


MUSINGS ON TECHNOLOGY ASSISTED REVIEW

CRAIG BALL



Whether you call it “Predictive Coding” or Technology Assisted Review,” the time is nigh to leave much of the heavy lifting of review to machines trained to find responsive documents. These tools won’t be heuristic marvels like HAL-9000; but on the plus side, they probably won’t try to kill us.

The 'Not Me' Factor

Craig Ball © 2013



I've been skeptical of predictive coding for years, even before I wrote my first column on it back in 2005. Like most, I was reluctant to accept that a lifeless mass of chips and wires could replicate the deep insight, the nuanced understanding, the sheer freaking brilliance that my massive lawyer brain brings to discovery. Wasn't I the guy who could pull down that one dusty box in a cavernous records repository and find the smoking gun everyone else overlooked? Wasn't it my rarefied ability to discern the meaning lurking beneath the bare words that helped win all those verdicts?

Well, no, not really. But, I still didn't trust software to make the sort of fine distinctions I thought assessing relevance required.

So, as others leapt aboard the predictive coding bandwagon, I hung back, uncertain. I felt not enough objective study had been done to demonstrate the reliability and superiority of predictive coding. I well knew the deep flaws of mechanized search, and worried that predictive coding would be just another search tool tarted up in the frills and finery of statistics and math. So, as Herb and Ralph, Maura and Gordon and Karl and Tom sung Hosannas to TAR and CAR from Brooklyn Heights to Zanzibar, I was measured in my enthusiasm. With so many smart folks in thrall, there had to be something to it, right? Yet, I couldn't fathom how the machine could be better at the fine points of judging responsiveness than I am.

Then, I figured it out: The machine's not better at fine judgment. I'm better at it, and so are you.

So why, then, have I now drunk the predictive coding Kool-Aid and find myself telling anyone who will listen that predictive coding is the Way and the Light?

It's because I finally grasped that, although predictive coding isn't better at dealing with the swath of documents that demand careful judgment, it's every bit as good (and actually much, much better) at dealing with the overwhelming majority of documents that don't require careful judgment—the very ones where keyword search and human reviewers fail miserably.

Let me explain.

For the most part, it's not hard to characterize documents in a collection as responsive or not responsive. The vast majority of documents in review are either pretty obviously responsive or pretty obviously not. Smoking guns and hot docs are responsive because their relevance jumps out at you. Most irrelevant documents get coded quickly because one can tell at a glance that they're irrelevant. There are close calls, but overall, not a lot of them.

If you don't accept that proposition, you might as well not read further; but if you don't, I question whether you've done much document review.

It turns out that well-designed and -trained software also has little difficulty distinguishing the obviously relevant from the obviously irrelevant. And, again, there are many, many more of these clear cut cases in a collection than ones requiring judgment calls.

So, for the vast majority of documents in a collection, the machines are every bit as capable as human reviewers. A tie. But giving the extra point to humans as better at the judgment call documents, HUMANS WIN! Yeah! GO HUMANS! Except....

Except, the machines work much faster and much cheaper than humans, and it turns out that there really is something humans do much, much better than machines: they screw up.

The biggest problem with human reviewers isn't that they can't tell the difference between relevant and irrelevant documents; it's that they often don't. Human reviewers make inexplicable choices and transient, unwarranted assumptions. Their minds wander. Brains go on autopilot. They lose their place. They check the wrong box. There are many ways for human reviewers to err and just one way to perform correctly.

The incidence of error and inconsistent assessments among human reviewers is mind boggling. It's unbelievable. And therein lays the problem: it's unbelievable. People I talk to about reviewer error might accept that some nameless, faceless contract reviewer blows the call with regularity, but they can't accept that potential in themselves. "Not me," they think, "If I were doing the review, I'd be as good as or better than the machines." It's the "Not Me" Factor.

Indeed, there is some cause to believe that the best trained reviewers on the best managed review teams get very close to the performance of technology-assisted

review. A chess grand master has been known to beat a supercomputer (though not in quite some time).

But so what? Even if you are that good, you can only achieve the same result by reviewing all of the documents in the collection, instead of the 2%-5% of the collection needed to be reviewed using predictive coding. Thus, even the most inept, ill-managed reviewers cost more than predictive coding; and the best trained and best managed reviewers cost much more than predictive coding. If human review isn't better (and it appears to generally be far worse) and predictive coding costs much less and takes less time, where's the rational argument for human review?

What's that? "My client wants to wear a belt AND suspenders?" Oh, PLEASE.

What about that chestnut that human judgment is superior on the close calls? That doesn't wash either. First—and being brutally honest—quality is a peripheral consideration in e-discovery. I haven't met the producing party who loses sleep worrying about whether their production will meet their opponent's needs. Quality is a means to avoid sanctions, and nothing more.


Moreover, predictive coding doesn't try to replace human judgment when it comes to the close calls. Good machine learning systems keep learning. When they run into one of those close call documents, they seek guidance from human reviewers. It's the best of both worlds.

So why isn't everyone using predictive coding? One reason is that the pricing has not yet shifted from exploitive to rational. It shouldn't cost substantially more to expose a collection to a predictive coding tool than to expose it to a keyword search tool; yet, it does. That will change and the artificial economic barriers to realizing the benefits of predictive coding will soon play only a minor role in the decision to use the technology.

Another reason predictive coding hasn't gotten much traction is that Not Me Factor. To that I say this: Believe what you will about your superior performance, tenacity and attention span (or that of your team or law firm), but remember that you're spending someone else's money on your fantasy. When the judge, the other side or (shudder) the client comes to grips with the exceedingly poor value proposition that is large-scale human review, things are going to change...and, Lucy, there's gonna be some 'splainin to do!

What are We Waiting For?

Craig Ball © 2012



Winston Churchill said that, “Democracy is the worst form of government except all those other forms that have been tried from time to time.” That famous quip neatly describes keyword search in e-discovery. It stinks, yet lawyers turn to keyword search again and again, because it seems like the best option out there. It’s the devil we know.

Though keywords serve us well when searching the web, they perform poorly finding “all documents touching, concerning or relating to” an issue in litigation. The failure is particularly pronounced when keyword search is

pursued in the usual fashion of opponents horse trading terms without testing them against sample data or adapting the list to ameliorate well-known flaws like misspellings, noise words and synonyms.

But that’s old news. Students of e-discovery know that keyword search is the worst form of search, and harbor no illusions that it’s better than the others that have been tried from time to time. Whether you call it advanced data analytics, predictive coding, concept search or whatever else leaps from the lips of marketing mavens, there exist techniques that, when implemented with care and judgment, do a better, less costly job than keyword search and linear review.

Yet whenever these techniques come up in conversations or articles, lawyers seem like kids inching toward the cookie jar, intently watching Mom’s face to see if it’s okay to snag some Mallomars. It may be better and cheaper, but nobody wants to give enhanced automated search much of a go until “it’s okay with Mom.”

What are we waiting for?

The answer seems to be some sort of authoritative court blessing of alternatives to keyword search. We’ve seen favorable mention of such techniques in footnotes to

decisions from the most influential judges writing on e-discovery issues, but nothing opining that use of enhanced search is “court approved.”

Again, what are we waiting for?

It’s not as though we held off using keyword search until a judge gave it the nod. We just did it. And, though keyword search never really got a judicial stamp of approval, neither was it summarily rejected. Again, we just did it, and in time it emerged as a standard.

Perhaps there will one day be a decision where a judge expressly cites enhanced search techniques as reliable proxies for human review or preferred alternatives to keyword search. I wouldn’t hold my breath waiting for it. The American justice system doesn’t favor advisory opinions. Courts expect genuine cases and controversies to drive our jurisprudence. New search techniques need to be used before they can be meaningfully addressed in reported decisions.

So, quit worrying about Mom and grab those Mallomars! If you believe enhanced automated search is better and cheaper, have the courage and wisdom to lead the way in its use.

Imagining the Evidence

Craig Ball © 2012



As a young lawyer in Houston, I had the good fortune to sip whiskey with veteran trial attorneys who never ran short of stories. One told of the country lawyer who journeyed to the big city to argue before the court of appeals. The case was going well until a judge asked, "Counsel, are you aware of the maxim, '*volenti non fit injuria*'?"

"Why, Your Honor," he answered in a voice as smooth as melted butter, "In the piney woods of East Texas, we speak of little else."

Lately, in the piney woods of e-discovery, the topic is technology-assisted review (TAR *aka* predictive coding), and we speak of little else. The talk centers on that sudsy soap opera, *Da Silva Moore v. Publicis Groupe*, and whether Magistrate Judge Andrew Peck of the Southern District of New York will be the first judge to anoint TAR as being "court approved" and a suitable replacement for manual processes now employed to segregate ESI.

TAR is the use of computers to identify responsive or privileged documents by sophisticated comparison of a host of features shared by the documents. It's characterized by methods whereby the computer trains itself to segregate responsive material through examination of the data under scrutiny or is trained using exemplar documents ("seed sets") and/or by interrogating knowledgeable human reviewers as to the responsiveness or non-responsiveness of items sampled from the document population.

Let's put this "court approved" notion in perspective. Dunking witches was court approved and doubtlessly engendered significant cost savings. Trial by fire was also court approved and supported by precise metrics ("*M'Lord, guilt is established in that the accused walked nine feet over red-hot ploughshares and his incinerated soles festered within three days*"). Whether a court smiles on a methodology may not be the best way to conclude it's the better mousetrap. Keyword search and linear review enjoy *de facto* court approval; yet both are deeply flawed and brutally inefficient.

The imprimatur that matters most is "opponent approved." Motion practice and false starts are expensive. The most cost-effective method is one the other side accepts without a fight, i.e., the least expensive method that affords opponents superior confidence that responsive and non-privileged material will be identified and produced.

Don't confuse that with an obligation to kowtow to the opposition simply to avoid conflict. The scenario I'm describing is a true win-win:

- Producing parties have an incentive to embrace TAR because, when it works, TAR attenuates the most **expensive** component of e-discovery: *attorney search and review*.
- Requesting parties have an incentive to embrace TAR because, when it works, TAR attenuates the most **obstructive** component of e-discovery: *attorney search and review*.

Producing parties don't just obstruct discovery by the rare and reprehensible act of intentionally suppressing probative evidence. It occurs more often with a pure heart and empty head as a consequence of lawyers using approaches to search and review that miss more responsive material than they find.

It's something of a miracle that documentary discovery works at all. Discovery charges those who reject the theory and merits of a claim to identify supporting evidence. More, it assigns responsibility to find and turn over damaging information to those damaged, trusting they won't rationalize that incriminating material must have had some benign, non-responsive character and so need not be produced. Discovery, in short, is anathema to human nature.

A well-trained machine doesn't care who wins, and its "mind" doesn't wander, worrying about whether it's on track for partnership. From the standpoint of a requesting party, an alternative that is both objective and more effective in identifying relevant documents is a great leap forward in fostering the integrity and efficacy of e-discovery. Crucially, a requesting party is more likely to accept the genuine absence of supportive ESI if the requesting party had a meaningful hand in training the machine.

Until now, the requesting party's role in "training" an opponent's machines has been limited to proffering keywords or Boolean queries. The results have been uniformly awful.

But the emerging ability to train machines to "find more documents like this one" will revolutionize requests for production in e-discovery. Because we can train the tools to find similar ESI using *any* documents, we won't be relegated to using seed sets derived from *actual* documents. We can train the tools with contrived examples—fabrications of documents like the genuine counterparts we hope to find.

I call this "imagining the evidence," and it's not nearly as crazy as it sounds.

If courts permit the submission of keywords to locate documents, why not entire documents to more precisely and efficiently locate other documents? Instead of

demanding “any and all documents touching or concerning” some amorphous litany of topics, we will serve a sheaf of dreams—freely forged smoking guns—and direct, “show me more like these.”

Predictive coding is not as linguistically fussy as keyword search. If an opponent submits contrived examples of the sorts of documents they seek, it’s far more likely a similar document will surface than if keywords alone were used. As importantly, it’s less likely that a responsive document will be lost in a blizzard of false hits. This allows us to rely less on our opponents to artfully construct queries. Instead, we need only trust them to produce the non-privileged, responsive results the machine finds.

There’s more to documents than just the words they contain, so mocking up contrived exemplars entails more than fashioning a well-turned phrase. Effective exemplars will employ contrived letterheads and realistic structure, dates and distribution lists to insure that all useful contextual indicia are present. And, of course, care must be taken and processes employed to ensure that no contrived exemplars are mistaken for genuine evidence.

The use of contrived examples may ruffle some feathers. I can almost hear a chorus of, “How dare they draft such a vile thing!” But the methodology is sound, and how we will go about “imagining the evidence” is likely to be a topic of discussion in the negotiation of search protocols once use of technology assisted review is commonplace.

Another “not as nutty as it sounds” change in discovery practice wrought by TAR will be affording requesting parties a role in training TAR systems. The requesting party’s counsel would be presented with candidate documents from the collection that the machine has identified as potentially responsive. The requester will then decide whether the sample is or is not responsive, helping the machine hone its capacity to find what the requester seeks. After all, the party seeking the evidence is better situated to teach the machine how to discriminate.

For this to work, the samples must first be vetted by the responding party’s counsel for privilege and privacy concerns, and the requesting party must be willing to undertake the effort without fretting about revealing privileged mental impressions. It’s going to take some getting used to; but the reward will be productions that cost less and that requesting parties trust more.

Volenti non fit injuria means “to a willing person, injury is not done.” When we fail to embrace demonstrably better ways of searching and reviewing ESI, we assume the risk that probative evidence won’t see the light of day and voluntarily pay too high a price for e-discovery.

Train, Don't Cull, Using Keywords

Craig Ball © 2012

I've been thinking about how we implement technology-assisted review tools and particularly how to hang onto the on-again/off-again benefits of keyword search while steering clear of its ugliness. The rusty flivver that is my brain got a kick start from many insightful comments made at the recent Carmel Valley E-discovery Retreat in Monterey, California. As is often the case when the subject is technology-assisted review (by whatever name you prefer, dear reader: predictive coding, CAR, automated document classification, Francis), some of those kicks came from lawyer [Maura Grossman](#) and computer scientist [Gordon Cormack](#). So, if you like where I go with this, credit them. If not, blame me for misunderstanding.

Maura and Gordon are the power couple of predictive coding, thanks to their thoughtful [papers](#) and presentations transmogrifying the metrics of NIST TReC into coherent observations concerning the efficacy of automated document classification. While they're spinning straw into gold. I'm still studying it all; but from where I stand, they make a lot of sense.

Maura expressed the view that technology-assisted review tools shouldn't be run against subset collections culled by keywords but should be turned to the larger collection of ESI (i.e., the collection/sources against which keyword search might ordinarily have been deployed). The gist was, 'use the tools against as much information as possible, and don't hamstring the effort by putting old tools out in front of new ones.' [I'm not quoting here, but relating what I gleaned from the comment].

At the same Monterey conference, [Judge Andrew Peck](#) reminded us of the perils of GIGO (Garbage In:Garbage Out) when computers are mismanaged. The devil is very much in the details of any search effort, but never more so than when one deploys predictive coding in e-discovery. *Methodology matters.*



If technology-assisted review were the automobile, we'd still be at the stage where drivers asked, "Where do I hook up my mules?" Our "mules" are keyword search.

When you position keyword search in front of predictive

coding; that is, when you use keyword search to create the collection that predictive coding “sees,” the view doesn’t change much from the old ways. You’re still looking at the ass end of a mule. Breath deep the funky fragrance of keyword search. Put axiomatically, no search technology can find a responsive document that’s not in the collection searched, and keyword search leaves most of the responsive documents out of the collection.

Keyword search can be very precise, but at the expense of recall. It can achieve splendid recall scores, but with abysmal precision. How, then, do we avail ourselves of the sometimes laser-like precision of keyword search without those awful recall in-laws coming to visit? Time-and-again, research proves that keyword search performs far less effectively than we hope or expect. It misses 30-80% of the truly responsive documents and sucks in scads of non-responsive junk, hiding what it finds in a blizzard of blather.

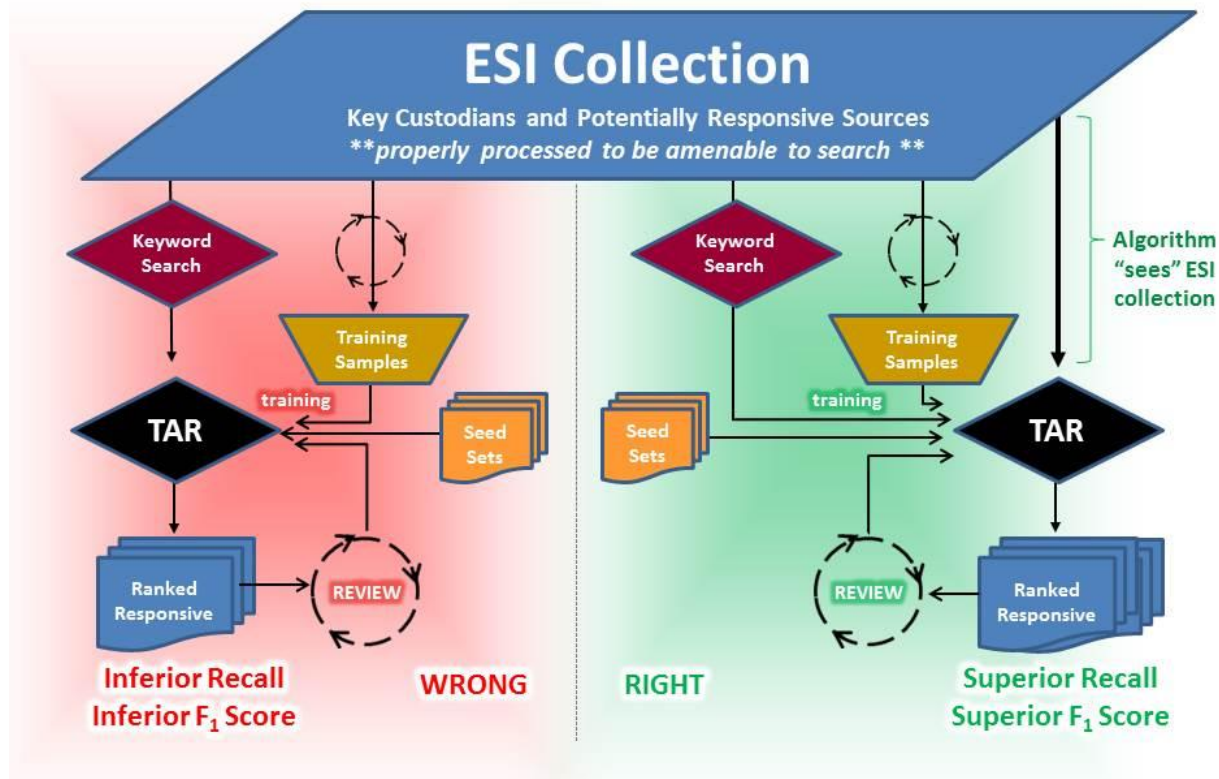
To be clear, that’s an established metric based on *everyone else in the world*. It doesn’t apply to *YOU*. *YOU* have the unique ability to frame fantastically precise and effective keyword searches like no one else. Likewise, all the findings about the laughably poor performance of human reviewers applies only to *other* reviewers, *not to YOU*. Tragically, not everyone has the immense good sense to employ *YOU*; so, let’s take *YOU* and what *YOU* can do out of the equation until human cloning is commonplace, okay?

For all their shortcomings, mules are handy. When your Model-T gets stuck in the mud, a mule team can pull you out. Likewise, keyword search is a useful tool to pull us out of the sampling swamp and generate training sets. Using keywords, you’re more likely to rapidly identify *some* responsive documents than using random sampling alone. These, in turn, increase the likelihood that predictive coding tools will find other responsive documents in the broader collection of ESI sources. *Good stuff in:good stuff out*.

With that in mind, I made the following diagram to depict how I think keyword search should be incorporated into TAR and how it shouldn’t. (George Socha is so much better at this sort of thing, so forgive my crude effort).

The Algorithm Must “See” the Collection

Train, don't cull, using keywords



I hope you'll agree that the interposition of keyword search to cull the collection before it's exposed to an automated document classification tool is wrong. But, in fairness, doing it the right way could come at a cost depending upon how you approach the assembly and processing of potentially responsive ESI. If you have to pay significantly more to let the tool “see” significantly more data, then quality will be sacrificed on the altar of savings. How it shakes out in your case hinges on how you handle keyword search and what you're charged for ingestion and hosting. Currently, many use keyword search via entirely separate tools and workflows to reduce the volume of information collected, processed and hosted. *Garbage In.*

Another caution I think important in using keywords to train automated classification tools is the requirement to elevate precision over recall in framing searches to insure that you don't end up training your predictive classification tool to replicate the shortcomings of keyword search. If only 20% of the documents returned by keyword search are responsive, then you don't want to train the tool to find more documents like the 80% that are junk. So when, in the illustration above, I depict keyword search as a

means to train technology-assisted review tools, please don't interpret the line leading from keyword search to TAR as suggesting that the usual guesswork approach to keyword search is contemplated and you'll just dump keyword results into the tool. That's like routing the exhaust pipe into the passenger compartment. The searches required need to be narrow—precise—*surgical*. They must jettison recall to secure precision...and may even benefit from a soupçon of human review.

For the promise of predictive coding to be fulfilled, workflows and pricing must better balance the quality vs. cost equation. Yes, a technology that is less costly when introduced at nearly any stage of the review process is great and arguably superior only by being no worse than alternatives. But if that is all we seek when quality is also within easy reach, we do a disservice to justice. The societal and psychic benefits of a more trusted and accurate outcome to disputes cannot be overvalued. "Perfect" is not the standard, but neither is "screw it."

The Streetlight Effect in Electronic Discovery

Craig Ball © 2012

In the wee hours, a beat cop sees a drunken lawyer crawling around under a streetlight searching for something. The cop asks, "What's this now?" The lawyer looks up and says, "I've lost my keys." They both search for a while, until the cop asks, "Are you sure you lost them here?" "No, I lost them in the park," the tipsy lawyer explains, "but the light's better over here."



I told that groaner in court, trying to explain why opposing counsel's insistence that we blindly supply keywords to be run against the e-mail archive of a Fortune 50 insurance company wasn't a reasonable or cost-effective approach e-discovery. The "Streetlight Effect," described by David H. Freedman in his 2010 book *Wrong*, is a species of observational bias where people tend to look for things in the easiest ways. It neatly describes how lawyers approach electronic discovery. We look for responsive ESI only where and how it's easiest, with little consideration of whether our approaches are calculated to find it.

Easy is wonderful when it works; but looking where it's easy *when failure is assured* is something no sober-minded counsel should accept and no sensible judge should allow.

Consider *The Myth of the Enterprise Search*. Counsel within and without companies and lawyers on both sides of the docket believe that companies have the ability to run keyword searches against their myriad siloes of data: mail systems, archives, local drives, network shares, portable devices, removable media and databases. They imagine that finding responsive ESI hinges on the ability to incant magic keywords like Harry Potter. *Documentum Relevantus!*

Though data repositories may share common networks, they rarely share common search capabilities or syntax. Repositories that offer keyword search may not support Boolean constructs (queries using "AND," "OR" and "NOT"), proximity searches (Word1 near Word2), stemming (finding "adjuster," "adjusting," "adjusted" and "adjustable") or fielded searches (restricted to just addressees, subjects, dates or message bodies). Searching databases entails specialized query languages or user privileges. Moreover, different tools extract text and index such extractions in quite different ways, with the upshot being that a document found on one system will not be found on another using the same query.

But the Streetlight Effect is nowhere more insidious than when litigants use keyword searches against archives, e-mail collections and other sources of indexed ESI,

That Fortune 50 company—call it All City Indemnity—collected a gargantuan volume of e-mail messages and attachments in a process called “message journaling.” Journaling copies every message traversing the system into an archive where the messages are indexed for search. Keyword searches only look at the index, not the messages or attachments; so, if you don’t find it in the index, you won’t find it at all.

All City gets sued every day. When a request for production arrives, they run keyword searches against their massive mail archive using a tool we’ll call *Truthiness*. Hundreds of big companies use *Truthiness* or software just like it, and blithely expect their systems will find all documents containing the keywords.

They’re wrong...or in denial.

If requesting parties don’t force opponents like All City to face facts, All City and its ilk will keep pretending their tools work better than they do, and requesting parties will keep getting incomplete productions. To force the epiphany, consider an interrogatory like this:

For each electronic system or index that will be searched to respond to discovery, please state:

- a. **The rules employed by the system to tokenize data so as to make it searchable;**
- b. **The stop words used when documents, communications or ESI were added to the system or index;**
- c. **The number and nature of documents or communications in the system or index which are not searchable as a consequence of the system or index being unable to extract their full text or metadata; and**
- d. **Any limitation in the system or index, or in the search syntax to be employed, tending to limit or impair the effectiveness of keyword, Boolean or proximity search in identifying documents or communications that a reasonable person would understand to be responsive to the search.**

A court will permit “discovery about discovery” like this when a party demonstrates why an inadequate index is a genuine problem. So, let’s explore the rationale behind each inquiry:

a. Tokenization Rules - When machines search collections of documents for keywords, they rarely search the documents for matches; instead, they consult an index

of words extracted from the documents. Machines cannot read, so the characters in the documents are identified as “words” because their appearance meets certain rules in a process called “tokenization.” Tokenization rules aren’t uniform across systems or software. Many indices simply don’t index short words (e.g., acronyms). None index single letters or numbers.

Tokenization rules also govern such things as the handling of punctuated terms (as in a compound word like “wind-driven”), case (will a search for “roof” also find “Roof?”), diacriticals (will a search for Rene also find René?) and numbers (will a search for “Clause 4.3” work?). Most people simply *assume* these searches will work. Yet, in many search tools and archives, they don’t work as expected, or don’t work at all, unless steps are taken to ensure that they will work.

b. Stop Words – Some common “stop words” or “noise words” are simply excluded from an index when it’s compiled. Searches for stop words fail because the words never appear in the index. Stop words aren’t always trivial omissions. For example, “all” and “city” were stop words; so, a search for “All City” will fail to turn up documents containing the company’s own name! Words like side, down, part, problem, necessary, general, goods, needing, opening, possible, well, years and state are examples of common stop words. Computer systems typically employ dozens or hundreds of stop words when they compile indices.

Because users aren’t warned that searches containing stop words fail, they mistakenly assume that there are no responsive documents when there may be thousands. A search for “All City” would miss *millions* of documents at All City Indemnity (though it’s folly to search a company’s files for the company’s name).

c. Non-searchable Documents - A great many documents are not amenable to text search without special handling. Common examples of non-searchable documents are faxes and scans, as well as TIFF images and some Adobe PDF documents. While no system will be flawless in this regard, it’s important to determine *how much* of a collection isn’t text searchable, *what’s* not searchable and whether the portions of the collection that aren’t searchable are of *particular importance* to the case. If All City’s adjusters attached scanned receipts and bids to e-mail messages, the attachments aren’t keyword searchable absent optical character recognition (OCR).

Other documents may be inherently text searchable but not made a part of the index because they’re password protected (i.e., encrypted) or otherwise encoded or compressed in ways that frustrate indexing of their contents. Important documents are often password protected.

d. Other Limitations - If a party or counsel knows that the systems or searches used in e-discovery will fail to perform as expected, they should be obliged to affirmatively disclose such shortcomings. If a party or counsel is uncertain whether systems or searches work as expected, they should be obliged to find out by, e.g., running tests to be reasonably certain.

No system is perfect, and perfect isn't the e-discovery standard. Often, we must adapt to the limitations of systems or software. But you have to know what a system *can't* do before you can find ways to work around its limitations or set expectations consistent with actual capabilities, not magical thinking and unfounded expectations.



Craig Ball, of Austin is a Board Certified Texas trial lawyer, law professor (University of Texas) and accredited computer forensics expert who has dedicated his career to teaching the bench and bar about forensic technology and trial tactics. Craig hung up his trial lawyer spurs to till the soils of justice as a court-appointed special master and consultant in electronic evidence, as well as to teach and publish on computer forensics, emerging technologies, digital persuasion and electronic discovery. Fortunate to supervise, consult or serve as Special Master in some of the world's largest and most prominent electronic discovery matters, Craig greatly values his role as an instructor in computer forensics and electronic evidence to the Department of Justice and other law enforcement and security agencies. Mr. Ball also serves on the faculty of the Georgetown University Law School E-Discovery Academy and sits on the CCE Certification Board of the International Society of Computer Forensic Examiners.