

:: The Metrics of E-Discovery

Proving What Your Document Review Is Worth

by Nicolas Economou of H5

Large-scale document review, by its very nature, presents a pressing question that is rarely asked and even more rarely answered: “Of all the documents we should have retained, produced or found, how many did we *actually* retain, produce or find?”

As disclosure requirements increase and litigation stakes grow higher, so does the risk of missing key documents. We all know from recent high-profile discovery abuse verdicts that the financial and legal risks of discovery mistakes can be tremendous, and the resulting harm can be irreversible.

Absent any reliable measurement of the performance of their document review systems, lawyers cannot manage such risks without resorting to some amount of guesswork. Without quality assurance protocols with which to appraise the effectiveness of review, lawyers and their review teams do not — and cannot — know whether or to what extent their document review results fall short of their expectations and of their legal obligations.

In order to be effective, any e-discovery process needs to be both self-evaluating and verifiable using widely accepted and open academic protocols. It is as true here as in any other critical business process: if you can't measure it, you can't improve it.

Fundamental Challenges to Document Review

Even the best lawyers with the best document review teams would be shocked to realize how many potentially relevant documents they may be missing. Several studies on the effectiveness of existing document review systems present an overwhelming consensus that in a typical large-scale review, the majority of relevant documents are missed.¹

Why are the results often so poor, even in top-priority cases where litigators can presumably devote all the resources they see fit to document review? To better understand this, let's consider three fundamental challenges underlying all large-scale document review processes:

Deficient Knowledge Transfer (a teaching challenge). Once a document review team grows beyond a few people, it is impossible to ensure that all the reviewers will understand *without any variance* what the lead litigator is looking for. Imperfect knowledge transfer becomes an even bigger problem when dealing with complex or changing case issues.

Lack of Consistency (a human nature challenge). Even with perfect knowledge transfer, a consistently correct application of knowledge is simply impossible. Humans are prone to fatigue, loss of focus, distraction, snap judgments and errors of reasoning or understanding. Not only will two different people sometimes evaluate the same document differently, the same person may do so at different times. Workforce turnover, variable levels of skill or dedication and tight deadlines only magnify the challenge.

Lack of Quality Assurance (a measurement challenge). As noted above, the need for externally *verifiable standards with which to measure the accuracy of your review results and benchmark the effectiveness of your review system is critical.* Without rigorous quality

assurance protocols that evaluate the accuracy of review, the first two problems cannot be appraised, much less addressed.

Unless each of these three challenges is addressed, even the most well-designed and well-staffed large-scale document review process remains rife with risk.

Addressing the Challenges — or Not?

Often, legal departments and law firms look to keyword and “concept” search tools as a solution. In response to this trend, the number of e-discovery vendors continues to grow — all promising a better, faster and cheaper way of doing document review, while further obscuring the challenges that lawyers must address. The comfort that search tools offer is largely illusory for two reasons:

First, even expert users reach only moderate levels of accuracy with current search tools. This is mainly because keywords and “concepts” are generally poor indicators of relevance. Numerous academic studies show that even highly experienced searchers, using the most advanced search tools on a relatively small collection, miss 50 percent or more of the documents that a human, upon review of the results, would have assessed to be relevant.²

Second, search software doesn't address the fundamental challenges outlined above. It doesn't tell reviewers whether or not they clearly understand what the lead litigator is looking for, it doesn't notify them when they are inconsistent in their document assessments and it doesn't apprise attorneys of the actual performance of the overall review. Therefore, heavy reliance on search software, in and of itself, may cause document reviewers to go faster in the wrong direction. This can put attorneys and their clients in the precarious position of believing that they have found more than they actually have, or that they know more than they actually do.

To be sure, search technology has its place. It would be inconceivable in large-scale document review or any other information retrieval or categorization endeavor involving more than a few thousand pages, to forgo the benefit that search tools offer. It is, however, equally important to understand the documented performance limitations of these technologies and the context and constraints within which they operate, which are often nontechnological and must be considered independently of the search technology used.

Search technology is part of the solution, but in itself is incomplete; on its own it simply cannot produce accurate results in document review.

Emerging Standards

With a confusing array of “solutions” being offered to address lawyers' large-scale document review woes, there is a growing recognition that the legal industry needs common standards for discovering and reviewing electronic documents. Two of the major initiatives to create a set of best practices for e-discovery are the Sedona Conference and the Electronic Discovery Reference Model (EDRM) Project.

The Sedona Principles, first published in late 2002, are among the earliest collaborative attempts at developing widely acceptable

guidelines for electronic document retention and production.³ These principles were developed by a group of lawyers, academics and e-discovery experts and vendors, under the aegis of the Sedona Conference, a nonprofit think-tank for advanced legal issues.

In May 2005, e-discovery consultants George Socha and Tom Gelbmann launched the EDRM Project to address a major concern that was identified in their preceding surveys of e-discovery vendors and consumers: the lack of a common, flexible and extensible framework for the development, selection, evaluation and use of electronic discovery products and services.⁴ Socha and Gelbmann are currently incorporating comments from all affected parties. The completed model will be placed in the public domain in May 2006.

These efforts are laudable and represent a significant first step toward creating effective document review systems that are embraced by the legal industry.

Yet no matter how well-developed, clearly-defined or well-implemented these proposed standards are, the need for validation of accuracy remains. These standards, in their current scope and intent, do not offer a mechanism that allows attorneys and review teams to answer the lingering question: “Of all the documents we should have retained, produced or found, how many did we *actually* retain, produce or find?” Legal practitioners still need to evaluate the effectiveness of their own review methods and protocols, after having incorporated such standards.

Anatomy of an Alternative Approach

So what sort of approach might get us closer to an “ideal” document review system?

First, it would address the three challenges discussed above: knowledge transfer, consistency and quality assurance (measurement of accuracy).

Second, it would recognize that evaluating documents about complex subjects requires the use of subject matter experts in those fields as well as process design experts.

Third, it would recognize the proper functions and uses of search technologies in the overall document review process.

Fourth, in order to be practical, it would not disrupt either the attorney's work or the firm's existing IT infrastructure.

Hence, the right approach to designing a markedly improved document review process would be conceptually “systemic” with respect to acknowledging and addressing the obstacles identified above, and it would be practicable. Because it would address the key underlying challenges of document review, this type of approach should achieve higher performance.⁵

The Secret to Success: A Self-Evaluating Approach to E-Discovery

There is more than one right way to conduct document review, but broadly speaking, an effective document review system must incorporate each of the following three components:

A detailed definition of the target with which to make relevancy determinations

Implementation of a review process (fully manual, fully automated, computer/search-assisted or some combination of these)

Measurement of performance using evaluation protocols

As discussed earlier, legal practitioners have placed much (and perhaps too much) emphasis on the review process (second item, above), and in particular, on search tools. In doing so, they often overlook the other two components, which are equally important, if not more so. Perhaps this is because the importance of the first and third items are so self-evident that they are presumed to be part of the review process. In fact, if you don't determine what you're looking for, you can't measure what you've found. And if you can't measure your results, you can't improve them.

Based on a proper definition of the target, the accuracy of the review, irrespective of whether it is conducted manually or with the support of a variety of technologies, can be assessed using academic protocols modeled on TREC (Text Retrieval Conference).

The need to utilize performance evaluation protocols is arguably the most important component for large-scale, complex document review. This is the secret ingredient that is missing in the overwhelming majority of document review processes. Reliable protocols can be used to benchmark not only accuracy, but also all-in cost, time-to-completion and exposure to risk. They are widely available, relatively easy to implement and can be applied at reasonable cost to representative samples of a document collection.

Get an Expert

Law firms and law departments do not need to be staffed with experts in search technology in order to choose the proper document review solution. But they would benefit by working with experts who can develop and implement protocols and processes by which they can measure — and improve — their performance and thereby more effectively manage their risk.

Endnotes

- 1 In an important study, attorneys leading a search of about 40,000 documents via Boolean searches believed they had found more than 75 percent of all relevant documents. In fact, they found only about 20 percent when their searches were checked against a manual review of the entire pool. (Blair & Maron, 1985; Dabney, 1986). See also *infra*, n. 2.
- 2 At acceptable levels of precision (where at least 7 out of 10 documents returned by a search are relevant), typically less than 50 percent of all relevant documents can be expected to be retrieved, as reported by the Text Retrieval Conference (TREC), an academic/industry initiative co-sponsored by the National Institute of Standards and Technology (NIST) and the U.S. Department of Defense. See <http://trec.nist.gov>.
A 1994 study conducted by Westlaw comparing manually constructed Boolean searches to automated document retrieval systems based on artificial intelligence found that neither method worked very well. The Boolean approach found only 24.4 percent of the relevant documents. Meanwhile, the automated system found 32.9 percent of all relevant documents, but each sample also turned up an average of 43 percent irrelevant documents. (*Turtle*, 1994)
- 3 See www.thesedonaconference.org.
- 4 See www.edrm.net.
- 5 In a controlled study, litigation management consultant Anne Kershaw evaluated such an offering. The “automated document review” was a combination of professional services, process design expertise and advanced artificial intelligence technology, intended to meet the criteria of being both conceptually systemic and practicable. The study demonstrated that the automated document review process reduced the risk of missing relevant documents by 90 percent or more vis-à-vis full manual review. See www.reasonablediscovery.com.