8-20-2024

The Sufficiency of Disclosure of AI Inventions

Mateo Aboy

Aparajita Lath

Timo Minssen

Kathleen Liddell

OXFORD

Journal of Intellectual Property Law & Practice, 2024, Vol. 00, No. 00

DOI: https://doi.org/10.1093/jiplp/jpae063

Article

The sufficiency of disclosure of AI inventions

Mateo Aboy[*], Aparajita Lath[†], Timo Minssen[‡] and Kathleen Liddell[§]

[*]Director of Research in Technology & Law, Biomedical Innovation, AI & Law, LML & CIPIL, Faculty of Law, University of Cambridge, Cambridge, UK. Email: ma608@cam.ac.uk.

[†]Assistant Professor of Law, Intellectual Property, National Law School of India University, Bangalore, India. Email: aparajita.lath@nls.ac.in.

[‡]Director of CeBIL & Professor, Faculty of Law, University of Copenhagen, Copenhagen, Denmark. Email: Timo.Minssen@jur.ku.dk.

[§]Professor of Intellectual Property and Medical Law & Director of the Centre for Law, Medicine and Life Sciences, Faculty of Law, University of Cambridge, Cambridge, UK. Email: k.liddell@law.cam.ac.uk.

- Mateo Aboy is Director of Research in Technology and Law (Biomedical Innovation, AI and Law) in the Faculty of Law at the University of Cambridge.
- Aparajita Lath is an Assistant Professor of Law, Intellectual Property at the National Law School of India University.
- Timo Minssen is the Director of the Centre for Advanced Studies in Bioscience Innovation Law (CeBIL) and Professor of Law in the Faculty of Law at University of Copenhagen.
- Kathleen Liddell is Professor of Intellectual Property and Medical Law and Director of the Centre for Law, Medicine and Life Sciences (LML) in the Faculty of Law at the University of Cambridge.

Abstract

- The complex and data-driven nature of artificial intelligence (AI) raises questions for the sufficient disclosure of patent applications in this field. What are the European patent disclosure requirements for AI inventions?
- One challenge is that, prior to training, AI systems can be considered generic models. But after training, they transform into specialized AI systems to solve a particular problem. This transformation requires training data, making it an integral part of the AI system's definition. But to what extent is the disclosure of the training data or training process necessary for patent disclosure?
- The Boards of Appeal of the European Patent Office (EPO) first dealt with this challenge in case T 0161/18, which involved a medical AI invention to calculate cardiac output. It held that the specialized artificial neural network (ANN) in the patent could not be carried out by a person skilled in the art due to insufficient disclosure of input data suitable for the training of the ANN or at least one data set suitable for solving the technical problem. Furthermore, without specialization, the invention lacked an inventive step.
- But, is it always necessary to disclose the input data or at least one data set suitable for solving the technical problem? Are there alternative ways for applicants to satisfy the disclosure requirements for AI inventions? And what evidence is there that patent applicants are disclosing specific details of the AI/machine learning (ML) training or specific AI/ML model architecture?
- In this article, we analyse case T 0161/18 and subsequent sufficiency of disclosure decisions (T 1539/20; T 0606/21; T 1526/20; T 1191/19) and consider these foundational questions for applicants drafting patent applications with claims directed to AI inventions. We also analyse the EPO's examination guidelines on sufficiency of disclosure for AI inventions, which were updated in early March 2024.

I. Introduction

Given the complex and data-driven nature of artificial intelligence (AI) technologies, ensuring adequate disclosure is particularly challenging. Part of the difficulty stems from intrinsic features of AI systems, which can solve problems in many different fields (meaning the mere idea of using AI is often obvious) and require training data and training processes to achieve such applications (raising questions about the necessity of disclosing such data and processes). It is precisely the training data and training process that transform a generic AI model into a particular AI system to solve a specific problem (ie a specialized AI).

The EPO Boards of Appeal's (BoA) decision in T 0161/18 (2020) addressed the sufficiency of disclosure for AI inventions, specifically AI inventions using artificial neural networks (ANNs). Several EPO Boards of Appeal cases have followed. Each deal with Article 83 of the European Patent Convention (EPC), which states: 'the European patent application shall disclose the invention in a manner sufficiently clear and complete for it to be carried out by a person skilled in the art.' For this to happen, the disclosure of the invention in the application must enable the person skilled in the art to reproduce the technical teaching in the claimed invention. A key question for AI-related inventions is therefore the following:

to what extent is disclosure of the training data or training process necessary for enablement?

Clear and complete disclosures are cornerstones in the *quid pro quo* of the patent system. The patent provides a 20-year exclusive right, creating artificial scarcity to incentivize R&D investment, in exchange for disclosing the invention to the public database. The theory of disclosure is that it increases social welfare by facilitating dissemination of information, competition at the end of the patent period and follow-on innovation for better substitutes or new products. However, if the disclosure requirements are too high, the incentive to disclose new inventions could be chilled.

As a business strategy, patentees may try to draft patents that minimize disclosure in order to file more quickly with the patent office and to obtain a degree of commercial advantage by keeping some information private. Furthermore, in the field of AI, it may be challenging to describe and explain the invention textually. This is due to some of the 'black box' characteristics of AI inventions and the fact that several parties are often involved in the ANN design, the training process and the data collection process.[1] Yet, the requirement of clear and complete disclosures is as relevant for AI inventions as for inventions in other fields.

T 0161/18 illustrates the difficulties of meeting the disclosure requirement in the context of medical AI inventions. The decision underscores the increasing scrutiny of AI-related patents, which in turn reflects broader legal and regulatory concerns about patents for AI technologies.

In this article, we analyse the T 0161/18 decision and subsequent case law on the sufficiency of disclosure of AI inventions. We tease out the emerging issues, in particular the extent to which patent applicants must disclose input data suitable for training AI. We also compare the emerging case law with recent updates to the EPO's Patent Examination Guidelines, suggesting several reforms.

II. Analysis of EPO's case law

A. Case T 0161/18: input data suitable for training the neural network or a data set (12 May 2020)

Case T 0161/18 concerned a European patent application that claimed an invention that uses machine learning (ML) involving an ANN to determine cardiac output from a peripherally measured arterial blood pressure curve (EPO patent application number 06804383.5).

In the final decision of T 0161/18 (Equivalent aortic pressure/ARC Seibersdorf) of 12 May 2020 the BoA of the EPO considered *a method for determining cardiac output* (EP1955228). The application was directed to a method and a device (T 0161/18, paragraph V), namely:

> Claim 1 'a method for determining the cardiac output from an arterial blood pressure curve measured at the periphery, in which the blood pressure curve measured at the periphery is mathematically transformed to the equivalent aortic pressure and the cardiac output is calculated from the equivalent aortic pressure, characterized in that the transformation of the blood pressure curve measured at the periphery is converted into the equivalent aortic pressure with the help of an artificial neural network, the weighting values of which are determined by learning' and
>
> Claim 8 - a device that implemented this method.

The EPO's Examining Division refused the application for lacking an inventive step under Article 56 EPC. The Examining Division relied on two prior art documents. The closest prior art disclosed the essential features of claim 1, except for the use of ANN for carrying out the calculation of cardiac output from a peripherally measured arterial blood pressure curve. The Examining Division observed that this distinguishing feature of the application was a technical contribution since it provided an alternate automated method of calculation, which would increase precision and efficiency. However, this technical contribution was rendered obvious in light of a document that disclosed the use of neural networks for similar purposes. Importantly, the Examining Division held that the application merely outlines an intention to use neural networks for this purpose and does not provide any technical teaching on the implementation of this technology (Examining Division's decision to refuse (grounds), paragraph 8.3.2). The application was therefore refused. The applicant appealed.

One of the main grounds for appeal was that the technical effect achieved through the use of ANN was independent of the specific technical implementation of the ANN. It was argued that the latter was an issue of enablement under Article 83 EPC which the Examining Division had not raised or considered. In its preliminary opinion, the BoA raised objections under both Articles 83 and 56 EPC. Its final decision turned on the interpretation of sufficiency of disclosure for inventions utilizing ANN (Article 83 EPC) and the implications of lack of sufficiency of disclosure for inventive step (Article 56 EPC). In this regard, the BoA held:

> Since in the present case the claimed method differs from the prior art only by an artificial neural network, the training of which is not disclosed in detail, the use of the artificial neural network does not lead to a special technical effect that could give rise to inventive step (T 0161/18).

The BoA offered the following reason for deciding that disclosure was insufficient under Article 83 EPC:

> […] With regard to the training of the neural network according to the invention, the present application only discloses that the input data should cover a wide range of patients of different ages, genders, constitution types, health conditions and the like, so that there is no specialization of the network (see page 5, last paragraph to page 6, first paragraph). However, the application does not disclose which input data is suitable for training the artificial neural network according to the invention, or at least one data set suitable for solving the technical problem at hand. The training of the artificial neural network cannot therefore be reworked by the person skilled in the art and the person skilled in the art can therefore not carry out the invention. The present invention, which is based on machine learning in particular in connection with an artificial neural network, is therefore not sufficiently disclosed, since the training according to the invention cannot be carried out due to a lack of corresponding disclosure. (T 0161/18)

A key issue for the BoA was that several types of information had not been disclosed. For instance, (i) the characteristics of the ANN such as type and architecture of the network were missing. And (ii) the application did not disclose specific input data suitable for training the ANN to achieve the invention (T 0161/18, paragraph 2.2). Instead, the application merely disclosed that input data should cover a 'wide range of patients', and mentioned general selection criteria for this purpose, such as age, genders, constitution types, health conditions, and the like.

[1] N Price, 'Big Data, Patents, and the Future of Medicine' (2017) 37 Cardozo Law Rev 1401.

In the absence of the characteristics of the neural network and specific details of the input data (or at least one data set suitable for solving the technical problem), the BoA was of the view that there was no specialization disclosed. Furthermore, persons skilled in the art would not be able to train their own specializsed ANN based on the disclosures in the application and common general knowledge. They could not realize the claimed invention because it claimed 'the weighting values of which are determined by learning', which left too much burden on the skilled person.

Without specialization, the invention also lacked an inventive step. The complainant argued that the use of ANN had the technical effect of determining the cardiac output reliably and precisely. However, since neither the claims nor the description contained details regarding training of the ANN, the BoA remained unconvinced on its technical effect. In its opinion, the mere reference to an 'ANN the weighting values of which are determined by learning' did not set the application apart from the prior art or from what experts already understood about how ANNs could be used. Therefore, the Board was of the view that the claimed neural network was not adapted to the specific claimed application and lacked inventive step (T 0161/18, paragraphs 3.5 and 3.6).

The decision showed the effect of a sufficiency-inventive step 'squeeze' in the field of AI inventions. It also showed how this can translate into a requirement to disclose training data or the training process in more detail. The basis of the squeeze was that—having omitted the details of the training data from the patent application—the applicant faced the prospect that either the critical feature for establishing inventive step was not sufficiently disclosed (ie ANN specialization) or else the patent was obvious because all that it taught was something the skilled person could and would do with common general knowledge and prior art (ie contemplate applying ANN to automate cardiac output). The net effect was that the patentee was 'squeezed' between the patent being insufficiently disclosed or lacking inventive step. The BoA made clear that the squeeze could have been avoided if the patent applicant had disclosed the missing information, such as the details of the specific input data suitable for training the ANN, or at least one data set suitable for solving the technical problem at hand.

A similar 'squeeze' will occur in other fact scenarios where the idea of applying AI (eg ML, deep learning or other types of AI learning) is close to being obvious. In those scenarios, disclosing a generic AI/ML/ANN may fall short. To an extent, the issue will depend on the specific claim language used in future cases. In T0161/18, the claim language was 'with the help of an artificial neural network, *the weighting values of which are determined by learning*' (emphasis added), and the BoA called for 'input data … suitable for training the ANN … invention, or at least one data set suitable for solving the technical problem' (ie to obtain the model weight values by learning). Nevertheless, T0161/18 signals that characterizing the input data necessary for training the ANN could often be important for a valid patent. One data set for solving the technical problem to enable replication of the 'specialized network' might suffice, depending on the circumstances.

1. *Training data*

This raises the question—where specialization of the AI is needed for inventive step, to what extent is disclosure of the training data or training process necessary for enablement? Does Article 83 EPC call solely for the disclosure of training input data or the training process? Or are there alternative ways for applicants to satisfy the disclosure requirements in AI inventions? T 0161/18 case was silent about alternative or complementary ways to satisfy the disclosure requirement for AI inventions. But, in practice, other alternatives could also be considered. For instance, in our view, instead of the input data the applicant could have disclosed the ANN model weights (ie 'the weighting values') to implement an embodiment of the claimed invention. This would also have resulted in a 'specialization of the network' to solve the technical problem—as a trained ANN is *specialised* and fully defined by its model weights. Alternatively, the patent specification could have disclosed a particular AI architecture (eg details of the particular type of neural network, number of layers, user-specified parameters, code listings, etc) to solve the technical problem, as well as the details of the training process so that those with ordinary skill in the art could carry out the training to replicate the claimed invention. Further, specialization of the network could also be achieved by disclosing the collection process of the training data set, trained weights, and in some cases validation data with examples.

While the BoA in this case was not exhaustive in detailing the various possible ways to satisfy the disclosure requirements for AI inventions, the decision makes it clear that the patent specification must disclose sufficient details so that a skilled person can make the AI that is no longer *generic* but *specialized* to solve the particular technical problem.

More guidance on the types of information that can meet the disclosure requirement emerges from subsequent cases discussed below.

B. Case T 1191/19: learning process, training data sets and validation data (1 April 2022)

The patent application in T 1191/19 concerned a medical AI invention. The invention related to an AI-implemented method for personalized therapy selecting between different neuroplasticity interventions for patients. To implement the method, the AI model had to be trained. This required a database of patients (classified based on a 'meta learning scheme') and patient outcomes with different interventions.

The BoA held that the application did not disclose how the AI-implemented method was applied to the problem (yet this was required to establish inventive step). Nor did it disclose training data or validation data that the AI-implemented method was trained on (this was required for sufficiency of disclosure).

On sufficiency, the BoA observed that details of the architecture of the ANN, its topology, activation functions, end conditions and learning mechanisms were not specified in the patent application. The BoA observed that the application also failed to disclose the minimum number of patients needed to compile the training data. The BoA specifically noted that '[given] the level of abstraction of the application, the available disclosure is more like an invitation to a research programme' (paragraph 4.1.). For these reasons, the patent application was refused for lack of sufficiency of disclosure (Article 83 EPC).

The BoA was also of the view that the patent application lacked inventive step. The BoA, referring to the case T 0161/18 stated that, 'the mere application of a known machine learning technique to problems in a particular field is a general trend in technology […] and cannot be inventive as such'. The BoA gave the applicant an opportunity to substantiate how the invention was specialized to solve the particular problem (ie predicting personalized interventions for a patient in processes of which the substrate is neuronal plasticity). However, since there was no non-obvious detail in the patent application, the application was refused on the ground of lack of inventive step.

Like the BoA's decision in T 0161/18, this decision is also an example of the sufficiency-inventive step 'squeeze' in the field of AI inventions. The BoA's findings on lack of disclosure also resulted in a finding of lack of inventive step. This was because the ML technique used in this application was generally known. Such a technique was also cited in prior art. Therefore, the idea of applying the ML technique to the problem at hand without information such as training data, architecture of the artificial neural network, its topology, activation functions, end conditions and learning mechanisms rendered the invention obvious. The critical feature for inventive step was specialization of the AI-implemented method.

C. Case t 1539/20: model definition, architectural details and steps (24 November 2022)

The patent application at issue in T 1539/20 related to monitoring the performance of an IT system with distributed network-connected nodes. The invention was cited as useful to businesses needing to monitor real-time behaviour and performance of applications. The invention claimed to enhance higher-level monitoring of application systems by understanding application performance through data messages passed between nodes. Although not expressly stated in the patent, the main focus of the application was on the process of correlating lower-level data messages between nodes to higher-level transactions/messages.

For this, the patent application disclosed steps of 'mapping the application system onto a hierarchical model for the application system according to a network topology of the application system'. Claim 1 defined the hierarchical model as having a plurality of levels with each level having a plurality of connected nodes, and each level of the model bearing some relationship to the application system.

The BoA found the patent application to be unclear since it did not disclose how the 'hierarchical model' relates to the 'application system'. Specifically, the application did not reveal how the number of levels, the nodes or groups of nodes of the application system relate to the hierarchical model and how the hierarchical model influences the messages. Therefore, they held that the meaning of 'mapping the application system onto a hierarchical model for the application system according to a network topology of the application system' was unclear under Article 84 EPC (this article specifies that inventions should be *clear and complete* for persons skilled in the art to carry it out).

The appellant argued that the term 'mapping' was clear to persons skilled in the art and was a common term of art. However, the BoA held that as per Article 84 EPC, which requires patent applications to be sufficiently disclosed, claims are to be clear from their wording alone. Not only were the claims unclear but also insufficiently disclosed since the patent application did not disclose how a person skilled in the art could implement the mapping process in software. The BoA noted that the mapping process generally referred to applying a 'general model' to a given application system to obtain an 'applied system'. Since the application did not disclose the relevant details of the distributed application system, and there was no explanation of how and based on what kind of input the software implementing the method was to obtain information about a given application system, the patent application was not enabling. In short, the patent application did not provide sufficient indication of how an automated mapping process could be implemented by a person skilled in the art.

Unlike the BoA's decision in T 0161/18, the BoA did not decide on inventive step. The refusal of the patent application was based only on lack of clarity and insufficiency of disclosure. The BoA clarified that information regarding training data and process for training were essential to persons skilled in the art to reproduce the invention. As observed, 'In the board's view, implementing such a complicated automated process in software, essentially mimicking the cognitive process of a human systems engineer, is well outside the normal abilities of the skilled person' (paragraph 3.2.4), indicating that reproducibility by persons skilled in the art required additional information.

D. Case T 0606/21: input data, testing and evaluation (28 February 2023)

The patent application in T 0606/21 pertained to a computer-implemented method for evaluating predictions of trajectories by autonomous driving vehicles. The method comprised generating a Deep Neural Network (DNN) model. The patent application disclosed the training process of the DNN model. This included defining process steps and inputs such as receiving certain data (pairs of trajectories of objects), analysing that data (extracting predicted features of the object and analysing similarity of trajectories) and generating outputs/similarity scores.

While the training process, as well as characteristics of input data, were broadly disclosed, the Examining Division refused the application for insufficiency of disclosure. The Examining Division held that the input data for training the DNN model was incomplete, which affected the reliability of the output. More specifically, the disclosed input data only included predicted trajectories, not actual trajectories (which are important for evaluating predicted trajectories).

While the applicant argued that the application clearly disclosed the invention, the BoA did not share the applicant's view. According to the BoA, the claimed invention consisted of two stages—one, generation and training of the DNN model by inputting pairs of trajectories in order to analyse the similarity of these trajectories and to improve the accuracy of this analysis. And second, testing and evaluation, in which a second predicted trajectory of the object was the input. The BoA held that there was a lack of disclosure in the second stage (ie testing and evaluation of the DNN model). The technical details of this stage were explained by the BoA as follows:

> This phase defines that a second predicted trajectory of the object is received, that features of points of the second predicted trajectory are extracted and that the trained DNN model is applied to the extracted features of the second predicted trajectory. According to claim 1, the trained DNN model then generates a second similarity score, the second similarity score indicating whether the second predicted trajectory is more likely to be close to an (unspecified) actual trajectory along which the object is likely to move in the near future. (paragraph 2.3)

The BoA took issue with the fact that the patent application did not disclose how the unspecified actual trajectory was generated and inputted. The BoA therefore agreed with the Examining Division that the application did not contain sufficiently clear and complete information that would enable persons skilled in the art to carry out the invention. Unlike the BoA's decision in T 0161/18, the BoA did not decide on inventive step in this case. The refusal of the patent application was based only on insufficiency of disclosure.

E. Case T 1526/20: validation and reliability (14 March 2023)

The patent application in T 1526/20 concerned a computer-implemented 'liveness' testing method for facial recognition. The

purpose of the invention was to use this method to detect impersonations using pictures of faces. The method relied on detecting differences in uniformity between real, live (3D) faces and 2D pictures of the person. The algorithm steps were disclosed. As per the process disclosed, uniformity was to be measured using a diffusion filter and by measuring diffusion speeds and illumination changes. The features were converted into binary numbers, and if the number of non-zero pixels was low, the imaged object was considered a fake. The BoA noted that the application discussed two distinct scenarios related to: first, liveness detection; and second, face recognition. In both scenarios, the first step was to apply image diffusion to the original image.

The Examining Division refused the application on the grounds that 'the claimed techniques, "relying only on diffusion filtering" were not robust against illumination changes and led to "totally unreliable results"'. The Examining Division was of the view that consistent results could be obtained only if further steps, as disclosed in identified scientific publications, were specified in the patent application. A similar objection was raised for computation of diffusion speeds. Scientific publications were cited by the Examining Division that indicated gaps in the claimed method. Non-inclusion of material steps and evidence that the invention actually worked (eg validation or performance results) rendered the patent application insufficiently disclosed. Specifically, additional evidence was required to show which conditions for image acquisition (eg which filters and filter parameters) were suitable to solve the problem.

The applicant argued that although the invention might not work in some hand-picked cases, it worked for the majority of use cases. It also argued that a patent application should not be assessed on strict disclosure criteria applicable to scientific publications. However, the BoA was not convinced. It held that additional evidence and steps were required to implement the invention; it was not sufficiently disclosed. While it did not expressly comment on the applicability of disclosure standards of scientific journals, it implied that such standards may be relevant to help assess the sufficiency of disclosure. Unlike the BoA's decision in T 0161/18, the BoA did not decide on inventive step. The refusal of the patent application was based only on insufficiency of disclosure and the inability of persons skilled in the art to implement the invention. The invention lacked reliability (or disclosure showing that the invention worked) and failed to teach persons skilled in the art the useful information required for its industrial applicability.

F. Reflections on the EPO case law

While a core AI invention advances the field of ML, an applied invention uses AI to solve specific problems in other technical areas. All of the cases considered above were for applied AI or computer-implemented inventions used to solve specific problems—determining cardiac output, automated IT systems performance monitoring, predictions for autonomous driving vehicles, liveness testing for facial recognition and personalization for neuroplasticity interventions.

When it comes to disclosure requirements, applied AI inventions can present some challenges given their 'black box' characteristics. However, these cases help identify the standards for future applicants. Collectively, these decisions show that typically with applied AI patent applications there must be sufficient detail to enable the skilled person to reproduce a *trained* AI. The AI must not be *generic* but *specialized* to solve the particular technical problem. The mere application of AI/ML to a particular field, to solve a particular problem, is not the basis of a valid patent (T1191/19). This is due to the requirements of inventive step, industrial application and sufficiency of disclosure. Patents must not be mere 'hunting licenses' for future research endeavours or 'armchair inventors' (G2/21;T0258/21) obtaining monopolies for something they have not actually invented.[2] They must not be directed to imaginary intentions or completely abstract conceptions of the AI invention. And, with this in mind, it is then also necessary to *disclose* sufficient detail for the skilled person to replicate the *trained* AI.

To replicate a trained AI, these cases show that various information could be disclosed to help the skilled person, and might actually be required. For example, (i) the structure and architecture of the network (T 1191/19), (ii) the complete training process (T 0606/21), (iii) input data for training, (iv) at least one data set suitable for solving the technical problem (T 0161/18; T 1191/19) or (v) validation data or performance results (T1526/20). Showing that the invention worked in a majority of, but not all, cases may not be sufficient (T 1526/20); consistently reliable outputs may be required. Furthermore, standards of disclosure applicable to scientific publications may apply to AI-related patent applications at least for disclosures related to the training process of the model. Relatedly, the completeness of training input data and information about the training process will be examined with an eye to the reliability of output; as these criteria affect meaningful reproducibility and whether the invention works to solve the technical problem (T 0606/21).

III Analysis of EPO's patent examination guidelines—2024 update

The EPO's examination guidelines (Guidelines) on sufficiency of disclosure were updated in early March 2024. They now state:

> If the technical effect is dependent on particular characteristics of the training data set 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. (Part G-II, paragraph 3.3.1.)

It is interesting to note that the EPO took close to 4 years to update its examination guidelines to reflect the standards established in T 0161/18 (May 2020). And yet, the Guidelines fail to capture the tone of the BoA's reasoning in T 0161/18 and subsequent cases. Those decisions indicate that details of training data will often be required for applied AI inventions in order to avoid insufficient disclosure, otherwise the patent is likely to suffer from lack of inventive step, industrial application or sufficiency of disclosure.

The updated Guidelines still do not incorporate the guidance from other Article 83 cases and the corresponding information that if not disclosed could result in insufficiency of disclosure under Article 83 EPC (see Part II F). For example, (i) the structure and architecture of the network (T 1539/20), (ii) the details of training process (T 0606/21 and T 1526/20), (iii) input data for training, (iv) at least one data set suitable for solving the technical problem (T 0161/18; T 1191/19) or (v) validation data (T 1191/19). The trained AI model weights could also assist the skilled person.

[2] P England, 'Patents and Plausibility' (2013) 9 JIPLP 22. See also; T Minssenand D Nilsson, 'The Industrial Application Requirement for Biotech Inventions in Light of Recent EPO and UK Case Law: A Plausible Approach or a Mere "Hunting Licence"?' in J Rosén (ed) *European Intellectual Property Law* (Critical Concepts in Intellectual Property Law) (Edward Elgar Publishing 2016) 164–78.

The Guidelines capture some of the additional nuance in the cases, but not all. For instance, the Guidelines make clear that showing that the invention worked in a majority of cases, but not all (ie not over the whole ranged claimed), may not be sufficient (T 1526/20).

> [I]n the field of artificial intelligence if 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.

On the other hand, the Guidelines do not mention that standards of disclosure applicable to scientific publications may apply to AI-related patent applications, at least for disclosures related to the training process of the model (T 1526/20). Addressing this issue would improve the EPO Guidelines.

Another improvement would be for the Guidelines to include a set of examples, and for the EPO to audit a sample set of granted patents in this field.[3] Certain fields of AI applications are maturing—image captioning models and object detection models—with metrics and benchmark data sets being established.[4] These would be useful for fine-tuning advice on disclosure requirements in the examination guidance.

A more substantial change (and improvement) would involve introducing an 'AI disclosure checklist' for patent examination.[5] The checklist could include a series of points, *inter alia*: (i) Does the patent specification include details of the AI architecture?; and (ii) Does the patent specification include details of the training process and data, of at least one data set suitable for training the claimed AI? If the answer to any of the checklist questions is negative, the examiner would be deemed to have satisfied the initial burden of proof to issue an Article 83 EPC objection—the burden would then shift to the patent applicant to rebut the presumption of insufficient disclosure. The examination checklist would thus help reduce prosecution time by encouraging patent applicants to disclose information pro-actively as part of their initial patent application, enabling patent examiners to evaluate the application more efficiently in their first office action with substantive examination and reducing the likelihood of an EPO objection. The patent applicant usually has the information suggested above for the checklist, and disclosure is unlikely to be strategically detrimental. This is because such disclosures can benefit patentees by raising the bar of patentability for potential competitors for subsequent inventions, reducing the risk of invalidity challenges and establishing earlier priority dates for subsequent continuation applications.[6]

Evolving legal perspectives and technological advancements raise questions about the person skilled in the art's use of AI assistance and generative AI. The shift now allowed from person skilled in the art to 'research or production team skilled in art' could potentially even move on to 'AI skilled in the art'. The significance of this for AI-related disclosures is that the person skilled in the art is taken to be the interpreter of information in the patent disclosure, and the benchmark for deciding if the disclosures clearly and completely disclose the technical teachings required for persons skilled in the art to carry out the invention *without burden*. If the hypothetical person's skill is pegged at the level of knowledge and processing power of an AI system, then lesser disclosures may be sufficient to enable the claims.[7] But it is unlikely that wholly omitting information about training data, AI architecture and the like would be acceptable for applied AI inventions in the near future.

For applicants in the process of disclosing their model weights or training data set, a remaining challenge is how to incorporate such disclosures in the patent document. This is not discussed in the EPO's examination guidelines. A couple of solutions might be to introduce a deposit system along the lines of the Budapest Treaty for Deposit of Microorganisms or the systems employed by scientific journals.[2]

In the microorganism context, the EPO allows applicants to deposit biological material to supplement their European patent applications in certain cases. In line with the Budapest Treaty for Deposit of Microorganisms, inventions for biological material that are not available to the public or cannot be described in a sufficiently enabling manner can be deposited with a recognized depository institution along with the patent application (Rule 31). The patent application then need only include details of the characteristics of the biological material. Further, the rules restrict onward disclosures and transfers. Material that has been deposited is available to only those requesters who make certain undertakings to the applicant regarding transmission and use of the material (Rule 33).

An analogous deposit mechanism could be considered in the AI context. Like the biological context where inner processes are not fully understood, model weights, example training data sets or other material could be deposited to satisfy teaching requirements for AI-related inventions.[8,9] Such a scheme would enhance AI patent disclosures especially when textual explanations are difficult to provide in patent documents. Also, inspiration may be taken from the procedure in place for biological material for transmission and use by the public of the deposited data for verifying sufficiency of disclosure, experimentation or for follow-on research.

Reporting guidelines used in scientific publishing could also be considered to determine standards of disclosure.[10] Scientific publication guidelines for AI-based medical devices may be especially useful for AI medical devices and other AI-enabled clinical support tools. Examples of these guidelines include the Consolidated Standards of Reporting Trials–Artificial Intelligence or the Minimum Information about Clinical Artificial Intelligence Modelling checklist.[11] These guidelines set forth disclosure standards for training, selection, handling, optimization, algorithms, performance, as well as hardware and software requirements of the AI's operational environment. They also set out a tiered system to evaluate transparency—from 'complete open sharing of all the software code and scripts' (Tier 1) to 'no sharing of the underlying model or codebase' (Tier 4). Depending on the context of use of the AI system, the sensitivity and risk involved, the EPO may consider borrowing from these guidelines to establish a calibrated disclosure standard.[12]

While scientific journals and a deposit system may be considered for overcoming practical challenges to improve AI disclosures, the EPO may also consider incentivizing such disclosures

[3] Rebeca Ferrero Guillén and Altair Breckwoldt Jurado, 'Vagueness in Artificial Intelligence: The "Fuzzy Logic" of AI-Related Patent Claims' (2023) 2 Digital Soc, 1–25.
[4] ibid.
[5] M Aboy and others, 'The Sufficiency of Disclosure of Medical Artificial Intelligence (AI) Patents' (2024) 42 Nat Biotechnol 839–45.
[6] ibid.
[7] Susan Y Tull and Paula E Miller, 'Patenting Artificial intelligence: Issues of Obviousness, Inventorship, and Patent Eligibility' (2018) 1 RAIL 313.
[8] Aboy and others (n 5).
[9] Tabrez Y Ebrahim, 'Artificial Intelligence Inventions & Patent Disclosure' (2020) 125 Penn State Law Rev 147.
[10] Matthew Chun, 'Artificial Intelligence for Drug Discovery: A New Frontier for Patent Law' (2024) 104 J Pat & Trademark Off Soc'y 5.
[11] ibid.
[12] ibid.

through process-related incentives such as prioritized examination and reduced maintenance fees.[13]

IV Empirical findings about the sufficiency of disclosure for AI inventions

The decisions in T0161/18 and T 1191/19 concerned medical AI patents. How many of these AI patents have been granted? Do they disclose the training data or other details? Are they vulnerable to invalidation/revocation? These sorts of questions shed light on the trends in the AI field and the magnitude of insufficiency issues.

Two recent empirical IP studies analysed these questions. The first study undertook a comprehensive review of medical AI trends.[14] It found a significant increase in the number of medical AI patents awarded by the EPO and USPTO. The number of patents containing medical AI/ML claims—similar to those in the T 0161/18 case—has been increasing since 2013 with a compound annual growth rate of 33.48 per cent.

The second study found that fewer than 70 per cent of the granted medical AI patents disclosed specific details of the AI/ML training or specific AI/ML model architecture. Based on T 0161/18 decision and subsequent decisions, a substantial number of these AI patents are likely to be insufficiently disclosed. The study also showed that, of the medical AI patents analysed, fewer than 30 per cent disclosed information needed to assess the performance of at least one embodiment of the claimed inventions. Less than 25 per cent of the disclosures included mathematical details, and fewer than 3 per cent included illustrative code listings.[15]

Although these empirical studies focussed on disclosures within medical AI patents, it is likely that similar trends exist with AI patent in other areas. Thus, there would appear to be a major issue with the validity of many granted AI patents, and patent drafting will need in future to change to avoid insufficiency.

VI. Conclusion

In light of these empirical results showing that there has been a marked increase in the number of AI patents, and that fewer than 70 per cent of a cohort of granted patents disclosed specific details of the AI/ML training or specific AI/ML model architecture, the BoA decisions in T 0161/18 and subsequent cases are highly significant. They emphasize that the mere idea of applying an AI system, a neural network, ML or deep learning to solve a technical problem may not meet inventive step requirements. Rather, the AI needs to be adapted to the specific claimed application and the patent application needs to disclose details regarding the specialization of the network.

While the decisions of the BoAs have not been exhaustive in detailing the various possible ways to satisfy the disclosure requirements for AI inventions, they have highlighted the importance of disclosing training details or at least one data set suitable for training the ANN according to the claimed invention. Other disclosures to support enablement include the AI architecture, the training process, model weights or validation data along with AI details.

The most recent revision of the EPO's examination guidelines includes some of the points emerging from the cases discussed in this article, signifying their importance to patent practitioners and patent examiners. But some of the tone and nuances have been overlooked. To improve the Guidelines, the EPO could also consider including additional examples and potentially adopting a 'AI disclosure checklist' for patent examination, as described above.

Getting the balance right with patent disclosures is very important for patent policy; not only domestically but also globally. Disclosures can help facilitate knowledge 'spillovers' for low and middle-income countries, enhance technological capabilities across the sector and lessen the technology gap between the Global North and the Global South.

This article forecasts that a large number of granted AI patents may be invalid due to a squeeze between disclosure, on the one hand, and inventive step, industrial application, and sufficiency of disclosure, on the other. But it also describes how patent attorneys and Examination Guidelines can handle the issues better in the future.

Funding

The research was supported, in part, by a Novo Nordisk Foundation Grant for a scientifically independent International Collaborative Bioscience Innovation & Law Programme (Inter-CeBIL programme—grant no NNF23SA0087056).

[13] Ebrahim (n 9).
[14] M Aboy, 'Mapping the Patent Landscape of Medical Machine Learning' (2023) 41 Nat Biotechnol 461–68.
[15] Aboy and others (n 5).

Journal of Intellectual Property Law & Practice, 2024, Vol. 00, No. 00 , DOI: https://doi.org/10.1093/jiplp/jpae063, , Article