

The Ohio State Technology Law Journal

THE eDISCOVERY MEDICINE SHOW

**MAURA R. GROSSMAN, J.D., PH.D. * AND GORDON V. CORMACK,
PH.D.***

* Maura R. Grossman, J.D., Ph.D., is a Research Professor, and Gordon V. Cormack, Ph.D., is a Professor in the David R. Cheriton School of Computer Science at the University of Waterloo, in Ontario, Canada. Professor Grossman is also Principal of Maura Grossman Law, an eDiscovery law and consulting firm in Buffalo, New York, U.S.A. Professor Grossman's work is supported, in part, by the National Sciences and Engineering Council of Canada ("NSERC"). The opinions expressed in this piece are the authors' own and do not necessarily reflect the views of the institutions, organizations, or clients with which they are affiliated. The authors wish to thank Jason R. Baron for his thoughtful comments on drafts of this article.

CONTENTS

I.	INTRODUCTION	3
II.	MISAPPLICATION OF EFFECTIVENESS MEASURES .	9
III.	MEASURING THE WRONG QUANTITY (A/K/A SEARCHING UNDER THE STREETLIGHT).....	10
IV.	NON-BLIND ASSESSMENT	11
V.	MISLEADING STATISTICS	11
VI.	THE FALLACIES OF INCRECULITY, SPECIAL PLEADING, AND ASSUMING THE CONVERSE	14
VII.	THE BALL IS IN THE COURTS' COURT	14

I. Introduction

As recently as 100 years ago, harmful practices such as bloodletting were still advanced as sound medical practice by expert practitioners.¹ Bloodletting gradually fell into disfavor as a growing body of scientific evidence showed its ineffectiveness and demonstrated the effectiveness of various pharmaceuticals for the prevention and treatment of certain diseases.² Basking in the reflected glory of such scientifically proven medicines, unscrupulous purveyors of magical elixirs promoted their wares using pseudo-scientific evidence and testimonials from quacks and charlatans, presented along with free entertainment.³ These medicine shows persisted until, among other things, the Food and Drug Administration was given the authority to prosecute unsubstantiated therapeutic claims in 1938.⁴

eDiscovery methods, like therapeutics, are amenable to scientific evaluation.⁵ But practitioners and their “experts,” vendors, and clients often ignore empirical evidence, citing instead existing or past practice to justify, for example, culling electronically stored information (“ESI”) using untested search terms, establishing neither their

¹ See WILLIAM OSLER, *THE PRINCIPLES AND PRACTICE OF MEDICINE: DESIGNED FOR THE USE OF PRACTITIONERS AND STUDENTS OF MEDICINE* 759 (9th ed. 1921). *But see* Charles S. Bryan, *New Observations Support William Osler’s Rationale for Systemic Bloodletting*, 32 BAYLOR UNIV. MED. CTR. PROC. 372, 372-76 (2019) (arguing that Osler’s prescribed indications might be considered rational within the context of 19th century medicine).

² Timothy M. Bell, *A Brief History of Bloodletting*, 11 J. LANCASTER GEN. HOSP. 119, 122 (2016); Mark S. Lesney, *Patents & Potions*, in PHARMACEUTICAL CENTURY (2000), <http://www3.uah.es/farmamol/Thepercent20Pharmaceuticalpercent20Century/Ch1.html> [<https://perma.cc/JU7N-Q3MR>].

³ Sarah Halter, *Ruining Snake Oil: A Brief History of Fraud*, IND. MED. HIST. MUSEUM (Mar. 16, 2021, 9:12 AM), <https://imhm.org/blog/10252600>. [<https://perma.cc/ADB6-RSKU>].

⁴ See Federal Food, Drug, and Cosmetic Act, 21 U.S.C. § 300 (1938).

⁵ See, e.g., Gordon V. Cormack & Maura R. Grossman, *Navigating Imprecision in Relevance Assessments on the Road to Total Recall: Roger and Me*, in PROC. 40TH INT’L ACM SIGIR CONF. RES. & DEV. INFO. RETRIEVAL, 5-14 (2017), <https://dl.acm.org/doi/10.1145/3077136.3080812> [hereinafter *Navigating Imprecision*] [<https://web.archive.org/web/20210414225531/https://dl.acm.org/doi/10.1145/3077136.3080812>]; Gordon V. Cormack & Maura R. Grossman, *Evaluation of Machine-Learning Protocols for Technology-Assisted Review in Electronic Discovery*, in PROC. 37TH INT’L ACM SIGIR CONF. RESEARCH. & DEV. INFO. RETRIEVAL, 153-62 (2014), <https://dl.acm.org/doi/10.1145/2600428.2609601> [hereinafter *Evaluation of Machine-Learning*] [<https://web.archive.org/web/20210314183349/https://dl.acm.org/doi/10.1145/2600428.2609601>].

necessity nor their efficacy.⁶ Or, they use pseudo-science to promote various potions marketed as “Artificial Intelligence,” “AI,” “technology-assisted review,” or “TAR.” Or, they employ pseudo-science and various logical fallacies to impugn scientific studies that contradict their claims.⁷ Or, they point to the oft-cited Sedona Principle 6⁸ as justification to do whatever they please.⁹ Or, sometimes, even all of the above. Trade shows and other “educational” activities sponsored by vendors promote their wares, complete with pseudo-scientific results, testimonials, sponsored receptions, prizes, and hospitality suites. The Continuing Legal Education (“CLE”)

⁶ Cf. *William A. Gross Constr. Assocs., Inc. v. Am. Mfrs. Mutual Ins. Co.*, 256 F.R.D. 134, 136 (S.D.N.Y. 2009) (“where 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 as to the words and abbreviations they use, 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, 257, 262 (D. Md. 2008) (noting “[t]he only prudent way to test the reliability of the keyword search is to perform some appropriate sampling of the documents,” and finding Defendants’ “failed to demonstrate that the keyword search they performed on the text-searchable ESI was reasonable, . . . the keywords selected nor the qualifications of the persons who selected them to design a proper search; [and that] they failed to demonstrate that there was quality-assurance testing.”); *United States v. O’Keefe*, 537 F. Supp. 2d 14, 24 (D.D.C. 2008) (“[w]hether search terms or ‘keywords’ will yield the information sought is a complicated question involving the interplay, at least, of the sciences of computer technology, statistics and linguistics. . . . Given this complexity, for lawyers and judges to dare opine that a certain search term or terms would be more likely to produce information than the terms that were used is truly to go where angels fear to tread.”); *In re Seroquel Prods. Liab. Litig.*, 244 F.R.D. 650, 660 n.6, 662 (M.D. Fla. 2007) (criticizing defendant’s search and failure to provide information “as to how it organized its search for relevant material, [or] what steps it took to assure reasonable completeness and quality control,” and observing that “while key word 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.”).

⁷ Thomas R. Holtz, Jr., *Argumentation and Logical Fallacies: Reasonable and Unreasonable Approaches to Thinking*, U. MD., <https://www.geol.umd.edu/sgc/lectures/fallacies.html> [<https://perma.cc/SM6S-AB9T>].

⁸ Sedona Principle 6 states that “[r]esponding parties are best situated to evaluate the procedures, methodologies, and technologies appropriate for preserving and producing their own electronically stored information.” The Sedona Conference, *The Sedona Conference Principles, Third Edition: Best Practices, Recommendations & Principles for Addressing Electronic Document Production*, 19 SEDONA CONF. J., 118-30 (2018), <https://thesedonaconference.org/sites/default/files/publications/Thepercent20Sedonapercent20Principlespercent20Thirdpercent20Edition.19TSCJ1.pdf>.

⁹ Craig Ball, *Sedona Principle Six: Overdue for an Overhaul*, CRAIGBALL (Oct. 10, 2014), <https://craigball.net/2014/10/10/sedona-principle-six-overdue-for-an-overhaul/> [<https://perma.cc/53LX-EJYC>].

industry and the trade press often echo these testimonials, failing to discriminate between practice and sound practice—let alone best practice—or between science and pseudo-science. So far, neither the courts nor any other authority has taken up the mantle, leaving parties to fend for themselves in the eDiscovery Wild West.

The eDiscovery task is one of information retrieval, which has been a discipline of scientific study for more than 70 years.¹⁰ Those 70 years have seen considerable advances in information retrieval methods, as well as ways to measure their effectiveness.¹¹ One of the earliest and most primitive approaches to information retrieval is the so-called Boolean search, which retrieves only documents containing certain search terms or combinations of search terms, as specified by a manually constructed query.¹² “Keyword culling” in eDiscovery almost always employs Boolean search, with minor enhancements that take into account the order and proximity of the terms within a document.¹³ Parties sometimes exchange “hit reports,” quantifying the number of times the search terms appear in the document collection, but these reports provide little, if any, useful information about the quality of the search terms.¹⁴

The most common—but certainly not the only—way to quantify the effectiveness of an information retrieval effort is to estimate recall and precision.¹⁵ These measures rely on the convenient fiction of binary relevance: that there is a “ground truth,” and that any particular document either is or is not relevant to the information need that occasioned the information retrieval effort.¹⁶ If we could somehow know with certainty the relevance of every document in the collection

¹⁰ Mark Sanderson & W. Bruce Croft, *The History of Information Retrieval Research*, 100 PROC. IEEE SPECIAL CENTENNIAL ISSUE 1444-51 (2012), <https://ieeexplore.ieee.org/iel5/5/4357935/06182576.pdf>.

¹¹ *Id.*

¹² CHRISTOPHER D. MANNING ET AL., INTRODUCTION TO INFORMATION RETRIEVAL 1-18 (Cambridge Univ. Press 2008).

¹³ *Id.* at 14-15.

¹⁴ *E-Discovery Update*, McGUIREWOODS (July 8, 2016), <https://www.mcguirewoods.com/client-resources/Alerts/2016/7/E-Discovery-Update-Producing-Search-Term-Hit-Reports> [<https://perma.cc/XGY6-6RM4>].

¹⁵ Manning et al., *supra* note 12, at 154-57.

¹⁶ *Id.* at 164-67.

to be searched, recall would be the percentage of all relevant documents that were retrieved, while precision would be the percentage of all retrieved documents that were relevant.¹⁷ But we cannot know with certainty—or even with high confidence—the relevance of every document, or even the number of relevant documents in the collection, and thus recall.¹⁸ At best, we can estimate them with varying—and typically large—margins of error.¹⁹ Any method whose operation depends on a purportedly precise estimate of recall involves pseudo-science and bogus statistics.²⁰

While some leading information retrieval scientists eschew the use of recall altogether,²¹ it is generally accepted that recall and precision can be used to gauge the *relative effectiveness* of competing information retrieval methods, provided that recall is measured under the same conditions, using the same information need, the same document collection, and the same independently derived relevance assessments.²² The National Institute of Standards and Technology’s (“NIST’s”) Text REtrieval Conference (“TREC”) was founded in 1992, precisely to mount a heroic collaborative effort, among academic, industry, and government researchers, to create such conditions—conditions that cannot possibly be reproduced within the context of a particular eDiscovery effort.²³ As such, efforts like TREC are useful to compare the effectiveness and reliability of different information retrieval methods, not to determine the effectiveness of any particular information retrieval effort.²⁴ Under TREC-like conditions, 65 percent recall and 65 percent precision represents a good result—approximately what we would expect if the entire

¹⁷ *Id.* at 154-57.

¹⁸ *Id.* at 164-67.

¹⁹ William Webber, *Approximate Recall Confidence Intervals*, 31 ACM TRANSACTION INFO. SYS. 1 (2013).

²⁰ Justin Zobel et al., *Against Recall: Is It Persistence, Cardinality, Density, Coverage, or Totality?*, 43 ACM SIGIR FORUM 3 (2009).

²¹ *Id.*

²² Ellen M. Voorhees, *Variations in Relevance Judgments and the Measurement of Retrieval Effectiveness*, 36 INFO. PROCESSING & MGMT. 697 (2000).

²³ ELLEN M. VOORHEES & DONNA K. HARMAN, TREC: EXPERIMENT AND EVALUATION IN INFORMATION RETRIEVAL, 63 (Ellen M. Voorhees & Donna K. Harman eds., 2005); *see also* TREC Proc., 1992–present, *available at* <https://trec.nist.gov/proceedings/proceedings.html> [<https://perma.cc/K4TR-7QLP>].

²⁴ Voorhees, *supra* note 22.

collection were assessed by one subject matter expert (“SME”), with recall and precision estimated using the independent assessments of a second SME.²⁵

65 percent recall and 65 percent precision are not necessarily the best possible results. If, for example, the majority vote of independent assessments by three SMEs were used, they might achieve on the order of 75 percent recall and 75 percent precision, according to the same evaluation conditions.²⁶ For similar reasons, it is possible for certain specific TAR methods—where relevance is determined using input from both an SME *and* artificial intelligence—to achieve better effectiveness than a single SME alone.²⁷

Much ink has been spilled over the entirely irrelevant question of whether or not Boolean search is a form of TAR. Boolean search, like TAR, is an information retrieval method.²⁸ Boolean search using manually constructed queries (as, for example, employed for keyword culling in eDiscovery), unlike certain TAR methods, generally cannot achieve 70 percent recall and 70 percent precision.²⁹ Perhaps one Boolean query can yield 70 percent recall and 20 percent precision, while another for the same information need can yield 20 percent recall and 70 percent precision.³⁰ But finding a single query that can

²⁵ *Id.*

²⁶ The majority vote of three independent assessments is more likely to agree with another independent assessment than any of the individual assessments. Thus, when another independent assessment is used to evaluate recall and precision, both of these measures will likely be higher for the majority vote than for any individual. This is an application of the well-known statistical phenomenon referred to as *regression to the mean*, “first noted by Sir Francis Galton that ‘each peculiarity in man is shared by his kinsmen, but on the average to a less degree.’” BRIAN S. EVERETT & ANDERS SKRONDAL, *THE CAMBRIDGE DICTIONARY OF STATISTICS* 363 (4th ed. 2010).

²⁷ See *Navigating Imprecision*, *supra* note 5, at 5-14; 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 (2011) [hereinafter *TAR in E-Discovery*].

²⁸ Maura R. Grossman & Gordon V. Cormack, *A Tour of Technology-Assisted Review*, in PERS. PREDICTIVE CODING & OTHER ADVANCED SEARCH METHODS L. PRAC. 63 (Jason R. Baron et al., eds., 2016).

²⁹ Eero Sormunen, *Extensions to the STAIRS Study—Empirical Evidence for the Hypothesised Ineffectiveness of Boolean Queries in Large Full-Text Databases*, 4 INFO. RETRIEVAL 257 (2001).

³⁰ *Id.*

yield both 70 percent recall and 70 percent precision is likely impossible.³¹

To make matters worse, eDiscovery medicine shows frequently promote the sequential use of two or more information retrieval methods, including Boolean search, TAR, and manual review.³² The net effect of this concoction is to achieve considerably lower recall than any of its constituent parts: when multiple information retrieval methods are used in sequence, overall recall is the product of the recall for each constituent method.³³ If keyword culling were to achieve 70 percent recall, the TAR tool were to achieve 80 percent recall, and manual review were to achieve 75 percent recall, the recall of a review effort combining them in sequence would be 70 percent \times 80 percent \times 75 percent = 42 percent. It is possible to quibble with the numbers presented here, but not with the fact that each constituent part is imperfect, and that overall or end-to-end recall is considerably less than the weakest link in the chain.³⁴

When applied sequentially, information retrieval methods—whether Boolean search, TAR, or manual review—will always yield inferior recall.³⁵ Yet the medicine shows would have us believe that we need to consider only the TAR-tool component in our recall calculations, ignoring relevant documents excluded by keyword culling and/or by post-TAR manual review. This is, at best, an extreme case of moving the goalposts, but more likely, a form of legerdemain.

No doubt purveyors and pundits will attack these arguments and examples with special pleading, *argumentum ad hominem*, appeals to common practice or to Sedona Principle 6, strawmen, testimonials, and more pseudo-science—anything but a valid scientific experiment published in a peer-reviewed venue.

³¹ *Id.*

³² Maura R. Grossman & Gordon V. Cormack, *Comments on “The Implications of Rule 26(g) on the Use of Technology-Assisted Review”*, 7 FED. CTS. L. REV. 285, 300-01 (2014) [hereinafter *The Implications of Rule 26(g)*].

³³ *Id.*

³⁴ *Id.*

³⁵ *Id.*

The following sections outline several of the methods of deception that have been used—and will continue to be used—to promote inferior methods and to cast aspersions on superior ones, until someone is ready to bell the cat.

II. Misapplication of Effectiveness Measures

Statistical measures are commonly used to evaluate the efficacy of medical treatments as well as information retrieval methods.³⁶ A cancer treatment might be judged by its five-year survival rate, while an antibiotic might be judged by the probability that it will cure a specific infection. Neither of these statistics can possibly be measured on a case-by-case basis at the time of treatment; they are instead used to establish which treatments have the best chance of success, so that a reasonable choice can be made when the need arises, considering not only efficacy but cost, availability, side effects, and so on. On a case-by-case basis, observations like symptom abatement, temperature reduction, and the absence of a rash are indicators—but not proof—of successful treatment. These indicators of success often occur within a day or two of antibiotic treatment, but it generally takes several more days to eliminate the infection. If treatment is discontinued prematurely, the infection will likely return, in a form more resistant to the treatment than before. It would cause genuine harm to continue antibiotics only so long as symptoms persist, and then to repeat the treatment when symptoms recur. It might be possible to conduct a laboratory blood test to see with reasonable certainty when the infection was gone, but the test itself would entail more burden, cost, and delay than simply completing the prescribed course of antibiotics.

Recall and precision were conceived by scientists to measure the *comparative effectiveness* of information retrieval methods, not the success of a particular information retrieval effort, and especially not as the sole determinant of when a particular information retrieval effort may be discontinued.³⁷ It is infeasible to compute a precise and accurate estimate of recall within the context of a particular

³⁶ See, e.g., DAVID G. KLEINBAUM & MITCHEL KLEIN, SURVIVAL ANALYSIS (Gail et al., eds., 3rd ed. 2010).

³⁷ See Voorhees, *supra* note 22.

eDiscovery effort,³⁸ and even if such an estimate were feasible, it would not inform the question of whether or not the review, if continued, would find enough additional relevant documents to justify the additional effort, a proportionality consideration that is included in the description of the scope of discovery set forth in Federal Rule of Civil Procedure 26(b)(1), and in the “reasonable inquiry” requirement of Federal Rule 26(g)(1).

The conflation of efficacy and success has led to the absurd propositions that no treatment or information retrieval method can be determined to be reasonable or unreasonable in advance, that no irreparable harm occurs when a treatment or search method fails in a particular case, and that success in a particular case can be solely determined by an estimate of a summary measure like body temperature or recall, without regard to symptoms or the quality of the production.

III. Measuring the Wrong Quantity (a/k/a Searching Under the Streetlight)

Recall and precision measure the *end-to-end effectiveness* of an information retrieval effort.³⁹ As noted above, the scientific literature indicates that 70 percent to 80 percent recall is achievable by certain TAR methods.⁴⁰ Many TAR methods employ a software tool in conjunction with manual review and/or prior Boolean keyword culling.⁴¹ The recall of the software tool, in isolation, must be considerably higher—perhaps 85 percent to 95 percent—to achieve an overall recall of 70 percent to 80 percent for the end-to-end search and review process, when correctly measured with respect to a blind independent assessment.⁴²

³⁸ David C. Blair, *STAIRS Redux: Thoughts on the STAIRS Evaluation, Ten Years After*, 47 J. AM. SOC. INFO. SCI. 4 (1996).

³⁹ *The Implications of Rule 26(g)*, *supra* note 32; Maura R. Grossman & Gordon V. Cormack, *Vetting and Validation of AI-Enabled Tools for Electronic Discovery*, LITIGATING ARTIFICIAL INTELLIGENCE 465 (2021) [hereinafter *Vetting & Validation*].

⁴⁰ *TAR in E-Discovery*, *supra* note 27.

⁴¹ *The Implications of Rule 26(g)*, *supra* note 32.

⁴² *Vetting and Validation*, *supra* note 39.

IV. Non-Blind Assessment

Several of the most egregious applications of pseudo-science involve the use of non-blind experiments. Clever Hans was a horse purported to be able to perform intellectual tasks including arithmetic.⁴³ In fact, Hans' handler was either deliberately or inadvertently telegraphing the correct answer to Hans through non-verbal cues.⁴⁴ A panel of thirteen experts debunked the handler's claims by conducting a blinded experiment in which Hans was isolated from the questioner and spectators—Hans literally wore blinders—using questions for which the questioner did and did not know the answers in advance.⁴⁵ The experiments concluded that Clever Hans could intuit the correct answer from the questioner, even if the questioner was not Hans' handler, but only if the questioner knew the answer in advance.⁴⁶ So Clever Hans was actually less clever than initially thought.

Notwithstanding this debunking, the Clever Hans show continued to tour Germany, attracting large and enthusiastic crowds.⁴⁷ Nowadays, recall is calculated from the non-blind relevance assessments of armies of Clever Hanses, to the delight of large and enthusiastic crowds at eDiscovery medicine shows.

V. Misleading Statistics

Statistics play several roles in information retrieval. First and foremost, statistics are used to calculate the efficacy and reliability of information retrieval methods over many information needs and/or document collections.⁴⁸ Second, summary measures like recall and precision can be estimated for a particular information retrieval effort from relevance assessments for only a statistical sample of the collection.⁴⁹ Third, many information retrieval methods rely on statistical algorithms to score and rank documents by their likelihood

⁴³ OSKAR PFUNGST, *CLEVER HANS (THE HORSE OF MR. VON OSTEN)*, at v (Carl H. Rahn trans., Henry Holt & Co. ed. 1911) (2010) (ebook).

⁴⁴ *Id.* at 1.

⁴⁵ *Id.* at 43.

⁴⁶ *Id.* at 35.

⁴⁷ *Id.* at 229.

⁴⁸ STEFAN BÜTTCHER ET AL., *INFORMATION RETRIEVAL: IMPLEMENTING AND EVALUATING SEARCH ENGINES* 412-40 (2016).

⁴⁹ Gordon V. Cormack & Maura R. Grossman, *Beyond Pooling*, 41 INT'L ACM SIGIR CONF. ON RES. & DEV. INFO. RETRIEVAL 1169, 1169-72 (2018).

of relevance.⁵⁰ While only the second use has any operational impact on eDiscovery, it is frequently conflated with the other two, and all are conflated with technology-assisted review.⁵¹

Measuring recall and precision—whether to establish the efficacy and reliability of a method or to assess the quality of a particular review—has little to do with the technology employed for the review, be it Boolean keyword culling and/or TAR and/or manual review.⁵² Regardless of what technologies are employed, recall is the fraction of relevant documents *in the entire collection* that are retrieved by the overall effort.⁵³ On the other hand, the mechanics of statistical information retrieval algorithms are of no more concern to the eDiscovery practitioner than the biochemistry of a pharmaceutical is to the patient.

Part of the reason that these uses are conflated is that TAR was the first eDiscovery technique whose efficacy was evaluated and compared to that of manual review using the tools of information retrieval.⁵⁴ TAR practitioners exposed to these methods naturally assumed that the use of statistics was unique to TAR, and that statistical expertise was needed to employ TAR.⁵⁵ Purveyors and pundits have harnessed this misconception to suggest that only the TAR tool should be subject to validation, while keyword culling and manual review should be exempt, as they have always been.

At the same time, purveyors and pundits have grossly exaggerated the accuracy with which recall and precision (even of the TAR tool alone)

⁵⁰ Amit Singhal, *Modern Information Retrieval: A Brief Overview*, 24 IEEE DATA ENGINEERING BULL. 35, 36 (2001).

⁵¹ *The Implications of Rule 26(g)*, *supra* note 32 at 287.

⁵² *Vetting and Validation*, *supra* note 39.

⁵³ MANNING ET AL., *supra* note 12 at 154-57.

⁵⁴ Douglas W. Oard et al., *Evaluation of Information Retrieval for E-Discovery*, 18 A.I. & L. 347, 375 (2010).

⁵⁵ See, e.g., Gabriela P. Baron, *Perspective: Crucial Steps When Using Analytics During E-Discovery*, BLOOMBERG L. (Mar. 25, 2016, 11:02 AM), <https://news.bloomberglaw.com/business-and-practice/perspective-crucial-steps-when-using-analytics-during-e-discovery> [<https://perma.cc/B2SS-A7P5>].

can be estimated using sampling.⁵⁶ One tell that is often seen in vendor and pundit publications is the use of the non-sequitur “statistically significant sample.”⁵⁷ A more subtle mangling of statistical concepts is the notion of a sample having a given margin of error at a given confidence level; the most common examples are “margin of error of ± 5 percent at a 95 percent confidence level,” and “margin of error of ± 2 percent at a 95 percent confidence level,” when a more direct statement of what is really meant is “a random sample of size 385,” and “a random sample of size 2400,” respectively.⁵⁸ Specifying sample size obliquely by the margin of error that can purportedly be derived from it serves to obfuscate the sampling process and to intimidate practitioners, while strongly implying that the estimates derived from such samples are more precise than they are, and that they are sufficient to yield estimates of *recall* with the stated margins of error.⁵⁹ They are not.⁶⁰

At best, samples of size 385—or even 2400—can give a coarse estimate of recall with a margin of error several times larger than implied, which may identify a major failure, but cannot guarantee success.⁶¹ Furthermore, if the samples are assessed by Clever Hanses who know or can intuit how the documents were previously coded, we can place no bound on the error—and therefore the over- or underestimation—that can result.

⁵⁶ Claims such as “using a randomly drawn sample with 95 percent confidence and 2 percent margin of error from the unreviewed documents in the TAR review set to determine with a high degree of statistical confidence that the targeted recall rate was actually achieved” are contradicted by *The Implications of Rule 26(g)*, *supra* note 32, at 305-10, and Webber, *supra* note 19.

⁵⁷ Significance relates to the result of a statistical hypothesis test. No particular sample can be determined in advance that will yield a statistically significant result. The term “statistically significant sample” is thus an oxymoron, and misleading to the extent that it implies that any hypothesis test derived from such a sample will necessarily yield a statistically significant result. Nonetheless, vendors, pundits, and even TAR protocols and court filings commonly use phrases like “review a statistically significant sample.”

⁵⁸ For example, a recent court filing stated that “[t]ypical test sets for a document review may be: *i.* 400 documents for a 95 percent [confidence level], ± 5.0 percent [margin of error]; *ii.* 1535 documents for a 95 percent [confidence level], ± 2.5 percent [margin of error]; *iii.* 2400 documents for a 95 percent [confidence level], ± 2.0 percent [margin of error].”

⁵⁹ *The Implications of Rule 26(g)*, *supra* note 32, at 305-10.

⁶⁰ *Id.*

⁶¹ *Id.*

VI. The Fallacies of Incredulity, Special Pleading, and Assuming the Converse

Pundits' and practitioners' incredulity that large numbers of relevant documents are excluded by keyword culling and by reviewer assessment errors does not refute the extensive body of scientific evidence showing that they are. Nor does a special pleading argument that their own unique assessors or quality assurance process or bloodletting instruments would have achieved a different result. Nor does pointing out limitations—whether perceived or real—in scientific studies constitute evidence for the converse of the studies' conclusions. There is a strong body of scientific evidence showing that keyword culling removes a substantial number of relevant documents, that certain TAR methods achieve as good or better recall and precision than human assessors, and that non-blind assessments provide unreliable results.⁶² If any of the issues raised in arguments to the contrary had merit, they could surely be demonstrated through scientific studies whose results were subject to peer review. That has not happened.

VII. The Ball is in the Courts' Court

Just as bloodletting came to be considered unreasonable in the face of mounting scientific evidence, so too should certain common eDiscovery practices. Just as the snake oils of the early 20th century were unreasonable from the outset, so too are various TAR tools and methods whose efficacy has not been established.

What is needed is a full-scale evaluation, and judicial recognition of which methods are reasonable and which are not. Unfortunately, an evidentiary hearing sufficient to establish reasonableness would likely entail burden and cost disproportionate to the needs of any particular case, as evidenced by the *Kleen Products* case,⁶³ in which then-Magistrate (now retired) Judge Nan R. Nolan, after two full days of evidentiary hearings and 11 status hearings and Rule 16 conferences⁶⁴

⁶² *Vetting and Validation*, *supra* note 39.

⁶³ *Kleen Prods. LLC v. Packaging Corp. of Am.*, Civ. Ac. No. 10-C-5711, 2012 WL 4498465, at *4 (N.D. Ill. 2012).

⁶⁴ *See* FED. R. CIV. P. 16(a).

with the parties, essentially coaxed them to work it out for themselves.⁶⁵

In the near-decade since *Kleen Products*, the scientific body of evidence regarding how to determine the effectiveness and reliability of TAR methods has grown substantially. It is time to revisit what is reasonable in light of this evidence.

Federal Rule of Civil Procedure 26(g)(1) requires that the producing party or their attorney personally certify that “after reasonable inquiry,” to the best of their knowledge, information, and belief, the production is not unreasonable, considering the burden or expense, the needs of the case, prior discovery in the case, the amount in controversy, and the importance of the issues at stake in the action.⁶⁶

Notwithstanding Sedona Principle 6’s prescription that “[r]esponding parties are best situated to evaluate the procedures, methodologies, and technologies appropriate for preserving and producing their own electronically stored information,” Federal Rule of Civil Procedure 26(g)(1)—which is the law—mandates that this evaluation must be consistent with a “reasonable inquiry,” and that the chosen procedures, methodologies, and technologies must be calculated to yield a reasonable production.⁶⁷

The courts need not presume that any method used in the past or any method chosen by the producing party is reasonable, especially if presented with clear and convincing evidence that it is unreliable or ineffective, compared to readily available alternatives. In the face of such evidence, the court need not wait until the method’s unreliability or ineffectiveness is manifested in the matter before it, which may occur only at the eleventh hour when proper remediation is no longer practical. At the same time, courts need not impose on producing parties the prescriptions of requesting parties unless they are similarly supported by clear and convincing evidence.

⁶⁵ *Kleen*, 2012 WL 4498465, at *4.

⁶⁶ FED. R. CIV. P. 26(g)(1).

⁶⁷ *Id.*

In assessing whether an inquiry, an eDiscovery method, or a production is reasonable, the courts may consider the *Daubert* factors:

Many considerations will bear on the inquiry, including whether the theory or technique in question *can be (and has been) tested*, whether it *has been subjected to peer review and publication*, its *known or potential error rate*, and the *existence and maintenance of standards controlling its operation*, and whether it has *attracted widespread acceptance within a relevant scientific community*. (emphasis added).⁶⁸

To be clear, we are not advocating for a full-blown evidentiary hearing every time a TAR or validation process is challenged. Rather, we are suggesting that the *Daubert* factors offer useful guidance in determining what a reasonable process is pursuant to Federal Rule of Civil Procedure 26(g)(1), and what is proper evidence to this effect.

The scientific literature contains a growing body of peer-reviewed and widely accepted evidence of the effectiveness and error rates of the information retrieval methods used in eDiscovery, including the error rates for Boolean search, manual review, specific TAR tools and protocols, and specific validation methods. eDiscovery and/or validation efforts can be considered tested only when the tools, methods, and standards of operation they employ are comparable to those of the research studies from which supporting evidence is derived.

⁶⁸ See Syllabus (c) to *Daubert v. Merrell Dow Pharm., Inc.*, 509 U.S. 579, 580 (1993). The *Daubert* factors are incorporated in Federal Rule of Evidence 702 governing expert testimony. While it is unlikely that Federal Rule of Evidence 702 per se governs eDiscovery representations, there can be no question that the effectiveness and reasonableness of eDiscovery methods is a matter of scientific inquiry. Indeed, in *Victor Stanley, Inc. v. Creative Pipe, Inc.*, 250 F.R.D. 251, 260 n.10 (D. Md. 2008), then-Magistrate (now District) Judge Grimm noted that “[i]t cannot credibly be denied that resolving contested issues of whether a particular search and information retrieval method was appropriate—in the context of a motion to compel or motion for protective order—involves scientific, technical or specialized information. If so, then the trial judge must decide a method’s appropriateness with the benefit of information from some reliable source—whether an affidavit from a qualified expert, a learned treatise, or, if appropriate, from information judicially noticed. To suggest otherwise is to condemn the trial court to making difficult decisions on inadequate information, which cannot be an outcome that anyone would advocate.” We agree.

Merely labeling an eDiscovery method as “TAR” does not establish that it has been tested. Nor does using a TAR tool that has been tested establish that it has been used according to a tested protocol. Nor does naming a particular protocol such as “Continuous Active Learning” or “CAL” establish that the protocol has been followed according to established standards of operation. Nor does proclaiming “75 percent recall” establish that recall has been calculated according to scientifically established methods with a known error rate, or that such a computation is evidence of either a reasonable inquiry or a reasonable production.

In contrast, producing parties should show—and the courts should demand that they show—the reasonableness of their eDiscovery search and review processes, as well as the resulting production, by hewing closely to tools, methods, and procedures that have been scientifically vetted and shown to be valid and reliable. Anything else belongs in a medicine show.