Intellectual Property Office

BL O/0057/25

22 January 2025

PATENTS ACT 1977

APPLICANT	SOPHOS LIMITED
ISSUE	Whether patent application GB2303441.6 complies with section 1(2) of the Patents Act 1977
HEARING OFFICER	Steven Evans

DECISION

Introduction

1 Patent application GB2303441.6, entitled “Natural language analysis of a command line using a machine learning model to generate a natural language description of the command line”, was filed on 9 March 2023. It was published on 8 November 2023 as GB 2618429 A.

2 A combined search and abbreviated examination report, dated 31 August 2023, reported under Sections 17(5)(b) and 18(3) that the invention as defined in claims 1 to 20 is excluded from patentability as a program for a computer, as such. The examiner maintained, over several rounds of examination, that the application does not comply with section 1(2) of the Patents Act 1977 (“the Act”).

3 The applicant was offered a hearing on the 7 August 2024, to which they responded with amendments and further arguments. The examiner remained unconvinced and the case was passed to me for a hearing. The examiner wrote to the applicant to inform them of this in the pre-hearing report dated 30 October 2024. The applicant subsequently requested that the decision be made based on the papers on file. I have considered all the correspondence on file in reaching this decision.

4 I note that a search has not been performed and full examination has been deferred. If I find in favour of the applicant, then the application will need to be remitted to the examiner to conduct a search and to complete substantive examination.

The invention

5 The application concerns the desire to understand computing commands, such as those input by a user to a command line interface. By way of non-limiting example, a cybersecurity analyst may need to understand what a particular command means so that they can determine whether it is malicious. However, even an experienced analyst may not be able to understand a given command to a desirable level.

6 The invention ultimately aims to provide an accurate natural language description of a command. The command is repeatedly input into a machine learning model to generate a plurality of natural language descriptions of the command. The natural language descriptions are then input into the machine learning model to generate a respective plurality of check commands. The check commands are compared to the initial command to determine which check command is most similar to the initial command. Finally, the natural language description that corresponds to the most similar check command is displayed.

7 The claims under consideration are those dated 25 September 2024, which are identical to the claims as originally filed (any amendments made during prosecution having been subsequently reversed by the applicant). There are three independent claims, which are reproduced below.

> 1. An apparatus, comprising:
>
> a memory comprising computer-executable instructions; and
>
> one or more hardware processors in communication with the memory that, having executed the computer-executable instructions, are configured to:
>
> repeatedly input, a predetermined number of times and to generate a plurality of different natural language (NL) descriptions, a command into a machine learning model;
>
> for each NL description from the plurality of different NL descriptions, input that NL description into the machine learning model to generate a check command (1) associated with that NL description, and (2) from a plurality of different check commands associated with the plurality of different NL descriptions;
>
> determine a plurality of distance values, each distance value from the plurality of distance values determined by comparing a unique check command from the plurality of different check commands to the command using a distance function;
>
> identify a selected check command from the plurality of different check commands based on a distance value from the plurality of distance values determined using the selected check command, the distance value indicating that the selected check command is more similar to the command than the remaining check commands from the plurality of different check commands; and
>
> cause display of a selected NL description from the plurality of different NL descriptions as a description of the command, the selected check command being generated based on the selected NL description.
>
> 7. A method, comprising:
>
> inputting, via one or more processors, a command into a machine learning model to generate a first natural language (NL) description;
>
> inputting, via the one or more processors, the command into the machine learning model to generate a second NL description;
>
> inputting, via the one or more processors, the first NL description into the machine learning model to generate a first check command;

inputting, via the one or more processors, the second NL description into the machine learning model to generate a second check command;

determining, via the one or more processors, a first similarity metric indicating similarity between the first check command and the command;

determining, via the one or more processors, a second similarity metric indicating similarity between the second check command and the command;

determining, via the one or more processors and based on the first similarity metric and the second similarity metric, that the first check command is more similar to the command than the second check command is to the command; and

causing, via the one or more processors, display of the first NL description, and not the second NL description, as a description of the command.

16. A processor-readable medium storing code representing instructions to be executed by one or more processors, the instructions comprising code to cause the one or more processors to:

repeatedly input a predetermined number of times, via a processor and to generate a plurality of different natural language (NL) descriptions, a command into a machine learning model;

input, via the one or more processors, the plurality of different NL descriptions into the machine learning model to generate a plurality of different check commands;

determine, via the one or more processors, a plurality of similarity metrics by comparing each check command from the plurality of different check commands to the command;

identify, via the one or more processors and based on the plurality of similarity metrics, a check command from the plurality of different check commands that is most similar to the command; and

cause, via the one or more processors, display of an NL description from the plurality of different NL descriptions, the NL description previously input into the machine learning model to generate the check command.

The law

8 The matter to be decided is that of exclusion under section 1(2) of the Act, specifically in relation to the program for a computer as such exclusion. Section 1(2) reads as follows:

> *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) a discovery, scientific theory or mathematical method;*
>
> *(b) a literary, dramatic, musical or artistic work or any other aesthetic creation whatsoever;*

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

(d) the presentation of information;

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 test for establishing whether a patent application relates to one of these excluded categories is set out in the Court of Appeal's judgment in *Aerotel*[1]. In *Comptroller v Emotional Perception AI Ltd*[2] the steps of the test were expressed as follows:

(1) Properly construe the claim;

(2) Identify the actual contribution (although at the application stage this might have to be the alleged contribution);

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

(4) If the third step has not covered it, check whether the actual or alleged contribution is actually technical.

10 In *Symbian*[3] the Court made clear that the question of whether a computer implemented invention is patentable has to be resolved by asking whether it reveals a technical contribution to the state of the art.

11 In *AT&T/CVON*[4], Lewison J (as he then was) set out five signposts that he considered to be helpful when considering whether a computer program makes a technical contribution. In *HTC*[5], the fourth signpost was reformulated slightly. I would stress that the signposts are only indicative. 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.

---

[1] *Aerotel Ltd v Telco Holdings Ltd and Macrossan's Application* [2006] EWCA Civ 1371
[2] *Comptroller General of Patents, Designs and Trade Marks v Emotional Perception AI Ltd* [2024] EWCA Civ 825
[3] *Symbian Ltd v Comptroller-General of Patents* [2009] RPC 1
[4] *AT&T Knowledge Venture/CVON Innovations v Comptroller General of Patents* [2009] EWHC 343 (Pat)
[5] *HTC v Apple* [2013] EWCA Civ 451

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.

Application of the *Aerotel* test

(1) Properly construe the claim

12 There are three independent claims, numbered 1, 7 and 16. The claims can be readily construed but there are some slight differences in terminology and scope that are worth noting.

13 Firstly, claims 1 & 16 both involve “repeatedly inputting” a command into a machine learning model “a predetermined number of times” to generate a “plurality” of different natural language descriptions. The generation of a “plurality” of check commands follows from this. In contrast, claim 7 involves inputting a command into a machine learning model two times to generate “first” and “second” natural language descriptions, with the generation of “first” and “second” check commands following. The terminology differs but there is nothing substantial to account for here. All the claims involve generating at least two natural language descriptions and their corresponding check commands.

14 Secondly, claim 1 involves determining a plurality of “distance values” by comparing each check command to the command using a “distance function” whereas claims 7 and 16 refer instead to determining “similarity metric(s)”. Paragraph 61 of the description explains that the similarity metric may be a distance value. Thus, claim 1 may be slightly narrower in that respect but nothing turns on it. The relevant point is that the claimed invention provides a way of measuring how similar the check commands are to the (initial) command so that the natural language description that was used to generate the most similar check command can be displayed.

15 Therefore, the claims will stand or fall together.

(2) Identify the actual or alleged contribution

16 The applicant and the examiner have reached agreement over this part of the test. The examiner’s pre-hearing report, sets out the contribution as follows (my emphasis added):

> “14. Again, I note we agree on the contribution of the invention and its advantages, the contribution is identified as:
>
> A method of generating natural language descriptions of [a] command (for example a command prompt command/command line) by inputting said command into a machine learning model a given number of times to generate a number of machine learning [natural language] descriptions of said command, these descriptions being passed into <u>another machine learning model</u> to generate check commands which are then compared against the input command to identify the most similar and thus find the most similar natural language description of the input command based on such for display to a user.

> "15. We agree that the advantages of the invention relate to increasing the accuracy of the natural language output to ensure that the selected command is the most accurate description of the command prompt (the one with the lowest distance value). As the command is repeatedly input to the machine learning model to generate a plurality of different natural language descriptions some being more accurate than others the invention measures the relevancy and accuracy of each in order to present the best description to the user. The use of natural language descriptions in this case allows for a less skilled user (see paragraph [0003]) to understand a given command."

17 I am content to accept this as the contribution, with two minor clarifications. In relation to the first underlined portion, it is the same machine learning model used to generate both the natural language descriptions and the check commands. That is what is defined by the wording of the claims, which only refer to a single machine learning model.[6] In relation to the second underlined portion, it appears this should read "the selected *check* command *corresponds to* the most accurate *NL* description of the command prompt". It is the check command which is most similar to the original command which is selected, and the corresponding NL description is displayed to the user.

(3) Ask whether it falls solely within the excluded subject-matter

(4) If the third step has not covered it, check whether the actual or alleged contribution is actually technical

18 The contribution plainly requires a computer program for its implementation. To determine whether the computer program nevertheless reveals a technical contribution, the examiner and applicant have both referred to the *AT&T* signposts. I will consider them in turn.

19 In relation to signpost i), the applicant notes that the invention objectively assesses the outputs of the natural language processing in order to identify the most accurate natural language description of the input command. They argue that this provides a technical effect of improved accuracy, reliability and robustness of generation of the natural language output which is provided outside of the computer. They further argue that this enables a user outside the computer to more reliably understand the input command.

20 The examiner, in their pre-hearing report, maintains their position that while they agree the above benefits are a likely outcome of working the invention, they are still unable to identify a relevant technical process carried on outside of the computer by the invention.

21 I accept that the contribution includes providing the user with what is determined to be the most accurate natural language description of the command in order to allow a less skilled user to understand that command. To the extent that reliably outputting an accurate natural language description and thus improving a user's understanding

[6] Paragraph 41 of the description discloses that multiple models could be used but this appears to be outside the scope of the claims as drafted.

can be viewed as a process that is outside of the computer, I do not believe there is any technical effect on this process.

22 The invention essentially acts as a translation tool for assisting a user without sufficient skill to interpret a command. Outputting translations of a command in a more understandable format for a user is not a technical effect. The fact that the invention uses a systematic approach of generating check commands and determining similarity measures or distance values to the initial command to ensure reliability and accuracy of the displayed natural language description does not alter this. Further, improving a user's knowledge of what a given command means by providing them with a natural language description is not a technical effect. The first signpost does not assist the applicant.

23 Signposts ii)-iv) have been considered by the examiner and found not to aid the applicant in this case. No arguments have been put forward by the applicant relevant to these signposts, and therefore I see no need to consider them in detail at this point, other than to say I am in agreement with the examiner that they are not satisfied here.

24 In relation to signpost v), it is important to clearly identify what the problem is that is being addressed. The problem, and the arguments put forward, are best understood with reference to paragraph 3 of the description:

> "Computing commands, such as a command prompt input into a command line interface, can sometimes be desirable to understand. For example, in a cybersecurity context, a cybersecurity analyst may desire to understand a computing command that was performed by a user to evaluate a potential risk associated with that computing command and/or user. For example, a cybersecurity analyst may want to determine a likelihood of maliciousness of a command line. Many cybersecurity analysts, however, do not have the knowledge to understand the computing command under analysis at a desirable level. Because commands can be complex and/or hard to decipher, even analysts having prior experience with computing commands may not have a desirable level of understanding of that command. As a result, a cybersecurity analyst may, for example, flag an event that may not need to have been flagged and/or not flag an event that should have been flagged."

25 Therefore, as the application itself sets out, the primary problem relates to a user (exemplified here as a cybersecurity analyst but equally could be a programmer, a student, a manager, etc) having insufficient understanding of what a given command means. This is not a technical problem. It is a problem of deficient user training or education in understanding what commands mean, which is a problem that is essentially non-technical in nature.

26 The applicant argues in their most recent correspondence that the problem relates to improving the accuracy of a natural language description, referring to the use of a "feedback loop" that they allege is "inherently technical".

27 This has been considered by the examiner in their pre-hearing report who maintains that, while the present system would benefit a user who is unfamiliar with the

commands, the process is merely automating what would be done by a skilled user and does not consider translating check commands into natural language to be a technical problem.

28 I agree that a further problem that the invention addresses is that of how to reliably provide an accurate natural language description of a command. However, this is a non-technical problem in my view. The advantage of improving the accuracy of the description of the check command, according to the agent's letter dated 23 September 2024, is that it "enables the analyst to have an improved understanding of the events". As discussed previously, the problem being addressed is ultimately insufficient user knowledge and/or training.

29 The "feedback loop" to which the applicant is referring is not clearly explained in the correspondence on file. Furthermore, the system as claimed does not appear to operate as a feedback loop – the output is not fed back in to the input at any point - it merely chooses the closest result from a set of results that have been generated. However, this may be a reference to the fact that the generated natural language descriptions are used to generate respective check commands, which in turn are each compared to the initial command in order to select the closest check command and ultimately display the corresponding natural language description. This process aims to ensure that the natural language description chosen for display to the user is the one which most accurately reflects the command.

30 Alternatively, the "feedback loop" may be referring to the retraining of the machine learning model using commands, check commands and/or quality metrics as input learning data and the NL descriptions as target learning data[7]. This aspect is not considered to be an essential part of your invention, nor does it form part of the identified contribution.

31 As discussed under signpost i) above, there is no relevant technical effect in such a solution. Even if I accept that the system makes use of a feedback loop, this in itself is not inherently technical. It depends entirely on what the feedback loop is used for. Here, it is alleged it improves the accuracy of a natural language description of a command for presentation to a user, which I have found is not a technical effect.

32 The applicant has also emphasised the context of cybersecurity when discussing the problem and solution. In particular, they argue that the invention provides a tool to improve the security of a system by enabling a person to understand a command, and the potential security risks associated with it, so that they can better identify malicious events.

33 This argument was addressed in the pre-hearing report by the examiner who considers that the problem does not directly contribute to the security of the computer, the computer is as secure as it was before running the allegedly better translation tool.

34 I am unconvinced by the applicant's argument. The discussion of cybersecurity in paragraph 3 of the description (quoted above) is merely illustrative of a field in which the invention might be employed. The claims which define the scope of the invention

[7] See paragraphs 35, 51, 64 of the description as filed

are not limited to cybersecurity – they are silent about what the commands relate to and about any action a user may or may not take in response to the natural language description being displayed. The contribution identified above is equally silent about any potential application to cybersecurity. For that reason, the argument that the invention is a tool to improve security and better identify malicious events is without basis. There is no contribution at that level.

35 The applicant has also argued that the invention provides users with an understanding of the computer's inner workings or internal processes, overcoming a technical problem. I do not accept that it does so in any relevant technical sense. The invention tells a user no more about the computer than the original command would tell a sufficiently skilled user. It is a tool for plugging a gap in the user's understanding of what a command means rather than a probe of what is going on inside the computer. The fifth signpost does not assist the applicant.

36 An analogy to *Lenovo*[8] has also been made by the applicant. They argue that the software of the invention "reduces the need for human intervention by eliminating the manual translation of a software command into a natural language description or assessing and ranking multiple natural language descriptions". They observe, correctly, that a computer-implemented invention may be patentable if it resolves a technical problem outside the computer by removing the necessity of a physical interaction.

37 However, in my view, *Lenovo* is not relevant here as there is no technical problem or process being solved in this instance. I would agree with the examiner that computer programs are routinely used to automate previously manual steps and that doing so does not lead automatically to a technical contribution. The fact that the invention potentially eliminates the need for a user to take further steps to translate a command or to rank potential translations is not enough to take it outside of the exclusion.

38 Stepping back, the invention provides a computer program operating as a tool for providing an unskilled user with the most accurate natural language description of a command to aid their understanding. That is nothing more than a computer program as such. The contribution falls entirely within the excluded subject matter.

Conclusion

39 I have found that the claimed invention does not comply with section 1(2) as it relates to a program for a computer as such. I therefore refuse the application under section 18(3).

[8] *Lenovo (Singapore) PTE Ltd v Comptroller General of Patents* [2020] EWHC 1706 (Pat)

Appeal

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

Steven Evans

Patent Examination Group Head