Dr. Pramod Kumar, 2026, 14:4
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

Digital Forensics of Large Language Models: Developing a Legal and Forensic Framework for Investigating AI-Generated Evidence

[1]Dr. Pramod Kumar,
Associate Professor, Faculty of Law, SAM Global University, Raisen, Madhya Pradesh, India[1]

[2]Ms. Ambika Tomar,
Assistant Professor, Faculty of Law, SAM Global University, Raisen, Madhya Pradesh, India[2]

[3]Mr. Saood Iqbal khan,
Assistant Professor, Faculty of Law, SAM Global University, Raisen, Madhya Pradesh, India[3]

[4]Mr. Lakshya Bhardwaj,
Assistant Professor, Faculty of Sciences, SAM Global University, Raisen, Madhya Pradesh, India[4]

Abstract- **We explored in this paper the performance of standard Markov chain Monte Carlo (MCMC) algorithms as well as of adaptive versions of these algorithms on a set of representative posterior distributions that can be encountered in variational applications to realistic target distributions. Target distributions were generated by a total variation mixture of Gaussian distributions, and they comprised multimodal distributions like the shifted mixture distribution, the Hartman 6D, the LB 4D and the Rosenbrock 6D. In this simulation study, we primarily examined a few basic algorithms, adaptive random walk Metropolis (RWM), adaptive independent Metropolis (AIM), adaptive Gaussian proposal Metropolis (AGPM), and generalized adaptive Metropolis (GAM), on the target distributions created. We observed that although the RWM based MCMC algorithms were capable of capturing most of the probability mass of the multimodal target distributions and generating a good degree of mixing. This however had a significant negative impact on it due to the choice of covariance structure. Further analyses on the performance of AIM, AGPM and GAM indicated that the GAM could yield desired coverage compared to RWM based MCMC but they are vulnerable to decrease the efficiency of sampling. Furthermore, less dimensional targets were required to get a good performance with the AIM than with the low dimensional tests. Overall, our simulation. This paper presents the doctrinal and technical problems that come with using LLM-generated content as an offered, challenged, or investigated piece of evidence. It discusses the existing framework of evidentiary procedures in the US, India and the European Union, and existing incidents of AI hallucination and machine-generated disinformation to bring to the fore the practical implications of the problem. The paper proposes a modified IDAP framework for identification, preservation, acquisition, analysis, and presentation, drawing on the probabilistic, non-deterministic, and cloud-based nature of large language model outputs, in light of existing digital forensic methodologies. Moreover, it argues that there should be a legislative framework that forces the logging of provenance and ensures that records produced by AI are qualified and that cross-border competence requirements for admissibility are harmonised. If there is no attempt at harmonization between the science of Forensics and the evidentiary law, then either AI-generated evidence will be discounted or bogus evidence will be accepted without question. The framework provides a rationale and a scientifically informed guide for legislatures, forensic practitioners, and the courts in a principled approach to this new type of evidence.**

Keywords— Digital forensics; large language models; artificial intelligence evidence; chain of custody; authentication of electronic records; hallucination; evidentiary reliability; comparative evidence law.

I. INTRODUCTION

In recent years, with the rise in the use of powerful and highly capable language models by NLP researchers in the last decade, institutions and society are developing their own innovative tools, services, and even content. Attorneys draft documents assisted by generative tools; businesses use AI-powered summaries of business transactions; and more and more, we are using platforms that add synthetic text, image or audio between speakers during human interactions. If someone disagrees, what they actually said, wrote or represented gets confused with another question that old evidence law wasn't designed for: was it what they actually said, wrote or represented, generated, altered or embellished by a statistical model that predicts likely sequences of words, rather than one which records or perceives events?

When there is evidence of authenticity, photos, audio recordings of conversations, server logs, and other such material constitute an event in the world and are important for their significance. The output of a large language model on the other hand, is created by a random process learned from a vast amount of text which can generate fluent, confident, and entirely fabricated text that is formally indistinguishable from real text. It is know as the phenomenon of the hallucination where the model produces citations, facts or quotations that do not exist. When this type of content is accidentally introduced into a court proceeding by counsel's use of an AI research tool, or intentionally introduced by a party seeking to produce evidence, it exacts a high cost on existing rules of authentication and reliability.

The study argues that, whether as a suspect, a tool in the commission of a crime or as a source of exonerating or incriminating information, AI-generated material will require the formulation of a specific forensic methodology as well as a forensic architecture. It is continued with six sections. Part II looks at large language models' technical features that are directly relevant for how they should be treated in court. Part III compares the laws concerning electronic evidence in the United States, India and the EU. In Part IV documented cases and regulatory experience are used to identify practical hazards that have already materialised. There is a suggestion of a 4-phase forensic model in Part V based on existing digital forensics practice. Part VI has legal reform which will make that framework realistic. Without a coordinated response from forensic science and evidentiary law, the courts are exposed to two types of failures. Excluding reliable evidence created or aided by AI, on the one hand, risks the unreliable failure to admit such evidence. Excluding reliable evidence created or assisted by AI, on the one hand, risks the unreliable failure to admit such evidence. On the other hand, courts risk the credible failure to deny access, by uncritically admitting synthetic material that simulates the appearance of authenticity.

II. LARGE LANGUAGE MODELS AS SOURCES AND SUBJECTS OF EVIDENCE

It's a large language model that predicts what's next in a series of tokens, and it's a statistical system trained on a very large text corpus. In contrast, a database query or a deterministic algorithm will return the same answer if you repeat the query or set the same parameters, but if a temperature or sampling parameter is used to generate the answer in a manner that is not completely deterministic, then the same prompt given to the same model twice can result in different answers. There are direct forensic consequences of this non-determinism. One of the fundamental principles of traditional digital forensic is that given a known input, using known software and hardware, will result in a reproducible and verifiable output. The researchers claim that for an existing model, at any point in time, the behaviour of the model can be assumed to be invariant to change. It cannot be the case for generative model, however, because generative model behavior is dependent on the model weights, which are updated periodically, and system prompts, which may not be communicated to end users, and also on stochastic sampling, which cannot be exactly reproduced.
The multi-layered architecture that most of the users of LLM are dealing with adds another layer of complexity. The end user provides a prompt via a web or an application interface, and this prompt is

sent to a remote server provided by the model provider. The prompt is passed to infrastructure which can be in several jurisdictions and a response is returned and served at the local level. The logs of client-side applications, the records of network transmissions, the logs of inferences on the server side and usage records on the part of the provider are each an artefact present from any one of these stages and have some degree of forensic potential. However, the control over various artefacts surrounding auto-generated data is not even – between the user, the platform operator, and the provider of the underlying model. All these actors might not have the data for long enough to meet a month or even years later investigator. This is in contrast to the personal computer or cellphone paradigm, in which one device designed to hold the record and a medium that is typically small has been the backbone of the chain-of-custody.

According to the exchange principle attributed to social scientist Edmond Locard in the field of forensics, forensic scientists believe that if two objects come in contact with each other, there will be mutual transfer of trace material. In the digital setting, this principle has generally been applied as the assumption that an action will leave a trail that can be found in a file system, registry or logs. This is disrupted by the outputs of large language models in two ways. First, a trace of the hallucinated fact exists in another coherent text, and is well-integrated, making the fabrication not stand out as an anomaly given the known forensic triage of a text. Moreover, a session log, which would be a discrete forensic artefact in much the same way as the model itself is a discrete artefact, and which would capture the generative process, the prompt and completion, is not retained as part of the model and thus is not known to AI models themselves.

Considering those qualities, it is appropriate to consider such content as a separate class of evidence rather than a subset of electronically stored information. Now the investigator was asking a more sophisticated question. No longer just if this toy is a toy. The artefact, of course, is genuine, and one question addressed by the orthodox doctrine of authentication is whether it is. The investigator then inquired whether it was generated (at any part of its generation) in any way by a system that does not distinguish between true and false outputs during the generation process itself.

III. A PARALLEL STUDY OF ARTIFICIAL INTELLIGENCE AND ELECTRONIC EVIDENCE

The Federal Rules of Evidence are the primary rules which control the admissibility of electronic evidence in the United States. Under the Rules, the person introducing evidence (called the proponent) must present sufficient evidence to allow for a ruling that the evidence is what the proponent alleges it is. The amendments took effect on December 2017, and introduced two new rules, Rule 902(13) and Rule 902(14). The former involves a self-authentication provision for records produced by a particular electronic process or system. The latter replicates documents that are created from a certain electronic device, storage medium or file. These provisions will be in effect while the proponent provides a certificate from a qualified individual that satisfies the requirements for a certified business record. These were designed to reduce the cost and delay of having to send for the production of a document simply to establish that it was a true copy and to enable the opponent to question whether that is so, and ask for the grounds for the accuracy claimed.

The standard set forth by the United States Supreme Court in Daubert v. Merrell Dow Pharmaceuticals, Inc., establishes that trial courts are gatekeepers who are required to evaluate the reliability of scientific or technical methodology, including forensic methodology, by considering whether the expert's opinion is based on a theory or technique that has been tested, subjected to peer review evaluation, has a known or knowable error rate, and is generally accepted in the scientific community of the subject matter.

The Supreme Court in Daubert is not worried about the self-authentication of records, but rather about expert testimony. However, its gate keeping logic is suggestive for AI-generated evidence. The reliability threshold is a requirement for any forensic method

that is developed to detect, attribute, and/or authenticate LLM output. It can be considered by a court only then. The recent case of Mata v. Avianca, Inc. in the United States District Court for the Southern District of New York demonstrates the ramifications of failing to exercise this standard gateway role. In this instance, the brief included citations to judicial rulings that had never actually existed, and were entirely the creation of the generative AI. The court imposed monetary sanctions pursuant to Fed. R. Civil Proc. 11 and his own inherent powers. It also directed the offending attorneys to reach out to every judge they had misled in their fake opinion about the judge's name. Since the case has been referenced a number of times, an illustration of the potential for generative AI output to be confused for or replaced by an actual legal authority has been provided.

Indian evidentiary practice has a very developed jurisprudence in the context of electronic records. It can be viewed from the introduction of sections 65A and 65B in the Indian Evidence Act of 1872 by the Information Technology Act of 2000. Any information contained in an electronic record (which includes any information in digital form, such as an e-mail message, a text file, an image, etc., that is produced by a computer) shall be admissible as a document without requiring proof of the original document, if the electronic record is the regular output of a computer operated by a party to the proceedings and further if the party relying upon the electronic record produces the certificate under 65B(4) identifying the electronic record, describing the manner in which it was produced and applicable particulars pertaining to the device, and the certificate is signed by a person in a responsible position in relation to the use of the computer operating the device.

The Supreme Court of India, 2014 Case Study - Anvar P. V. v. V.P.K. The Supreme Court has held in Basheer case that Sections 65A and 65B are a complete code with regard to the admissibility of electronic record. If the secondary electronic evidence is a compact disc and it does not include a Section 65B(4) certificate, then it will not be admissible. This overturned a prior decision, which was more lenient, that the electronic record could be presented through oral evidence alone. A subsequent division bench in Shafhi Mohammad v. State of Himachal Pradesh held that the requirement of issuing certificate was not mandatory when the party seeking records was not in control of the device generating the records. This led to a period of doctrinal confusion as to the admissibility of electronic evidence which was cleared in Arjun Panditrao Khotkar v. Kailash Kushanrao Gorantyal (2020) by a three-judge bench. Anvar P.V. This is because the decision made in Shafhi Mohammad was correct, and the other remarks in the decision were also incorrect with the qualification that if the original electronic device is produced before the court, the root identifier or operator identifies it, there is no need for a certificate. In more than a decade, this trilogy of cases has highlighted the challenges facing the legal system in recognizing a "static" electronic record. The challenge is growing due to the introduction of generative, non-deterministic AI output. The Section 65B(4) certificate (certificate to the regular and accurate functioning of a computer system) issuer is not easily able to do so. They are not able to endorse a model that produces likely text, rather than actual text or a reproduction of it.

The EU policymakers decided to regulate AI providers and deployers to tackle the issue of AI liability. Content produced using AI or deep fakes would have to be marked as AI-generated or manipulated. But this is subject to certain exceptions. The Regulation stipulates documentation, logging and recording obligations for all providers of high-risk AI systems. These requirements are primarily directed towards compliance with the Regulation and monitoring by market competitiveness and do not seek to generate litigation, but rather a source of retained technical documentation. Sometimes, such technical retention documentation can be relevant when a high risk system is involved in disputes in a forensic investigation. These duties are on top of the existing data protection regime established by the General Data Protection Regulation, Regulation (EU) 2016/679, which regulates the lawfulness of processing personal data in the process of training, fine-tuning and operation of language models. It

further imposes its own accountability and record-keeping obligations which may overlap with a forensic investigation on the origin of AI-generated outputs.

Table 1: Comparative Framework for Authentication of Electronic and AI-Generated Evidence

Jurisdiction	Governing Provisions	Core Evidentiary Requirement
United States	Federal Rules of Evidence 901, 902(13)-(14); Daubert v. Merrell Dow Pharmaceuticals, Inc.; Federal Rule of Civil Procedure 11	Proponent must show the item is what it is claimed to be; certification by a qualified person may substitute for live testimony; expert reliability is subject to judicial gatekeeping
India	Indian Evidence Act, 1872, ss. 65A-65B; Anvar P.V. v. P.K. Basheer; Arjun Panditrao Khotkar v. Kailash Kushanrao Gorantyal	A signed certificate under Section 65B(4) is a condition precedent to admissibility of secondary electronic records, unless the original device is produced in court
European Union	Regulation (EU) 2024/1689 (Artificial Intelligence Act); Regulation (EU) 2016/679 (GDPR)	Provider-level transparency, disclosure, documentation and logging obligations; admissibility remains governed by national procedural law, informed by regulatory records

Source: Compiled by the author from the Federal Rules of Evidence, the Indian Evidence Act 1872, the cited case law, and Regulation (EU) 2024/1689 and Regulation (EU) 2016/679.

IV. THE DIFFICULTIES EMERGING: LESSONS FROM PRACTICE.

The case of Mata v. Avianca is not only about the wrongdoings of a lawyer, but also the journey the content went through as it went through multiple levels of professionals and courts. The attorney asked the same tool whether the cases were real, reciting the title to each case, the number and internal citations, and leaving the attorney with the impression that that was the case, rather than telling him the cases in question were fake.

It is apparent that the episode has a structural weakness, which the forensic framework needs to directly confront. Verification done with the same generative system that generated the challenged content is not independent verification as the generative system lacks ground truth to point out its own output. Thus, any forensic protocol for AI-generated evidence should involve external verification of the evidence, not just from the AI, but from outside the AI process. For instance, case law databases, business records for the time, or human review.

AI-generated or manipulated audio, image and video content is not just text. This is a development that poses a threat to an important default presumption which is this: that, when a recording is presented as being whole, it may be treated as a true account of what the recording shows. Where a litigant can plausibly argue that an inconvenient recording is itself AI-generated, courts are confronted with what is elsewhere in the literature referred to as a liar's dividend: good evidence can be called into question where something is AI-generated. The pressure is thus to use forensic methods that will detect statistical anomalies of AI generation, including: unexpected compression patterns, biometric implausibility, or the absence of expected sensor noise. It also means that this detection needs to be reliable to a certain degree. And this threshold is similar to that required of any other forensic analysis before a court will accept what it concludes.

Another practical challenge relates to the conservation of artefacts that can be used to recreate the way an AI-generated artefact has been created. For most inferences, they are done on the infrastructure of the model provider, not the end-user. In this respect, the investigator seeking to determine how a specific output was generated is dependent upon retention and disclosure practices and co-operation not typically seen in forensics using a device. Where a provider can only retain logs for a short time, or is able to retain logs, but only through a formal legal process that targets an entity outside the jurisdiction of the investigator, the practical time limit for forensic reconstruction may be reached long before the dispute becomes litigious. This reality provides strong support for the choice of preservation (as opposed to "reconstruction" after the event) for future generations that is discussed in Part V.

V. TOWARDS A FORENSIC FRAMEWORK FOR EVIDENCE GENERATED BY LLMS

NIST (National Institute of Standards and Technology) Special Publication 800-86 and International Standard ISO/IEC 27037:2012 guidance on the identification, collection, acquisition and preservation of digital evidence, outlines the methodology of digital forensic investigation in sequential phases that will maximise the integrity and defensibility of the digital evidence recovered. The framework still works well for organising principles but each stage needs to be adapted to the unique features of generative AI output mentioned above.

The investigator needs to identify whether an artefact has been created, synthesized or materially altered by an AI system, and which system, which version, and in what deployment context? Questions are applicable in scoping which do not have an equivalent in classic device forensic. For instance, if it was created using a consumer facing chat interface or embedded application feature, or if it was created using an automated pipeline that called a model via an application programming interface. Whether the provider stores any server-side information of the transaction.

Where there is a risk of dispute, it is recommended that the parties and their advisors keep the full history of prompts and responses, including the system prompts, model identifiers, times and any configuration parameters. If possible, cryptographic hash values and/or provenance metadata added by the platform, e.g., content credentials that follow the Coalition for Content Provenance and Authenticity specification, which some platforms are now implementing for AI-generated content, such as images.

Two simultaneous levels of assessment must be advanced. The first one is in the traditional sense of authenticity, that is, is it the artefact that it says it is, and was it not altered since being acquired? This is checked with hash comparison and metadata review. The second, something unique to AI-generated evidence, is the question of what is the evidence's provenance and reliability. Was the content created as a result of a generative process? Has been checked by someone else against outside material? In addition, can the information, a quotation or a citation that is claimed in it be substantiated by information from a source that is outside the generative system? If this can be done, then analysts should note the specific version of a model used; as model weights are updated periodically and production from one version of a nominal product will not be reproduced with another version.

The final phase is for the investigator to present his results to provide a court with a comprehensible and cross-examinable report. It should also give an honest opinion about how far it is possible to go with attribution when the server logs are not available, or the non-deterministic nature of the system makes it impossible to reproduce the disputed output. Judges, according to Daubert, are the gatekeepers of admissibility. Like the Federal Rules of Evidence norm 902(13)-(14) and the Indian Evidence Act section 65B(4), a certification attached to the evidence generated by artificial intelligence must include the information of the person making the certification and the description of the method it was acquired from must be enough to allow an independent test. It should not include any limitation of the methodology; the opposing party should be

able to present their counter-expert testimony if qualified, in some way.

Table 2: Proposed Four-Phase Forensic Model for LLM-Generated Evidence

Phase	Traditional Digital Forensics Focus	Adaptation for LLM-Generated Evidence
Identification	Locate devices, media, and files relevant to the investigation	Determine whether an AI system generated or altered the content, and identify the model, version, and access channel involved
Preservation and Acquisition	Create forensic images and maintain an unbroken chain of custody for physical or storage media	Preserve prompts, responses, system instructions, timestamps, and configuration parameters; obtain provider-side logs and provenance metadata through cooperation or legal process
Analysis	Examine acquired data for relevant artefacts using validated tools	Verify conventional authenticity through hashing and metadata review; separately assess factual reliability by corroborating content against sources external to the generative system
Presentation	Report findings in a form suitable for court, supported by validated methodology	Certify the acquisition method, disclose the non-deterministic and version-dependent limits of attribution, and distinguish authenticity from factual reliability in the final report

Source: Author's adaptation of the four-phase digital forensic model reflected in NIST Special Publication 800-86 and ISO/IEC 27037:2012, informed by the evidentiary requirements discussed in Part III.

VI. RECOMMENDATIONS FOR A LEGAL FRAMEWORK

1. Amend Controlled Authentication Rules to Statutes.

The provisions of self-authentication such as the Federal Rules of Evidence Rule 902(13)-(14) and Section 65B(4) of the Indian Evidence Act are based on the premise that the qualified person making the certificate is capable of establishing the normal and correct operation of the deterministic system. There should be a new certification classification for generative AI products within the legislature. This would necessitate that the person certifying the output, explain what model and version they were using, whether they have retained logs or not, and whether they have available independent verification. The provision would not relax current authentication standards, but rather would give courts a structured way to assess evidence that "genetically differs significantly" from a standard business record.

2. Compliance with tracking and logging regulations

The EU's AI Act requirements for disclosure of AI-generated or modified content, and documentation requirements for providers of high-risk systems, could be used as a model for other jurisdictions seeking to require documentation for AI systems. Requiring prompt and response logs to be retained for at least the specified period, and providing a statutory mechanism for them to be disclosed on showing that they are relevant in litigation would significantly improve the likelihood of forensic reconstruction and decrease the need to rely on one party's claim of the other party simply saying so.

3. Experts and Courts may certify.

It would be incorrect to presume that any such "technique" for detecting or attributing AI-generated content will prove reliable in a court proceeding; the court must follow Daubert or a similar reliability standard when deciding whether to allow a technique to be introduced into evidence. Accreditation bodies and professional associations should create validated protocols and proficiency testing for AI-content forensics which are similar to

the ones in place for traditional Digital Forensics examination to ensure that courts are not forced to adjudicate new methods without any external tool to assess their reliability.

4. Global Standardisation.

There is a possibility that competing national rules on authentication/retention could result in "forum shopping" outcomes when disputes are handled in "functionally similar" situations with the providers and/or users of the model. In the past, requests for AI system logs and provenance records have been difficult to process because they were not accessible in a timely fashion via international instruments or mutual legal assistance mechanisms. The delays have marred the preservation of volatile electronic evidence.

Table 3: Taxonomy of Forensic Artefacts Relevant to LLM Investigations

Artefact Category	Description	Typical Custodian
Prompt and completion records	The text or multimodal input submitted by the user and the corresponding model output, together with timestamps	End user device or application; occasionally the platform operator
System and configuration parameters	System prompts, temperature and sampling settings, and model version identifiers governing generation behaviour	Platform operator or enterprise deployer of the AI system
Server-side inference logs	Provider-maintained records of API calls, session identifiers, and infrastructure metadata associated with a request	Model provider
Provenance and content credentials	Cryptographically signed metadata indicating that content was generated or edited by an AI system, where implemented	Model provider or content publishing platform
Network and application logs	Records of data transmission between client and server, including timing and volume of requests	Internet service provider, enterprise network administrator, or platform operator

Source: Compiled by the author with reference to NIST Special Publication 800-86 and the transparency and documentation obligations of Regulation (EU) 2024/1689.

VII. CONCLUSION

Large language models have arrived in the everyday professional and personal setting, creating a type of evidence that feels hard to place in either of the previous two categories: the static "digital" evidence that has been established by law and practice, or the expert 'evidence' that is settled. The recommendations proposed in this article for the establishment of a new certification class for generative content, the requirement for AI providers to log and provide evidence of the origin of the content created, judicial oversight of detection methods, and stronger international collaboration to preserve and disclose logs of content creation across borders are less revolutionary than radical. The Indian Evidence Act, European Union rules and Federal Rules of Evidence already have doctrinal bases on which they rest and not offer to replace such legal framework. What they need is recognition in forensic practice and evidentiary law that a system that is capable of producing plausible text is not the same type of witness as one that is capable of producing a record of the events in question, and that law of evidence must be adjusted so as to maintain its function of distinguishing trustworthy evidence from persuasive fabrication. This adaptation isn't confined to the courtroom. “The First Law of the Generative A.I. Revolution” will be published in the Medical Law Review 2024 and will be written from a ‘think-tank' perspective covering practical regulation requirements, as well as foundational thinking around product, research and medical applications. – Read more here. The era of AI fabrication, which is relatively new, is likely to present a better opportunity for principled and

technically sound adjudication than will the next media-wacked frenzy of AI fabrication.

Declarations

Ethics Approval and Consent to Participate

Not applicable. This literature does not involve any new studies with human or animal participants conducted by the authors.

Consent for Publication

Not applicable. The manuscript does not contain any person's data in any form.

Competing Interests

The authors declare that they have no competing interests.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Authors' Contributions

All authors contributed equally to the conceptualization, literature review, analysis, and writing of the manuscript. All authors read and approved the final version.

Acknowledgements

The authors thank the institutional support provided by SAM Global University, M.P., India, for facilitating this research.

REFERENCES

1. Statutes and Regulatory Instruments.

- Federal Rule of Civil Procedure 11.
- Federal Rules of Evidence, Rules 901 and 902(13)–(14).
- Indian Evidence Act, 1872, ss. 65A–65B (as inserted by the Information Technology Act, 2000).
- Information Technology Act, 2000 (India).
- Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 (General Data Protection Regulation), OJ L 119/1.
- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on Artificial Intelligence (Artificial Intelligence Act), OJ L 1689.

2. Case Law

India

- Anvar P.V. v. P.K. Basheer, (2014) 10 SCC 473.
- Arjun Panditrao Khotkar v. Kailash Kushanrao Gorantyal, (2020) 7 SCC 1.
- Shafhi Mohammad v. State of Himachal Pradesh, (2018) 2 SCC 801.
- State (NCT of Delhi) v. Navjot Sandhu, (2005) 11 SCC 600.

United States

- Daubert v. Merrell Dow Pharmaceuticals, Inc., 509 U.S. 579 (1993).
- Mata v. Avianca, Inc., 678 F. Supp. 3d 443 (S.D.N.Y. 2023).

Books

- Casey, Eoghan, Digital Evidence and Computer Crime: Forensic Science, Computers, and the Internet (4th edn, Academic Press 2011).
- Carrier, Brian, File System Forensic Analysis (Addison-Wesley Professional 2005).
- Mason, Stephen and Daniel Seng (eds), Electronic Evidence and Electronic Signatures (6th edn, Institute of Advanced Legal Studies 2021).
- Nelson, Bill, Amelia Phillips and Christopher Steuart, Guide to Computer Forensics and Investigations (7th edn, Cengage Learning 2023).
- Rogers, Marcus and Matt Bishop (eds), The Science of Digital Forensics (Addison-Wesley Professional 2006).

3. Journal Articles

- Biasiotti, Maria Angela et al., 'Legal Issues in Digital Forensics' (2009) 25(4) Computer Law & Security Review 315–322.
- Carrier, Brian and Eugene H. Spafford, 'Getting Physical with the Digital Investigation Process' (2003) 2(2) International Journal of Digital Evidence 1–20.

- Casey, Eoghan, 'Digital Evidence and the Crime Scene' (2004) 1(Suppl. 1) Digital Investigation S1–S6.
- Casey, Eoghan, 'What Does "Forensically Sound" Really Mean?' (2007) 4(2) Digital Investigation 49–50.
- Garfinkel, Simson L., 'Digital Forensics Research: The Next Ten Years' (2010) 7(Suppl. 1) Digital Investigation S64–S73.
- Horsman, Graeme, 'Formalising Investigative Processes within Digital Forensics' (2019) 30 Digital Investigation 10–16.
- Horsman, Graeme, 'Tool Testing and Reliability in Digital Forensics' (2020) 33 Forensic Science International: Digital Investigation 301021.
- Karie, Nickson M., Hossein S. Venter and Karim A. Elloumi, 'The State of the Art of Digital Forensic Investigation' (2019) 4(2) Forensic Sciences Research 63–79.
- Pollitt, Mark M., 'A History of Digital Forensics' (2010) 7(Suppl. 1) Digital Investigation S3–S10.
- Quick, Darren and Kim-Kwang Raymond Choo, 'Digital Forensic Intelligence: Data Subsets and Open Source Intelligence (DFINT+OSINT): A Timely and Cohesive Mix' (2015) 50 Future Generation Computer Systems 558–566.