



Intellectual
Property



A Primer to AI and Patents



AI/ML is no longer confined to the fields of maths, stats and computer science – advisors must be able to identify the relevant issues

INTRODUCTION

Recent years have seen renewed interest in the world of artificial intelligence and machine learning. Tools such as TensorFlow, PyTorch, MXNet and AWS have moved machine learning from academic niches to the masses. AI/ML is being used today in fields as diverse as drug discovery, immunotherapy, automotive, finance, telecommunications and manufacturing — the list is too extensive to feasibly list here. This crossover means that it is essential for advisors to up-skill to ensure they are able to properly advise applicants. AI/ML is no longer confined to the fields of mathematics, statistics and computer science. Advisors must be able to identify the relevant issues so that they know when to seek specialist advice.

In this paper we have set out some useful guidance points for practitioners, regardless of their technical background. We hope it is relevant for both machine learning experts as well as complete novices — let's call it a 'primer'.

We begin by saying what we will not cover. This paper is not concerned with crystal ball gazing or other policy considerations but is instead a pragmatic guide to the state of play in 2019.

Ideas on non-human inventorship or giving personhood to AI are better suited to academic writing. Similarly we will not touch on the ethics of AI or the philosophical role that IP can have in improving public trust in machines; neither will we touch on matters closer to home such as changes that will (or should) occur to the hurdle for determining inventiveness as a result of the incorporation of AI into the skilled team.

As AI becomes more widespread and invades the public consciousness there is a mistaken supposition that systems are being developed to replicate the thinking of human specialists. This anthropocentric view of 'intelligent' systems is dangerous when considering AI patents. Systems of today are increasingly out-performing human experts, not by copying high-performing people but by exploiting the distinctive capabilities of new technologies such as massive data-storage capacity and brute-force processing. As we go through the current AI issues we caution the reader to be wary of this kind of bias. AI inventions do not currently occupy a special category, and can be considered on the same basis as others in the field of computer technology.



AI/ML 'PRIMER'

So what do we mean by AI? Generally when most people use the term "AI" they really mean machine learning and we will predominantly use this term here, but let's start with some terminology:

Artificial Intelligence

The modern definition of artificial intelligence (or AI) is "the study and design of intelligent agents" where an intelligent agent is a system that perceives its environment and takes actions which maximizes its chances of success.

Artificial General Intelligence

Artificial general intelligence is a machine that could successfully perform any intellectual task that a human being can. In effect, true intelligence and not data analysis. Most academics believe we are decades from this level of technology.

Machine Learning

Machine learning (ML) is a field of AI but is usually what people mean when they refer to AI. Machine learning is the programming of a digital computer to behave in a way which, if done by human beings or animals, would be described as involving the process of learning.

Machine learning uses statistical techniques to give computer systems the ability to "learn" (e.g., progressively improve performance on a specific task) from data, without being explicitly programmed.

Typically the process of machine learning is building a mathematical model from a set of input training data and then applying that model to a set of test data to provide a prediction or output.

Often the most important part of a machine learning development process is the feature analysis. This is the selection of a subset of relevant features (variables, predictors) for use in model construction. For example, the programmer that chooses which features of a set of cars should be given to a ML model to help the model identify which car will be fastest.

Deep Learning

Deep learning is a subset of machine learning and is based on learning data representations, as opposed to task-specific algorithms. With deep learning it is generally not possible to identify "how it did it". Since there is generally no standard definition of the term we encourage readers to avoid its use where possible.

PATENTABILITY

As far as patents are concerned there are generally three categories of inventions — inventions on AI, inventions using AI as a tool, and truly AI-created inventions. Most applications will fall into the second category as the use of machine learning becomes more prevalent. In the majority of jurisdictions, the patentability of an invention in any one of these categories is considered to be similar to applying the general principles of patent eligibility as the inventions are considered to be computer implemented inventions (or mathematical methods).

Patentability in Europe

The EPO held a conference on patenting AI in 2018. The message that the EPO was trying to get across is that it believes its main challenge is likely to be rapid growth of AI across a range of technical fields — what analysts call “the fourth industrial revolution” —and that the EPO already allows technically mixed examining divisions which may be useful for inventions that apply ML in diverse technologies. The EPO is particularly keen to emphasise that simply claiming the use of ML in a particular application will be considered obvious¹.

On patentability, the EPO has identified three types of patentable inventions to do with AI:

AI patents: type 1 — ‘Core AI’

These are patents which relate to algorithms, as such. These inventions relate to the design of AI as a tool, rather than its application to a particular problem.

These inventions are likely to be a rarely encountered by most patent practitioners. More commonly, applicants will be using an existing AI tool in a particular application; indeed most currently used AI algorithms are based on academic papers—some in fact are fairly old. Applicants will need to demonstrate the technical character of the invention and that the invention is more than a mathematical method.

These techniques are likely to be difficult to protect, since they are unlikely to be tied to a particular technical advantage. Therefore, the EPO is likely to view this category of inventions as excluded from patentability on the grounds that they relate to no more than a mathematical method.

One point of note is the so-called ‘safe harbour’ where a type 1 invention can gain a technical effect by virtue of the invention being adapted



in some way to function more effectively on a computer (see point (ii) in the following section). In this case, application of the algorithm to a particular technical field is not necessary to gain technical character. Our view is that inventions capable of making use of the safe harbour will be encountered very infrequently in practice.

AI patents: type 2 — Generating a training set or training a model

The EPO considers that the steps of generating a training set may contribute to the technical character of the invention if they support achieving a technical purpose. An objection of lack of inventive step owing to an absence of technical character could arise if an examiner does not see a clear link between the training set and a technical result.

AI patents: type 3 — AI as a tool

This type of patent is the use of AI in an applied field, defined by way of technical effects. This is the most likely avenue for success for applicants. In fields such as autonomous vehicles and healthcare, AI might be claimed as a tool for using a training set to provide a solution, which yields a technical advantage.

¹(T 1510/10; see in particular Reasons point 8)

2018 EPO Guidelines

In November 2018 the EPO updated their examination guidelines to specifically include a passage on machine learning methods. They outline that the computational models and algorithms used in AI and ML are considered to be mathematical methods, therefore the general principles regarding examination of these methods apply.

Mathematical methods must be claimed together with technical means, to avoid the exclusion under Art. 52(2)(a) EPC and must either:

- (i) serve a technical purpose by their application to a field of technology, and/or
- (ii) be adapted to a specific technical implementation.

The EPO Guidelines give some examples of AI and ML techniques that find applications in various fields of technology, such as:

- Use of a neural network in a heart monitoring apparatus for the purpose of identifying irregular heartbeats. Such a technique would make a technical contribution.
- The classification of digital images, videos, audio or speech signals based on low-level features (e.g. edges or pixel attributes for images). These are typical technical applications of classification algorithms.
- Classifying text documents solely in respect of their textual content. This is unlikely to be is however not regarded as to be a technical purpose, rather the purpose is linguistic (T 1358/09), which is non-technical.
- Classifying abstract data records or even “telecommunication network data records” without any indication of a technical use of the resulting classification. This is also not a technical purpose (T 1784/06), according to the EPO.

At the EPO, AI is generally considered a mathematical method and these tests apply

Finally, the guidelines also cover the type 2 patents identified above, where a classification method serves a technical purpose, the steps of generating the training set and training the classifier may also contribute to the technical character of the invention if they support achieving that technical purpose.

Patentability in the UK

The UK IPO has yet to issue guidance on the patentability of AI inventions and we are unlikely to see any in the short term. The UK IPO tends to provide guidance following high profile cases in the Courts, and so we await guidance from the Courts when an AI-based patent is litigated— there is no indication that this will happen any time soon.

The best guidance remains the ‘5 signposts’ and guidance from the Courts in the Halliburton decision. In Halliburton the modelling/simulation of a drill bit was considered patentable as it was for a technical purpose. Applying this approach to machine learning, it is likely that machine learning inventions will be patentable if it can be demonstrated they are for a technical purpose (similar to type 3 at the EPO identified above). Type 1 and 2 at the EPO are less likely to be looked upon favourably at the UK IPO, but to be considered patentable by examiners at the UK IPO a machine learning invention will most likely have to satisfy one of the criteria below. It is particularly helpful to be able to map an invention to one of signposts (i) to (iv), and to draft a patent application with this in mind — for example, it is helpful to explicitly limit a claim to the use of an output of an AI invention in a technical context so as to fall within signpost (i).

The signposts are:-

- (i). whether the claimed technical effect has a technical effect on a process which is carried on outside the computer
- (ii). whether the claimed technical effect operates at the level of the architecture of the computer; that is to say whether the effect is produced irrespective of the data being processed or the applications being run
- (iii). whether the claimed technical effect results in the computer being made to operate in a new way
- (iv). whether the program makes the computer a better computer in the sense of running more efficiently and effectively as a computer
- (v). whether the perceived problem is overcome by the claimed invention as opposed to merely being circumvented.



will be considered obvious as will other similar 'problem-inventions'. Applicants may have to demonstrate there is a prejudice against using machine learning in the art for some reason. It is likely to be a challenge to obtain patent protection for this category of invention, since the use of machine learning is likely to be considered obvious for any problem that is capable of being analysed on a computer.

Step 2: Decide upon the data you need

Patentable inventions could lie in the non-obvious selection of a particular parameter ('feature') in the data used to train a model or the structuring of that data to achieve a technical result. In practice it may be considered obvious to consult 'domain experts' to gather the data.

Step 3: Select a type of machine learning model

While patentability is possible here, it will likely be important to tie the use of a model to a particular technical process. One should be careful of providing an enabling disclosure and linking any claims to the particular technical process.

Step 4: Gather the data

It may be possible to demonstrate an invention in data gathering if there was a particular technical constraint involved. The requirements specification analogy used by the EPO could be useful here; that is, what tasks would the business person ask the engineer to perform? If there is a technical challenge being addressed to meet the requirements set by the business person then this may indicate the presence of a technical invention.

Step 5: Train the model

This is likely to be considered the application of a model and is likely to be considered routine if there are no unique aspects to the application implementation.

Step 6: Use the trained model to make predictions

Technical character may be conferred by this step of the process. There will be no invention for the use of a trained model for what it is intended for but the technical purpose of a method is important for determining if there is an invention. For example, use of model output to control technical process X.

IDENTIFYING THE INVENTION

It is clear from the above that identifying patentable inventions will not be straightforward where inventions involve the application of machine learning techniques. The following section sets out pointers which splits the process of implementing a machine learning process into 6-steps:

Step 1: Identify the problem you want to solve

Step 2: Decide upon the data you need

Step 3: Select a type of machine learning model

Step 4: Gather the data

Step 5: Train the model

Step 6: Use the trained model to make predictions

Step 1: Identify the problem you want to solve

To find a patentable invention the applicant will need to show that they have identified a problem that is not obvious to solve with machine learning. Since machine learning is versatile it is likely these

DRAFTING CONSIDERATIONS

Attempting to describe all the considerations that go into drafting the 'perfect' machine learning patent application will probably be the subject of a comprehensive text book to be updated monthly as practice changes and technology moves on. However, we've tried to set out a few points to look out for — by no means an exhaustive list.

Clarity

One of the main challenges the EPO identified in their 2018 conference is clarity due to the prevalence of buzzwords. Poorly drafted patent applications are likely to fail if they merely amount to a sprinkling of buzzwords. As an example, identifying the scope of a claim reciting "applying deep learning" will be difficult and impractical — is a neural network implicitly essential to this claim or is the use of classifiers excluded?

Location of infringement

As can be seen from the 6 steps above, a machine learning process can be categorised into discrete sub-processes. It is common for those sub-processes to be carried out in a distributed manner and potentially across borders. For example, data might be gathered by a smartphone in the UK and sent to a cloud server to apply a model to the data at a server in Ireland, where the model was previously trained in the US. In a claim to the whole process steps would be carried out in different jurisdictions. Wherever possible claims should be drafted to cover steps carried out by discrete entities as well as methods and computer program products. The EPO has also indicated that it may object to the clarity of claims that do not explicitly identify where a particular operation is performed within a distributed system, increasing the importance of considering infringement location up-front when drafting a specification.

Technical effect/output/purpose

To infer technical character it may be necessary to include a technical purpose into a claim but this may unnecessarily limit an invention to a particular use where it is potentially more widely applicable. Similarly the specification is likely to need to be clear about how the algorithm could be used for a technical benefit and how the output is used. Some applicants may be reticent to include this information, but practitioners will need to push back on this to ensure that they can draft a comprehensive description.

Failure to include this information will typically result in an inability to support an inventive step, as both the UK IPO and EPO will dismiss

algorithms not tethered to a particular technical purpose. When claim drafting, some time should be taken to drafting an appropriate claim preamble that includes a purposive statement to frame the algorithm in a technical context, with support in the description (both EPO and UK IPO) and explicit claim language (UK IPO, EPO optional) linking the output of the algorithm to achieving the technical purpose set out in the claim preamble.

In early 2021 the EPO provided further guidance on how technical character should be inferred through context of the invention. Though this was in relation to simulations, many of the lessons are still useful when drafting an ML application.

The EPO clarified that it is the context of a simulation that is most important when considering technicality. For example, simulating a weather system would not be technical if used to forecast a financial product, but could be if the outcome of that same simulation was used to automatically open and close window shutters on a building. We note that it is not necessary to directly claim the full context (i.e. steps defining controlling shutters as a result of the forecast output) but the purpose of the claim should be limited and at least be implicit from the wording.

This is similar to the current approach taken for ML applications and an indication that the EPO is happy to continue with it - a clear technical purpose is still essential.

In the case of an algorithm applied in a field having some non-technical considerations, it is recommended to emphasise the technical aspects of the invention. For example, detection of credit card fraud using ML can be described as an ML algorithm configured to detect unauthorised network events, with the output being some automatic response such as raising an alarm to a network administrator.

We also recommend not to rely on an output that is subject to user interpretation. The EPO in particular is not receptive to inventions that produce an output that is reviewed by a human before a decision is taken, even where the human is able to make a more informed decision that can result in an advantageous technical effect.

In the context of ML this means providing support for the ML output being automatically processed by a machine and a subsequent action being taken. At a minimum, we recommend that applications claim and describe the automatic generation of an alarm or similar as a result of the ML output, such that an argument of improved alerting of a user can be made.

Enablement/plausibility

Continuing the theme, it may be important to detail aspects of the algorithm that some applicants may prefer to keep secret. In some circumstances it may be enough to simply say “inputting the data to an artificial neural network”, as long as the skilled person would understand how to put the wider invention into practice. However, it is important to consider the level of detail disclosed carefully as more detail is likely to be required depending on the interplay between the ML algorithm and the application in question.

Neither the EPO nor the UK IPO is particularly receptive to arguments that ML inventions reduce processing time, consume less processing resources, use less memory, and the like. The EPO in particular has adopted the position that such improvements are relative to other algorithms for performing the same task, and it is always possible to conceive of a less efficient algorithm over which the subject invention will always offer an improvement. Inventive step arguments resting on these advantages are often deemed implausible and unconvincing by EPO examiners. Therefore, if seeking to rely on such improvements for technical character, consider including supporting data in the specification.

If the invention only achieves the technical effect when executed by a particular piece of hardware, e.g. a GPU rather than any generic processor, ensure that the claims reflect this fact².

Pseudo-code and/or detailed mathematics

We are seeing increased approval from patent offices of the inclusion of pseudo-code and/or detailed mathematics into applications. There are significant pros and cons of doing so and it is important to balance the disclosure requirements, the potential for limitation and the release of trade secrets when considering how much to include in an application. A more detailed discussion on this debate is out of scope here.

Detectability

One may want to consider carefully how possible it will be to detect the use of an invention by third-parties in filing strategies. For example, if the invention is in the way the model is trained and only a model or coefficients are shared with third parties, will it be possible to reverse engineer how the model was trained?

Replicability

One may want to consider how the model is trained when preparing an application. Hypothetically, where an invention is trained once on a training dataset and then a set of coefficients or a model applied repeatedly, if the invention lies in how the model is trained, is the value of the application affected?

Physical Hardware

Practice in computer implemented inventions has changed dramatically over the years when it comes to the amount of hardware that must be described and/or claimed in an application. This is true of machine learning applications, particularly if one considers the distributed nature of many implementations. It is important to balance the need to include hardware into a claim and the likely implementation.

Lengthy recitations of conventional computer hardware in the description or claims of an application will not assist in obtaining an inventive step at the UK IPO or EPO. If relying upon a technical effect tied to particular hardware, it is recommended to ensure that the details of this hardware that give rise to the technical effect are clearly described in the context of their interaction with the ML algorithm.

²(T 1125/17; see in particular Reasons 6.5 & 6.6.)

CONCLUSION

Providing effective advice to applicants using machine learning is rapidly becoming an essential part of a patent attorney's skillset and we can see no evidence of that trend abating.

The realities of machine learning mean that patent protection cannot always be a silver bullet. While copyright, trade secrets and database rights are out of scope of this article, it is important to mention that these are fundamentally important assets to businesses utilising machine learning techniques. The EU trade secrets directive and copyright protection afforded to software implementations are not to be overlooked. Maintaining proper version control, repositories, documentation and a trade secrets register is absolutely essential to businesses.

Hopefully we have provided some useful tips and pointers to enable discussion. While the approaches of patent offices around the world is changing and settling there are common threads running through the approach of each. We are always on hand to discuss any issues further.

Going forward we look forward to the harmonisation efforts and a settled approach as the law gets tested by applications over time and by applicants willing to test the boundaries.

GET IN TOUCH

If you would like to know more about how our expertise in the artificial intelligence/machine learning sector can benefit your business then please get in touch.

We welcome enquiries via computertech@gje.com and will be delighted to explain in more detail why you should consider GJE for your IP needs.



Gill Jennings & Every LLP



The Broadgate Tower
20 Primrose Street
London EC2A 2ES
T +44 (0)20 7655 8500
F +44 (0)20 7655 8501
gje.com