



PATENTS ACT 1977

APPLICANT	Google LLC
ISSUE	Patents Act 1977, section 1(2)(c)
HEARING OFFICER	Peter Mason

DECISION

Introduction

- 1 This decision concerns application GB 1720784.6, published as GB 2566764 A on 27th March 2019. The application is entitled “Assessing accuracy of a machine learning model” and the decision concerns whether the invention, as defined in the claims, is excluded from patentability under Section 1(2)(c) of the Patents Act 1977.
- 2 The application is the GB national phase of international application PCT/US2017/014783 which has an earliest priority date of 30th December 2016. There have been several rounds of correspondence, but the applicant has thus far been unable to persuade the examiner of the patentability of the claims. In their letter of 20th June 2022, the examiner invited the applicant to request a hearing to settle the issue. Although no request was forthcoming, unconvinced by the ensuing amendments and arguments, the examiner has forwarded the application for a decision by a hearing officer to be issued on the papers on file. I will therefore make a decision based on the papers presently available on file.

Preliminary matters

- 3 The only substantive matter before me is whether the invention is excluded from patentability under section 1(2)(c) of the Patents act 1977. I note however that the search is incomplete and all other matters, including novelty and inventive step, have been deferred . Therefore, if I find that the claimed invention is allowable I will return the application to the examiner to update the search and complete the examination.
- 4 The Section 20 date expired 9th April 2023.

The invention

- 5 The invention relates, generally, to the field of machine learning models wherein the machine learning model receives a request and produces a decision. Prior art machine learning models typically determine a decision based, at least in part, on previous decisions. Therefore, depending on the value of those previous decisions,

the quality of any subsequent decision may tend to be increasingly better or increasingly worse.

- 6 The invention attempts to ensure that the machine learning model does not iterate away from producing high quality decisions by comparing the machine learning output with a baseline assurance value. The baseline assurance value is achieved by comparing the output of the machine learning model and an output of a statistical model each having a similar input, wherein the result of the comparison influences subsequent input into an updated machine learning model.
- 7 The claims have been amended since filing and are now as presented, as filed on 21st October 2022. There are seven independent claims; 1, 10, 16, 22, 27, 33, 39. Claims 1, 10 and 27 relate to a computer implemented methods, whilst claims 16, 22 and 33, 39 relate to a computer system and a non-transitory computer readable medium with respect to the method claims 1, 10 and 27. The method claims 1, 10 and 27 are recited below, the system and non-transitory computer readable medium claims are substantially identical to the following method claims.

Claim 1 A computer-implemented method of assessing accuracy of a machine learning model, comprising:
receiving multiple requests for a first content item;
assigning a first proportion of the received requests for the first content item to a first group, and
assigning the remaining ones of the received requests for the first content item to a second group;
for a request assigned to the first group, using a machine learning model to select, from multiple candidate variants, a variant of the first content item for presentation in response to the request;
for a request assigned to the second group, using a statistical model to select, from the multiple candidate variants, a variant of the first content item for presentation in response to the request; and
obtaining information about the accuracy of the machine learning model by comparing performance data for variants of the first content item selected by the machine learning model with performance data for variants of the first content item selected by the statistical model,
determining that performance data for variants of the first content item selected by the machine learning model indicates a performance level that is below a benchmark threshold; and
responsive to the determination, assigning a second proportion of received requests for the first content item to the first group, the second proportion being lower than the first proportion.

Claim 10 A method comprising:
receiving, by a data processing system including one or more servers, a request from a client device for a content item;
assigning, by the data processing system, the received request to an audience pool based on a distribution factor, the audience pool selected for assignment from a set of audience pools including a first audience pool that receives content parameterized using a machine learning model and a

second audience pool that receives content parameterized using a statistical model;

selecting, by the data processing system, a content variation responsive to the received request using a selection mode corresponding to the assigned audience pool;

transmitting, by the data processing system, the selected content variation to the client device;

obtaining, by the data processing system, an acceptance indicator representing acceptance of the delivered content variation; and

updating, by the data processing system, the distribution factor based on the obtained acceptance indicator and the assigned audience pool,

determining that the obtained acceptance indicator, combined with previously collected acceptance data, indicates that an acceptance rate for content selected by the machine learning model is lower than an acceptance rate for content selected by the statistical model; and

updating the distribution factor to direct a decreased percentage of new requests to the first pool that receives content parameterized using the machine learning model.

Claim 27 A method comprising:

receiving, by a data processing system including one or more physical servers, a first plurality of data requests including a first data request from a first client device and a second data request from a second client device;

assigning, by the data processing system, the first data request to a first response pool selected by the data processing system from a plurality of response pools based on an allocation policy that allocates a first portion of the plurality of data requests to the first response pool and a second portion of the plurality of data requests to a second response pool, the first response pool corresponding to a machine learning model and the second response pool corresponding to a statistical model;

responding, by the data processing system, to the first data request with a first response selected by the data processing system using the machine learning model;

assigning, by the data processing system, the second data request to the second response pool selected by the data processing system from the plurality of response pools based on the allocation policy;

responding, by the data processing system, to the second data request with a second response selected by the data processing system using the statistical model;

obtaining, by the data processing system from the first client device, first feedback information indicating a performance level of the machine learning model;

obtaining, by the data processing system from the second client device, second feedback information indicating a performance level of the statistical model;

comparing, by the data processing system, an aggregate performance level of the machine learning model determined using the first feedback information to an aggregate performance level of the statistical model determined using the second feedback information;

updating the allocation policy, by the data processing system responsive to the comparison;

receiving, by the data processing system, a second plurality of data requests subsequent to updating the allocation policy; and assigning, by the data processing system, a first subset of the second plurality of data requests to the first response pool and a second subset of the second plurality of data requests to the second response pool in accordance with the updated allocation policy,

determining, by the data processing system, that the aggregate performance level of the machine learning model is less than the aggregate performance level of the statistical model by at least a threshold amount; and

updating, by the data processing system, responsive to the determination, the allocation policy to decrease the first portion of requests allocated to the first response pool corresponding to the machine learning model

- 8 The first invention defined by claim 1, the second invention defined by claim 10, and the third invention defined by claims 27 not only differ in language but also differ in scope. The examiner has deferred consideration of plurality.

The law

- 9 The examiner raised an objection under Section 1(2) of the Act that the invention is not patentable because it relates to one or more categories of excluded matter. The relevant provisions of this section of the Act are shown with added emphasis below:

1(2) It is hereby declared that the following (amongst other things) are not inventions for the purpose of the Act, that is to say, anything which consists of

(a)

(b)

(c) a scheme, rule, or method for performing a mental act, playing a game or doing business, or a program for a computer;

(d)

but the foregoing provisions shall prevent anything from being treated as an invention for the purposes of the Act only to the extent that a patent or application for a patent relates to that thing as such.

- 10 The assessment of patentability under Section 1(2) is governed by the judgment of the Court of Appeal in *Aerotel*¹, as further interpreted by the Court of Appeal in *Symbian*². In *Aerotel*, the court reviewed the case law on the interpretation of Section 1(2) and set out a four-step test to decide whether a claimed invention is patentable:

(1) Properly construe the claim;

¹ *Aerotel Ltd v Telco Holdings Ltd and Macrossan's Application* [2006] EWCA Civ 1371; [2007] RPC 7

² *Symbian Ltd v Comptroller-General of Patents* [2009] RPC 1

(2) identify the actual contribution;

(3) ask whether it falls solely within the excluded subject matter;

(4) check whether the actual or alleged contribution is actually technical in nature.

11 The Court of Appeal in *Symbian* made it clear the four-step test in *Aerotel* was not intended to be a new departure in domestic law; it was confirmed that the test is consistent with the previous requirement set out in case law that the invention must provide a “*technical contribution*”. Paragraph 46 of *Aerotel* states that applying the fourth step of the test may not be necessary because the third step should have covered the question of whether the contribution is technical in nature. It was further confirmed in *Symbian* that the question of whether the invention makes a technical contribution can take place at step 3 or step 4.

12 The case law on computer implemented inventions has been further elaborated in *AT&T/CVON*³ which provided five helpful signposts to apply when considering whether a computer program makes a relevant technical contribution. In *HTC v Apple*, Lewison LJ reconsidered the fourth of these signposts and felt that it had been expressed too restrictively. The revised 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 make the computer a better computer in the sense of running more efficiently and effectively as a computer; and

v) whether the perceived problem is overcome by the claimed invention as opposed to merely being circumvented.

13 The relevance of the legislation and legal precedent has gone uncontested throughout the proceedings.

Applying the *Aerotel* test

Step 1 - Properly construe the claim

14 There has been no contention in regard to how the examiner has construed the claims, which I have repeated below;

³ *AT&T Knowledge Venture/CVON Innovations v Comptroller General of Patents* [2009] EWHC 343 (Pat), paragraph 8.

4. *Claim 1 defines a method of assigning a proportion of received requests for a content item to a group where a variant of the content item is selected using a machine learning model, and the remaining requests are assigned to a group where a variant of the content item is selected using a statistical model. Performance data for the variants selected in both groups are compared. If it is determined that the performance of the variants selected by the machine learning model is below a threshold value, then the proportion of requests assigned to the machine learning group is lowered.*

5. *Claim 10 defines a method of assigning a request for a content item into a pool that receives content parametrised using a machine learning model or into a pool that receives content parametrised using a statistical model. The assignment is based on a distribution factor. A content variation is selected using rules corresponding to the respective pool and transmitted to the user. An acceptance indicator is received from the user, and the distribution factor is updated using the acceptance indicator and assigned pool. If it is determined that an acceptance rate for content selected by the machine learning model is lower than an acceptance rate for content selected by the statistical model, then the distribution factor is updated to direct a lower proportion of requests to the machine learning pool.*

6. *Claims 16 and 22 respectively define the system and the computer-readable medium holding instructions that perform the method of claim 10.*

7. *Claim 27 defines a method of assigning one request for data into a pool that responds according to a machine learning model and another request for data into a pool that responds according to a statistical model. The assignment is based on an allocation policy. Responses to the requests are selected, and feedback is received relating to each response. The performance level of the machine learning model is compared to the performance level of the statistical model using the feedback, and the allocation policy is updated based on the comparison. Further requests for data are assigned to the pools based on the updated allocation policy. If it is determined that the aggregate performance level of the machine learning model is less than the aggregate performance level of the statistical model by at least a threshold amount, then the allocation policy is updated to direct a lower proportion of requests to the machine learning group.*

8. *Claims 33 and 39 respectively define the system and the computer-readable medium holding instructions that perform the method of claim 27.*

- 15 The examiner has observed that, whilst different terminology is used, there is no difference between the term 'groups' used in claim 1 and 'pools' used in claims 10 and 27. Similarly, there is no difference between the selected portion of claim 1, and 'distribution factor' or 'allocation policy' of claims 10 and 27 respectively. On this I agree, and I note that there has been no contention on this matter from the applicant.
- 16 Despite their length, and complexity, I find the independent claims clear, and I have no difficulty construing them. The claims clearly have overlapping scope and whilst I generally agree with the construction given by the examiner I consider each independent claim to be further distinguished by the following;

Claim 1 applies the allocation policy to a limited dataset input and uses values determined using that dataset to update the allocation policy to that dataset.

Claim 10 updates the allocation policy based on historical data and applies the allocation policy to subsequent data inputs.

Claim 27 applies an allocation policy to a first and second plurality of data requests originating from a first and second client device respectively, and the feedback received in respect to the performance of the machine learning model and performance of the statistical model originates from the first and second client device respectively.

- 17 Claims 16 and 22 respectively define a system and a computer-readable medium holding instructions that require all the features of claim 10. Similarly, claim 33 and 39 respectively define a system and a computer-readable medium holding instructions that require all the features of claim 27. The examiner has primarily found objection in claim 1 and has observed, throughout their correspondences, that a similar reasoning applies to later claims *mutatis mutandis*. There is no contention from the applicant and throughout their correspondences both the applicant and the examiner have considered the seven independent claims together. I see no reason to depart from this approach, particularly in light of the similarity between the claimed subject matter of each group of claims.

Step 2 – Identify the actual contribution

- 18 In paragraphs 43 and 44 of *Aerotel*, Jacob LJ outlined some factors to consider when identifying the contribution made by the claims:

*The second step – identify the contribution – is said to be more problematical. How do you assess the contribution? Mr Birss submits the test is workable – it is an exercise in judgment probably involving the **problem said to be solved, how the invention works, what its advantages are**. What has the inventor really added to human knowledge perhaps best sums up the exercise. The formulation involves looking at substance not form – which is surely what the legislator intended.*

- 19 Paragraph 43 of *Aerotel* suggests that the contribution can be assessed from the point of view of the problem to be solved, how the invention works and what the advantages are, stating “*What has the inventor really added to human knowledge perhaps sums up the exercise*”. Knowledge of the prior art plays a role in assessing the contribution, and as Lewison J noted⁴, the examiner should have some notion of the state of the art. This does not necessarily mean however that the contribution is defined by what is new and inventive in the claim. That said, proper consideration of the prior art has been shown to assist in establishing the contribution with regard to what is deemed common general knowledge or conventional.
- 20 The examiner sets out in the pre-hearing hearing report what they contend the contribution be;

⁴ AT&T Knowledge Venture/CVON Innovations v Comptroller General of Patents [2009] EWHC 343 (Pat), paragraph 8.

15. *The contribution of claim 1 is:*

A method of assigning a proportion of requests for a content item to a group where a content item variant is selected using a machine learning model; assigning the remaining requests to a group where a content item variant is selected using a statistical model; comparing the performance data of variants selected in the two groups; determining if the performance of the variants selected by the machine learning model is below a threshold value, and lowering the proportion of requests assigned to the machine learning group if it is.

16. *The contribution of claims 10, 16 and 22 is:*

A method of assigning, using a distribution factor, a request for a content item into either a group that receives content parametrised using a machine learning model, or into a group that receives content parametrised using a statistical model; selecting a content item variant from the relevant group and transmitting it to a user; receiving feedback from the user and using the feedback to update the distribution factor; determining if an acceptance rate for content selected by the machine learning model is lower than an acceptance rate for content selected by the statistical model, and updating the distribution factor to direct a lower proportion of requests to the machine learning pool if it is.

17. *The contribution of claims 27, 33 and 39 is:*

A method of assigning, using an allocation policy, a request for data into a group that selects data using a machine learning model, and another request for data into a group that selects data using a statistical model; transmitting responses to the requests to users; receiving feedback from the users; using the feedback to update the allocation policy; using the updated allocation policy to assign further requests for data to one of the two groups; determining if a performance level of the machine learning model is less than a performance level of the statistical model by a threshold amount, and updating the allocation policy to direct a lower proportion of requests to the machine learning group if it is.

21 The applicant appears to set out what they consider to be the contribution in their letter dated 10th April 2022 which reads;

“In summary, the claimed method assigns content item requests to a first group or a second group. For requests assigned to the first group, a variant of the requested content item is selected using a machine learning model. For requests assigned to the second group, a variant of the requested content item is selected using a statistical model. The selected variant is delivered to the client device and information about the accuracy of the machine learning model is obtained by comparing performance data for variants selected by the machine learning model with performance data for variants selected by the statistical model.”

22 With regard to the contribution set forth by the applicant; in claim 1 a variant of a first and second content item *suitable* for presentation is selected, and performance data

for each content item is compared. There is no additional restriction relating to the delivery of either variant to a client device. However, this limitation is clearly presented in respect to claims 10 and 27. To my mind the contribution alleged by the applicant is not valid across all claims.

- 23 Additionally, the examiners assessment of the contribution of claim 1 is silent with respect to delivering the selected variant to the client device, however it seems that they apparently acknowledge this feature in their assessment regarding claims 10 and 27. Furthermore, both examiner and the applicant disregard the restriction of claim 27 wherein the first and second data requests are received from a first and second client device, and feedback pertaining to the performance level of the machine learning model and statistical learning model are respectively received from the first and second client device.
- 24 That said, the examiner adeptly argues that the claimed computing hardware, including the data processor/computer implementing the respective methods and the manner in which the hardware is networked is entirely conventional. Consequently, the examiner is explicit that the hardware and the standard communication between the hardware do not form part of the contribution. There is no contention on this point from the applicant therefore I am inclined to accept this.
- 25 Therefore, I understand the contribution common to claims 1, 10 and 27 to be;

A computer implemented method to assess the accuracy of a machine leaning model comprising receiving at least one request for a content item; assigning the at least one request to a first group or a second group wherein the first group returns a first variant using a machine learning model and the second group returns a second variant using a statistical model; comparing performance data of the first and second variant in order to influence the assignment of requests for content items between the first group and the second group.

- 26 The additional contribution provided by each independent claim is as follows;

Claim 1 includes receiving multiple data requests and assigning the received data request using the comparative data.

Claim 10 includes transmitting a selected content variation, receiving feedback on the selected content variation, and assigning future data request using the comparative data based on said feedback and historical feedback.

Claim 27 includes receiving at least two data requests and assigning at least one data request to the first group, and at least one data request to the second group; receiving feedback on a first content variant from the first group, and receiving feedback on a second content variant from the second group; and assigning future data request using the comparative data based on said feedback and historical feedback.

Steps 3 and 4 Ask whether it falls solely within the excluded matter and check whether the actual or alleged contribution is actually technical.

- 27 The third and fourth steps of the Aerotel test involve considering whether the contribution falls solely within excluded categories, and then checking whether the contribution is technical in nature. It is appropriate to consider these two steps together because whether the contribution is technical in nature will have a direct impact on whether it falls solely within excluded matter.
- 28 Although the contribution is implemented using a computer program running on a network of computers, that does not mean that it should immediately be excluded as a computer program as such. In *Symbian*⁵, the Court of Appeal stated that a computer program may not be excluded if it makes a technical contribution. In order to determine if the contribution is technical in nature I will consider the AT&T signposts as set out in paragraph 12 above. That said, the applicant has not provided any argument with respect to the second, third and fourth signpost, nor do I think it is necessary for me to consider them in this decision.

The first signpost - whether the claimed technical effect has a technical effect on a process which is carried on outside the computer;

- 29 The applicant argues that the present invention may be used to quickly detect whether the machine learning model is no longer performing effectively, consequently this alleviates irrelevant data being generated by the machine learning model which reduces any necessary repeat requests for content or data. Consequently the applicant further argues that the invention, due to the continual validation of output, reduces network traffic. The applicant maintains that this reduction in network traffic demonstrates that the technical effect is on a process outside the computer and therefore the first signpost is met.
- 30 The contribution concerns a computer system potentially including a data processing system and client devices arranged on a network, and how the data processing system receives, handles, and outputs data wherein the output is applied within the data processing system. There is clearly no technical effect in reducing repeat requests for data. Nonetheless, there may be a technical effect in reducing network traffic. However, this technical effect is clearly exclusively carried on within the computer system. Therefore this alleged technical effect is not implemented outside the computer. It is therefore not necessary for me to consider whether or not this effect is technical in nature.
- 31 I am unable to establish any other effect, technical or otherwise, on a process that is carried on outside the computer. Claims 1, 10 or 27 do not meet signpost 1.

The fifth signpost - whether the perceived problem is overcome by the claimed invention as opposed to merely being circumvented.

- 32 In *Lantana Birss J* stated “[i]t makes sense to think of something which is a solution to a technical problem as itself having technical character because it takes that character from the technical nature of the problem to be solved. But if a thing is not solving the technical problem but only circumventing it, then that thing cannot be said to have taken any technical character from the problem.”. Therefore whilst providing a solution to a technical problem an invention may receive some technicality from

⁵ *Symbian Ltd v Comptroller-General of Patents [2009] RPC 1*

that solution, if the invention merely circumvents the problem the same cannot be said.

- 33 In their latest correspondence dated 21st October 2022, the applicant argues that the problem being addressed is technical, implying that some technical merit must be assumed from the problem it is attempting to solve. The agent alleges that the solution provided by application alleviates irrelevant data being generated by the machine learning model which reduces any necessary repeat requests for content or data and reduces network traffic.
- 34 It must be noted that the problem that the application is attempting to solve is set out clearly at paragraphs [002]-[003] of the application and surmised at paragraph 5 above; here the problem is related to validation of a machine model output rather than reducing network traffic.
- 35 It seems to me that providing a system that consistently validates a machine model output requires greater network traffic, during the validation phase, than a similar system that is not consistently validated. I do accept that an increase in traffic due to repeat requests with a more periodically validated machine model output could be envisaged. However I am mindful that I have already concluded, at paragraph 30, that there is clearly no technical effect in reducing repeat requests for data. Therefore, in respect to network traffic during the validation phase, the alleged problem is not resolved, and there is no technicality that can be arrived from merely reducing repeat requests by the user.
- 36 I have found that the actual contribution falls solely within excluded subject matter and does not provide a relevant technical effect. The contribution is not technical in nature because it does not provide a contribution in a non-excluded field or overcome a technical problem. I have found that the invention does not provide the required technical contribution to satisfy section 1(2).

37 **Conclusion**

- 38 I find the invention claimed in GB 1720784.6 falls solely within matter excluded under Section 1(2) as a program for a computer as such. I therefore refuse the application under Section 18(3).

Appeal

- 39 Any appeal must be lodged within 28 days after the date of this decision.

Peter Mason

Deputy Director, acting for the Comptroller