



Technology-Assisted Review: Advice for Requesting Parties

While best practices are well-established for traditional search and review methods used in discovery, parties remain uncertain of how to navigate the technology-assisted review (TAR) process and what transparency entails in the TAR context. Requesting parties can better trust the results of TAR if they ask the right questions of producing parties and seek to gain more insight at the outset of the TAR process.



LEA MALANI BAYS

OF COUNSEL
ROBBINS GELLER RUDMAN & DOWD LLP

Lea focuses her practice on a range of e-discovery issues, from preservation through production. She is a member of The Sedona Conference Working Group on Electronic Document Retention and Production, and served as a Team Leader for the Sedona Conference's recently published *TAR Case Law Primer*.



LEXI J. HAZAM

PARTNER
LIEFF CABRASER HEIMANN & BERNSTEIN, LLP

Lexi represents plaintiffs in mass torts, *qui tam*, and complex class actions. She serves as Chair Elect of the American Association for Justice's Section on Toxic, Environmental, and Pharmaceutical Torts, and was a member of The Sedona Conference Working Group 1 Drafting Team for the *Primer on Technology-Assisted Review*.

The term technology-assisted review, or TAR, is now ensconced in the legal lexicon. But for many litigators, TAR remains poorly understood and haphazardly implemented. Spurred by ever-growing volumes of electronically stored information (ESI), TAR goes beyond traditional search methods used in discovery. TAR tools can rank, prioritize, and learn to classify ESI by extrapolating features from an attorney-coded subset of documents that are characteristic of relevance or non-relevance, and applying those decisions to an entire universe of documents.

It is vital for litigators to be familiar with the various TAR methodologies and platforms, and know how to assess the adequacy of a TAR process and measure its success. For requesting parties, it is also essential to understand:

- The benefits of TAR when compared with manual review.
- Unsettled questions about TAR.
- The current debate over the appropriate degree of transparency regarding the use of TAR.

- The key issues to address with producing parties proposing the use of TAR.



Search [Predictive Coding: A Primer](#) and [Predictive Coding: It's Here to Stay](#) for more on the technology behind TAR and common TAR tools.

BENEFITS OF TAR

Studies have demonstrated that, when properly used, TAR processes yield better results than traditional search methods, such as manual review and keyword searches (see, for example, Maura R. Grossman & Gordon V. Cormack, *Technology-Assisted Review in E-Discovery Can Be More Effective and More Efficient Than Exhaustive Manual Review*, 17 Rich. J.L. & Tech. 11, at 1-2 (2011); Herbert L. Roitblat and others, *Document Categorization in Legal Electronic Discovery: Computer Classification vs. Manual Review*, 61 J. Am. Soc'y for Info. Sci. & Tech. 70, 79 (2010); see also Box, *Continuous Active Learning*).

TAR has the potential to mitigate both under-inclusiveness and over-inclusiveness, thereby enhancing completeness and accuracy. In tests performed at the Text REtrieval Conference (TREC) to measure the effectiveness of TAR against traditional methods, TAR consistently outperformed manual review in both:

- **Recall.** This refers to the percentage of documents identified by the protocol as relevant within the entire document universe.
- **Precision.** This refers to the percentage of truly relevant documents within the set of documents identified by the protocol as relevant.

If a party using TAR can review only a fraction of documents and reach the same levels of recall and precision as manual review, TAR has the potential to dramatically reduce the hours spent on review to achieve the same result (Grossman & Cormack, 17 Rich. J.L. & Tech. 11, at 52-55), and lower costs for the producing party.

For example, many TAR tools give each document a relevance score based on its similarity to the training documents. The party can use the scores to prioritize the production of those documents the tool deems most likely to be relevant and eliminate the need for an exhaustive manual review of the documents least likely to be relevant. TAR might help the parties to agree to limit or even forgo further review of documents identified as relevant, if adequate privilege screens and claw-back provisions are in place.

This conservation of resources reduces the burden on the producing party, a key factor considered by courts in determining the appropriateness and proportionality of discovery requests under the Federal Rules of Civil Procedure (FRCP). Indeed, the 2015 amendments to the FRCP seem to contemplate, if not encourage, the use of advanced technologies to reduce the burden of discovery. The advisory committee notes acknowledge that as TAR methods continue to develop, and particularly in cases involving large volumes of ESI, courts and parties “should be willing to consider the opportunities for reducing the burden or expense of discovery as reliable means of searching [ESI] become available.” (2015 Advisory Committee Notes to FRCP 26.)



Search [APB to Requesting Parties: Prepare for Proportionality](#) for more on proportionality under the amended FRCP.

Additionally, requesting parties may consider using TAR to more efficiently review the documents they receive from producing parties. This can also help requesting parties understand and gain comfort with the use of this type of technology.

UNSETTLED QUESTIONS ON TAR USE

Perhaps the most persuasive argument in favor of using TAR is that the judiciary has embraced it. Since the seminal decision by Judge Peck of the Southern District of New York in *Da Silva Moore v. Publicis Groupe*, 287 F.R.D. 182 (S.D.N.Y. 2012), various other courts have encouraged or countenanced TAR use.

However, important unsettled questions remain regarding:

- The recall threshold that any review process, whether based on TAR or more traditional methods like search terms, should meet. Although at least one court has suggested that TAR should not be held to a higher standard than search terms (see *Rio Tinto PLC v. Vale S.A.*, 306 F.R.D. 125, 129 (S.D.N.Y. 2015)), a party's use of search terms might be held to a higher standard in the future and parties may be required to undertake some level of validation, especially with TAR available as an alternative (see, for example, *In re Lithium Ion Batteries Antitrust Litig.*, 2015 WL 833681, at *2-3 (N.D. Cal. Feb. 24, 2015)).
- Whether a court may compel a producing party to use TAR (see, for example, *Hyles v. New York City*, 2016 WL 4077114, at *3 (S.D.N.Y. Aug. 1, 2016) (declining to order a producing party to use TAR but noting that “[t]here may come a time when TAR is so widely used that it might be unreasonable for a party to decline to use TAR”).
- Whether a producing party may combine TAR with more traditional search methods, such as keyword filtering (compare, for example, *Bridgestone Ams., Inc. v. Int'l Bus. Machs. Corp.*, 2014 WL 4923014, at *1 (M.D. Tenn. July 22, 2014) (permitting plaintiffs to use TAR to further cull documents after keyword searching) with *Progressive Cas. Ins. Co. v. Delaney*, 2014 WL 3563467, at *10-11 (D. Nev. July 18, 2014) (declining to permit plaintiffs to use TAR to further cull documents after keyword searching)).

Additionally, as discussed in more detail below, an ongoing area of contention in discovery practice and in the case law relates to the level of transparency the producing party should accord the requesting party when using TAR, particularly concerning seed sets and control sets.

TRANSPARENCY IN THE TAR PROCESS

The need for transparency in discovery is not a novel concept, and the rise of vast volumes of ESI made a degree of transparency between the parties regarding electronic search parameters a practical necessity.

For example, courts have recognized that search terms are not privileged and should be disclosed to the other side (see,

for example, *Romero v. Allstate Ins. Co.*, 271 F.R.D. 96, 110 (E.D. Pa. 2010) (finding that information such as search terms, date ranges, and key custodians are not protected from disclosure by the work product doctrine or the attorney-client privilege)). Indeed, courts commonly expect that parties will discuss and negotiate search parameters, including specific search terms (see, for example, *Brown v. Tellermate Holdings*, 2014 WL 2987051, at *16, *21 (S.D. Ohio July 1, 2014); *Apple Inc. v. Samsung Elecs. Co.*, 2013 WL 1942163, at *2-3 (N.D. Cal. May 9, 2013); *DeGeer v. Gillis*, 755 F. Supp. 2d 909, 929 (N.D. Ill. 2010)).

BENEFITS OF TRANSPARENCY

As courts have observed, transparency about TAR engenders comfort with the process and helps to minimize future disputes (see, for example, *Da Silva Moore*, 287 F.R.D. at 192; *Burnett v. Ford Motor Co.*, 2015 WL 4137847, at *8-10 (S.D. W. Va. July 8, 2015)). A transparent process also increases the chances that the use of TAR will be accepted by the requesting party and the court (see *Progressive Cas. Ins. Co.*, 2014 WL 3563467, at *10-11 (indicating that the lack of transparency was a factor in denying a producing party's proposal to use TAR)).

Transparency about TAR at the outset of a case can be a good preventive strategy for both parties. If the receiving party does not believe the proposed TAR method will be adequate after thorough explanatory disclosures by the producing party, the parties can address the dispute with the court before incurring any expenses. Conversely, if the producing party refuses to provide disclosures on its TAR approach, there is an increased risk of costly do-overs.

Just as a producing party should be aware of the benefits of reasonable transparency regarding TAR, the requesting party should carefully consider the level of disclosure it seeks. Providing some documents for the seed set and reviewing a sample of documents for validation purposes is a reasonable objective. By contrast, requiring equal participation in every step of the TAR process might cause discovery to become bogged down in micro-disputes, delaying production and consuming resources to the point that the efficiency gains of TAR are defeated.

OBJECTIONS TO TRANSPARENCY

From a requesting party's perspective, there is little difference between disclosing the terms and other parameters used to retrieve relevant documents and disclosing the training documents used to retrieve relevant documents. Indeed, to the extent a producing party has concerns about potential misuse of non-relevant information disclosed during validation of TAR, that party can shield itself through confidentiality or attorneys-eyes only provisions in protective orders and orders under Federal Rule of Evidence 502(d). Further, counsel can negotiate an agreement not to use, print, or save information from training or validation documents for purposes beyond evaluating the efficacy of the TAR process and log highly sensitive documents in the same manner as privileged documents.




Search [FRE 502\(d\) Order](#) and [Privilege Waiver Clause with Claw-Back Provision](#) for more on protecting against waiver of the attorney-client privilege and work product protection.

However, in the TAR context, some commentators have objected to these disclosures based on work product or relevance issues, despite the availability of potential remedies.

Work Product

Some commentators have argued that the selection of seed documents used in training should be protected by the work product doctrine (see generally Hon. John M. Facciola & Philip J. Favro, *Safeguarding the Seed Set: Why Seed Set Documents May Be Entitled to Work Product Protection*, 8 Fed. Cts. L. Rev. 1 (2015)).

However, TAR seed sets reflect no more litigation strategy than search terms. Rather, they reveal simply that the documents are responsive to the document requests and that similar documents would likely also be responsive. This is a far cry from the narrow application of the document selection privilege, which some courts recognize to protect "something as subtle as the act of selecting or ordering documents" (*In re Trasyol Prods. Liab. Litig.*, 2009 WL 936597, at *3-4 (S.D. Fla. Apr. 7, 2009) (requiring a party seeking application of this privilege to "come forward with some evidence that disclosure of the requested



Achieving a sufficient understanding of the manner in which TAR will be used for the production of documents may require detailed discussions during the meet and confer process with opposing counsel, as well as their technologists.

documents creates a real, non-speculative danger of revealing counsel's thoughts").

Relevance

Some commentators have suggested that disclosing non-relevant documents in the training, validation, or control sets is wholly inconsistent with the discovery obligations imposed by other search methods. Although there are differing views on these matters, the disclosure of some non-relevant information to preserve the context surrounding relevant information is generally an accepted practice for many parties. For example, some courts have:

- Precluded producing parties from redacting non-relevant parts of documents on relevance grounds, and have held that all individual emails or attachments in a family must be produced even where only some are relevant (see, for example, *Families for Freedom v. U.S. Customs & Border Prot.*, 2011 WL 4599592, at *5 (S.D.N.Y. Sept. 30, 2011); *Howell v. City of New York*, 2007 WL 2815738, at *2 (E.D.N.Y. Sept. 25, 2007)).
- Required producing parties to disclose a sampling of documents not identified by agreed-on search terms to ensure that a significant percentage of relevant documents are not being missed, and allowed refinement of the terms to achieve better results (see, for example, *In re Lithium Ion Batteries Antitrust Litig.*, 2015 WL 833681, at *2-3).
- Instructed parties claiming that search terms are unduly burdensome to provide document hit counts and specific examples of irrelevant documents captured by a search (see, for example, *Finisar Corp. v. Nistica, Inc.*, No. 13-3345, slip op. at 4 (N.D. Cal. Dec. 12, 2014)).

These practices counterbalance the expectation that requesting parties should accept that, in light of resource and technology limitations, some (often many) relevant documents will not be identified or produced.

TAR DISCUSSION TOPICS

While TAR may be more efficient than the more traditional approaches, it requires parties to ask more questions at the outset regarding both the technology chosen and how the producing party will use it. Many attorneys are reluctant to acknowledge the deficiencies in traditional and more familiar approaches and may take an "ostrich" response to TAR. Yet understanding how a producing party will use TAR is as important as understanding the search terms to be used and the proper use of connectors, proximity searches, root expanders, case sensitivity, Boolean operators, wildcards, and similar techniques (see *L-3 Commc'ns Corp. v. Sparton Corp.* 313 F.R.D 661, 667 (M.D. Fla. Feb. 12, 2015) (emphasizing the importance of attorneys understanding the mechanics of keyword searching); *Nat'l Day Laborer Org. Network v. U.S. Immigration & Customs Enft Agency*, 877 F. Supp. 2d 87, 106-07 (S.D.N.Y. 2012) (noting the impossibility of assessing the adequacy of a search without knowing the terms used since seemingly minor decisions regarding the terms will have major consequences)).

Achieving a sufficient understanding of the manner in which TAR will be used for the production of documents may require detailed discussions during the meet and confer process with opposing counsel, as well as their technologists. Where a producing party proposes using TAR to find relevant documents for production, the parties should address certain issues early in the process. Productive discussions will facilitate informed decision-making about the TAR protocol and any anticipated court intervention. Topics to address include:

- The vendor and software to be used.
- The propriety of any pre-TAR culling.
- The proposed seed sets and training workflow.
- The individuals who will train the TAR algorithm.
- The documents that will be used to train the TAR algorithm.
- The criteria that will be used to determine whether the TAR algorithm is sufficiently trained.
- The metrics that will be used to measure the TAR tool's performance.

The level of transparency appropriate under the circumstances should be determined by the parties at the outset. As with all issues regarding search methodology, the parties should discuss these issues during the Rule 26(f) meet and confer process to avoid unnecessary motion practice and costly do-overs.



Search [Rule 26\(f\) Conference Checklist](#) and [Rule 26\(f\) Report and Discovery Plan](#) for more on key issues parties should address before and during the Rule 26(f) conference.

VENDORS AND SOFTWARE

Requesting parties can learn a great deal about a producing party's TAR process simply by knowing which vendor and software opposing counsel proposes to use.

Some TAR vendors provide a significant amount of information about their TAR abilities on their websites, including case studies, white papers, blog posts, online demonstrations, product briefs, and sample workflows. This publicly available information can provide enough insight for a requesting party to develop targeted questions regarding the proposed TAR process.

The parties should consider having a joint call with the vendor, during which the vendor can describe the TAR process and answer questions. This simple step can often save time and avoid miscommunications. Given the significant range in the types, quality, and cost of both TAR tools and the vendors offering them, the parties may consider enlisting the help of an outside expert to evaluate a TAR proposal.



Search [Considerations When Selecting an E-Discovery Vendor Checklist](#) for more on the issues a producing party should consider when engaging an e-discovery vendor.

PRE-TAR CULLING

Actions taken before any documents are loaded into the chosen TAR platform can often impact the final result. There are

Continuous Active Learning

A relatively new TAR method known as continuous active learning (CAL) might reduce certain disputes over transparency. CAL is different from other TAR models in that it does not use a control set or a discrete training set, but instead learns from each coding decision, continuously using the relevance feedback to rank and prioritize documents until there are fewer and fewer relevant documents identified for review.

However, a producing party's use of CAL should be driven by a confidence in its methodology rather than a desire to avoid transparency. Indeed, even with a CAL protocol, requesting parties will likely seek transparency on issues such as:

- Determining when to stop training, which is usually dependent on a drop in the proportion of relevant documents being returned in the training sets.
- Calculating recall, which is dependent on document sampling.
- Deciding on a discrete set of seed documents to begin training the algorithm, which the requesting party will likely want to review or supplement.



Search [Continuous Active Learning for TAR](#) for more on CAL protocols.

benefits and risks associated with filtering, which vary depending on the specific TAR tool employed. Understanding the culling proposed by the producing party will help the parties avoid pitfalls and analyze the final results. Some types of pre-TAR culling that the parties should discuss include:

- **Traditional culling.** Depending on the nature of the litigation and the amount of data available, some traditional filters, such as date ranges, custodians, and ESI sources, might be reasonable. This type of culling typically reduces the number of documents that are loaded into the TAR platform, which can decrease costs and increase the richness of the document set. The parties might also consider whether certain categories of relevant information are easily segregable and can be removed and produced without the use of TAR. For example, in a pharmaceutical product liability case, there is often a segregated file containing every submission to the US Food and Drug Administration, all of which would be relevant.
- **File type filtering.** This type of culling can help ensure certain file types are viewed properly. Because TAR uses language-based text to identify relevant documents, some file types might be less conducive to the use of TAR and require manual review, such as:
 - number-heavy spreadsheets;
 - photos and images;

- hard-copy documents;
- graphs; and
- password-protected or corrupted files.

If potentially problematic file types are not handled separately outside of the TAR process, the parties should consider whether the tool can identify the documents that have not been categorized and determine whether manual review is needed.

- **Search term narrowing.** Any form of culling runs the risk of filtering out relevant information. However, a producing party's use of pre-TAR search terms is especially problematic and might cause the results to fall below an acceptable level. To calculate a recall rate, the parties should separately consider the percentage of relevant documents missed by the search terms and by the TAR tool. The TAR tool might have a recall rate of 80%, a level many would consider acceptable. But if the recall rate of the search word culling done first was only 30%, then the overall recall rate would be 24%, far below an acceptable level.

SEED SETS AND WORKFLOWS

As discussed above, many TAR workflows include the use of an initial seed set to begin the training. The parties might have previously identified this discrete group of relevant documents through investigations or litigation, custodians, or targeted searches. The extent to which the TAR process relies on, or is impacted by, this initial set of relevant documents depends on the chosen tool, subsequent workflow, and the number of seed documents.

Workflows that rely heavily on a seed set to train the system might bias the system towards finding documents that are similar to the seed documents while ignoring relevant yet dissimilar documents. Parties can address the potential bias of a seed set by finding an appropriate number of training documents through random sampling or other means. Some models, such as CAL, ultimately train with significantly more documents, which might lessen the impact of each individual training document (see *Box, Continuous Active Learning*). Nonetheless, a seed set is often the first step in an iterative process and will initially impact the next set of training documents prioritized for review. The parties should therefore consider:

- Jointly selecting seed documents.
- Agreeing on criteria for a producing party to select seed documents.
- Disclosing non-privileged seed documents.

There might also be discrete sets of training documents beyond the seed set. The manner in which subsequent training documents are selected varies based on the chosen tool and workflow, and the methods a producing party uses can significantly impact the effectiveness and efficiency of the training. For example, training sets might consist of:

- A purely random or stratified sample, to counterbalance seed set bias.
- Prioritized documents that are most similar to the seed documents, so that these documents are coded first.

- Grey area documents, to help a producing party draw a more distinct line between relevant and non-relevant documents.
- Sample documents from categories of documents the system has not been trained on yet.
- Additional judgmental samples using search terms or other targeted selection criteria.
- Documents where the human coding differs from the system's predictive coding, to resolve coding discrepancies.

The parties should also discuss how additional or rolling collections will be handled.

INDIVIDUALS RESPONSIBLE FOR TRAINING

Documents should always be reviewed by individuals who clearly understand the issues in the case. If the proposed TAR workflow uses only a sampling of documents to train the system, the knowledge and competence of the individuals coding the training documents are even more important. Any incorrect coding or narrow definition of relevance is amplified because these decisions are extrapolated to the entire document universe. The number of reviewers can also impact the results of the trainings because inconsistent coding of training documents can lead to less effective training. Reviewers do not always agree on the relevance of the same documents and may have more or less inclusive interpretations of relevance that impact coding decisions.

Notably, proponents of CAL assert that the level of expertise and the number of reviewers are less important because the greater number of documents used to train a CAL algorithm lessens the impact of inconsistencies (see *Box, Continuous Active Learning*). Yet this distinction seems to undercut what has been one of the greatest selling points for TAR: achieving near-perfect coding by replicating coding decisions made by a subject matter expert across an entire document universe. Notwithstanding these concerns, research indicates that CAL can be an effective and efficient TAR tool.

TRAINING DOCUMENTS

In selecting a seed set or documents used to train the TAR system, parties must consider whether there should be any set limits or criteria for training documents and whether an experienced project manager should assess the chosen documents from a technology perspective.

In particular, the parties should consider asking a TAR vendor questions such as:

- Is there a recommended minimum amount of relevant text in the document for the tool to adequately recognize concepts or relationships in the language?
- How does the ratio of irrelevant text to relevant text impact the training, and is there a way to emphasize only the relevant text in the document?
- Is there a recommended minimum or maximum amount of text for each training document?
- Is metadata, such as date or email participants, used in the tool's analysis and, if so, to what extent? For example, what is the impact of training documents coded as irrelevant based purely on the date or author of the document?
- How does the tool treat numbers, and what role do numbers in the text play in determining relevance?
- How does the tool resolve potential problems with repeated text, such as email footers?
- Can the tool be trained to review foreign language documents, or should these documents be segregated for another type of review?

TRAINING COMPLETION CRITERIA

The parties should discuss the criteria that a producing party will use to determine when to stop the training and where to draw the line on which documents can safely be excluded without further review. Some TAR tools have a stabilization point, meaning that the recall and precision metrics have reached a



If the proposed TAR workflow uses only a sampling of documents to train the system, the knowledge and competence of the individuals coding the training documents are even more important. Any incorrect coding or narrow definition of relevance is amplified because these decisions are extrapolated to the entire document universe.

point where they are not improving despite additional training. Once a review has reached the stabilization point, there may be no reason for the producing party to continue coding training documents unless the parties determine how to change to the workflow to enhance the training.

Some TAR tools use percentages or rankings to demonstrate likely relevance, rather than a strict binary classification. Parties often agree that documents below a certain threshold will not be reviewed (other than as part of sampling) or produced due to the low likelihood of relevance. Similarly, the parties should consider whether documents that are above an agreed-on relevance threshold can be produced without further review. Because no tool is perfect, this might entail the production of some non-relevant documents, just as the tool might withhold some relevant documents. However, with appropriate protective orders, privilege screens, and non-waiver and claw-back provisions in place, the parties might agree that this is an acceptable consequence when weighed against the burden of additional review.

Because CAL models have no control set or distinct training phase, they are more likely to consider a sharp decline in the proportion of relevant documents in the batches prioritized for review as an indicator that an acceptable number of relevant documents has been identified. If the parties have a mutual understanding on how the system was trained and how the recall rate is calculated, it is more likely that the parties can come to an informed agreement on whether to stop training or what additional measures might enhance the results.

PERFORMANCE METRICS

There are several metrics parties commonly use to measure TAR performance. For requesting parties, the recall rate is usually the most important metric because it represents the proportion of relevant documents that have been identified as relevant. While the parties might agree on a minimum recall rate at the outset of the process, they should remain flexible on the calculation method to achieve a mutually agreeable outcome in a reasonable amount of time.

Understanding how the recall rate is determined is vital, particularly given the considerable debate on how recall should be calculated and how precise it should be (see John Tredennick, *Measuring Recall in E-Discovery Review, Part One: A Tougher Problem Than You Might Realize*, Catalyst (Oct. 15, 2014), available at catalystsecure.com). Understanding recall rates and the arguments on why one method might be better than another may require statistical analysis beyond many attorneys' skill sets. Yet parties should keep in mind that while recall is an important measurement, it is one part of the overall process that the parties must have confidence in.

The method used to calculate recall can vary. Control sets are the primary method parties use to calculate recall for more traditional forms of TAR. For CAL systems, parties use alternative measurements. Any method used to calculate recall should involve a large enough sample so that the margin of error is at an acceptable level. For example, a 95% confidence

level and a 2% margin of error is usually considered an effective sample size. The parties can use a sample size calculator to determine the appropriate sample (see, for example, Creative Research Systems, *Sample Size Calculator*, available at surveysystem.com/sscalc.htm). The sample set should be coded as perfectly as possible, because any recall metrics will be based on this coding, at least partially.

Control Sets

A control set is a sample of the document universe that has been manually coded. Parties can use the control set as the gold standard to test how well the TAR system has been trained to recognize the relevant control set documents as being relevant.

Control set documents are not used to train the TAR tool, because they would bias the testing. If a producing party creates a control set at the beginning of the TAR process, the party can use the set at any time, with any frequency, to measure recall.

There must be an adequate number of relevant documents in the set to adequately calculate recall. Therefore, the parties should understand that they may need to increase the number of documents in the control set if the richness (the proportion of relevant documents in the document universe) is low. Additionally, if the control set is created at the beginning of the process, the documents in the set may need to be updated if there are rolling collections.

Alternative Measurements

CAL systems do not use a control set and therefore employ other means to calculate recall. A common method is to take a sample of the documents that a CAL tool deemed non-relevant and calculate the proportion of relevant documents that the tool failed to detect. By itself, this analysis is not always particularly meaningful. For example, if a sample reveals that 2% of the discarded documents are relevant, this does not mean that the system only missed 2% of the relevant documents. Instead, the 2% figure is meaningful if the party already knows the percentage of relevant documents in the entire document universe. In that case, the party can use a calculation sometimes referred to as eRecall, which calculates recall based on both elusion (proportion of relevant documents left behind) and the richness of the document universe.

Instead of using a sample to estimate the number of relevant documents in the document universe in order to calculate recall, some CAL users propose sampling only the discard pile and adding the number of unretrieved relevant documents estimated from that sample to the number of documents coded as relevant during the review. This necessarily assumes that all of the coding decisions made by various reviewers over an extended period of time are correct, and might reflect an overstatement of relevant documents found.

The views expressed in this article are those of the authors and should not be attributed to Robbins Geller Rudman & Dowd LLP, Lieff Cabraser Heimann & Bernstein, LLP, or the firms' clients.