

Technology assisted review approved for use in English High Court litigation

By **Clive Freedman**

In complex litigation it is now not unusual for there to be millions of documents which need to be assessed for relevance to the issues prior to disclosure being given to the other party. However human beings, including experienced lawyers, are not good at reaching consistent decisions as to whether individual documents are or are not relevant to the issues in a case. Computers on the other hand are now able to analyse vast numbers of documents very quickly in order to identify the probability of each document being relevant to the issues in the case, based on a comparison between on the one hand the contents of the document and on the other hand a sample set of documents which an expert review has classified as being relevant or irrelevant to the issues. The computer can be relied on to reach consistent decisions over a very large pool of documents, and to work continuously without being interrupted or getting bored. There are potentially very considerable savings to be made by using computers in this way.

In 2012, in the US Federal Court case of *Moore v Publicis Groupe, SA*, 2013 WL 4483531, 11 Civ 1279 (ALC)(AJP), Magistrate Judge Peck approved the use of technology assisted review using predictive coding in giving discovery. In 2015 the Irish High Court also approved the use of predictive coding in *Irish Bank Resolution Corporation Limited v Quinn* [2015] IEHC 175; [2015] 3 JIC 0306, notwithstanding the opposition of the defendants in that case.

It was not until February 2016 that an English judge was asked to approve the use of predictive coding, in *Pyrrho Investments Ltd v MWB Property Ltd* [2016] EWHC 256 (Ch). The parties in that case were agreed, subject to the approval of the court, that the case was a case in which it was appropriate for predictive coding to be used. Master Matthews, co-author of *Disclosure* (4th edn, Sweet and Maxwell, 2011), delivered a reserved judgment in which he described predictive coding, set out the advantages of using it, and gave his reasons for approving its use in the case

before him. In *Brown v BCA Trading Ltd* [2016] EWHC 1464 (Ch), the respondents proposed to use predictive coding to identify potentially relevant documents, but the petitioner opposed this. Mr Registrar Jones approved the use of predictive coding notwithstanding the petitioner's opposition.

In *Pyrrho*, Master Matthews described the proposed use of predictive coding in the case before him as follows:

“17. Mr Spencer explains in his statement that the term ‘predictive coding’ is used interchangeably with ‘technology assisted review’, ‘computer assisted review’, or ‘assisted review’. It means that the review of the documents concerned is being undertaken by proprietary computer software rather than human beings. The software analyses documents and ‘scores’ them for relevance to the issues in the case. This technology saves time and reduces costs. Moreover, unlike with human review, the cost does not increase at the same rate as the number of documents to be reviewed increases. So doubling the number of documents does not double the cost.

...

19. In modern times, as I understand it, the predictive coding process runs more or less like this. First of all, the parties will settle a predictive coding protocol, setting out the process in more detail, including definition of the data set, sample size, batches, control set, reviewers, confidence level and margin of error. Then criteria (perhaps agreed, perhaps unilateral) must be decided upon for inclusion of documents in the process. Those criteria will include who had the documents (“custodians”) and the date range, but perhaps also whether the documents contained any of the keywords chosen. Certain types of documents, not having any or any sufficient text, will be excluded (they will

have to be considered manually). The resulting documents are 'cleaned up', by removing repeated content (*eg* email headers or disclaimers) and words that will not be indexed (*eg* because not useful in assessing relevance).

20. Then a representative sample of the 'included' documents is used to 'train' the software. In the present case, Mr Spencer suggests that it will comprise 1600-1800 documents (a size set by the size and variety of the entire document set). A person who would otherwise be making the decisions as to relevance for the whole document set (*ie* a lawyer involved in the litigation) considers and makes a decision for each of the documents in the sample, and each such document is categorised accordingly. It is essential that the criteria for relevance be consistently applied at this stage. So the best practice would be for a single, senior lawyer who has mastered the issues in the case to consider the whole sample. Where documents would for some reason not be good examples, they should be deselected so that the software does not use them to learn from. The software analyses all of the documents for common concepts and language used. Based on the training that the software has received, it then reviews and categorises each individual document in the whole document set as either relevant or not.

21. The results of this categorisation exercise are then validated through a number of quality assurance exercises. These are based on statistical sampling. The sampling size will be fixed in advance depending on what confidence level and what margin of error are desired. The higher the level of confidence, and the lower the margin of error, the greater the sample must be, the longer it will take and the more it will cost. (These quality assurance exercises are clearly "additional techniques" contemplated by paragraph 27 of Practice Direction B to Part 31.)

22. The samples selected are (blind) reviewed by a human for relevance. The software creates a report of software decisions overturned by humans. The overturns are themselves reviewed by a senior reviewer. Where the human decision is adjudged

correct, it is fed back into the system for further learning. (It analyses the correctly overturned documents just as the originals were analysed.) Where not correct, the document is removed from the overturns. Where the relevance of the original document was incorrectly assessed at the first stage, that is changed and all the documents depending on it will have to be re-assessed.

23. The process of sampling is repeated as many times as required to bring the overturns to a level within agreed tolerances, and so as to achieve a stability pattern. This is usually not less than 3, making 4 rounds in total. In his statement, Mr Spencer says that he understands that in fact it should involve review of some 8 to 12 batches of documents. The trend of overturns should be lower from round to round. Ultimately there will be a final overturn report within the agreed tolerance, so that the expense of further rounds of review will not be justified by the reduced chance of finding further errors, and the list of relevant documents can be produced.

24. Although the number of documents that have to be manually reviewed in a predictive coding process may be high in absolute numbers, it will be only a small proportion of the total that need to be reviewed in the present case. Thus - whatever the cost per document of *manual* review - provided that the exercise is large enough to absorb the up-front costs of engaging a suitable technology partner, the costs overall of a predictive coding review should be considerably lower. It will be seen that, because the software has to be trained for every case, each use of the predictive coding process is bespoke for that case."

As has been pointed out,¹ it is important to note that variations to the workflows described by Master Matthews are possible. In different circumstances those variations may or may not be more cost-effective.

¹ Celina McGregor, senior associate at Herbert Smith Freehills, quoted at <http://www.legaltechnology.com/latest-news/high-court-says-yes-to-predictive-coding-the-reaction/>. See also the comments of Chris Dale on *Pyrrho* at <https://chrisdale.wordpress.com/2016/02/23/reacting-to-the-reactions-to-the-pyrrho-predictive-coding-judgment>.

The Master quoted a passage from *Irish Bank Resolution Corporation Ltd v Quinn* which set out statistics showing that predictive coding was more accurate than human reviewers in identifying which documents were relevant, and more accurate than human reviewers in excluding documents which were not relevant.

Master Matthews' reasons for approving the use of predictive coding including the following:

(a) There was no evidence to show that the use of predictive coding software leads to less accurate disclosure being given than, say, manual review alone or keyword searches and manual review combined, and indeed there is some evidence (referred to in the US and Irish cases) to the contrary.

(b) There would be greater consistency in using the computer to apply the approach of a senior lawyer towards the initial sample (as refined) to the whole document set, than in using dozens, perhaps hundreds, of lower-grade fee-earners, each seeking independently to apply the relevant criteria in relation to individual documents.

(c) The number of electronic documents which must be considered for relevance and possible disclosure in the present case was huge, over 3 million.

(d) The cost of manually searching these documents would be enormous, amounting to several million pounds at least, and a full manual review of each document would therefore be unreasonable.

(e) The costs of using predictive coding software would depend on various factors, including whether the number of documents was reduced by keyword searches, but the estimates given in the case before him varied between £181,988 plus monthly hosting costs of £15,717, to £469,049 plus monthly hosting costs of £20,820. This was obviously far less expensive than the full manual alternative, though of course there might be additional costs if manual reviews still needed to be carried out when the software had done its best.

Brown v BCA Trading Ltd concerned a section 994 petition by a minority shareholder. The majority of the documents which might be relevant were in the hands of the respondents, who wished to use predictive coding to identify potentially relevant documents. They estimated that the costs would be in the region of £132,000 using predictive coding, as opposed to an estimated cost of at least £250,000 for a disclosure exercise using keywords. The petitioner opposed this, but did not submit factual or expert evidence in support of his position. Mr Registrar Jones approved the use of predictive coding, adopting similar reasons to Master Matthews' reasons in *Pyrrho*.

The decisions in *Pyrrho* and *Brown* are unsurprising. The Civil Procedure Rules, including CPR 31 and Practice Direction 31B dealing with disclosure, encourage the parties to use cost-effective methods to minimise costs, and it is not difficult to demonstrate that predictive coding would be cost-effective in a case involving very large numbers of documents and large sums of money. Even if predictive coding is not used for the initial selection of documents which are relevant to the issues, it may be advantageous to use it as a method of quality control in supervising human reviewers, and in analysing the documents disclosed by opposing parties. The judgments in *Pyrrho* and *Brown* will provide support to legal advisers seeking to persuade the opposing party to agree to the use of predictive coding for document review, and can only help to encourage the wider use of predictive coding in litigation in England.

The judgments in *Pyrrho* and *Brown* do not address contentious issues which may arise when parties are engaged in trying to reach agreement on a protocol for the use of predictive coding. An example of such an issue would be the approach to be taken to disclosure of documents which have been assessed as being irrelevant. When predictive coding is used, the results of analysing a small set of documents are extrapolated over a large pool of documents. It is in the nature of such an exercise that an initial coding error by a human reviewer in relation to one or more documents dealing with a particular subject-matter may lead to the failure to disclose a substantial number of documents which deal with that subject-matter. In US cases such as *Rio Tinto Plc v Vale S.A.*, 306 F.R.D. 125 (S.D.N.Y. 2015) and in the Irish case of *Quinn* the protocols included provision for disclosure to the receiving party of a number of documents

which had been assessed as being not relevant, and provision for the initial assessment of non-relevance to be challenged. In London it is rarely suggested (let alone agreed) that non-relevant documents should be disclosed in this manner, since the disclosure of documents which a party is not compelled to disclose is almost always resisted. This may be a point which in due course comes before a court for consideration.

© Clive Freedman, 2016

Clive Freedman, barrister, whose website dealing with the English case law on electronic disclosure is at www.edisclosure.uk.com.

[Note – the author is not Clive Freedman QC, who appeared for the Second Defendant in *Pyrrho Investments*]