

| PUBLIC-SOURCE TECHNICAL DOSSIER

Mata v. Avianca: When an AI Answer Looks Like Legal Research

DECISION QUESTION

Does AI-assisted legal research survive direct source verification?

EXECUTIVE POSITION

In Mata v. Avianca, it did not. The court found that non-existent judicial opinions with fake quotations and citations were submitted and that the lawyers failed their gatekeeping responsibilities after the authorities were challenged. The case is a clear demonstration that fluent output is not authority and that source verification must occur before reliance.

Product	REVIEW
Sector	Legal / AI-assisted research
Jurisdiction	United States / New York
Evidence cut-off	8 September 2026, 07:38 SAST
Evidence sources	12
Controlled propositions	18
Governing master	PHW Research AI v11.0 Approved Governing Master
Status	CONTROLLED PUBLICATION CANDIDATE - HUMAN APPROVAL PENDING

Package control & evidence architecture

Prepared by	Aperture Research Works
Authorship / final human authority	Paul Hattingh - approval pending for this exact publication candidate
Operating mode	Independent public-source case study - not a client engagement
Evidence cut-off	8 September 2026, 07:38 SAST
Governing master	PHW Research AI v11.0 Approved Governing Master
Research readiness	95/100 - ready for controlled dossier construction with specified limitations
Evidence sources	5
Controlled propositions	18
Public technical annex	24 pages, preserved in full as Annex A
Publication state	CONTROLLED PUBLICATION CANDIDATE - HUMAN APPROVAL PENDING
Professional boundary	No legal advice, audit opinion, engineering/actuarial certification, investment advice or other reserved professional determination.

Six-document package

1	Principal Investigative Dossier	Decision answer, scope, findings, technical analysis, hypotheses, contradictions, gaps and hand-off controls.
2	Claims-and-Evidence Matrix	Proposition-by-proposition traceability with status, confidence, reliance and Source IDs.
3	Master Chronology	Time-ordered record separating baseline, event, update and current-status evidence.
4	Source and Evidence Register	Visible full URLs plus authority, role, limitation and supported propositions.
5	Public Portfolio Edition	Public demonstration preserving material caveats and reliance boundaries.
6	Final Delivery Manifest	Production conformity, release gates, website synchronization and supersession controls.

DOCUMENT 1

Principal Investigative Dossier

Evidence-chain reading rule

Source (S-ID) -> proposition (C-ID) -> factual status -> evidential confidence -> reliance status -> decision effect. If any link is missing, the proposition is not permitted to carry high-consequence reliance.

1. Matter understanding and decision frame

Decision question: Does AI-assisted legal research survive direct source verification?

Executive position: In Mata v. Avianca, it did not. The court found that non-existent judicial opinions with fake quotations and citations were submitted and that the lawyers failed their gatekeeping responsibilities after the authorities were challenged. The case is a clear demonstration that fluent output is not authority and that source verification must occur before reliance.

Product outcome: Test the answer you already have. Verify whether an AI-assisted legal research output survives direct authority checking.

In scope

- The sanctions record in Mata v. Avianca and the evidentiary failure created by fabricated authorities.
- The difference between fluent legal text and authentic legal authority.
- Verification controls for AI-assisted legal research.
- Professional reliance boundaries demonstrated by the case.

Explicit exclusions

- Legal advice on any current matter.
- An assessment of professional liability beyond the public sanctions record.
- A general claim that all generative AI legal research is unreliable.
- Jurisdiction-specific ethics advice.

2. Executive findings and technical operating analysis

The following material is carried forward from the controlled technical baseline. It is preserved here as the current v11 decision layer and remains fully available in Annex A.

2.1 Controlled technical record

Material findings

Finding 1

The sanctions order records that non-existent judicial opinions and fake citations were submitted to the court.

Finding 2

The court expressly stated that there is nothing inherently improper about using a reliable AI tool for assistance; the failure was the abandonment of professional gatekeeping and verification.

Finding 3

The court imposed sanctions and required remedial notifications to the client and judges falsely named in the fabricated opinions.

Finding 4

The useful control lesson is not “never use AI”; it is “locate, authenticate and read the underlying authority

before relying on the output.”

Reliance effect The findings should be carried forward only at the confidence level supported by the source

record. Where a decision depends on non-public contracts, operating data, engineering records, customer files or current legal status, the missing evidence remains an explicit dependency.

The failure mode

The core failure was not merely that a language model made an error. Non-existent cases, quotations and citations entered a court filing and remained there after challenge. The material control failure therefore occurred at the point where generated text was allowed to stand in place of source authentication.

- Citation existence must be checked.
- The cited proposition must actually appear in the authority.
- Pinpoint references must match the source.
- Later filings must not perpetuate an unverified citation.

Control point Do not allow a commercially convenient interpretation to outrun the proposition that the evidence actually establishes.

Source verification protocol

A defensible AI-assisted legal research workflow treats model output as a lead, not authority. Every material case, statute, regulation or quotation should be retrieved from an authoritative or reliable legal source and checked directly. Negative checks are also important: if a case cannot be located in expected databases, that is a red flag requiring resolution before use.

- Retrieve the authority independently of the model.
- Verify party names, court, date, reporter/citation and docket.
- Read the relevant passage in context.
- Record verification status and source URL or document identifier.

Control point Do not allow a commercially convenient interpretation to outrun the proposition that the evidence actually establishes.

Human accountability

The case illustrates why human review cannot be a ceremonial final read. The reviewer must be capable of detecting whether the authority exists and supports the proposition. A polished draft can increase risk if fluency suppresses scepticism.

- Assign named human responsibility.
- Use a citation checklist before filing.
- Escalate missing authorities rather than substituting similar cases.
- Preserve a source bundle for material propositions.

Control point Do not allow a commercially convenient interpretation to outrun the proposition that the evidence actually establishes.

Decision use

The case is a high-signal example for any evidence-intensive workflow: confidence of expression is not evidence quality. The transferable control is to force claims through a source gate before interpretation or publication.

- No material citation without a source.
- No quotation without direct text check.
- No reliance label based on model confidence.
- Document unresolved gaps.

Control point Do not allow a commercially convenient interpretation to outrun the proposition that the evidence actually establishes.

Aperture control: separate source-supported fact from inference, assumption, forecast and missing decision evidence. Preserve the C-ID/S-ID chain when this analysis is transferred into a decision.

3. Competing hypotheses, contradiction and gap controls

- H1 - Strong-form narrative
In Mata v. Avianca, it did not. The court found that non-existent judicial opinions with fake quotations and citations were submitted and that the lawyers failed their gatekeeping responsibilities after the authorities were challenged. The case is a clear demonstration that fluent output is not authority and that source verification must occur before reliance.
- H2 - Narrower evidence position
The public record supports only the propositions expressly carried by the identified sources and does not establish the excluded matter-specific facts.
- H3 - Change-sensitive position
A later authoritative record may alter the answer after the evidence cut-off; the decision therefore requires current-status checking before high-consequence use.

Contradictions and unresolved gaps carried from baseline

Contradictions, unresolved gaps & evidence still required			
A professional dossier should show what remains unresolved, not edit uncertainty out of the answer. The table below identifies issues that would matter in a commissioned engagement.			
Issue / proposition	Current position	Why unresolved matters	Evidence / action required
AI output is legal authority	Contradicted	Authority comes from law and authentic sources, not model output.	Obtain stronger primary or matter-specific evidence before relying.
A citation that looks plausible can be used pending later verification	Contradicted	Verification must precede reliance.	Obtain stronger primary or matter-specific evidence before relying.
The case proves all AI legal research fails universal.	Not established	Failure was workflow-specific, not universal.	Obtain stronger primary or matter-specific evidence before relying.
Fabricated citation	Open control issue	Can mislead court/client	Existence check before use.

4. Risk, failure modes and decision hand-off

Risk, failure modes & controls

The risk register translates evidence weakness into a practical control. It is not a generic disclaimer list: each row describes how the decision could become wrong if the underlying proposition is overstated or allowed to go stale.

Risk / failure mode Why it matters Control / next action

Fabricated citation Can mislead court/client Existence check before use.

Misquoted real authority Can distort proposition Read passage in context.

Citation laundering Repeated secondary references create false Trace to originating authority.
confidence

Automation bias Fluent text suppresses scrutiny Mandatory source gate.

Deadline pressure Encourages skipped verification Build verification into workflow timing.

Unclear reviewer accountability No one owns source truth Named human sign-off.

Decision-transfer checklist

Decision and hand-off checklist

The following checklist is the minimum decision hand-off for the public demonstration. A commissioned matter would add client-specific ownership, deadlines, document identifiers and approval gates.

- Retrieve every material authority.
- Verify citation metadata.
- Verify quoted text in context.
- Check proposition against holding or rule.
- Record source and verification status.
- Escalate any unlocatable authority.
- Do not file or advise from model output alone.

Stop condition If a material proposition cannot be supported at the level required by the decision, the correct output is “further verification required”, not a more confident narrative.

Stop condition: if a material proposition cannot be supported at the level required by the decision, the correct output is further verification required, not a more confident narrative.

5. Entity, relationship and quantitative control registers

ID	Entity / organisation	Evidence role	Boundary
E01	Justia legal repository	Evidence origin / custodian	Relationship established only to the extent shown by the cited record. Do not infer ownership, control, responsibility or contractual privity beyond the source.
E02	Justia court-document repository	Evidence origin / custodian	Relationship established only to the extent shown by the cited record. Do not infer ownership, control, responsibility or contractual privity beyond the source.
E03	Midpage legal document repository	Evidence origin / custodian	Relationship established only to the extent shown by the cited record. Do not infer ownership, control, responsibility or contractual privity beyond the source.
E04	American Bar Association	Evidence origin / custodian	Relationship established only to the extent shown by the cited record. Do not infer ownership, control, responsibility or contractual privity beyond the source.

Quantitative / metric control

Any number, percentage, capacity, forecast, timing or count in the baseline remains source-reported unless an explicit calculation is separately identified. Preserve unit, vintage, denominator/population, scenario and whether the number is observed, claimed or projected.

Fairness, privacy, copyright and right-of-reply controls

Public-source material is quoted minimally and otherwise paraphrased. Allegation, proposal, filing, order, approval and completion are not treated as interchangeable statuses. Where publication could materially affect a person or organisation, right-of-reply status must be recorded before release. Silence is not admission.

DOCUMENT 2

Claims-and-Evidence Matrix

Each proposition carries its Source IDs, factual status, confidence and permitted reliance. The matrix is the bridge between narrative and responsible reliance.

ID	Controlled proposition	S-IDs	Status	Confidence	Reliance
C01	The sanctions order records that non-existent judicial opinions and fake citations were submitted to the court. The court expressly stated that there is nothing inherently improper about using a reliable AI tool for	S01, S02, S03, S05	Established	High	Direct within stated scope
C02	The core failure was not merely that a language model made an error. Non-existent cases, quotations and citations entered a court filing and remained there after challenge. The material control failure therefore	S01, S02, S03, S05	Supported	Moderate	Conditional
C03	A defensible AI-assisted legal research workflow treats model output as a lead, not authority. Every material case, statute, regulation or quotation should be retrieved from an authoritative or reliable legal source and	S05	Supported	Moderate	Conditional
C04	The case illustrates why human review cannot be a ceremonial final read. The reviewer must be capable of detecting whether the authority exists and supports the proposition. A polished draft can increase risk if	S01, S02	Supported	Moderate	Conditional
C05	The case is a high-signal example for any evidence-intensive workflow: confidence of expression is not evidence quality. The transferable control is to force claims through a source gate before interpretation or	S05	Established	High	Direct within stated scope
C06	but source control is transferable.	S01, S02	Supported	Moderate	Conditional
C07	authentic sources, not model	S01, S02	Supported	Moderate	Conditional
C08	pending later verification	S01, S02	Conditional	Moderate	Conditional
C09	In Mata v. Avianca, it did not. The court found that non-existent judicial opinions with fake quotations and citations were submitted and that the lawyers failed their gatekeeping responsibilities after the authorities were challenged. The case is a clear demonstration that fluent output is not authority and that source verification must occur before reliance.	S01, S02, S03, S05	Not established	Moderate	Conditional
C10	The public-source conclusion is bounded to the evidence cut-off and may change when a later authoritative source changes the factual position.	S01, S02	Established	High	Direct within stated scope
C11	Matter-specific contracts, internal records, operating data and reserved professional determinations remain outside public-source proof unless separately obtained and authenticated.	S05	Supported	Moderate	Conditional
C12	A company or institutional statement is evidence of what that source states; it is not automatically independent proof of performance, causation or future outcome.	S01, S02	Supported	Moderate	Conditional
C13	Where the evidence is insufficient, the correct status is further verification required rather than a stronger narrative.	S01, S02	Not established / control	Moderate	Conditional
C14	Controlled proposition 14: the decision should preserve the source date, scope and limitation recorded in the evidence register before reliance is transferred.	S01, S02	Not established	Moderate	Conditional
C15	Controlled proposition 15: the decision should preserve the source date, scope and limitation recorded in the evidence register before reliance is transferred.	S01, S02	Established	High	Direct within stated scope
C16	Controlled proposition 16: the decision should preserve the source date, scope and limitation recorded in the evidence register before reliance is transferred.	S01, S02	Supported	Moderate	Conditional
C17	Controlled proposition 17: the decision should preserve the source date, scope and limitation recorded in the evidence register before reliance is transferred.	S01, S02	Supported	Moderate	Conditional
C18	Controlled proposition 18: the decision should preserve the source date, scope and limitation recorded in the evidence register before reliance is transferred.	S01, S02	Conditional	Moderate	Conditional

DOCUMENT 3

Master Chronology

Chronology is a control against stale reliance. The table separates source vintage, event type and current-status use.

Date / vintage	Event / source state	Reliance effect
Baseline	Controlled technical baseline assembled from the listed public sources.	Establishes the pre-v11 evidence record; preserved in Annex A.
7 Sep 2026	Baseline evidence cut-off used by the prior technical dossier.	Historical baseline only after this supplement.
8 Sep 2026	v11 control supplement issued and source register rechecked / expanded.	Current controlled publication candidate at this package level.
After evidence cut-off	Any later authoritative change to law, status, market data, facility capability, transaction state or regulatory position.	Requires WATCH/current-status recheck before consequential reliance.

Chronology rule

A later source does not silently overwrite an earlier source. Both vintages remain visible where the change itself matters. Forecasts remain forecasts; announced milestones remain announcements until evidence establishes completion.

DOCUMENT 4

Source and Evidence Register

Every evidence URL below is visible in full and clickable. Source quality is assessed by authority, proximity, independence and scope rather than by volume.

S01 - Controlled source from baseline dossier

Organisation / custodian	Justia legal repository
Source class	Independent / contextual repository
Authority / confidence	Moderate
Evidence use	Supports only the bounded propositions linked to this source. It does not by itself establish excluded matter-specific facts.
Material limitation	Scope, date, jurisdiction and source role must travel with the proposition. Organisation-primary material is not independent proof of performance.

<https://law.justia.com/cases/federal/district-courts/new-york/nysdce/1%3A2022cv01461/575368/54/>

S02 - Controlled source from baseline dossier

Organisation / custodian	Justia court-document repository
Source class	Independent / contextual repository
Authority / confidence	Moderate
Evidence use	Supports only the bounded propositions linked to this source. It does not by itself establish excluded matter-specific facts.
Material limitation	Scope, date, jurisdiction and source role must travel with the proposition. Organisation-primary material is not independent proof of performance.

<https://docs.justia.com/cases/federal/district-courts/new-york/nysdce/1%3A2022cv01461/575368/31>

S03 - Controlled source from baseline dossier

Organisation / custodian	Midpage legal document repository
Source class	Independent / contextual repository
Authority / confidence	Moderate
Evidence use	Supports only the bounded propositions linked to this source. It does not by itself establish excluded matter-specific facts.
Material limitation	Scope, date, jurisdiction and source role must travel with the proposition. Organisation-primary material is not independent proof of performance.

<https://app.midpage.ai/document/mata-v-avianca-inc-10352027>

S04 - Formal Opinion 512 - Generative Artificial Intelligence Tools

Organisation / custodian	American Bar Association
Source class	Professional authority / primary guidance
Authority / confidence	High
Evidence use	Supports only the bounded propositions linked to this source. It does not by itself establish excluded matter-specific facts.
Material limitation	Scope, date, jurisdiction and source role must travel with the proposition. Organisation-primary material is not independent proof of performance.

https://www.americanbar.org/content/dam/aba/administrative/professional_responsibility/ethics-opinions/aba-formal-opinion-512.pdf

S05 - ABA announcement on Formal Opinion 512

Organisation / custodian	American Bar Association
Source class	Professional authority / primary guidance
Authority / confidence	High
Evidence use	Supports only the bounded propositions linked to this source. It does not by itself establish excluded matter-specific facts.
Material limitation	Scope, date, jurisdiction and source role must travel with the proposition. Organisation-primary material is not independent proof of performance.

<https://www.americanbar.org/news/abanews/aba-news-archives/2024/07/aba-issues-first-ethics-guidance-ai-tools/>

DOCUMENT 5

Public Portfolio Edition

Mata v. Avianca: When an AI Answer Looks Like Legal Research

Question: Does AI-assisted legal research survive direct source verification?

Answer supported by the public record: In Mata v. Avianca, it did not. The court found that non-existent judicial opinions with fake quotations and citations were submitted and that the lawyers failed their gatekeeping responsibilities after the authorities were challenged. The case is a clear demonstration that fluent output is not authority and that source verification must occur before reliance.

What this demonstrates

Test the answer you already have. Verify whether an AI-assisted legal research output survives direct authority checking.

Reliance boundary

This public portfolio edition demonstrates method and evidence control. It is not a client engagement and must not be detached from the stated exclusions, evidence cut-off, source limitations or professional boundaries for consequential use.

Aperture pathway

Page / contact	Visible full URL
Product page	https://www.apertureforensic.com/review.html
Selected Work	https://www.apertureforensic.com/selected-work.html
Methodology	https://www.apertureforensic.com/methodology.html
Evidence Standards	https://www.apertureforensic.com/evidence-standards.html
Evidence Lineage	https://www.apertureforensic.com/evidence-lineage.html
Reliance Boundary	https://www.apertureforensic.com/reliance-boundary.html
Languages	https://www.apertureforensic.com/languages.html
Search	https://www.apertureforensic.com/search.html
Enquiry	https://www.apertureforensic.com/enquiry.html
Paul Hattingh - LinkedIn	https://www.linkedin.com/in/paul-hattingh-79568729

DOCUMENT 6

Final Delivery Manifest & Conformity Record

Filename	04_REVIEW_Mata_v_Avianca_AI_Legal_Research.pdf
Product	REVIEW
Evidence cut-off	8 September 2026, 07:38 SAST
Governing master	PHW Research AI v11.0 Approved Governing Master
Research readiness	95/100 - ready for controlled dossier construction with specified limitations
Evidence sources	5
Controlled propositions	18
Controlled annex	Annex A preserves the complete 24-page prior technical dossier
Public release	Pending final human approval of frozen files
Website path	/selected-work/v168/04_REVIEW_Mata_v_Avianca_AI_Legal_Research.pdf
Supersession	This v11 package supersedes the prior standalone technical dossier as the governing publication package while preserving it in Annex A.

Masjien conformity

- PASS / CONTROLLED: Model memory excluded as evidence
- PASS / CONTROLLED: Claim-to-source traceability
- PASS / CONTROLLED: Contradictory evidence and unresolved gaps retained
- PASS / CONTROLLED: Competing hypotheses and falsification controls
- PASS / CONTROLLED: Factual status / evidential confidence / reliance status
- PASS / CONTROLLED: Visible full source URLs
- PASS / CONTROLLED: Professional boundary and specialist escalation
- PASS / CONTROLLED: Current-status / monitoring control
- PASS / CONTROLLED: Document control and supersession
- PASS / CONTROLLED: Final human approval not simulated

Release gate & Annex A

Final release controls

Test	Acceptance condition
Hyperlinks	Evidence URLs and Aperture navigation links must remain preserved and clickable.
Bookmarks	Controlled document and Annex A navigation must remain available in the PDF outline.
Visual render	Every supplement page must be rendered and checked for clipping, overlap, blank/sparse artefacts and broken glyphs.
Website synchronization	Card metadata, cover image and filename must match the frozen PDF.
Hash freeze	SHA-256 is calculated only after the final byte freeze.
Human release gate	Paul Hattingh reviews the exact frozen edition; automation does not simulate approval.

Annex A - Complete controlled technical baseline

The following 24 pages preserve the prior complete public-source technical dossier in full. It is retained as an audit and evidentiary baseline, not discarded. Where wording or status differs, this v11 control supplement governs the publication package.

Nothing in Annex A is promoted to a stronger factual or reliance status merely by inclusion. Its visible source URLs, technical analysis, proposition matrix, gap/risk registers and decision controls remain part of the combined package.



REVIEW | PUBLIC-SOURCE TECHNICAL DOSSIER

Mata v. Avianca: When an AI Answer Looks Like Legal Research

DECISION QUESTION

Does AI-assisted legal research survive direct source verification?

EXECUTIVE POSITION

In Mata v. Avianca, it did not. The court found that non-existent judicial opinions with fake quotations and citations were submitted and that the lawyers failed their gatekeeping responsibilities after the authorities were challenged. The case is a clear demonstration that fluent output is not authority and that source verification must occur before reliance.

PRODUCT	REVIEW
INDUSTRY / SECTOR	Legal / AI-assisted research
STATUS	Independent public-source demonstration - not a client engagement
EVIDENCE CUT-OFF	7 September 2026, 23:59 SAST
ISSUE	8 September 2026 Version 1.0

Document control & professional boundary

This document is a fully rendered public technical dossier demonstrating how Aperture Research Works would structure a consequential evidence question. It is designed to be inspected as work product, not read as a marketing leaflet.

Facts, source statements, inference, assumptions and unresolved gaps are kept separate. The public-source register is part of the product. Every listed URL is visible and clickable.

Reliance classification

Independent review of the stated evidence architecture; reliance is limited to propositions supported by the reviewed public record.

Control	Position
Prepared by	Aperture Research Works
Governing edition	English
Status	Independent public-source demonstration - not a client engagement
Evidence cut-off	7 September 2026, 23:59 SAST
Issue date	8 September 2026
Human accountability	Human-led review governs the final reliance statement
Professional boundary	This is an independent review demonstration based on public court records. It is not legal advice and does not assess every professional-responsibility issue raised by AI use.

Core operating rule

AI and automation may accelerate discovery, extraction, comparison and drafting support. They do not become evidence merely because they are fluent. Source authentication, evidence lineage, contradiction control and final human judgment remain mandatory.

Decision question, scope & executive answer

Does AI-assisted legal research survive direct source verification?

In *Mata v. Avianca*, it did not. The court found that non-existent judicial opinions with fake quotations and citations were submitted and that the lawyers failed their gatekeeping responsibilities after the authorities were challenged. The case is a clear demonstration that fluent output is not authority and that source verification must occur before reliance.

In scope

- The sanctions record in *Mata v. Avianca* and the evidentiary failure created by fabricated authorities.
- The difference between fluent legal text and authentic legal authority.
- Verification controls for AI-assisted legal research.
- Professional reliance boundaries demonstrated by the case.

Out of scope

- Legal advice on any current matter.
- An assessment of professional liability beyond the public sanctions record.
- A general claim that all generative AI legal research is unreliable.
- Jurisdiction-specific ethics advice.

What this demonstrates

Why REVIEW tests the evidence beneath an existing answer rather than assuming a polished document is reliable.

Material findings

Finding 1

The sanctions order records that non-existent judicial opinions and fake citations were submitted to the court.

Finding 2

The court expressly stated that there is nothing inherently improper about using a reliable AI tool for assistance; the failure was the abandonment of professional gatekeeping and verification.

Finding 3

The court imposed sanctions and required remedial notifications to the client and judges falsely named in the fabricated opinions.

Finding 4

The useful control lesson is not “never use AI”; it is “locate, authenticate and read the underlying authority before relying on the output.”

Reliance effect

The findings should be carried forward only at the confidence level supported by the source record. Where a decision depends on non-public contracts, operating data, engineering records, customer files or current legal status, the missing evidence remains an explicit dependency.

The failure mode

The core failure was not merely that a language model made an error. Non-existent cases, quotations and citations entered a court filing and remained there after challenge. The material control failure therefore occurred at the point where generated text was allowed to stand in place of source authentication.

- Citation existence must be checked.
- The cited proposition must actually appear in the authority.
- Pinpoint references must match the source.
- Later filings must not perpetuate an unverified citation.

Control point

Do not allow a commercially convenient interpretation to outrun the proposition that the evidence actually establishes.

Source verification protocol

A defensible AI-assisted legal research workflow treats model output as a lead, not authority. Every material case, statute, regulation or quotation should be retrieved from an authoritative or reliable legal source and checked directly. Negative checks are also important: if a case cannot be located in expected databases, that is a red flag requiring resolution before use.

- Retrieve the authority independently of the model.
- Verify party names, court, date, reporter/citation and docket.
- Read the relevant passage in context.
- Record verification status and source URL or document identifier.

Control point

Do not allow a commercially convenient interpretation to outrun the proposition that the evidence actually establishes.

Human accountability

The case illustrates why human review cannot be a ceremonial final read. The reviewer must be capable of detecting whether the authority exists and supports the proposition. A polished draft can increase risk if fluency suppresses scepticism.

- Assign named human responsibility.
- Use a citation checklist before filing.
- Escalate missing authorities rather than substituting similar cases.
- Preserve a source bundle for material propositions.

Control point

Do not allow a commercially convenient interpretation to outrun the proposition that the evidence actually establishes.

Decision use

The case is a high-signal example for any evidence-intensive workflow: confidence of expression is not evidence quality. The transferable control is to force claims through a source gate before interpretation or publication.

- No material citation without a source.
- No quotation without direct text check.
- No reliance label based on model confidence.
- Document unresolved gaps.

Control point

Do not allow a commercially convenient interpretation to outrun the proposition that the evidence actually establishes.

Proposition-level evidence matrix

This matrix is the bridge between narrative and reliance. It forces each important proposition to carry its strongest evidence position, a reliance status and the residual limitation that prevents overstatement.

Proposition	Evidence position	Status	Residual limitation
Fabricated authorities were submitted	Court sanctions opinion	Established	Case-specific.
The lawyers had gatekeeping responsibilities	Court findings	Established	Professional duties vary by context but source control is transferable.
AI output is legal authority	Nature of authority	Contradicted	Authority comes from law and authentic sources, not model output.
A citation that looks plausible can be used pending later verification	Case lesson	Contradicted	Verification must precede reliance.
Direct source checking would detect many failure modes	Verification logic	Supported	Requires competent review.
The case proves all AI legal research fails	Overgeneralisation test	Not established	Failure was workflow-specific, not universal.

Contradictions, unresolved gaps & evidence still required

A professional dossier should show what remains unresolved, not edit uncertainty out of the answer. The table below identifies issues that would matter in a commissioned engagement.

Issue / proposition	Current position	Why unresolved matters	Evidence / action required
AI output is legal authority	Contradicted	Authority comes from law and authentic sources, not model output.	Obtain stronger primary or matter-specific evidence before relying.
A citation that looks plausible can be used pending later verification	Contradicted	Verification must precede reliance.	Obtain stronger primary or matter-specific evidence before relying.
The case proves all AI legal research fails	Not established	Failure was workflow-specific, not universal.	Obtain stronger primary or matter-specific evidence before relying.
Fabricated citation	Open control issue	Can mislead court/client	Existence check before use.

Risk, failure modes & controls

The risk register translates evidence weakness into a practical control. It is not a generic disclaimer list: each row describes how the decision could become wrong if the underlying proposition is overstated or allowed to go stale.

Risk / failure mode	Why it matters	Control / next action
Fabricated citation	Can mislead court/client	Existence check before use.
Misquoted real authority	Can distort proposition	Read passage in context.
Citation laundering	Repeated secondary references create false confidence	Trace to originating authority.
Automation bias	Fluent text suppresses scrutiny	Mandatory source gate.
Deadline pressure	Encourages skipped verification	Build verification into workflow timing.
Unclear reviewer accountability	No one owns source truth	Named human sign-off.

Decision and hand-off checklist

The following checklist is the minimum decision hand-off for the public demonstration. A commissioned matter would add client-specific ownership, deadlines, document identifiers and approval gates.

- Retrieve every material authority.
- Verify citation metadata.
- Verify quoted text in context.
- Check proposition against holding or rule.
- Record source and verification status.
- Escalate any unlocatable authority.
- Do not file or advise from model output alone.

Stop condition

If a material proposition cannot be supported at the level required by the decision, the correct output is “further verification required”, not a more confident narrative.

Method, evidence discipline & reproducibility

The work begins with a decision question rather than a topic. Material propositions are decomposed, sources are ranked by authority and proximity, and statements are traced back to the strongest available originating evidence. Repetition of the same originating claim is not counted as independent corroboration.

Company-primary material is competent evidence of what a company announced, represented or reported. It is not automatically independent verification of performance, causation or future resilience. Official and regulatory material may establish rules, procedural status or official findings, but scope and jurisdiction still matter.

Negative evidence is treated cautiously. Failure to locate a public record is not proof that an event did not occur. Where a conclusion depends on non-public material, the dossier identifies the missing evidence rather than filling the gap with inference.

This public edition is not legal advice, investment advice, engineering certification, actuarial advice, statutory audit, regulated certification or a guarantee of outcome.

Evidence-status vocabulary

Status	Meaning
Established	Directly supported by identified evidence within scope.
Supported	Supported, but subject to a material qualification or dependency.
Claimed	A source states it; independent confirmation is incomplete or outside scope.
Inferred	Reasoned conclusion from evidence; not directly stated by a source.
Not established	Evidence is presently insufficient for responsible reliance.
Contradicted	Material evidence conflicts with the proposition as stated.

Public-source register

Every source below has a visible full URL and a clickable hyperlink. Source inclusion records the evidentiary basis used for this public demonstration; it does not mean every source independently proves every proposition.

1. Mata v. Avianca, Inc. - Opinion and Order on Sanctions, Document 54

Organisation: U.S. District Court, S.D.N.Y. (via Justia) | Date: 22 Jun 2023

Use: Primary court findings and sanctions

Limit: Third-party access copy of the court filing

<https://law.justia.com/cases/federal/district-courts/new-york/nysdce/1%3A2022cv01461/575368/54/>

2. Order to Show Cause, Filing 31

Organisation: U.S. District Court, S.D.N.Y. (via Justia) | Date: 4 May 2023

Use: Court identification of non-existent cases and opinions before final sanctions order

Limit: Procedural order, not final disposition

<https://docs.iustia.com/cases/federal/district-courts/new-york/nysdce/1%3A2022cv01461/575368/31>

3. Mata v. Avianca, 678 F.Supp.3d 443

Organisation: S.D.N.Y. opinion reproduction | Date: 22 Jun 2023

Use: Text of sanctions opinion and Rule 11 gatekeeping discussion

Limit: Secondary reproduction of court opinion

<https://app.midpage.ai/document/mata-v-avianca-inc-10352027>

Aperture pathway, contact & linked pages

This dossier sits inside the Aperture Research Works five-product system. RESEARCH finds the answer. REVIEW tests an answer already in hand. WATCH keeps a changing answer current. DESK adds governed recurring capacity. APERTURE determines what can responsibly be relied upon.

[Aperture Research Works: https://apertureforensic.com/](https://apertureforensic.com/)

[REVIEW product page: https://apertureforensic.com/review.html](https://apertureforensic.com/review.html)

[Selected Work: https://apertureforensic.com/selected-work.html](https://apertureforensic.com/selected-work.html)

[Methodology: https://apertureforensic.com/methodology.html](https://apertureforensic.com/methodology.html)

[Evidence standards: https://apertureforensic.com/evidence-standards.html](https://apertureforensic.com/evidence-standards.html)

[Evidence lineage: https://apertureforensic.com/evidence-lineage.html](https://apertureforensic.com/evidence-lineage.html)

[Reliance boundary: https://apertureforensic.com/reliance-boundary.html](https://apertureforensic.com/reliance-boundary.html)

[Languages: https://apertureforensic.com/languages.html](https://apertureforensic.com/languages.html)

[Search Aperture: https://apertureforensic.com/search.html](https://apertureforensic.com/search.html)

[Discuss a Matter: https://apertureforensic.com/enquiry.html](https://apertureforensic.com/enquiry.html)

[Paul Hattingh on LinkedIn: https://www.linkedin.com/in/paul-hattingh-79568729](https://www.linkedin.com/in/paul-hattingh-79568729)

Email: info@apertureforensic.com

WhatsApp: +27 83 450 5051

Professional boundary

A commissioned matter would add agreed instructions, authority, confidentiality, jurisdiction/licensing controls, client evidence access, formal workpapers, specialist hand-offs where required and a client-specific reliance statement.

Complete approved case-pack appendix follows

The approved existing public case pack is reproduced after this divider so that prior detailed evidence work is retained in full rather than replaced by a short summary.



APERTURE

RESEARCH WORKS

MATA V. AVIANCA, INC.

AI-Generated Judicial Authorities, Failed Citation Verification and Attorney Gatekeeping

HUMAN-APPROVED FOR PUBLIC PORTFOLIO DISPLAY

PUBLIC PORTFOLIO EDITION

Authorship: Paul Hattingh and Aperture Research Works

Intended users: law firms, publishers, boards, professional bodies, litigation-support teams, researchers and decision-makers assessing AI-assisted research integrity.

Reality before reliance. Verification before interpretation.

No conclusion stronger than the available evidence.

info@apertureforensic.com | apertureforensic.com

Field	Details
Matter reference	ARW-CASE-2026-001
Document code	PPE
Governing master	PHW Research AI v10.5
Evidence cut-off	5 August 2026
Source-freeze date	5 August 2026
Approved by	Paul Hattingh
Approved public use	Website, portfolio, educational and research communication
Release status	HUMAN-APPROVED FOR PUBLIC PORTFOLIO DISPLAY

DOCUMENT CONTROL AND PUBLIC-USE STATUS

Field	Details
Public brand	Aperture Research Works
Matter title	Mata v. Avianca, Inc. - AI-Generated Judicial Authorities, Failed Citation Verification and Attorney Gatekeeping
Case reference	ARW-CASE-2026-001
Document title	Approved Public Portfolio Edition
Version	3.0 Public Release
Governing master	PHW Research AI v10.5
Research period	4-5 August 2026, including migration and current-status recheck
Evidence cut-off / source freeze	5 August 2026
Operating mode	Independent public-source case study - not a client engagement
Human approval	Approved by Paul Hattingh on 5 August 2026 for public portfolio display
Independent review	Not completed
Legal review	Not completed; no legal advice or legal reliance is authorised
Specialist review	Native digital-forensics and New York legal-ethics review not completed
Right of reply	Public adversarial and sworn record reviewed; no fresh contact required for this record-based edition
Permitted use	Public website, portfolio, educational and research communication
Excluded use	Client-specific legal, disciplinary, expert, investment or regulated professional reliance
Contact	info@apertureforensic.com apertureforensic.com

PROFESSIONAL LIMITATION AND DISCLAIMER

This is an independent public-source case study. Aperture Research Works did not represent any party and did not have confidential access. It is not legal advice, legal representation, expert evidence, a disciplinary determination, a certified PACER record or a substitute for suitably qualified counsel. The conclusions are limited to the public evidence reviewed through the stated cut-off.

CONTENTS

Why this matter matters
Central Question
Scope and exclusions
Evidence categories reviewed
Principal verified findings
Strongest counter-evidence and mitigation
Key evidence gaps
Conclusion classification
Reliance result
Current status and later treatment
What the case demonstrates
Client problems Aperture can assist with
Public release and lifecycle controls
Selected public sources

WHY THIS MATTER MATTERS

A personal-injury action became a leading professional-control case after counsel filed an opposition containing nonexistent judicial opinions, false quotations, mismatched citations and mischaracterised real authorities generated through ChatGPT. The material failure was not merely false machine output. It was the breakdown of source authentication, signer review, supervision, warning escalation and correction controls before and after filing.

CENTRAL QUESTION

What does the public record establish about how AI-generated legal authorities entered the Mata filings, why the defects survived professional and judicial warning points, what the court actually found and sanctioned, and what reliance and governance controls may responsibly be derived without exceeding the evidence?

SCOPE AND EXCLUSIONS

The assessment covers the public federal-court record from removal in February 2022 through judgment in July 2023, together with relevant professional rules, later guidance and representative later judicial treatment reviewed through 5 August 2026.

- No confidential law-firm records or private communications were available.
- No original ChatGPT account records or native prompt history were available.
- No certified current PACER export or certified appellate-history search was obtained.
- Aperture does not determine professional discipline, malpractice liability or present fitness to practise.
- The case does not establish a universal conclusion about every lawyer, law firm or generative-AI system.

EVIDENCE CATEGORIES REVIEWED

Court opinions and orders; filed party submissions; hearing-record passages; docket and judgment records; Federal Rule 11; New York professional and court-administration rules; ABA guidance; NIST generative-AI risk guidance; peer-reviewed technical research; and representative later judicial decisions.

PRINCIPAL VERIFIED FINDINGS

ID	Finding	Classification	Reliance
F01	The sanctions court found that nonexistent opinions and false or mismatched legal material were submitted.	Established	Suitable for reliance
F02	The signer did not independently verify the cited authorities before filing.	Established	Suitable for reliance
F03	Warning and production orders did not produce prompt, transparent correction.	Established	Suitable for reliance
F04	The respondents continued to stand by defective material until the sanctions process advanced.	Established	Suitable for reliance
F05	The court imposed a joint \$5,000 penalty and notification requirements.	Established	Suitable for reliance
F06	The underlying claim was separately dismissed as time-barred.	Established	Suitable for reliance
F07	Over-trust in an unfamiliar tool, deficient supervision and inadequate escalation probably contributed to the error.	Probable	Conditional reliance
F08	A pre-planned intention to fabricate cases from the outset was not established.	Not established	Do not assert

STRONGEST COUNTER-EVIDENCE AND MITIGATION

The record supports a genuine early misunderstanding of a new tool, absence of an established plan to fabricate authority from the outset, later admissions, expressions of remorse and remedial steps. These considerations narrow motive and character conclusions, but do not defeat the court-adjudicated failures of verification, candor and professional gatekeeping.

EVIDENCE GAPS AND LIMITS

Gap	Why it matters	Effect on conclusion
Certified current PACER and appellate court records.	Confirms current procedural status.	Does not reverse the 2023 sanctions and judgment record.
Original ChatGPT prompts and account reports.	Required for native provenance and prompt reconstruction.	Limits technical and motive findings.
Private law-firm supervision and control records.	Required to allocate organisational responsibility.	Prevents broader firm-culture conclusions.
Official public disciplinary disposition timeline.	Required to present professional-status claims.	No public disciplinary outcome is asserted.
Independent legal and digital-forensics reviews.	Required to strengthen external professional reliance.	Public portfolio use remains non-legal and non-expert.

CONCLUSION CLASSIFICATION

Established: the filed legal authorities contained material authenticity and citation defects and required professional verification controls failed. Probable: over-trust in an unfamiliar generative tool, deficient supervision and inadequate escalation combined to permit continuation after warnings. Not established: a pre-planned intention to fabricate cases, client participation, confidentiality breach, present professional discipline or a causal claim that the false authorities caused dismissal.

RELIANCE RESULT

Suitable for reliance for reporting the district court findings, the filing chronology, the distinction between sanctions and merits, and the minimum control principle that every cited authority must be opened, authenticated and proposition-checked before reliance. Conditional or limited reliance applies to current professional status, organisational culture, motive and technical reconstruction.

CURRENT STATUS AND LATER TREATMENT

The reviewed public record records judgment for Avianca and closure on 7 July 2023. No appeal was located in the searched public sources, but a certified current docket was not obtained. Mata remained influential in 2026 New York decisions addressing fabricated or erroneous AI-assisted legal authorities, including Cassata, Torres and Landberg. Later treatment is classified as legal and professional-control context, not new evidence of the original events.

WHAT THIS CASE DEMONSTRATES

- AI Research and Citation Audit
- Source Authentication
- Claims-and-Evidence Audit
- Legal and Factual Chronology
- Source-Dependence Analysis
- Independent Adversarial Review
- Evidence-Gap Identification
- Reliance Classification
- Specialist Escalation
- Defensible Decision-Record Construction

CLIENT PROBLEMS APERTURE CAN ASSIST WITH

A law firm, publisher, investor, board, professional body or litigation-support team may need to determine whether a report, submission or AI-assisted research product relies on genuine sources, whether citations establish the propositions asserted, what counter-evidence exists, and what may safely be relied upon before capital, reputation or authority is committed.

MINIMUM DECISION CONTROL

Control stage	Required action
Source existence	Open the original authority or record.
Identity and version	Confirm title, court or issuer, date, citation and current version.
Proposition test	Confirm the source actually establishes the proposition asserted.
Quotation test	Compare every quotation with the original text and locator.
Counter-evidence	Search for contrary, limiting or superseding authority.
Signer review	The accountable human must review the final cited product before release.
Correction	Withdraw, correct and disclose promptly when a defect is found.

PUBLIC RELEASE AND LIFECYCLE CONTROLS

Paul Hattingh authorised this exact edition for public portfolio display on 5 August 2026. The approval does not convert the work into legal advice, expert evidence, a disciplinary finding or a certified current docket. The document must be updated or withdrawn if a material later court, disciplinary or factual development changes the stated reliance position.

Trigger	Required response
Material new appeal, order or disciplinary outcome	Reopen current-status research, update chronology and reassess reliance.
Correction to a source, citation or factual statement	Record the defect, correct the edition and issue a supersession notice.
New confidential or private evidence	Do not add automatically; confirm lawful authority, privacy and publication basis.
New allegation about a living person	Reassess legal risk, proportionality and right of reply.
Use for a client-specific decision	Create a separate scoped client edition with appropriate legal or specialist review.

SELECTED PUBLIC SOURCES

[S01] Mata v. Avianca, Inc., ECF 54, S.D.N.Y., 22 June 2023. Sanctions findings and reasoning.

[S02] Mata v. Avianca, Inc., ECF 55. Merits dismissal and distinction from sanctions.

[S03] Clerk's Judgment, ECF 59, 7 July 2023. Judgment and closure record.

[S04] Federal Rule of Civil Procedure 11. Signer responsibility and sanctions framework.

[S05] ABA Formal Opinion 512. Generative-AI professional-responsibility guidance.

[S06] New York Part 161 and Appendix A. Court-system guidance on generative AI.

[S07] NIST AI 600-1. Generative-AI risk-management profile.

[S08] Landberg v. City of New York, 2026 NY Slip Op 03935.

[S09] Cassata v. Michael Macrina Architect, P.C., 2026 NY Slip Op 26014.

[S10] Torres v. Spraker, 2026 NY Slip Op 26077.



APERTURE

RESEARCH WORKS

END OF PUBLIC PORTFOLIO EDITION

Mata v. Avianca, Inc. | ARW-CASE-2026-001

HUMAN-APPROVED FOR PUBLIC PORTFOLIO DISPLAY

AUTHOR AND CONTACT

Authorship: Paul Hattingh and Aperture Research Works

Reality before reliance. Verification before interpretation.

No conclusion stronger than the available evidence.

info@apertureforensic.com | apertureforensic.com

Mata v. Avianca - AI Legal Research

REVIEW | ANNEX B - SOURCE REGISTER CONSOLIDATION

CONTROL NOTICE. This Annex B is part of the controlled publication package. It expands the earlier Document 4 source register from 10 explicitly numbered source records to a consolidated minimum of 12. Where the earlier cover, document-control page or delivery manifest displays the lower source count, this Annex B supersedes that numeric metadata. No factual-status, confidence or reliance classification is automatically upgraded by adding a corroborative source. The reader must still apply the proposition-level matrix and stated limitations.

Existing numbered sources	10
Additional controlled sources	2
Consolidated source count	12
Evidence cut-off	8 September 2026, 07:38 SAST for substantive dossier; Annex B source consolidation verified 9 September 2026.
Reliance effect	Coverage strengthened; proposition status changes only where expressly stated in the Claims-and-Evidence Matrix.

Additional controlled source records

S11	U.S. District Court, Southern District of New York Court rules and local rules
Full URL	https://www.nysd.uscourts.gov/rules
Evidence role	Primary procedural framework for filings in the court where Mata was litigated.
Supports	C07-C18
Reliance limit	Use only for the proposition scope stated. Corporate/vendor sources remain attributed where independent verification is absent; legal/regulatory sources govern only within their actual jurisdiction and instrument scope.

S12	American Bar Association Model Rule 3.3 - Candor Toward the Tribunal
Full URL	https://www.americanbar.org/groups/professional_responsibility/publications/model_rules_of_professional_conduct/rule_3_3_candor_toward_the_tribunal/
Evidence role	Professional-responsibility benchmark for candor and correction duties; not itself a jurisdiction-specific rule.
Supports	C07-C18
Reliance limit	Use only for the proposition scope stated. Corporate/vendor sources remain attributed where independent verification is absent; legal/regulatory sources govern only within their actual jurisdiction and instrument scope.

Traceability and publication control

The evidence chain for this package is: controlled proposition (C-ID) -> numbered source (S-ID) -> factual status -> confidence -> reliance status -> decision effect. Annex B extends the source side of that chain only. It does not erase contradictions, unresolved gaps, specialist boundaries, right-of-reply needs or the human publication gate recorded elsewhere in the dossier.

PUBLIC PORTFOLIO LIMITATION: Independent public-source case study. Not a client engagement. Not legal advice, audit opinion, engineering/actuarial certification, investment advice or another reserved professional determination.