

UKIPO Case Studies Open the Black Box on Examining Patent Applications Relating to AI Inventions

On 22 September 2022, the United Kingdom Intellectual Property Office (IPO) published Enhanced Guidance on Examining Patent Applications Relating to Artificial Intelligence (AI) Inventions (referred to hereafter as the Guidance). The Guidance may be seen as an effort to address the concerns of certain respondents to the IPO's Call for Views on AI and intellectual property (IP) earlier in the year, many of whom expressed that it is difficult to predict the outcome of IPO decisions on patentability of AI inventions, and that more favourable results are often obtained at the European Patent Office (EPO).

Along with the Guidance, the IPO published 18 case studies each involving a brief description and claim to an AI invention, along with reasoning on whether the AI invention is excluded from patentability. For this analysis, the IPO distinguished between "applied AI inventions" and "core AI inventions".

Summary of Principles

The primary obstacles to the patentability of AI inventions in the UK are the exclusions under Section 1(2) of the Patents Act, and in particular those relating to mathematical methods and programs for computers as such. The Guidance broadly confirms that these exclusions are applied for AI inventions in the same way as for any other computer-implemented invention. Put simply, a computer-implemented invention is no more likely, or less likely, to fall foul of the exclusions by virtue of involving an AI technique as

opposed to, say, a rules-based algorithm performing an equivalent function.

p2

The relevant test as to whether the exclusions bite is the “technical contribution” test established by the Court of Appeal in *Aerotel/Macrossan*, along with the five “signposts to a relevant technical effect”, set out in their current form in *HTC v Apple*, indicating that the following factors can be used to determine whether computer-implemented invention makes a technical contribution:

- 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.

Applied AI Inventions

The IPO defines an applied AI invention as one in which an AI technique is applied to a field other than the field of AI. An applied AI invention may perform a process or solve a problem lying outside the computer on which it runs, or may perform a process or solve a problem relating to the internal workings of the computer itself.

From the case studies, it can be inferred that an applied AI invention is unlikely to be excluded from patentability if it involves an AI algorithm processing measurements of a physical entity, separate from the computer on which the AI algorithm is run, in order to provide information about that physical entity or to generate a control signal. Case studies falling into this category include (1) an Automatic Number Plate Recognition (ANPR) system, (2) monitoring a gas supply for faults, (3) analysing and classifying movement from motion sensor data, (4) detecting cavitation in a pumping system, (5) controlling a fuel injector in a combustion engine in dependence on measurements of engine characteristics, and (6) measuring percentage of blood leaving a heart. In all of these cases, signposts (i) and/or (v) come to the rescue.

On the other hand, an applied AI invention is likely to be excluded from patentability if it involves processing data that is not derived from a physical system (generally corresponding to what the EPO refers to as administrative data). Case studies in this category include (7) automated financial instrument trading, and (9) identifying junk email based on its semantic content. Furthermore, little hope is held for AI inventions which perform a purely administrative function, such as (8) analysing patient health records to group patients into risk groups. Although in this example the input data may include measurements of physical entities (patients), the invention is not directed towards specific processing of these measurements, but instead towards generic processing that is agnostic to the physical nature of the data (in contrast with case studies (1) to (6) above). Such inventions are likely to be excluded as programs for a computer as such (and potentially also as business methods as such).

Those familiar with IPO practice in relation to computer-implemented inventions will probably consider most, if not all, of the above examples to be relatively clear-cut and predictable. Murkier waters are encountered when applying the exclusions to an applied AI invention that performs a process or solves a problem relating the internal workings of a computer. While many would not be surprised that the example of (10) cache management using a neural network manages to escape exclusion via signposts (ii), (iv) and (v), two further case studies are presented in which the exclusions are avoided, and these warrant careful consideration:

Case study (11): continuous user authentication. In this example, a determination is made of whether a current user of a computer is a malicious user, based on a comparison between behaviour scores calculated by a machine learning model for the current user and a known genuine user. The Guidance describes this functionality as monitoring the internal workings of the computer, and characterises the detection of a malicious user as solving a technical problem lying within the computer system. Therefore, the invention escapes exclusion by virtue of signposts (ii) and (v).

Case study (12): virtual keyboard with predictive text entry. In this example, a machine learning model predicts and ranks next words to be typed by a user of a virtual keyboard on a touch-screen device. The predicted words are displayed on the touch screen to alleviate the burden of manually typing the words. The Guidance indicates that this solves a technical problem concerning the operation of the device by improving the speed and accuracy of text entry, which in turn leads to a more efficient and effective device. Therefore, the invention escapes exclusion by virtue of signposts (iv) and (v).

Case studies (11) and (12) are instructive, and suggest that an applied AI invention that solves a problem relating to the internal functioning of the computer has a good chance

of avoiding exclusion from patentability, where the definition of “the internal functioning of the computer” may be applied quite generously. It is worth noting that in the case studies presented as positive examples, the applied AI inventions are not limited to specific software applications (as may be the case for e.g. identifying junk email), but improve more generic aspects of the computing system, namely memory usage, security, and a user interface.

Core AI Inventions

The IPO defines a core AI invention as an advance in the field of AI itself (e.g. an improved AI model, algorithm, or mathematical method).

Put bluntly, core AI inventions are much more likely to be excluded from patentability than applied AI inventions. In particular, innovation which results in an improved AI algorithm or method, irrespective of the hardware on which it is run, will likely run into trouble. Case studies which did not escape the exclusions involve (13) optimising a neural network, (14) avoiding unnecessary processing of data by a neural network, and (15) active training of a neural network.

By contrast, inventions in which specific hardware is leveraged to facilitate or improve the functioning of an AI algorithm, may avoid exclusion. Case studies in this category include (16) processing a neural network on a heterogeneous computing platform, (17) a special purpose processing unit for machine learning computations, and (18) a multiprocessor topology adapted for machine learning. In these examples, signpost (iii) may be invoked, along with possibly signpost (iv) if the innovation represents an improvement in the efficiency of existing hardware (there appears to be an error in the IPO’s publication here, which points to signpost (v) instead).

Sufficiency

The Guidance briefly discusses the extent to which a dataset used in training an AI invention must be disclosed to satisfy the requirement of sufficiency of disclosure in the UK. In summary, an application relying upon features of a dataset to achieve a technical contribution should teach the details of the dataset in such a manner that the invention can be worked across the entire scope of the claim without undue burden. We expect that this can be achieved either by identifying at least one suitable publicly-available dataset, or by disclosing a procedure for collecting data to generate a suitable dataset. In the latter case, the burden of actually collecting the data would probably not be considered an undue burden, whereas having to guess or work out how to collect the data may well

be an undue burden.

p5

Conclusions

Unsurprisingly, one thing that remains unchanged is the elusive nature of the word “technical”. Nevertheless, the IPO’s case studies will undoubtedly prove useful in assessing whether to file a UK patent application for an AI invention, and if so in arguing for the patentability of the invention. Applicants will be on firmer ground to argue for patentability if they can draw analogy to one or more of the case studies, suggesting that the case studies could end up playing an important role in prosecution similar to the way that the Subject Matter Eligibility Examples published by the United States Patent and Trademark Office (USPTO) can play an important role in the examination of US patent applications. The case studies are particularly illuminating in relation to AI inventions addressing problems relating to the internal functioning of a computer.

It is likely to remain difficult to obtain patent protection for core AI inventions in the UK, other than in the rare situations where the invention leverages a specific arrangement of hardware. In many instances, we do not see this to be a problem because the value in protecting core AI using patents at all is debatable, given the speed and incremental nature of advances in the field. If patent protection is desired, then the EPO may be a more favourable route provided that some innovative technical effect can be identified.

As an interesting aside, the Guidance suggests that an AI invention claimed as a hardware-only implementation (i.e. not relying on a programmable device), would automatically be exempt from exclusion under Section 1(2). This could provide an alternative route for protecting AI inventions (core or applied), if the reality of a hardware-only implementation has any commercial value.