
SCP/38/4
ORIGINAL: ENGLISH
DATE: SEPTEMBER 28, 2026

Standing Committee on the Law of Patents

Thirty-Eighth Session
Geneva, November 2 to 6 2026

BACKGROUND DOCUMENT ON PATENTS AND EMERGING TECHNOLOGIES
(UPDATE OF SCP/30/5)

Document prepared by the Secretariat

Table of Contents

I.	INTRODUCTION	3
II.	INTRODUCTION TO AI TECHNOLOGY: NEURAL NETWORK, DEEP LEARNING and Generative AI	4
A.	Machine Learning	5
B.	Neural Networks	6
C.	Deep Learning	8
D.	Generative AI	11
E.	Current Limits of Deep Neural Networks, Generative AI and Agentic AI	20
F.	Current Technical Landscape of AI and Innovation	23
III.	PATENT PROTECTION OF AI-RELATED INVENTIONS	24
A.	General Considerations	24
B.	Patentable Subject Matter	28
C.	Novelty and Inventive Step	38
D.	AI-Generated Prior Art and PHOSITA in the Context of Assessing Patentability	40
E.	Sufficiency of Disclosure and Claims	46
F.	Industrial Applicability	52
G.	Inventorship and Ownership	52
IV.	AI TECHNOLOGY AS A TOOL IN THE PROSECUTION AND ADMINISTRATION OF PATENT SYSTEMS.....	56
A.	Tools for the IP Authorities	56
B.	Tools for Applicants, Third Parties and IP Professionals	57

ANNEX [pages 1-3]

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. Consequently, the Secretariat submitted document SCP/30/5 to the thirtieth session of the SCP for discussion.

2. 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 artificial intelligence (AI) as an issue to be discussed by the Committee. Therefore, while the term “emerging technologies” might have a broad meaning, covering various new technologies, including artificial intelligence (AI) and machine learning, blockchain, synthetic biology and gene editing, document SCP/30/5 covered background information on patents and AI.

3. At the thirty-fifth session of the SCP, the Committee agreed that the Secretariat would update document SCP/30/5 by compiling current laws and practices relating to the patentability of AI-related inventions, based on the information received from Member States and regional patent offices to be submitted to the thirty-sixth session of the SCP. Pursuant to that decision, Part III, Sections A to E of document SCP/30/5 [were updated and submitted at the thirty-sixth session \(SCP/36/5\)](#).

4. In addition, to maintain the integrity and readability [of the document](#), the Secretariat added, in Part II of document SCP/30/5, an overview of Generative AI. Furthermore, Parts III, Section F, and IV of document SCP/30/5 were updated on the basis of the activities of the SCP related to the topic of patents and [submitted for the Committee’s consideration at SCP/36 \(SCP/36/5\)](#).

5. This document maintains the structure of document SCP/30/5, [along with relevant updates made to the document \(SCP/36/5\)](#), which consists of three parts. The first part of the document provides background information about AI technology. It illustrates the basic technological concept of AI, particularly on machine learning technology, which is the core of the current AI development, and Generative AI. Such an introductory description of the technology is considered necessary since the implication of a particular technology to the patent system requires at least the basic understanding of the technology itself. [In doing so, the latest technological advancements in AI since the last update \(SCP/36/5\) have been briefly captured. They relate to recent changes in efficiency and deployment of deep learning systems, multi-modal generative AI, and the changes in foundation models and post-training, as well as agentic AI. A new Section F that crystallizes the AI and innovation landscape based on WIPO’s AI and Innovation Toolkit has been added to the current update in Part II of this document. 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 AI technology \(or AI-related inventions\) as the subject of patent protection, and the third part discusses use of AI technology as a tool for the authorities and users of the patent system. Changes reported by WIPO Member States and based on independent research conducted by the Secretariat have been recorded. A new Section D in Part III discusses emerging issues in AI-generated prior art and PHOSITA in the context of assessing patentability as these challenges are increasingly being discussed in the context of patent law and policy. Part III is accordingly renumbered.](#)

6. 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 as in the invention for which the patent is granted meeting the substantive criteria; 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 relating to the patent protection of AI-related

inventions touch upon the first aspect of patent quality, while the issues about the improvement of patent procedures using AI technology relate to the second aspect of patent quality.

7. In addition, the document contains an Annex, which lists webpages of the World Intellectual Property Organization (WIPO) and IP Offices, dedicated to AI and patents. For the benefit of the readers of this update document, paragraphs (or additions within a paragraph) and footnotes which are new or have been amended are shown in blue. To support accessibility for visually impaired users, these paragraphs are 3 to 5 and 7, 8, 9, 15, 17, 23, 24, 26, 34, 36 to 39, 54, 59, 65 to 77, 79, 82, 84 to 86, 89, 90, 93-96, 101, 103, 104, 108, 113, 115 to 118, 120 to 124, 126 to 128, 130, 131, 133 to 145, 150, 153, 156 to 173, 182 to 184, 188 to 200, 209 to 211, 213, 214, 216 to 219, 223 to 225, 227 to 231, and Annex pages 1-3.

II. INTRODUCTION TO AI TECHNOLOGY: NEURAL NETWORK, DEEP LEARNING AND GENERATIVE AI

8. While there is no single, universally accepted definition of artificial intelligence (AI), legal and policy instruments that do define AI systems generally adopt broad and technology-neutral formulations, reflecting the evolving and diverse nature of AI technologies. In this context, the EU AI Act and the OECD definition both describe an AI system as a machine-based system that, with varying levels of autonomy and adaptiveness, infers from inputs how to generate outputs such as predictions, content, recommendations, or decisions that influence physical or virtual environments. For the purposes of this technical introduction,¹ AI systems can therefore be viewed primarily as systems that learn or infer patterns from data.

9. Currently, only a few jurisdictions have offered definitions on AI and related concepts. For example, the Republic of Korea has definitions of “AI”, “AI Systems” and “AI Technology” in its 2025 Framework Act on the Development of Artificial Intelligence and the Creation of a Foundation for Trust. The term “artificial intelligence system” is defined as “an artificial intelligence-based system that infers outputs such as predictions, recommendations, and decisions that affect real and virtual environments for a given goal with various levels of autonomy and adaptability”, while AI itself is defined as “the electronic implementation of human intellectual abilities, such as learning, reasoning, perception, judgment, and language comprehension”.² While these definitions have primarily developed in the regulatory realm rather than patent law contexts, they illustrate a broad global convergence and understanding that covers the latest developments in AI technologies. However, it is pertinent to note that while regulatory definitions can be a good starting point to understand the nature of AI technologies, specific terminologies around AI inventions for the purposes of patent law require a deeper and nuanced classification (see Section 2(F) of Part II below).

10. The first part of the document introduces the most important technical concepts around Neural Network (NN), Deep Learning (DL), and Generative AI, which are today the paramount technologies in AI.³ It provides an understanding of how these emerging technologies work in

¹ See Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonized rules on artificial intelligence (Artificial Intelligence Act), Article 3(1); OECD, Explanatory Memorandum on the Updated OECD Definition of an AI System, OECD Artificial Intelligence Papers, No. 8, OECD Publishing, 2024, pp. 4-5. Both definitions use machine-based, technology-neutral terms and refer to autonomy/adaptiveness and outputs such as predictions, content, recommendations or decisions.

² See, Framework Act on the Development of Artificial Intelligence and the Creation of a Foundation for Trust (Act No. 20676, Jan. 21, 2025), available at: https://elaw.klri.re.kr/eng_service/lawView.do?hseq=71019&lang=ENG

³ 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.

an accessible manner for non-computer specialists in order to assist better understanding about the intersection between AI technology and patents.

A. Machine Learning


11. 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.

12. 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.

13. 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.

14. 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



$$\text{hand size} = 0.098 * \text{height} + 2.23$$

$$\begin{aligned} \text{if height} &= 180 \text{ cm,} \\ \text{hand size} &= 0.098 * 180 + 2.23 \\ \text{hand size} &= 19.9 \text{ cm} \end{aligned}$$

15. 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 her or his hand size is 19.9 cm (see the right box). **In this simplified example, the two learned numbers in the equation, the slope and the intercept, are the model's parameters. More realistic ML**

systems also learn parameters from data, although in far larger numbers and with more complex mathematical functions.


16. 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's 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.

17. Among those more advanced methods, Neural Networks (NN) offer a universal predictor, able to accept any kind of input. Each input is multiplied by a weight and sum-up for solving tasks involving unstructured data as input, such as image, speech, [text, code, video, audio, and biological or chemical data after numerical encoding](#). As an advanced type of NN, Deep Learning continues to be the core technique in AI patent applications [and in current generative AI systems](#). WIPO's 2024 GenAI patent landscape report notes that advances in deep learning and computing power have spurred GenAI, and that GenAI patent families increased by over 800% since the introduction of the transformer architecture in 2017.

B. Neural Networks

18. The fundamental building block of a NN is the 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 .

19. 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 follows:




20. 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.

21. 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 is called ReLU (Rectified Linear Unit).⁴

22. An artificial neuron is a relatively simple function. It can be programmed in less than 25 lines of code. 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.



23. 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. Thus, data flows forward through the network, from the input layer to the output layer. For common feed-forward networks (for example, multilayer perceptrons used for classification or prediction), inference is a one-way computation through the layers. Other architectures process sequence information differently.

24. 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. In current terminology, this adjustment is usually calculated by an algorithm known as backpropagation, combined with an optimization method such as gradient descent or one of its variants.


25. 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

⁴ 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”.

(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.⁵

C. Deep Learning

26. Although neural networks are known since the fifties, the usual number of hidden layers remained in practice only *one* until the years 2000s. Improved computational power, larger datasets and enhanced training methods have been made possible since the 2010s, led to increase (so "deepen") the number of layers 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 follows:



27. 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. In the example below, the image of a cat is transformed into a matrix of size 18×40 , with each pixel defined by its level of grey, expressed as an integer between 0 (white) and 256 (black). For a realistic cat and dog image classification, we use typically images of size 128×128 pixels, with each pixel defined by three values for red, green, blue levels, i.e., 49,152 input nodes, and consequently 49,152 input features for *each* following.

⁵ "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>.

how a human perceives an image how a computer perceives an image



```

5  4  6  5  5  5  6  7  9  8  9  8  4  6  11 40 136
6  8  5  5  5  5  6  7  8  8  9  4  6  5  27 18  61
5  9  7  4  4  4  4  5  5  6  6  7  7  7  123 86  9 26
4  82 187 6  5  4  4  4  4  5  6  7  5  85 98 111 6  10
4  183 132 77  3  4  4  4  5  6  10 26  91 98 127  5  7
8 140 139 93 18 11  5  3  5  5  6  6  66 92 133 149  4  9
18 149 182 78 73  6  1  4  5  5  11  8  91 120 122 168  9 16
15 155 149 85 94  5 10  5  6  8  17 30 86 115 143 121  9 28
22 147 149 142 80 101 27 18 23  7 148 85 119 137 157 129 13 46
15 144 146 126 113 126 135 68 156 97 182 129 121 158 95 152 18 61
18 150 149 99 135 144 217 212 176 119 198 174 112 83 94 142 36 81
16 144 147 85 181 135 179 223 196 187 201 151 131 87 161 123 41 185
43 110 114 115 68 93 147 187 218 197 229 165 182 136 99 98 74 117
27 114 133 83 75 117 179 168 168 184 234 161 181 142 149 86 95 112
11 84 185 121 84 153 188 183 215 171 228 158 169 161 124 56 84 189
3  66 118 112 116 135 136 233 235 218 127 196 174 97 118 164 144 118
2  48 84 185 113 148 149 238 183 146 74 195 187 159 173 284 168 182
6  21 136 129 114 153 98 192 144 172 61 124 128 138 189 188 176 181
9  11 188 140 121 122 92 244 133 287 187 78 182 117 192 147 172 112
6  9 179 168 121 53 48 141 116 154 119 141 238 135 216 189 188 138
8  9 165 185 188 125 242 184 133 117 138 193 284 126 223 191 193 175
8  33 183 176 139 212 171 221 123 132 63 280 192 65 195 195 285 283
7  12 181 233 289 127 218 188 51 164 48 197 192 184 280 217 211 216
6  14 191 229 158 42 195 286 82 132 189 74 187 212 161 223 211 218
8  22 182 191 185 139 181 65 178 185 136 78 187 199 174 178 287 218
10 29 168 163 178 193 211 62 211 75 188 118 171 280 217 173 213 195
18 21 184 137 158 158 153 193 95 156 161 134 280 184 171 286 281
21 35 98 195 234 163 86 111 131 112 156 72 68 139 172 168 286 281
31 46 181 161 159 155 186 151 53 197 195 93 67 114 198 167 197 284
32 29 87 188 124 87 38 192 49 225 112 57 68 133 148 154 195 142
48 26 66 184 194 67 36 97 48 152 127 186 119 283 178 72 185 181
55 39 72 145 166 198 183 114 91 174 147 114 176 125 179 189 194 228
97 72 89 173 163 185 163 91 182 114 94 98 164 156 163 114 184 224
135 111 111 187 168 143 169 94 63 76 84 115 144 155 169 89 128 215
163 151 115 164 184 174 126 135 74 86 98 132 158 143 123 158 187 216
195 173 88 131 193 183 112 189 186 184 128 128 114 131 158 175 191 213
198 177 178 185 158 127 112 154 181 184 112 123 167 166 183 197 187 287
139 157 161 116 126 122 124 89 92 89 118 97 142 169 189 218 174 194
121 112 148 122 141 118 92 98 85 88 94 91 156 214 211 152 139 185
97 184 181 178 153 114 186 95 88 96 84 129 217 221 157 156 178 285

```

(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 million of weights and bias parameters to be set during the training, requiring ten thousand of labeled images.

28. Surprisingly, with an existing Open Source (OS) framework, such as Keras,⁶ 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).

29. 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.

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

Discovery of features representations

31. 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 some features based on their 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.


32. For the first time in machine learning, Deep Neural Networks show a practical ability to automatically discover 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 automatically identify the ones to exploit for the

⁶ Keras: The Python Deep Learning library, François Chollet and others, 2015-2024. <https://keras.io/>.

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

33. 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.



34. Mathematically, artificial neural network models can be understood as just a set of matrix operations and finding derivatives.⁷ With the increase of computational power, deep learning can surpass any other ML approaches, as long as a massive amount of training data and computing power is available.

35. How far can we scale such neural networks and continue to improve performance and capacities of Deep Learning models? The experiments that took place to explore and improve this scaling capacity led to the current Generative AI outstanding success.

Efficiency and deployment of deep learning models

36. In recent years, a second complementary line of progress has focused on making deep learning models smaller, cheaper, and easier to deploy. Techniques such as pruning, quantization, distillation, sparsely activated models and parameter-efficient fine-tuning can reduce the amount of computation or memory required for a model. Pruning, for example, removes less important weights, neurons or other structures from a neural network so that the resulting model requires less memory or computation at inference time.⁸

37. Deep learning systems are therefore not always deployed as one large model running in a central computing center. Depending on the task, a system may use a large cloud model, a smaller specialized model, an on-device model, or a combination of models and tools. These

⁷ 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.

⁸ See [Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding](#) (2015); Tailin Liang et al., [“Pruning and Quantization for Deep Neural Network Acceleration: A Survey”](#) (2021).

deployment choices affect use cases, cost, speed, accuracy, confidentiality, energy use, and the practical accessibility of AI systems.

D. Generative AI

38. Generative AI can be defined as a technology that can create new, meaningful and human-like content, including text, images, audio, or video, when prompted by a user. "Prompts" here correspond to textual instructions in natural language, usually produced by the human users, although in recent systems prompts may also include images, audio, video, code, retrieved documents, or outputs produced by other software tools or AI components. By combining text with images, audio, video, or other inputs, the prompt can give the model a broader context and more closely approximate human perception and data processing, since people understand the world by integrating signals from multiple sources.⁹

39. The next section describes the main concepts of Generative AI:

- the distinction between discriminative and generative tasks and models;
- the transformer architecture and the attention mechanism, which are the foundation behind the modern rise of Generative AI; and
- the main characteristics of Large Language Models (LLMs) and foundation models;
- foundation models, post-training and the role of instruction and preference tuning;
- retrieval, tool use and agentic workflows; and
- multimodal generative AI systems capable of processing and generating different types of data.


Discriminative versus generative tasks

40. Deep neural networks can usually be adapted to two different kinds of tasks:

- *Discriminative* tasks are tasks involving a decision on the input data, such as classification, identifying names in texts or segmenting an image. *Discriminative* models are models adapted and trained to separate input data into these different classes.
- *Generative* tasks are tasks involving the creation of new data given some input data. *Generative* models are models adapted and trained to create such new data. They are typically used to translate text, generate images, summarize text or answer questions.


41. The two following diagrams illustrate these fundamental types of machine learning tasks.

⁹ See, e.g., Tadas Baltrušaitis, Chaitanya Ahuja and Louis-Philippe Morency, "Multimodal Machine Learning: A Survey and Taxonomy" (2017), arXiv:1705.09406; Yuan Yuan, Zhaojian Li and Bin Zhao, "A Survey of Multimodal Learning: Methods, Applications, and Future", ACM Computing Surveys, Vol. 57, No. 7, Article 167 (2025).



Source: Painting visuals from Wikimedia Commons under public domain, originally from WIPO Patent Landscape Report on *Generative Artificial Intelligence*, 2024. <https://doi.org/10.34667/tind.49740>

42. Above: A discriminative task for classifying whether an image is the Mona Lisa painting or not. The model must learn how to distinguish two classes: the painting *Mona Lisa* and *other* paintings. For this sort of models, the learning process focuses on the criteria to distinguish the classes. So, expressed as a space of painting characteristics, the model concentrates on representing the boundary between the two classes of paintings.



Source: Painting visuals from Wikimedia Commons under public domain, originally from WIPO Patent Landscape Report on *Generative Artificial Intelligence*, 2024. <https://doi.org/10.34667/tind.49740>

43. Above: **A generative task for producing new painting samples.** For the generative task, the model must learn the global aspect of every painting to be able to generate coherent new paintings. For this sort of models, the learning focuses on representing the global distribution of the characteristics of the painting. The two generated images were produced using the original Stable Diffusion¹⁰ model.

44. Discriminative models excel in classification but cannot generate new data. In contrast, generative models can also address discriminative tasks, but usually with lower accuracy than discriminative models. Generative models have more parameters, are computationally more expensive and require more training data than discriminative models. For this reason, discriminative tasks have received initially more interest. As computing power increased, and more training data became available, generative tasks became more feasible and attracted more development effort in the last years.


¹⁰ Rombach R., Andreas Blattmann A., Dominik Lorenz D., Patrick Esser P., Björn Ommer B. (2021). "High-Resolution Image Synthesis with Latent Diffusion Models." arXiv:2112.10752, <https://arxiv.org/abs/2112.10752>.

The Transformer architecture

45. Transformer is a type of Deep Neural Network designed specifically for natural language processing (NLP) tasks. It was introduced first in 2017¹¹ and has been developed to better scale during training than other traditional Deep Learning methods. Transformers are based on two important concepts: encoder-decoder and self-attention.

46. An **encoder-decoder** model is based on three parts: encoder, code and decoder. The encoder is a neural network that learns how to encode and compress input data into an intermediary representation, the code, which is basically a sequence of numbers. The code is then used by the decoder, another neural network, which has learnt how to decompress and reconstruct data into the expected input format. Beyond data compression, the objective of the encoder-decoder is to learn how to represent the nature of some data, so that modifications of this internal representation can still be re-constructed into a new meaningful output.

47. This approach was in particular developed to address machine translation, where an input in a source language is encoded into a compressed numerical representation, which is then decoded into a target language, based on a large volume of translation input/output examples.



48. Above: Encoder-decoder Neural Network architecture. An input, such as text or image, is encoded into vectors through several layers of neural networks. The *code* is an intermediary compressed numerical representation, which can be then used by the decoder to generate an output in a desired target language or image representation. Note that the number of layers in modern models can be much higher than represented here.

49. For a system dedicated to discriminative tasks, the Decoder part is usually removed, because the output is normally just a classification, not newly created data. This is the case for the BERT (Bidirectional Encoder Representations from Transformers) model family, which was the center of interest of the Machine Learning field between 2018 and 2022. For a system dedicated to generative tasks, the Encoder might be removed, because the output should be newly generated data, and the encoder might not be helpful. This is the case for the OpenAI GPT (Generative Pre-trained transformers) models.

50. **Attention** is a neural network mechanism that makes it possible for the learning to focus on different parts of the text simultaneously. It allows to capture long-range dependencies in

¹¹ Vaswani A., Shazeer N., Parmar N., Uszkoreit J., Jones L., Gomez A., Kaiser L., Polosukhin I. (2017). "Attention Is All You Need." Advances in neural information processing systems 30.

text, which is important for understanding and phrasing complex language. In practice, *attention* assigns weights to input tokens based on importance so that the model gives more emphasis to relevant tokens. However, how are these important tokens selected? Depending on this selection criteria, different types of attention mechanisms exist.

51. The original transformer architecture uses *self-attention*: following this approach, the correlations between every possible pair of words in all the inputs of the training data is evaluated. To give a quick insight into this process: if two tokens appear together more frequently in the training data than by random chance, they will be assigned stronger weights than non-correlated tokens.

52. In conclusion, the **Transformer** is an encoder-decoder deep neural network that takes advantage of attention to maintain faster training time than previous deep neural network as its number of layers increases. The capacity to scale is directly related to this number of layers. Stacking more layers means multiplying the global number of neurons and thus the capacity to store more learned parameters. In Transformers, each layer uses attention. As a result, the learning process does not loop over the input multiple times, which is expensive in term of computation. Instead, the Transformer passes the input through multiple attention layers.

53. As visible in the table below, the number of layers of Transformers has continuously increased over time, allowing to train on larger datasets and to learn more parameter information.

Model Name	Developer	Year	Number of layers	Number of parameters	Training data size (estimates)
Transformer (original)	Google	2017	6	110 millions	800 million words
BERT	Google	2018	12	110 millions	3.3 billion words
GPT	OpenAI	2018	12	110 millions	600 billion words (40GB)
GPT-2	OpenAI	2019	48	1.5 billions	8 millions web pages
GPT-3	OpenAI	2020	96	175 billion	570 GB of text data
GPT-4	OpenAI	2023	120	1.8 trillions(est.)	13 trillion tokens

54. The table is useful to illustrate the historical increase in the scale of transformer models. It should not, however, be read as a current list of leading models. For many recent models, the number of parameters, layers, or training tokens is not publicly disclosed. Moreover, current capabilities are also affected by other factors, such as context-window length, supported modalities and tool access.


55. We mentioned previously that the generative capacity of a model increases as the size of the model and the training increase. We understand here why the transformer architecture made possible the rise of Generative AI, while discriminative tasks were before the most successful AI systems.

From Large Language Models to ChatGPT

56. As illustrated by the previous table, LLMs are very large models trained with a massive amount of training data. Their development requires large and modern computing infrastructure. Even with a large computing center of hundred thousand processors, the training of such LLMs take typically between two weeks and three months. These models are the basis of modern assistants such as the well-known OpenAI ChatGPT, Microsoft Copilot or Google Gemini. They share a particular generation mechanism called *autoregression*.

57. Instead of generating a complete text or image in once, **autoregressive models** generate data one element at a time, based on previously generated elements. For example, with text, the generation becomes the result of an iterative process where the next word is predicted considering the previously generated words, until the end of a text is predicted. This can be viewed as an auto-completion function, which mimics human language. This is also similar to the well-known next word prediction feature of smartphone keyboards. However, models considered here are considerably larger than next word prediction of a phone application, typically by a factor 1000 or more.

58. One key advantage of this approach is that the prediction of a next token or word can be learned just with ordinary text. As the text is scanned during training to set the internal parameters of the model, these models have the capacity, given a text context, to predict the most likely next word. Such training does not require manually labeled data. The computer can learn to identify patterns without a human to provide guidance along the way, a type of learning called *unsupervised* learning. However, a massive amount of text is required to reach the current performance of these models.



59. Above, an example of text generated via autoregression, illustrating the attention mechanism. As the model generates the text word by word (in bold), it can focus on previous words that are relevant to each newly generated word via the attention mechanism (blue arrows). In practice, the window size of words that can be reached is very large, *over a million in leading models*.

60. Two interesting properties emerge in such models:

- they not only learn the general language, but also how to generate text about a variety of facts concerning entities and events in the world, when seen in the training data,
- they are able to memorize so much text content that the completion becomes surprisingly precise and relevant if enough initial context is provided.

61. These two properties make possible versatile assistants like ChatGPT, which can go beyond the generation of grammatically correct text. These assistants can further answer questions and realize various tasks such as summarization, classification, rephrasing, translating, creation of stories, etc. However, given that these models are designed to predict one next word at the time for completing an input text, how can they achieve such sophisticated tasks?

62. This capacity to assist and dialogue with users is provided by tuning these models with thousands of examples of **instructions and expected results**. The instructions correspond to the “prompts” that would be entered by users to express their request. The model will learn to

generate text to match the expected results, as a completion of these instructions. While the base LLMs captures the competence of generating a huge variety of human language texts, this additional tuning teaches the model how to use this competence for performing tasks following instructions, enabling fluent and versatile dialogues.

63. The previous traditional "supervised" machine learning models require a large amount of task-specific annotated training data. In contrast, these LLMs can generate new content just by writing natural language "prompts." Therefore, using these Generative AI tools does not require technical skills anymore. Modern cutting-edge AI becomes directly accessible to the general public, who just need to prompt the models with natural language instructions.

64. Considerable research work has been conducted for addressing various tasks and assistance in multiple business areas, as well as generating other form of language, such as programming code or structures describing proteins. A reference is made to the recent WIPO Patent Landscape Report on Generative AI¹² for a comprehensive description of applications and examples.

Foundation models and post-training

65. The term "foundation model" is commonly used for a model trained on broad data, generally using self-supervision at scale, that can be adapted to a wide range of downstream tasks. Not every generative AI system is built from a foundation model, but many current systems use a broadly trained base model and then adapt it to particular tasks or domains.¹³

66. It is important to note that distinctions exist between proprietary or closed models where weights and training details are not publicly available and open models where model weights are made available to the public for inspection, fine tuning or for the purposes of independent deployment. For non-specialist purposes, the development of many current generative AI systems can be described in three stages. First, a base model is pre-trained on broad data. Second, the model is post-trained to make its behavior more useful, controllable and aligned with deployment objectives, using techniques such as supervised instruction tuning, reinforcement learning from human feedback (RLHF), reinforcement learning from AI or automated feedback (RLAIF/RLAF) and direct preference optimization (DPO), through which developers use demonstrations, rankings, rewards or other feedback signals to steer model behavior. This means the deployed product is not simply the base model with neutral outputs from pre-training data; post-training reflects deliberate design choices and is often used to improve accuracy, safety, transparency and accountability and to mitigate bias or other ethical risk.

67. Third, the deployed system may be connected to external data, tools or human review processes. These stages are not always strictly separated, and the terminology differs across developers, but the distinction helps explain why two systems built on similar base models may behave differently in practice.¹⁴

¹² World Intellectual Property Organization (WIPO) (2024). *Generative Artificial Intelligence*. Patent Landscape Report. Geneva: WIPO. <https://doi.org/10.34667/tind.49740>.

¹³ R. Bommasani et al., "On the Opportunities and Risks of Foundation Models" (Stanford Center for Research on Foundation Models, 2021); NIST AI 600-1 (July 2024),

¹⁴ See L. Ouyang et al., "Training Language Models to Follow Instructions with Human Feedback" (NeurIPS, 2022) (supervised fine-tuning and RLHF); R. Rafailov et al., "Direct Preference Optimization: Your Language Model is Secretly a Reward Model" (2023) (DPO); Y. Bai et al., "Constitutional AI: Harmlessness from AI Feedback" (2022) and H. Lee et al., "RLAIF: Scaling Reinforcement Learning from Human Feedback with AI Feedback" (ICLR 2024 submission) (RLAIF). On post-training recipes and transparency of training data/recipes, see N. Lambert et al., "Tulu 3: Pushing Frontiers in Open Language Model Post-Training" (2025).

68. Recent technical discussion also refers to iterative refinement or model-assisted refinement. In such approaches, generated examples, critiques, evaluation or other feedback signals may be used to improve later versions of a model or to improve outputs within a workflow. Such refinement does not necessarily mean that the model autonomously changes its own architecture; in most practical settings, developers or a controlled system use feedback, additional examples or evaluation results to update model behavior or select better outputs. A related technique is knowledge distillation, in which outputs or other signals from a larger or more capable “teacher” model are used to train a smaller “student” model. Distillation can help transfer aspects of the teacher model’s behavior into a more compact model, reducing deployment costs, latency, memory requirements and other resource demands, although the student model may also inherit limitations, errors or biases from the teacher model.¹⁵

Retrieval, tool use and agentic workflows

69. A common way to adapt a generative AI system without retraining the model is Retrieval Augmented Generation (RAG). In RAG, the system retrieves information from an external source, such as a search index, database or collection of patent and scientific documents, and provides that information to the model as part of the context for generation. The retrieved information can make outputs more specific to a domain and easier to check, but the quality of the result still depends on the quality of the retrieval, the instructions and the model response.¹⁶

70. Fine-tuning is also undertaken as a process to further train the model for a specific task where the weights assigned are based on prior knowledge. RAG and fine-tuning are complementary adaptation strategies since RAG is better equipped to handle tasks requiring regular update to domain specific factual content, while fine-tuning is better suited for adapting format or specialized reasoning patterns.

71. As the need to address more complex tasks than simple text and document generation appeared, LLM reasoning was developed as a new type of specialized model. These models have the ability to make logical connections to solve complex, multi-step problems, rather than just predicting the next word. Instead of generating directly an answer, these models generate intermediary prompts corresponding to a sequence of workflow steps. They can also self-correct and evaluate their own logic to revisit intermediary prompts before outputting a result.

72. Many recent systems also connect the model to software tools, such as search engines, calculators, code execution environments, databases, design tools or application programming interfaces. AI models become “agents,” because they can perform actions on external services rather than only generating a static answer. The term “agentic AI” is used for systems in which a reasoning model decomposes an objective into steps, selects or calls tools, observes the results and continues the task through one or more iterations. An orchestration layer or a scaffolding software infrastructure surrounds the model and manages the tools, stores intermediate results and routes the information between different components. In multi-agent workflows, several model-based components may coordinate or delegate subtasks. These systems go beyond a standalone chat interface, but their behavior remains shaped by the models, the tools made available to it, the instructions, the safeguards and any human approvals built into the workflow. Unlike traditional generative AI where human prompting is necessary at every stage, agentic AI is capable of autonomously acting based on goals and objectives to be achieved, planning,

¹⁵ See Chuanpeng Yang et al., “Survey on Knowledge Distillation for Large Language Models: Methods, Evaluation, and Application,” *ACM Transactions on Intelligent Systems and Technology*, 2025, DOI: 10.1145/3699518; Jianping Gou, Baosheng Yu, Stephen J. Maybank and Dacheng Tao, “Knowledge Distillation: A Survey,” *International Journal of Computer Vision*, 129(6), 1789–1819, 2021, DOI: 10.1007/s11263-021-01453-z. For the foundational formulation, see Geoffrey Hinton, Oriol Vinyals and Jeff Dean, “Distilling the Knowledge in a Neural Network,” 2015.

¹⁶ P. Lewis et al., “Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks,” *Advances in Neural Information Processing Systems* 33 (2020).

reasoning, memory to achieve complex and multi-step goals without a continuous need for human oversight.

LLM customized for generating invention

73. One of the most relevant applications for the SCP is **AI-generated inventions**. To enable this capacity, LLMs are typically fine-tuned with different or additional instructions and expected results relative to invention. They can also be further pre-trained with patent documents, scientific texts or other technical content, [and combined with RAG to retrieve such literature](#).

74. Such customized models can be used in different ways:

- (i) A model is used to help the drafting of a patent application, based on the outline of a human-based invention, including suggestion of the relevant prior art.
- (ii) A human identifies inventions and uses the model to identify other possible problems that may be solved by these inventions.
- (iii) A human identifies a problem and uses the model to generate possible solutions.
- (iv) A human uses the model to identify potentially valuable problems and to generate possible solutions to these problems.
- (v) [A human uses a tool-connected workflow to retrieve prior art, generate candidate solutions, test or simulate alternatives, and rank the results for further human and AI-based review.](#)
- (vi) [A human uses a multimodal system to combine text, diagrams, experimental data, or design files when exploring technical alternatives.](#)

75. [These different usages of a customized LLM](#) can be viewed as gradually higher levels of inventorship, from assistance and insights to human inventor, to more significant generations of novel ideas. However, other factors determine the nature and extent of human effort involved in the generation of inventive LLM outputs:

- (i) the creation of the prompts for querying the model,
- (ii) the technical principles and features provided for guiding of the model,
- (iii) the choice of model parameters,
- (iv) the number of iterations involved in prompting the LLM.
- (v) [the external documents retrieved by the system and the tools connected to the model,](#)
- (vi) [the human decisions made when selecting, rejecting, testing or modifying generated alternatives, and](#)
- (vii) [the records of prompts, retrieved sources, model versions, tool calls, intermediate outputs and final outputs.](#)


76. [From a technical perspective, such records can be important for reproducibility and auditability, although the completeness and reliability of such logs will depend on the system design and the practices of its users.](#)¹⁷

Generative AI for images and other modalities

77. We have focused in this document on text generation. However, mature models for image generation appeared before LLM, using different type of models. Novel and meaningful images are constructed based on models trained similarly on hundred thousand images, controlled by additional machine learning mechanisms. Textual descriptions of the training

¹⁷ NIST, [Secure Software Development Practices for Generative AI and Dual-Use Foundation Models: An SSDF Community Profile](#), NIST SP 800-218A (July 2024).

images make possible to combine text and image modalities. For example, diffusion models can produce high-resolution images from short textual description, as illustrated by the famous Stable Diffusion model, released in 2022. [More recently, multimodal generative AI has become a standard direction of development, with systems capable of processing or generating text, image, audio, video, code or other data types within a single workflow.](#)



Source: Stable Diffusion model with the prompt “a cat reading a patent,” originally from WIPO Patent Landscape Report on *Generative Artificial Intelligence*, 2024. <https://doi.org/10.34667/tind.49740>

78. Above, the image generation is using the original Stable Diffusion model with the prompt “a cat reading a patent.” The generation process corresponds to: (i) supplying an image with random noise (random pixels) and (ii) iteratively use the neural network to remove noise. The supplied textual prompt is used to drive the generation. As the noise is progressively removed, a novel and meaningful image is constructed, using figurative elements statistically related to the prompt, as learned with a large amount of training data. Other modalities successfully covered by such models include video (an image being successively translated to another image), speech, music, and 3D scenes built from one or several 2D pictures.

79. [In multimodal systems, the prompt is not limited to text: an image, speech recording, video clip, diagram or other input may also guide the generation. In the most integrated systems, sometimes described as "any-to-any" systems, the same workflow may accept more than one type of input and produce more than one type of output, for example responding to spoken instructions about an image by producing text, code, a diagram or a further audio response.](#)

E. Current Limits of Deep Neural Networks, Generative AI [and Agentic AI](#)

Deep neural networks are black box

80. 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.

81. 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 billion 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 and lack transparency. The training itself is realized by the NN *on its own* and the resulting network is enormously complex.

82. [Transparency can be further complicated when a model is combined with retrieval databases, external tools, or other AI components, because the observed result may depend not only on internal parameters but also on the prompt, retrieved materials, tool outputs, model](#)

version and deployment settings. For that reason, logging and documentation of the system workflow have become an important practical complement to model explainability.

Deep learning requires a lot of data

83. 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 and generative AI 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*.

84. 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 data (e.g., processing rare human languages, drug discovery for rare diseases, etc.) or domains with legal restrictions. [Modern generative AI may reduce some barriers by reusing foundation models or by supplementing scarce datasets with retrieval or synthetic data, but these techniques do not eliminate the need for high-quality, lawful and contextually representative data.](#)

Real world data is biased

85. 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.¹⁸ 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 can also arise from the over-representation of high-resource languages, regions or online communities in internet-scale data, and may be less visible in aggregated evaluations.](#)
- *Bias amplification*: By nature, the machine learning training methods tend to identify discriminant patterns in data for quickly increasing the quality of prediction or generation. Consequently, they not only learn our actual bias, but oftentimes, they also amplify our bias. [In generative systems, bias amplification may appear as stereotyped images or text, exclusion of minority viewpoints, or homogenization of outputs, especially when synthetic data is reused for further training.](#)
- *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.

Trustworthiness

86. It is not uncommon for LLMs to generate inaccurate content expressed in a convincing manner. Such untruthful outputs are referred to as *hallucinations*. Although a lot of work has been dedicated recently to reduce the frequency of such factual errors, they remain significant, especially in specialized domains. [Similar failures may occur in multimodal systems, for example when a system misinterprets an image, chart, sound or video and expresses the error](#)

¹⁸ 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>).

confidently. The problem is not limited to false facts; it can include flawed reasoning, inaccurate citations or unjustified confidence.

87. A recent study estimated that the frequency of hallucinations per output in GPT-4.0 is around 28.6% on simple task¹⁹, which is likely higher than usually perceived by users. When applied to the legal domain for example, a study from Stanford researchers²⁰ measured that the frequency of hallucinations per response of GPT-4.0 is at least 58%. Even with customized LLM such as the ones commercialized by LexisNexis and Thomson Reuters, specifically pre-trained with legal texts and using legal databases to enrich and control the prompts (a technique known as Retrieval Augmented Generation or RAG), the rate of hallucinations per generated response was still estimated between 17% and 33%, according to a study from the same researchers²¹.

88. The evaluation of LLM over time appears difficult. In classical Machine Learning evaluation, systems should not be tested and evaluated on the data they were trained on. This principle ensures that different systems can be benchmarked and compared over time in a fair manner using open-access datasets. LLMs are trained on large portions of the Internet, which means LLMs might be trained with these existing evaluation datasets commonly published online, possibly multiple times. This phenomenon is called *data contamination*. The evaluation of a LLM model is then not anymore measuring its capacity to answer correctly unseen questions, but to simply memorize published solutions.

89. Some researchers are also describing a future gradual loss of performance related to the quality of available training data, called *Model Collapse*. As the output of Generative AI systems is flooding the internet, future models might themselves be trained on online content that earlier models have created, including all their biases and errors, resulting in a progressive decrease of accuracy and capacity. However, the risk is not that synthetic data is always harmful; rather, indiscriminate or poorly curated use of model-generated content can reduce data diversity and amplify errors. Synthetic data can be useful when it is generated, filtered and validated with attention to quality and representativeness.

90. As generative AI systems are connected to external tools, databases and application programming interfaces, trustworthiness also includes security. A user or third party may attempt to manipulate a system through prompt injection, including indirect prompt injection hidden in documents or web pages retrieved by the system. Other risks include leakage of sensitive information, insecure handling of generated outputs, data or model poisoning, and excessive agency, where a system is given more autonomy or tool permissions than are necessary for the task. These risks are heightened in agentic AI because hallucinations, manipulated instructions or indirect prompt-injection attacks can produce consequential actions, such as sending communications, deleting or modifying files, initiating purchases or other transactions, changing records or calling external services, that may be difficult or impossible to reverse. Accordingly, deployers should limit tool functionality, permissions and autonomy and

¹⁹ Chelli M, Descamps J, Lavoué V, Trojani C, Azar M, Deckert M, Raynier JL, Clowez G, Boileau P, Ruetsch-Chelli C. Hallucination Rates and Reference Accuracy of ChatGPT and Bard for Systematic Reviews: Comparative Analysis. *J Med Internet Res*. 2024 May 22;26:e53164. doi: 10.2196/53164.

²⁰ Matthew Dahl, Varun Magesh, Mirac Suzgun, Daniel E Ho, Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models, *Journal of Legal Analysis*, Volume 16, Issue 1, 2024, Pages 64–93, <https://doi.org/10.1093/jla/laa003>.

²¹ Varun Magesh, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D. Manning, & Daniel E. Ho, Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools, *22 Journal of Empirical Legal Studies* 216 (2025).

require human confirmation for high-impact actions, with logging, monitoring and clear fallback or recovery procedures.²²

Offensive and sensitive content

91. The Generative AI models are trained with a massive amount of data, text and images. It is not possible to filter out in a perfect manner all inappropriate, morally and ethically challenging content from the pre-training data. The consequence is that these models can reproduce and recompose similar offensive content in the generated output.

92. Various techniques try to prevent this. For example, the predicted words or phrases are continuously filtered through what are known as *guardrails* to remove offensive content (guardrails can be blacklisted words for examples). However, there is today no guarantee to fully neutralize such generated output, especially when toxic prompts are used.

AI resources and infrastructure

93. LLM pre-training is expensive, and the exploitation of LLM at scale requires similarly large computing centers, [dedicated hardware such as Graphical Processing Unit \(GPU\) whose production depends on very few manufacturers, and growing electricity power capacity](#). Generative AI relies also on data at internet scale. Combined with the investment capacity and the availability of local digital skills, concerns about a gap between high income nations and low- and medium-income countries have been addressed in various fora.²³ At the same time, AI technologies have also been deployed to improve, for example, education, healthcare and financial inclusion in developing countries.²⁴

94. [Some recent technical developments may partly reduce these barriers. Smaller specialized models, model compression, distillation, quantization, caching and on-device deployment can reduce cost, latency and energy use for particular applications. Nevertheless, access to high-quality data, skilled personnel, compute infrastructure and reliable electricity remains uneven, and the increasing use of tool-connected and multimodal systems can add new operational and security costs.](#)

F. Current Technical Landscape of AI and Innovation

95. [WIPO's AI policy toolkit distinguishes between four types of AI inventions depending on the role that AI plays in the innovation process. This taxonomy is useful because AI may be the subject of innovation itself, a tool used by human inventors, a component incorporated into an](#)

²² See NIST, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1 (2024), sections 2.2 and 2.9 (confabulation and prompt injection); OWASP, "LLM06:2025 Excessive Agency," OWASP GenAI Security Project, identifying excessive functionality, excessive permissions and excessive autonomy as root causes of damaging LLM-agent actions; Partnership on AI, Prioritizing Real-Time Failure Detection in AI Agents (2025), noting that agent failures can result in unauthorized actions, irreversible changes or real-world harm and may arise during planning, tool use or execution; Kai Greshake et al., "Not What You've Signed Up For: Compromising Real-World LLM-Integrated Applications with Indirect Prompt Injection," AISec '23 (2023).

²³ See, for example, United Nations and International Labour Organization report: "Mind the AI Divide: Shaping a Global Perspective on the Future of Work", July 26th, 2024, ISBN: 9789211066524, <https://www.ilo.org/media/581631/download>. A research paper notes that the income level, share of youth population, digital infrastructure, specialization in high-skill tradable services, English proficiency, and human capital are strongly correlated with higher uptake of generative AI. See "Who on Earth is Using Generative AI?," Policy Research Working Paper 10870, World Bank Group at: <https://documents1.worldbank.org/curated/en/099720008192430535/pdf/IDU15f321eb5148701472d1a88813ab677be07b0.pdf>.

²⁴ See for example, "Tipping the scales: AI's dual impact on developing nations", World Bank Blogs at: <https://blogs.worldbank.org/en/digital-development/tipping-the-scales--ai-s-dual-impact-on-developing-nations>.

invention, or, in a future scenario, a system that generates an invention with little or no material human input.²⁵

- The first category is an **AI model**, meaning a new AI model or algorithm. This may include, for example, a new model architecture, optimization method, training technique, or an algorithmic improvement that makes an AI system more accurate, efficient or adaptable.
- The second category is an **AI-assisted invention**, where humans use AI as a tool in the inventive process. In this category, AI may help screen options, identify patterns, select candidates, optimize a design or test possible solutions, but the invention remains human-directed. An example is the use of AI in drug discovery to identify or select drug candidates.
- The third category is an **AI-based invention**, where the invention incorporates AI, and the AI forms part of the basis of the product, process, or service. In these cases, AI is not merely a research tool used in the background as an aid in the inventive process; it is part of how the invention operates. Examples may include a translation device using AI deep learning, AI-based image enhancement, diagnostic tools or other systems whose functionality depends on AI.
- The fourth category is an **AI-generated invention**, which WIPO describes as a future scenario in which an invention is generated autonomously by AI without material human input. This category is useful as a conceptual boundary for policy discussion because it helps distinguish current uses of AI from possible future cases where AI may play a more independent role in the inventive process.

III. PATENT PROTECTION OF AI-RELATED INVENTIONS

96. This Part of the document looks at patent protection of AI-related inventions. As discussed in Part II.F, the term “AI-related inventions” is used in this document as an umbrella term covering AI models, AI-assisted inventions, AI-based inventions and the possible future category of AI-generated inventions. The implications of AI technology for patent law may differ depending on whether AI is the subject of the claimed invention, part of how the invention operates, or a tool used in the inventive process.

A. General Considerations

97. 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 benefit 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.

98. 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

²⁵ World Intellectual Property Organization (WIPO), *Getting the Innovation Ecosystem Ready for AI: An IP Policy Toolkit* (Geneva: WIPO, 2024).

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.

99. 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 design (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.

100. 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.

101. For AI technology, this document follows the taxonomy introduced in Part II.F rather than setting out a separate classification. In broad terms, the computer-technology categories above correspond to AI models, AI-based inventions and AI-assisted inventions, while Part II.F also identifies the possible future category of AI-generated inventions.

102. 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. However, as the AI technology develops, the nature of human intervention in the creation process might change relative to increased autonomous performance of an AI system.

103. Within the taxonomy introduced in Part II.F, AI-related inventions, particularly AI-assisted inventions and the possible future category of AI-generated inventions, may be further understood by focusing on the degree and nature of human contribution to the creation of the core inventive concept. From that perspective, the following scenarios may be considered:

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

104. In the second scenario, the relevance of the AI technology in the invention creation process may be from minimum to determinative. The fourth scenario, i.e., Artificial General

Intelligence (AGI) or superintelligence,²⁶ is not something that the current technology permits.²⁷ Due to definitional ambiguity, while there is no consensus on what must constitute a fully developed AGI, there is a view that Artificial Narrow Intelligence (ANI) has arrived (e.g., natural language processing, image recognition), and recent industry reports warn that reaching the level of AGI can be predicted on the short term. 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.

105. 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 techniques²⁸ for numerous AI functional applications²⁹ in a variety of AI application fields³⁰. Open source (or open innovation) approaches are also popular among AI developers.³¹ For the detailed patent landscape data on AI-related inventions, reference is made to the said WIPO publication. In conjunction with its rising public prominence and the growth in attention in the scientific literature, recent years have also shown an exponential rise in patents relating to generative artificial intelligence.³² In terms of the number of patents filed, the most important generative artificial intelligence models are (i) generative adversarial networks (GANs); (ii) variational autoencoders (VAEs); and (iii) decoder-based large language models (LLMs). These patents are not concentrated in a specific field or industry and have been filed in a vast array of key applications areas.³³

106. How the AI technology affects patent laws has not been determined yet. Many states have not established a special procedure for examining AI-related applications.³⁴ 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;
- (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

²⁶ 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.

²⁷ WIPO Technology Trends 2019 – Artificial Intelligence, p.19.

²⁸ For example, machine learning, fuzzy logic and logic programming.

²⁹ For example, computer vision, natural language processing and speech processing.

³⁰ For example, transportation, telecommunication and life and medical science.

³¹ WIPO Technology Trends 2019 – Artificial Intelligence, p.109.

³² World Intellectual Property Organization (WIPO) (2024). *Generative Artificial Intelligence*. Patent Landscape Report. Geneva: WIPO. <https://doi.org/10.34667/tind.49740>

³³ For example, software, life sciences, publishing, business solutions, industry, transportation, security, and telecommunications. As a concrete example in Tanzania, the Patent Act, Cap. 217 does not expressly address emerging technologies such as artificial intelligence or biotechnology, but the general patentability requirements of novelty, inventive step and industrial applicability apply to all types of inventions, including those in emerging fields - see comments received from Tanzania in response to C.9304.

³⁴ For example, the Russian Federation, see comments received from the Russian Federation in response to C.9199.

on one of the fundamental principles of the patent system, that is, dissemination of new technological knowledge.

107. 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?

108. 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 be 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 the categories described in Part II.F, as further discussed from the perspective of human contribution above.](#) At this point, there are very few official guidance that specifically address patent law questions applied to AI-related inventions. As AI is 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 needs to be seen whether AI-related inventions *per se* would bring additional complication to such already complex questions.

109. 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.

110. Recent advancements in AI, particularly in areas like deep learning and generative models, have led to increasingly sophisticated AI systems capable of performing complex tasks with minimal human intervention. These developments have sparked significant discussions regarding the implications of AI-generated inventions on the current patent system.³⁵ The debate over whether AI can be considered an inventor under existing patent laws has risen in prominence, with some jurisdictions issuing guidance on this matter.

111. Moreover, the rapid integration of AI into various sectors has underscored the need for clearer guidelines and standards concerning the patentability of AI-related inventions for legal certainty, consistency and fairness in the patenting process. The unique nature of AI, characterized by its ability to learn and evolve autonomously, challenges traditional patentability requirements, such as novelty, inventive step, and sufficient disclosure. The black-box nature of

³⁵ For example, see comments received from Australia and Brazil in response to C.9199.

some AI models, especially in deep learning, complicates the requirement for a clear and reproducible description of the invention.

B. Patentable Subject Matter

112. 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 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.³⁶ 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,³⁷ 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 law³⁸ 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.³⁹

113. Using the taxonomy introduced in Part II.F, this Section concerns primarily patentable subject matter issues for AI models and AI-based inventions, while AI-assisted inventions are addressed separately where relevant. 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 such, many Member States consider AI-related inventions to be a subset of computer implemented inventions.⁴⁰ 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.

114. Patent eligibility of computer-implemented inventions (CII) or software-implemented inventions has already been one of the areas where it has been 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.”

115. AI is generally considered as a branch of computer science as it involves the study of computation and information comprising algorithmic data structures, programming languages, computer architecture and software engineering, including the ways and means to represent,

³⁶ See “Certain Aspects of National/Regional Patent Laws – Exclusions from patentable subject matter” at: https://www.wipo.int/scp/en/annex_ii.html.

³⁷ The United States of America.

³⁸ Section 2(1) and (4) of the Japan Patent Act.

³⁹ 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).

⁴⁰ For example, see comments received from Chile, Lithuania, and Portugal in response to C.9199

store, process and transmit computational outputs and information in an efficient manner.⁴¹ However, there is significant debate over whether AI models and algorithms can readily qualify as CII.⁴² AI's statistical, probabilistic, and non-deterministic nature introduces new challenges and complexities that existing CII frameworks may not fully address for a variety of reasons. Computer programs consist of programming and instruction-based output, while AI models and machine learning systems function through an internal logic based on training rather than programming. Therefore, unlike conventional computer programs, which involve explicit and deterministic functional steps where the relationship between informational inputs and outputs are clear, trained neural networks are defined by numerical parameters whose relationship to network behavior may remain unclear even after adequate examination by humans to trace the AI system's thought process or the reasons for its decisions (i.e. commonly known as the "black box" problem).⁴³ Thus, differences between CII and AI inventions raise some fundamental doctrinal issues for patent law and policy in the context of subject matter eligibility.

116. While the related doctrinal issues have not been widely discussed, except in academic literature, judicial developments in at least one jurisdiction have raised such concerns, offering instructive policy insights. In *Emotional Perception AI v Comptroller-General of Patents, Designs and Trademarks*, the High Court of the United Kingdom held that AI inventions that rely on Artificial Neural Networks (ANNs) could not be excluded based on the application of the existing CII framework, on the assumption that AI models are computer programs as such and hence non-patentable subject matter.⁴⁴ The discussions in the context of the case illustrate that the use of Artificial Intelligence (i.e. ANNs for media file recommendations) does not fit well within the traditional concept of a "program for a computer". At least two key observations from the case highlight the doctrinal issues on subject matter eligibility of AI models. First, the Court distinguished traditional program instructions prescribed by a human from the operation of a trained ANN, whose operative parameters are produced through learning and are not implemented as code given to it by a human. It therefore concluded, in the circumstances of the case, that the trained emulated ANN could be distinguished from the underlying software and was not itself a "program for a computer" for the purposes of the exclusion. Second, the case illustrates that applying the computer program exclusion to a self-trained ANN requires careful identification of the relevant computer, the relevant program, what is actually claimed and, if necessary, whether the claimed invention produces a relevant technical effect. Although a computer program was involved in initiating or achieving the training, the Court held that it was only a subsidiary part of the claim and did not represent the invention claimed.⁴⁵ Although this case was later overturned in appeal (see para 131 below), the observations made by the Court provide a framework of analysis to understand the points of differences between CII and AI inventions.

117. The CII framework presents a further difficulty with respect to "gray-zone" AI technologies, most notably AI simulations. AI simulations employ data-driven algorithms to detect patterns, generate outputs, and improve processes without explicit programming. Simulations are needed to provide a real-world scenario for training AI models. For example, for short-mobility AI devices used in different uses cases on factory floors, it would be important to simulate the real-world environment for training the AI model and for further testing it based on synthetic data created for such purposes using AI models there by complementing and optimizing the deployment of the models. As mathematical functions converting input into output through

⁴¹ Is AI Computer Science, CRLN Team (2025), available at: <https://www.crln.org/is-ai-computer-science/>

⁴² WIPO AI and IP Toolkit, p. 13.

⁴³ See paras 80 and 81 above. Also see, AI's mysterious 'black box' problem, explained, University of Michigan-Dearborn (2023), available at: <https://umdearborn.edu/news/ais-mysterious-black-box-problem-explained>

⁴⁴ [2023] EWHC 2948 (Ch). The case involved the question of patentability of use of Artificial Neural Networks (ANNs) for providing an improved system of offering media file recommendations and sending the file and a message to an end user based on human perception and emotion irrespective of the genre of music and conventional classifications.

⁴⁵ Ibid. paras 32, 54, 56-61, 63 and 76-78.

learned parameters, AI simulations primarily produce informational outputs (namely predictions, classifications, or recommendations) without necessarily causing concrete changes to hardware. A patent may accordingly be refused on the grounds that the output constitutes informational content based on an abstract mathematical concept/methods. Applying the existing case law to the question of computer-implemented simulation involving a method implemented through a computer program and an apparatus used for simulating the movement of a pedestrian crowd through an environment like a railway station for crowd management purposes, the EPO Enlarged Board of Appeals concluded in G1/19 – Pedestrian Simulation (2021) that the downstream technical effect must relate to a physical reality.⁴⁶ Therefore, an AI simulation used to assist in design or verification of other solutions may retain technical character, provided the claimed technical effect is clearly established and claims are appropriately limited.

118. The survey below of various jurisdictions indicates that no jurisdiction has enacted amendments to its patent legislation specifically addressing the patent eligibility of AI inventions. However, legislative proposals have emerged in certain jurisdictions in the context of broader debates on subject matter eligibility reforms. Such proposals, if adopted, could materially affect the treatment of AI inventions by limiting judicially developed subject matter exclusions and replacing them with a closed, statutory list of ineligible categories. Under such proposals, eligibility assessments would consider the claimed invention as a whole, as understood by a person having ordinary skill in the art, focusing on whether the invention exists in nature independently of human activity or solely in the human mind.⁴⁷

119. As to the patent eligibility requirement applied to AI-related inventions, some patent offices have issued guidance pertaining to AI-related inventions which are described below.

120. In Australia, IP Australia considers that AI-related inventions are generally a subset of computer-implemented inventions for the purpose of subject matter eligibility. Although Australian courts have not yet considered the patentability of an invention including or using AI specifically, recent developments concerning computer-implemented inventions are relevant to AI-related applications. In *Aristocrat Technologies Australia Pty Ltd v Commissioner of Patents* [2025] FCAFC 131, the Full Court of the Federal Court confirmed that a computer-implemented invention should be properly characterized as a matter of substance, and that the relevant question is whether the claimed subject matter is merely an abstract idea manipulated on a computer or an abstract idea implemented on a computer to produce an artificial state of affairs and a useful result. The High Court refused special leave to appeal in *Commissioner of Patents v Aristocrat Technologies Australia Pty Ltd* [2026] HCADisp 15.

121. Following *Aristocrat*, IP Australia updated its Patent Manual of Practice and Procedure. The Manual now states that computer-implemented inventions may constitute patentable subject matter where, as properly characterized, they involve the implementation of an abstract idea on a computer to produce an artificial state of affairs and a useful result, and provides a decision-making framework considering, inter alia, the physical elements of the claim, whether the computerized product or system works in an improved or altered manner, whether a relevant technical effect occurs outside the computerized product or system, and whether a technical problem is solved with the computerized product or system. Under these principles, patentability may be found for AI-related claims where the AI or computer implementation,

⁴⁶ G 0001/19 (Pedestrian simulation) of 10.03.2021, EPO Enlarged Board of Appeals

⁴⁷ Congressional Research Service, Patent-Eligible Subject Matter Reform: An Overview (2026). Available at: https://www.congress.gov/crs_external_products/IF/PDF/IF12563/IF12563.4.pdf

properly characterized, is not merely the manipulation of an abstract idea on generic computer technology, but contributes to such a patentable implementation.⁴⁸

122. In Brazil, the National Institute of Industrial Property (INPI) updated their Computer Implemented Inventions (CII) Guidelines in 2020 to address AI-related patent applications.⁴⁹ Patent eligibility is determined by whether the claimed subject matter falls within one of the exclusions from patentability in INPI's examination guidelines. The revised guidelines also emphasize that artificial intelligence techniques, including machine learning and deep learning tools, when applied to solving technical problems, may be considered inventions.⁵⁰ Here, AI models or algorithms, when taken detached from applications in a specific technical field, are considered to be a mathematical method or algorithm and thus not patent eligible. In contrast, eligible subject matter includes inventions related to changes in the AI core techniques, such as alterations to the training process or the development of a new neural network architecture, when such changes are justified considering the specific technical issue being addressed, in addition to being incorporated into a technical field and generating a technical effect. In Brazil, neural network architectures, when viewed alone, are considered mathematical methods. In regard to hardware based artificial intelligence systems, to ensure patent eligibility, it is essential to present the specific characteristics and details on how the hardware implementation is carried out, it is insufficient to merely mention a possible hardware interpretation.⁵¹ In 2025, INPI opened Public Consultation No. 03/2025 on draft Guidelines for the Examination of Patent Applications related to Artificial Intelligence; the draft treats AI-related inventions as a subgroup of computer-implemented inventions and retains the requirement that a patentable AI-related invention solve a technical problem, present a technical solution and produce a technical effect.⁵²

123. In China, the Patent Examination Guidelines were revised with effect from January 20, 2024 to address, among other matters, the examination of invention patent applications involving artificial intelligence and big data. Under those Guidelines, AI-related inventions may constitute patentable subject matter where the claimed solution, considered as a whole, uses technical means to solve a technical problem and produce a technical effect. For example, an AI or big-data algorithm improvement may constitute a technical solution where the algorithm has a specific technical association with the internal structure of a computer system and improves hardware computing efficiency or execution results. On December 31, 2024, CNIPA issued the Guidelines for Patent Applications for AI-Related Inventions (Trial Implementation), which provide further AI-specific guidance and examples. Those Guidelines confirm that AI-related claims must not merely fall within rules and methods for mental activities and must also satisfy the technical-solution requirement; they further clarify that merely adding an AI-related technical feature or generic computer hardware will not suffice where the substance of the claim remains an abstract mathematical algorithm or model-training method without a technical problem or technical effect. The Patent Examination Guidelines were further amended by Order No. 84, effective January 1, 2026, including by adding AI and big-data-related provisions

⁴⁸ See comments received from Australia in response to C.9199; Aristocrat Technologies Australia Pty Ltd v Commissioner of Patents [2025] FCAFC 131; Commissioner of Patents v Aristocrat Technologies Australia Pty Ltd [2026] HCADisp 15; IP Australia, "Commissioner of Patents guidance update regarding computer-implemented inventions", official notice, October 2025; IP Australia, "Update: Aristocrat Technologies Australia Pty Ltd v Commissioner of Patents", official notice, February 6, 2026; IP Australia, Patent Manual of Practice and Procedure, section 5.6.8.6, "Computer Implemented Inventions, Mere Schemes, and Business Methods", updated March 16, 2026; IP Australia, "IP Australia Patents Manual Update – March 2026".

⁴⁹ Resolution INPI/PR No. 411, of 2020.

⁵⁰ See comments received from Brazil in response to C.9199.

⁵¹ *Ibid.*

⁵² INPI, Consulta Pública INPI/PR No. 03/2025, "Minuta das Diretrizes de Exame de Pedidos de Patente relacionados à Inteligência Artificial", consultation period August 19 to October 19, 2025; INPI, Minuta das Diretrizes de Exame de Pedidos de Patente relacionados à Inteligência Artificial, August 2025, sections 1.1, 1.3–1.4 and 2.1–2.2.

concerning claims containing algorithmic features or business rules and methods, and by providing that applications involving data collection, label management, rule setting or recommendation decisions that violate law, social morality or public interests cannot be granted patent rights.⁵³

124. In France, the National Institute of Industrial Property (INPI) applies, in its examination guidelines, the analytical framework based on Article L. 611-10 of the Intellectual Property Code for determining patentable subject matter. Certain subject matter is not considered to be an invention, in particular mathematical methods, computer programs and business methods, but only when they are claimed as such. In the case of inventions using artificial intelligence, INPI does not apply a separate regime: AI, insofar as it is based on calculations, is considered to be a mathematical method implemented by computer and, when claimed as such, cannot be considered an invention. For an invention using AI to be patentable, the AI must contribute to the technical character of the invention by providing a technical solution to a technical problem through non-generic technical means or through the processing of measured technical data. Examples considered technical include the processing, recognition or classification of images or videos, real-time control in robotics or monitoring and control processes, classification of IP traffic using machine learning to improve traffic management, and the use of a neural network in a cardiac monitoring device to perform predictive analysis and detect an irregular heartbeat.⁵⁴

125. To date, there is no established case law in Germany specific to the patentability of AI-related inventions. However, as these inventions are conceptually similar to computer-implemented inventions, eligibility is thus generally addressed by applying the three-stage examination approach for program-related inventions which was established by the German Federal Court of Justice (BGH) on the basis of Sections 1, 3, and 4 of the Patent Act of Germany.⁵⁵

126. In India, the 2025 CRI Guidelines contain a dedicated chapter on inventions related to AI, ML and DL. The Guidelines state that such technologies are fundamentally built on mathematical models, learning algorithms, computational models and abstract principles. An abstract idea, mathematical formula, or theoretical AI/ML/DL construct is not patentable for lack of practical application, but an AI/ML/DL innovation may be patentable when it transforms abstract principles into a real-world tangible application and provides an implementable technical solution to a problem through technical means. The Guidelines also identify technical implementations, such as efficient training for machine-learning models or inventive neural-network architectures that improve performance or reduce computational cost for a specific technical task, as examples that may avoid exclusion as a computer program per se.⁵⁶

127. Starting with the Delhi High Court's decision in *Ferid Allani v. Union of India* in 2019 which allowed patentability for computer-related invention when a "technical effect" was shown, Indian courts have also subsequently addressed AI/ML-related patent eligibility under section 3(k). In *Dr. Vandana Parvez v. Controller of Patents*, the Madras High Court set aside a refusal of an ML-related online advertising and personalization application, noting prima facie technical advancement and technical effect. Conversely, in *Caleb Suresh Motupalli v. Controller of Patents*, the Madras High Court upheld refusal of an AI, augmented-reality and mixed-reality-

⁵³ CNIPA, Interpretation of the 2023 Amendments to the Patent Examination Guidelines (IV) – Examination of Patent Applications for Inventions Involving Computer Programs, January 18, 2024; CNIPA, Guidelines for Patent Applications for AI-Related Inventions (Trial Implementation), December 31, 2024; CNIPA, FAQ concerning the Guidelines for Patent Applications for AI-Related Inventions (Trial Implementation), December 31, 2024; CNIPA, Decision on Revising the Patent Examination Guidelines (Order No. 84), November 13, 2025, effective January 1, 2026.

⁵⁴ See comments received from France in response to C.9304.

⁵⁵ See comments received from Germany in response to C.9199.

⁵⁶ Office of Controller General of Patents, Designs and Trade Marks, Guidelines for Examination of Computer Related Inventions (CRIs) - 2025, sections 4.5 and 5.1-5.2

related application, holding that the claimed invention lacked working examples and technical effect and did not surmount section 3(k).⁵⁷

128. Examples relating to AI inventions are also included in the Examination Handbook for Patent and Utility Model, issued by the Japan Patent Office (JPO).⁵⁸ These examples illustrate the criteria for determining whether *AI Invention is a creation of a technical idea utilizing the laws of nature*, which is essential for patent eligibility. In the Republic of Korea, the Ministry of Intellectual Property (MOIP) has prepared a detailed examination practice guideline in the artificial intelligence field.⁵⁹ Similar to other offices, the guidelines state that the criteria on patent eligibility for AI-related inventions are in principle the same as those for computer software-related inventions. In general, AI inventions will be patent eligible when: (i) information processing is performed by a “combination of software and hardware”; and (ii) the claims do not involve a “human mental activity or offline activities.” The *Examination Practice Guidelines prepared by MOIP* also include examples relating to AI inventions.

129. Singapore similarly provides specific guidance for determining the patent eligibility of AI-related inventions.⁶⁰ According to the Examination Guidelines, AI and machine learning methods, such as neural networks, support vector machines, discriminant analysis, decision trees, and k-means are considered mathematical methods and are not regarded as inventions on their own.⁶¹ Merely implementing an AI method using conventional computer hardware is unlikely to meet the eligibility criteria unless the actual contribution goes beyond the underlying mathematical method. However, AI-related inventions that apply AI to solve specific problems, like recognizing human speech or images, may be considered patent-eligible. The claim should be functionally limited to solving the specific problem, either explicitly or implicitly, by establishing a sufficient link between the problem and the steps of the mathematical method.⁶² For instance, specifying how the input and output of the sequence of mathematical steps relate to the problem ensures that the method is causally linked to solving it. Furthermore, AI methods claimed with reference to their implementation on a computer or using computer hardware must demonstrate that they solve a specific problem. If the application involves merely using conventional hardware to implement a machine learning method, it is unlikely to be patent-eligible, unless the hardware interaction addresses a specific problem to a material extent.

130. Additionally, in October 2024, the Intellectual Property Office of Singapore (IPOS) introduced the Supplemental Guidance for the Examination of AI-related Patent Applications. This resource complements the existing patent examination guidelines by providing greater clarity on the patentability of AI inventions. It sets out current examination practices and includes hypothetical examples illustrating the approach to assessing patentable subject matter for inventions in the AI field.⁶³

131. In the United Kingdom, the courts have considered the subject matter eligibility of artificial neural networks (ANNs) in *Emotional Perception AI Limited v Comptroller General of Patents, Designs and Trade Marks*. The High Court had held that the use of ANNs for media file

⁵⁷ Madras High Court, *Dr. Vandana Parvez v. Controller of Patents*, CMA(PT) No. 33 of 2023 (Feb. 23, 2024); Madras High Court, *Caleb Suresh Motupalli v. Controller of Patents*, C.M.A. (PT) No. 2 of 2024 (Jan. 29, 2025)

⁵⁸ Annexes A and B 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.

⁵⁹ *Examination Practice Guidelines by Technology Field*, Ministry of Intellectual Property, available at: https://www.kipo.go.kr/upload/en/download/EXAMINATION_PRACTICE_GUIDELINES_BY_TECHNOLOGY_FIELD.pdf; see comments received from the Republic of Korea in response to C.9304.

⁶⁰ Examination Guidelines for Patent Applications at IPOS, paras 8.22-8.27.

⁶¹ *Ibid.*

⁶² *Ibid.*

⁶³ See comments received from Singapore in response to C.9304.

recommendations were outside the scope of computer program exclusion since AI driven ANNs were conventionally not considered as computer programs (see the discussion in para 116 above). The Court of Appeal held that an ANN, whether implemented in hardware or software, fell within the computer program exclusion, and that the claimed recommendation system did not produce a relevant technical effect. In February 2026, the UK Supreme Court allowed the appeal. The Supreme Court held that an ANN is a program for a computer within the meaning of the computer program exclusion, but that the claimed subject matter was not a computer program “as such”, because the claimed method involved technical means, including the hardware on which the ANN was implemented and, in the relevant claims, a database, communications network and user device. The Supreme Court rejected the *Aerotel* approach and adopted the same interpretation of Article 52 EPC as the EPO Enlarged Board in G 1/19, under which the “invention” requirement is a low threshold and is assessed separately from novelty, inventive step and industrial applicability. The application was remitted to the UKIPO to apply the “intermediate step” and the remaining requirements for patentability.⁶⁴ Following the judgment, the UKIPO withdrew its AI examination guidelines and accompanying scenarios; the UKIPO Manual of Patent Practice notes that its guidance is subject to change pending its review of the Supreme Court judgment.⁶⁵

132. In the United States of America, the 2019 Revised Patent Subject Matter Eligibility Guidance issued by the United States Patent and Trademark Office (USPTO) includes one example that specifically discusses patent eligibility of a computer implemented method of training a neural network for facial detection comprising a series of steps for such training.⁶⁶ The USPTO subsequently updated its eligibility guidance for AI inventions in 2024. Central to the updated guidance is the application of the framework from the Supreme Court’s decision in *Mayo*, *Myriad*, and *Alice*. Step one of the USPTO’s subject matter eligibility analysis addresses whether the claimed invention is a new and useful process, machine, manufacture, composition of matter or improvement thereto (i.e., statutory subject matter). If an invention is not a process, machine, manufacture, or composition of matter, the invention is statutorily not patentable. Step 2A prong one involves determining whether a claim is directed to a judicial exception to patentability such as an abstract idea, law of nature, or natural phenomenon. If the claim does not recite a judicial exception, it is considered eligible, and the eligibility analysis ends. If a judicial exception is identified, step 2A prong two assesses whether the claim elements, individually, or in combination, add significantly more to the judicial exception, thus transforming it into a patent-eligible application. If instead, the claim is “directed to” the judicial exception, the analysis will proceed to step 2B to evaluate whether the claimed additional elements amount to significantly more than the recited judicial exception itself. Step 2B includes a consideration of whether the additional element is a well-understood, routine, conventional activity. A claim may be found to lack significantly more (and thus be ineligible) based on one or more of these judicial considerations (e.g., a conclusion that the additional limitation(s) are insignificant extra-solution activity or more instructions to apply an exception). In contrast, if USPTO personnel determine in Step 2B that the additional elements do amount to significantly more than the judicial exception, the claim is patent eligible.

133. In relation to AI inventions, the updated guidelines provide additional clarity in particular to two areas of the above-mentioned test: (1) the evaluation of whether a claim recites an abstract idea in Step 2A prong one; and (2) the evaluation of the improvement consideration in step 2A prong two. The updated guidelines provide considerable commentary and hypothetical

⁶⁴ [Emotional Perception AI Limited v Comptroller General of Patents, Designs and Trade Marks \[2026\] UKSC 3](#), UK Supreme Court, judgment of February 11, 2026.

⁶⁵ UKIPO, Guidelines for examining patent applications relating to artificial intelligence (AI), withdrawn on February 11, 2026 following the Supreme Court judgment in *Emotional Perception AI Limited v Comptroller General of Patents, Designs and Trade Marks [2026] UKSC 3*; UKIPO, Scenarios applying the guidelines for examining patent applications for AI, withdrawn on February 11, 2026; UKIPO, Manual of Patent Practice, Sections 1 and 3, noting that guidance is subject to change pending review of that judgment.

⁶⁶ 2019 Revised Patent Subject Matter Eligibility Guidance, Example 39.

examples of how these steps in the test are applied to AI related technologies.⁶⁷ While the new guidelines emphasize that the *Alice/Mayo* test for analyzing the subject matter eligibility has not changed, the new guidance introduce three new examples that either pass or fail the test. These examples are related to “AI for Anomaly Detection Using Neural Networks,” “AI for Analyzing Speech Signals,” and “AI for Personalized Medicine.”⁶⁸ The United States further indicated that the July 2024 guidance built on existing guidance by providing further clarity on how the USPTO and applicants should evaluate subject matter eligibility of claims in patent applications and patents with inventions related to AI technology, and that it did not treat AI inventions differently from other inventions.⁶⁹

134. At the EPO, artificial intelligence and machine learning are addressed in the Guidelines for Examination under the guidance on mathematical methods. A dedicated subsection on artificial intelligence and machine learning was introduced in the 2018 Guidelines and updated in the 2025 Guidelines. The current Guidelines state that artificial intelligence and machine learning are based on computational models and algorithms, such as artificial neural networks, genetic algorithms, support vector machines, k-means, kernel regression and discriminant analysis, which are per se of an abstract mathematical nature, irrespective of whether they can be trained using training data. However, the use of artificial intelligence or machine learning does not by itself render an invention non-patentable. If a claim directed to an AI- or machine-learning-related invention is directed either to a method involving the use of technical means, such as a computer, or to a device, its subject matter has technical character as a whole and is not excluded from patentability under Article 52(2) and (3) EPC. In such cases, the computational models or algorithms themselves contribute to the technical character of the invention where they contribute to a technical solution to a technical problem, for example by being applied in a field of technology or by being adapted to a specific technical implementation.⁷⁰

135. The EPO Guidelines give examples of technical applications, including use of a neural network in a heart-monitoring apparatus to identify irregular heartbeats, and classification of digital images, videos, audio or speech signals based on low-level features. Conversely, classifying text documents solely according to their textual content, or classifying abstract data records without any indication of a technical use of the resulting classification, is not per se considered a technical purpose.⁷¹

136. In April 2025, in *Recentive Analytics, Inc. v. Fox Corp.*, the United States Court of Appeals for the Federal Circuit reaffirmed this approach and confirmed that the two-part test commonly used to evaluate patent subject matter eligibility should also be applied to AI inventions. In August 2025, the USPTO released a memorandum to patent examiners in certain Technology Centers that provided targeted reminders on evaluating subject matter eligibility for software-related inventions, including artificial intelligence and machine learning technologies; the memorandum reinforces existing USPTO practice and clarifies common issues in applying the patent eligibility analysis.⁷² In a September 2025 decision in *Ex parte Desjardins*, the USPTO found a set of AI machine learning claims patent subject matter-eligible after finding that, even though the claims are directed to an abstract idea in the form of a mathematical algorithm, they

⁶⁷ 2024 Guidance Update on Patent Subject Matter Eligibility, Including on Artificial Intelligence, available at: <https://www.federalregister.gov/documents/2024/07/17/2024-15377/2024-guidance-update-on-patent-subject-matter-eligibility-including-on-artificial-intelligence>.

⁶⁸ July 2024 Subject Matter Eligibility Examples: <https://www.uspto.gov/sites/default/files/documents/2024-AI-SMEUpdateExamples47-49.pdf>.

⁶⁹ See comments received from the United States of America in response to C.9304.

⁷⁰ Guidelines for Examination in the European Patent Office, Part G, Chapter II, sections 3.3 and 3.3.1, “Artificial intelligence and machine learning”; EPO, List of sections amended in the 2025 revision, Part G, G-II, 3.3.1, “Clarification of definition of artificial intelligence and machine learning”; EPO, Index for Computer-Implemented Inventions, noting that G-II, 3.3.1 was introduced in GL 2018 and updated in GL 2025.

⁷¹ Ibid.

⁷² See comments received from the United States of America in response to C.9304.

recite additional elements that integrate the abstract idea into a practical application. Specifically, for purposes of determining subject matter eligibility, the claims were found to sufficiently claim an improvement in how the AI machine learning model operates, not merely the underlying mathematical calculation; thus, the *Desjardins* decision confirmed that AI machine learning inventions are not categorically barred from patentability based on their recitation of a mathematical algorithm.⁷³

137. In December 2025, the USPTO also issued two memoranda providing additional clarification regarding the use of Subject Matter Eligibility Declarations (SMEDs). Under USPTO rules, applicants may submit a SMED to provide factual evidence relevant to the eligibility inquiry, such as evidence of technological improvement, the state of the art at the time of filing, or information demonstrating how a judicial exception is integrated into a practical application. The first memorandum, directed to the Examining Corps, explains that SMEDs are a voluntary option and clarifies that, when a SMED is properly submitted, examiners must consider it as part of the evidentiary record and evaluate it using the preponderance-of-the-evidence standard; the second memorandum, directed to applicants and practitioners, notes that SMEDs should provide objective evidence tied to the claimed invention and should not be used to supplement the original disclosure.

138. Together, the above two memoranda reinforce that SMEDs do not alter existing USPTO procedures, but clarify how applicants may provide factual evidence relevant to the eligibility analysis and how examiners should evaluate such submissions under longstanding practice.⁷⁴

139. In the Philippines, the current Intellectual Property Code (IP Code) does not provide a prohibition as to the extent of the use of AI in the generation of an invention or design, and such use cannot, by itself, be considered as a ground for non-patentable inventions. However, extensive use of AI may have its own legal consequences under other relevant laws, such as copyright, unfair competition and infringement.⁷⁵

140. Regarding AI-assisted inventions, the consideration of patentable subject matter depends on the nature of the resulting claimed 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.

141. In the context of AI models, distinct concerns have emerged about patent eligibility relating to training data, data classification, methods of training data, data structures, and trained models. Claims directed to training data “per se” are generally considered mere presentations of information and do not qualify as patentable subject matter, notwithstanding the critical role of training data in the performance of the AI model. Examples provided in the Japan Patent Office’s Guidelines clearly illustrate that claims directed at the training data itself are not patentable. Claim 1 of example 3, cited as examples as part of the JPO AI patent guidelines relates to “sugar content data of apples” where the claim is directed at the data itself. Accordingly, as per the JPO, “claim 1 does not have technical features in the presentation of information (i.e. presentation per se, means for presentation or method of presentation), its feature resides solely in the content of the information, and its main object is to present information”. Similarly, claim 1 of example 5 of the same guidelines state that the use of training data “to train a discriminator” is something that is very normal operation in machine learning of discriminators. As per the JPO, such a claim “does not bring any technical features to learning means or methods for discriminators” and is considered to be “characterized only in the content of information that it consists of images for training data” generated by combining several data

⁷³ See comments received from the United States of America in response to C.9304.

⁷⁴ See comments received from the United States of America in response to C.9304.

⁷⁵ See comments received from the Philippines in response to C.9304.

inputs.⁷⁶ The EPO guidelines and case law also clearly highlight that patent claims directed at simple classification of data or processing of such data, including data that is technically sourced, are not patentable since there are no linkages between the classification result and a technical purpose or technical application.⁷⁷

142. Methods of generating training data, or methods that concretely perform information processing based on technical properties of an object, may satisfy eligibility requirements where software and hardware are found to operate together. Claim 3 of example 5 cited as examples as part of the JPO AI patent guidelines allows claims directed radiological images of the human body when combined with images of surgical instruments to generate new images for use as training data. Such training data may be generated through defined information-processing operations performed by software using hardware resources.⁷⁸ Similarly, as per example 47, claims 2 of the AI patent subject matter eligibility guidelines of the USPTO, a claim relating to the training and use of an artificial neural network (ANN) was found not to constitute patent-eligible subject matter. The claimed method involved receiving and discretizing training data, training the ANN using backpropagation and gradient descent algorithms, and using the trained ANN to detect and analyze anomalies in a data set. The relevant steps were considered to encompass mathematical concepts and mental processes, while the use of a generic computer and a trained ANN did not integrate those abstract ideas into a practical application or provide an inventive concept. Accordingly, the claim was considered ineligible for patent protection.⁷⁹

143. Data structures may qualify where the data structure is functionally equivalent to a program satisfying the software-hardware interaction requirement. Example (Case 2-15) dealing with 3D Printing Data cited as example as part of the JPO AI patent guidelines states that structured data may, in certain circumstances, be considered patent-eligible subject matter where its structure determines the processing performed by a computer or other technical system. For example, 3D printing data may contain, for each layer of a three-dimensional object, pointers to model-material data and supporting-material data. These pointers may enable the control unit of a 3D printer to retrieve the relevant data from a memory unit and perform the corresponding operations during the printing process. In such a case, the data structure defines the information processing performed by the control unit of the 3D printer and is used to control the operation of the device. The structured data may therefore have characteristics comparable to those of a computer program, insofar as its structure determines the processing carried out by the technical system. Under the applicable national law, such structured data may consequently be regarded as patent-eligible subject matter.⁸⁰

144. Trained models characterized as programs may satisfy the software-hardware interaction requirement. Case 39, cited by the USPTO as an example in the context of AI patent eligible subject matter relates to an example that concerns a computer-implemented method for training a neural network for facial detection. The method generates training data by applying transformations, such as mirroring, rotation, smoothing or contrast reduction, to existing facial images and combining the resulting images with facial and non-facial images to form a training dataset. Following an initial stage of training, non-facial images incorrectly classified as facial images are added to the training dataset, and the neural network is trained again. Thus, the claimed method was considered patent-eligible subject matter because it constituted a process and did not itself recite a judicial exception (i.e. abstract idea), notwithstanding that certain aspects of the method could be based on mathematical concepts.⁸¹ Example from the JPO,

⁷⁶ Annex A of the Examination Handbook for Patent and Utility Model, Japan Patent Office.

⁷⁷ T 1784/06 (Classification method/COMPTTEL) of 21.09.2012.

⁷⁸ Example 5 of the Annex A Cases pertinent to Eligibility for Patent and Industrial Applicability, JPO

⁷⁹ USPTO, 2024 Guidance Update on Patent Subject Matter Eligibility, Including on Artificial Intelligence (July, 2024 subject matter eligibility examples).

⁸⁰ Annex A of the Examination Handbook for Patent and Utility Model, Japan Patent Office [Case 2-15] 3D Printing Data.

⁸¹ Example 39 of USPTO Subject Matter Eligibility Examples: Abstract Ideas.

(Case 2-14) concerns a trained model where the claim was considered patent-eligible subject matter where it functioned as a program that caused a computer to perform specific information processing. In the example, the trained neural-network model used trained weights and input data concerning the frequency of particular words to calculate and output quantified values relating to the reputation of accommodations. Since the trained model defined specific processing performed through the interaction of software and hardware resources and thereby caused the computer to perform a particular function, it was considered to constitute an invention, notwithstanding that the claimed subject matter was described as a model rather than a program.⁸²

145. The examples above specifically in the context of AI patent eligibility of claims related to training data, data structures, and trained models suggest that different approaches to patent eligibility converge on some key elements. The characterization of AI-related data or models does not, by itself, determine their eligibility for patent protection. Consideration is given to their functional or technical contribution in the context of the claimed invention. Training data or classifications of data, claimed merely as information or defined principally by their informational content, may not constitute patent-eligible subject matter. By contrast, methods of generating or processing training data, structured data and trained models may be eligible where they contribute to a technical purpose or application, determine the operation or processing of a technical system, or implement specific information processing through the interaction of software and hardware. At the same time, the mere use of a computer, mathematical methods, an ANN or other AI techniques may not be sufficient where the claimed subject matter remains directed to an abstract idea without further technical implementation or application. Accordingly, patent eligibility in the AI context may depend less on whether the claimed subject matter is labelled as data, a data structure, a training method or a trained model, and more on the nature and function of the claimed subject matter, its relationship with a technical system or purpose and the applicable requirements under national or regional patent law.

C. Novelty and Inventive Step

146. It is said that inventive step analysis is the most difficult requirement in the patentability criteria to assess.⁸³ 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. Questions have emerged regarding how to analyze the inventive step and define the concept of a person skilled in the art for inventions involving the use of AI.⁸⁴ 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).⁸⁵

147. 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 scope 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 matures, common interpretations and standard practices have gradually emerged in many technology areas.

⁸² Annex B of the Examination Handbook for Patent and Utility Model, Japan Patent Office.

⁸³ 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.

⁸⁴ See comments received from Chile in response to C.9199.

⁸⁵ WIPO Technology Trends 2019 – Artificial Intelligence, pp.115 to 117.

148. 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.⁸⁶ 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.⁸⁷ 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.”⁸⁸

149. In China, when conducting a novelty examination of an invention patent application that involves the algorithmic features or business rules and methods, all the features described in the claims should be considered. These features include both the technical features and algorithmic features or business rules and methods features.⁸⁹ Furthermore, if an algorithm, which has a specific technical relationship with the internal structure of a computer system, enhances the system’s internal performance, then it should be considered when evaluating inventiveness. Similarly, if a solution improves user experience through technical features, or through the interaction of technical features, algorithmic features, and business rules or methods, this enhancement in user experience should be taken into account when assessing inventiveness.⁹⁰ For example, if the algorithm in the claim is applied to a specific technical field, and can solve a specific technical problem, then it can be considered that the algorithmic features and the technical features functionally support each other and have an interactive relationship, and the algorithmic features become part of the technical means adopted.⁹¹

150. In France, patent applications in the field of artificial intelligence are examined in the same way as patent applications concerning computer-implemented inventions. For the assessment of inventive step in mixed-type inventions, namely claims that include both technical and non-technical features, only features that contribute to the technical character of the invention are taken into account. This may include features that are not technical when considered in isolation but that, in the context of the invention, contribute to producing a technical effect serving a technical purpose; conversely, features that do not contribute to the technical character of the invention cannot support the existence of an inventive step.⁹²

151. Annex A of the Examination Handbook for Patent and Utility Model, issued by the JPO, contains several examples relating to assessment of inventive step for AI-related invention.⁹³ At the EPO, in regard to inventive step, AI-related inventions are assessed in a similar way to other computer implemented inventions. Here, it is made sure that only features contributing to the technical character of the invention are considered for the assessment of inventive step. In particular, “non-technical” features, understood in this context as features which, on their own, would fall within a field excluded from patentability, can only be considered for assessment if they contribute to solving a technical problem. Even technical features may be ignored with regard to assessment of inventive step if they do not contribute towards solving a technical

⁸⁶ See document SCP/22/3.

⁸⁷ Document SCP/22/3, paragraphs 34 and 35.

⁸⁸ See document SCP/28/4.

⁸⁹ CNIPA Guidelines for Examination Section 2, Chapter 9, Article 6.1.3

⁹⁰ The Fourth Interpretation of the Revisions to the Patent Examination Guidelines (2023) – Examination of Patent Applications for Invention Involving Computer Programs.

⁹¹ JPO - CNIPA Comparative Study on AI-Related Inventions.

⁹² See comments received from France in response to C.9304.

⁹³ Annex A of the Examination Handbook for Patent and Utility Model, Examples 31 to 40, JPO.

problem. There have been several EPO Board of Appeal Decisions that pertain to the assessment of inventive step in relation to inventions pertaining to AI and machine learning.⁹⁴

152. 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.⁹⁵

153. In the Philippines, despite the absence of a prohibition on AI-assisted inventions or AI-generated designs, an application aided by AI should be examined with extra diligence or prudence, including as to whether the materials from which the AI may have based its output may be considered as prior art.⁹⁶

154. 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.

155. 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.⁹⁷ Such a policy objective may guide the determination of the inventive step for each case, including the AI-related inventions.

D. AI-Generated Prior Art and PHOSITA in the Context of Assessing Patentability

156. AI systems are capable of generating technical disclosures continuously and at a scale and speed that may exceed the capacity of existing prior art search, classification and retrieval systems.⁹⁸ Such disclosures may include, for example, predictions of protein structures with potential relevance to structural biology, drug discovery, molecular biology and biotechnology.⁹⁹ The resulting volume of information may make it more difficult for applicants, examiners and third parties to identify disclosures that are technically relevant, publicly accessible and sufficiently reliable. The materials discussed below do not indicate that jurisdictions have adopted legislative or administrative reforms specifically governing AI-generated prior art.

157. The European Patent Office (EPO), however, has clarified that AI-generated material may form part of the state of the art if it was made available to the public before the relevant date and contains an enabling technical teaching that the person skilled in the art can put into

⁹⁴ Artificial Intelligence and Machine Learning Case Law, available at: https://www.epo.org/en/legal/case-law/2022/clr_i_d_9_2_11_e.html.

⁹⁵ All Prior Art project (<https://allpriorart.com/about/>).

⁹⁶ See comments received from the Philippines in response to C.9304.

⁹⁷ Document SCP/22/3, paragraph 3.

⁹⁸ Intellectual Property Expert Group, Navigating the Challenges of AI-Generated Prior Art, available at: <https://www.ipeg.com/navigating-the-challenges-of-ai-generated-prior-art/>

⁹⁹ Zhidian Zhang, Chenxi Ou, Yehlin Cho, Yo Akiyama, Sergey Ovchinnikov, Artificial intelligence methods for protein folding and design, Current Opinion in Structural Biology, Volume 93, 2025.

practice.¹⁰⁰ Under that approach, the origin, geographical location or form of a disclosure is not, in itself, determinative. The principal considerations are the technical information made available to the public, the relevant date and the enabling character of the disclosure. Existing prior art frameworks may therefore apply to AI-generated material. Nevertheless, the extent to which those frameworks can respond effectively to the scale, quality and reliability of AI-generated disclosures remains an open practical question.

158. Other patent offices, including the United States Patent and Trademark Office (USPTO), have initiated consultations concerning the possible impact of AI-generated disclosures on prior art and patentability determinations.¹⁰¹ These discussions have identified a number of questions that may also be relevant in other jurisdictions:

- whether applicable patent law presupposes that a prior art disclosure is authored or published by a human, and whether the absence or degree of human contribution affects the availability of an AI-generated disclosure as prior art;
- which forms of AI-generated disclosure may be relevant to the assessment of patentability, how such disclosures are currently made available to the public and whether additional mechanisms for ensuring public accessibility may be appropriate;
- whether a party submitting a publication or other evidence known to be AI-generated should notify the patent office of that fact, and whether that party should have any duty to determine how the disclosure was generated; and
- whether AI-generated disclosures should be treated differently from human-generated disclosures, particularly in view of possible inaccuracies or hallucinations and their implications for operability, enablement, evidentiary reliability and public accessibility.

159. The volume and speed at which AI systems may produce and disseminate technical information could have significant practical implications for the patent system. Large quantities of automatically generated material increase the time and resources required by applicants to assess the relevant state of the art and by examiners to search, classify and evaluate potentially relevant patent and non-patent literature. This risks burying relevant prior-art in the sea of irrelevant AI-generated prior art in many instances¹⁰². Accordingly, one view is that such AI-generated disclosures may have limited value and lacking in evidentiary weight as technical literature for engineers and scientists and may not reflect their actual work and that the scale of such AI-generated prior art may burden the patent system by attempting to preclude future inventions.¹⁰³

160. The accumulation of duplicative, speculative or technically unreliable material also makes pertinent disclosures more difficult to locate and could, in certain circumstances, create practical barriers to patentability notwithstanding the formal public availability of the material. In the absence of AI-specific rules, however, general principles governing internet-based disclosures appear applicable.¹⁰⁴ Accordingly, online material may be regarded as publicly accessible where, before the relevant date, members of the public were able to gain knowledge of its content without being subject to an obligation of confidentiality or another restriction on its use

¹⁰⁰ EPO, patentability of digital inventions, available at: <https://www.epo.org/en/news-events/in-focus/digital-innovations/patentability-digital-inventions>.

¹⁰¹ Request for Comments Regarding the Impact of the Proliferation of Artificial Intelligence on Prior Art, the Knowledge of a Person Having Ordinary Skill in the Art, and Determinations of Patentability Made in View of the Foregoing 89 FR 34217, APR 30, 2024, DOCKET NO. PTO-P-2023-0044.

¹⁰² Comments from BIO on USPTO's Request for Comments (89 FR 34217, APR 30, 2024, DOCKET NO. PTO-P-2023-0044).

¹⁰³ Ibid.

¹⁰⁴ Tarrere et. al. (2026), Ideas for Challenging AI-generated prior art at the EPO, available at: <https://www.santarelli.com/en/document-genere-par-ia-contester-oeb/>

or dissemination.¹⁰⁵ Whether a particular AI-generated disclosure satisfies that standard remains a fact-specific determination, including with respect to its date of availability, retrievability and accessibility to the relevant public.¹⁰⁶

161. A central issue is whether an AI-generated disclosure is enabling prior art.¹⁰⁷ For the purposes of novelty, a prior art document generally prejudices the novelty of a claimed invention only to the extent that it contains sufficient information, considered together with the common general knowledge at the relevant date, to enable the person skilled in the art to carry out the disclosed technical teaching. A mere reference to a concept, result, compound or structural feature may therefore be insufficient where the disclosure does not provide a practicable means of obtaining or implementing it. Actual reduction to practice or prior empirical testing is not necessarily required. A theoretical or predictive AI-generated disclosure may still be relevant to novelty if the disclosed subject matter can be put into practice by the person skilled in the art. In the assessment of inventive step, the inquiry is different and depends on the applicable national or regional approach. The relevant questions include whether the person skilled in the art would have regarded the disclosed teaching as a credible and technically meaningful starting point, or taken it into account, when seeking to arrive at the claimed invention.

162. AI-generated disclosures present particular difficulties in determining enablement and evidentiary reliability.¹⁰⁸ An AI-generated output does not reveal the model parameters, algorithms, training data, prompts or other conditions that contributed to its generation, which makes it difficult to establish whether the disclosed result is reproducible or to identify the reason for any failure to reproduce it. In addition, AI outputs may contain hallucinations, factual inaccuracies, internally inconsistent statements or imprecise technical descriptions.¹⁰⁹ These features do not necessarily mean that an AI-generated disclosure is non-enabling; each disclosure would need to be assessed on its content and in light of the common general knowledge of the person skilled in the art. They may, however, be particularly significant where large volumes of speculative, prophetic, unverified or technically incoherent material are automatically placed in the public domain. This raises the unresolved policy question of whether existing rules concerning public accessibility, evidentiary weight and enabling disclosure provide sufficient means of distinguishing technically meaningful prior art from material that merely describes an unsupported or inoperable possibility, or whether further clarification of those rules may be useful.

163. The available approaches suggest that the origin of a disclosure, whether human- or AI-generated, is generally not the sole determinant of its status or relevance as prior art. Subject to the applicable law, relevant considerations include whether the disclosure date can be established, whether the information was publicly accessible before the relevant date and whether the technical teaching is sufficiently enabling. An AI-generated disclosure that provides a technically reproducible teaching may therefore prejudice novelty or be relevant to the assessment of inventive step, whereas a bare concept, unsupported prediction or technically inoperable disclosure may be insufficient for those purposes. Although certain legal frameworks require the labeling of AI-generated content in other contexts, the materials reviewed do not indicate that such labeling requirements have generally been incorporated into patent law rules

¹⁰⁵ Ibid.

¹⁰⁶ Ibid.

¹⁰⁷ Holman et. Al., (2024), Top 5 Potential Implications of AI-Generated Prior Art on Patent Law, available at: <https://www.sterneckessler.com/news-insights/insights/top-5-potential-implications-of-ai-generated-prior-art-on-patent-law/>

¹⁰⁸ Tarrere et. Al. (2026), Ideas for Challenging AI-generated prior art at the EPO, available at: <https://www.santarelli.com/en/document-genere-par-ia-contester-oeb/>

¹⁰⁹ Mehta (2024), Prior art problems: counsel mull AI-generated disclosures, available at: <https://www.managingip.com/article/2d6pfurgj9wil10moq2o0/patents/prior-art-problems-counsel-mull-ai-generated-disclosures>

governing the identification or submission of prior art.¹¹⁰ Nevertheless, parties and practitioners submitting AI-generated references remain subject to applicable duties of reasonable inquiry, candor, good faith or professional responsibility.

164. The increasing availability and sophistication of AI tools raise questions concerning the legal construct of the person skilled in the art, also referred to in some jurisdictions as the person having ordinary skill in the art (PHOSITA). This hypothetical standard is relevant across several aspects of patent examination, including the interpretation of claims and prior art for the purposes of novelty, the assessment of inventive step or non-obviousness, and the determination of whether an invention has been disclosed sufficiently for the skilled person to carry it out.¹¹¹ In particular, the growing use of AI in research and development may affect which activities are regarded as routine, and which continue to require inventive skill. The principal question is therefore not whether AI makes inventions generally more or less inventive, but whether particular AI-assisted activities have become ordinarily available to, and routinely performed by, the skilled person in the relevant technological field at the applicable date¹¹².

165. The available national and regional practices do not indicate that an AI-specific reconstruction of PHOSITA, or of the common general knowledge (CGK) attributed to that person, has been generally adopted.¹¹³ Current approaches largely retain a human-centric legal construct and treat AI as a tool that is available to the person skilled in the art, rather than as a constituent element of that notional person. Whether the use of a particular AI tool forms part of the skilled person's ordinary capabilities will therefore depend, among other considerations, on its accessibility and customary use in the relevant field at the applicable date. Accurate characterization of those capabilities has horizontal implications. It may affect inventive step, the interpretation of claims, the assessment of anticipation or novelty and, where relevant under applicable law, written description, enablement and sufficiency of disclosure.

166. Against this background, a distinction may be drawn between, on the one hand, the legal identity or character of the notional person skilled in the art and, on the other hand, the knowledge, tools and routine capabilities attributed to that person. While no jurisdiction has issued PHOSITA guidelines in the context of challenges posed by AI, the USPTO has initiated consultations, which are discussed below.¹¹⁴ Such questions are conceptually relevant to the patent systems generally and could pose similar challenges in other jurisdictions. They include

- whether the person skilled in the art should continue to be understood as a human-centric construct or, where the nature of the technology so requires, as an interdisciplinary team, and whether either construct may be assumed to use AI tools without treating AI itself as the legal person;
- how the widespread availability and routine use of AI tools affect the CGK attributed to the skilled person, the interpretation of claim terms and the identification of analogous art for the purposes of inventive step;

¹¹⁰ China's AI Regulation: Interim Measures to the 2025 Labeling Rules, available at: <https://casrai.org/guides/china-ai-regulation-explained>

¹¹¹ Roberto Dini, Mario Franzosi, The so-called 'person skilled in the art': a donkey or a genius?, *Journal of Intellectual Property Law & Practice*, Volume 19, Issue 5, May 2024, Pages 387–389 and Roberto Dini, AI innovation and a 'person skilled in the art', *Journal of Intellectual Property Law & Practice*, Volume 20, Issue 10, October 2025, Pages 643–644.

¹¹² WIPO AI and IP Policy Toolkit, p. 15.

¹¹³ Ibid.

¹¹⁴ Request for Comments Regarding the Impact of the Proliferation of Artificial Intelligence on Prior Art, the Knowledge of a Person Having Ordinary Skill in the Art, and Determinations of Patentability Made in View of the Foregoing 89 FR 34217, APR 30, 2024, DOCKET NO. PTO-P-2023-0044.

- how the probabilistic and non-deterministic operation of certain AI systems affects the assessment of predictability and reasonable expectation of success;
- whether, and to what extent, AI-assisted search capabilities should be taken into account when identifying prior art for the purposes of novelty and inventive step;
- how the version, capabilities, training data and information available to an AI system at the effective filing date should be aligned temporally with the characterization of the skilled person; and
- how access to AI-assisted experimentation and optimization tools affect enablement and sufficiency of disclosure, including whether the experimentation required to carry out the full scope of a claimed invention would have been reasonable or undue for the person skilled in the art.

167. In the context of the European Patent Organization, while the approach to PHOSITA remains human-centric, materials show that similar concerns have been recognized in discussions among EPO Contracting States and the policy question of whether a machine can be recognized as a skilled person remains open. It is recognized that the use of AI may increase the knowledge of the skilled person, which is difficult to establish. Furthermore, it is also recognized that the skilled person must be an interdisciplinary team able of use AI.¹¹⁵

168. In certain AI-related inventions, the person skilled in the art may appropriately be represented by a team combining expertise in AI technology with expertise in the relevant application domain. For example, the assessment of an AI-based image-recognition system used for security and surveillance requires knowledge of both machine-learning techniques and surveillance technologies. This does not necessarily amount to recognition of an AI-augmented legal person; rather, it reflects the established possibility that the skilled person may, depending on the technical problem, be represented by a multidisciplinary research or production team.¹¹⁶ As standard AI tools become routinely used in particular fields, a degree of AI literacy or the ability to use such tools increasingly form part of the ordinary capabilities or CGK attributed to that person. Any such conclusion should, however, be based on evidence of the actual accessibility and adoption of the relevant tools in the field at the applicable date, rather than on their subsequent development or use in other technological sectors.¹¹⁷

169. A further difficulty arises from the rapid and continuous development of AI models and tools.¹¹⁸ Where an inventive-step argument assumes that the person skilled in the art would have used AI to identify, select or combine teachings from two or more prior art references, it would be necessary to establish what capabilities were actually available to that person at the priority date or effective filing date, as applicable. Reliance on the capabilities of a later or materially improved version of an AI tool could introduce hindsight into the analysis, particularly where successive versions differ substantially in their search, reasoning, prediction or optimization capabilities. Evidence concerning release dates, model or version information, conditions of access, training-data cutoffs and contemporaneous use may therefore assist the assessment. Digital records, logs or time-stamped information might also be relevant in an

¹¹⁵ EPO, Legal aspects of patenting inventions involving artificial intelligence (AI): Summary of feedback by EPC contracting states, available at: https://link.epo.org/web/AI_inventorship_summary_of_answers_en.pdf

¹¹⁶ Steadman et. al.,(2026), The AI-augmented POSITA: How artificial intelligence is reshaping the obviousness inquiry in US patent law, available at: <https://www.dlapiper.com/en/insights/publications/intellectual-property-news/2026/ai-augmented-posita-how-artificial-intelligence-is-reshaping-obviousness-inquiry>

¹¹⁷ Block and Ratz (2021), Killed In The Art? How Artificial Intelligence Challenges The Fictional Concept Of The Skilled Person In Patent Law, available at: <https://lesi.org/article-of-the-month/killed-in-the-art-how-artificial-intelligence-challenges-the-fictional-concept-of-the-skilled-person-in-patent-law/>

¹¹⁸ Ibid.

appropriate case.¹¹⁹ However, their possible evidentiary value and the circumstances in which they might be requested remain matters for further consideration and should not be understood as establishing a general recordkeeping obligation for applicants.¹²⁰

170. Under prevailing approaches, the PHOSITA generally remains a human-centric legal construct, with AI treated as a tool available to that person or, where appropriate, to an interdisciplinary team. Accordingly, the mere use of AI or software to implement a function previously performed by hardware, the transfer of a common software procedure to another field for the same function, or the computer implementation of a known human-performed business method will fall within ordinary skill where only routine system analysis or adaptation is required.¹²¹ For example, mere application of AI to a field of technology, patent offices that offer guidance by way of examples generally consider that an attempt to implement a function performed by hardware such as a circuit by means of software is within the normal craftsmanship of a person skilled in the art. Similarly, applying a common software procedure or means from one specific field to another will, where the function or action is the same, be regarded as falling within the ordinary capabilities of the skilled person. Where a cited invention discloses a human-performed business service or method, but not its systemization, implementing it by computer will still fall within the ordinary creative abilities of the skilled person if it can be achieved through routine system analysis and design.¹²² This assessment nevertheless depends on the capabilities and tools reasonably available in the relevant field at the effective filing date, since the increasing routine availability of AI affects the baseline knowledge and capabilities attributed to PHOSITA without necessarily altering the applicable legal standard.¹²³

171. With respect to the modification of training data, the assessment may distinguish between routine choices and configurations that produce an unexpected technical effect. The selection of variables that would ordinarily be expected to correlate with the desired output, including for improving the reliability or accuracy of a machine-learning system, will form part of common general knowledge and therefore will not support an inventive step. By contrast, inventive step may be present where input data are combined on the basis of a previously unknown correlation and specific training parameters produce an unexpected technical effect.¹²⁴ The EPO has allowed claims related to transformation of data (e.g. invention processed one image into many when compared to prior art involving processing of many images into one) as having an inventive step.¹²⁵ In making an inventive step assessment from a PHOSITA perspective, account may also need to be taken of the probabilistic and non-deterministic nature of AI systems, particularly when determining whether the skilled person would have had a reasonable expectation of achieving the claimed result.¹²⁶

172. Similarly, although improving the accuracy of a neural network through the pre-processing of training data constitutes a routine objective, a particular pre-processing modification should

¹¹⁹ Ibid.

¹²⁰ Ibid.

¹²¹ JPO Examination Handbook for Patent and Utility Model, Annex B, Chapter 1, Section 2.2.3.3 see, examples 1, 2, 4, 5 and Example 5] and Examination Handbook for Patent and Utility Model, Annex A, see Case 33, 34, 37, 38 and 40. Also, see MOIP, Republic of Korea, Examination practice guide by technology field [Chapter 6, Case 8.11.14] [Case 34(Claim 1)] [Case 37], [Case 38], [Case 40].

¹²² Ibid.

¹²³ Stephen How, A Hou, Stephen M., International Patent Office Guidance on Artificial Intelligence Inventions (2024) available at: <https://www.solveintelligence.com/blog/post/international-patent-office-guidance-on-artificial-intelligence-inventions>.

¹²⁴ JPO Examination Handbook for Patent and Utility Model, Annex A [Case 34(Claim 2)] and [Case 35] and MOIP, Korea Examination practice guide by technology field [Chapter 1, Case 4] [and [Chapter 6, Case 10.12.13].

¹²⁵ T 1286/09 (Image classifier/INTELLECTUAL VENTURES) of 11.06.2015 (EPO).

¹²⁶ Stephen How, A Hou, Stephen M., International Patent Office Guidance on Artificial Intelligence Inventions (2024) available at: <https://www.solveintelligence.com/blog/post/international-patent-office-guidance-on-artificial-intelligence-inventions>.

not automatically be regarded as obvious. Inventive step will arise where the modification is neither disclosed in the prior art nor part of common general technical knowledge and produces a significant and unexpected technical effect extending beyond routine optimization.¹²⁷ The assessment will therefore proceed by considering, in sequence, the nature of the modification, the relevant prior art, the common general knowledge of the skilled person, whether the specific training configuration would have been obvious and the technical effect obtained, while taking into account the tools and experimentation capabilities reasonably available to PHOSITA at the relevant date.

173. Finally, where deep-learning training and estimation using trained neural networks are well known, and the prior art addresses the same problem and performs the same function, the skilled person will be expected to apply that teaching. In such circumstances, a claimed invention whose contribution consists solely in using a known deep-learning method to obtain a predictable improvement in estimation will not involve an inventive step.¹²⁸

E. Sufficiency of Disclosure and Claims¹²⁹

174. 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.

175. 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).¹³⁰ 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.

176. 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.¹³¹ The assessment of sufficiency of disclosure of AI-inventions is thus a new challenge for which many countries have yet to develop case law on this point.¹³² 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.

¹²⁷ JPO Examination Handbook for Patent and Utility Model, Annex A [Case 36] and Examination practice guide by technology field [Chapter 6, Case 16]

¹²⁸ JPO, Examination Handbook for Patent and Utility Model, Annex A [Case 32] and MOIP Republic of Korea Examination practice guide by technology field [Chapter 1, Case 2.3.5]

¹²⁹ See also document SCP/34/5 (Further Study on the Sufficiency of Disclosure (PART I)), Part IV, AI-related inventions.

¹³⁰ 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.

¹³¹ See comments received from Chile and Germany in response to C.9199.

¹³² For example, see comments received from Germany in response to C.9199.

177. 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 algorithm 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.

178. 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.

179. 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. The 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.

180. 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).

181. 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).¹³³ 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 America¹³⁴ 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”.¹³⁵ 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.

182. In China, the Guidelines for Patent Examination were revised in 2025, clarifying that the relevant requirements for sufficiency of disclosure are satisfied in two scenarios involving artificial intelligence: “Where the invention relates to the construction or training of an AI model, it is generally necessary to clearly describe in the description the necessary modules, layers or connection structures of the model, as well as the specific steps and parameters essential for training. Where the invention relates to the application of an AI model or algorithm in a specific field or scenario, it is generally necessary to clearly describe in the description how the model or

¹³³ See document SCP/22/4.

¹³⁴ Section 112(a) of Title 35 of the United States Code. See document SCP/22/4.

¹³⁵ *Ibid.*

algorithm is integrated with the specific field or scenario, and how the input and output data of the algorithm or model are arranged to reflect their intrinsic relationship, so that a person skilled in the art can carry out the technical solution of the invention based on the description.” China further indicated that the revised Guidelines include two cases illustrating the applicable standards of the relevant provisions.¹³⁶

183. In France, the guidelines concerning the sufficiency of disclosure of patent applications in the field of artificial intelligence are currently being drafted and are expected to be available online no later than the end of 2026.¹³⁷

184. In India, the 2025 CRI Guidelines give specific disclosure guidance for AI/ML/DL inventions. They emphasize that AI/ML/DL claims must be supported by sufficient details enabling the person skilled in the art to reproduce the technical solution without undue experimentation. Depending on the claimed contribution, applicants should disclose the neural-network architecture and structural details, training data characteristics and labelling methodology, preprocessing steps, training details such as loss functions, optimizer settings and epochs, validation metrics and real-world test results, input-output correlations and, where relevant, the data pipeline or algorithmic architecture. Where inventiveness depends on particular traits of a training dataset, the Guidelines state that those traits should be disclosed unless they can be identified without undue experimentation, although the specific dataset itself need not ordinarily be disclosed.¹³⁸

185. 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 JPO, contains several concrete examples.¹³⁹ In the Republic of Korea, the description should thoroughly explain the specific means, technical problems addressed, and solutions provided by the AI-related invention, ensuring that a skilled person can easily understand and reproduce it. An invention is considered inadequately disclosed if it does not specifically provide for a correlation between input data and output data of a trained model as a certain means for implementing an AI-related invention. Correlation between input data and output data of a trained model means: (i) training data is specified; (ii) a correlation for solving a technical problem of the claimed invention between characteristics exists; (iii) a learning model to train by using training data or a training method is specifically described; and (iv) a trained model is generated for solving a technical problem of the claimed invention by means of such training data and training method. However, if a person skilled in the art presumes or understands such a correlation on the basis of the model(s) described in the description of the invention, as taking common general knowledge at the time of the filing of the application into account, the enablement requirement is deemed to be met. Additionally, merely using functional block diagrams or general flow charts without providing specific implementation details is generally insufficient. The application must describe how the hardware or software implements the invention's functions in a manner that is understandable and reproducible by someone skilled in the art. This detailed disclosure ensures that the invention can be reproduced and utilized effectively, satisfying the enablement requirement, and supporting the innovation's practical applicability.¹⁴⁰

186. In March 2024, the Guidelines for Examination at the EPO were updated and included substantive changes relating to the sufficiency of disclosure requirement for AI-related inventions.¹⁴¹ One of the significant revisions in the updated Guidelines is that inclusion of detailed information about the training data for a claimed AI model in a patent application may

¹³⁶ See comments received from China in response to C.9304.

¹³⁷ See comments received from France in response to C.9304.

¹³⁸ Office of Controller General of Patents, Designs and Trade Marks, Guidelines for Examination of Computer Related Inventions (CRIs) - 2025, section 5.2.2

¹³⁹ Annex A of the Examination Handbook for Patent and Utility Model, Examples 46 to 55, JPO.

¹⁴⁰ Examination Guide in the Artificial Intelligence Field, Korean Intellectual Property Office.

¹⁴¹ Guidelines for Examination in the European Patent Office (EPO), Part G, Chapter II, 3.3.1.

be necessary to meet the sufficiency of disclosure requirement. It aims to ensure sufficient disclosure and serves as evidence of technical effect. Specifically, the updated guidelines state that a patent disclosure is insufficient when *“the mathematical methods and the training datasets are disclosed in insufficient detail to reproduce the technical effect over the whole range claimed. Such a lack of detail may result in a disclosure that is more like an invitation to a research programme.”*¹⁴² The new guidelines also state that *“the technical effect that a machine learning algorithm achieves may be readily apparent or established by explanations, mathematical proof, experimental data or the like [...] If the technical effect is dependent on particular characteristics of the training dataset used, those characteristics that are required to reproduce the technical effect must be disclosed unless the skilled person can determine them without undue burden using common general knowledge. However, in general, there is no need to disclose the specific training dataset itself.”*¹⁴³ While the Guidelines currently state that full disclosure of the underlying training dataset is not generally required, instances where it is necessary have not yet been fully defined.

187. There have been two decisions at the EPO pertaining to sufficiency of disclosure and AI. In T 0161/18, the EPO Board of Appeal determined that a patent application pertaining to an artificial neural network under the case failed to meet the sufficiency of disclosure requirement. It noted that the application did not disclose which input data are suitable for training the artificial neural network or at least one dataset suitable for solving the technical problem.¹⁴⁴ Similarly, in T 1191/19, a patent application was found not to be in compliance with the sufficiency of disclosure requirement. In reaching this conclusion, the Board noted that the application did not provide any example of training data or even the amount of training data required to enable the invention to provide a meaningful prediction.¹⁴⁵ These cases stress the importance for applicants to ensure that their disclosures are enabling, potentially including detailed training data and other pertinent information to comply with the EPO's requirements.

188. [The Intellectual Property Office of the Philippines \(IPOPHL\) has also identified unreliable sources or made-up sources as among common problems in AI-assisted inventions. Where such sources are identified, examiners, while not required to check each cited reference in the specification, should raise the matter in an Office Action for removal or correction of such resources in order to maintain the integrity and quality of inventions granted in the Philippines.](#)¹⁴⁶

189. [The above survey of different jurisdictions shows that as a new technological area, the application of the sufficiency of disclosure requirement to AI-related inventions is still evolving and has not yet been clearly set. Some jurisdictions are in the process of developing reports and rules to address these unique challenges.](#)¹⁴⁷ Continued development of legislation, guidelines, and court rulings in this area will help to establish more precise standards and practices. In a general sense, the black box phenomenon inherent to many AI algorithms may not pose a problem regarding the sufficient disclosure of the invention as long as sufficient details are given about which AI algorithm to use and how to train it.¹⁴⁸

190. [It is worth noting that, in many AI models and AI-based inventions, the essence of the claimed invention lies in a technical relationship linking a particular category of input data, for instance medical images, sensor readings, or textual data, to a specific output or result, for instance a diagnosis, prediction, or classification.](#)¹⁴⁹ Where the correlation would not be

¹⁴² Guidelines for Examination in the European Patent Office (EPO), Part F, Chapter III, Insufficient Disclosure.

¹⁴³ Guidelines for Examination in the European Patent Office (EPO), Part G, Chapter II, 3.3.1.

¹⁴⁴ EPO Board of Appeal Decision T0161/18.

¹⁴⁵ EPO Board of Appeal Decision T1191/19.

¹⁴⁶ [See comments received from the Philippines in response to C.9304.](#)

¹⁴⁷ [See comments received from Brazil in response to C.9199.](#)

¹⁴⁸ [See comments received from Germany and Portugal in response to C.9199.](#)

¹⁴⁹ [See SCP/34/5, pp. 27, 32–33.](#)

understood from common general knowledge, the description should provide a specific basis enabling the skilled person to understand it.¹⁵⁰ In other words, the description must provide sufficient grounds for a reader to conclude that the correlation being claimed genuinely exists and is technically well-founded, rather than being merely asserted.¹⁵¹

191. Simply stating that a correlation exists, without offering any supporting reasoning, is not sufficient to satisfy this requirement. The description needs to substantiate the claimed correlation, and this may be achieved by explaining the underlying theoretical basis for why the correlation holds true, or by presenting empirical data, statistical results, experimental findings, or other forms of technical justification embedded within the specification. This requirement concerns the adequacy and quality of the disclosure itself, and is a separate consideration from questions relating to the scope of the claims.¹⁵²

192. A closely related issue concerns whether the specification includes actual empirical performance data obtained from AI models that have genuinely been built and tested, such as accuracy rates, precision and recall figures, or comparisons against recognized benchmarks, in order to substantiate the claimed correlation. The presence of such data helps distinguish an AI system that has only been conceptually described from one that has been shown to actually perform as claimed. This, in turn, helps prevent speculative or purely prophetic AI claims in which the underlying model has never actually been implemented or tested.¹⁵³

193. The pre-processing of input data, which may involve data cleaning, normalization, feature extraction, transformation, and augmentation, forms a technically important part of many AI-based inventions. If these pre-processing steps are described only in vague or generic terms, the specification risks failing the enablement requirement, since a person skilled in the art would not be able to reproduce the invention on that basis. In the same way, the specification needs to describe, in sufficiently concrete terms, the learning model being used, including details such as the type of neural network, the number of layers, the activation functions, and the relevant hyper-parameters, as well as the training methodology, including the learning algorithm, the loss function, and the optimization strategy employed.¹⁵⁴ A generic reference may be insufficient where it does not enable the skilled person to reproduce a model specialized to the claimed technical problem.¹⁵⁵

194. A number of jurisdictions have clarified that, depending on the nature of the invention concerned, disclosure requirements for AI inventions should identify the essential model structures, the training steps involved, the relevant parameters, and the relationship connecting the input data to the output data, so that a person skilled in the art is able to carry out the technical solution without needing to undertake undue experimentation.¹⁵⁶ Functional diagrams or general flow charts, on their own, have been regarded as insufficient for this purpose. Case law from at least one jurisdiction has confirmed that a patent application relating to an artificial neural network did not meet the sufficiency requirement, because it failed to disclose which types of input data would be suitable for training the network, or to identify at least one dataset capable of solving the technical problem at hand. In a similar vein, an application that provided no example of training data, and gave no indication of the minimum quantity of training data

¹⁵⁰ Ibid.

¹⁵¹ JPO Examination Handbook for Patent and Utility Model, Annex A [Case 46] [Case 47] and [Case 48] and MOIP, Republic of Korea Examination practice guide by technology field [Chapter 1, Case 1] [Chapter 6, Case 2.6.7].

¹⁵² JPO Examination Handbook for Patent and Utility Model, Annex A [Case 49].

¹⁵³ JPO Examination Handbook for Patent and Utility Model, Annex A [Case 50].

¹⁵⁴ MOIP, Republic of Korea Examination practice guide by technology field [Chapter 1, Case 1] [Chapter 6, Case 4].

¹⁵⁵ SCP 34/5 p.27.

¹⁵⁶ SCP/34/5, p. 34 and Mateo Aboy et. al., The sufficiency of disclosure of AI inventions, *Journal of Intellectual Property Law & Practice*, Volume 19, Issue 11, November 2024, Pages 836 and 838.

needed to produce meaningful predictions, was found not to satisfy the enablement requirement.

195. Training data raises a distinct set of disclosure-related challenges. Where the technical performance of the invention depends on a particular dataset, questions naturally arise as to how much information about that dataset needs to be disclosed within the application, and whether third parties should be given access to the dataset itself so that they may verify that the claimed invention actually works as described. Requiring the full disclosure of training data poses considerable practical difficulties, given that such data is often proprietary and treated as commercially confidential. A further practical complication may arise in connection with privacy and data protection law, particularly in situations where the relevant data is available to the inventor or applicant only under a limited license for training purposes, but not for the purpose of being disclosed in connection with the claimed AI model that has been trained using that specific dataset. At present, these questions remain largely unresolved within existing patent office guidance.¹⁵⁷

196. Although full disclosure of training data is not presently required, a finding of insufficiency may nonetheless arise where an application fails to disclose even a single suitable training dataset, or fails to indicate the minimum amount of training data necessary for the invention to function as claimed. Where the technical effect produced by the invention depends on particular characteristics of the dataset used, those characteristics must be disclosed, unless a skilled person would be able to determine them independently without undue effort. That said, disclosure of the dataset itself in its entirety is not ordinarily considered necessary.

197. One possible mechanism that has been identified for validating claimed AI performance involves the partial or full disclosure of training data through deposit with designated institutions, or through verification by third parties following grant of the patent. This approach, however, gives rise to further questions concerning cross-border access to, and validation of, deposited data across different jurisdictions. While comparable issues relating to the deposit of micro-organisms have already been resolved through international cooperation under existing treaty frameworks, the situation concerning training data is somewhat different, since the subject matter involved is intangible and can, in principle, be made universally accessible, meaning that physical deposit is not strictly necessary. The central challenge in this context therefore lies not in physical accessibility, but rather in validating the training data for the purpose of assessing the enablement of the claimed invention.

198. While not all jurisdictions have addressed the issue of hallucinated prior-art submitted by applicant and implications for sufficiency of disclosure, at least some patent offices have raised concerns regarding presence of unreliable or fabricated sources, including so-called AI-generated hallucinations appearing within specifications, as a recurring issue in patent applications that have been prepared with the assistance of AI tools.¹⁵⁸ Where such issues come to light, examiners are expected to raise them during examination, through appropriate correspondence, and to request that the applicant correct or remove the problematic content, so as to preserve the overall quality of the patent and the integrity of the evidence contained within it.¹⁵⁹ In the United States, for example, the USPTO has indicated that existing professional obligations apply to AI-assisted filings, while not imposing a general obligation to disclose the use of a particular AI tool unless specifically requested.¹⁶⁰

¹⁵⁷ WIPO AI and IP Policy Toolkit, p.15.

¹⁵⁸ Hughes (2025), Use of AI in the patent industry: The spectre of hallucination, available at: <https://ipkitten.blogspot.com/2025/10/use-of-ai-in-patent-industry-spectre-of.html>

¹⁵⁹ See comments received from the Philippines in response to C.9304.

¹⁶⁰ Guidance on Use of Artificial Intelligence-Based Tools in Practice Before the United States Patent and Trademark Office [Docket No. PTO-P-2024-0013].

199. Considering the ambiguity of the requirements to disclose model weights or training data sets and that jurisdictions are still in the process of providing guidance on how applicants should incorporate them, academic commentators have suggested introducing an international deposit system similar to framework of the Budapest Treaty for Deposit of Microorganisms or the local systems employed in depositing scientific journals.¹⁶¹

200. Taken together, it is evident that sufficiency standards applicable to AI inventions remain in a state of ongoing development. That said, the black-box nature that characterizes many AI systems need not, in itself, prevent an application from meeting the sufficiency requirement, provided that the application adequately explains the underlying algorithm and the training process involved. It is anticipated that further developments, whether through legislation, examination guidelines, or case law, will continue to progressively refine the standards applicable in this area.¹⁶²

F. Industrial Applicability

201. 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.^{163 164}

202. A key aspect of industrial applicability is the reproducibility and consistency of the claimed invention. For AI systems, especially those employing deep learning and neural networks, demonstrating consistent performance can be challenging due to their inherent complexity and the often opaque nature of their decision-making processes.

203. As an example, an AI system used for predictive maintenance in industrial machinery may be considered to be industrially applicable only if it can consistently predict failures accurately. This may require detailed documentation of the AI model, including its training data, algorithms, and validation results, to ensure that it can be reliably implemented and produce consistent results in an industrial setting.

G. Inventorship and Ownership

204. As AI becomes increasingly integrated into the invention process, important questions arise regarding the ownership of AI created inventions and the criteria that make a natural person the inventor when assisted by AI.¹⁶⁵ In 2023, document SCP/35/7 was created to address the evolving topic of inventorship and ownership in the context of AI-related inventions. That document presents an in-depth overview of the interaction between humans and AI in the invention process and offers a thorough analysis of current legal frameworks and policy questions surrounding AI inventorship.¹⁶⁶ Therefore, this Section only provides an overview of the inventorship/ownership issues regarding the AI-related inventions.

¹⁶¹ Mateo Aboy et. al., *The sufficiency of disclosure of AI inventions*, *Journal of Intellectual Property Law & Practice*, Volume 19, Issue 11, November 2024, p. 839.

¹⁶² SCP/35, p. 34, and also see Mateo Aboy et. al., *The sufficiency of disclosure of AI inventions*, *Journal of Intellectual Property Law & Practice*, Volume 19, Issue 11, November 2024, Pages 838-840.

¹⁶³ 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.

¹⁶⁵ See comments received from Chile in response to C.9199.

¹⁶⁶ At the thirty-fifth session of the SCP, the Committee agreed that the Secretariat would update Sections V and VI of document SCP/35/7, and would submit it to the thirty-seventh session of the SCP.

Inventorship and Ownership under Patent Law

205. 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.¹⁶⁷ If more than one inventor jointly create an invention, they are joint inventors.

206. 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, inaccurate indication of inventors may have serious legal consequences. Incorrect designation of inventors can lead to various legal consequences, with different jurisdictions offering remedies for such inaccuracies.

207. 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).¹⁶⁸ 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.¹⁶⁹

208. 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?

209. In India, the 2025 CRI Guidelines expressly distinguish AI-generated inventions from AI-assisted inventions. They state that AI-generated inventions - those created autonomously or with very limited human intervention - are not patentable because AI cannot be termed a “person” claiming to be the true and first inventor under section 6 of the Patents Act, 1970. By contrast, AI-assisted inventions are not categorically non-patentable under section 3(k) provided they meet patentability criteria and demonstrate technical effect through tangible applications.¹⁷⁰

¹⁶⁷ Guide to the Application of the Paris Convention for the Protection of Industrial Property, G. H. C. Bodenhausen (WIPO Publication No. 611).

¹⁶⁸ 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. In Tanzania for example, the term inventor is recognized as only pertaining to humans and as such inventions generated solely by artificial intelligence are therefore not recognized under the current legal framework - see comments received from Tanzania in response to C.9304.

¹⁶⁹ See comments received from the Philippines in response to C.9304, noting that, under Philippine law, an AI system cannot be recognized as a natural or juridical person and lacks the civil personality and juridical capacity needed to hold, transmit or enforce IP rights; accordingly, AI cannot be named as an inventor, maker, designer or applicant in an IPOPHL application, including in cases of joint inventorship or joint designship, where the relevant persons must likewise be persons recognized under Philippine law; France indicated that, to date, there is no French case law of the Court of Appeal or the Court of Cassation concerning the possibility of an AI system being designated as inventor, or INPI, the inventor must be a natural person, designated when the application is filed – see comments received from France in response to C.9304.

¹⁷⁰ Office of Controller General of Patents, Designs and Trade Marks, Guidelines for Examination of Computer Related Inventions (CRIs) - 2025, section 5.2.1; Indian Patent Office, Application No. 202017019068, Order under Section 15 (Apr. 15, 2026).

210. At the same time, the Philippines notes that AI may be used as a tool to aid inventors, makers or designers, subject to ethical considerations; where an individual uses prompts or instructions for an AI system to assist in generating an invention or design, that individual may still be considered the inventor and/or applicant, because inventions and designs remain linked under current Philippine law to human creativity and intellectual contribution, while the AI system itself must be excluded from being named in the application.¹⁷¹

211. In November 2025, the USPTO published Revised Inventorship Guidance for AI-Assisted Inventions, rescinding its February 2024 guidance on this topic in its entirety. The new guidance underscores longstanding precedent that the same legal standard for determining inventorship applies to all inventions, regardless of whether AI systems were used in the inventive process, and that no new, separate or modified standard is created for or applies to AI-assisted inventions. In the United States, only natural persons can be properly named as inventors on patent applications: 35 U.S.C. § 100(f) provides that an “inventor” means the individual who invented or discovered the subject matter of the invention, and the “invention” in the Patent Act refers to the inventor’s conception. The USPTO presumes that the inventors named on the application data sheet or oath/declaration are the actual human-being inventors, and AI systems, including generative AI and other computational models, are tools used by human inventors that do not qualify for, or elevate such assistance to, inventor status.¹⁷²

212. 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.

213. Under current approaches, AI-assisted inventions are generally considered within existing inventorship frameworks, although the applicable standards and thresholds vary among jurisdictions. The determination of inventorship therefore requires an assessment, under the relevant law, of whether one or more natural persons made a legally sufficient contribution to the claimed invention. As AI tools assume a greater role in the inventive process, identifying the nature and extent of such contributions, including indirect but invention-specific contributions to the development, training or operation of an AI system, will become increasingly complex and may warrant further clarification.

214. The legal consequences of an inaccurate designation of inventors differ across jurisdictions and may depend on whether the error concerns a mere formality, a correctable mistake, or a defect in entitlement. Where an applicant fails to identify the inventor, patent offices may invite the applicant to correct the omission within a prescribed period, failing which the application may be rejected. Many patent laws also permit correction of an incorrectly designated inventor, either during prosecution or after grant. In many countries, an error in the indication of the inventor is not, by itself, a ground for invalidation of the patent. However, where the inaccurate designation indicates that the applicant or patentee did not derive title from the true inventor, or where the invention was obtained through usurpation, the consequences may be more significant, including invalidation or revocation of the patent, unenforceability in cases involving intent to deceive, transfer or assignment of the application or patent to the person entitled, rectification of the patent register, damages or other remedies available under the applicable law.¹⁷³

¹⁷¹ See comments received from the Philippines in response to C.9304.

¹⁷² See comments received from the United States of America in response to C.9304.

¹⁷³ WIPO, Artificial Intelligence (AI) and Inventorship (Update of SCP/35/7), document SCP/37/5.

215. 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. Therefore, the concepts of inventorship (i.e., a person who invented the invention) and ownership (i.e., a person who has the right to file a patent application or obtain a patent) are separate but related. 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.¹⁷⁴ Therefore, the inventorship/ownership issues may be part of the essential policy questions for the designing of a patent system.

216. El Salvador has indicated that its Law for the Promotion of Artificial Intelligence and Technologies (Legislative Decree No. 234 of February 26, 2025) contain specific provisions on ownership and protection of intellectual property resulting from AI research, development or innovation carried out within the territory of El Salvador. The submission states that all such intellectual property, including patents, copyrights, trademarks and industrial secrets, belongs exclusively to the developing entity or person, while ownership rights in collaborative projects involving several entities are determined by prior contractual agreements between the parties. The law also provides that proprietary algorithms, datasets and AI systems developed in El Salvador are protected by law and may not be disclosed or used without the explicit consent of the owner, and that the State shall facilitate access to international intellectual property protections applicable to AI or similar technologies, including by assisting developers in registering innovations under global treaties such as the PCT. The submission does not address the separate question of whether an AI system may be named as an inventor.¹⁷⁵

AI Inventorship and the DABUS Case

217. As AI systems are increasingly used in research and development, patent systems have had to consider how existing rules on inventorship and entitlement apply where an invention is said to have been generated by, or with the assistance of, an AI system. This question is distinct from the patentability of AI-related inventions. It concerns whether an AI system can be named as an inventor and how the right to apply for or obtain a patent should be established when the asserted inventive contribution is attributed to AI.

218. The DABUS applications have served as the principal test case for these issues. They were filed on the basis that the AI system DABUS had generated the relevant inventions autonomously and should therefore be identified as the inventor. The resulting proceedings have focused on a threshold issue: whether, under existing patent laws, the concept of an “inventor” can extend to a non-human machine. The overall position emerging from the DABUS proceedings is that there has not been a surviving final substantive determination accepting that an AI system may be named as an inventor under current patent law. This conclusion does not necessarily preclude patent protection for inventions developed with the assistance of AI, but it underscores the continuing importance of identifying the relevant human contribution and establishing entitlement in accordance with applicable law.

219. For the purposes of this study, the DABUS matter is therefore best understood as an illustration of how existing inventorship, designation-of-inventor and entitlement requirements operate where no human inventor is identified. It does not resolve all legal and policy questions concerning AI-assisted invention, including the level or nature of human contribution that may be required where AI is used as part of the inventive process. A more detailed account of the decisions of IP Offices and courts concerning the DABUS matter is provided in document SCP/38/5.

¹⁷⁴ 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.

¹⁷⁵ See comments received from El Salvador in response to C.9304.

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

220. 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, or to resolve disputes by judiciaries etc.

A. Tools for the IP Authorities

221. 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 Offices¹⁷⁶ is an on-line portal on which such use of the AI technology is listed 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) patent examination management; (iii) helpdesk services; (iv) image search; (v) machine translation; (vi) patent classification; (vii) patent prior art search; (viii) trademark classification; (ix) copyright registration; and (x) data analysis.

222. 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.¹⁷⁷ 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.¹⁷⁸ As a follow-up to the Meeting, WIPO established a dedicated web page on AI¹⁷⁹ 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.¹⁸⁰

223. More recently, in April 2026, WIPO launched the WIPO ICT Leadership Dialogue (WILD), which provides a forum for ICT leaders and senior management officials from IP offices to exchange practical experience and explore cooperation on digital solutions and emerging technologies for IP business and service delivery. WILD focuses on areas such as artificial intelligence, cybersecurity, and data management, and operates as a voluntary, non-binding platform for practical cooperation among IP offices. Following its first session, WILD established, among other things, an AI Tool Collaboration Group, bringing together AI experts and product owners from IP offices to collaborate on the development, evaluation and implementation of AI-powered tools for IP administration, coordinate pilot projects and share AI solutions.

224. For patent examination, the public WILD materials show that IP offices are exchanging experience on AI-assisted patent search, patent classification, examiner and formalities support, prior-art identification and related workflow tools. The second session of WILD included, for example, discussion of supporting examiners and formalities officers with AI, presentations on AI-powered solutions for patent examination, and exchanges on AI tools currently in use by IP

¹⁷⁶ https://www.wipo.int/about-ip/en/artificial_intelligence/search.jsp.

¹⁷⁷ Documents and presentations of the meeting are available at: https://www.wipo.int/meetings/en/details.jsp?meeting_id=46586.

¹⁷⁸ Document WIPO/IP/ITAI/GE/18/5 (Summary by the Facilitator).

¹⁷⁹ https://www.wipo.int/about-ip/en/artificial_intelligence/.

¹⁸⁰ Document CWS/6/3.

offices and opportunities for collaboration. WILD therefore complements the WIPO Index of AI Initiatives in IP Offices and the SCP sharing sessions by providing a practical cooperation channel through which offices can compare examination-related AI use cases, lessons learned and possibilities for joint development, while leaving adoption of any particular tool to each office.¹⁸¹

225. In addition, since February 2019, the WIPO Conversation on Intellectual Property and Artificial Intelligence (later renamed to the WIPO Conversation on IP and Frontier Technologies) provides an open, inclusive forum to engage with, and facilitate discussion and knowledge building among, the widest possible set of stakeholders on the impact on IP of frontier technologies, including artificial intelligence (AI).¹⁸² Recent sessions have continued to address AI-related issues, including AI inventions, the use of AI by IP Offices and questions raised by generative AI.

226. 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.¹⁸³

227. The International Bureau of WIPO has also used AI for its work in order to enhance functions and processes at the Organization. It also develops and provides a range of AI-powered services and tools to assist users and stakeholders. WIPO currently uses and provides AI-powered tools and services in a number of areas, including speech-to-text, image similarity search within the Global Brand Database, automatic patent classification through IPCCAT, Nice classification support through the Global Goods & Services Explorer, the Vienna Classification Assistant, and machine translation through WIPO Translate.¹⁸⁴

228. During the sharing sessions held at thirty-first, thirty-third and thirty-fifth and thirty-sixth sessions of the SCP, IP offices and the International Bureau of WIPO presented their use of AI in patent administration and patent examination.¹⁸⁵ The thirty-eighth session of the SCP also includes a sharing session on the use of various tools, including AI, for effective patent examination procedures and related challenges faced by IP offices. Member States will also share their experiences on new avenues for cooperation in this area, and the Secretariat will share its initiatives on such tools.

B. Tools for Applicants, Third Parties and IP Professionals

229. AI tools are also increasingly used by applicants, patent owners, third parties and IP professionals in activities connected with patent filing, prosecution, portfolio management, enforcement and validity assessment. These uses are not limited to tools developed by IP Offices. They include publicly available search, classification and translation tools made

¹⁸¹ WIPO, WIPO ICT Leadership Dialogue (WILD) webpage; WIPO ICT Leadership Dialogue (WILD), First Session, Geneva, April 14 to 16, 2025; WIPO ICT Leadership Dialogue (WILD), Second Session, Geneva, April 14 to 16, 2026; WIPO, WILD/2/1, Program, April 21, 2026.

¹⁸² https://www.wipo.int/about-ip/en/frontier_technologies/frontier_conversation.html.

¹⁸³ WIPO Index of AI Initiatives in IP Offices. See also documents SCP/32/4, and 4 Corr. as well as SCP/34/4; see, for example, comments received from Portugal in response to C.9304.

¹⁸⁴ For detailed information, please visit the WIPO website at: <https://www.wipo.int/en/web/ai-tools-services/>.

¹⁸⁵ Presentations are made available on the respective SCP meetings webpages. Documents SCP/32/4 and 4 Corr. as well as SCP/34/4 contain reports of the sharing sessions held during SCP/31 and SCP/33, respectively. A similar sharing session will also be held during the thirty-sixth session of the SCP.

available by WIPO and IP Offices, as well as commercial and general-purpose AI systems used by applicants, practitioners and other users of the patent system.

230. From a functional perspective, such tools may assist with document automation, process automation, and AI-enabled analysis. For example, they may support prior art searching, classification, translation, proofreading, consistency checks, monitoring of legal status, patent landscaping, portfolio analysis and the organization or summarization of technical information. The increasing availability of generative AI and semantic search tools may further assist users in formulating search strategies, reviewing large volumes of patent and non-patent literature, comparing technical features and preparing first drafts or working documents for human review.

231. At the same time, AI tools should be understood as aids to human professional judgment, rather than substitutes for it. Outputs generated by AI systems may be incomplete, inaccurate or unsupported, and users may need to consider confidentiality, data security, professional responsibility and any applicable procedural requirements when relying on such tools in patent-related matters.

[Annex follows]

Webpages and publications hosted by WIPO and IP offices dedicated to AI

Artificial Intelligence and Intellectual Property

<https://www.wipo.int/en/web/frontier-technologies/artificial-intelligence/index>

Intellectual Property for Frontier Technologies

<https://www.wipo.int/en/web/frontier-technologies>

AI Tools and Services

<https://www.wipo.int/en/web/ai-tools-services>

Index of AI initiatives in IP offices

<https://www.wipo.int/en/web/ai-tools-services/ipos-initiatives>

WIPO Conversation on Intellectual Property and Artificial Intelligence

<https://www.wipo.int/en/web/frontier-technologies/artificial-intelligence/conversation>

Artificial Intelligence Infrastructure Interchange (AII)

<https://www.wipo.int/en/web/frontier-technologies/artificial-intelligence/ai-infrastructure-interchange>

Getting the Innovation Ecosystem Ready for AI - An IP Policy Toolkit (2024)

<https://www.wipo.int/edocs/pubdocs/en/wipo-pub-2003-en-getting-the-innovation-ecosystem-ready-for-ai.pdf>

WIPO Technology Trends 2019 - Artificial Intelligence

https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf

Patent Landscape Report - Generative Artificial Intelligence (GenAI) (2024)

<https://www.wipo.int/web-publications/patent-landscape-report-generative-artificial-intelligence-genai/index.html>

Generative AI: Navigating Intellectual Property (2024)

<https://www.wipo.int/edocs/pubdocs/en/wipo-pub-rn2024-8-en-generative-ai-navigating-intellectual-property.pdf>

Artificial Intelligence (AI) and Inventorship (Update of SCP/35/7)

https://www.wipo.int/edocs/mdocs/scp/en/scp_37/scp_37_5.pdf

Learning Machines – An Introduction to AI and IP for Small and Medium-sized Enterprises

<https://wipo-dpp.gateb.com/tempDownload/120DC6C1B13095019A0B0A57974086D223573983>

Australia

Generative AI and the IP system

<https://www.ipaustralia.gov.au/temp/Generative-AI-and-the-IP-System.html>

<https://www.ipaustralia.gov.au/tools-and-research/professional-resources/data-research-and-reports/publications-and-reports/generative-ai-and-the-ip-system>

Patent Manual, 5.4.3 Artificial Intelligence - Inventorship and Entitlement

<https://manuals.ipaustralia.gov.au/patent/5.4.3-inventions-produced-by-artificial-intelligence>

Brazil

Final Report IP BRICS Offices, IP BRICS-INPI

<http://www.ipbrics.net/secondpage/project/Patent%20Processes%20and%20Procedure%20-%20AI%20Study%20Report.pdf> ‘

Canada

CIPO, Artificial intelligence and intellectual property

<https://ised-isde.canada.ca/site/canadian-intellectual-property-office/en/artificial-intelligence-and-ip>

China

Examination of Patent Applications for Inventions Involving Computer Programs

https://www.cnipa.gov.cn/art/2024/1/18/art_2199_189877.html

Decision on amending the Patent Examination Guidelines (Order No. 84

https://www.cnipa.gov.cn/art/2025/11/13/art_74_202560.html?utm_source

France

INPI, Le role du brevet dans les inventions utilisant l'intelligence artificielle

<https://www.inpi.fr/actualites/dossiers-thematiques/role-du-brevet-dans-inventions-utilisant-lintelligence-artificielle>

Japan

“AI-related Inventions” webpage

<https://www.jpo.go.jp/e/system/patent/gaiyo/ai/index.html>

Recent Trends in AI-related Inventions (2026)

https://www.jpo.go.jp/e/system/patent/gaiyo/ai/ai_shutsugan_chosa.html

JPO - CNIPA Comparative Study on AI-Related Inventions (2023)

https://www.jpo.go.jp/e/news/kokusai/cn/document/ai_report_2023_e/cn_ai_report_en.pdf

Comparative Study on Computer Implemented Invention/Software Related Inventions between JPO and EPO (2021)

https://www.jpo.go.jp/e/system/laws/rule/guideline/patent/document/ai_jirei_e/01_en.pdf

Patent Examination Case Examples pertinent to AI-related technologies
https://www.jpo.go.jp/e/system/laws/rule/guideline/patent/ai_jirei_e.html

Philippines

IPOP HL, Examination Guidelines of Inventions, Utility Model, and Industrial Design Applications
Relating to Artificial Intelligence
<https://www.ipophil.gov.ph/services/patent-examination-guidelines/>

Republic of Korea

https://www.kipo.go.kr/upload/en/download/Examination_Guide.pdf KIPO, AI and patentability
https://www.kipo.go.kr/en/HtmlApp?c=92012&catmenu=ek03_08_05

KIPO, Examination Practice Guidelines by Technology Field - Chapter I, Artificial Intelligence
Field (2026)
https://www.kipo.go.kr/upload/en/download/EXAMINATION_PRACTICE_GUIDELINES_BY_TECHNOLOGY_FIELD.pdf

United States of America

“Artificial Intelligence” webpage
<https://www.uspto.gov/initiatives/artificial-intelligence>

AI-related resources
<https://www.uspto.gov/initiatives/artificial-intelligence/artificial-intelligence-resources>

European Patent Office

“Artificial intelligence” webpage
<https://www.epo.org/en/news-events/in-focus/ict/artificial-intelligence>

Guidelines for Examination in the EPO, G-II, 3.3.1 Artificial intelligence and machine learning
(2026)
https://www.epo.org/en/legal/guidelines-epc/2026/g_ii_3_3_1.html

[End of Annex and of document]