ISSN: 2755-6190

Review Article

Open Access Journal of Artificial Intelligence and Technology

AI: Stare Non Decisis

John McClellan Marshall

Senior Judge, Fourteenth Judicial District of Texas Honorary Professor of the University, UMCS, Lublin, Academician, International Academy of Astronautics, Poland

Corresponding author
John McClellan Marshall, Senior Judge, Fourteenth Judicial District of Texas Honorary Professor of the University, UMCS, Lublin, Academician, International Academy of Astronautics, Poland.

Received: June 25, 2026; **Accepted:** July 02, 2026; **Published:** July 10, 2026

ABSTRACT

AI: Stare non Decisis identifies a new danger arising from the use of artificial intelligence in legal practice. The problem is not simply that large language models may generate fictitious citations or hallucinated authorities. The deeper concern is that lawyers, judges, and institutions may gradually delegate the responsibility for verification to machines. If unverified machine-generated authority enters the legal record, the integrity of precedent itself may be undermined. This article argues that stare non decisis is more than a legal problem. It is an intellectual example that reveals a broader crisis of verification affecting universities, hospitals, and other institutions of knowledge. AI-generated legal citations, student essays, and diagnostic recommendations all expose a common challenge: the production of plausible information is becoming easier than the verification of its reliability. The central question of AI in the modern era is therefore shifting from who governs institutions to who governs the infrastructures that organize knowing itself. The article concludes that preserving the Enlightenment commitment to sapere aude may well require the redesign of legal and educational institutions around verification, accountability, transparency, and stewardship. The future challenge is not whether AI participates in knowledge production, but whether present-day human institutions retain the capability for responsibly determining what counts as trustworthy knowledge.

Keywords: Authenticity, Hallucination, Dataverse, LLM, Discipline

"What is Truth?", said jesting Pilate and would not stay for an answer [1]. The problem is that the court must provide the answer. It is axiomatic in the legal profession that the judicial process, no matter how it is structured, is intended to be a mechanism for determining the "truth" of the matter before the court. To that end, it is presumed that a witness will testify honestly to the best of his or her ability as to what he or she saw, smelled, felt, or heard about a given event.

Normally, this testimony, whether corroborated by other witnesses or documents, becomes the cornerstone of the determination by the judge or jury as to what the facts actually were. Obviously, if a witness intentionally [or through a disability of some sort] cannot accurately recount the event, then the system tends to break down. Even so, if there should be more than one witness or physical evidence to substantiate what is claimed, then the process works very well. In the Anglo-American judicial tradition, this is the first leg of a three-legged stool upon which a judgment will be based.

In the past couple of years, the advent of AI [artificial intelligence] has impacted the judicial process insofar as establishing a factual base. Some of the more interesting cases involve the use of chatbots to perform research into how to establish facts and what cases exist to support these methods. The principal lawyerly justification for this method of legal research is that it saves time. Unfortunately, the Large Language Model [LLM] upon which most, if not all, AI algorithms are currently based, tends to cast a broad net across the dataverse and then synthesizes the result into an answer that, if not carefully phrased at the outset, may be the result of cyber-hallucination. This reality that is at the base of AI in fact mandates that the lawyer go through the answer and verify all of the "facts" and sources that it cites in response. If the lawyer should fail to do so, then he or she may well have committed malpractice in relation to the client and fraud upon the court as well. In the latter case, the consequences can be dire, more than mere money, but potential loss of a license. Against this has to be measured whether the employment of the chatbot is, in fact, a time-saving device. Further, at its root, the LLM is not capable of qualitatively evaluating what it finds and providing an estimate of factual certainty that is superior to an evaluation derived from the "real world" experience of the

Citation: John McClellan Marshall. AI: Stare Non Decisis. Open Access J Artif Intel Tech. 2026. 2(3): 1-4. DOI: doi.org/10.61440/OAJAIT.2026.v2.32

lawyer. At best, it can only provide an opinion. Put another way, even if the output were one that appears to support the lawyer's position, it is essential to reexamine the premise. Indeed, it is this characteristic of AI that defines it as "artificial".

To overcome this flaw in the LLM and the problems that it generates in the legal process, it will likely be necessary to develop an algorithmic model that does allow for mathematical evaluation, a Large Mathematical Model [LMM]. This would clearly be of great value in determining such things as the long term physical and mental impact on an individual from a physical injury or even the proper design of a building. If the data can be expressed in numbers, it may well be a fact, but if it is only expressed in words, it may lead to no more than an opinion. While this would still be a form of intelligence that is "artificial", by narrowing the scope of the algorithmic inquiry an LMM-based algorithm would allow for the answer to be more precise and, therefore, more useful in defining the "fact" of an event.

At the root of the judicial process is the evidentiary base upon which the judge or jury will render a verdict as to the facts and then for the judge to apply the law. It is at this point where the advent of AI platforms, such as Microsoft CoPilot, ChatGPT, and Claude, is having the greatest impact on the judicial process. There is little question that the LLM-driven bots allow for extensive information gathering across the dataverse. The output of these platforms potentially is so great that the user can be overwhelmed by the sheer quantity of information. In such a case, the result can be a major cost factor in the ability of the client to pay for the services of the lawyer to take the time to review the materials. Yet, the ethical requirements of the legal profession may well mandate that the lawyer take that time or face possible malpractice consequences.

An important consideration in the decision to include AI-bots in any phase of the legal/judicial process is to understand that the relationship between the lawyer and the AI is not merely one of human-to-machine. It is a conversation in which the attorney must frame the inquiry in a language that the machine understands and in which it can then formulate a response. Put another way, the LLM model has to be constrained by the inquirer, or else the response is not likely to be accurate or useful. It is as if the attorney is speaking Greek, while the machine hears and responds based on Sanskrit. To that extent, the "conversation" moves toward the limit defined by the "Turing test", if it is assumed by the lawyer that the machine understands the question. Quite simply, LLM-driven platforms cannot provide factually reliable responses; hence, for the AI-bot to be of genuine utility, it must have the LMM analytical capability that allows it to understand the wider nuances of the language that is in use by both user and machine. At the moment, an LMM model with that capability does not exist, so the utility of AI-bots is necessarily limited with the result that a lawyer must, in conformity with the disciplinary rules and to avoid a charge of malpractice, review and check the response of the bot in detail. For example, in Texas, "the signatures of attorneys or parties constitute a certificate by them that they have read the pleading, motion, or other paper; that to the best of their knowledge, information, and belief formed after reasonable inquiry the instrument is not groundless and brought in bad faith or groundless and brought for the purpose of harassment."[1] This specifically imposes upon the lawyer a professional and ethical duty to review the material generated by and AI-bot before presenting it to the court. A failure to do so, no matter how time-consuming, will likely result in serious and negative consequences in court. Foor example, in a recent case, two attornies who not only used fabricated case citations in appellate briefs, but denied having used generative AI in their production were subjected to monetary sanctions and the suspension of their right to practice law in the Ninth Circuit for a period of six months.[2] Clearly, the ability of an AI-bot to manipulate or hallucinaate "evidence" in response to inquiries by the lawyer could skew the "facts" as presented in court beyond recognition. In such a situation, justice simply cannot be served.

In the search for the truth, particularly in cases in which a jury is empaneled to determine the "facts" from the evidence brought forward, the court has to adhere to certain guidelines that are derived from history, legislation, and the Constitution. In the process of applying the law to the facts, the judge likely will have consulted prior judicial decisions from similar cases as being useful, if not binding, in the current case. "It is a maxim among these lawyers, that whatever hath been done before, may legally be done again. . . ."[3] Put another way, stare decisis is the preferred basis for the judicial resolution of a dispute because it promotes the evenhanded, predictable, and consistent development of legal principles, fosters reliance on judicial decisions, and contributes to the actual and perceived integrity of the judicial process.[4] This is perhaps a somewhat overbroad definition of the legal principle known as stare decisis, but it will suffice as an introduction to the topic that is the second leg of the stool. . .legal precedent.

What is not generally considered is that there are actually several levels of stare decisis that are "in play" as part of the judicial process. In state or federal trial courts, the most commonly applied is "vertical stare decisis" that is what a trial court considers as part of the process. It characteristically involves research on the part of the judge or the lawyers into earlier cases from higher courts that are analogous to the case under consideration. The judge reviews the cases that are brought forward and determines whether they are persuasive or not on a given point.[5] Based upon that decision, the judge then formulates a judgment that, hopefully, will survive the action of an intermediate appellate court that would likely review the same cases. At the appellate level, the reviewing court would likely be confronted with "horizontal stare decisis" in which cases from appellate courts of equal dignity would be brought

[1]See Texas Rules of Civil Procedure, Rule 13 for the responsibilities of an attorney as to the authenticity of what is presented in court; see also Federal Rules of Civil Procedure, Rule 11 (b) as to the authenticity of presentations to the court.
[2]See LNU, Blance v. The case specifically cites to a lack of candor to the Court as a factor. 2026. 24-4790
[3]Jonathan Swift. Gulliver's Travels, Ignatius Press .2010. 279.
[4]Payne v. Tennessee, 501 U.S. 808, 827, 111 S.Ct. 2597, 115 L.Ed.2d 720 .1991.
[5]For a discussion of "vertical stare decisis", see Swilley v. McCain, 347 S.W.2d 871. 1964.

forward to persuade the reviewing court that it should follow its fellow jurists if it has not already considered the particular point. In such a situation, the appellate court would also be guided by decisions of the supreme court, if there should be any that are useful.

It is at the supreme court level where the principle of stare decisis encounters influences that can dilute it or even set it aside. This is because the supreme court has the authority to pronounce public policy in its opinions, which, in some instances, means overruling its own earlier cases, regardless of their antiquity, due to such considerations as social change. The case of Brown v Board of Education of Topeka in 1954 is a preeminent example.[6] In that case, the Supreme Court overruled the precedent set, and followed extensively, in 1896 in the case of Plessey v. Ferguson and did away with racial segregation in public schools.[7] At first glance, the potential for changes in the public policy pronounced by the supreme court, which would necessarily become "vertical stare decisis", may dilute the principal objectives of stability and predictability of the legal system. Even so, as stated by Roscoe Pound, "May we not conclude that in the part of the law which has to do immediately with conduct complete justice is not to be attained by the mechanical application of fixed rules"?[8] For example, in a recent series of changes to the rules governing appellate procedure, while the Texas Supreme Court indicated an awareness of the use of computer-generated pleadings, it did not specifically require the lawyers to certify to the authenticity of citations in the petitions.[9] By contrast, Rule 15 §4 of the Rules of the Supreme Court of the United States, require that the contents of petitions "present with accuracy, brevity, and clarity whatever is essential to ready and adequate understanding of the points."[10] It is at this admittedly imprecise intersection of technology and legal precedent that the third factor in the fact-finding process is introduced: the judge.

While it is presumed that the judge "knows all the law", this particular fiction is especially vulnerable to the impact of technology. In some respects, it does not matter what the subject of a given case is that is presented to the judge. What is of great importance is the preparation, both in terms of academic preparation and legal experience, of the person who presides over a case. In most instances, judges through the end of the 20th Century were products of law schools firmly grounded in the undergraduate liberal arts that prepared the student to enter the law. There was little room in the curriculum for such things as computer science or other tools that might be of value to a lawyer in research or evaluation of a case. The emphasis was on the "case method" in which the cases that defined certain areas of the law were collected and then discussed in classes focused on the "Socratic method" of interaction between professor and student. The result, almost by definition, is a legal frame of reference that is both narrowly focused and can be relatively rigid, almost mechanical, in its ability to engage in flexible thinking about a legal problem.

In order to provide the court system with an opportunity for some flexibility, the judge in the Anglo-American common law system sometimes can exercise discretion in the application of the law to the facts of a particular case as found by the jury. The obvious limit to that discretion comes from a combination of the wording of statutes or "vertical stare decisis". There is, however, another factor in the exercise of judicial discretion born of common sense. As stated by Ronald Dworkin, "Discretion, like the hole in a doughnut, does not exist except as an area left open by a surrounding belt of restriction. It is therefore a relative concept."[11] It is expressed judicially as "whether the court acted without reference to any guiding rules and principles. . .arbitrary and unreasonable"[12] If review of the decision of a lower court finds that it was arbitrary or unreasonable, the appellate court or supreme court has the ability to reverse the trial court's decision, which can have the effect of maintaining stare decisis at least in that case.

At the threshold, the question as to whether or not a given piece of evidence, whether the result of AI analysis or AI generation, will be allowed in court resolves itself into one word: authenticity. When such evidence is offered, the judge must exercise the "gatekeeper" function that is imposed by the Daubert standards.[13] The credentials of the expert, in this case the creator of the AI platform that is at the base of the evidence, and the methodology by which the results are reached must be established for it to be allowed. It is in this phase of the trial process that the testing of the evidence tends to reach a peak due to the right of cross-examination by the opposing counsel. Further, as a simple example, in a case that involves a computer-generated printout that is offered by a records custodian, on cross-examination, if it should be shown that the custodian has no personal knowledge of how it was created, then it will be disallowed. Similarly, the creator of an AI platform would need to show evidence of the objectivity and comprehensive analysis of the data that it produced to provide a sense of "authenticity". In March 2024, in the case of State of Washington v. Poluka, the trial court rejected cellphone recordings that had been AI-enhanced. The court held that this did not even meet the much lesser Frey standard and rejected the evidence.[14]

On the other hand, if it is assumed that the AI-generated evidence meets the authenticity standards of Daubert and is presented to a jury, there is the issue of how the jury sees the evidence. For example, if AI is used to re-create an automobile accident in order to illustrate how certain injuries occurred, the presentation would need to have various effects in order to be complete in the

[6]347 U.S. 483,74 S. Ct. 686; 98 L. Ed. 873.
[7]163 U.S. 537,16 S. Ct. 1138; 41 L. Ed. 256.
[8]Roscoe Pound, An Introduction to The Philosophy of Law (Yale University, 1922).at 140. 1922.
[9]See Texas Supreme Court MIsc. Docket Order 25-9104 (December, 23, 2025). 2025.
[10]Rules of the Supreme Court of the United States, 14 (effective March 16, 2026). 2026.
[11]Ronald Dworkin, Taking Right Seriously (Harvrd University Press 1977). 1977.
[12]Downer v. Aquamarine Operators, Inc, 701. S.W.2d 238, 241-2 (Tex. 1985). 1985.
[13]Daubert v. Merrell Dow Pharmaceuticals, Inc., 509 U.S. 579, 113 S. Ct. 2786, 125 L. Ed. 2d 469. 1993.
[14]Kevin J. Quilly, "Washington Court Rejects Novel Use of AI-enhanced Video", 16 The National Law Review #154 (June 4, 2024). See also Frye v. United States, 293 F. 1013 (D.C. Cir. 1923).

sensory stimulation of the jurors [e.g. relative size of vehicles, traffic conditions, sunlight or dark, speed, sight, and sound of vehicular impact and body movements of driver and passengers]. Each component of the presentation would need to be supported by testimony or physical evidence as a predicate to the AI-presentation itself. It must be admitted that such a presentation could well be of significant aid to a jury in clarifying the details of how such an event occurred, with the result of a higher quality of verdict than might otherwise happen. Of course, the decision of the judge as to whether to allow such a presentation would be subject to the Daubert review and the discretion of the court as to whether the AI adds anything to the case for the jury. As the judicial system presently operates, it is more likely than not that the AI evidence would be allowed on the theory that the jury will evaluate it anyway and give it the weight it deserves.

The problem, in the LLM-driven algorithm, is that cross-examination of the creator might disclose inherent biases or oversights that render the output suspect at best no matter the precision of the mathematics. For example, in a recent New York case, the expert witness testimony as to real estate values was found by the court to be "inherently unreliable. . .[and] based on speculation". It was determined that the expert had used AI [specifically Microsoft CoPilot] to check his calculations, but "could not recall what input or prompt he used to assist him." Indeed, he could not even explain how CoPilot works.[15] The impact on "authenticity" of a given precedent is unmistakable in the ability of the judge to exercise the "gatekeeper" function with any certainty. In addition, it would need to be shown that the output was not an "hallucination" created by the platform to show an answer, any answer, to the user. It is this potential that has created "real world" problems for lawyers who depend upon the AI-bot to draft papers presented to the court.[16] In the event that the AI-bot creates citations or quotes from fictitious cases that are used to persuade a judge to follow that "precedent", the result is a skewing, not to say "perversion", of the law that dilutes the concept of stare decisis at least at the "horizontal level". If not detected in the appellate process, this error potentially becomes "vertical stare decisis". In such a situation, the use of the AI-bot unknowingly can create "stare non decisis", a situation in which the machine has taken control of a portion of the justice system away from the humans who created it and who it is meant to serve. . .and neither the lawyers nor the courts may know that it happened. Over time, the cumulative effect of such errors would likely serve to skew the entire precedential picture and undermine confidence in precedent itself. The result could be a fundamental change in how the Anglo-American common law system operates and evolves.

As a practical matter, at the appellate or supreme court level of review as to the admission or exclusion of AI-generated evidence, there would likely be very little "second guessing" as to authenticity if the Daubert protocols were reflected in the record to have been satisfied. The reviewing courts would tend to yield to the jury findings unless there were evidence that the judge had simply "gone off the rails" one way or the other in the rulings. Only if post-judgment motions in the trial court based upon hitherto unknown evidence that the AI-generated evidence was the product of "hallucination" or deliberate manipulation sufficient to call into question whether the admission of the evidence denied a fair trial to one party or the other or both would likely persuade the reviewing court to order a new trial. Such a procedural situation could well prevent the creation of an inappropriate precedent for future litigation.

The current tendency of lawyers to employ AI-bots as more than a mere dataverse search tool must be carefully monitored by judges whose own preparation for the task is academically and experientially narrow and must be expanded. Put another way, "We have met the enemy, and he is us." This is not to suggest that the judiciary should become a community of technogeeks as that would be just as dehumanizing as would be a complete surrender to machine decision-making. It is to venture that they need to be alert that, in the case that is before them, as well as cases upon which they rely to make decisions, there may well be evidence that has been subject to manipulation by AI-bots that dilute or otherwise diminish the value of precedent. In such a situation "stare decisis" may well conceal "stare non decisis".

[15]Matter of Weber, 2024 NY Slip Op 24258 [85 Misc 3d 727 at 740-41] (Surrogate's Court, Saratoga County, 2024). 2024.
[16]See Mata v. Avaianca, Case No. 1:22-cv-01461 (US District Court for the Southern District of New York) June 2023, cited in Law 360, June 23, 2023.