



Neutral Citation Number: [2023] EWHC 2948 (Ch)

Case No: CH-2022-000144

IN THE HIGH COURT OF JUSTICE
CHANCERY DIVISION
BUSINESS AND PROPERTY COURTS OF ENGLAND AND WALES
ON APPEAL FROM THE UK INTELLECTUAL PROPERTY OFFICE

Royal Courts of Justice, Rolls Building
Fetter Lane, London, EC4A 1NL

Date: 21/11/2023

Before :

SIR ANTHONY MANN

Between :

Emotional Perception AI Ltd **Appellant**
- and -
Comptroller-General of Patents, Designs and Trade
Marks **Respondent**

Mark Chacksfield KC and Henry Edwards (instructed by **Hepworth Brown**) for the
Appellant
Anna Edwards-Stuart (instructed by **The Comptroller-General of Patents, Designs and**
Trade Marks) for the **Respondent**

Hearing dates: 5th and 6th July 2023; further submissions received on 22nd and 27th September
2023

Approved Judgment

This judgment was handed down remotely at 10.00am on 21st November 2023 by circulation
to the parties or their representatives by e-mail and by release to the National Archives.

.....
SIR ANTHONY MANN

Sir Anthony Mann :

Introduction

1. The Patents Act 1977 section 1(2)(c) excludes from patent protection “a program for a computer ... as such”. The courts have had to grapple from time to time with the difficulties of this concept in relation to what I can call traditional computers and software. This appeal raises new questions because it involves deciding whether the use of an aspect of Artificial Intelligence, namely an Artificial Neural Network (“ANN”), in the circumstances of this case, engages the exclusion. I am told that this issue has not yet arisen in any of the authorities. A hearing officer in the UK Intellectual Property Office, Dr Phil Thorpe, decided that it did and therefore refused grant of the proposed patent in question, in a decision dated 22nd June 2022 (BL/O/542/22) (“the Decision”). The applicant for the patent, Emotional Perception AI Ltd, appeals that decision. On this appeal Mr Mark Chacksfield led for the appellant; Ms Anna Edwards-Stuart appeared for the respondent, the Comptroller General of Patents.
2. It needs to be understood that this appeal concerns the exclusion only. It does not, for example, deal with any sufficiency points which might arise, or any other questions going to validity save for briefly referring to the “mathematical method” exclusion in the same sub-section. I also point out that in this judgment I use the expression “computer program” as a synonym for “a program for a computer”, because it is shorter, but at all times I have the actual statutory wording in mind.

The field of the patent and the claims

3. In this section I shall describe the invention. In the next section I describe what an ANN is and how it works. For the purposes of understanding this section it can be envisaged as a black box which is capable of being trained as how to process an input, learning by that training process, holding that learning within itself and then processing that input in a way derived from that training and learning.
4. The applied for patent is said to provide an improved system for providing media file recommendations to an end user, including sending a file and message in accordance with the recommendation. A typical field of use is music websites, where a user may be interested in receiving music similar to another track of which he/she knows or already has. Existing websites are capable of offering similar pieces in, say, the same category (rock, heavy metal, folk, classical and so on) but the categorisation tends to be limited to types of music. The categorisation is derived from human tagging, the playlists of others and the like, which tend ultimately to derive from a human being’s

classification. The advantage of the proposed patent is said to be that it is able to offer suggestions of similar music in terms of human perception and emotion irrespective of the genre of music and the apparently similar tastes of other humans, and to arrive at such suggestions by passing music through a trained ANN which does the categorisation in that respect.

5. The two principle claims of the patent are set out in the Appendix to this judgment.
6. Claim 1 is a product by process claim, and claim 4 is the process. It was common ground that for these purposes there was no material difference between them and that I can work from either claim for the purposes of this judgment so far as necessary. The Hearing Officer worked from Claim 4.
7. Some simplified explanation is required at this stage of the way in which the claim works. Since there was no dispute about that I can reduce it to more everyday terms in the following manner without resort to the wording of the rest of the application. The invention is said to be applicable to various media, including music, images and text, but for the purposes of explanation the parties have considered music as a typical, perhaps the most likely, usage example. I shall do the same. The explanation which follows is generalised and not complete, but it is sufficient to provide the context for the issues in the case.
8. A useful starting point is the training of the ANN and to assume for the moment that the ANN itself is a hardware system (as opposed to a software emulation). A pair of music files is taken, each of which is accompanied by a natural language description of some sort of the type of music in its file in terms of how that music is perceived by a human. At its simplest the music might be described as happy, or sad, or relaxing, though the descriptions will be more complicated and wordy than that. The descriptions are in word form (hence the use of the word “semantic” and its derivatives which are used to describe this sort of feature of the music) and are to be analysed by an ANN via natural language processing software. An ANN is given instructions which enable it to assimilate the characterisation of the tracks and produce a vector or co-ordinates in a notional space (the semantic space) based on the type of music for each of the items in the pair. The similarity or difference between the semantic types of music is reflected by the distance between those two vectors (co-ordinates) in the semantic space. Two tracks of music which are semantically similar will have co-ordinates closer together; the farther apart they are in similarity the farther apart their vectors (co-ordinates) will be.

9. At the same time the same two tracks are analysed in another ANN (via parameters set by a human) for what are described as its physical properties - tone, timbre, speed, loudness and a lot of other characteristics set by the human (I am deliberately avoiding calling him/her a programmer for the moment). That analysis produces vectors (co-ordinates) in a notional “property space” (or “property embedding space” in claim 1), again with the differences or similarities in the music thus assessed reflected in the proximity of the co-ordinates. This is the ANN which will be the final operative ANN in the system.
10. The next step is a significant “trick” in the invention. The second ANN is trained to make the distances between pairs of the property co-ordinates converge or diverge in alignment with the distancing between them in the semantic space. Thus if the property space co-ordinates are farther apart than those in the semantic space, they are moved closer together, and conversely if the distancing is too close together in the property space to reflect semantic dissimilarity. This training is achieved by a process called back-propagation in which the “error” in the property space is corrected in order to make the results coincide with the training objectives. The back-propagation is done via an algorithm provided by a human and the correction is achieved by the ANN’s adjusting its own internal workings in such ways as adjusting weighting and bias in its nodes and levels of assessment. It learns from the experience without being told how to do it by a human being.
11. This training process is repeated many times with many pairs of tracks and the ANN learns, by repetitive correction, how to produce property vectors whose relative distances reflect semantic similarity or dissimilarity. This goes on until it is assessed that the ANN is getting it right, at which point it is “frozen” and ready to perform its intended function. It can now provide a single vector in property space for any given track of music which will have a degree of semantic similarity to other files (tracks) which is reflected in their relative property vector proximity. Similar semantic styles will be reflected in the property vectors being closer together; dissimilar styles will be reflected in the vectors being farther apart. The ANN has learned how to discern semantic (dis)similarity from physical properties. It has not done so because any human (programmer) has told it how to do it. It has done it by producing results, being provided with information reflecting its degree of error, adjusting its own internal assessment parameters, reprocessing the files to reduce the error and repeating this process until it gets it sufficiently right sufficiently often.
12. The ANN is now ready to take any given track of music provided or proposed by a remote user, determine its physical properties and attribute a property or physical vector to it. It can then relate that vector to the vectors of files in an overall database from which it is to make recommendations, and can ascertain music which is semantically similar by looking for tracks with proximate physical vectors and make a recommendation of a similar track from those nearby vectors. It does this by sending a message and a file to the remote user.

13. The advantage of this over other systems for providing recommendations of similar music to users is described in the Decision in the following terms, which are not disputed on this appeal:

“49. At this point it is helpful to turn to the main piece of prior art identified by the examiner on the basis of the searching conducted so far, US 2018/0349492 A1. Much discussion of this document was provided in the skeleton arguments, in Professor Pardoe’s report, and again at the hearing. It generally discloses training an ANN-based system to label media items with relevant contexts which can be used to generate playlists themed around those contexts. Several differences between this document and the claimed invention are identified by the applicant, not least of which is the lack of pairwise comparisons of the property and semantic vectors of files to provide convergence of semantically similar files in property space during the ANN training stage. Further, the prior art requires a larger number of ANNs in both the training and inference stages as compared to the claimed invention. The claimed invention is said to be simpler and faster as a result. I am willing to accept these alleged differences and advantages over the prior art.”

The Professor Pardoe referred to is a witness who gave expert evidence about common general knowledge and technical features relevant to the application.

The nature of an ANN

14. The nature of an ANN and its mode of operation is critical to some of the issues that arise in this case. Some of its relevant features were summarised in various parts of the Decision. Others appear in the report of Professor Pardoe and were not, at least for present purposes, disputed. The following is an abbreviated account extracted from those sources. I start with a hardware ANN, that is to say a physical box with electronics in it.
15. An ANN consists of layers of neurons which, anthropomorphising somewhat, are akin to the neurons in the brain. They are arranged in layers and connected to each other, or at least some others, and to layers below. Each neuron is capable of processing inputs and producing an output which is passed on to other neurons in other layers, save that

the last layer produces an output from the system and not to another layer. The processing is done according to internal instructions and further processes such as weights and biases applied by the neurons. Thus one feeds data in at the “top” and it is processed down through the layers in accordance with the states of the neurons, each applying its weights and biases and passing the result on, until the result of the processing is reflected in an output at the bottom.

16. The ANN is capable of learning how to process by training. As Professor Pardoe said:

“34. The training stage is looking to alter the internal parameters of the neural network. This will use an iterative learning process to determine the changes to these parameters. The learning process will use a training dataset, a validation dataset, and a loss (or cost) function. It will repeatedly present the training dataset to the ANN and determine how to modify the network parameters to reduce the error in its classification (predicted output).

35. The job of the loss function is to determine the difference (or error) between the desired outputs (often called targets) and the actual output generated by the network. The learning process will then proportionally use this error to make small changes to the network parameters [ie what the nodes do with data when received]. This process is done repeatedly for every example in the training dataset. One very common approach is called back-propagation ...”

17. The “network parameters” referred to are the weights and biases applied by each neuron. The training process is one in which those weights and biases are adjusted by the ANN itself as part of the learning or training process. Once the training process is deemed to be complete, the structure or the topology is then frozen and it can be used for real, as opposed to training, data. No more adjustments take place. It is significant to the appellant’s case that at this point no activity which might be called programming activity takes place in relation to the ANN or the data. The data which is passed through the ANN is subjected only to the processing provided by the ANN via its nodes. Furthermore, the state of the nodes is not determined by any human being programming those nodes. The state of the nodes, in terms of how they each operate and pass on data, is determined by the ANN itself, which learns via the learning process described above.
18. What I have just described is a hardware ANN. That is to say, it is a piece of hardware which can be bought off the shelf and which contains the nodes and layers in hardware

form. However, an ANN can also exist in a computer emulation. In this scenario a conventional computer runs a piece of software which enables the computer to emulate the hardware ANN as if it were a hardware ANN. Professor Pardoe explains that this is slower than a dedicated hardware ANN. However, it has the same effect. It should be appreciated that there is (or is said by Emotional Perception to be) a distinction between the underlying software on the computer which “creates” the emulated ANN, and the emulated ANN itself. I will have to return to this point.

The Decision

19. After introductory material the Decision considered points relating to added matter, sufficiency and clarity. None of those points is relevant to this appeal and I can ignore them. At paragraph 27 the Hearing Officer turned to the excluded matter point.

20. At paragraph 31 he held that it was appropriate to follow the 4 stage approach in *Aerotel Ltd v Telco Holdings Ltd* [2007] RPC 7 in order to determine whether an alleged invention falls foul of the exclusion, though not to follow it blindly. He then set out those stages, and the 5 signposts for assessing technical contribution proposed by Lewison J in *AT&T Knowledge Venture v Comptroller of Patents* [2009] FSR 19. It will be convenient to set those matters out here. The four *Aerotel* stages are:
 - (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 matter.
 - (4) If the third step has not covered it, check whether the actual or alleged contribution is actually technical.

21. The *AT & T* 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.

22. His next section deals with the first *Aerotel* step (construction) and finds that he can focus his decision on claim 4. He deals with relatively minor points. The only point of significance to this appeal is his conclusion in paragraph 41 that:

“In summary, the method of claim 4 is computer implemented and the ANN can be implemented in software or hardware as is conventional in the art.”

23. That was not disputed on this appeal. As I have described above, the ANN can take the form of a dedicated hardware unit, or it can be emulated on a computer, via appropriate software. He then concludes with his summary of how claim 4 is to be construed:

“ 42. Claim 4 can otherwise be construed straightforwardly. It defines a method in which an ANN is trained on pairs of modally identical files (for example pairs of songs) in order to map the distance between their property vectors in property space towards the distance between their semantic vectors in semantic space. When an input (for example representative of a user selected song) is presented to the ANN, it generates a file vector in property space for said input. This file vector is then compared to the property vectors of reference files (representative of a library of songs, say) to identify those files having similar vectors in property space to that of the input. Such files will be semantically similar to the input by virtue of the ANN having been trained to map/converge distances in property space towards distances in semantic space (i.e. so that it generates similar property vectors for semantically similar files). A file can then be sent to, and output by, a user device. In this way, it provides a tool for recommending semantically similar files.”

Again, no issue was taken with that.

24. He then turns to the step 2 question of identifying the contribution, and after observing that:

“The contribution does not reside in NLP [natural language processing] or the extraction of measurable properties from files per se ...” (para 51) ...

he accepted the summary proposed to him by Mr Chacksfield:

“53. “...the invention of the Application is an ANN-based system for providing improved file recommendations. The invention may be hardware or software implemented. The fundamental insight is in the training of the ANN which analyses the physical properties of the file by pairwise comparisons of training files. In these pairwise comparisons the distance in property space between the output (property) vectors of the ANN is converged to reflect the differences in semantic space between the semantic vectors of each pair of files. The result is that in the trained ANN, files clustered close together in property space will in fact have similar semantic characteristics, and those far apart in property space will have dissimilar semantic characteristics. Once trained the trained ANN can then be used to identify, swiftly and accurately, files from a database which correspond semantically to a target file, and to provide - against [sic] swiftly and accurately - file recommendations to a user device (over a communication network).”

Ms Edwards-Stewart did not take issue with that as a description of the contribution.

25. Then the Hearing Officer took *Aerotel* steps 3 and 4 together. He rejected a submission that the fact that the ANN was, or could be, hardware implemented meant that the computer program exclusion was not engaged at all. He observed, correctly, that this was not a case of a contribution to hardware *per se*, and that the case of the applicant was more “nuanced”. The applicant relied on the fact that no programmer needed to code all the detailed logical steps of the training model; the ANN adjusted itself through its training to produce a model which satisfied the training objective. The operator defined the problem and the training approach. Nonetheless, the Hearing Officer held that there was computer programming:

“61. In terms of the present invention, the applicant’s key insight involves training using pairwise comparisons of files and performing a backpropagation process to adjust weights and biases such that distances between output property vectors are converged towards the corresponding distances in semantic space. I do not believe that they are suggesting that this is a process performed entirely independently of any instruction

from the programmer. The programmer defines the problem and the training approach, and the ANN operates within those boundaries to build a suitable model. This is still no more than a computer program in my opinion.

“62. ...the key to the contribution is to specify the training method (pairwise comparison) and objective (converging distances), and this is no more than a computer programming activity.”

26. In paragraph 63 the Hearing Officer said something about emulations which Ms Edwards-Stuart sought to rely on in further written submissions made by her:

“63. I am not persuaded that the ANN can truly be decoupled from the software platform that supports it in the way Mr Chacksfield suggests. However, even if it can, it is important to consider what an ANN is at this level of generality. It is an abstract model which takes a numerical input, applies a series of mathematical operations (applying weights, biases and an activation function), and outputs a numerical result at successive layers. A claim to an ANN or the algorithm by which it is trained, in a general and abstract sense, relates wholly to a mathematical method and it fails at step 3. Even if there is something more than a mathematical method present, I cannot see how it is technical in nature and so it would not satisfy step 4.”

He went on to reject parallels with other cases where there was a definable output (*VLSI Chip Design* T0453/91, *Infineon* T1227/05 and *Re Halliburton Energy Services Inc* [2011] EWHC 2508 (Pat)) because he considered that an application to recommend semantically similar files was not a technical process (para 65).

27. He then went on to consider a further argument based on *Halliburton* to the effect that a wider technical contribution was made than the ANN merely functioning internally in that it was providing file recommendations. He posed a question and answer:

“68. What task then is the program performing? It performs an ANN training stage using pairwise comparisons of files, an ANN inference stage where an input file is analysed and semantically similar files are identified from a database, and it finishes by sending the file to the user over the network. At its core, this is a data analysis and information retrieval task which involves the

processing of data within the computer or the computer network.”

28. Ms Edwards-Stuart suggested that the word “program” was a mistake and he really meant “system”. At first I thought that was right, but on reflection I do not think that it is. I think he meant “program”. That is a point going to the question of identifying the “program for a computer”, which I consider in detail below.
29. At paragraph 69 he rejected the suggestion that the provision of a file over the network (the end result of the use of the system in practice) was a relevant external effect, saying:

“It is external to the computer in the sense that there is a beneficial effect on the end user in being provided with a better recommendation, such as a song they are likely to enjoy. However, such a beneficial effect is of a subjective and cognitive nature and does not suggest there is any technical effect over and above the running of a program on a computer.”

30. A parallel with the provision of an image in *Vicom* was rejected because in that case the end-product image was a changed image, whereas in this case the provided musical file was not changed in the process; and he rejected an analogy with the alert system in *Re Protecting Kids The World Over (PKTWO) Ltd* [2011] EWHC 2710 (Pat) on the footing that in the present case, unlike *Protecting Kids*, there was not a relevant technical process. Nor did a comparison with *Gemstar-TV Guide International Inc v Virgin Media Ltd* [2010] RPC 10 assist because the output was merely characterised by the content of the information, ie more semantically relevant file recommendations (para 73). Nor did the signposts assist. His final standing-back conclusion was:

“79 ... The ANN-based system for providing semantically similar file recommendations is not technical in nature.”

The basis of the appeal

31. Emotional Perception challenge the Decision on 2 main bases:
- (a) The computer program exclusion is not engaged at all; one does not get as far as finding a relevant computer program.
 - (b) The reasoning of the Hearing Officer fails to acknowledge a line of cases which Mr Chacksfield described as the “patentable ignoring a computer program” line of cases.

(c) If there is a computer program and the exclusion is prima facie engaged, it does not apply because the claim reveals a technical contribution and the claim is not to a program for a computer “as such”.

How, if at all, is the exclusion engaged - where is the computer and where is the program?

32. This question arose only after the hearing and as a result of my considering the parties’ oral submissions. Having considered the submissions made by the parties it seemed to me that those submissions, and the Decision, did not clearly address what seemed to me to be two fundamental questions - what is the computer, and where is the program which is said to engage the exclusion? In most, if not all, of the cases which deal with this exclusion there has been no apparent difficulty in identifying the computer and the program in issue. However, in the present case the fact that a self-trained ANN is involved means that the issue is something which requires more focus than it had been given in order to resolve the questions that arise, and it transpired that asking those questions revealed significant positions on the part of UKIPO which had not hitherto clearly emerged (and indeed an arguable change of position was revealed). The manner in which the matter was presented to the Hearing Officer does not seem to have thrown up those questions, or at least not clearly, which doubtless explains why they were not addressed clearly in the Decision. In those circumstances I either have to address the matter myself, or send it back to be addressed by the Office. It seemed to me to be more appropriate to take the former course, so far as possible.
33. I therefore invited further written submissions on those two points, and the result was a significant degree of helpful clarification of what the UKIPO’s case was together with an apparent change of tack on the point. In her further written submissions Ms Edwards-Stuart submitted the following:
- (a) In the case of a hardware ANN, the “computer” was the hardware itself.
 - (b) In the case of a hardware ANN, “there is no relevant computer program to which the exclusion applies.” Accordingly, if the claim had involved only a hardware ANN “it is unlikely” that the exclusion would be engaged. By this I took her to mean that it was her case that it would not have been engaged. This was not a position which she clearly adopted at the hearing, and is not consistent with some of her submissions. However, she did make that point clearly in her additional written submissions, so it is now her stance.
 - (c) Insofar as the invention was implemented in software, the computer is that which permits the implementation of the ANN, ie the “computer comprising code that, when executed by processor intelligence, performs the method of the various aspects recited herein, and, particularly, in the claims” (Application at p11 lines

18-20). The “code” is said by Ms Edwards-Stuart to be the code by which the input file is taken in, its vectors compared with the vectors of other known files, and the end file recommendation is sent out. In other words, code which is equivalent to the workings of a trained ANN.

(d) However, she went on to say that there was no material distinction between the computer program used to implement the trained ANN (ie the end product) and the computer program used to train the ANN. It was not clear from her final submissions, but it seems that she thereby sought to apply the computer program exclusion to the training software. That would be consistent with submissions that she made orally. She submitted that the “training instructions” were within the exclusion and the “key contribution” is the training, which she said was a computer program.

(e) The ANN itself, in a general and abstract sense (ie stripped of the software by which it was trained and/or implemented) was simply an algorithm and, as such, related wholly to a mathematical method and was excluded from patentability on that basis. She relied on the finding by the Hearing Officer in paragraph 63 of the Decision.

(f) Insofar as the invention is implemented in software form, the trained ANN is not a computer, but the trained ANN is supported by a computer.

(g) She submitted that the method of training the ANN and the operation of the trained ANN (including sending out the file recommendation) were each a computer program to which the exclusion applied. She submitted that the Hearing Officer found that the former was a computer program (Decision paragraph 61).

34. Mr Chacksfield’s position, predicably, was that a hardware ANN was not a computer, and its internal workings was not a computer program. So far as an emulated computer is concerned, he relied on a distinction between the enabling software on the one hand and the emulated ANN which it enabled on the other. If the former involved a computer program, the operation of the ANN itself did not because that was not the nature of what was going on internally.
35. I will consider the two key elements separately.

Where is the computer for the purposes of the exclusion?

36. First, I should address the question of whether a hardware ANN is a computer. Mr Chacksfield said it was not. He pointed out that it was an odd position to have something which was a computer but which did not execute a computer program, the absence of a program being conceded by Ms Edwards-Stuart. A computer was a machine which was intended to operate a computer program. He would agree with Ms Edwards-Stuart that the ANN was not operating a program, and pointed to the evidence of Professor Pardoe and various dictionary definitions of a computer program:

“Cambridge Dictionary:

“a set of instructions that makes a computer do a particular thing.”

Collins English Dictionary:

“a set of instructions for a computer to perform some task.”

The Macmillan dictionary:

“a set of instructions stored inside a computer that allows the user to do a particular thing, for example produce a document or play a game. Someone who writes computer programs is called a computer programmer.”

The Free Dictionary:

“(computer science) a sequence of instructions that a computer can interpret and execute; ‘the program required several hundred lines of code”

Ms Edwards-Stuart did not quibble with any of those definitions.

37. Professor Pardoe contrasted that sort of activity with machine learning:

“22. ... Machine learning eliminates the need to define complex hand-crafted rules that strictly follow a defined specification written by the programmer [as occurs in the development of computer programs] since the abstract machine in ML/AI technology is not processing data on a step-by-step instructional basis, but instead uses training data to learn the logic to solve a specific problem and thereby reconfigures the machine. Machine learning does not therefore follow an 'if-then' statement approach.”

38. The Hearing Officer did not address this point, no doubt because it was not raised before him.

39. I received no submissions on what a computer is, and there was no elaboration as to whether a hardware ANN is a computer. Neither side produced authorities. However, it seems to me that Ms Edwards-Stuart is likely to be correct as to whether it is a computer. I can start with the Oxford English Dictionary definition:

“An electronic device (or system of devices) which is used to store, manipulate, and communicate information, perform complex calculations, or control or regulate other devices or machines, and is capable of receiving information (data) and of processing it in accordance with variable procedural instructions (programs or software); esp. a small, self-contained one for individual use in the home or workplace, used esp. for handling text, images, music, and video, accessing and using the internet, communicating with other people (e.g. by means of email), and playing games.”

40. The first part of that definition is capable of describing a hardware ANN. The “variable procedural instructions” are, while it is learning, the elements by which it learns and back-propagates, and the frozen state contains biases, weighting and so on which it has learnt for itself and which one might call instructions. However, the key is that it is processing data. I consider that in everyday parlance it would be regarded as a computer, and ought to be treated as one within the exclusion.
41. In the light of the concession by Ms Edwards-Stuart about a hardware ANN’s internal workings, it may not matter whether it is a computer or not. What is more important is where any program is, and in the case of a hardware ANN it is not to be found in the state of the frozen ANN itself (according to the concession), but in the interests of trying to produce some sort of overall consistency, at least, I will express my views on it. A traditional computer is a computer because of its functions and activities. It is not defined by the fact that it runs things called programs. That puts things the wrong way round. So if the ANN is new thing which does not run a program in the normal sense that does not prevent it from being a computer if its functions and activities justify that description. I consider that they do justify it.
42. So far as a software emulation is concerned a computer, as normally understood, is involved (to put the matter neutrally). The emulation has to run on what can undoubtedly be viewed as a computer. That much, at least, is clear. It is also clear that the Hearing Officer treated this computer as a relevant computer for the purposes of the exclusion.

Where is the program?

43. In the case of a hardware ANN Ms Edwards-Stuart concedes that there is no program to which the exclusion applies and if the application had been confined to a hardware ANN it would not have been excluded. Mr Chacksfield does not dispute that, for obvious reasons. I am not myself minded to consider the correctness of the concession,

though I think that the debate would have been interesting had the concession not been made.

44. The Hearing Officer did not consider this question separately in relation to each of the hardware ANN and the emulated ANN, remarking in paragraph 41:

“In summary, the method of claim 4 is computer-implemented and the ANN can be implemented in software or hardware as is conventional art.”

45. The thrust of the Decision seems to be that he did not really discriminate between the two types of ANN for the purposes of the Decision. His reasoning then contains two strands. The first treated the relevant program as being the training element. In paragraph 61 he identified the key insight as being training using pairs of files, and observed that this did not take place independently of instructions from a programmer. As set out above he said (paragraph 61):

“The programmer defines the problem and the training approach, and the ANN operates within those boundaries to build a suitable model. This is still no more than a computer program in my opinion.”

46. That would seem to treat both the parameters given to the ANN and its learning progress as being the program. His next paragraph is a little more limited, apparently confining the relevant program to the provision of training objectives:

“even if this is so, key to the contribution is to specify the training method (pairwise comparison) and objective (converging distances), and this is no more than a computer programming activity.”

47. His paragraph 68, however, seems to treat the whole system as being a programming activity – see above. In the light of his varied approach (which is probably not surprising in the light of the way the matter was probably presented to him) I will have to consider it myself.

48. So far as the emulated ANN is concerned Ms Edwards-Stuart’s case was that there was a program involved in two parts of the system. The first is at the training stage, when a program was necessary to provide the training. The second is in the computer which operates the trained ANN. Her case as it finally evolved is as set out above - the operation of the trained (emulated) ANN involves the operation of a computer program. At times she also sought to combine the two as if part of one overall whole.

49. Taking those two elements separately, it does seem to be the case that the training stage involves a computer program. Mr Chacksfield accepted that that is correct at that stage. So that is a program which is going to have to be considered.

50. The knottier problem is whether the internal training and the subsequent operation of the trained emulated ANN is a computer program at all for the purposes of the exclusion. Mr Chacksfield's case was that in the case of a software emulation the emulated ANN existed at a layer above the software platform which enabled the computer to carry out the emulation. There was no program at that point because no person had given a set of instructions to the computer to do what it does - the ANN had trained itself. What it was operating was not a set of program instructions at all. It was applying its own weights, biases and so on to produce relevant vectors or co-ordinates. It was emulating a piece of hardware which had physical nodes and layers, and was no more operating or applying a program than a hardware system was.
51. As identified above, Ms Edwards-Stuart's case was that the emulated ANN was "supported" by a computer, and the manner in which it operated was "code" (to use a word used in the application) which permits the process of taking in a queried file, assesses its vectors, compares them to vectors in the reference database and spits out a file recommendation. She pointed out that the "code" is referred to in the application (page 11) in the following terms:

"In a further aspect of the present invention there is provided a computer program comprising code that, when executed by processor intelligence, performs the method of various aspects as recited herein and, particularly, in the claims."

52. The Hearing Officer dealt with Mr Chacksfield submissions in paragraph 63 of the Decision. He simply said that he was not persuaded that the emulated ANN could be decoupled from the software platform that supports it, without giving reasons.
53. The evidence before him about how emulation worked was very limited. In terms of evidence there is just a small part of Professor Pardoe's report. At paragraph 42 he says:

"42. Although AI and ANNs are often discussed in the context of a software emulation, as I have mentioned above, that is not necessarily the case."

And he goes on to discuss the availability of specialist hardware, the use of which has speed advantages. However, this sentence suggests that software emulation is more common than the use of hardware. Then at paragraph 43 he says:

"43. In software emulations the same architecture is simulated (or emulated), operating in the same manner. Software and

hardware implementations are the same in terms of the architecture, weights and biases, and the outputs produced. It is just a question of which is more convenient or efficient to use in any particular scenario.”

I take this to mean an emulation contains virtual nodes and layers which are the equivalent of the physical nodes and layers in the physical ANN.

54. Ms Edwards-Stuart’s concession about the operation of a hardware ANN was not accompanied by reasons, but presumably it is because the hardware is not implementing a series of instructions pre-ordained by a human. It is operating according to something that it has learned itself. That, at any rate, would be one justification even if it is not hers. I do not see why the same should not apply to the emulated ANN. It is not implementing code given to it by a human. The structure, in terms of the emulation of uneducated nodes and layers, may well be the result of programming, but that is just the equivalent of the hardware ANN. The actual operation of those nodes and layers inter se is not given to those elements by a human. It is created by the ANN itself.
55. I do not consider that the single sentence from the application which is relied on by Ms Edwards-Stuart is sufficient for her purposes. It appears in the middle of a number of paragraphs which refer to ANNs. It seems to refer to a different method of achieving the results of the invention which does not involve an ANN. It does not seem to be referring to an emulated ANN - it seems to be referring to something different.
56. In the light of all this I am not convinced by the Hearing Officer’s lack of conviction. It seems to me that it is appropriate to look at the emulated ANN as, in substance, operating at a different level (albeit metaphorically) from the underlying software on the computer, and it is operating in the same way as the hardware ANN. If the latter is not operating a program then neither is the emulation.
57. I should deal with Ms Edwards-Stuart’s submission that there is no difference between what she said was the computer program used to implement the trained ANN and the computer program used to train the ANN, even though the process or method implemented by the computer during training is slightly different. I do not accept this submission. First, they seem to be clearly very different things. Second, it is inconsistent with her submission that in the case of an emulated ANN the relevant “program” is that identified earlier in this section, that is to say the internal workings of a trained ANN.
58. I therefore consider that the “decoupling” can be achieved and is correct and the emulated ANN is not a program for a computer for these purposes.

59. The only remaining candidate computer program is therefore the program which achieves, or initiates, the training. That was found to be a programming activity by the Hearing Officer and, as I have observed, Mr Chacksfield agreed with that. There is no other program involved. In my view it is not correct to view the whole thing as some sort of overall programming activity for the purposes of the exclusion, which it might be thought the Hearing Officer did. One needs to be a bit more analytical than that. And if it were right to take that view then it ought to apply to the situation of a hardware ANN, and Ms Edwards-Stuart conceded that it did not. That, therefore, is not the correct approach to the questions which I have been considering so far.

Is the invention a claim to a computer program at all?

60. I therefore turn to consider whether the presence of a computer program at the training stage means that the patent claims “a program for a computer ... as such”. The programming involves setting the training objectives in terms of the structure of the ANN (if in software) and the training objectives. It is not possible to define the programming any further than that. As I have indicated, it was accepted by Mr Chacksfield that this involves some computer programming activity.
61. So far, then, there is a computer program involved, at least at that level. However, it does not seem to me that the claim claims that program. What is said to be special is the idea of using pairs of files for training, and setting the training objective and parameters accordingly. If that is right, and I consider it is, then the actual program is a subsidiary part of the claim and is not what is claimed. The claims go beyond that. The idea of the parameters itself is not necessarily part of the program. On this footing as a matter of construction the claim is not to a computer program at all. The exclusion is not invoked.
62. In the next section I consider the effect of my being wrong in this conclusion.

Technical contribution

63. If the training process does not involve a claim to a program for a computer, and if nothing else qualifies as a program, then it is strictly speaking unnecessary to consider the question of technical contribution which would arise if there were a claim which could be said to be a claim to a computer program. However, I will consider the question of technical contribution, which arises if I am wrong about the claim not being a claim to a program for a computer, or if Ms Edwards-Stuart is correct about there being a program for a computer elsewhere. The point arises in the following manner.
64. It is plain on the authorities that the mere involvement of a computer or a computer program does not by itself invoke the statutory exclusion. There are many cases which demonstrate that - for example, *Protecting Kids all over the World (PKTWO) Ltd's Application* [2012] RPC 13 and *Halliburton Energy Services Inc's Patent Application* [2012] RPC 12. Were the claim to involve a claim to a computer program it would be necessary to consider steps 3 and 4 of the *Aerotel* steps in the light of the identified contribution, and in particular whether the computer program made a technical contribution outside itself. Those steps are:

- (3) ask whether it falls solely within the excluded subject matter;
- (4) check whether the actual or alleged contribution is actually technical in nature.

65. The contribution, in what are essentially agreed terms, is that identified by the Hearing Officer and set out above.
66. Again there was no dispute between the parties as to the law in this area. If there is a technical effect (contribution) which lies outside the excluded subject matter, then the invention is unlikely fall foul of the computer program exclusion because it is not a claim to a program “as such”, but it still has to be a technical effect and one which does not itself fall within any of the statutory exclusions. There are various cases which deal with this point and which provide for a relevant external technical effect to allow an escape from the clutches of the exclusion. Several of them were relied on by Mr Chacksfield. Ms Edwards-Stuart criticised a lot of the reliance on these authorities as being “fact-matching” rather than a proper deployment of legal principle, but I consider that to be an unjustified criticism. The cases set out principles and other matters which can be relied on, and the decision on the facts of each case helps towards understanding the operation of the principles. Furthermore, consistency of approach, and therefore result, is as important in this field as in other areas of the law, and a comparison with the facts is necessary in order to keep an eye on consistency.
67. A useful starting point in this case is *Halliburton*. In that case the invention involved the process of a computer designing a drill bit by a process of simulation and alteration of a drill bit parameter. Although the words “computer” and “program” were not mentioned in the claims, it was “perfectly obvious” to Birss J that a computer was to be involved (see paragraph 20). In that context he had to consider what further questions needed to be addressed. He said:

“32. Thus when confronted by an invention which is implemented in computer software, the mere fact that it works that way does not normally answer the question of patentability. The question is decided by considering what task it is that the program (or the programmed computer) actually performs. A computer programmed to perform a task which makes a contribution to the art which is technical in nature, is a patentable invention and may be claimed as such. Indeed (see *Astron Clinica* [2008] RPC 14) in those circumstances the patentee is perfectly entitled to claim the computer program itself.

33. If the task the system performs itself falls within the excluded matter and there is no more to it, then the invention is not patentable ...

...38. What if the task performed by the program represents something specific and external to the computer and does not fall within one of the excluded areas? Although it is clear that that is not the end of the enquiry, in my judgment that circumstance is likely to indicate that the invention is patentable. Put in other language, when the task carried out by the computer program is not itself something within the excluded categories then it is likely that the technical contribution has been revealed and thus the invention is patentable. I emphasise the word "likely" rather than "necessarily" because there are no doubt cases in which the task carried out is not within the excluded areas but nevertheless there is no technical contribution at all."

68. Mr Chacksfield sought to apply the above principles by reference to what he said was found to be the ultimate task performed by any program in this case. That appears from the contribution - it is the provision of improved file recommendations via a sophisticated learning process and operation of the ANN. Mr Chacksfield relied on the steps and signposts in *Aerotel* and *A T & T*. The technical effect relied on by Emotional Perception in this respect is the sending of an improved recommendation message. It is said that the case is effectively on all fours with *Protecting Kids* and the Hearing Officer was wrong to hold that the external end result of the invention was not a "relevant technical effect". It was a relevant technical effect and it prevented the exclusion from applying.
69. There have been a number of cases which deal with this aspect of the argument. It is not necessary to consider all of them, but some are informative as to what is a relevant external technical effect within the fourth *Aerotel* step and the first *AT&T* signpost, acknowledging, as I do, that the meaning of the word technical is elusive (see *Symbian v Comptroller-General of Patents* [2009] RPC 1 at paragraph 51) and that boundaries are imprecise (paragraph 50).
70. In *Vicom* T208/84 a patent was granted for a computer-controlled process which generated an altered or improved digital image. There was an external effect which was technical and which meant that the program which achieved that was not a computer program *as such*. As already pointed out, in *Halliburton* (supra) Birss J had no difficulty in finding that a "a computer implemented method of designing drill bits" (paragraph 67) made a technical contribution, presumably in the resulting information about designs which was external to the computer and definitely a technical matter.

71. In *Gemstar -TV Guide International Inc v Virgin Media Inc* [2010] RPC 10 the court had to consider 3 patents which involved computer systems. Two of them did no more than produce displays, and that was held to be an insufficient technical effect to prevent the inventions falling foul of the exclusion. The third, however, had as its effect the facilitation of the moving of a file from one apparatus to another and that was held to be a sufficient “real world” or external effect to be technical and sufficient to enable the patent to escape the exclusion. Since this case turned on the moving of a file, and the application in suit in this case does the same, Mr Chacksfield said that this case helped him.
72. The case which Mr Chacksfield said was closest is the *Protecting Kids* case, to which I have already referred. The facts of that case involved a “data communication analysis engine” which was capable of detecting the undesirable use of computers by children (or others) by “packet sniffing”. The contents of the packets were assessed by the computer for an alert level, and if that level was reached an alert was sent digitally and electronically to an appropriate adult so that an appropriate response could be sent back to the computer. Floyd J came to the conclusion that there was the necessary technical contribution to enable the invention to escape the clutches of the exclusion. He said:

“ 34. I am unable to accept [submissions that there was no relevant technical contribution]. I start with the proposition that the generation and transmission of an alert notification to the user/administrator is not a relevant technical process. I accept that in many cases this may be correct. Plainly it was correct in the case of two out of the three patents considered by Mann J in *Gemstar*, where information was simply displayed on a screen. But what is in play in the present case, namely an alarm alerting the user, at a remote terminal such as a mobile device, to the fact that inappropriate content is being processed within the computer, is in my judgment qualitatively different. First of all, the concept, although relating to the content of electronic communications, is undoubtedly a physical one rather than an abstract one. In that respect it was more akin to the third of the three patents considered by Mann J in *Gemstar*. Secondly, the contribution of claim 33 does not simply produce a different display, or merely rely on the output of the computer and its effect on the user. The effect here, viewed as a whole, is an improved monitoring of the content of electronic communications. The monitoring is said to be technically superior to that produced by the prior art. That seems to me to have the necessary characteristics of a technical contribution outside the computer itself. “

73. Mr Chacksfield urges on me that the same approach favours him in this case. He submits that there is the moving of data outside the computer system in the form of the file that is transferred, which provides an external (outside world) effect. That is as technical as the alert in *Protecting Kids*. It is not a disqualification that the result may be to facilitate user enjoyment.

74. I agree with Mr Chacksfield's case on this point. The Hearing Officer found, in this respect, that the sending of the file to the end user is a matter external to the computer (acknowledging earlier that this makes it more likely, but not necessarily inevitable, that there is a technical effect) but said that that was achieved in a standard fashion within a conventional computer network. He went on to find:

“69. ... It is external to the computer in the sense that there is a beneficial effect on the end user in being provided with a better recommendation, such as a song they are likely to enjoy. However, such a beneficial effect is of a subjective and cognitive nature and does not suggest there is any technical effect over and above the running of a program on a computer.”

75. He distinguished *Vicom* in that in that case the file (a photograph) was changed by the process, whereas in this case the output file is not altered in any way; the assessment of the emotional nature of the file “is not a technical process” (paragraph 70). He repeated this basis in paragraph 72:

“There is nothing at the level of improved monitoring of the content of electronic communications [as in *Protecting Kids*]. There is only the improved identification and recommendation of files based on their semantic similarity, which is not a relevant technical effect.”

And again in paragraph 76:

“An effect on the end user by way of receiving a semantically similar file, such as a song they might enjoy, is not a relevant technical effect.”

76. A decision of an expert tribunal such as the IPO Hearing Officer in this case is entitled to respect in relation to technical matters, and in respect of judgments such as were

made in relation to technical effect, but I am afraid that in this instance I consider the judgment to be flawed and disagree with the assessment. The Hearing Officer seemed to consider that a subjective appreciation of the output of the system was just that, subjective and in the user, and therefore not a technical effect. I do not consider that to be the correct analysis. The Hearing Officer was right to acknowledge that the result of the invention was an effect external to the computer in the transmission of a chosen file. That is usefully analogous to the file that was moved in the third *Gemstar* patent. The correct view of what happened, for these purposes, is that a file has been identified, and then moved, because it fulfilled certain criteria. True it is that those criteria are not technical criteria in the sense that they can be described in purely technical terms, but they are criteria nonetheless, and the ANN has certainly gone about its analysis and selection in a technical way. It is not just any old file; it is a file identified as being semantically similar by the application of technical criteria which the system has worked out for itself. So the output is of a file that would not otherwise be selected. That seems to me to be a technical effect outside the computer for these purposes, and when coupled with the purpose and method of selection it fulfils the requirement of technical effect in order to escape the exclusion. I do not see why the possible subjective effect within a user's own non-artificial neural network should disqualify it for these purposes. To adapt the wording of Floyd J in *Protecting Kids*, the invention is not just one depending on the effect of the computerised process on the user. There is more than that. There is a produced file with (it is said) certain attributes. The file produced then goes on to have an effect on the user (if the thing works at all) but one cannot ignore the fact that a technical thing is actually produced. It would not matter if the user never listened to the file. The file, with its similarity characteristics, is still produced via the system which has set up the identification system and then implemented it. This effect is qualitatively different from the first two instances in *Gemstar*, and qualitatively similar to the effect in *Protecting Kids*.

77. The preceding reasoning looks to the end result as being something which helps to take the case away from being a case of a computer program "as such". However, there is another way of approaching the matter if one is assuming for these purposes that the computer program is either the training program or the overall training activity (which the Hearing Officer seems to have considered to be a computer program – see his paragraph 61 – and which was relied on by Ms Edwards-Stuart).
78. If, contrary to my findings, one were considering those two program candidates, it seems to me that the resulting ANN, and particularly a trained hardware ANN, can be regarded as a technical effect which prevents the exclusion applying. At the hearing Ms Edwards-Stuart seemed to accept, in argument, that a trained ANN could be a technical advance for these purposes, but proposed that it had to be defined in terms of the actual function of each of its nodes so as to be identifiable as a particular ANN, or be determined by reference to the training that it received. The first of those is obviously not part of the application, but I do not see why the second, or something very close to it, has not been achieved. I therefore consider that, insofar as necessary, the trained hardware ANN is capable of being an external technical effect which prevents the exclusion applying to any prior computer program. There ought to be no difference between a hardware ANN and an emulated ANN for these purposes.

Other exclusions

79. Section 1(2)(a) of the Act excludes a “mathematical method” ... as such. In paragraph 63 the Hearing Officer said that:

“A claim to an ANN or the algorithm by which it is trained, in a general and abstract sense, relates wholly to a mathematical method and it fails at step 3 [of *Aerotel*]. Even if there is something more than a mathematical method present, I cannot see how it is technical in nature and so it would not satisfy step 4.”

80. In the conclusions in her post-hearing written submissions Ms Edwards-Stuart invited me to find that a software-implemented ANN, or its training algorithm, was excluded as a mathematical method, based on that apparent determination by the Hearing Officer.

81. Mr Chacksfield resisted such a finding on both procedural and on substantive bases. He pointed out that the Hearing Officer did not make mathematical method an alternative basis of his decision and Ms Edwards-Stuart did not file a respondent’s notice to resurrect it. This submission is based on paragraph 78 of the Decision:

“78. On the issue of exclusion as a mathematical method, although an ANN and a method of training an ANN per se is no more than an abstract mathematical algorithm, its specific application here as part of a file recommendation engine is, in my opinion, enough to dispense with the mathematical method as such objection.”

82. That procedural objection seems to me to be justified. If Ms Edwards-Stuart wished to run it as an alternative she should have served such a notice. This is a case in which that matters. It had the result that the point was not properly addressed at the appeal hearing even though it was (Mr Chacksfield tells me) a live argument at the hearing below. Ms Edwards-Stuart did not propound it in her skeleton argument or in her oral submissions. Mr Chacksfield did refer to the point in his skeleton argument but did not fully develop it. Substantively he pointed out what he said was an inconsistency between paragraphs 63 and 78, and made the point briefly that every computer program is a mathematical method in that it relies on a combination of 0s and 1s and was therefore mathematical, saying that therefore the point did not go anywhere in these proceedings.

83. Since the point was not in play in the appeal I do not think it right to consider it further and I shall not do so.

Conclusion

84. For the reasons appearing above, I would allow the appeal. The order I should make can be the subject of debate if it is not agreed.

Appendix - the principal claims

1 . A system for providing semantically relevant file recommendations, the system containing:

a) an artificial neural network “ANN” having an output capable of generating a property vector in property space, the ANN trained by subjecting the ANN to a multiplicity of pairs of training data files sharing a content modality and where for each pair of training data files there are two independently derived separation distances, namely:

a first independently derived separation distance that expresses a measure of relative distance between a first pair of training data files in semantic embedding space, where the first independently derived separation distance is obtained from natural language processing “NLP” of a semantic description of the nature of the data associated with each one of the first pair of training data files; and

a second independently derived separation distance that expresses a measure of relative distance similarity between the first pair of training data files in property embedding space, where the second independently derived separation distance is a property distance derived from measurable properties extracted from each one of the first pair of training data files, and

wherein training of the ANN by a backpropagation process uses output vectors generated at the output of the ANN from processing of said multiplicity of pairs to adjust weighting factors to adapt the ANN during training to converge distances of generated output vectors, in property embedding space, towards corresponding pairwise semantic distances in semantic space, and wherein shared content modality is: (i) video data files; or alternatively (ii) audio data files; or alternatively (iii) static image files; or alternatively (iv) text files; and

b) a database in which is stored a multiplicity of reference data files with content modality with target data and a stored association between each reference data file and a related individual

property vector, wherein each related individual property vector is obtained from processing, within the trained ANN, of file properties extracted from its respective reference data file and each related individual property vector encodes the semantic description of its respective reference data file;

c) a communications network;

d) a network-connected user device coupled to the communications network;

e) processing intelligence arranged:

in response to the trained ANN receiving target data as an input and for which target data an assessment of relative semantic similarity of its content is to be made, and the ANN producing a file vector (V_{piie}) in property space for the target data based on processing within the trained ANN of file properties extracted from the target data;

to access the database;

to compare the file vector of the target data with individual property vectors of the multiplicity of reference data files in the database to produce an ordered list which identifies relevant reference data files that have property vectors measurably similar to the property vector and thus to identify relevant reference files that are semantically similar to the target data; and

to send, over the communications network, relevant reference files to the user device;

wherein

the user device is arranged to receive the relevant reference files and to output the content thereof.

....

4. A method of providing semantically relevant file recommendations in a system including an artificial neural network "ANN" having an output capable of generating a property vector in property space, the method comprising:

a) training the ANN by subjecting the ANN to a multiplicity of pairs of training data files sharing a content modality and where for each pair of training data files there are two independently derived separation distances, namely:

a first independently derived separation distance that expresses a measure of relative distance between a first pair of training data files in semantic embedding space, where the first independently derived separation distance is obtained from natural language processing "NLP" of a semantic description of the nature of the data associated with each one of the first pair of training data files; and

a second independently derived separation distance that expresses a measure of relative distance similarity between the first pair of training data files in property embedding space, where the second

independently derived separation distance is a property distance derived from measurable properties extracted from each one of the first pair of training data files,

and wherein shared content modality is: (i) video data files; or alternatively (ii) audio data files; or alternatively (iii) static image files; or alternatively (iv) text files;

b) in a backpropagation process in the ANN, using output vectors generated at the output of the ANN from processing of said multiplicity of pairs to adjust weighting factors in the ANN, thereby adapting the ANN during training to converge distances of generated output vectors, in property embedding space, towards corresponding pairwise semantic distances in semantic space, and

b) storing, in a database, a multiplicity of reference data files with content modality with target data and a stored association between each reference data file and a related individual property vector, wherein each related individual property vector is obtained from processing, within the trained ANN, of file properties extracted from its respective reference data file and each related individual property vector encodes the semantic description of its respective reference data file;

c) in response to the trained ANN receiving target data as an input and for which target data an assessment of relative semantic similarity of its content is to be made, and the ANN producing a file vector (V_{nie}) in property space for the target data based on processing within the trained ANN of file properties extracted from the target data;

d) comparing the file vector of the target data with individual property vectors of the multiplicity of reference data files in the database to produce an ordered list which identifies relevant reference files that are measurably similar to the property vector and thus identifying relevant reference data files that are semantically similar to the target data;

e) sending, over the communications network, relevant reference data files to the user device; and

f) at the user device, receiving the relevant reference files and outputting the content thereof.