



SUGGESTED TOPICS AND QUESTIONS FOR CONSIDERATION
by the
ADVISORY COMMITTEE ON EVIDENCE RULES
on
DRAFT RULES OF EVIDENCE 707 and 901(c)

September 3, 2026

Introduction

The Advisory Committee on Evidence Rules’ upcoming Fall 2026 mini-conference (“symposium”) presents an important opportunity to evaluate whether draft Federal Rules of Evidence 707 and 901(c) are necessary, and if so, how they can be written to provide meaningful guidance to courts while remaining workable for litigants and companies as technology evolves.¹ The symposium is well positioned to address foundational issues concerning the admissibility, reliability, authentication, and procedural treatment of AI-generated evidence—issues that will increasingly affect litigation as artificial intelligence becomes more widely integrated into society and the administration of justice. That discussion should include not only whether AI-generated evidence can satisfy an admissibility threshold, but also whether evidence that clears that threshold can be meaningfully tested at trial without a sponsoring expert available for cross-examination. Lawyers for Civil Justice (“LCJ”)² previously submitted a public comment on proposed Rule 707³ and testified before the Advisory Committee in January 2026. In the spirit of assisting the Advisory Committee and conference participants in this important effort, LCJ offers the following topics and questions for consideration at the symposium.

¹ These comments relate to the redrafted text and Committee Notes as conveyed in the Report of Judge Furman to the Standing Committee on May 17, 2026, at pp. 7-11 and 12-14.

² LCJ is a national coalition of corporations, law firms, and defense trial lawyer organizations that promotes excellence and fairness in the civil justice system to secure the just, speedy, and inexpensive determination of civil cases. Since 1987, LCJ has been closely engaged in reforming federal procedural rules in order to: (1) promote balance and fairness in the civil justice system; (2) reduce costs and burdens associated with litigation; and (3) advance predictability and efficiency in litigation.

³ Lawyers for Civil Justice, *Matching Text to Intent: Revising Proposed Rule 707 to Close the Rule 702 Gap and Define Reliability Standards for Machine Opinions* (Jan. 5, 2025) (“LCJ Comment”), <https://www.regulations.gov/comment/USC-RULES-EV-2025-0034-0013>.

I. RULE 707 DRAFT

A. AI-Specific Reliability Standards in the Rule Text

The version of Rule 707 published for public comment cross-referenced to Rule 702(a)-(d) and imported language meant for human experts (“expert’s knowledge,” “testimony,” “expert’s opinion”) that does not translate well to AI systems. It also failed to identify the specific criteria—training data representativeness, validation methodology, error rates, explicability—that courts have to assess to determine whether an AI output is reliable.

The revised rule reproduces the four Rule 702 factors directly in the text of Rule 707, eliminating the cross-reference. But it retains language pertaining to human expert witnesses contained in Rule 702, rather than referencing AI-specific considerations. The current draft does not explain how the Rule 702 factors should apply to AI generated material. For example, how should courts interpret “sufficient facts and data” in the context of training datasets? Similarly, how should courts interpret “sufficient facts and data” in the context of a prompt or multiple prompts? Who or what “applies” principles and methods in AI systems? How should courts evaluate whether AI output is “helpful to the trier of fact”? The answer to these questions could help appropriately define the scope of any rule.

Some AI-specific considerations are contained in the draft Committee Note: training data sufficiency, validation in similar circumstances, records and auditability, independent evaluation, and disclosure analogous to expert disclosure. That is not how any rule should be structured. As LCJ observed in its comment, the requirements of the rule should be in the rule’s operative text; not in a Note.⁴ Committee Notes are required by the Rules Enabling Act to explain (to Congress and the public) why the amendment is being promulgated.⁵ Notes also can provide guidance, but do not have the force of law—courts and parties can and do disregard them. Reliability factors are essential to the rule’s purpose and should be in the rule itself.

Accordingly, the symposium should examine what AI-specific reliability factors should be included in the rule’s text, and how those factors should be drafted. Experts in AI systems can help the Advisory Committee understand which factors are truly essential, which are case-specific, and how to express them in rule language that is durable as the technology evolves. The discussion should also address whether unexplainable (*i.e.*, “black box”) outputs should ever be admissible, with or without the sponsorship of qualified experts.

B. The Scope and Applicability of Any Rulemaking

Consideration should also be given to the draft’s scope. As a threshold issue, should the draft rule be limited to machine-generated “opinions or otherwise,” as reflected in Rule 702, or should

⁴ LCJ Comment at 5, 7-8.

⁵ 28 U.S.C. § 2703 (“In making a recommendation under this section or under section 2072 or 2075, the body making that recommendation shall provide a proposed rule, an explanatory note on the rule, and a written report explaining the body’s action, including any minority or other separate views.”).

it apply more broadly to all forms of output produced by artificial intelligence? If the latter, should the rule define the scope of machine-generated evidence? Should any rule also define “artificial intelligence” separately, or simply use terminology like “Machine Opinions” or “Machine-Generated Opinions,” which could help narrow the scope the rule? In thinking about the scope of the rule, should it apply only to inferential outputs such as predictions, recommendations, and conclusions, or reach all AI-generated output? And how, if at all, should a rule distinguish purely mechanical outputs, such as measurements, from analytical outputs? Because the current draft seeks to apply Federal Rule of Evidence 702 to the “product of artificial intelligence,” consideration may want to be given to limiting the proposal to “machine opinions” to more closely track the application of Rule 702.

C. The Expert Testimony Requirement—Admissibility vs. Trial

Expert testimony should be the default for establishing admissibility. An expert who provides the foundation for admissibility should be available for cross-examination at trial. A reliability determination made pre-trial does not solve the problem if the AI output goes to the jury with no witness who can be questioned about how the system reached its specific conclusion in the case.

LCJ’s comment warned against allowing parties to establish AI admissibility without an expert, because validation data is often proprietary, vendors control what they disclose, and meaningful reliability review in the absence of any expert who can be cross-examined is very difficult to achieve.⁶

The revised draft rule establishes expert testimony as the “presumptive” requirement for admissibility—a significant improvement over the original proposed rule’s open-ended pathway. But the draft Committee Note contains a caveat: “while an expert is ordinarily required to provide the foundation for admissibility, it would be possible for a proponent to offer that evidence at trial without an expert.” No authority for such a significant deviation from established procedures is offered. This vague suggestion about an undefined “possibility” would undoubtedly encourage novel attempts to reduce AI admissibility standards—contrary to what the Advisory Committee recently sought to accomplish by amending Rule 702. The Note merely acknowledges that the inability to cross-examine is “a concern” and recommends a limiting instruction but does not require a testifying expert at trial.

This gap between pre-trial admissibility and trial presentation is the critical unresolved issue. Pre-trial scrutiny of AI systems is valuable but only addresses general reliability—the overall functioning of the system. Cross-examination at trial addresses specific reliability—whether the system was correctly applied to the facts of this case, what data it relied on, what alternatives it considered, and whether the output accounts for case-specific circumstances. A limiting instruction cannot substitute for that inquiry.

The symposium should address whether the rule should require AI evidence that clears the general admissibility threshold to have a sponsoring expert at trial, someone to address the technology’s case-specific issues and who may be cross-examined. Technology and legal experts should be asked: what meaningful adversarial testing of AI output is possible at trial without a

⁶ LCJ Comment at 6-7.

cross-examinable witness? Is there a workable middle ground—for example, allowing non-expert trial presentation only in cases where the AI system’s application to the specific facts is not in dispute?

D. The “Exceptional Circumstances” Exception

Under the current Rule 707(b) draft, an expert witness would “ordinarily” be required to establish the admissibility of AI-generated evidence, but “exceptional circumstances” may justify the use of other proof to establish reliability. The revised rule’s “exceptional circumstances” language is narrower than the original rule’s open pathway but still preserves the possibility of establishing admissibility through “other proof” such as validation tests and objective studies. The current draft Committee Note suggests this exception might become more significant in the future if AI becomes “self-validating.”

The “exceptional circumstances” exception is currently undefined. It should be cabined with sufficient precision, including supporting citations, to ensure that it does not swallow the rule. Without guardrails, proponents will argue that published accuracy studies, vendor certifications, or widespread prior use constitute the “other proof” needed to satisfy the rule—precisely the kind of evidence that courts and parties are least equipped to evaluate critically without expert assistance and the ability to cross-examine.

The symposium should examine what “exceptional circumstances” means in practice, both today and as technology develops. What types of evidence should—and should not—qualify? How should courts evaluate vendor-supplied validation studies without expert assistance? Who bears the burden when the proponent claims “exceptional circumstances” apply, and what standard of proof applies? Technology experts can help the Advisory Committee understand the difference between rigorous independent validation and self-serving accuracy claims and can help draft language that limits the exception to genuinely exceptional cases.

E. The Rule 703 Interaction—The Expert Reliance Escape Hatch

The revised draft does not address Rule 703’s interaction with draft Rule 707. The Advisory Committee’s report to the Standing Committee⁷ discusses Rule 703 only as a separate issue, focusing on the Confrontation Clause implications of *Smith v. Arizona*, 602 U.S. 779 (2024), and does not contemplate any interaction between Rule 703 and draft Rule 707.

LCJ’s comment raised the concern that Rule 703, which allows experts to rely on facts and data of the type reasonably relied upon in their field, could become an excuse for an end run around Rule 707.⁸ An expert relying on AI output in forming an opinion could proffer that output under Rule 703, bypassing Rule 707’s reliability requirements entirely.

⁷ Report of the Advisory Committee on Evidence Rules, Agenda Book, Committee on Rules of Practice and Procedure (June 3-4, 2026), 525-39, <https://www.uscourts.gov/sites/default/files/document/2026-06-standing-committee-meeting-agenda-book.pdf>.

⁸ LCJ Comment at 9-10.

If an expert radiologist can rely on an AI diagnostic tool under Rule 703 without triggering Rule 707, or if any party can avoid Rule 707 simply by having an expert adopt the AI output as the basis for their opinion, then the reliability protections of Rule 707 are easily circumvented. The Rule’s scope needs to account for this pathway.

Accordingly, the symposium should address whether Rule 707 should expressly govern all litigation use of AI output, including output forming the basis of an expert’s opinion under Rule 703, or whether Rule 702, as applied in the wake of the 2023 amendments, is sufficient for that scenario. The discussion should include all pathways by which AI evidence reaches the jury and ensure that Rule 707’s reliability requirements apply consistently to all such pathways.

F. Procedure and Litigation Mechanics

The revised rule includes a notice provision (Rule 707(c)) requiring “reasonable notice” of intent to offer AI evidence so that the adverse party has a “fair opportunity to meet it.” The Committee Note describes the provision as “intentionally flexible” and directs parties to consult the Civil and Criminal Procedure rules for additional pretrial obligations.

“Reasonable notice” in a vacuum, without specification of timing or content, replicates the vagueness that made the original rule’s reliance on Committee Note guidance inadequate. The January hearing testimony documented in detail how vagueness leads to forum shopping, inconsistent treatment across courts, and ad hoc litigation over scope and timing—exactly the problems a uniform rule is designed to prevent.

The symposium should address “reasonable notice” concretely. What must the notice contain—model identity, training data description, validation studies, error rates? When must this notice be provided, relative to expert disclosure deadlines and the close of fact discovery? The symposium should also consider whether the disclosure issues raised by AI-generated evidence are fundamentally procedural rather than evidentiary. Unlike admissibility determinations, the timing, scope, and content of disclosures are traditionally governed by the Federal Rules of Civil Procedure. Participants should therefore discuss what amendments to Federal Rule of Civil Procedure 26 or related procedural rules the Evidence Rules Committee should propose to provide adequate notice, avoid surprise, and promote consistent nationwide treatment.

G. Is any Rule Necessary?

The symposium should provide a clearer empirical picture of where and how AI evidence is actually being offered in federal courts today, and what problems courts are experiencing in managing it under existing rules.

LCJ’s comment recommended either revising the rule text to match the Advisory Committee’s restrictive intent or withdrawing the rule for further development, explaining that a rule that does not do what it says is worse than no rule at all.⁹

⁹ LCJ Comment at 2-4, 10-11.

The Advisory Committee’s report to the Standing Committee reflects genuine ambivalence. It notes the revised rule would require re-publication, and the Advisory Committee scheduled the Fall symposium before deciding whether to proceed—explicitly reserving the question of whether any rule is necessary.¹⁰ The Advisory Committee noted that while some commentators believe that no problem currently exists, others, particularly from the criminal defense and civil liberties communities, believe that problematic AI evidence is already being offered, especially in criminal cases.

The symposium should therefore consider the threshold question: is there a demonstrated need for Rule 707 now? Should AI evidence initially be addressed through a criminal rule? This discussion should draw on testimony from civil litigators, practitioners in criminal cases (where AI issues appear most acute), and judges who have encountered AI evidence issues at the trial level. The symposium should also address the counterfactual: if Rule 707 is not adopted, what happens? Does the status quo—where AI output can come in through Rule 901(9) or 902(13) with no 702 scrutiny—potentially produce worse outcomes? Is it enough to rely on, as many courts have done, the authority of the gatekeeper to borrow from Rule 702 principles, relying on Rule 102?

Finally, given the continuing rapid development of AI technology, the symposium should consider how, if at all, Rule 707 should account for that reality. Would the Rule benefit from either more flexibility to account for technological developments or from greater specificity to provide practitioners and judges with more clarity and predictability? Is there a way to structure the Rule to provide both? The symposium should consider whether there are ways to draft Rule 707 to allow for future innovation without destabilizing the rule.

* * *

One of LCJ’s core points in its prior submission—that rules must work on their face and cannot rely on Committee Notes to do the work the text should do¹¹—applies with equal force to the revised draft. The revised rule is meaningfully better than the original, but several of its most important protections are relegated to the Note rather than the text, and the undefined “exceptional circumstances” exception and the absence of a trial-level expert requirement leave significant gaps that the symposium should directly address.

II. RULE 901(c) DRAFT

The Advisory Committee’s working draft of Rule 901(c) addresses a genuine evidentiary problem: the emergence of generative AI tools that make it easy to fabricate convincing audio, video, image, and text evidence, and increasingly difficult for courts and jurors—and even experts—to detect the fabrication. The draft reflects careful thinking about the two-sided risk: requiring too little of opponents who claim deepfake fabrication could paralyze trials with unfounded authenticity challenges, while requiring too little of proponents once a credible challenge is raised could allow fabricated evidence to be entered into evidence and relied upon by a court or a jury in deciding a matter.

¹⁰ *Supra* n. 7.

¹¹ LCJ Comment at 5-6.

The following topics are offered as issues the Advisory Committee should examine at the symposium to ensure the rule is workable, consistent with existing authentication doctrine, and durable as the technology continues to evolve. LCJ has not submitted any prior comment on the Rule 901(c) draft.

A. The Threshold Showing—What Is “Evidence Sufficient to Support a Finding of Fabrication”?

Draft Rule 901(c)(1) requires an opponent claiming deepfake fabrication to present “evidence sufficient to support a finding of fabrication” before the court is required to conduct an inquiry. This is designed to prevent bare accusations of deepfakes from disrupting trials without a genuine factual basis. The threshold “evidence sufficient to support a finding” standard is borrowed from the conditional relevance framework of Rule 104(b), but it is not obvious that this standard is well-suited for deepfake claims. Three problems deserve attention.

First, the standard may be too easy to meet. As deepfake technology becomes more widely accessible and understood, an expert could testify that fabrication is “possible” given the characteristics of a piece of digital evidence. “Evidence sufficient to support a finding” of possibility is a low bar. That is appropriate for Rule 104(b) and is consistent with the traditional authentication framework of Rule 901(a). What is novel is using that threshold showing as the trigger for a heightened authenticity inquiry under Rule 104(a).

Second, the standard may be too hard to meet in cases where it matters most. A deepfake specifically designed to evade detection—precisely what sophisticated actors are likely to produce—may leave insufficient forensic evidence of fabrication to satisfy even a *prima facie* standard. The opponent who most needs the rule’s protection may be the least able to invoke it.

Third, the standard gives no guidance on what types of evidence qualify. Does expert forensic analysis suffice? Does circumstantial evidence, such as the fact that the proponent had motive to fabricate, or that the item appeared suddenly without provenance, count? Should the analysis differ depending on the type of content (video vs. static image vs. audio vs. text)?

The symposium should examine how the threshold standard would function in practice and seek ways in which to create more certainty in its application. Technology experts can help the Advisory Committee understand what evidence of fabrication is realistically available to opponents, how quickly detection technology is evolving, and whether the standard needs to be calibrated differently for different types of content or different litigation contexts (criminal vs. civil, trial vs. pretrial).

B. The Proponent’s Burden—Preponderance of the Evidence—Is Higher than Rule 901(a), But Is It Workable?

Once the opponent makes the threshold showing, the proponent must demonstrate authenticity by a preponderance of the evidence under Rule 104(a), a standard the Committee Note acknowledges is higher than the *prima facie* standard ordinarily required under Rule 901(a). The

Note justifies this standard on the grounds that deepfakes are now easy to make and hard to detect.

The preponderance standard raises several practical questions. First, how does a proponent prove a negative—that something has *not* been fabricated—by a preponderance? For digital content with a clear chain of custody, this may be straightforward. For content obtained from third parties, social media, or surveillance systems, establishing non-fabrication by a preponderance is likely to require substantial forensic analysis that is beyond the means of many litigants.

Second, the standard applies to all challenged items once the threshold is met, regardless of the sophistication of the claimed fabrication concern. A proponent defending a text message screenshot faces the same burden as a proponent defending contested surveillance video, even if the realistic risk of deepfake fabrication differs dramatically.

Third, the relationship between the Rule 104(a) preponderance standard and the existing Rule 901(b) illustrations is unclear. The Committee Note states that satisfying a Rule 901(b) illustration will not ordinarily suffice—but “ordinarily” is vague for the reasons discussed, above, in section I(D), generating litigation and inconsistent application.

Accordingly, the symposium should examine whether the preponderance standard is appropriately calibrated, or whether a sliding scale—tied to the strength of the opponent’s threshold showing or the type of content at issue—would be more workable. The Advisory Committee should also hear from practitioners about its potential impact on litigants who lack resources for forensic experts and consider whether implementing the rule might require cost-allocation or appointment of neutral experts in appropriate cases.

C. Scope—Generative AI Only, or All AI-Assisted Fabrication?

Draft Rule 901(c) is expressly limited to fabrication “by generative artificial intelligence,” defined as a computer system that emulates the structure and characteristics of input data to create derived synthetic content. The Committee Note acknowledges this is an open question and flags for symposium discussion whether the rule should cover “all AI as opposed to generative AI.”

The generative AI limitation has both over- and under-inclusion problems. On the under-inclusive side, some of the most concerning evidence manipulation does not involve generative AI in the technical sense—it involves selective editing, splicing, speed manipulation, or context removal that can be accomplished with basic software. A video that has been deceptively edited to remove exculpatory context is not a “deepfake” in the generative AI sense, but it raises identical authenticity concerns. The rule as drafted would not cover it.

On the over-inclusive side, “generative AI” as a technical matter encompasses a wide range of tools, including basic applications like photo filters, autocorrect, and grammar-correction software that “generate” content in a technical sense but raise no serious fabrication concerns.

The statutory definition from former Executive Order 14110¹² may not draw the line in the right place for evidentiary purposes.

There is also a definitional durability problem. Generative AI is evolving so rapidly that a definition anchored to current technology may be obsolete within the rule’s lifespan. The Advisory Committee has already confronted this problem with Rule 707’s “artificial intelligence” definition, borrowed from the National AI Initiative Act, and should consider whether a statutory definition is the right approach for 901(c) as well.

The symposium should also address the scope question, with input from technology experts who can help the Advisory Committee understand the distinction between generative AI fabrication and other forms of digital manipulation, and whether that distinction is meaningful for evidentiary purposes. The Advisory Committee should also consider whether a technology-neutral definition—focused on the nature of the fabrication rather than the tool used—would be more durable.

D. The Notice Requirement

Draft Rule 901(c)(4) would require an opponent pressing a deepfake argument to provide “reasonable pretrial notice” of the intent to present evidence of fabrication. The draft Committee Note explains that no specific time periods are set because the issue may arise in different contexts and acknowledges that a court may excuse late notice if the opponent discovers the deepfake concern after trial has begun.

The notice requirement raises the question whether Rule 901(c)(4), as opposed to the Federal Rules of Civil Procedure, should define notice in civil litigation. In the opposite direction is a strategic use concern. A well-resourced party could use the notice requirement as a tactical tool—forcing the opponent to commit early to a deepfake theory, revealing its forensic strategy before trial.

The symposium should examine whether the notice requirement should be defined by the Federal Rules of Civil Procedure and, if so, whether the Advisory Committee should propose specific rule language to that committee. The Advisory Committee should also consider whether Rule 901(c) should address the resources needed to evaluate deepfake claims, including whether courts should have authority to appoint neutral forensic experts or order cost-sharing in appropriate cases.

E. Coordination with Existing Authentication Doctrine

Draft Rule 901(c)(3) provides that the rule applies to items offered under either Rule 901 or Rule 902. The draft Committee Note addresses the interaction with Rule 901(b) illustrations, concluding that satisfying an illustration “ordinarily” will not suffice to meet the preponderance standard once a deepfake challenge is raised. The same vagueness issue referenced, above, in Part I(D) exist here, as well.

¹² E.O. 14110 was revoked by E.O. 14148 on January 20, 2025.

The relationship between Rule 901(c) and existing authentication doctrine is not fully explained, and the symposium should examine several specific interactions. Rule 902(13) and 902(14), which allow self-authentication of computer-generated records by certificate, present a particular problem. These rules were designed to streamline authentication of digital records, but a certificate of authenticity says nothing about whether the underlying content has been generated or manipulated by AI. A deepfake video accompanied by a proper Rule 902(14) certificate is authenticated only as to its digital form, not as to the content’s genuineness. The rule and its Note should clarify how this interaction works.

Similarly, the interaction with Rule 901(b)(9)—which authenticates the output of a “process or system” that produces an accurate result—needs clarification. That provision was designed for contexts like computerized business records. It cannot apply to deepfakes, which by definition involve technology used to generate inaccurate results. It is not apparent that this part of Rule 901 has any applicability to deepfakes.

The symposium should discuss the interactions between Rule 901(c) and the existing Rule 901 and 902 framework, with particular attention to the self-authentication provisions. The Advisory Committee also should consider whether those provisions need corresponding amendments, or whether the Note to Rule 901(c) can adequately address the interactions without additional rule changes.

* * *

Draft Rule 901(c) addresses a real and growing problem. The two-step framework—threshold showing by opponent, then preponderance by proponent—is conceptually sound. But several of the rule’s most important operational questions are left to “reasonable” standards and flexible notice provisions that could generate the inconsistency and ad hoc litigation that uniform rules are designed to prevent.

Conclusion

The questions presented by draft Rules 707 and 901(c) implicate foundational issues concerning the role and impact of AI in the evidentiary process, the ability of courts to assess reliability and authenticity of such evidence, and the procedural safeguards necessary to ensure fairness to litigants. The upcoming symposium provides an important opportunity to develop a more complete record regarding both the need for rulemaking and the appropriate scope of any proposed amendments. LCJ respectfully offers these topics and questions to assist that effort.