine vision could use as many as 15 million images for training.

---

5 Keras: The Python Deep Learning library, François Chollet and others, 2015-2019. https://keras.io/.

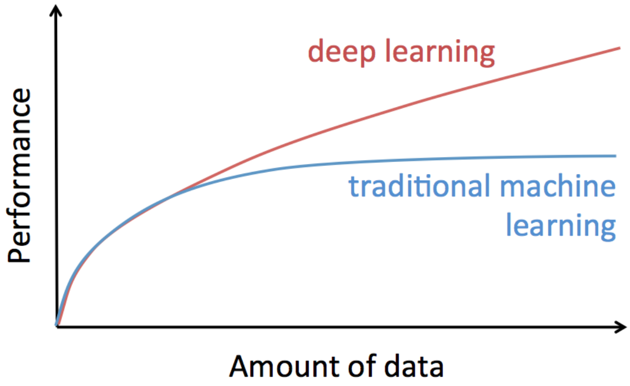

31. Mathematically, artificial neural network models can be understood as just a set of matrix operations and finding derivatives.[6] With the increase of computational power, deep learning can surpass any other ML approaches, as long as a massive amount of training data is available.

D. The Limits of Deep Neural Networks

Deep neural networks are black box

32. Contrary to more classical algorithms, the decision process captured by a neural network during the learning process cannot be explicitly expressed in a comprehensible form for a human. As mentioned earlier, a Deep Neural Network could learn itself useful features in data. For instance, for the dog and cat classification task, the network could identify prototypical ears or muzzle of a cat. But in practice, most of the time, these features are not interpretable by a human. These patterns emerge from the numerical optimization process in the hidden layers, and are not accessible to our interpretation.

33. In addition, it is not possible to exhibit an equation or the coefficients defining a relation between an input and an output in term of standard mathematics. The network is the final equation of the relation, possibly involving hundred million parameters. Such a complex decision process cannot be illustrated with a flowchart or any kind of traditional methods to represent algorithms. This explains why it is often said that the deep neural networks are "ultimate" black box. The training itself is realized by the NN *on its own* and the resulting network is enormously complex.

Deep learning requires a lot of data

34. A surprising observation is that neural networks and deep learning are amongst the simplest machine learning models in terms of involved mathematical modeling. It is often said that the underlying mathematics is accessible to a good high school student. Still, they provide today by far the best results. The reason is that they are the most adapted to take advantage of very large training dataset. The success of deep learning is today much less related to theoretical progress than pure increase of computational power and availability of massive human behavioral data: something often called *brute force*.

35. The immediate limits of DL are related to the cases where brute force is not possible. This covers in particular tasks with no or limited training data (e.g., processing rare human languages, drug discovery for rare diseases, etc.) or domain with legal restrictions.

---

[6] Running such mathematical calculations can be highly optimized for vector processors (doing the very same calculations on large amounts of data points over and over again) and speed up by magnitudes using GPUs (Graphical Processing Unit, the same used for speeding-up video games) or new dedicated hardware.

Real world data is biased

36. The success of Deep Learning depends on the availability of a large volume of data, but this dependency on massive datasets also creates several issues:

– *Data bias*: Data collection at scale is often not neutral, some groups in relation to age, gender, and ethnic origins being under or over-represented.[7] Bias can come from the data collection technique, from existing social bias, or from the fact that people who create datasets and the models were not a diverse group.

– *Bias amplification*: By nature, the machine learning training methods tend to identify discriminant patterns in data for quickly increasing the quality of prediction. Consequently, they not only learn our actual bias, but oftentimes, they also amplify our bias.

– *Lack of reproducibility*: As a model depends on a unique composition of training data, reproducing some claimed results is only possible in the very rare case of open data.

Deep learning still require a lot of human efforts

37. Although Deep Learning is able to learn which features should be used as explained in section C, human efforts are still necessary for creating a Deep Neural Network model in many areas. For example:

– creation of the network architecture (which type of layers, which order of layers, etc.);

– determination of the best parameters (number of neurons per layer, size of input, etc.);

– selection of the resulting classes; and

– decision on how to encode the input in numerical format.

However, the largest effort is by far the creation of the training data.

38. The most common form of machine learning today is *supervised learning*. The examples previously presented all belong to this category. The training data is a set of input and output pairs, where the output pair is the expected response to the input. The learning is thus guided, as a supervisor teaching the algorithm what conclusions it should come up with. Labeling manually thousands or millions of examples is a considerable effort often necessary to reach good accuracy. Another consequence is that supervised learning works only on narrow problems involving very limited decisions, for example, detection of melanoma on medical images. However, it cannot easily address more open tasks, such as diagnostic implying a larger variety of input and tailored decisions, or involving some general reasoning.

39. In contrast, with *unsupervised learning*, a computer can learn to identify processes and patterns without a human to provide guidance along the way, assigning its own new labels to the data groups it created. Unsupervised learning reduces the role of the human by avoiding

[7] Amazon scraps secret AI recruiting tool that showed bias against women, Jeffrey Dastin. Reuters Business News, Oct. 2018 (https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G).

both the selection of labels and the very costly labeling of examples in the training data. This is closer to human learning, which is largely unsupervised: humans discover the structure of the world by observing and interacting with it, not by being told the name of every object.

40. Deep Neural Networks can learn both in a supervised or unsupervised way. However, unsupervised learning performs today significantly worse than supervised learning. Using a small amount of manually labeled data is enough to surpass an unsupervised learning even with a huge amount of unlabeled data. Some recent methods require less supervision (active learning, transfer learning, reinforcement learning), and in the longer term, unsupervised learning is expected to become more important.

E. Where Does Innovation Take Place Today in Deep Neural Networks?

41. While the principles of Deep NN are relatively simple and generic, innovative works in Deep Learning today embrace a larger scope than the core NN aspects:

– *Training data*: As training volume is the most impactful factor in DL, innovating on how to best create, exploit or reduce datasets for particular applications is a major challenge;

– *Computational power*: More computational power leads in practice to better models. A lot of efforts in ML are focusing on hardware and software optimization;

– *Application*: AI technologies can be applied to multiple fields for performing various functions. What are the problems and new functional tasks where Deep Learning can be successful? How to integrate efficiently these techniques into larger applications?

– *Neural network architecture*: In practice, different types of hidden layers exist, with different properties, like Recursive Neural Networks adapted to sequential data (speech recognition, translation, etc.), or Convolutional Neural Networks more adapted to object recognition in images. Designing the best Deep NN architecture is complex, because it depends on the task, the nature of the data, the domain and on the amounts of available training data.

– *Robustness*: Deep Neural Network can be relatively easily fooled by adversarial attacks,[8] where a second Deep NN is competing against a first one to identify its weaknesses. The safety and reliability of such ML systems will be critical in the next years.

42. Overall, the organizations with the largest datasets and computing power have a considerable advantage for developing the leading AI systems, independently from core technical innovations. In general, the core technical innovations are made available very early in Open Source software distribution.

III PATENT PROTECTION OF AI-RELATED INVENTIONS

43. This Part of the document looks at patent protection of AI-related inventions. The "AI-related inventions" may take different forms. Innovation may occur in the improvement of AI techniques, while they may take place through integration of the AI technology in existing

[8] Researchers design patch to make people 'virtually invisible' to AI detectors, April 2019. https://www.computerworld.com.au/article/660283/researchers-design-patch-make-people-virtually-invisible-ai-detectors/.

devices in order to improve its functionality or add a new feature. In addition, the AI technology can be used as a tool for R&D to create a new invention. Implication of the AI technology to patent law may not necessarily be the same among those different forms of the AI-related inventions.

A. General Considerations

44. It is widely recognized that the patent system should contribute to the promotion of technological innovation as well as to the transfer and dissemination of technology, for the benefits of the society at large, through balanced rights and obligations of technology producers and users of technological knowledge. To this end, each country provides a legal framework and enacts laws and regulations, which are interpreted by courts and supplemented by practical guidance developed by the administrative body.

45. As the patent system is technology neutral, whenever a new technology emerges, a question is often raised as to whether the purposes of the patent system could continue to be served. It has been the case for semiconductor technology, computer software, information technology and biotechnology: the debates continue as technology develops. It is therefore not surprising that the emergence of AI has raised similar questions and debates, scrutinizing the readiness of the current patent system to accommodate the AI technology.

46. For decades, computer technology, covering both hardware and software, has been utilized to assist the invention creation process of humans in many fields of technology. For example, developments in the mechanics and electronics have been assisted by computer-aided designs (CAD), bioinformatics has facilitated researchers to analyze and interpret biological data, and computational chemistry has helped chemists to find new chemical substances. Computers have also been integrated into devices and apparatus, to perform a specific function.

47. In the case of computer technology, new inventions relating to that technology may be categorized into three types:

(i) new inventions that improve the computing functions of computers as such;

(ii) new inventions (a device, an apparatus etc.) that incorporate computers to carry out a specific function; and

(iii) new inventions created through the assistance of computers, which can be in any field of technology.

48. A similar kind of categorization may be possible for AI technology:

(i) new inventions on the core AI technology itself;

(ii) new inventions that incorporate the AI technology (for example, a translation device incorporating AI deep learning, and a medical device for diagnosing a specific disease); and

(iii) new inventions created with the assistance of the AI technology (for example, a new material found with the assistance of the AI technology).

49. At the current stage of the technological development of AI, instructions and interventions by humans are still an important part in the creation process of those inventions. As explained in Part II, unsupervised learning performs significantly worse than supervised learning so far. However, as the AI technology develops,[9] the level of necessity or relevance of human intervention in the creation process might be diminished relative to increased autonomous performance of an AI system.

50. Therefore, the AI-related inventions may be understood from another angle, focusing on the creation of a core inventive concept. From that perspective, AI-related inventions may be categorized as follows:

(i) identification of a problem and conception of a solution are made by humans, while the AI technology is used for mere verification, automation, adaptation or generalization of the human solution;

(ii) identification of a problem is made by humans and conception of a solution is assisted, guided or led by the AI technology; and

(iii) identification of a problem and conception of a solution are made by the AI technology without any human intervention.

In the second scenario, the relevance of the AI technology in the invention creation process may be from minimum to determinative. The third scenario, i.e., artificial general intelligence or superintelligence,[10] is not something that the current technology permits.[11] Nevertheless, the possibility of such development marks a significant difference from the conventional computer technology. Such a difference leads to new questions of a different nature when it comes to AI patenting.

51. Since the emergence of the AI technology, innovators and researchers have filed patent applications, and patents have been granted, on those inventions. As illustrated in the "WIPO Technology Trends 2019 − Artificial Intelligence", they cover various AI techniques[12] for numerous AI functional applications[13] in a variety of AI application fields[14]. Open source (or open innovation) approaches are also popular among AI developers.[15] For the detailed patent landscape data on AI-related inventions, reference is made to the said WIPO publication.

52. How the AI technology affects patent laws has not been determined yet. However, certain characteristics of the AI technology seems to hint the areas in the patent laws that might be impacted by this emerging technology in the future, if not immediately. Thoughts may be given to the points, such as:

(i) Since the AI technology is primarily implemented by software, current patent law issues surrounding the computer implemented inventions and inventions using software may continue to be relevant to the AI technology;

---

[9] Increase of computation power has allowed AI machines to manage a large search space: for example, the chess game involves 10^{47} possibilities (Deep Blue, February 10, 1996) and the go game involves 10^{170} possibilities (AlphaGo, March 2016).

[10] It means that AI systems are able to successfully perform any intellectual tasks that could be undertaken by the human brain, or the hypothetical ability of a machine far surpasses the human brain.

[11] WIPO Technology Trends 2019 − Artificial Intelligence, p.19.

[12] For example, machine learning, fuzzy logic and logic programming.

[13] For example, computer vision, natural language processing and speech processing.

[14] For example, transportation, telecommunication and life and medical science.

[15] WIPO Technology Trends 2019 − Artificial Intelligence, p.109.

(ii) The cognitive characteristics of the AI technology call for further thoughts on how this technology might be integrated into the human innovation processes, and on its implication to the assumption of “human-made” inventions under the patent system and patent law;

(iii) The inherent technical limitations in fully reproducing and describing the processes carried out in the deep learning neural network draws our attention to their potential impact on one of the fundamental principles of the patent system, that is, dissemination of new technological knowledge.

53. So long as the rationale of the patent system is to contribute to the promotion of technological innovation as well as to the transfer and dissemination of technology, patent system needs to continue providing incentives for innovation and mechanisms for sharing new knowledge in the field of AI as well (unless there are other legal/social/economic tools that sufficiently address these matters). At the policy level, the main considerations could be: with a view to the objective of the patent system, would the development of AI technology distort the balance sought by the patent system? If so, how that could be restored? Is it useful to update patent laws and practices in light of the development of AI technology? Are there, or will there be, any gaps between the existing legal concepts of the patent system and the emergence of the AI?

54. In order to answer those questions, there is a need to understand the technical specificity of AI compared with the conventional computer technology, and to evaluate how the current law and practice might possibly applied to the AI technology today, and beyond. This background document does not attempt to describe the full set of issues in a comprehensive manner. However, the following paragraphs provide a sample of patent law issues that may be relevant, where patent protection is sought, and patents are granted, on the AI-related inventions. The term “AI-related inventions” refers to various kinds of inventions as described in paragraphs 48 and 50, above. At this point, there are very few official guidance that specifically address patent law questions applied to AI-related inventions. As AI being a new technology, case law has not been fully developed, and a few patent offices have issued guidance, clarifying its practices in this field. Enforcement and licensing of AI patents against the backdrop of claim interpretation, might also be part of the future discussion items, along with more commercialization of AI-integrated products in the market. In general, negotiating licensing agreements and solving patent disputes require complex and multi-faceted considerations. It still need to be seen whether AI-related inventions *per se* would bring additional complication to such already complex questions.

55. The current patent system is built on the assumption that certain incentive mechanisms would promote creative activities by humans. From the high-level policy perspective, potentials shown by the development of the AI technology pose a legal philosophical question on the incentive theory of the patent system. Although it is still a science fiction, this may be particularly so once an AI-machine is capable of comprehensively processing various data (not only scientific and technological data but also personal and behavioral data as well as social and legal data), identifying a problem, solving the problem with a new invention and producing new products to the market to satisfy humans, all being done autonomously. While it may be an intellectually interesting question, it is well beyond the scope of this document.

B. Patentable Subject Matter

56. In general, patents shall be available for any inventions, whether they are products or processes, in all fields of technology, provided that they meet all the legal requirements, including the requirement that the inventions do not fall under the excluded subject matter. There is no international mandatory definition of the term “invention” and a national laws define

the scope of the excluded subject matter, in line with the international treaties to which the country is a party. Consequently, there are differences in the scope of patentable subject matter from one country to another.[16] Many countries exclude from the patentable subject matter mathematical methods, schemes, rules and methods for performing mental acts, business rules and methods and programs for computers. Some of them clarify that those subjects are excluded from the patentable subject only to the extent that a patent application relates to such subject matter as such. In one jurisdiction,[17] the case law establishes that claims directed to law of nature, natural phenomena and abstract ideas are excluded from patent protection. In another jurisdiction, its patent law[18] defines the term “invention” as “the highly advanced creation of technical ideas utilizing the laws of nature” and the category of a product invention includes a computer program and any other information that is to be processed by an electronic computer equivalent to a computer program.[19]

57. Beyond the improvements on hardware components that run the AI functions, inventions relating to AI techniques and AI functional applications are mostly directed to software. As in the case of conventional computer technology, AI applications can also be used in non-technology areas, such as finance, insurance, e-commerce etc. In addition, machine learning is based on computation models and algorithms for classification, clustering, regression and dimensionality reduction, which may be considered mathematical techniques. Furthermore, while the importance of training data for the performance of machine learning cannot be denied, data *per se*, which is mere information, is not a patentable invention.

58. Patent eligibility of computer-implemented inventions or software-implemented inventions has already been one of the areas that are difficult to draw a clear-cut line between eligible and non-eligible subject matter. For example, in many countries, the “technicality” of the claimed invention is considered important for the determination of patent eligibility. In those countries, case law and office practices have been developed to clarify the concepts such as the “technical problem”, “technical means”, “technical effects” and “technical purpose”. In the United States of America, in order to apply the U.S. Supreme Court’s decision to the evaluation of patent eligibility (the *Alice/Mayo* test), the United States Patent and Trademark Office (USPTO) issued the 2019 Revised Patent Subject Matter Eligibility Guidance in January 2019 with a view to increasing the clarity of the methodology.[20] The patent eligibility of software-implemented inventions, however, involves complex questions, which may continue to evolve with further technological development.

59. As to the patent eligibility requirement applied to AI-related inventions, some patent offices issued guidance pertaining to AI-related inventions. The USPTO’s 2019 Revised Patent Subject Matter Eligibility Guidance includes one example that specifically discuss patent eligibility of a computer implemented method of training a neural network for facial detection comprising a series of steps for such training.[21] In the November 2018 edition of the Guidelines for Examination in the European Patent Office (EPO), under the sections in respect of the patentability of mathematical methods and schemes, rules and methods for performing mental acts, playing games or doing business, new sub-sections relating to, *inter alia*, artificial

---

16 See “Certain Aspects of National/Regional Patent Laws – Exclusions from patentable subject matter” at: https://www.wipo.int/scp/en/annex_ii.html.

17 The United States of America

18 Section 2(1) and (4) of the Japan Patent Act.

19 For more information about exclusions from patentable subject matter and patent eligibility of computer-implemented inventions, see SCP/13/3 and SCP/15/3 (as regards computer programs as excluded patentable subject matter, see, in particular, Annex II of SCP/15/3).

20 2019 Revised Patent Subject Matter Eligibility Guidance, available at: https://www.uspto.gov/patent/laws-and-regulations/examination-policy/subject-matter-eligibility.

21 2019 Revised Patent Subject Matter Eligibility Guidance, Example 39.

intelligence and machine learning have been created in order to define the relevant patentability criteria more precisely.[22] The Examination Handbook for Patent and Utility Model, issued by the Japan Patent Office (JPO), also includes examples relating to AI inventions.[23]

60. Regarding the inventions created with the assistance of AI technology, the consideration of the patentable subject matter obviously depends on the nature of the final invention and how it is claimed. For example, in countries where plants are excluded from patentable subject matter, patent claims defining a new and innovative plant, created by the assistance of an AI tool, would not be patentable.

C. Novelty and Inventive Step

61. It is said that inventive step analysis is the most difficult requirement in the patentability criteria to assess.[24] Among the rejected patent applications, many of them are rejected on the ground of lack of inventive step. When the validity of patents is challenged by third parties, they often base their arguments on non-compliance with the inventive step requirement. Patent applications and patents in the field of AI appear to be the same. Although the available data is limited, among the oppositions filed by third parties in relation to AI-related applications/patents, many of them are on the grounds of lack of inventive step (obviousness).[25]

62. Oftentimes, when new technology emerges, assessment of inventive step faces a particular challenge. This is because prior art references are scarce, and the determination of the exact scopes of the hypothetical person skilled in the art, and of the general common knowledge in that particular art, have not fully been established. Lack of case law and official guidance makes it difficult to assess inventive step in a consistent manner. However, as the technology manures, common interpretations and standard practices have gradually emerged in many technology areas.

63. Since the assessment of inventive step is made by a person skilled in the art, the determination of the level of knowledge and skill possessed by this hypothetical person is one of the cornerstones of the inventive step assessment.[26] The exact level of such knowledge and skill needs to be defined for each concrete individual case. It also changes with the technological development. In general, the capacity and knowledge of a hypothetical person skilled in the art can, where appropriate, correspond to those of a team of persons working in various relevant fields.[27] Therefore, it is expected that the more an AI tool is used in the relevant art, the less innovative such use would become, since a person skilled in the art, i.e., an interdisciplinary team able to use the AI tool, would turn to the usage of such a tool in its research. The similar consideration applies to the notion of the “common general knowledge”.[28]

---

22 Guidelines for Examination in the European Patent Office (EPO), Part G, Chapter II, 3.3.1. In essence, the Guidelines state that Artificial intelligence and machine learning are based on computational models and algorithms for classification, clustering, regression and dimensionality reduction, which are *per se* of an abstract mathematical nature, irrespective of whether they can be "trained" based on training data. However, if Artificial intelligence and machine learning find applications in various fields of technology, making a technical contribution and supporting the achievement of a technical purpose, such invention may be considered patentable subject matter.

23 Annex A of the Examination Handbook for Patent and Utility Model. As regards the patent eligibility, the examples discussed are: claims directed to data that is mere presentation of information; a data structure that enables information processing, which can be performed in voice interactive systems; a trained model for analyzing reputation of accommodations.

24 For more information about how the inventive step requirement is implemented in different countries, see SCP/22/3, SCP/28//4, SCP/29/4 and SCP/30/4.

25 WIPO Technology Trends 2019 − Artificial Intelligence, p.115 to 117.

26 See document SCP/22/3.

27 Document SCP/22/3, paragraphs 34 and 35.

28 See document SCP/28/4.

64. Annex A of the Examination Handbook for Patent and Utility Model, issued by the Japan Patent Office (JPO), contains several examples relating to assessment of inventive step for AI-related invention.[29] For example:

– lack of inventive step, because the invention merely systematized human operations in an AI system (Example 33);

– lack of inventive step because of a mere modification of a method for estimating output data from input data (Example 34);

– involvement of inventive step, because adding certain training data presents a significant effect (Example 34);

– lack of inventive step, because a modification of training data for machine learning is a mere combination of known data, without any significant effect (Example 35); and

– involvement of inventive step due to certain pre-processing of training data (Example 36).

65. In relation to inventions "invented" by AI machines, concerns about massive creation of "new inventions" by AI machines have been raised, with a fear that it would lead to the situation where everything would be invented by the machine and patented. Somewhat mirroring the above, there are projects to generate "prior art" using the AI technology by publishing the outputs of AI machines, so that any of such output would no longer be patentable by others.[30] As to the new inventions, the enabling disclosure requirement and industrial applicability (utility) requirement would prevent patenting of, for example, a mere combination of known chemical elements without any description of how such a compound can be produced and how it can be used. Similarly, information described in a published reference can only be regarded as having been made available to the public, and thus an eligible prior art reference, if the information is described in sufficient detail to enable a person skilled in the art to practice the teaching. A chemical structure disclosed merely in the form of a chemical formula, for example, is most likely not considered as an eligible prior art reference to deny the novelty/inventive step of the corresponding chemical compound.

66. The rationale of the inventive step (non-obviousness) requirement is that patent protection should not be given to an invention that could be deduced as an obvious consequence of what is already known to the public, since it would contribute very little to the society.[31] Such a policy objective may guide the determination of the inventive step for each case, including the AI-related inventions.

D. Sufficiency of Disclosure and Claims

67. Similar to the inventive step assessment, new technologies pose particular challenges to disclose inventions in a clear and complete manner, and to draft clear and concise claims that adequately cover the scope of legitimate protection. Lack of case law and official guidance makes it also difficult for the IP offices and users of the patent system to assess the compliance with the disclosure requirements.

---

29 Annex A of the Examination Handbook for Patent and Utility Model, Examples 31 to 36, JPO.
30 All Prior Art project (https://allpriorart.com/about/).
31 Document SCP/22/3, paragraph 3.

68. Regarding the description of the claimed invention, in general, national/regional patent laws require that an applicant for a patent shall disclose the invention in a manner sufficiently clear and complete for the claimed invention to be carried out by a person skilled in the art (enabling disclosure requirement).[32] It is through this requirement that the patent system facilitates the dissemination of information and access to technological knowledge contained in patent applications and patents. This results in the expansion of public stocks of technological knowledge and an increase in the overall social benefits, for example, inducing the technology transfer and avoiding a duplicative R&D.

69. In relation to the AI technology, a question may be to what extent an AI algorithm, a training model, a neural network architecture, a learning process, training data, hardware components etc. should be disclosed in a patent application in order to meet the enabling disclosure requirement. One of the challenges may come from the fact that, under the current deep-learning technology, it is problematic for humans to identify each process step taken in a deep learning neural network and to explain exactly how the neural network arrives at the final result. When a system has several ten millions of weights that contribute to a classification, it is too complex to express it in a human comprehensible form. In certain cases, it may be more difficult to rationalize the AI output (i.e., to provide reasoning in a credible way) without having a real-world experimental data.

70. At the same time, the extent of the disclosure of the claimed invention in the description part of a patent application obviously depends on what is claimed in the claims part of the application. For example, in case where an invention relates to the application of the AI technology to solve a problem by training a deep learning algorism with a specific dataset, if the claimed invention encompasses broader application, not one type of dataset but all dataset types that are necessary for a person skilled in the art to carry out the broad scope of the claimed invention may be required in the description.

71. In this regard, the notion of a person skilled in the art is also important for the assessment of the enabling disclosure. For example, if an AI technology is applied to an invention in a specific field (for example, an image recognition neural network applied to an invention in the field of security and surveillance), a team of persons skilled in the art in the AI technology and in the surveillance area may constitute a hypothetical person skilled in the art for the assessment of such invention.

72. Another issue might arise from the fact that deep learning technologies are non-deterministic: they involve some randomized initialization. Therefore, even the same training data and the same neural network architecture might lead to slightly different performance of machine learning. Two training of a model with the same training data and same neural network architecture will result in two slightly different training behavior. Similar to the cases of biological materials where biological variability is unavoidable, a consideration might be given to the so-called reproducibility or plausibility of the claimed inventions based on the disclosure in a patent application.

73. In relation to the training data, solving a problem with one particular AI-technique might require a particular dataset. The important role that a training dataset plays in the performance of the deep machine learning might raise questions as to the extent of its disclosure in a patent application and to the availability of such a dataset with a view to verify the claimed invention by third parties (i.e., whether the claimed invention actually works or not).

---

[32] See document SCP/22/4. See also “Certain Aspects of National/Regional Patent Laws – Sufficiency of Disclosure” at: https://www.wipo.int/scp/en/annex_ii.html.

74. As regards the claims, many national laws stipulate that the claims shall be clear and concise. In addition, the claims shall be supported by the description (support requirement).[33] In general, the rationale of this requirement is that the claimed invention should not exceed the scope of the invention disclosed publicly in the description. Similarly, the essential policy goals of the written description requirement provided under the law of the United States of America[34] is "to clearly convey the information that an applicant has invented the claimed subject matter and to put the public in possession of what the applicant claims as the invention".[35] Accordingly, those requirements point to the fundamental principle that patent protection shall not be accorded to what has not been invented by the applicant as of the filing date and what has not been shared with the public through the disclosure in the patent application as of the filing date. Since the AI-related inventions are mostly computer-implemented inventions, as to the techniques of claiming AI-related inventions, applicants may face similar challenges in properly covering their inventions in the claims.

75. Regarding the application of the disclosure requirements to AI-related inventions, Annex A of the Examination Handbook for Patent and Utility Model, issued by the Japan Patent Office (JPO), contains several concrete examples.[36] The examples primarily illustrate the cases where the AI technology is applied to inventions in various fields of technology, and thus the machine learning generally requires multiple types of training data. They discuss the importance of showing a certain relationship (such as a correlation) among those data in order to fulfill the disclosure requirements. In addition, one example discusses the case where the AI technology is presumed to provide a certain function to a product invention claimed. The claimed invention does not meet the disclosure requirement, since the description only provide the AI inference data (no experimental data of the product) in the description, and neither prior art nor the general common knowledge suggest that the AI inference data be able to substitute the experimental data.

E. Industrial Applicability

76. In relation to the reproducibility and plausibility of the claimed inventions, in some countries, the compliance with the industrial applicability requirement may also necessitate the claimed invention to be reproducible with the same characteristics, whenever necessary.[37]

F. Inventorship and Ownership

77. Article 4*ter* of the Paris Convention states that the inventor shall have the right to be mentioned as such in the patent. This provision refers to what is commonly called the "moral right" of the inventor to be named as such in the patent granted for his invention in all countries of the Paris Union. It is generally understood that the inventor can waive such right, unless national legislation prescribed otherwise. As the Paris Convention does not define the term "inventor", the identification of an inventor/inventors as well as the procedure for the exercise of such moral right is regulated by each Member State in its applicable law.[38]

---

33 See document SCP/22/4.
34 Section 112(a) of Title 35 of the United States Code. See document SCP/22/4.
35 Ibid.
36 Annex A of the Examination Handbook for Patent and Utility Model, Examples 46 to 51, JPO.
37 SCP/5 Informal Paper (The Practical Application of Industrial Applicability/Utility Requirements under National and Regional Laws). See also the Case Law of the Boards of Appeal of the European Patent Office, Part I.E.2.
38 Guide to the Application of the Paris Convention for the Protection of Industrial Property, G. H. C. Bodenhausen (WIPO Publication No. 611).

78. Although the patentability requirements (such as the patentable subject matter, novelty, inventive step (obviousness), industrial applicability (utility) and disclosure requirements) are independent from the question of inventorship, false indication of inventors may have serious legal consequences.

79. While not all national legislations define the term “inventor”, considering the rationale of the patent system and the moral right being one of the fundamental rights associated to patent rights, there might have been a general presumption that an inventor(s) under patent law is presumed to be a person(s).[39] If this presumption is valid, the logical consequence might be that regardless of the level of contribution by the AI machine to the conception of the invention, the machine is not an inventor.

80. Where the invention creation process involves the use of an AI system, as long as a person (or persons) in that process qualifies as an “inventor” under the applicable law – broadly speaking, contributing to the conception of the claimed invention – that person (or persons) would be an inventor (or inventors) of that invention, be it an AI programmer, an AI developer, an AI user or otherwise. A question, theoretical at this point, is if no person would qualify as an inventor under the applicable law, who has the right to a patent?

81. While it is expected that AI machines would possess higher cognitive abilities with the technological advancement, the evolution of technology is often incremental. In addition, the AI technology might play a different role in the invention creation process, depending on each case, i.e., any role within the range from a mere assisting tool to a means that is instrumental for the perception of the inventive concept. Therefore, setting “inventions by humans” against “inventions by a machine” appears to be too simplistic for the complex discussion on inventorship issues.

82. In general, the right to a patent belongs to an inventor (or inventors) at the first place, while the inventor(s) can assign the right to another natural or legal person. In many countries, where an invention is made under employment, the right to a patent, in principle, belongs to an employer, often under certain conditions.[40] Therefore, the inventorship/ownership issues may be part of the essential policy questions for the designing of a patent system.

IV. AI TECHNOLOGY AS A TOOL IN THE PROSECUTION AND ADMINISTRATION OF PATENT SYSTEMS

83. The AI technology solutions may be used in patent proceedings and beyond, i.e., as a tool to file patent applications by applicants, to process patent applications by patent offices, to enforce patents by patentees, to invalidate patents by third parties, to resolve disputed by judiciaries etc.

---

[39] According to 35 U.S.C. §100(f), an “inventor” is “the individual or, if a joint invention, the individuals collectively who invented or discovered the subject matter of the invention”. In the United States of America, the inventor, or each individual who is a joint inventor, of a claimed invention must, in principle, execute oath or declaration directed to the application.

[40] For the sake of completeness, it should be also added that the right to a patent may also be transferred to another person through inheritance.

A. Tools for the IP Authorities

84. IP offices have already started to use AI technology to facilitate IP administration and delivery of their service. The WIPO Index of AI Initiatives in IP Offices[41] is an on-line portal on which such use of the AI technology is searchable by country/territory and by business application of AI. The categories of business applications in the Index, which are the major business areas of IP Offices' work facilitated by the AI technology, are: (i) digitization and process automation; (ii) examination; (iii) helpdesk services; (iv) image search; (v) machine translation; (vi) patent classification; (vii) patent prior art search; and (viii) trademark classification.

85. During the WIPO Meeting of Intellectual Property Offices (IPOs) on ICT Strategies and Artificial Intelligence (AI) for IP Administration, held in Geneva, from May 23 to 25, 2019, one of the main themes was how the applications of AI and other advanced technologies had been, and could be, used by IP Offices.[42] The discussions at the Meeting indicated the progress that had been made in various offices to harness the potential of AI in the IP administrative systems, and demonstrated the desire from offices for an ongoing exchange of information and experience in AI, which would also avoid, *inter alia*, a duplication of efforts.[43] As a follow-up to the Meeting, WIPO established a dedicated web page on AI[44] and an electronic forum for the discussion of ICT strategies and AI for IP Administration, which is restricted to the experts nominated by IP Offices. Furthermore, the Committee on WIPO Standards (CWS) established a Task Force on ICT Strategy and Standards, which, *inter alia*, reviews the Recommendations presented at the Meeting.[45]

86. In the field of patent administration, national and regional patent offices have developed (or have been developing) the AI application tools for: classification of patent applications; formality check; prior art search; machine translation of relevant documents; assistance to substantive examination (for example, automatic annotation of patent literature and automatic detection of exclusions from patentable subject matter); and more generally, data conversion and document management.[46]

87. The International Bureau of WIPO has also used AI for its work in order to enhance functions and processes at the Organization. WIPO currently uses AI in three main areas: machine translation (WIPO Translate); image search within the Global Brand Database; and automatic patent classification.[47]

B. Tools for Applicants, Third Parties and IP Professionals

88. Considering the ever-increasing amount of publicly available information generated through the patent system, the AI technology may also assist applicants, third parties and IP professionals for achieving higher quality and efficiency in their respective activities.

---

41 https://www.wipo.int/about-ip/en/artificial_intelligence/.
42 Documents and presentations of the meeting are available at: https://www.wipo.int/meetings/en/details.jsp?meeting_id=46586.
43 Document WIPO/IP/ITAI/GE/18/5 (Summary by the Facilitator).
44 https://www.wipo.int/about-ip/en/artificial_intelligence/.
45 Document CWS/6/3.
46 WIPO Index of AI Initiatives in IP Offices.
47 For detailed information, please visit the WIPO website at: https://www.wipo.int/about-ip/en/artificial_intelligence/.

89. AIPPI, AIPLA and FICPI consider that the applications of AI in IP practices can be grouped into three categories: (i) document automation; (ii) process automation; and (iii) AI-enabled insights.[48] They predict that AI-document automation would be able to look at language in context, and assist, for example, application drafting and proofreading. AI-based process automation would leverage patent data for search purposes, and would be used for docketing, generating office action shells and creating and managing information disclosure statements. AI-enabled insights would provide users of the patent system with insights and predictions, which they may use to make better-informed decisions.

[Annex follows]

48 The AIPLA/AIPPI/FICPI AI Colloquium Primer, available at: AIPPI/AIPLA/FICPI Joint Colloquium on Artificial Intelligence, March 28 and 29, 2019 https://ficpi.org/colloquium.

REFERENCES TO CONFERENCES ORGANIZED BY WIPO AND/OR IP OFFICES AND THEIR WEBPAGES AND PUBLICATIONS DEDICATED TO AI

WIPO

WIPO Technology Trends 2019 − Artificial Intelligence (WIPO Publication No. 1055E/19)
https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf.

"Artificial Intelligence and Intellectual Property" webpage
https://www.wipo.int/about-ip/en/artificial_intelligence/

Meeting of Intellectual Property Offices (IPOs) on ICT Strategies and Artificial Intelligence (AI) for IP Administration, May 23 to 25, 2018
https://www.wipo.int/meetings/en/details.jsp?meeting_id=46586

WIPO Conversation on Intellectual Property (IP) and Artificial Intelligence (AI), September 27, 2019
https://www.wipo.int/meetings/en/details.jsp?meeting_id=51767

Argentina

Seminario : Inteligencia Artificial Y Patentes, May 9, 2019
https://eventos.udesa.edu.ar/evento/seminario-inteligencia-artificial-y-patentes-0

Estonia

Conference on Artificial Intelligence and Smart Economy, May 23, 2019
https://www.epa.ee/en/news/tomorrow-100th-anniversary-estonian-patent-office

Finland

IP Rights as Key Success Factors for AI Driven Businesses, February 5, 2019
https://ipruc.fi/koulutus-tapahtuma/ip-rights-as-a-key-success-factors-for-ai-driven-businesses/

Israel

International Conference on Emerging Technologies and Intellectual Property – Connecting the Bits, July 16, 2019

Singapore

IP/IT Issues in Artificial Intelligence, July 23, 2018
https://docs.wixstatic.com/ugd/55329f_a9a5de07b0a546818c345078331ae8a5.pdf

Russian Federation

International Conference "Digital Transformation: Focus on IP", April 23 and 24, 2019
https://rupto.ru/en/news/anons-international-conference-focus-on-ip-en

United Kingdom

AI: decoding IP – Exploring the Commercial, Economic and Legal Implications, June 18 and 19, 2019
https://orcula.com/ipo

United States of America

Artificial Intelligence: Intellectual Property Policy Considerations, January 31, 2019
https://www.uspto.gov/about-us/events/artificial-intelligence-intellectual-property-policy-considerations

European Patent Office

"Artificial intelligence" webpage
https://www.epo.org/news-issues/issues/ict/artificial-intelligence.html

Patenting Artificial Intelligence, May 30, 2018
https://www.epo.org/learning-events/events/conferences/2018/ai2018.html

[End of Annex and of document]