

Forthcoming in Jill Presser, Jesse Beatson, and Gerard Chan (eds.), Artificial Intelligence & the Practice of Law: A Playbook for Litigators (Edmond Publishing 2020/2021)

Part V – AI-Enabled Litigation Tools

17. Vetting and Validation of AI-Enabled Tools for Electronic Discovery

Maura R. Grossman and Gordon V. Cormack

17.1 Introduction

The process of document discovery in litigation has been almost entirely supplanted by electronic discovery (“eDiscovery”) in which responsive¹ electronically stored information (“ESI”) must be identified, and unless withheld for privilege, produced to the requesting party. While diverse sources and formats of ESI—such as email, text messages, social media, and scanned images subject to optical character recognition (“OCR”)—have replaced filing cabinets and banker’s boxes filled with hard-copy documents, the term “document” has survived to describe any unit of ESI subject to production, if it is relevant to the claims and defences of the matter at hand, and can be found and produced without disproportionate burden and cost.

Concordant with the advent of ESI has been an explosion in the number of potentially relevant “documents,” along with the cost and burden of examining them to determine which are responsive, and among those that are responsive, which may be withheld for privilege. It is typically infeasible to review every email for a single individual, let alone email for the key custodians of an organization, along with other sources of ESI such as text messages and social media. Unlike filing cabinets and banker’s boxes, email folders and other electronic documents are seldom well enough organized that one can identify just a few folders that are highly likely to contain the documents of interest.

Let us imagine a hypothetical “ideal” situation in which cost and timeliness were not an issue, and an army of competent reviewers could examine every potentially responsive document, and label each as responsive or not. Artificial Intelligence (“AI”) methods—collectively dubbed technology-assisted review (“TAR”)—can effectively emulate this ideal, even when there are many thousands or millions of potentially responsive documents to be labeled. In 2011, the authors published the results of an experiment showing that two different TAR methods were at least as effective as such an exhaustive manual review, with a tiny fraction of the effort.² The

¹ The terms “relevant” and “responsive” are often used interchangeably. For these purposes, we will use “relevant” to refer to documents containing information that is pertinent to a subject matter of interest, *e.g.*, the claims or defences in a litigation, and “responsive” to refer to documents containing information that, unless protected by privilege or work-product protection, is subject to production in response to a request for production or subpoena.

² Maura R. Grossman and Gordon V. Cormack, *Technology-Assisted Review in E-Discovery Can Be More Effective and More Efficient Than Exhaustive Manual Review*, 17 *Richmond J. Law & Tech.* 1 (2011) (hereinafter “JOLT study”).

two TAR methods we found to be more effective than exhaustive manual review were a *rule-based* method and a *supervised machine-learning* method.³

Citing our work either directly or by reference, courts in the United States, Ireland, the United Kingdom, and Australia have approved the use of TAR for eDiscovery,⁴ and its use is now generally accepted in principle in these and other jurisdictions, including Canada. Controversy remains, however, as to precisely what constitutes a reasonable TAR process and, after following such a process, whether a reasonable result has been achieved.

In a perfect world, a recognized body would set standards for the application of TAR technology in eDiscovery and would certify particular tools and protocols for adherence to those standards. The practitioner could then enjoy a reasonable probability of success by properly applying a certified tool, and could confirm success by applying a certified validation protocol. At the time of this writing (June 2020), however, no organization had taken up the mantle. Therefore, for now—and for the foreseeable future—practitioners must take it upon themselves to identify reasonable tools, procedures, and validation protocols, based on the best available evidence coupled with their own investigations.

This chapter summarizes the available evidence concerning how AI-based TAR (and some non-TAR) eDiscovery tools and processes work, which tools and processes have (and have not) been shown to be effective, and how to measure their effectiveness. It offers practical guidance and model protocols covering both how to select and use a TAR tool, and how to evaluate its efficacy—before, during, and after a particular eDiscovery review effort.

17.2 TAR and Non-TAR Methods and Their Effectiveness

Exhaustive manual review typically entails too much time, effort, and cost to be a viable review strategy for all but the smallest matters. In order to reduce this burden, keyword culling has traditionally been used to reduce the size of the collection, excluding from review all documents that do not match a particular set of search terms. The excluded documents (the “null set”) are presumed to be non-responsive, while the matching documents are subject to manual review and labeling for responsiveness, privilege, confidentiality, and sometimes specific issues related to the claims or defences in the litigation.

³ *Id.* at 29-34. A *rule-based* TAR method relies on a set of rules created by one or more experts to emulate the human decision-making process for the purposes of classifying documents. A *supervised machine-learning* TAR method is one in which an algorithm learns to distinguish between relevant and non-relevant documents using a training set and labeled human judgments. See Maura R. Grossman and Gordon V. Cormack, *The Grossman-Cormack Glossary of Technology-Assisted Review*, 7 Fed. Cts. L. Rev. 1, 28, 31 (2013) (hereinafter “TAR Glossary”). Supervised machine learning methods are discussed more fully, *infra*, at Section 17.3.

⁴ E.g., *Da Silva Moore v. Publicis Groupe*, 287 F.R.D. 182 (S.D.N.Y. 2012); *Irish Bank Resolution Corp. & Ors. v. Sean Quinn & Ors.* [2015] IEHC 175 (Mar. 3, 2015); *Pyrrho Inv. Ltd. v. MWB Prop. Ltd.* [2016] EWHC (Ch) (Feb. 16, 2016); *McConnell Dowell Constructors (Aust) Pty Ltd. v. Santam Ltd. & Ors. (No. 1)* [2016] VSC 734 (2 Dec. 2016).

The *end-to-end recall* of such a review effort—as for any review effort—is the fraction of responsive documents in the collection that are correctly labeled as responsive, which can happen only when they are matched by the search terms *and* correctly labeled by the reviewer. The *end-to-end precision* of such a review effort is the fraction of documents labeled responsive that are in fact (*i.e.*, truly) responsive.⁵ Again, as for any review effort, it is possible to compute only an *estimate* of recall and precision, by comparing the results to a *gold standard*.⁶

It is important to distinguish the end-to-end recall of the review effort with the recall of the keyword culling alone (*i.e.*, “search-term recall”). Search-term recall is the fraction of responsive documents that are matched by the keywords, regardless of whether they are correctly labeled by the reviewer. Similarly, search-term precision is the fraction of documents matched by the keywords that are actually responsive, irrespective of human review. A search-term recall estimate of 70% would result in a much lower end-to-end recall estimate, because a competent reviewer would be expected to identify *at most* about 70% of the 70% responsive documents remaining after keyword culling—49% recall overall (*i.e.*, 70% of 70%).⁷

In unusual circumstances, the results of a keyword search may be produced without subsequent manual review. Only in such circumstances would the search-term recall and the end-to-end recall be one and the same; only in such circumstance is 70% (or better) search-term recall consistent with an effective result.

Beware of keyword culling.

In a seminal 1985 study, Blair and Maron⁸ had lawyers and paralegals employ iterative searches and then review the resulting documents, until they believed they had achieved at least 75% recall. Fifty-one different information needs (*i.e.*, topics or requests for production (“RFPs”))

⁵ See TAR Glossary, *supra* n.3, at 27 and 25 for definitions of *recall* and *precision*, respectively.

⁶ A *gold standard* is the best available determinant of relevance or non-relevance of all (or a sample) of a document collection, used as a benchmark to evaluate the effectiveness of a search and review effort. It is also referred to as “ground truth.” TAR Glossary, *supra* n.3, at 18.

⁷ See JOLT study, *supra* n.2, at 37 tbl. 7 (showing recall estimates for human reviewers ranging from 25.2% to 79.9%, with an average of 59.3%); Herbert L. Roitblat et al., *Document Categorization in Legal Electronic Discovery: Computer Classification vs. Manual Review*, 61 J. Am. Soc’y for Info. Science & Tech. 70, 76 tbl. 2 (2010) (showing recall estimates for human reviewers ranging from 45.8% to 53.9%, with an average of 50.3%). See also Scott M. Cohen et al., *The Tested Effectiveness of Equivio>Relevance in Technology Assisted Review*, Winston & Strawn White Paper, at 6 fig. 5, 7 (2014), available at <http://www.equivio.com/files/files/White%20Paper%20-%20Winston%20and%20Strawn.pdf> (showing a recall estimate of 52.4% for human reviewers). An example of the confusion between search-term recall and end-to-end can be seen in *In re Biomet M2A Magnum Hip Implant Prods. Liabil. Lit. (MDL 2391)*, Case No. 3:12-MD-2391, 2013 WL 1729682, at *2 (Apr. 18, 2013), where the Court concluded that “a comparatively modest number of documents” had been missed by the TAR process when the search-term recall estimate was approximately 60%, and the end-to-end recall estimate (including TAR and manual review) would necessarily have been substantially lower than that. See William Webber, *What is the maximum recall in re Biomet?*, Evaluating E-Discovery (Apr. 24, 2013), available at <http://blog.codalism.com/index.php/what-is-the-maximum-recall-in-re-biomet/>.

⁸ David C. Blair and M. E. Maron, *An Evaluation of Retrieval Effectiveness for a Full-Text Document-Retrieval System*, 28 Commc’ns ACM 289 (1985).

relating to a Bay Area Rapid Transit (“BART”) train accident were employed in the study.⁹ These search and retrieval efforts achieved, on average, only 20% recall, at the same time achieving 79% precision.¹⁰ From this result—which has not been contradicted in 35 years—we can conclude that it is difficult to achieve high recall using search terms, that it is difficult to know whether or not high recall has been achieved, that searchers overestimate the recall of their searches,¹¹ and that high precision is often mistakenly construed to be an indicator of high recall.

Keyword culling yields dramatically inferior results to exhaustive manual review, but prior to the advent of TAR was—and sometimes still is—considered a necessary evil. Even so, it was commonly recognized that some keywords were better than others, and that some sort of evaluation of search terms was necessary.¹² This observation, in part, gave impetus to the National Institute of Standards and Technology’s (“NIST’s”) Text REtrieval Conference (“TREC”) Legal Track, which, from 2006 through 2011, measured the effectiveness of various search methodologies by comparing their results to a gold standard created for the purpose.¹³ The TREC Legal Track showed results for keyword culling consistent with those of Blair and Maron,¹⁴ and dramatically superior results (as compared to manual review) for two specific TAR

⁹ *Id.* at 291.

¹⁰ *Id.* at 293.

¹¹ See also Maureen Dostert and Diane Kelly, *Users’ Stopping Behaviors and Estimates of Recall*, in Proceedings of the 32nd Int’l ACM SIGIR Conference on Research & Dev. in Info. Retrieval (“SIGIR 2009”) 820 (2009) (showing that most subjects in an interactive information retrieval experiment reported they had found 51-60% of the relevant documents when, on average, recall was only 7%).

¹² See, e.g., *William A. Gross Constr. Assocs., Inc. v. American Mfrs. Mutual Ins. Co.*, 256 F.R.D. 134, 136 (S.D.N.Y. 2009) (“This Opinion should serve as a wake-up call to the Bar in this District about the need for careful thought, quality control, testing . . . in designing search terms or ‘keywords’ to be used to produce emails or other electronically stored information (‘ESI’)”; “[W]here counsel are using keyword searches for retrieval of ESI, they at a minimum must carefully craft the appropriate keywords, with input from the ESI’s custodians . . . and the proposed methodology must be quality control tested to assure accuracy in retrieval and elimination of ‘false positives.’”); *Victor Stanley, Inc. v. Creative Pipe, Inc.*, 250 F.R.D. 251, 260, 262 (D. Md. 2008) (“While keyword searches have long been recognized as appropriate and helpful for ESI search and retrieval, there are well-known limitations and risks associated with them, and proper selection and implementation obviously involves technical, if not scientific knowledge.”; “Selection of the appropriate search and information retrieval technique requires careful advance planning by persons qualified to design effective search methodology. The implementation of the methodology selected should be tested for quality assurance. . . .”); *In re Seroquel Prods. Liab. Litig.*, 244 F.R.D. 650, 662 (M.D. Fla. 2007) (“[W]hile keyword searching is a recognized method to winnow relevant documents from large repositories . . . [c]ommon sense dictates that sampling and other quality assurance techniques must be employed to meet requirements of completeness.”).

¹³ “The goal of the Legal Track at the Text Retrieval Conference (TREC) is to assess the ability of information retrieval techniques to meet the needs of the legal profession for tools and methods capable of helping with the retrieval of electronic business records, principally for use as evidence in civil litigation. In the USA, this problem is referred to as ‘e-discovery.’ Like all TREC tracks, the Legal Track seeks to foster the development of a research community by providing a venue for shared development of evaluation resources (‘test collections’) and baseline results to which future results can be compared.” <https://trec-legal.umiacs.umd.edu/>.

¹⁴ See Douglas W. Oard et al., *Overview of the TREC 2008 Legal Track*, in NIST Special Publication: SP 500-277, The Seventeenth Text REtrieval Conference (TREC 2008) Proceedings, at 4, 8-9 (2008), available at <https://trec.nist.gov/pubs/trec17/papers/LEGAL.OVERVIEW08.pdf> (showing an average recall for search terms of 24% over 45 topics); S. Tomlinson et al., *Overview of the TREC 2007 Legal Track*, in NIST Special Publication: SP 500-274, The Sixteenth Text REtrieval Conference (TREC 2007) Proceedings, at 4, 10 (2007) (showing an average recall for search terms of 22% over 50 topics), available at

methods.¹⁵ To be clear, it also showed inferior results for the majority of TAR methods evaluated at TREC, and inconsistent results for others.¹⁶

Not all “TAR” is effective.

Myriad service providers were quick to repurpose (or simply relabel) their existing search and analytics offerings as “TAR,” pointing to TREC and the authors’ 2011 JOLT study as (specious) evidence of their efficacy. Along with these offerings came a number of misconceptions and prescriptions—many of which still persist today—about how TAR should be conducted and how its results should be validated.¹⁷ Many practitioners were disappointed by the results of using these “TAR” offerings and following this advice, and the adoption of TAR was hindered, either because inferior results were obtained or because the prescribed rituals for the use of these tools were too complicated and burdensome.

More recently, the machine-learning method that the authors have found to be most effective, which has been trademarked as “Continuous Active Learning” or “CAL,”¹⁸ and known generically in the eDiscovery industry as “TAR 2.0,”¹⁹ has come to be recognized as the industry standard and the method of choice. A number of service providers have claimed to implement similar methods to CAL, typically branding their products either as “TAR 2.0,” or with two of the three words, “continuous,” “active,” or “learning.” At the time of this writing, no independent evaluation of these commercial offerings had been conducted.

The rule-based method that was found to be effective at the TREC Legal Track is a trade secret of H5.²⁰ We do know that it involves extensive interplay among subject-matter experts, linguists, and statisticians.²¹ Some service providers have claimed to implement rule-based

<https://trec.nist.gov/pubs/trec16/papers/LEGAL.OVERVIEW16.pdf>; Jason R. Baron et al., *TREC-2006 Legal Track Overview*, in NIST Special Publication: SP 500-272, The Fifteenth Text REtrieval Conference (TREC 2006) Proceedings, at 4-5, 11 (2006), available at <https://trec.nist.gov/pubs/trec15/papers/LEGAL06.OVERVIEW.pdf> (showing an average recall for search terms of 57% over 43 topics).

¹⁵ See Bruce Hedin et al., *Overview of the TREC 2009 Legal Track*, in NIST Special Publication: SP 500-278, The Eighteenth Text REtrieval Conference (TREC 2009) Proceedings, at 17 tbl. 6 (2009), available at https://trec.nist.gov/pubs/trec18/t18_proceedings.html.

¹⁶ See *id.*

¹⁷ See, e.g., Karl Schieneman and Thomas C. Gricks III, *The Implications of Rule 26(g) on the Use of Technology-Assisted Review*, 7 Fed. Cts. L. Rev. 239 (2013). But see Maura R. Grossman and Gordon V. Cormack, *Comments on “The Implications of Rule 26(g) on the Use of Technology-Assisted Review,”* 7 Fed. Cts. L. Rev. 285 (2013) (criticizing Schieneman and Gricks) (hereinafter “Comments paper”).

¹⁸ “Continuous Active Learning” and “CAL” are registered trademarks of Maura R. Grossman and Gordon V. Cormack. See <https://trademarks.justia.com/866/34/continuous-active-86634255.html> and <https://trademarks.justia.com/866/34/cal-86634265.html>, respectively.

¹⁹ See John Tredennick, *Ask Catalyst: What is the Difference Between TAR 1.0 and TAR 2.0?*, E-Discovery Search Blog (Oct. 26, 2016), available at <https://catalystsecure.com/blog/2016/10/ask-catalyst-what-is-the-difference-between-tar-1-0-and-tar-2-0/>.

²⁰ H5 is a San Francisco-based eDiscovery service provider. See <https://www.h5.com/about-us/>.

²¹ See JOLT study *supra* n.2, at 29-31 & n.145. See also Christopher Hogan et al., *H5 at TREC 2008 Legal Interactive: User Modeling, Assessment & Measurement*, in NIST Special Publication: SP 500-277, The Seventeenth Text REtrieval Conference (TREC 2008) Proceedings (2008), available at <https://trec.nist.gov/pubs/trec17/papers/h5.legal.rev.pdf>.

methods that are just as good if not better than H5's version, but, again, as of the time of this writing, there has been no independent verification of such claims. Overall, rule-based TAR methods have attracted less attention in eDiscovery than supervised machine-learning TAR methods, perhaps because they are more time and resource intensive, and can be applied only by a highly trained team.

Is keyword culling obviated by TAR?

TAR plays a similar role to keyword search in that it identifies a subset of documents that are likely to be responsive. Typically, all documents so identified by TAR are labeled by reviewers, in which case TAR fulfills the same role as keyword culling—to identify for review only documents that are likely to be responsive. *TAR recall* is the fraction of responsive documents identified by TAR for review. It should be distinguished from end-to-end recall, which is the fraction of responsive documents identified by TAR *that are also correctly labeled by reviewers*.

Like search-term recall, estimated TAR recall must substantially exceed 70% for end-to-end recall to approach 70%.

When using TAR, the primary reason for keyword culling—to save manual review effort by excluding documents from review—no longer exists because TAR eliminates the need to review the vast majority of non-responsive documents. Despite this fact, many eDiscovery practitioners believe it may still be necessary to use keyword culling to reduce the size of the collection prior to TAR, either because search is the only practical way to retrieve documents for review from the cloud or elsewhere, or because of capacity limits or the costs of processing and ingesting documents into the TAR platform. In any event, the number of documents that can be reviewed by TAR is orders of magnitude greater than what can be manually reviewed. If keyword culling must be used before TAR, one can exploit the recall-precision trade-off²² by using very broad keywords to achieve higher recall at the expense of lower precision than would be practical for a manual review. When search terms, TAR, and manual review are used together, the estimated end-to-end recall is equal to the product of the search-term recall, the TAR recall, and the reviewer recall. If it is possible to achieve 70% search-term recall, 70% TAR recall, and 70% reviewer recall, the estimated end-to-end recall will be $70\% \times 70\% \times 70\% = 34\%$. On the other hand, if it were possible to achieve 85% search-term recall, 85% TAR recall, and 85% reviewer recall, then the estimated end-to-end recall would be $85\% \times 85\% \times 85\% = 61\%$. Under the same assumptions, but without the keyword culling, estimated end-to-end recall for the TAR plus manual review would be $85\% \times 85\% = 72\%$. This is why, setting aside cost considerations, it is the better practice not to use keywords before TAR.

²² It is well known that in conducting searches, there is a trade-off between recall and precision; a broader search returning many documents will tend to have higher recall and lower precision, while a narrower search returning fewer but more targeted documents, will tend to have lower recall and higher precision. See TAR Glossary, *supra* n.3, at 27, 26.

Regardless of whether keyword culling or TAR is used, end-to-end recall should be estimated with respect to the gold standard of an exhaustive manual review of the entire collection had that been performed. To do so, the authors recommend using a *stratified sample*,²³ with separate strata for (i) documents labeled responsive by reviewers, (ii) documents labeled non-responsive by reviewers, (iii) documents excluded from review by TAR, and (iv) documents excluded from review by keywords. In this way, the gold standard provides an independent assessment to gauge effectiveness regardless of the method of search and review. The steps for taking and using stratified samples is described in detail in the Appendices to this chapter.

17.3 Machine-Learning Methods

The term “machine learning” encompasses two vastly different mechanisms with different purposes: (i) *supervised* machine learning, in which an algorithm is “taught” by showing it examples of human decision making, from which it learns to emulate that decision making on new examples; and (ii) *unsupervised* machine learning, in which the algorithm, without being taught, learns to organize or to draw inferences from the data without examples.

Logistic regression and support vector machines are state-of-the art TAR algorithms.

While certain supervised machine-learning algorithms and protocols have been shown to be effective for TAR,²⁴ certain unsupervised machine-learning methods have been referred to as (or included with) “TAR” offerings, but their effectiveness has not been established.

Within the context of TAR, a supervised machine-learning algorithm is “taught” using documents that have been labeled as responsive or non-responsive by a human reviewer. The algorithm learns to emulate the coding decisions of that reviewer, either by predicting categorically whether the reviewer would label an as-yet-unlabelled document as responsive or not, or by predicting the likelihood that the human reviewer would label the document responsive, and then ranking the documents in the collection from most to least likely to be responsive. When using supervised machine learning for TAR, the most important choices are (i) which particular machine-learning algorithm to use, and (ii) the protocol for selecting the

²³ A *stratified sample* is a sample formed as the aggregate of separate random samples for distinct sub-populations. For example, a stratified sample of Canadians might be formed by combining samples from each of the 13 Canadian provinces and territories in proportion to their respective populations.

²⁴ See JOLT study, *supra* n.2; Gordon V. Cormack and Maura R. Grossman, *Evaluation of Machine Learning Protocols for Technology-Assisted Review in Electronic Discovery*, Proceedings of the 37th Int’l ACM SIGIR Conference on Research & Dev. in Info. Retrieval (“SIGIR 2014”) 253 (2014) (hereinafter “SIGIR 2014”). See also Maura R. Grossman et al, *TREC 2016 Total Recall Track Overview*, in NIST Special Publication: SP 500-321, The Twenty-Fifth Text REtrieval Conference (TREC 2016) Proceedings (2016), available at <https://trec.nist.gov/pubs/trec25/papers/Overview-TR.pdf>; Adam Roegiest et al., *TREC 2015 Total Recall Overview*, in NIST Special Publication: SP 500-319, The Twenty-Fourth Text REtrieval Conference (TREC 2015) Proceedings (2015), available at <https://trec.nist.gov/pubs/trec24/papers/Overview-TR.pdf>; Bruce Hedin et al., *Overview of the TREC 2009 Legal Track*, in NIST Special Publication: SP 500-278, The Eighteenth Text REtrieval Conference (TREC 2009) Proceedings, at 17 tbl. 6 (2009), available at https://trec.nist.gov/pubs/trec18/t18_proceedings.html.

documents from which the algorithm will learn to make distinctions between responsive and non-responsive documents. Common supervised machine-learning algorithms for TAR include *support vector machines* (“SVM”), *logistic regression*, and *nearest neighbour* (“NN” or “1-NN”).²⁵

Supervised machine-learning methods first decompose each document into *features*,²⁶ which may be words, phrases, or word fragments, or *latent features* (e.g., concepts) generated using an unsupervised machine-learning method. Based on a (human) labeled set of training documents, the supervised machine-learning algorithm constructs a *model*²⁷ to predict whether a document would likely be labeled responsive or not, based on its features. Within the context of TAR, the two algorithms that have been demonstrated to be most effective are SVM and logistic regression, using words, phrases, or word fragments as features. A common, but much less effective method applies NN to latent features (i.e., concepts). A number of other methods, including “random forests,” “k-nearest neighbour (“k-NN”),” “neural networks,” and “deep learning,” have shown promise in the scientific literature, but have yet to be applied successfully to TAR.

Support vector machines employ a geometric interpretation of the document features, representing each document as a point in hyperspace—a space with many more than three dimensions—where each dimension represents the presence or absence of a feature. Responsive documents will tend to be nearer to one another in this space than to non-responsive documents, and vice versa. SVM attempts to find a separating hyperplane in this space that best separates responsive from non-responsive documents. The distance from the hyperplane indicates the likelihood that a document is responsive: If the document is far from the hyperplane on the side with mostly responsive documents, it is likely to be responsive; if the document is far on the opposite side it is very unlikely to be responsive; if the document is near the hyperplane, its responsiveness is uncertain.

Logistic regression employs a probabilistic interpretation of the document features, determining how well each feature distinguishes responsive from non-responsive documents, and combining these estimates for all of the features in a document to determine the probability that it is responsive or not. Logistic regression has been used for nearly a century as a data analysis tool; for example, to estimate the risk of cancer given a combination of environmental factors. For text categorization in general, and for TAR in particular, the scientific literature shows its effectiveness to be comparable or superior to the newer SVMs.

Nearest Neighbour is an inferior TAR algorithm.

²⁵ See TAR Glossary, *supra* n.3, at 31, 22, and 24, respectively, for definitions of *Support Vector Machine*, *Logistic Regression*, and *Nearest Neighbor*.

²⁶ See *id.*, at 17 for a definition of *features*.

²⁷ See *id.*, at 30, for a definition of *model*, which is also referred to as a “statistical model.”

In eDiscovery vernacular, “latent semantic indexing” or “latent semantic analysis” are synonyms for nearest neighbour.

Nearest neighbour (“NN”), like SVM, considers each document to be a point in hyperspace. The label for each document is deemed to be the same as the label of the nearest document in the training set. Because responsive documents tend to be more similar to each other than to non-responsive documents, and vice versa, chances are that the correct label for a document is the same as its nearest neighbour. NN does not work well for the high-dimensional spaces arising from the use of words, phrases, or word fragments as features,²⁸ and is generally applied to latent features derived from unsupervised machine-learning algorithms, such as latent semantic indexing or analysis (“LSI” or “LSA”), probabilistic latent semantic analysis (“PLSA”), or latent Dirichlet allocation (“LDA”).²⁹ In the eDiscovery literature, the method of NN with latent features is often referred to by the unsupervised latent-feature algorithm, meaning that when an article refers to “LSA” as a TAR method, it almost certainly means “NN with LSA features.” In eDiscovery, latent feature analysis is also sometimes referred to as *clustering*³⁰ or “concept clustering.”

NN with LSA features falls short of the state of the art for TAR, and its shortcomings have engendered “learned wisdom” that does not apply to state-of-the-art TAR methods; for example, that TAR cannot be applied to short documents or spreadsheets, that not all labeled documents should be used as training examples, that the training set must be randomly selected, that the training set must be carefully hand-picked by a senior lawyer, or that one or a few mislabeled training documents can result in a cascade of TAR errors (*i.e.*, “garbage in, garbage out”).

Other supervised machine-learning methods, notably those involving deep learning, have received much attention of late. At the time of this writing, none of these methods have been shown to be effective for TAR, or to be efficient enough for the size of the document collections that typically arise in eDiscovery.

17.4 TAR Protocols

Continuous Active Learning (“CAL” or “TAR 2.0”) is state of the art in circumstances in which every document is to be reviewed before production.

Simple Active Learning is state of the art in circumstances in which documents are to be produced without prior review.

Simple Passive Learning is inferior, whether or not every document is to be reviewed before production.

²⁸ See Rosalind B. Marimont and Monte B. Shapiro, *Nearest Neighbour Searches and the Curse of Dimensionality*, 24 IMA J. Appl. Math 59 (1979).

²⁹ See TAR Glossary, *supra* n.3 at 22 and 26, respectively, for definitions of “Latent Semantic Analysis”/“Latent Semantic Indexing,” and “Probabilistic Latent Semantic Analysis.”

³⁰ See *id.*, at 11 for a definition of *clustering*.

The generic term “TAR 1.0” encompasses both Simple Active Learning and Simple Passive Learning.

Supervised machine-learning *protocols* determine the method by which training examples are selected and labeled, how a supervised machine-learning algorithm is applied, and how the results are harvested. The primary protocols in use today are *Simple Passive Learning* (“SPL”), *Simple Active Learning* (“SAL”), *Continuous Active Learning* (“CAL”), or a hybrid approach involving two or more of these protocols.³¹ The machine-learning literature in the domain of information retrieval (*i.e.*, outside of the field of eDiscovery) refers to SPL as “supervised learning” and to SAL as “active learning,” without qualification. Little non-TAR literature concerning CAL exists; however, in the information-retrieval literature the term *relevance feedback* generally refers to a similar protocol used to refine keyword searches.³²

A TAR tool may be used for one of two distinct purposes:

1. Where the TAR tool identifies substantially all potentially responsive documents, which are then manually reviewed and labeled by human reviewers. Documents not identified by the TAR tool are deemed non-responsive following a validation process.
2. Where the TAR tool classifies every document as responsive or not based on a limited training set, and then labels the remaining documents according to the TAR tool’s ranking or classification process, without any further human review.

In the traditional context of a review for production in litigation, the first purpose is the most common: The producing party typically wants to review each and every document to be produced in order to label those that are confidential, to avoid producing non-responsive or privileged documents, and to gain an understanding of their case. In some instances, (*e.g.*, “second requests,”³³ or matters where parties are cost sensitive or the collection is likely to contain few privileged or confidential documents), the burden of reviewing every document outweighs the risk of producing non-relevant or privileged/confidential documents, occasioning the second approach.

For the first purpose, CAL is the most effective protocol of which the authors are aware. CAL is quite simple. At the outset, one or more responsive documents (typically real but possibly

³¹ See generally Comments paper, *supra* n.17 and SIGIR 2014, *supra* n.24. See also Maura R. Grossman and Gordon V. Cormack, *Continuous Active Learning for TAR*, Practical Law J. (Apr./May 2016), at 36 tbl.

³² See, *e.g.*, Christopher D. Manning et al., *Introduction to Information Retrieval* 178 (Cambridge University Press 2009 online version), available at <https://nlp.stanford.edu/IR-book/html/htmledition/relevance-feedback-and-pseudo-relevance-feedback-1.html>; see also TAR Glossary, *supra* n.3, at 28 (defining “relevance feedback” as a process in which the documents with the highest likelihood of relevance are labeled by a reviewer and added to the training set).

³³ “Second requests” are document requests propounded by the Antitrust Division of the U.S. Department of Justice or the U.S. Federal Trade Commission to investigate mergers and acquisitions that may have anticompetitive effects. Typically, second requests seek to gather information about the markets, sales, facilities, assets, and structures of the businesses that are parties to the intended transaction.

synthetic or contrived) are used to train the algorithm, which scores the remaining documents as to their likelihood of responsiveness. Some of the top-scoring documents are reviewed and coded as responsive or not, and used to train the algorithm. This process continues until very few of the top-scoring remaining documents are responsive, and those few responsive documents that remain are neither novel nor important.³⁴ Once this threshold is met, post-hoc validation will very likely show high recall.³⁵ In the unlikely event that validation uncovers that a substantial number of novel or important responsive documents have been missed, those documents can be used to train the algorithm, and the CAL process can be resumed and continued for as long as necessary.

For the second purpose, SAL is the most effective commercially available protocol of which we are aware. SAL starts out like CAL, using one or more responsive documents or synthetic documents to initiate the process. Then the machine-learning algorithm chooses from all the yet-to-be reviewed documents the ones from which it will likely learn the most. Two strategies that have been shown to be effective for SAL are *uncertainty sampling*³⁶ and *Scalable Continuous Active Learning* (“S-CAL”).³⁷ The selected documents are reviewed and labeled, and used to train the machine-learning algorithm. This process continues until a fixed number of documents (typically several thousand) have been reviewed and labeled, or until “stabilization” occurs, indicating that further training examples would not substantially improve the model constructed by the algorithm. At this point, the model is used to label all the as-yet-unlabeled documents. Post-hoc validation is conducted, as for any review effort. If validation shows that the estimated recall is too low, it can typically be increased (at the expense of decreased precision) by altering the threshold score above which the documents are presumptively labeled responsive.

While we do not recommend SAL for the first purpose, it has been and continues to be used for this purpose. Its use for this purpose follows the same course as for the second purpose, but the documents labeled responsive by the TAR tool (said to constitute the “review set”) are reviewed and labeled by a human reviewer. When used in this manner, estimated recall and precision must be calculated based on the final labels, not the presumptive labels assigned by the TAR tool. Otherwise, they do not reflect the recall and precision of the end-to-end review effort, but

³⁴ For a description of the CAL process, see Maura R. Grossman and Gordon V. Cormack, *Continuous Active Learning for TAR*, Practical Law J. (Apr./May 2016). See also Comments paper, *supra* n.17, at 289-91; SIGIR 2014, *supra* n.24, at 154.

³⁵ See Gordon V. Cormack and Maura R. Grossman, *Multi-Faceted Recall of Continuous Active Learning for Technology-Assisted Review*, in Proceedings of the 38th Int’l ACM SIGIR Conference on Research & Dev. in Info. Retrieval (“SIGIR 2015”) 555 (2015).

³⁶ With *uncertainty sampling*, the machine-learning algorithm selects the documents as to which it is least certain about relevance for review by a human reviewer. See TAR Glossary, *supra* n.3, at 33-34.

³⁷ For a description of S-CAL see Gordon V. Cormack and Maura R. Grossman, *Scalability of Continuous Active Learning for Reliable High-Recall Text Classification*, in Proceedings of the 25th Int’l ACM Conference on Info. & Knowledge Mgmt. (“CIKM 2016”) 1039 (2016). S-CAL is a patented method of Gordon V. Cormack and Maura R. Grossman. See *Systems and Methods for a Scalable Continuous Active Learning Approach to Information Classification*, US20160371262 (allowed 2020), available at <http://www.freepatentsonline.com/20160371262.pdf>.

only that of a single phase. (*i.e.*, the TAR component, not including the manual review that follows it).³⁸

In the academic literature, the term “supervised learning” typically refers to the process the authors denote as Simple Passive Learning or “SPL,” to distinguish it from SAL and CAL. In SPL, training examples are chosen not by the machine-learning algorithm, but solely by random selection and/or by a human reviewer. SPL is not an efficient training strategy, especially when responsive documents are sparse in the collection, and has yielded results inferior to SAL and CAL in studies of TAR-tool effectiveness.³⁹ The use of SPL has engendered much controversy as to how to choose the training examples, with requesting parties often insisting on micromanaging the process because of the belief that this method can easily be gamed. There is no empirical evidence that justifies ad-hoc efforts to choose the “best” training documents, or to use heroic efforts to ensure agreement on the labels used to teach the TAR algorithm. Nonetheless, if parties insist on using SPL (or any other TAR method for which evidence of efficacy is lacking), time is likely better spent ensuring valid computation of recall and precision estimates rather than in arguing about the mechanics of the process.

17.6 Measures of Effectiveness

Relevance is in the eye of the beholder.

One of the most difficult challenges of estimating review effectiveness is that *relevance* is a subjective concept.⁴⁰ Regardless of how clearly the criteria for responsiveness are set forth in a request for production (“RFP”) or in a subpoena, equally knowledgeable and well-intentioned expert reviewers—whether lawyers, state archivists, or intelligence analysts—will disagree on the relevance of a surprisingly large number of documents.⁴¹ Suppose that two separate, equally competent teams—Team A and Team B—were to independently review the same set of potentially responsive documents for responsiveness to an RFP or a subpoena, labeling each document as responsive or not. Scientific studies indicate that no more than about 70% of documents labeled “responsive” by Team A will also be labeled as “responsive” by Team B, and no more than about 70% of the documents labeled “responsive” by Team B will also be labeled

³⁸ See Comments paper, *supra* n.17, at 300-01.

³⁹ See, *e.g.*, SIGIR 2014, *supra* n.24.

⁴⁰ See Tefko Saracevic, *The Notion of Relevance in Information Science: Everybody Knows What Relevance is. But What is it Really?*, 8 Synthesis Lectures on Info. Concepts, Retrieval, and Services (2016); Peter Bailey et al., *Relevance Assessment: Are Judges Exchangeable and Does It Matter?*, in Proceedings of the 31st Int’l ACM SIGIR Conference on Research & Dev. in Info. Retrieval (“SIGIR 2008”) 667 (2008).

⁴¹ See, *e.g.*, Gordon V. Cormack and Maura R. Grossman, *Navigating Imprecision in Relevance Assessment on the Road to Total Recall: Roger and Me*, in Proceedings of the 40th Int’l ACM SIGIR Conference on Research & Dev. in Info Retrieval (“SIGIR 2017”) 5 (2017) (senior state archivist); JOLT study, *supra* n.2, at 10-14; Herbert L. Roitblat et al., *Document Categorization in Legal Electronic Discovery: Computer Classification vs. Manual Review*, 61 J. Am. Soc’y for Info. Science & Tech. 70 (2010) (lawyers); Ellen M. Voorhees, *Variations in Relevance Judgments and the Measurement of Retrieval Effectiveness*, 36 Info. Processing & Mgmt. 697 (2000) (intelligence analysts).

“responsive” by Team A.⁴² This observation is a consequence of the fact that relevance is subjective, not that either labeling is necessarily defective.

Effectiveness is quantified by “recall” and “precision” measured with respect to an independent gold standard.

Notwithstanding the subjectivity of relevance, the results of any competent review effort (e.g., Review A), can be used as a “gold standard” or “ground truth” against which to evaluate another review effort (e.g., Review B), provided that *the review to be evaluated is independent of the review from which the gold standard is derived*. The two most common measures of effectiveness in the science of information retrieval are recall and precision, which are predicated on the convenient fiction that the gold standard is infallible and therefore denotes “true relevance.” In theory, recall is the proportion of “truly relevant” documents that are identified by a review effort out of all of the “truly relevant” documents in the collection (i.e., a measure of the *completeness* of the review effort), while precision is the proportion of documents identified by the review effort that are “truly relevant,” (i.e., a measure of the *accuracy* of the review effort). In practice, recall is estimated as the proportion of documents labeled “relevant” in the gold standard that are also labeled “relevant” in the review effort being evaluated, while precision is estimated as the proportion of documents labeled “relevant” in the review effort being evaluated that are also labeled “relevant” in the gold standard. Assuming that both the gold standard and the review effort being evaluated are conducted by “ideal” but independent, exhaustive manual reviews, we would not expect either estimate to exceed the typical *positive agreement*⁴³ between two independent reviews, or about 70%.⁴⁴

Exact recall and precision scores can never be known, because they rely on an infallible gold standard which cannot be achieved.

Recall and precision estimates fall short of the exact values.

Recall and precision estimates are incomparable unless calculated according to precisely the same independent gold standard.

70% recall (or any other recall score) is meaningless in isolation, and inappropriate as an absolute standard of review effectiveness.

The scientific observation—that the recall and precision of a high-quality review effort, estimated with respect to an independent gold-standard, are typically in the neighborhood of 70%—has been widely misconstrued in the legal community to imply that 70% recall (without specifying the recall of what and how it is measured) is necessary and sufficient to establish the adequacy of an eDiscovery review effort. This mantra is espoused by requesting and producing

⁴² See *id.*

⁴³ *Positive agreement* is the proportion of documents labeled “relevant” by one review that are also labeled “relevant” by a second review. See TAR Glossary, *supra* n.3, at 25.

⁴⁴ See *supra* n.41.

parties alike, either to set an unreasonably high and potentially unachievable recall target, or to provide cover for inadequate or incomplete results.

As a preliminary matter, “70% recall” can never be established; at best a party can demonstrate an *estimated recall* of 70%. If the estimate is conducted with respect to an ideal, independent, gold standard, 70% estimated recall is indeed consistent with—but not proof of—an effective review effort.

A recall estimate is meaningless unless the gold standard is independent of the review effort for which it is calculated; an estimate of about 70% is merely the best that is likely to be achieved by *a single qualified reviewer, unassisted by quality control or by technology*.

Taking the following as an absurd example, an exact copy of the gold standard would achieve an estimated recall of 100% and an estimated precision of 100%, when evaluated with respect to itself as the gold standard. But it would achieve no more than 70% estimated recall and 70% estimated precision, when correctly evaluated with respect to an independent gold standard. Similarly, a copy of the gold standard as to which 30% of the documents labeled responsive were re-labeled non-responsive, would achieve an estimated 70% recall as compared to the original gold standard—and might be offered as evidence of an adequate production—but recall correctly estimated *with respect to an independent gold standard* would show the recall estimate of the revised gold standard to be at most $70\% \times 70\% = 49\%$.

Recall and precision can be used to compare review efforts only if calculated according to a common gold standard that is independent of the review efforts being compared.

We can compare the effectiveness of Reviews B and C by calculating their estimated recall and precision according to independent gold standard, Review A. If Review B yields substantially higher estimated recall and precision than Review C, according to gold-standard Review A, we can say with some confidence that Review B is more effective than Review C. But what if Review B yields higher estimated recall, while Review C yields higher estimated precision? Which is more effective depends on the relative importance of finding *all* responsive documents (*i.e.*, achieving high recall) versus finding *only* responsive documents (*i.e.*, achieving high precision). In eDiscovery, recall is generally the more important of the two measures, considering the aspirational goal of most productions: to find “all” responsive documents. However, if precision is very low (*e.g.*, much less than 50%), the responsive documents that are produced can be overwhelmed and therefore obscured by the non-responsive documents.⁴⁵

There is a trade-off between recall and precision, in both the conduct and evaluation of a review effort.

⁴⁵ For example, this was the case in *In re Domestic Airline Travel Antitrust Litig.*, MDL No. 2656, Misc. No. 15-1405 (CKK), 2018 WL 4441507 (D.D.C. Sept. 13, 2018), where one of the defendants’ productions had a recall of 97%, but a precision of only 17%, meaning that only about 600,000 of the 3.5 million documents produced were responsive, or that five out of every six documents produced were non-responsive.

The human reviewer (or the TAR tool) can be instructed to construe responsiveness broadly, thus increasing recall at the expense of precision, or to construe responsiveness narrowly, increasing precision at the expense of recall. For the purposes of eDiscovery, Review B is generally considered superior to Review C if it achieves substantially higher recall, while still achieving similar precision, or if it achieves similar recall, and substantially higher precision.

F₁ quantifies the extent to which a review effort achieves high recall and also high precision.

F₁ offers no insight into the recall-precision trade-off, and should not be used as the sole measure to evaluate review effectiveness.

F_1 is the harmonic mean of recall and precision.⁴⁶ Unlike the traditional arithmetic mean, F_1 emphasizes the lesser of the two, so that a high F_1 score is achieved only when recall and precision are both high, and approaches zero as either recall or precision approach zero. When recall and precision are equal, so is F_1 . The same F_1 score can be achieved if recall is slightly higher and precision is slightly lower, or vice versa; but if either is substantially lower than the other, F_1 will fall. These properties justify the observation that an exhaustive manual review is unlikely to yield F_1 greater than 70%, which would be achieved if recall and precision were both 70%, or thereabouts, as measured with respect to an independent gold standard.

17.5 The Use and Misuse of Statistics

Statistics play an essential role in many human endeavours, including eDiscovery, but are easily miscalculated, misapplied, or misinterpreted. A common use of statistics is to estimate a quantity that is too large to be counted, such as the number of people in a country that hold a particular opinion. Another common use is to estimate a quantity that can only be known in the future, such as how many people will contract a particular disease.

Numbers or percentages?

Statistical estimates may be expressed as numbers; *e.g.*, “approximately 760,000 Canadians believe that the world is flat,” or “about 617 Canadians will be diagnosed with cancer tomorrow.”⁴⁷ Or, they may be expressed as percentages; *e.g.*, “approximately 2% of Canadians believe the world is flat,”⁴⁸ or “about 0.0017% of Canadians will be diagnosed with cancer tomorrow.” To convert from one to the other, it is necessary to have an estimate of the size of the population of Canada, which, as of 2020, was about 38 million. Thus, the estimated number of Canadian flat-earthers is $2\% \times 38 \text{ million} \approx 760,000$, while the estimated daily percentage of Canadians that will be diagnosed with cancer is $617 \div 38 \text{ million} \approx 0.0017\%$. More generally, statistical estimates—whether numbers or percentages—tell an incomplete story absent an estimate of the size of the population from which they are derived.

⁴⁶ See TAR Glossary, *supra* n.3, at 16 for a definition of F_1 .

⁴⁷ See <https://www.cancer.ca/en/cancer-information/cancer-101/cancer-statistics-at-a-glance>.

⁴⁸ See <https://www.pollara.com/wp-content/uploads/2017/12/Pollara-Beliefs2019-RptF2.pdf>.

Counting found versus missed responsive documents.

In eDiscovery, the principal application of statistics is to estimate *how many responsive documents have been identified for production, and how many have been missed*. If the number missed is large compared to the number found, the review is incomplete; if the number missed is small compared to the number found, the review may be complete, depending on the novelty and importance of the missed responsive documents.

From these estimates we can derive an estimate of recall, which is the estimated number of responsive documents that are identified for production, as a percentage of the estimated overall number of responsive documents in the collection.⁴⁹ Recall substantially greater than 50% indicates that the number found is greater than the number missed. For example, an estimate of 75% recall indicates that three times as many responsive documents were found than were missed.

Precision is the estimated number of truly responsive documents identified for production, as a percentage of the total number of documents labeled responsive by the review effort. Precision greater than 50% indicates that the production contains more responsive than non-responsive documents. For example, an estimate of 75% precision indicates that a three quarters of the documents identified as responsive were actually responsive.

Remember that valid estimates of recall and precision rely on an independent gold standard of relevance; that is, they must be based on comparison to a *ground truth*. It is generally impractical to conduct an independent review of the entire collection of documents in order to form such a gold standard. *Statistical sampling methods* allow us to estimate, based on an independent review of many fewer documents, what the result would be, were all the documents to be independently reviewed to establish the ground truth.

Statistical sampling and blind review yield an independent gold standard.

To estimate the number of truly responsive documents identified for production, it is necessary to draw a random sample of the production set. To estimate the number of responsive documents that were missed, it is necessary to draw one or more random samples of all documents excluded from the production set for any reason—whether by keyword culling, by TAR, or by manual review. The documents labeled non-responsive by a search and review process are typically referred to as the “null set” (or sets). These different samples must be scrambled together and labeled by an independent reviewer, *who can have no knowledge of whether a document has been included or excluded from the production, or why*. This is

⁴⁹ The estimated overall number of responsive documents in the collection is simply the estimated number of responsive documents found plus the estimated number of responsive documents missed.

typically referred to as a *blind review*.⁵⁰ The results of this independent review can be used to properly estimate end-to-end recall, precision, and other measures.

Sample size, margin of error, and confidence level.

Estimates may differ from the hypothetical but unknowable true values for a variety of reasons. The first reason—*uncertainty in the conception of relevance*—is inherent in any recall or precision calculation, whether statistical or not. This error is mitigated by the use of an independent gold standard labeled diligently by one or more knowledgeable reviewers. It could be—but in practice rarely is—mitigated by labeling each document according to the majority vote of a panel of three (or more) reviewers.

Compounding the error arising from uncertainty in relevance is the *error arising from statistical estimation*. The terms of art used in information retrieval to describe this error are *bias*—the extent to which the estimate systematically misses the mark—and *random error*—the extent to which the estimate differs from the true value by chance. Bias is mitigated by random sampling and blind independent review. Random error is mitigated by increasing sample size: The larger the sample, the smaller the random error.

The terms of art used to describe the magnitude of random error are *margin of error* and *confidence interval*, which are always qualified by a *confidence level*.⁵¹ For example, we might say that “760,000 Canadians believe the earth is flat, with a margin of error of $\pm 76,000$, and a confidence level of 95%.” Equivalently, we might say that “the number of Canadians who believe the earth is flat falls in the confidence interval of 684,000 to 836,000 people, with a 95% confidence level.” Either statement may be interpreted as strong evidence (but never proof) that the true value is *somewhere between 684,000 and 836,000, and likely nearer to 760,000 than the extremes of the interval*. The confidence level of 95% has been widely, but somewhat arbitrarily, adopted as a standard for reporting results in the scientific literature.

Extreme caution must be used when expressing margin of error as a percentage, and combining it with a numerical value. If we say “760,000 Canadians believe that the earth is flat, with a margin of error of $\pm 10\%$,” it is ambiguous whether we mean $\pm 10\%$ of the 760,000 Canadians, or $\pm 10\%$ of all Canadians. When we translate these to numbers, we see that the first interpretation means $760,000 \pm 76,000$ people, while the second means $760,000 \pm 3.8$ million people (which would be a rather useless estimate). Yet eDiscovery practitioners often commit the fallacy of plugging the numbers into a statistical calculator that uses the latter interpretation and representing the result as if it were the former.

⁵⁰ In a *blind review* certain information is withheld from the reviewer to reduce bias and error, *e.g.*, the population from which a sample was drawn, the prior responsiveness label given to a document, etc. See Robert MacCoun and Saul Perlmutter, *Blind Analysis: Hide Results to Seek the Truth*, 526 *Nature* 187 (2015). Blind review is discussed further in Appendix A para. F.

⁵¹ See TAR Glossary, *supra* n.3, at 22-23, 12, and 12, respectively, for definitions of *margin of error*, *confidence interval*, and *confidence level*.

Margin of error is commonly miscalculated and misinterpreted.

For estimates of a small percentage of the population, margin of error cannot be predetermined.

For estimates of a very small percentage of the population, no reasonable estimate can be achieved by sampling.

The fallacy of misinterpreting margins of error has engendered a pervasive school of specious “learned wisdom” in eDiscovery circles, which can be identified by its vernacular. The typical eDiscovery practitioner is likely to encounter articles, briefs, and protocols calling for a “sample with a 2% margin of error at a 95% confidence level,” or simply a “2% sample,” or a “statistically significant sample.” Depending on what statistical calculator is used, the size of such a sample turns out to be 2,395, 2,399 or 2,405. This sample size ensures a margin of error of no more than $\pm 2\%$ of the Canadian population; in our flat-earth example, that would be $\pm 760,000$ Canadians (*i.e.*, $2\% \times 38 \text{ million} \approx 760,000$ Canadians). Thus, the margin of error would be 100% of the number of Canadian flat-earthers ($760,000 \pm 760,000$), not 2% of the flat-earthers as implied. To achieve a margin of error of at most $\pm 2\%$ of Canadian flat-earthers (15,200 people or 0.04% of all Canadians) would, according to the same assumptions and the same statistical calculator, require an absurdly large sample size of 518,000.

A more sophisticated analysis using a *binomial calculator*⁵² yields a somewhat smaller margin of error, which can be calculated *only after the sample has been reviewed*. Suppose that in a sample of 2,400 Canadians, 48 were to believe that the earth was flat. From the fact that 2% of the sample believed that the earth was flat, we can estimate that about 2% of all Canadians (*i.e.*, 760,000 Canadians) also believe that the earth is flat. Using a binomial calculator, we derive a confidence interval of 1.48% to 2.65% (or between 562,000 and 1,007,000) Canadians that believe the earth is flat, with 95% confidence.

Suppose that 2,400 Canadians were surveyed on a particular day to determine whether or not they had received a cancer diagnosis that day. In all likelihood, none of them would respond affirmatively. From this result one can conclude that being diagnosed with cancer on a given day is unlikely, but not much else. According to the binomial calculator, this result yields a confidence interval of 0% to 0.15% (or between zero and 57,000) Canadians who receive such a diagnosis on a given day. The sample provides strong evidence that fewer than 57,000 Canadians received a positive diagnosis on that particular day, but conveys no other information about the number of diagnoses on that particular day, or on any future day. Clearly, a method other than simple sampling of a single day’s data would be necessary to yield a meaningful

⁵² A *binomial calculator* is a statistical method used to calculate confidence intervals based on the binomial distribution (as opposed to the Gaussian or normal distribution). When there are few relevant documents in the sample, the binomial estimation is more accurate than the Gaussian or normal estimation; when there are many relevant and non-relevant documents in the sample then the binomial and the Gaussian (or normal) estimates are nearly identical. See TAR Glossary, *supra* n.3, at 9-10. An example of a binomial calculator can be found at <https://statpages.info/confint.html>.

estimate in this example. The reader is urged to eschew the received wisdom in eDiscovery that any quantity can be estimated “ $\pm 2\%$ ” with a sample of size approximately 2,400. When necessary, hire a statistical expert to assist in calculating margins of error, confidence intervals, confidence levels, and appropriate sample sizes; do not do this yourself.

Statistics for vetting and for validation.

The purpose of *vetting* a TAR tool is to estimate, in advance, how well it will work for a given review effort, by examining how well it has worked for one or more past review efforts. For these purposes, it is feasible to spend considerably more time and effort in constructing a gold standard, and in sampling to achieve a small margin of error, as compared to evaluating a single review effort. It is also feasible to repeat the process for many past reviews, so as to aggregate the results. For example, the TREC 2016 Total Recall Track simulated 34 reviews for different information needs using the email collection from Jeb Bush’s administration as Governor of Florida.⁵³ Six professional reviewers were employed for six weeks to create substantially complete gold standard labels for each of the information needs.

Using such test collections, it is possible to assess the relative effectiveness and reliability of the methods employed by the participants at TREC 2016, or to assess the effectiveness and reliability of new methods, relative to those employed at TREC 2016. Such an assessment yields a reasonable prediction as to how effective the methods would likely be when applied to a new matter. At the time of this writing, very few commercially available TAR tools have been vetted in this manner.

In contrast, the purpose of *validation* is to estimate the effectiveness of a particular review effort. The sample size that is used is limited by time and cost; a sample size of about 2,400 is often chosen for the specious reasons described above. While a sampling strategy involving the review of 2,400 documents cannot possibly yield an estimate of recall or precision with a margin of error of $\pm 2\%$, it can yield a sufficiently precise estimate that can, along with other insights, help to confirm the effectiveness of a particular review effort.

The validation strategy that the authors propose is to use *blind review of a combined stratified sample* to compute separate estimates of (i) the number of truly responsive documents identified for production, (ii) the number of responsive documents incorrectly excluded from production by reviewers, (iii) the number of responsive documents incorrectly excluded from review by the TAR process, and, if employed, (iv) the number of responsive documents incorrectly excluded by keyword culling, or by any other culling method that may have been used. *When combined, these estimates can provide an end-to-end estimate of the recall and precision of the entire search and review effort*, rather than just search-term recall or TAR recall, which can be

⁵³ Maura R. Grossman, Gordon V. Cormack, and Adam Roegiest, *TREC 2016 Total Recall Track Overview*, in NIST Special Publication: SP 500-321, The Twenty-Fifth Text REtrieval Conference (TREC 2016) Proceedings (2016), available at <https://trec.nist.gov/pubs/trec25/papers/Overview-TR.pdf>.

misleading because they both make it appear as if the recall of the production set is higher than it actually is.⁵⁴

The number of truly responsive documents identified in the production set will be a fairly large proportion—perhaps 70%—of all documents identified for production.⁵⁵ A sample of 400 documents is sufficient to estimate this proportion $\pm 5\%$ with 95% confidence, which is sufficient for these purposes. The number of responsive documents incorrectly excluded by reviewers will be a smaller but still substantial proportion—perhaps 10%—of the total number of documents excluded by the reviewers. If it is 10%, a sample of 400 documents is sufficient to estimate this proportion $\pm 3\%$ with 95% confidence, which is also sufficient for these purposes. The number of responsive documents incorrectly excluded by an effective TAR process is likely to constitute a very small proportion of the total number excluded; that is, the vast majority of excluded documents will be non-responsive. A sample of 1,600 may well reveal no responsive documents. In this event, we can conclude that, with 95% confidence, no more than 0.23% of the documents excluded by TAR are responsive. If search terms are used before TAR, an additional sample of 1,600 documents excluded by the keywords should be drawn. As with the TAR sample, this will, with high confidence, reveal whether an inordinate number of responsive documents have been missed by the keywords, but it can only provide a coarse upper limit of that number.

Assuming that no keyword culling was performed before TAR, we employ a combined sample size of 2,400 documents for the purposes of validation because that number is the one most often used in eDiscovery today to evaluate review efforts (albeit for the wrong reasons as described above), and because it should be sufficient and proportionate for most matters. The total size of the three samples— $400+400+1,600=2,400$ —(or $2,400+1,600=4,000$, if keywords were used)—can be adjusted on a case-by-case basis to balance the tension between validation review effort and the precision of the estimates achieved through the validation process. The estimates may be combined to yield estimates of recall and precision, which summarize the effectiveness of the end-to-end review effort. However, the separate estimates, and the samples themselves, can also offer deeper insights into potential shortcomings of various aspects of the review effort, and opportunities for mitigation. A model protocol for determining the effectiveness of an individual review effort using this method is provided in Appendix A.

Statistics like accuracy, elusion, and F_1 do not tell the whole story.

Some statistics may convey the illusion of—but no actual insight into—review effectiveness. *Accuracy* is simply the overall proportion of documents that are correctly labeled as either

⁵⁴ A step-by-step guide for taking a blind stratified sample to compute end-to-end recall and precision for an individual review effort is provided in Appendix A.

⁵⁵ It will not be *all* (*i.e.*, 100%) of the documents identified for production, because the reviewer and the gold standard will sometimes disagree, and because often, document families are produced in whole, such that non-responsive attachments are produced along with their responsive parent email (or vice versa).

responsive or non-responsive by the review effort (when combined together).⁵⁶ Suppose that 1% of the documents in a collection are responsive and that a vacuous review labels every one of them non-responsive. The accuracy of this review is 99%, although it does not identify a single responsive document. Accuracy is an uninformative measure of review effectiveness.

Elusion is the percentage of excluded documents that are responsive (*i.e.*, the percent of responsive documents found in the null set).⁵⁷ A small elusion number (*e.g.*, 1%) is commonly touted as evidence of the effectiveness of a review effort. But our vacuous review that identifies no documents would achieve an elusion of 1% even though no responsive documents were identified. *By itself, elusion conveys no useful information.* Combined with the size of the excluded set, elusion can be used to estimate the number (as opposed to the percentage) of excluded responsive documents. This number can be compared to the number of produced responsive documents, but only if both are sampled and measured with respect to the same gold standard, derived from a blind review, as proposed above and in the Appendices to this chapter.

F_1 combines recall and precision. Arguably, a high F_1 score is evidence that both precision and recall are high, but a low F_1 score gives no indication as to which of the two is low, or how to remedy the problem.

17.6 Establishing the Effectiveness of TAR Tools and Review Efforts

As noted in the previous sections, estimates of effectiveness scores such as recall, precision, and F_1 can be used to compare the efficacy of review efforts, given an independent gold standard. For the purpose of establishing the reasonableness of a particular review effort, it would be desirable to establish beforehand that the effort would be likely to achieve—and afterward that it did achieve—recall and precision comparable to or surpassing that of accepted practice or, better still, the hypothetical “ideal” of exhaustive manual review.

The effectiveness of TAR tools and protocols should be established in advance; the effectiveness of review efforts should be validated after the fact. Neither is a substitute for the other.

The same approach may be used to establish the effectiveness of a surgical procedure. First and foremost, it is necessary to use tools, procedures, and surgeons whose outcomes have been validated for similar patients with similar conditions; second, it is necessary to verify for every case that post-operative tests yield results consistent with a successful surgery. Unfortunately, few eDiscovery service providers have subjected their tools, procedures, or experts to anything resembling a “clinical trial,” or to any sort of rigorous evaluation. In the future, a consortium of service providers, regulators, and/or practitioners should conduct such product testing. In the meantime, practitioners are on their own to vet the TAR tools and methods they choose to use, and to establish the effectiveness of their individual review efforts after the fact.

⁵⁶ See TAR Glossary, *supra* n.3, at 8.

⁵⁷ See *id.*, at 15.

TREC offers standard test collections that can be used as a benchmark to assess the effectiveness and reliability of TAR methods.

The Text REtrieval Conference (“TREC”) of the National Institute of Standards and Technology (“NIST”) is an annual conference that evaluates information retrieval methods using *test collections* consisting of a common set of documents, information needs (*i.e.*, topics or RFPs), and independent gold-standard relevance assessments.⁵⁸ Each year, academic, government, and industry participants test various approaches to search and review on the test collections, and the results are reported in the TREC proceedings. The test collections are also made publicly available, so that they can be used to conduct experiments on methods and tools that were not represented at TREC.

Of particular interest here are the Legal Track, an evaluation campaign that ran at TREC from 2006 through 2011, and the Total Recall Track, which ran at TREC in 2015 and 2016. From 2006 through 2008, the Legal Track employed a collection of seven million documents from the Tobacco Litigation⁵⁹; from 2009 through 2011 it employed a collection of about 700,000 documents captured from Enron at the time of its collapse.⁶⁰ Each year, information needs were specified using a mock complaint and mock RFPs. A gold standard was created for a statistical sample of the document collection, and was used to estimate recall, precision, *F₁*, and other effectiveness measures for the TREC participants’ retrieved results. Results are published in the TREC proceedings⁶¹; the collections and tools, which can be used to evaluate the results of future experiments, are available from NIST, subject to a usage agreement.⁶²

The Jeb Bush collection from the TREC 2015 and 2016 Total Recall Tracks offers 290,000 documents, 44 topics, and independent gold standards for research into the effectiveness of TAR tools and methods.

⁵⁸ TREC was initiated in 1992. Its purpose is to support research within the information-retrieval community by providing the infrastructure necessary for large-scale evaluation of text-retrieval methodologies. In particular, the TREC workshop series has the following goals: (i) to encourage research in information retrieval based on large test collections; (ii) to increase communication among industry, academia, and government by creating an open forum for the exchange of research ideas; (iii) to speed the transfer of technology from research labs into commercial products by demonstrating substantial improvements in retrieval methodologies on real-world problems; and (iv) to increase the availability of appropriate evaluation techniques for use by industry and academia, including development of new evaluation techniques more applicable to current systems. For more information about TREC, see <https://trec.nist.gov/>.

⁵⁹ David D. Lewis et al., *Building a Test Collection for Complex Document Information Processing*, in Proceedings of the 29th Int’l ACM SIGIR Conference on Research & Dev. in Info. Retrieval (“SIGIR 2006”) 665 (2006).

⁶⁰ Gordon V. Cormack, *TAR Evaluation Toolkit Release 1.0.0* (copyright 2013), available at <https://cormack.uwaterloo.ca/tar-toolkit/>.

⁶¹ Nat’l Inst. of Standards and Tech., *TREC Proceedings*, available at <https://trec.nist.gov/proceedings/proceedings.html>. See also *TREC Legal Track*, available at <https://trec-legal.umiacs.umd.edu/>.

⁶² See Nat’l Inst. of Standards and Tech., *TREC Data*, available at <https://trec.nist.gov/data.html> (for most TREC datasets, including the Legal Track) and <https://plg.uwaterloo.ca/~gvcormac/total-recall/TR-group.html> (for the TREC Total Recall datasets).

The TREC Total Recall Track employed six different collections: (1) a set of 290,000 emails from Jeb Bush’s tenure as Governor of Florida; (2) a set of 500,000 postings from two “black hat” web forums; (3) a set of 900,000 news articles from the Columbia region of North America; (4) a set of 90,000 clinical records from a hospital intensive care unit; (5) a set of 400,000 emails from Tim Kaine’s tenure as Governor of Virginia; and (6) a set of 300,000 emails from Rod Blagojevich’s tenure as Governor of Illinois.⁶³ Of these collections, the first three are publicly available for research purposes. Of most interest to practitioners is the Jeb Bush collection, for which 44 information needs (*i.e.*, topics or RFPS) and full gold-standard relevance assessments are available, as well as three alternate, mutually independent sampled sets of gold-standard assessments for the 34 TREC 2016 Total Recall topics.

The authors have used these collections—and similar collections from other evaluation campaigns—in their research. Our paper comparing the effectiveness of rule-based and supervised machine-learning methods for TAR to exhaustive manual review was based on a retrospective analysis of the TREC 2019 Legal Track results.⁶⁴ Our paper comparing the effectiveness of different supervised machine-learning protocols for TAR reported a simulation study using the TREC 2019 Legal Track collection.⁶⁵ Our paper measuring the reliability of TAR methods used the TREC Total Recall collections, as well as several others.⁶⁶ Practitioners may use the same collections to evaluate commercial TAR tools and protocols.

The results of a prior review effort can be used as the basis for a test collection to evaluate new tools and methods.

Many practitioners also have access to the results of a prior review that they believe was competently performed. The documents, RFPs, and relevance assessments for that review can be used to compare the effectiveness of two new methods; for example, TAR versus a traditional manual review, or two TAR methods. For this purpose, the prior relevance assessments are used as the gold standard; in other words, the results of the prior review (presumably representing accepted practice) may be used to assess the effectiveness of the two new methods.

In the alternative, it is possible to compare the effectiveness of the prior review with that of a new review, by creating a new independent gold standard. In this case, it is necessary to create the gold standard by conducting a *blind review of a stratified sample of documents* so as to

⁶³ See *supra* n.53. See also Adam Roegiest et al., *TREC 2015 Total Recall Overview*, in NIST Special Publication: SP 500-319, The Twenty-Fourth Text REtrieval Conference (TREC 2015) Proceedings (2015), available at <https://trec.nist.gov/pubs/trec24/papers/Overview-TR.pdf>.

⁶⁴ See JOLT study, *supra* n.2.

⁶⁵ See SIGIR 2014, *supra* n.24.

⁶⁶ See Gordon V. Cormack and Maura R. Grossman, *Engineering Quality and Reliability in Technology-Assisted Review*, in Proceedings of the 39th Int’l ACM SIGIR Conference on Research & Dev. in Info. Retrieval (“SIGIR 2016”) 75 (2016).

populate a *confusion matrix*,⁶⁷ from which recall, precision, F_1 , and other measures can be estimated.

To compare review methods, A and B, the stratified sample must contain a number of documents selected at random from each of these four strata, as detailed in Appendix B para. B:

1. Documents labeled responsive by Review A and responsive by Review B;
2. Documents labeled responsive by Review A and non-responsive by Review B;
3. Documents labeled non-responsive by Review A and responsive by Review B; and
4. Documents labeled non-responsive by Review A and non-responsive by Review B.

The documents in the samples are shuffled together and labeled by an expert reviewer who is *given no information regarding the stratum from which each document was derived*. This blinding process is essential to the validity of any estimate derived from the gold standard. It is well established that, regardless of the motivation of the reviewer, they would be influenced by knowing from which stratum each document was derived, or even by knowing some factor associated with the stratum; for example, the order in which the documents are presented for review.⁶⁸

From the samples and the gold-standard assessments, one can estimate the number of documents labeled responsive by Review A that are also labeled responsive according to the gold standard (*i.e.*, precision), and the number of documents that are labeled responsive according to the gold standard that are also labeled responsive according to Review A (*i.e.*, recall). Precision and recall for Review B can be estimated in a similar manner, and compared to that of Review A. A similar strategy can be employed to compare more than two methods, using strata to capture all

⁶⁷ A *confusion matrix* (also sometimes referred to as a “contingency table”) is a two-by-two table listing the values for the number of true positives (*i.e.*, documents labeled responsive that are truly responsive), true negatives (*i.e.*, documents labeled non-responsive that are truly non-responsive), false positives (*i.e.*, documents labeled responsive that are truly non-responsive), and false negatives (*i.e.*, documents labeled non-responsive that are truly responsive). Virtually all of the standard evaluation measures in information retrieval are algebraic combinations of the four values in the confusion matrix. For an example of a confusion matrix and the formulas for the information-retrieval measures that may be computed from it, see TAR Glossary, *supra* n.3 at 12.

⁶⁸ Adam Roegiest and Gordon V. Cormack, *Impact of Review-Set Selection on Human Assessment for Text Classification*, in Proceedings of the 39th Int’l ACM SIGIR Conference on Research & Dev. in Info. Retrieval (“SIGIR 2016”) 861 (2016); Falk Scholer et al., *The Effect of Threshold Priming and Need for Cognition on Relevance Calibration and Assessment*, in Proceedings of the 36th Int’l ACM SIGIR Conference on Research & Dev. in Info. Retrieval (“SIGIR 2013”) 623 (2013); Mark D. Smucker and Chandra P. Jethani, *Human Performance and Retrieval Precision Revisited*, in Proceedings of the 33rd Int’l ACM SIGIR Conference on Research & Dev. in Info. Retrieval (“SIGIR 2010”) 595 (2010); William Webber et al., *Assessor Error in Stratified Evaluation*, in Proceedings of the 19th ACM Int’l Conference on Info. & Knowledge Mgmt. (“CIKM ’10”) 539 (2010); Mu-hsuan Huang and Hui-yu Wang, *The Influence of Document Presentation Order and Number of Documents Judged on Users’ Judgments of Relevance*, 55 J. Amer. Soc. Info. Science & Tech 970 (2004); Michael Eisenberg and Carol Barry, *Order Effects: A Study of the Possible Influence of Presentation Order on User Judgments of Document Relevance*, 39 J. Amer. Soc. Info. Sci. & Tech. 293 (1988).

combinations of relevance (and non-relevance) labels among the systems. Such an approach was used for the TREC 2008 through 2011 Legal Tracks.⁶⁹

One information need for one collection is sufficient as a proof of concept; multiple information needs and multiple collections are necessary to show reliability.

Each TREC collection entails massive volunteer effort to construct a gold standard for each of many information needs. Given the results of a method over all of these information needs, it is possible to estimate *reliability*—the percentage of the time that the method succeeds, for some definition of success—in other words, a measure of consistency. We might say, for example, that a method is reliable if its estimated recall equals or exceeds that of standard practice for at least 95% of all information needs. Further evidence of reliability can be demonstrated by showing that a method is reliable using different test collections.

While the literature provides evidence of the effectiveness and reliability of some search and review methods,⁷⁰ it is largely silent on the effectiveness and reliability of most commercially available TAR tools. The authors strongly advocate standardized testing to address these questions, but at the present time and for the foreseeable future, the practitioner's best recourse is to consider the available empirical evidence regarding the effectiveness of various methods, to consider critically service providers' claims as to which tools and methods they employ, to conduct proofs of concepts where feasible, to monitor the progress of their ongoing review efforts, and to evaluate their own results after the fact using stratified sampling and blind review, as set forth in the Appendices to this chapter.

At the very least, after-the-fact validation of a putative review effort should be based on the blind review of a stratified sample of documents, as detailed in the model validation protocol set forth in Appendix A, including both:

1. Documents labeled responsive by the review effort; and
2. Documents labeled non-responsive by the review effort.

From this blind review, estimates of end-to-end recall and precision—according to the gold-standard blind review—can be derived. As a rule of thumb, an estimate of 70% end-to-end recall and 70% end-to-end precision is consistent with an adequate review *when estimated under these circumstances*; however, cases of disagreement between the review result and the gold standard should be examined carefully to determine whether they indicate a systemic problem with the review effort. Moreover, responsive documents in the sample that were missed by the review effort should be examined for their novelty and importance to determine whether an identifiable class of responsive documents has been left behind.

A more convincing, but heretofore unused method for after-the-fact validation would be to have the sample reviewed by two independent, equally qualified, teams—one serving as the gold

⁶⁹ See *supra* n.61.

⁷⁰ See *supra* n.24.

standard and one as a control. If the estimated recall and precision of the actual review effort equals or exceeds that of the control (when judged according to the gold standard), one can reasonably conclude that the review was as effective as the control, which represents exhaustive manual review.

Finally, we note that recall and precision are not the sole determinants of effectiveness. For a typical information need, the majority of responsive documents—perhaps 70% or more—will be similar and perhaps only marginally important. A minority of the documents may be dissimilar to the majority, but very important. A review that systematically excludes these different-but-important documents is deficient, regardless of what estimate of recall and precision it achieves. As a preliminary matter, documents for which the review effort and the gold standard disagree should be examined for instances of excluded categories of responsive documents. Better still, auxiliary keyword searches of the collection and/or of the null set(s) should be conducted to try to find as many diverse examples of responsive documents as possible, to determine if any have mistakenly been excluded by the review effort. This approach is akin to audit sampling in accounting. It is not the same as random sampling, but it can be very effective when the prevalence or richness of the collection is low (*i.e.*, when there are few responsive documents to be found), and the nature of difficult-to-find responsive documents can be anticipated.

17.7 A TAR Checklist and Parting Words

What follows is a checklist of our key take-aways from this chapter.

- ***Not all search and review efforts are created equal.*** In choosing a TAR tool and protocol, the user should rely first and foremost on independent, controlled studies as an indicator of which algorithms and protocols are most effective and reliable. To the degree that such independent product evaluations are unavailable, users should test the efficacy and efficiency of TAR tools and protocols themselves, in advance of their use, so they are not first learning about their (in)effectiveness after they complete a review effort. The TREC test collections, the user’s own prior competent review efforts, and the procedures described in Appendix B provide the user with all they need to vet search and review methods. Users should not settle for scripted searches demonstrated by the service provider’s marketing team to make purchasing decisions. Most service providers use some version of the Enron dataset for the purposes of their product demonstrations. At the very least, in vetting search tools, users can consult the TREC 2009-2011 Legal Track Overview papers⁷¹ for possible information needs (*i.e.*, topics or RFPS) to use to test the system beyond the canned searches used in provider product demonstrations.
- ***Search-term recall, TAR recall, and reviewer recall are not the same things as end-to-end recall.*** Measuring only a single component of a review effort and presenting it as an estimate of the quality of the production set misrepresents the recall estimate so that it

⁷¹ See *supra* n.61.

appears better than it actually is. Each culling or review step introduces its own error and reduces the overall (*i.e.*, end-to-end) recall estimate; the combined effect of which can be substantial. Therefore, unless it is infeasible or too costly to do so, do not use keywords to cull before TAR. Your end-to-end recall will be considerably higher that way.

- ***Know your TAR algorithms.*** If at all possible, use an SVM or logistic regression. Do not confuse concept-clustering methods like LSI and LSA with TAR.
- ***Know your TAR protocols.*** Use a CAL protocol when you intend to review all potentially responsive documents before production; use a SAL or S-CAL protocol when you do not intend to do so. If you are using a SAL protocol, make sure that the two components that make up the F_1 score (*i.e.*, recall and precision) are *both* acceptable before you stop training.
- ***Manage expectations about outcome.*** Relevance is subjective and two equally qualified reviewers will label the same documents differently about 30% of the time. Thus, without the use of quality control measures and technology, it is challenging to obtain recall estimates above 70% for manual review, since that is the level at which reviewers typically agree. All review efforts entail a trade-off between recall and precision; the higher of one that is achieved, the lower the other will be. F_1 summarizes the two, but does not show which is lower or what can be done about that.
- ***Always measure recall and precision using a blind review of stratified samples that include examples of documents labeled responsive for any reason, and examples of documents excluded or labeled non-responsive for any reason.*** Only when measured in this manner, as described in Appendix A, and according to an independent gold standard, can an end-to-end recall on the order of 70% to 80% be said to demonstrate an adequate review effort (assuming that no novel or important documents are identified during the validation process). In Appendices A and B, the authors provide model protocols with step-by-step guidelines for evaluating a TAR Tool or review method (Appendix B) and for evaluating an individual review effort (Appendix A).
- ***Statistics and online statistical calculators are not toys.*** Beware of “learned wisdom” regarding margins of error, confidence intervals, confidence levels, and sample sizes; it is often wrong. For example, there is no such thing as a “statistically significant training set.” There is no such thing as a guarantee that you have found 95% of all relevant documents, or that you will find all of the relevant documents 95% of the time.

Statistical terms like “margin of error,” “confidence interval,” and “confidence level,” apply only to the estimate of a particular quantity from a particular sample; for example,

the *number* of relevant documents found, or the *number* of relevant documents missed. The mantra that a sample of 2,395 (or any other magic number of) documents will give you an estimate of these numbers $\pm 2\%$ is misleading. The margin of error for the estimate of a *number* should be expressed as a *number*; the margin of error for the estimate of a *percentage* should be expressed as a *percentage*. The notion that a *recall* estimate with “a margin of error of $\pm 2\%$, with a confidence level of 95%” can be derived from a simple sample of 2,395 from the collection is simply false. Such a sample can yield nothing more than a coarse estimate of the proportion of a certain type of document (*i.e.*, the “richness” or “prevalence” of responsive documents) in a collection, which, like the proportion of Canadian flat-earthers, is typically low. Compared to this low richness, a margin of error of $\pm 2\%$ is enormous.

Appendix A

A Model Validation Protocol for Examining the Effectiveness of an Individual Review Effort⁷²

- A. The review effort should be conducted using tools and protocols whose effectiveness has been established prior to the review effort under consideration. The review process should incorporate quality control and quality assurance measures to ensure that the selected review tool and protocol are performing as expected during the course of the review effort. Once the producing party reasonably believes that it has identified for production substantially all responsive non-privileged documents, it should conduct validation according to the sampling protocol described below. This Validation Protocol should apply to the review process regardless of whether keyword culling, manual review, or TAR (or any combination of them) was used by the producing party.
- B. The Document Collection (“Collection”) is defined as including all documents identified for review for responsiveness. This Validation Protocol assumes that the completeness or adequacy of the Collection has already been established.
- C. The Collection shall be partitioned into the following Subcollections as appropriate:
 1. Documents identified by the review effort as responsive to at least one request for production (“RFP,”) including any privileged documents, but not including family members of responsive documents, unless those family members are deemed to be responsive in their own right (“Subcollection C(1)”);
 2. Documents coded as non-responsive by a human reviewer, regardless of how the documents were selected for review (“Subcollection C(2)”);
 3. Documents excluded from manual review as the result of a TAR process (“Subcollection C(3)”). If the review process did not employ TAR, the Collection will not include Subcollection C(3).
 4. Documents excluded from manual review by keyword culling (“Subcollection C(4)”). If the review process did not employ keyword culling, the Collection will not include Subcollection C(4).

⁷² This Validation Protocol was adapted from the Order Regarding Search Methodology for Electronically Stored Information in *In re Broiler Chicken Antitrust Litig.*, Case No. 1:16-cv-08637, 2018 WL 1146371 (N.D. Ill. Jan. 3, 2018 (Special Master Maura R. Grossman)).

D. A sample shall be drawn consisting of the following:

1. 400 documents selected at random from Subcollection C(1) (“Subsample D(1)”);
2. 400 documents selected at random from Subcollection C(2) (“Subsample D(2)”);
3. 1,600 documents selected at random from Subcollection C(3), if TAR was used (“Sample D(3)”). If TAR was not used, there will be no Subsample D(3).
4. 1,600 documents selected at random from Subcollection C(4), if keyword culling was used (“Subsample D(4)”). If keyword culling was not used, there will be no Subsample D(4).

E. The sample sizes in paragraph D above are chosen so as to estimate the proportion of responsive documents in Subcollections C(1) and C(2) with a margin of error of no more than $\pm 5\%$, and to determine if collections C(3) and C(4) contain more than 0.23% responsive documents, at the 95% confidence level.

Should smaller margins of error be required, sample sizes may be increased as follows. Quadrupling the size of D(1) and D(2) from 400 to 1,600 will halve the size of the confidence interval, from $\pm 5\%$ to $\pm 2.5\%$. Doubling the size of D(3) and D(4) from 1,600 to 3,200 will halve the proportion of responsive documents that can be detected, from 0.23% to 0.12%.

Conversely, quartering the size of D(1) and D(2) from 400 to 100 will double the margin of error from $\pm 5\%$ to $\pm 10\%$, while halving the size of D(3) and D(4) from 1,600 to 800 will double the proportion of responsive documents that can be detected, from 0.23% to 0.46%.

F. The sample of documents comprised of the documents from Subsamples D(1), D(2), and, if present, D(3) and/or D(4), shall be combined into a single Validation Sample.

The Validation Sample shall be reviewed and coded by a subject matter expert (“SME”) who is knowledgeable about the subject matter of the litigation. This should be a lawyer who is familiar with the RFPs and the scope of relevance (*i.e.*, the claims and defences at issue in the litigation), but need not be a senior partner. The documents shall be presented to the reviewer in random or arbitrary order (*e.g.*, by MD5 hash value). During the course of the review of the Validation Sample, the SME shall not be provided with or have available to them any information concerning the Subcollection or Subsample from which any document was derived, the prior coding of any document, or the score

afforded to any document by the TAR tool. The intent of this requirement is to ensure that the review of the Validation Sample is blind, as necessary to form an independent gold standard for the purpose of evaluating the production.

G. Once the coding in Paragraph F has been completed, the producing party shall prepare a table listing each of the documents in the Validation Sample. For each document, the table shall include:

1. the Bates number of the document (for any documents previously produced), or a control/identification number (for any non-produced documents);
2. the Subsample from which the document came (*i.e.*, D(1), D(2), and, if present, D(3) or D(4));
3. the SME's responsiveness coding for the document (*i.e.*, responsive or non-responsive);
4. the SME's privilege coding for the document (*i.e.*, privileged or not privileged). If the document is coded as non-responsive, a privilege determination need not be made for that document. All documents in the Validation Sample coded as responsive and privileged may be withheld from production (subject to logging or other stipulated requirements), but shall be counted as responsive for the purpose of this Validation Protocol.

H. The following items shall be provided to the requesting party:

1. the table described in Paragraph G;
2. a copy of each responsive, non-privileged document in the Validation Sample that was identified for production but was not previously produced to the requesting party;
3. the statistics and recall estimate detailed in Paragraph J below.

I. The parties shall meet and confer to determine whether or not they agree that the recall estimate, and the quantity and nature of the responsive documents identified through the Validation Protocol, indicate that the review is substantially complete. If the recall estimate and the samples indicate that Subcollections C(2) and/or C(3) and/or C(4) still contain a substantial number of non-marginal, non-duplicative responsive documents as compared to Subcollection C(1), the review and quality control processes should

continue, and the Validation Protocol should be repeated, as warranted. If the parties are unable to agree on whether the review is substantially complete, or whether the Validation Protocol should be repeated, the parties shall seek the Court's intervention.

J. Method for Estimating Recall and Precision:

An estimate of recall and precision shall be computed to inform the decision-making process described in Paragraph H above; however, the absolute numbers in their own right shall not be dispositive of whether or not a review is adequate or substantially complete. Also of concern is the novelty and materiality (or conversely, the duplicative or marginal nature) of any responsive documents identified in Subsamples D(2) and/or D(3) and/or D(4). The estimates of recall and precision shall be derived as described below. It should be noted that, when conducted by an SME pursuant to Paragraph F of this Validation Protocol, a recall estimate on the order of 70% to 80% is consistent with, but not the sole indicator of, an adequate (*i.e.*, high-quality) review. A recall estimate somewhat lower than this does not necessarily indicate that a review is inadequate, nor does a recall in this range or higher necessarily indicate that a review is adequate; the final determination also will depend on the quantity and nature of the documents that were missed by the review effort.

Recall Estimation Calculation:

The number of responsive documents found \approx the size of Subcollection C(1) \times the number of responsive documents found in Subsample D(1) \div 400.

The number of responsive documents coded incorrectly \approx the size of Subcollection C(2) \times the number of responsive documents found in Subsample D(2) \div 400.

The number of responsive documents excluded by TAR \approx size of Subcollection C(3) \times the number of responsive documents found in Subsample D(3) \div 1,600 (if TAR is used; otherwise 0).

The number of responsive documents excluded by keyword culling \approx size of Subcollection C(4) \times the number of responsive documents found in Subsample D(4) \div 1,600 (if keyword culling is used; otherwise 0).

Estimated recall \approx the number of responsive documents found \div (the number of responsive documents found + the number of responsive documents coded incorrectly + the number of responsive documents excluded by TAR + the number of responsive documents excluded by keyword culling).

Precision Estimation Calculation:

Estimated precision \approx the number of responsive documents found \div the size of Subcollection C(1).

Appendix B

Model Validation Protocol for Vetting a TAR Tool or Comparing Two Different Review Methods

- A. One method to establish the effectiveness of a new review tool or protocol is to compare it to the results of a prior review that is known to be of high quality. To this end, it is necessary to estimate the recall and precision of both the prior review (Review Effort A) and the new review (Review Effort B), for the same information need(s), according to the same independent gold standard.
- B. The first step in the process is to identify the following four Subcollections:
1. Documents identified as responsive by both Review Effort A and Review Effort B (Subcollection C(1));
 2. Documents identified as responsive by Review Effort A but not Review Effort B (Subcollection C(2));
 3. Documents identified as responsive by Review Effort B but not Review Effort A (Subcollection C(3)); and
 4. Documents not identified as responsive by Review Effort A or by Review Effort B (Subcollection C(4)).
- C. From each of these Subcollections, a random sample of 600 documents is drawn, comprising Subsamples D(1), D(2), D(3), and D(4), respectively. These Subsamples are combined and reviewed blind by a subject-matter expert, as described above in Appendix A para. F.
- D. A sample size of 600 yields an estimate of each proportion with a margin of error of no more than $\pm 4\%$, with 95% confidence. Quadrupling the sample size will halve the margin of error, while quartering the sample size will double the margin of error.
- E. ***Recall, Precision, and F_1 Estimation:***

The number of responsive documents found by Review Effort A and Review Effort B \approx the size of Subcollection C(1) \times the number of responsive documents found in Subsample D(1) \div 600.

The number of responsive documents found by Review Effort A but not by Review

Effort B \approx the size of Subcollection C(2) \times the number of responsive documents found in Subsample D(2) \div 600.

The number of responsive documents found by Review Effort B but not Review Effort A \approx size of Subcollection C(3) \times the number of responsive documents found in Subsample D(3) \div 600.

The number of responsive documents found by neither Review Effort A nor Review Effort B \approx size of Subcollection C(4) \times the number of responsive documents found in Subsample D(4) \div 600.

Estimated Recall for Review Effort A \approx (the number of responsive documents found by Review Effort A and Review Effort B + the number of responsive documents found by Review Effort A but not Review Effort B) \div (the number of responsive documents found by Review Effort A and Review Effort B + the number of responsive documents found by Review Effort A but not Review Effort B + the number of responsive documents found by Review Effort B but not Review Effort A + the number of responsive documents found by neither Review Effort A nor Review Effort B).

Estimated Recall for Review Effort B \approx (the number of responsive documents found by Review Effort A and Review Effort B + the number of responsive documents found by Review Effort B but not Review Effort A) \div (the number of responsive documents found by Review Effort A and Review Effort B + the number of responsive documents found by Review Effort A but not Review Effort B + the number of responsive documents found by Review Effort B but not Review Effort A + the number of responsive documents found by neither Review Effort B nor Review Effort A).

Estimated Precision for Review Effort A \approx (the number of responsive documents found by Review Effort A and Review Effort B + the number of responsive documents found by Review Effort A but not Review Effort B) \div (the number of documents—whether responsive or not—in Subcollections C(1) + C(2)).

Estimated Precision for B \approx (the number of responsive documents found by Review Effort A and Review Effort B + the number of responsive documents found by Review Effort B but not Review Effort A) \div (the number of documents—whether responsive or not—in Subcollections C(1) + C(3)).

Estimated F_1 (for Review Effort A or Review Effort B) $\approx 2 \times$ (estimated recall \times estimated precision of the review effort) \div (estimated recall + estimated precision of the

review effort).

- F. If Review Effort B (the new tool or protocol) has comparable or superior recall and comparable or superior precision to Review Effort A, it is reasonable to conclude that Review Effort B is at least as effective as A. F_1 combines recall and precision into a single effectiveness measure, which is commonly reported, but may obscure the fact that Review Effort B has inferior recall or inferior precision, to Review Effort A.