

Jon Polenberg
Shareholder
Phone: 954.364.6037 Fax: 954.985.4176
jpolenberg@beckerlawyers.com

Becker

1 East Broward Blvd.
Suite 1800
Ft. Lauderdale, FL 33301

January 21, 2026

VIA ELECTRONIC SUBMISSION

Committee on Rules of Practice and Procedure
Administrative Office of the United States Courts
One Columbus Circle, NE
Washington DC 20024

Public Comment on Proposed Rule of Evidence 707

The following comment is from me, who is a commercial litigation trial lawyer, serving as Vice Chair of my firm's Litigation Practice Group and a member of Working Group 13 of The Sedona Conference, which focuses on artificial intelligence and the law. I am also Past Chair of The Florida Bar's Business Law Section. But this comment reflects only my individual views and not the views of any organization, working group, client, or institution.

I. Introduction and Summary of Position

This comment responds to the Committee considering proposed Rule of Evidence 707. The proposal raises important questions about how the Rules should address emerging technology while preserving doctrinal coherence, proportionality, and judicial discretion. Those questions implicate not only technical accuracy but also the structure courts use with evidentiary decision-making and allocating responsibility between judges and juries.

Proposed Rule 707 rests on a flawed premise. AI-generated outputs do not constitute testimony and do not express human opinion. They represent machine-generated information that existing evidentiary doctrine already governs with rigor and flexibility.

The Rules regulate admissibility through relevance, authentication, foundation, proportionality, and adversarial testing. Courts have applied those principles to successive generations of machine-generated evidence without creating technology-specific regimes. Proposed Rule 707 would depart from that discipline. By elevating AI outputs into a special admissibility category, the proposal would distort litigation incentives, invite unnecessary technical disputes, and narrow judicial discretion without improving accuracy or fairness.

Five points support that conclusion. First, the Rules preserve a clear distinction between evidence and expert testimony. Rule 707 blurs that boundary. Second, courts have managed

complex machine-generated outputs, including polygraph examinations, through existing doctrine. Third, courts tolerate substantial reliability risk in human evidence, including eyewitness identification, and rely on adversarial testing rather than categorical exclusion. Fourth, parties already possess effective tools to challenge AI outputs under the existing Rules. Fifth, the proposal would predictably convert ordinary evidentiary disputes into routine technical litigation.

II. First Principles: Evidence Versus Expert Testimony

The Rules of Evidence begin with a simple proposition: evidence consists of information offered to prove or disprove a fact with legal consequence. That information appears in many forms, including documents, recordings, images, physical artifacts, data compilations, and machine-generated outputs. The Rules regulate admissibility through relevance under Rules 401 and 402, authentication and foundation under Rule 901, personal knowledge where Rule 602 requires it, hearsay limits when statements appear, and proportional balancing under Rule 403.

Courts apply those principles every day to machine-generated and machine-assisted information. Radar readings, breathalyzer results, GPS and telematics data, cell-site location records, software logs, automated financial systems, and medical imaging devices generate information modern litigation depends upon. Polygraph examinations provide a useful illustration. A polygraph system records physiological signals, processes those signals through algorithms, and produces numerical outputs that carry scientific authority's appearance. Courts approach polygraph evidence with caution. Some jurisdictions exclude it as policy. Others admit it only by stipulation or subject it to careful Rule 403 balancing and foundation. Courts have never treated the polygraph machine as a witness or an expert. The machine generates data. Human witnesses explain administration and interpretation. Existing evidentiary doctrine governs admissibility. See *United States v. Scheffer*, 523 U.S. 303 (1998).

That practice reflects a stable institutional judgment. Machines generate information. They do not perceive events, form beliefs, exercise judgment, or make assertions. Courts authenticate outputs, examine reliability, and assign weight through adversarial testing. AI-generated outputs fit squarely within that tradition. An AI system processes inputs through programmed statistical mechanisms and produces an output. However complex the computation, the system does not possess intent, memory, perception, or reasoning in the human sense. It generates data. The evidentiary inquiry therefore turns on authentication, relevance, and weight, not quasi-testimonial status.

A different analysis governs human opinion testimony. When a witness offers conclusions derived from specialized knowledge jurors cannot readily evaluate, Rule 702 and the Daubert framework impose gatekeeping to protect the fact-finding process. Gatekeeping examines qualifications, methodology, application, and fit. The rule seeks to prevent powerful opinion testimony from reaching juries when reasoning fails to satisfy minimum reliability standards. See *Daubert v. Merrell Dow Pharmaceuticals, Inc.*, 509 U.S. 579 (1993); *Kumho Tire Co. v. Carmichael*, 526 U.S. 137 (1999).

Concerns about bias, selective inputs, manipulation of assumptions, and opaque reasoning animate that function. Courts confront those risks routinely in expert testimony. Parties

challenge data selection, testing protocols, model construction, analytical choices, and litigation-driven tuning. Courts address those risks through disclosure, cross-examination, competing experts, and exclusion when methodology fails Rule 702's requirements.

Those same concerns now appear in discussions about AI-generated outputs. Prompt engineering can shape results. Input selection can embed bias. Iterative prompting can steer outputs toward preferred conclusions. None of those risks transforms machine-generated information into human opinion testimony. They describe provenance, transparency, and weight questions. Courts already possess tools to expose manipulation through adversarial testing and to exclude evidence that misleads or confuses juries under Rule 403.

Absent a technology-specific rule, discovery into AI output generation will often proceed under ordinary civil discovery principles rather than under structured disclosure limits governing expert methodology. Parties will seek prompts, inputs, configurations, procedures, version histories, and vendor documentation as ordinary electronically stored information. Courts will manage scope through proportionality, burden, confidentiality protections, and protective orders. That reality may increase transparency and litigation cost. It does not justify reclassifying machine-generated outputs as expert testimony or importing *Daubert* gatekeeping into domains where authentication, foundation, and adversarial testing already govern admissibility. Discovery management belongs to discovery doctrine, not embedded in evidentiary admissibility rules.

The proper response to bias, manipulation, and steering lies in disciplined application of existing evidentiary and procedural controls rather than reclassification of machine outputs. Courts may admit expert testimony to explain system behavior when assistance proves necessary. Courts should not treat the output itself as expert opinion. Doing so expands gatekeeping beyond its institutional purpose and blurs the boundary between human inference and mechanical production.

Courts have preserved that boundary for decades. Judges permit experts to explain complex data machines generate. Judges do not transform machines into autonomous expert actors subject to qualification and credibility analysis. That restraint preserves doctrinal coherence and prevents unnecessary expansion of expert gatekeeping into domains authentication and adversarial testing already regulate.

Proposed Rule 707 would unsettle that balance by elevating AI-generated outputs into a category associated with expert testimony. The existing Rules already provide sufficient tools to address bias, manipulation, transparency, and reliability without restructuring admissibility doctrine. Treating machine outputs as quasi-expert evidence would expand gatekeeping into areas historically governed by authentication, relevance, and adversarial testing, and would erode the conceptual boundary that separates human inference from machine generation.

III. Polygraphs and Institutional Continuity

Courts have confronted technologies promising objective measurement yet carrying disputed reliability and risk that juries overvalue precision. Polygraph examinations illustrate the point. A polygraph system records physiological responses, processes signals through algorithms, and generates numerical outputs that appear scientific and authoritative.

Judicial experience reflects restraint rather than technological exceptionalism. Many jurisdictions exclude polygraph evidence as policy. Others admit it only by stipulation or under careful foundation and Rule 403 balancing. *Scheffer* confirms rule makers may categorically exclude a machine-assisted test based on reliability concerns without offending constitutional protections. Courts scrutinize examiner qualifications, testing conditions, equipment calibration, and methodological limits. Courts also recognize danger that juries overvalue machine-generated precision.

Even with those concerns, courts have not recast polygraph machines as expert witnesses or created a separate evidentiary regime. The machine generates data. Human witnesses explain administration and interpretation. Existing doctrine governs admissibility and weight.

That pattern reflects broader institutional commitment. Courts adapt to new technologies by applying established evidentiary principles rather than creating technology-specific rules that displace settled doctrine. That approach preserves coherence, predictability, and judicial discretion.

AI-generated outputs present no reason to depart from that tradition. Courts can address reliability concerns through authentication, foundation, proportionality, expert explanation when helpful, and adversarial testing. Proposed Rule 707 would abandon that discipline.

IV. Reliability Risk Does Not Justify Special Evidentiary Regimes: Eyewitness Identification

Human testimony often carries substantial and well-documented reliability risk. Eyewitness identification provides the clearest example. Empirical research over decades has demonstrated that memory reconstructs rather than records experience. Stress, lighting, distance, weapon focus, cross-racial perception, and post-event suggestion distort recall. Confidence often increases as accuracy declines.

Courts acknowledge those limits. Appellate doctrine, jury instructions, and constitutional safeguards reflect sustained concern about suggestive procedures and memory fallibility. Courts nevertheless admit eyewitness testimony and rely on adversarial testing rather than categorical exclusion. *See Manson v. Braithwaite*, 432 U.S. 98 (1977); *Perry v. New Hampshire*, 565 U.S. 228 (2012).

Trial courts manage eyewitness evidence through cross-examination, cautionary instructions, constitutional screening, and Rule 403 balancing. Those mechanisms calibrate risk without disabling factfinding.

That institutional choice matters. The Rules tolerate known error rates because adversarial testing and judicial supervision provide adequate safeguards. The law does not require perfection as admissibility's condition.

Against that backdrop, imposing stricter structural barriers on machine-generated outputs than on human memory lacks justification. If courts trust juries to evaluate eyewitness testimony

despite documented fallibility, courts can trust the same institutional mechanisms to evaluate AI-generated outputs. Reliability risk alone does not justify a separate evidentiary regime.

V. Existing Rules Provide Effective Tools to Challenge AI Outputs

Courts already possess a comprehensive framework to assess admissibility, reliability, and probative value without elevating machine outputs into expert testimony. The Rules allocate responsibility across authentication, relevance, hearsay, probative balancing, and adversarial testing. Each Rule performs a distinct function, and together they provide calibrated judicial control that preserves flexibility and proportionality.

Rule 901 requires the proponent to produce sufficient evidence to support a finding that the item is what the proponent claims it is. That general proposition governs authentication across all categories of evidence, including machine-generated information. The rule frames the inquiry around whether a reasonable factfinder could conclude that the exhibit genuinely reflects the source and process the proponent asserts.

Courts authenticate machine-generated evidence by requiring a foundational showing sufficient to permit a reasonable factfinder to conclude that the identified system generated the proffered output, operating in the claimed condition, using the claimed inputs, and preserved without alteration. Courts apply that approach to breath-testing instruments, radar devices, automated business systems, forensic software, and digital logs. The proponent typically identifies the generating system, establishes operating condition and configuration, demonstrates input integrity where inputs affect the output, and accounts for preservation and chain of custody.

This showing does not require proof of technical perfection or exhaustive validation. It requires a foundation adequate to support authenticity under Rule 901(a) and nothing more. Where a party cannot establish provenance or system integrity, courts may exclude the evidence or limit admissibility to narrower purposes.

Where disputes remain, courts may permit focused discovery or testimony to resolve foundational uncertainty. Courts also retain discretion to sequence proof, impose protective limits, or require corroboration where appropriate. These practices enforce transparency and accountability without converting authentication into reliability gatekeeping.

AI-generated outputs follow the same logic. A proponent must identify the system that generated the output, describe the configuration and inputs that produced it, and establish that the output reflects the system's operation rather than alteration or corruption. The system's complexity may affect the foundation a court requires, but it does not change the governing rule or analytical structure.

Rules 401 and 402 require evidence to make a fact of consequence more or less probable. Rule 104 assigns courts responsibility for resolving preliminary questions affecting admissibility. Together, these rules require a proponent to establish a logical and factual connection between the proffered evidence and the disputed issue.

Courts apply these principles routinely to machine-generated evidence. Judges examine what the output measures, what variables influence the result, and whether the output meaningfully bears on the factual proposition the proponent advances. When a system's assumptions, inputs, or operating context disconnect the output from the disputed issue, courts may exclude the evidence as irrelevant or may require additional foundation to establish fit.

That inquiry does not demand abstract judgments about general system reliability. It asks whether the specific output, generated under specific conditions, advances proof of a material fact in the case. Courts may resolve close questions through Rule 104 hearings, targeted foundational testimony, or limited discovery.

AI-generated outputs present the same inquiry. Courts may examine what the output reflects, what inputs and prompts shaped it, and whether the output meaningfully informs the factual dispute. Where the connection proves attenuated or speculative, Rules 401, 402, and 104 supply authority to exclude or limit the evidence without converting admissibility into a generalized technical audit.

Rule 801 defines hearsay as a statement made by a declarant offered for its truth. When an exhibit reflects mechanical or computational processing rather than a human assertion, no declarant exists and the hearsay rule does not apply. Courts have long recognized that machine-generated data falls outside the hearsay definition for that reason. See *United States v. Washington*, 498 F.3d 225, 231–32 (4th Cir. 2007).

When a machine output reproduces or quotes human statements, ordinary hearsay analysis governs in the same manner as if a human witness offered the material. Courts may require redaction, limiting instructions, or exclusion when no exception applies. These safeguards ensure that machine outputs do not bypass traditional evidentiary limits on human assertions.

Parties may attempt to avoid hearsay limits by suppressing quotations while generating synthesized outputs from inadmissible sources. That maneuver does not evade the Rules. Courts may examine foundation, relevance, and jury-misleading risk under Rules 104 and 403 when the output adds no independent analytical value and merely repackages excluded material.

Opposing parties may expose dependency on inadmissible sources through cross-examination and competing proof. Courts apply similar controls to summaries, compilations, and expert reliance on inadmissible material. These mechanisms preserve parity between machine-generated presentations and traditional evidentiary forms.

Rule 403 authorizes exclusion when probative value substantially outweighs the risk of unfair prejudice, confusion, or misleading presentation. Courts use that authority to regulate how evidence reaches the jury rather than only whether it enters the record. The rule supplies flexible authority calibrated to the context of each case.

In cases involving technical or machine-generated evidence, judges may limit presentation format, restrict demonstrative overlays, sequence witnesses to provide context before technical exhibits appear, and instruct jurors regarding appropriate weight and limitations.

Courts may narrow admissible scope or exclude cumulative material when complexity threatens distortion rather than illumination. These tools protect the jury's role without suppressing relevant proof.

AI-generated outputs present similar concerns. Where computational complexity risks overvaluation or misunderstanding, Rule 403 permits courts to calibrate presentation and scope. The rule avoids categorical exclusion while preserving proportional safeguards.

Rules 611 and 607 empower courts to control witness examination and permit impeachment. Those rules structure the adversarial process through which parties test credibility, foundation, and reliability. Judicial control over mode and order of proof preserves clarity and efficiency.

In cases involving machine-generated evidence, parties may replicate inputs, test alternative systems, compare outputs across versions, and demonstrate variability under controlled conditions. Cross-examination may probe input selection, stability, sensitivity, and limitations. Courts may supervise testing protocols to preserve proportionality and fairness.

AI-generated outputs fit within that same adversarial framework. Concrete testing allows parties to present empirical demonstrations rather than abstract critique. Rules 611 and 607 give judges tools to manage scope, prevent abuse, and preserve clarity for the jury.

Together, these Rules distribute responsibility across authentication, relevance, attribution, proportionality, and adversarial testing. The framework preserves judicial discretion, protects fairness, and maintains doctrinal coherence. It avoids unnecessary expansion of expert gatekeeping into domains the Rules already govern.

VI. Predictable Litigation Distortions Under Proposed Rule 707

Proposed Rule 707 would alter how courts resolve foundational disputes by transforming technical reliability from a question that ordinarily informs weight and case-specific admissibility into a threshold gatekeeping determination. Rather than evaluating reliability through authentication, relevance, proportional balancing, and adversarial testing, courts would confront reliability as a preliminary admissibility barrier that the court must resolve before the jury can consider the evidence at all. That structural shift carries predictable consequences for how parties litigate.

Once admissibility turns on a formal reliability determination, parties will have strong incentives to litigate the technical characteristics of the generating system rather than the probative value of the specific output offered. Parties will retain competing experts to analyze architecture, training composition, processing pathways, update histories, and claimed error rates even when the output bears on narrow or peripheral factual issues. What might otherwise remain a matter for cross-examination or contextual weight will migrate into pretrial admissibility disputes.

That migration will expand discovery and motion practice independent of the evidentiary stakes in the case. Parties will pursue system documentation, model updates, training

methodologies, vendor relationships, audit materials, and version histories to build or defeat admissibility records rather than to illuminate disputed facts. Discovery will widen because admissibility will depend on technical pedigree rather than on relevance and proportional fit. Motions practice will multiply because each contested output will invite threshold litigation over system reliability. Courts will confront recurring technical disputes that consume time and resources disproportionate to the evidentiary value the output contributes.

The proposal also introduces structural asymmetry. Courts will continue admitting fallible human testimony subject to adversarial testing, cross-examination, and Rule 403 balancing, even when empirical research documents persistent error rates. At the same time, the rule would impose heightened procedural barriers on machine-generated evidence regardless of context, materiality, or cumulative value. That imbalance lacks doctrinal justification and undermines the Rules' commitment to evenhanded evidentiary treatment.

Finally, fixed admissibility standards would narrow judicial discretion at precisely the point where flexibility matters most. Technology evolves rapidly. System architectures, training methods, and deployment practices change on short cycles. A rigid admissibility framework risks entrenching outdated assumptions and forcing courts to litigate technical compliance rather than exercise contextual judgment. The Rules traditionally preserve adaptability by entrusting trial courts with calibrated discretion rather than prescribing technology-specific admissibility regimes.

VII. Conclusion

The Rules already provide a coherent framework for evaluating machine-generated information. Courts authenticate evidence, examine foundation, assess relevance, balance probative value against risk, and permit expert testimony when it assists factfinders. That structure has guided evidentiary doctrine across successive waves of technological change and has preserved stability without sacrificing flexibility.

AI-generated outputs fit within that framework. They represent machine-generated information rather than testimony or human opinion. Courts can test provenance, expose limitations, and assign appropriate weight through adversarial testing and calibrated judicial supervision.

Proposed Rule 707 would displace that balance. It would invite routine technical litigation, expand discovery by necessity rather than judgment, distort proportionality, and narrow judicial discretion without corresponding benefit. The rule would shift evidentiary disputes away from case-specific relevance and toward generalized system compliance.

For these reasons, the Committee should decline to adopt proposed Rule 707 and continue relying on the existing Rules to govern admissibility and evaluation of AI-generated evidence. That course preserves doctrinal coherence, institutional neutrality, and proportional adjudication. It also maintains the flexibility courts require to adapt to technological change without premature codification.

Sincerely,

A handwritten signature in blue ink, appearing to read "John".

Jon Polenberg