Intellectual Property Office

BL O/0705/23

24 July 2023

PATENTS ACT 1977

APPLICANT	Capital One Services, LLC
ISSUE	Whether patent application GB2103723.9 complies with Section 1(2) of the Patents Act 1977
HEARING OFFICER	J Pullen

DECISION

Introduction

1 Patent application GB2103723.9 entitled 'Adaptive artificial intelligence systems and methods for token verification' was filed on 17 March 2021 claiming priority from two earlier US applications filed on 17 March 2020 and 8 March 2021. It was published on 10 November 2011 as GB2594789A.

2 The combined search and examination report, dated 27 August 2021, reported under Sections 17 & 18(3) that the invention, as claimed in claims 1, 17 & 18 is excluded from patentability as a method of doing business and a program for a computer as such. Subsequent rounds of communication have reiterated these objections. An offer of a hearing was made in the examiner's letter of 19 December 2022 highlighting that if the agent responded but did not request a hearing then the application may, nonetheless, be passed for a decision on the papers on file. The agent responded on 3 February 2023 with amendments and arguments, requesting a hearing with a hearing officer as a precautionary measure, but the examiner remained unconvinced, and the case was passed to me for a hearing. The examiner also wrote to the applicant to inform them of this in a letter dated 30 March 2023. The agent has subsequently filed a skeleton argument dated 12 April 2023.

3 The hearing took place on 12 May 2023 where Mr Jason Pelly of Boult Wade Tennant LLP represented the applicant. The matter to be decided is whether the claimed invention is excluded as a program for a computer and/or a method for doing business as such.

The invention

4 The application relates to the field of token verification. Tokenisation may be used to protect data, wherein "tokens" as used in the application, may include substitute data (e.g., non-sensitive data substituting for sensitive data). Tokens may have no intrinsically exploitable value and are used to protect sensitive data such as bank account or credit card numbers. Conventional token verification approaches may be

unable to detect a malicious actor possessing a genuine token and mimicking system identifiers of a legitimate system.

5 The present invention proposes an adaptive approach to verifying tokens. More particularly, the present invention relates to computer-implemented method and systems for adaptive Artificial Intelligence (AI) token verification that obtains training data for a token verification model via three different processes. The method and systems verify whether instructions to perform operations are malign or genuine based the verification model, either block or grant the request respectively, followed by updating the verification model based on the client responding from a verified client device to a verification prompt, or a challenge associated with a verification being received from a client device, or determining that no verification challenge and no verification data have been received from the client device during a predetermined time interval. As such, there are three processes of verification data acquisition, the latter allows obtaining model verification data in absence of client feedback. Updating the verification model is said to involve using any of a large number of conventional machine-learning models such as neural network, random forest, and deep learning, amongst others.

6 The current claim set, as amended 3 February 2023, comprises three independent claims: claims 1 and 16 to an adaptive token verification system, and claim 15 to a computer-implemented method for adaptive token verification. Claims 1, 15, 16 overlap in scope, for example, claim 1 requires a tokenised credit card number (whilst claims 15 and 16 require a more generalised tokenised payment number), claims 15 specifies the model output is a similarity match, while claim 16 lists several options for what the model output is, but the form of the claims is such that they will stand or fall together. Therefore, for brevity, I have included only claim 1, as currently amended, below.

7 Claim 1 reads:

> *An adaptive token verification system, the system comprising:*
> *at least one memory storing instructions; and*
> *one or more processors configured to execute the instructions to perform operations comprising:*
> *training a verification model to verify tokenized requests based on respective system identifiers of the tokenized requests;*
> *receiving a tokenized request from an external system, the tokenized request comprising a system identifier of the external system, a payment request, and a tokenized credit card number;*
> *generating output of the verification model based on the system identifier; and based on the output, performing one of:*
> *granting the tokenized request, the granting comprising retrieving an account number linked to the tokenized credit card number and processing the payment request using the account number; or*
> *blocking the tokenized request;*
> *transmitting a verification prompt to a client device associated with the tokenized request, the verification prompt indicating whether the tokenized request was granted or blocked; receiving, from the client device, verification data associated with the verification prompt; and updating the verification model based on the verification data;*

receiving a verification challenge from the client device, the verification challenge being associated with the tokenized request, and updating the verification model based on the verification challenge; and

determining that no verification challenge and no verification data have been received from the client device during a predetermined time interval, and updating the verification model based on the determination.

The Law

8 The examiner has objected that the invention is excluded from being patented as a program for a computer and a method for doing business. The relevant section of the Act is s.1(2), the most relevant provisions of which are shown below with my emphasis added:

> *1(2) It is hereby declared that the following (among other things) are not inventions for the purposes of this Act, that is to say, anything which consists of-*
> *(a) ...;*
> *(b) ...;*
> *(c)* ***a… method for... doing business, or a program for a computer;***
> *(d) ...; but the foregoing provision shall prevent anything from being treated as an invention for the purposes of this Act only to the extent that a patent or application for a patent relates to that thing as such.*

9 The Court of Appeal has said that the issue of whether an invention relates to subject matter excluded by Section 1(2) must be decided by answering the question of whether the invention reveals a technical contribution to the state of the art. The Court of Appeal in *Aerotel/Macrossan*[1] set out the following four-step approach to help decide the issue:

> *(1) Properly construe the claim;*
> *(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.*

10 The operation of the approach is explained at paragraphs 40-48 of the judgment. Paragraph 43 confirms that identification of the contribution is an exercise in judgment involving the problem said to be solved, how the invention works and what its advantages are; essentially, what it is the inventor has really added to human knowledge, looking at substance, not form. Paragraph 47 adds that a contribution which consists solely of excluded matter will not count as a technical contribution.

11 In *Symbian*[2] the Court of Appeal reaffirmed the *Aerotel* approach while considering a question of "technical contribution" as it related to computer programs emphasising the need to look at the practical reality of what the program achieved, and to ask whether there was something more than just a "better program".

---

[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

12 The case law on computer implemented inventions was 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*[4], Lewison LJ reconsidered the fourth of these signposts and felt that it expressed too restrictively. 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.*

13 The examiner's reports also refer to the guidance in *Merrill Lynch*[5], *Kapur*[6] and several Hearing Officer decisions[7].

Assessment

(1) Properly construe the claim

14 In their letter of 30 March 2023, paragraph 7, the examiner notes a possible internal contradiction within claims 1, 15, 16, it being potentially unclear whether or not the invention is required to be capable of updating the model in all of the three different ways and the examiner has adopted this construction. Mr Pelly confirms the examiner's understanding that the independent claims are intended to cover an invention capable of performing the model update in the three different ways, although for any given scenario, only one may be carried out. I agree with Mr Pelly that it is how the claims should be construed.

15 By way of a summary, the claims are considered to require receiving a request that has been tokenised for security purposes, the tokenised request being received from an external system such as a merchant's website, the request having an identifier of the external system, a payment request, and a tokenised credit card / payment number. A machine learning model trained to verify received requests is provided. Based on the output of the model, the request is either granted or blocked. In order for the model to improve its accuracy, or perhaps adapt to changing patterns, the model goes through a continual update or retraining process. This continual process then comprises three mechanisms for acquiring the training data. First – the system sending a prompt to a client device (known source of accurate information). Second – unsolicited data coming from a particular client device associated with that account (perhaps challenging the verification outcome later). Third – where no verification

[3] AT&T Knowledge Ventures/CVON Innovations v Comptroller General of Patents [2009] EWHC 343 (Pat)
[4] HTC v Apple [2013] EWCA Civ 451
[5] Merrill Lynch's Application [1989] RPC 561, 569
[6] Kapur v Comptroller General of Patents, Designs & Trade Marks [2008] EWHC 649 (Pat)
[7] BL O/291/16, BL O/221/21 and BL O/083/12

data in response to a prompt, or no verification challenge was received within a time period – that outcome is taken as an indication that the verification outcome was a correct one.

(2) Identify the actual contribution

16 In their letter of 30 March 2023, paragraph 12, the examiner identifies the contribution to be:

> A *method of verifying tokenized requests, where an adaptive model is used to decide whether to grant or block payment requests based on a system identifier in the request and where three feedback processes aim to obtain or infer accurate information about whether the decision was correct based on verification data/challenges received from a client device associated with the payment request (or a lack of within a given time period) in order to update the model, thus improving the accuracy with which future payment requests are granted or blocked.*

17 Mr Pelly is in general agreement with what the examiner identifies the contribution to be, albeit offering that the emphasis, or the main issue at play, is the gathering of the training data by three specific mechanisms, how they work together, and that training data can be gathered in absence of receiving a prompted response, or an unprompted challenge, and thus generating accurate training data in absence of direct user feedback (without which training data would otherwise likely be quite sparse).

18 I agree with Mr Pelly's assessment. What the applicant has added to human knowledge are mechanisms for collecting data to be used with conventional machine-learning models to train and update a verification model to verify tokenized payment requests.

(3) Ask whether it falls solely within the excluded subject matter and (4) Check whether the actual or alleged contribution is actually technical in nature

19 I will consider steps (3) and (4) together.

20 The examiner states that the contribution is realised as a computer program. They then go on to determine that the contribution does not solve a technical problem within the computer or have a technical effect on a technical process outside the computer with reference to the signposts discussed in *AT&T/CVON* and *HTC v Apple*. I will also take this approach.

> **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

21 In their letter of 30 March 2023, the examiner briefly discusses signposts (i)-(iii) in paragraph 21, concluding that none of the first three signposts is met. I agree, and Mr Pelly had no objection to this assessment. Instead, the emphasis of the discussion is placed on signposts (iv)-(v) below, which is where the disagreement between the examiner and Mr Pelly occurs.

iv. whether the program makes the computer a better computer in the sense of running more efficiently and effectively as a computer

22 The examiner proposes that the computer is not operating more efficiently or effectively as a computer. In their letter of 30 March 2023, the examiner proposes that the benefits of obtaining or inferring accurate training data with which to update the model to make more accurate decisions about grant or blocking payment requests are confined to the program, and that the computer itself is no better as a result of this.

23 Mr Pelly proposes that signpost (iv) is relevant. Mr Pelly proposes that the computer is a better computer because it is able to continuously adapt and update itself and converge to a more accurate system more quickly by collecting data via the three mechanisms, and that the process of collecting training data, updating the model, and using the updated model improves the verification system because the verification system is executing the model. Mr Pelly asserts that the verification system is a better verification system because it is able to effectively collect and use training data to adapt and improve itself.

24 During the hearing, on 12 May 2023, Mr Pelly reflected on this point, stating that the purpose of the system is to perform verification, and the system does that more quickly or efficiently due to a combination of the three mechanisms used to acquire training data on which to train the verification model, in particular the third mechanism that enables constant ongoing training information irrespective of whether or not a user provides feedback (both upon a prompt, or by way of a challenge later on).

25 Mr Pelly noted that signpost (iv) would be nonsense as an exclusion, if it required to look at the computer independently of what the program was doing, and insofar as it is a verification system, the computer wants to perform verification, and the program is making that verification system run more efficiently and effectively by continuously updating itself, and adapting to new and emerging data patterns (e.g. changing practice of malicious actors).

26 In *AT&T/CVON*[8] at paragraph 34, Lewison J stated:

> "*In Symbian itself, the invention was patentable because it resulted in a faster and more reliable computer. The increase in speed and reliability was not, as I understand the invention, dependent of the type of data being processed or*

[8] AT&T Knowledge Ventures/CVON Innovations v Comptroller General of Patents [2009] EWHC 343 (Pat), paragraph 34.

the particular application being used to do the processing. The invention operated at a much higher level of generality within the computer."

By contrast, in the present invention, any increase in efficiency and effectiveness is achieved at the level of the payment token verification system becoming more effective or efficient at verifying transactions (processing a very particular set of data, and either approving or rejecting requests). The invention does not make changes at a level of generality within the computer, rather any performance improvements arise directly out of the specific application and the specific data being processed, while the reliability of the computer itself is not improved, and the computer is not made a more efficient computer in a general sense. I conclude the invention does not make any changes to the underlying computer system, nor does it enable the computer to operate more efficiently as a computer. I therefore do not consider signpost (iv) to be met.

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

27 In their letter of 30 March 2023, at paragraph 18, the examiner agrees that the problem of obtaining accurate training data has been overcome (rather than merely circumvented), but proposes that the problem is non-technical, therefore concluding that signpost (v) doesn't apply. The reasoning put forward by the examiner is that obtaining accurate information with which to update a machine learning model for verifying tokenised requests and granting/blocking the associated payment request is not a technical problem. The examiner also proposes that the solution of receiving feedback on the decision from a client device, or making an assumption that the decision was correct when no such feedback is received is proposed, also has no technical effect beyond the running of a program on a computer to update a model for conducting a business task, and the problem of legitimate payment requests being blocked and illegitimate ones being granted is purely an administrative problem.

28 During the hearing on 12 May 2023, Mr Pelly suggested that the problem is how a system that has got a verification system in it, is adapted (via the aforementioned three mechanisms) so that accurate and reliable training data can be obtained quickly and effectively, proposing that it is technical, and therefore not excluded. Mr Pelly suggests that the emphasis should be placed on the system having a verification system interacting with an external system and client devices, wherein the underlying issue is how to obtain the training data in the first place based on the components of this system. This, Mr Pelly proposes, is a technical problem that needs addressing, which is addressed in a technical manner in claim 1.

29 Mr Pelly has also turned to discussion of the *AI guidance document*[9] published by the Intellectual Property Office on 22 September 2022, in particular the document outlining scenarios. The scenarios discussed are numbered as scenario 4 and scenario 9.

[9] Examining patent applications relating to artificial intelligence (AI) inventions: The Scenarios, published 22 September 2022. https://www.gov.uk/government/publications/examining-patent-applications-relating-to-artificial-intelligence-ai-inventions/examining-patent-applications-relating-to-artificial-intelligence-ai-inventions-the-scenarios.

30 Scenario 4 relates to a method of training a neural network classifier to detect cavitation in a pump system (using physical data for a specific technical purpose). Scenario 9 relates to an AI system for classifying emails into junk, not junk, or unsure categories. The training data is obtained from the user. The examiner and Mr Pelly do not agree on the relevance of the above scenarios. Mr Pelly noted that neither scenario maps well to the current invention, and I am inclined to agree. Furthermore, I note that neither of the two scenarios relates to either of signpost (iv) or (v).

31 During the hearing on 12 May 2023, Mr Pelly additionally discussed scenario 11, noting that whilst it does not relate to training as such, it involves checking whether a user is malicious based on their behaviour. The guidance document proposes it is not excluded under at least signpost (v). Mr Pelly noted that the issues and terminology are similar with the present invention, and whilst the scenario did not involve training data acquisition, it was similar to the present invention at a metalevel in that it used a model to work out malicious actors based on activities in the system, noting that the scenario indicated that a technical effect is achieved under signpost (v), and suggesting that detection of a malicious user may be enough for signpost (v) to be satisfied. Whilst scenario 11 does share some terminology with the present application, and is similar to the present application insofar as it relates to detecting malicious actors using a machine learning algorithm, the scenario does not map well to the current invention beyond that point, as it relies on monitoring characteristic usage of a computer (its internal workings) by a user, and it is that monitoring of the internal workings (such as the manner of typing, application usage, mouse movements) that is technical in nature in scenario 11. In other words, in scenario 11, it is not detection of a malicious user *per se* that meets signpost (v), but rather the particular contribution to that generic problem. By contrast, in the present invention, the particular contribution to this generic problem is different (I.e. obtaining reliable training data for a payment authorisation system, by way of user feedback, or absence thereof.

32 The aforementioned scenarios provide helpful guidance, but as emphasised in paragraph 4 of *AI guidance document*, the document is not a source of law, and the opinions on the patentability of the scenarios shall not be binding for any purpose under the Patents Act 1977. There is limited benefit to be gained by drawing analogies to different inventions, as they propose that such things may be patentable in some cases but does not show that the invention in this case is patentable. With that in mind, I need to take a step back to consider whether the invention in this case is patentable.

33 The contribution in this case concerns mechanisms for collecting data to be used with conventional machine-learning models to train and update a verification model to verify tokenized payment requests. Mr Pelly emphasised that the method provides for collecting accurate training data and that that would in turn facilitate more accurate verification of tokenised payment requests. Whilst payment verification systems can have technical character, the emphasis on collecting data about the validity of payments is a step removed from that. Any improvement in accuracy of verification of payment requests as a result of training using the collected data is a matter of hope not design. The updates to the verification model using the data collected by the described mechanisms might lead to more accurate verification of tokenised payment requests but equally they might not. Collecting data about the

validity of decisions to grant or deny tokenised payment requests is not technical in nature. I therefore consider the application to be excluded from being patented under Section 1(2) as a program for a computer as such.

Business method

34 The examiner argues that the contribution represents an activity that is also excluded as a method for doing business. Mr Pelly disagrees. However, in view of my conclusion above that the application is excluded as a program for a computer, I do not need to consider this in more detail.

Conclusion

35 I find the application to be excluded from being patented under Section 1(2) as a program for a computer as such. I therefore refuse the application under Section 18(3).

Appeal

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

J Pullen
Deputy Director, acting for the Comptroller