

THE AI INSTITUTE / OPERATING INTELLIGENCE

EXECUTIVE CASE NOTE · 03 OCT / 2026



Keep AI Reconstructions Separate from Evidence

A convincing presentation must not turn an inferred view into a person's own words.

Boards, C-suite and senior management
Research period: September 2026 – October 2026
First published: 3 October 2026
Current to 3 October 2026

PUBLICATION RECORD

An Institute thesis and decision framework

Before using AI-generated representations in a consequential decision, separate recorded evidence, attributed interpretation and invented illustration—and make the decision defensible without the synthetic presentation.



WHAT LEADERS SHOULD TAKE FROM THIS PAPER

- | | |
|----|---|
| 01 | Disclosure does not establish that a statement is reliable. |
| 02 | Keep a person's recorded words separate from another person's reconstruction. |
| 03 | Review the source before the persuasive presentation. |
| 04 | Record which evidence actually supports the decision. |

Research method

Evidence base: the Arizona Court of Appeals opinion of 30 September 2026, read in full through its Justia reproduction, and selected research from the 2 October arXiv batch. Business controls are original Institute analysis. No private conversation or personal anecdote is quoted.

First published 3 October 2026 · Updated 3 October 2026 · Research period September 2026 – October 2026 · Version 1.0 · Suggested citation: The AI Institute, *Keep AI Reconstructions Separate from Evidence* (2026).

CONTENTS

Inside this paper

01	The label did not settle the question The case
02	Keep interpretation attached to its author Institute analysis
03	Approve the evidence before the performance Before presentation
04	Make the recommendation survive without the avatar During the decision
05	Transfer the control, not the legal conclusion Across markets
06	Record the basis, not just the approval The next review
M	Method and limitations
R	References

How to use this brief

Start with the decision and use the framework to challenge the current plan. Figures retain their original populations and definitions. Institute analysis is labelled, and limitations are recorded before the references.

The label did not settle the question

Before allowing an AI reconstruction into a consequential decision, ask what it is evidence of. A convincing voice can make an interpretation feel like a direct statement. A realistic scene can make a proposed future feel observed. Neither becomes better evidence because the production is impressive.

On 30 September 2026, the Arizona Court of Appeals affirmed Gabriel Paul Horcasitas's manslaughter conviction but vacated his sentence and ordered resentencing. At sentencing, an AI video had depicted the deceased victim speaking. The video disclosed its artificial origin and included genuine footage. The court distinguished that footage from imagined statements presented in the victim's own likeness and voice. It found the synthetic material unreliable and prejudicial; the precise effect on the sentence was not quantified.¹

This is a particular US sentencing decision, not a worldwide rule against synthetic media. The Institute's business lesson is narrower: knowing how material was made does not establish whether it should influence the decision in front of you.

Keep interpretation attached to its author

Organisations already use interpretation legitimately. A manager explains what a customer complaint may mean. An engineer proposes how a failure happened. A researcher describes a possible user response. The problem starts when the presentation removes the distance between that interpretation and the person or event being represented.

Consider an illustrative customer avatar in an investment presentation. It can help a team explore a service idea. It cannot establish that customers asked for the service, accepted its price or understood its risks. Those claims need interviews, observed behaviour or another suitable source. Calling the avatar synthetic does not supply the missing demand evidence.

The control is attribution, not a ban on explanation. Keep the researcher's inference in the researcher's voice. If a reconstruction helps discussion, name its purpose and assumptions. Do not give it the apparent authority of a real customer, employee or independent expert whose views it merely predicts.

Approve the evidence before the performance

Ask the decision owner to review the underlying material before the polished version. Separate direct records from summaries and illustrative additions. A source pack should make it easy to find the original statement, the surrounding context and any omissions that could change its meaning.

This does not require directors to inspect every raw file. It requires a responsible reviewer to be able to explain the path from source to claim. For a consequential assertion, preserve the source, identify the transformation and show why the transformed version remains suitable for the decision.

Test that path with a small challenge. Replace a quantity, remove a qualifying sentence or change who is speaking. Does the review notice that the conclusion should change? New research on numerical claim verification reports that small numerical changes can expose weaknesses in language-model checking, while targeted training can help on the tested datasets. It is supporting evidence for testing the reviewer, not a guarantee about a particular product.²

EXHIBIT / THREE DIFFERENT ROLES FOR DECISION MATERIAL

01 Record What was actually said, observed or measured? Institute recommendation	02 Interpretation Who draws this conclusion, and from which record? Institute recommendation	03 Illustration What is invented to explore or explain a possibility? Institute recommendation
--	--	--

Make the recommendation survive without the avatar

Run the decision once without the synthetic presentation. Use the source evidence and a plain-language account of uncertainty. If the recommendation changes after the performance is added, ask what new information the performance supplied. If it added none, the change deserves examination rather than automatic acceptance.

This is an Institute recommendation, not a claim that every emotional response is irrational. Experience and human testimony can convey important consequences that a spreadsheet misses. The aim is to keep the person who experienced those consequences distinct from someone else's model of that person.

Provide a route to challenge attribution before approval. An employee representative may question whether a simulated worker speaks for colleagues. A customer researcher may reject a composite persona as evidence of willingness to pay. A technical reviewer may identify animation that hides an unresolved failure. Give those objections an owner and a recorded answer.

Transfer the control, not the legal conclusion

The Arizona ruling concerns a US criminal sentencing process. Boards elsewhere should not copy its legal conclusion into employment, procurement or investment policy. Ask local counsel which evidentiary, privacy and likeness rules apply to the actual use. The management practice can travel without pretending the law is identical.

Language also changes the review. A translated or dubbed statement can alter emphasis, uncertainty or social meaning while retaining a recognisable face. Review the original and the translation with someone competent in the relevant language. Do not substitute fluent output for that competence.

A new Japanese-language evaluation tested eleven models on three benchmarks with five types of typing error. Effects varied by error type. That is not a study of synthetic-video persuasion, but it reinforces a useful limit: performance on clean text is not enough to establish reliability in local input conditions. Evidence for other languages and institutions must be gathered, not assumed. ³

Record the basis, not just the approval

For the next material AI-assisted proposal, ask its sponsor to mark each decisive claim as a record, an interpretation or an illustration. Start with claims that affect funding, employment, customer rights or safety. Resolve ambiguous attribution before the meeting.

After the decision, retain the evidence that justified it, the uncertainties accepted and who authorised reliance on the transformed material. Preserve the synthetic version too where retention rules permit, so a later review can see what decision-makers actually encountered. Access and retention should follow the sensitivity of the underlying records.

Good synthetic media can make a difficult idea easier to discuss. Its value is explanatory. Do not let that strength quietly change the status of the evidence. A decision should remain defensible when the compelling face and voice are gone.

RESEARCH RECORD

Method and limitations

Method

Evidence base: the Arizona Court of Appeals opinion of 30 September 2026, read in full through its Justia reproduction, and selected research from the 2 October arXiv batch. Business controls are original Institute analysis. No private conversation or personal anecdote is quoted.

Limitations

One appellate case does not establish a global legal rule or a general rate of AI-induced decision error. The official PDF endpoint was inaccessible; the reproduced opinion and official case listing were checked. Numerical-verification work is listed as accepted to AACL-IJCNLP 2026 Findings; the Japanese evaluation is a preprint. Neither studies persuasion by synthetic video.



Sources and limitations

Source facts link to the original publication. Institute analysis reconciles definitions and turns the material into a management proposition. Frameworks are decision aids, not claims of universal causality. Read every important statistic with its population, date, definition and caveat.

Research cutoff: 3 October 2026. Review cycle: six months, or earlier after a material change in the underlying facts, standards or policy.

REFERENCES

Evidence behind the thesis

01 Arizona Court of Appeals: State v. Horcasitas, 1 CA-CR 25-0191

Filed 30 September 2026; opinion reproduced by Justia. Paragraphs 33–54 concern the AI video. Conviction affirmed; manslaughter sentence vacated and remanded.

02 Aarnes and Setty: Towards Robust Numerical Claim Verification

v1 submitted 30 September, listed 2 October 2026. Authors report acceptance to ACL-IJCNLP 2026 Findings. Controlled numerical perturbation tests, not a universal verification guarantee.

03 Mibayashi and Ohshima: Japanese LLM robustness to typing errors

v1 submitted 1 October, listed 2 October 2026. Preprint: eleven models, three Japanese benchmarks, five typo categories. No video-persuasion or cross-country population study.

THE AI INSTITUTE

Make the decision.
Then measure what changes.

Research · Readiness · Capability

theai.institute