

Defending the Algorithm™ Newsletter

Edition 3 By Henry M. Sneath, Esq.

Artificial Intelligence (AI) in Civil Litigation Document Review and Production: Technology Assisted Review (TAR 2.0)

Welcome back to *Defending the Algorithm*™, a LinkedIn newsletter helping defense attorneys, insurance professionals, corporate counsel and clients to navigate the intersection between artificial intelligence and the law. This newsletter was written and edited with assistance from Claude Opus 4.6 from Anthropic and Google Gemini 3.0 Pro and with research confirmation by Westlaw Precision AI. It is a companion to our podcast and blog series, available at: [Defending the Algorithm](#)™. Thanks to Houston Harbaugh attorney [Eric Spada](#) for his research, writing and editing contributions to this post.

In Edition 1, we examined the insurance discovery showdown in *Estate of Lokken v. UnitedHealth Group*, where plaintiffs are fighting to crack open the algorithmic “black box” behind AI-driven claims processing. In Edition 2, we took a sharp turn into trademark territory—and examined what happens when Silicon Valley’s biggest players choose brand names that collide with established marks in the rapidly emerging field of AI-powered, screenless computing. Now, in Edition 3, we shift gears to the latest in the ever-evolving, AI-assisted litigation e-Discovery process – specifically technology-assisted review (“TAR”), and its most recent iteration: TAR 2.0

What Is Technology-Assisted Review (TAR)?

Before diving into the case law, some context for readers who may not yet have encountered TAR in practice. Technology-Assisted Review (“TAR”) is a process by which machine learning algorithms are used to identify relevant documents in litigation discovery, significantly reducing the costs and time associated with manual linear review—by some estimates, 40–60%. In a world where large commercial and IP litigation can generate millions of electronically stored documents, TAR has become an indispensable tool in the modern litigator’s arsenal.

TAR 1.0 v. TAR 2.0

The Sedona Conference has provided this succinct description of the differences between TAR 1.0 and TAR 2.0:

The terms “TAR 1.0” and “TAR 2.0,” . . . refer to contrasting TAR workflow methodologies. The earlier of the TAR workflows to emerge, often known as TAR 1.0, refers to the use of

discrete training sets within the entire review population. Then, counsel may or may not engage in further responsiveness review of the categorized documents. By contrast, TAR 2.0 refers to a workflow where, generally, every document the TAR model identifies as most likely to be responsive is prioritized for review by human reviewers, and their coding further trains the algorithm. *The Sedona Conference, The Sedona Conference Tar Case Law Primer, Second Edition A Project of the Sedona Conference Working Group on Electronic Document Retention and Production (Wg1)*, 24 *Sedona Conf. J.* 1, 16–17 (2023). [Ed. Note: In other words, human reviewers following the lead of technology review.]

TAR 2.0, also known as Continuous Active Learning (“CAL”), represents the current state of the art. Rather than relying on a fixed training set, TAR 2.0, using AI tools, continuously updates and refines its model with every coding decision made by human reviewers during the review process itself. Platforms such as Relativity® employ active learning protocols that prioritize the most likely relevant documents for human review, with each reviewer’s coding decision further training the algorithm in real time.

The TAR workflow generally proceeds as follows: counsel sets review goals, trains the system (either through a seed set or through continuous coding), validates the model’s performance through statistical sampling—measuring both recall (the percentage of relevant documents found) and precision (the percentage of retrieved documents that are actually relevant)—and then reviews only the documents the algorithm identifies as most likely relevant. A “control set” or “gold standard” sample is often used to test the TAR system’s accuracy without biasing the training data.

The benefits of TAR are well established: increased consistency compared to manual review, faster review timelines, and significantly lower costs. But TAR 1.0 is not without its challenges. It requires initial manual training by knowledgeable reviewers, its output depends heavily on the quality of input, and the “black box” nature of machine learning algorithms can raise defensibility concerns—making thorough documentation of the process essential. In the TAR 2.0 protocol, documents already culled by some application of TAR can be reviewed for example, at what we call the 1L level by less case knowledgeable reviewers (like less expensive contracted document review lawyers) and then spot checked for accuracy by 2L reviewers. The 2L’s are more substantively aware of issues in the case (and generally have higher billing rates) and can do a quality check assessment of the 1L review. The cost savings come from lower rate 1L reviewers conducting initial review of large document batches, with higher rate lawyers handling 2L review of smaller batches of documents identified and coded first by TAR and then by 1L reviewers. Machine learning can further train the review process by 2L review and corrective coding if needed. Courts have generally embraced TAR as a valid and often superior method for document review compared to traditional linear review.

The seminal case *Da Silva Moore v. Publicis Groupe*, 287 F.R.D. 182 (S.D.N.Y. 2012), established early judicial approval of TAR, and courts since then have continued to affirm that TAR methodologies can be not only acceptable, but preferable to manual review when properly implemented with transparency and adequate validation protocols.

Practice Tip: The key concept for the TAR review process is the “defensibility” of the process if challenged by opposing counsel. If the process is challenged in court, there will likely be a need for affidavit support for the process, either by counsel or vendor, so documentation of the process is fundamental. Create a protocol (if agreed to by all parties then so much the better) and document the compliance with that protocol as the process moves forward.

With that foundation in place, we turn to a recent and instructive example of how courts are engaging with the next generation of TAR technology.

TAR 2.0: An Exemplar Case Study

While technically optional, the use of a negotiated TAR protocol between the parties is good practice and is consistent with the principles of cooperation. *See* § 33:43. TAR protocols and negotiation, 4 Bus. & Com. Litig. Fed. Cts. § 33:43 (5th ed.). Given the exponential speed with which AI tools are advancing, it feels inevitable that the use of TAR in large commercial and IP cases will go from optional to mandatory in the not too distant future. And in that regard, there is bad news for practitioners who have just started to feel comfortable with technology-assisted review TAR 1.0 as TAR 2.0 is already here. As illustrated by the court’s decision in *In re Insulin Pricing Litigation*, No. 23-md-3080 (BRM) (RLS), 2025 WL 1112837 (D.N.J. Apr. 11, 2025), and as initially unsettling as it may be, we lawyers must become accustomed to artificial intelligence increasingly being the mechanism that guides our large document reviews, but which reviews certainly require followup human review, certification and validation.

Insulin Pricing Litigation

In re Insulin Pricing, Magistrate Judge Rukhsanah L. Singh largely permitted the Pharma defendants to utilize the latest TAR 2.0 techniques in document production. Addressing a number of technological disagreements, each of which are summarized below, Judge Singh repeatedly decided that allowing the technology to guide the review process was proper and consistent with the mandates in the Federal Rules that the court reach a “just, speedy, and inexpensive determination” using a “reasonable inquiry” standard. Fed. R. Civ. P. 1, 26. There was general agreement by all parties as to the structure of the TAR process, which agreement the Court commended, but there were also 3 main areas of disagreement within the components of that structure.

1) Training Methodology

The parties disputed “the precise question of whether [defendants] may choose to limit training of the TAR model ‘towards the end of the review’ to only those records that have gone through quality control review.” *Insulin Pricing* at 3. In our description above, we referred to that as 2L review. The court allowed the defendants to proceed as they proposed, stressing that “[f]lexibility in any review process, whether TAR or otherwise, particularly in proceeding with complex ESI discovery can be critical.” *Id.* at 4.

2) Criteria for Stopping TAR Review

The plaintiffs sought pre-set stopping of review criteria which expressly specified when defendants could or must pause review and begin validation. *See id.* Defendants, by contrast, argued that they should determine the stopping point “at the point in time with the facts of the actual in-process workflow and its progress in hand, rather than dictated by *ex ante* projections or rules of thumb.” *Id.* In TAR review there is a review stopping point that is dictated by various case features including the volume of documents, the complexity of the search and the measurement and estimate of probabilities as to the overall quality of the document review. In other words, there is a point at which the database controllers feel that the probability is high that a vast majority of responsive, non-privileged documents have been found through AI assistance. Yet again, the court agreed with the defendants: “Prior court experiences with the application of TAR models . . . reflect that a reasonable and proportional stopping point is not solely quantitatively based, but also qualitative. . . . [T]he better course is to determine the stopping point in real time as the process proceeds.” *Id.* at *5 (cleaned up).

3) Validation Sampling Methodology

This is the only matter on which the court adopted the plaintiff’s proposal. Validation is the process of a final QC check by all parties and the protocol varies by case and complexity. It can be done on different sets of selected documents either by agreement or court order. Here it was by court order. For example, the subset of documents to be validated might include documents reviewed by the TAR process and in the following hypothetical sample categories: (1) documents marked responsive, (2) documents marked non-responsive by human reviewers, and (3) documents excluded by TAR from manual review.

In deciding on a protocol, the court premised its decision as follows: “To be clear, the Court is reluctant to force a responding party to adopt validation metrics imposed by a requesting party; the premise that a responding party is best equipped to determine its search methodology remains true even as technology evolves.” *Id.* at *6. Nevertheless, the court required defendants to validate their TAR model as requested by plaintiff, with certain statistical details, including a sampling from the full TAR document universe and not the proposed and more targeted subsets. *Id.*

Related Issues Presented to Judge Singh: Disclosure Requirements

The court declined to require defendants to disclose “information concerning the stratum” from which each document was drawn, instead allowing defendants to simply disclose non-privileged responsive documents from the validation samples. *See id.* at *7.

Related Issues Presented to Judge Singh: Validation Recall Target

The Validation Recall Estimate is a key performance metric used to measure the success of a TAR process and to determine a “completeness” estimate that allows document review to cease. It is done by the database controllers, with input from, and collaboration with the producing counsel and uses algorithms to determine the approximate percentage of total relevant,

non-privileged documents which have been identified and are to be produced. Here again, the court declined to adopt plaintiff's proposed pre-set validation recall target, instead opting to allow defendants to target a reasonable and proportional recall rate and potentially adjusting as the process required. *Id.* at *7. The recall estimate is a fluid number and again is case specific. Common validation techniques include:

- **Elusion Test:** Measuring the accuracy of documents categorized as not relevant.
- **Control Sets:** Testing the algorithm's performance on a small, known sample.
- **Checking Seed Set/Topic Clusters:** Ensuring the model has covered all relevant subject areas.

Conclusion - Takeaways

With each issue, Judge Singh largely erred on the side of the technology, demonstrating a judicial willingness to engage with technical aspects of TAR 2.0 methodology and to trust the technology – really, the AI – to guide the review. Unlike TAR 1.0, which often required the use of a sample set of training documents, TAR 2.0 utilizes continuous learning, to create a streamlined, and hopefully time saving method of large volume document review. It is largely a result of the creation of emails which typically comprise a large percentage of the documents downloaded from the drives and network of custodians. Accordingly, as the implementation of TAR 2.0 becomes more common and eventually mandatory, it is crucial to have counsel who understand the ins-and-outs of TAR 2.0 and can negotiate with opposing counsel and engage with judges from a position of knowledge and strength so that the resultant protocol is fair, efficient, and sufficiently protects client interests.

Connect With Us

Defending the Algorithm™ Podcast & Blog: [Click Here.](#)

Henry M. Sneath, Esq.

Houston Harbaugh, P.C. Pittsburgh, Pennsylvania

sneathhm@hh-law.com and 412-288-4013

In cooperation with the DRI Center for Law and Public Policy — AI Task Force.

This newsletter provides general information and does not constitute legal advice. The views expressed are those of the author and do not necessarily reflect the views of Houston Harbaugh, P.C. or its clients.

© 2026 Henry M. Sneath for Houston Harbaugh, P.C. All rights reserved.

Defending the Algorithm™ is a mark used by Houston Harbaugh, P.C. for which federal registration is in progress