

From Recollection to Reconstruction

Machine-Mediated Police Narratives and the Transformation of Evidentiary Reliability

Author: Jeffrey Meyers

Version: 0.7 (Final Publication Master)

Status: Policy White Paper

About the Author

Jeffrey Meyers is an investigator, founder, and strategist operating at the intersection of criminal justice, forensic integrity, and sovereign technological architecture. Drawing upon an extensive foundational background in law enforcement and organized crime investigation with the Royal Canadian Mounted Police (RCMP), his career centers on defending operational data provenance and structural accountability.

As the Founder of VRFY, Jeffrey Meyers leverages this deep domain expertise to pioneer advanced forensic verification and secure infrastructure solutions explicitly tailored for law enforcement, public safety, and national security environments. His operational insights directly inform his advocacy for robust digital provenance and technical safeguards, ensuring that the next generation of artificial intelligence systems meets the rigorous, uncompromised evidentiary standards demanded by democratic courts worldwide.

Executive Summary

The emergence of generative artificial intelligence (AI) within policing represents a fundamental shift in the evidentiary assumptions underpinning democratic criminal justice systems. Historically, police reports derived legitimacy from interconnected principles: human authorship, identifiable perception, contemporaneous independent recollection, and direct accountability. Generative AI disrupts this framework by introducing machine-mediated narrative formation directly into the evidentiary chain.

Public debate has largely focused on algorithmic bias, privacy, and automated surveillance. While critical, these concerns obscure a deeper evidentiary challenge: the transformation of contemporaneous human recollection into post-event synthesized reconstruction. Unlike transcription technologies that preserve spoken words verbatim, generative AI actively organizes chronology, summarizes interactions, prioritizes relevance, and constructs narrative flow. In doing so, it moves from evidence preservation to evidence interpretation.

This shift has profound implications for evidentiary reliability, disclosure obligations, procedural fairness, and public trust. Crucially, generative AI systems risk artificially suppressing the ambiguity, inconsistency, fragmentation, and uncertainty that traditionally form the baseline of a reasonable doubt analysis. This paper examines how machine-mediated narrative construction challenges the core tenets of common law evidence, cognitive psychology, and international legal frameworks, arguing that future systems must enforce a strict separation between evidence preservation and machine-generated interpretation.

1. Introduction

Artificial intelligence is rapidly reshaping modern policing through the deployment of systems capable of automatically producing narrative police reports from body-worn camera (BWC) footage, audio recordings, and computer-aided dispatch (CAD) data. Marketed as assistive efficiency tools designed to reduce administrative burdens, these systems position officers as reviewers who edit and approve machine-generated drafts.

However, the legal and evidentiary implications of this transition remain poorly understood. While legal scholarship frequently critiques algorithmic bias and surveillance, the deeper risk lies in the mediation of narrative formation.

Historically, the legal system accepted that human memory is imperfect and evolved doctrines—such as contemporaneous note-taking and refreshed memory—to preserve accountability while maintaining identifiable human authorship. Generative AI alters this dynamic. A verbatim transcript captures what was said; a generative narrative constructs an interpretive representation of what allegedly occurred.

This introduction of machine-mediated narrative construction creates a new evidentiary category: **post-event synthesized reconstruction**. This paper examines whether existing evidentiary doctrines remain conceptually adequate when a machine assumes the role of narrative constructor, ultimately outlining the governance mechanisms required to preserve procedural legitimacy and public trust.

2. Foundations of Recollection in Common Law Evidence

The foundational doctrines governing witness testimony and police notes evolved around deeply human assumptions: that individuals perceive events directly, form memories imperfectly, and record those impressions close in time to preserve reliability.

2.1 Human Recollection and Testimonial Reliability

The legitimacy of testimonial evidence derives from identifiable perception, personal cognition, and the ability to test reliability through adversarial scrutiny. Police reports are not merely objective logs; they contain threat assessments, behavioral interpretations, and contextual analyses that cannot be captured by recording devices alone. A BWC documents a physical field of view, but it cannot capture peripheral awareness, fear perception, or tactile sensations. Evidentiary value relies on the assumption that the narrative reflects direct human cognition challengeable through cross-examination (*R. v. Fliss*, 2002 SCC 16).

2.2 Present Memory Revived and Refreshed Recollection

Under the doctrine of present memory revived, a witness may review a document or recording to refresh their memory before testifying (Federal Rule of Evidence 612). However, the evidence remains the witness's present recollection, not the document itself; the document functions strictly as a stimulus. If a witness merely adopts external wording or a machine-generated narrative structure, the reliability foundation shifts from active recollection to passive adoption.


2.3 Past Recollection Recorded and the Wigmore Framework

Where independent recollection is entirely lost, the document itself may become admissible under past recollection recorded if it meets the traditional Wigmore criteria: personal observation, contemporaneous creation or adoption while memory was fresh, and a past acknowledgment of accuracy. The legal system tolerates these exceptions precisely because the human witness remains the originating cognitive source. Generative AI complicates this by pre-structured the narrative before the witness independently formulates their recollection.

3. Human Memory, Cognitive Reconstruction, and Narrative Contamination

3.1 Reconstructive Memory and Evidentiary Reliability

Modern cognitive science demonstrates that human memory does not function like a video recording; it is dynamic, reconstructive, and highly susceptible to post-event information and narrative framing (Loftus & Palmer, 1974).



When an officer reviews an AI-generated report, they face a severe risk of **unconscious memory contamination**. Generative systems are optimized to produce fluid, persuasive, and internally consistent text. Because human cognition strongly prefers coherent narrative structures, an officer may inadvertently absorb machine-generated chronology, causal framing, and synthesized language into their own subjective recollection. Wording never independently selected by the officer gradually becomes cognitively familiar, creating several distinct evidentiary risks:

- Post-event memory alignment with probabilistic model outputs.
- The artificial reinforcement of certainty regarding ambiguous events.
- An accelerated cognitive dependency on synthesized narrative templates.

3.2 The Reality of the Baseline Report

Critically, this risk is amplified by institutional realities. Traditional human-authored police reports are rarely pristine reflections of pure memory; they are constrained by administrative pressure and boilerplate templates. However, while human reporting has always suffered from individual imperfections and systemic biases, generative AI scales these limitations exponentially. It replaces localized, human inconsistency—which defense counsel can actively probe under cross-examination—with a centralized, sanitized, and authoritative machine-generated veneer.

4. Body-Worn Cameras and the Limits of Objective Recording

4.1 The Myth of Complete Objectivity

While BWCs provide critical audiovisual preservation, they capture a technologically mediated perspective, not human perception. They record directional audio and a constrained visual frame; they do not document an officer’s attention allocation, adrenaline response, or evolving threat interpretation under stress. BWC footage functions as evidence of an event, never as definitive evidence of an officer's subjective perception.

4.2 Source Evidence Versus Narrative Interpretation

A critical conceptual divide must be maintained: **raw recordings preserve source evidence, whereas generative narratives interpret it.**

Evidentiary Medium	Operational Function	Primary Legal Value	Risk Profile
Source Evidence (BWC / Audio)	Verbatim capture, immutable recording, and sensory data preservation.	Objective baseline of environmental facts.	Technological constraints (angles, lighting, direction).
Narrative Interpretation (Generative AI)	Summarization, chronological sequencing, and causal framing.	Administrative efficiency and standardized formatting.	Cognitive overwriting, automation bias, and artificial certainty.

Real-world incidents are defined by environmental chaos and overlapping communications. When a generative AI system summarizes these interactions, it does not merely copy data; it assigns statistical weight, omits perceived noise, and establishes a linear chronology. It transitions from preservation to active interpretation.

5. Generative AI and Machine-Mediated Narrative Formation

5.1 From Information Processing to Narrative Construction

Unlike prior law enforcement technologies that merely stored or transmitted human-inputted data, generative large language models (LLMs) are active narrative constructors. They generate outputs based on probabilistic prediction, semantic inference, and statistical language modeling. Every generated report reflects underlying model architecture, training data influences, and prompt weighting, creating a seamless appearance of objective neutrality that masks interpretive choices.

5.2 The Adoption Problem and Automation Bias

Vendors seek to maintain legal defensibility by relying on "human-in-the-loop" verification, requiring officers to review and sign off on drafts. However, this safety valve is fundamentally compromised by **Automation Bias**—the well-documented cognitive tendency for humans to over-trust automated systems, particularly when fatigued, facing immense administrative backlogs, or operating under strict disclosure deadlines.

Verification is not equivalent to authorship. An officer suffering from automation bias is highly likely to treat a highly coherent AI draft as "accurate enough," implicitly adopting sequencing, characterizations, and vocabulary they never independently formulated.

6. From Recollection to Reconstruction

The integration of generative AI signals a transition from evidentiary recollection to evidentiary reconstruction. Under traditional models, the observing officer was the sole author of the narrative. Under an AI-mediated workflow, the machine acts as an **interpretive intermediary**, synthesizing a polished narrative structure before the officer can independently formulate their recollection.

This creates **artificial narrative certainty**. In dynamic use-of-force or crisis encounters, human memory is naturally fragmented. Generative AI eliminates this organic instability by design, smoothing over hesitation, rectifying inconsistencies, and imposing absolute chronology. This polished coherence creates an artificial inflation of evidentiary reliability, presenting an unpolished, chaotic human experience as a highly certain, linear sequence of events.

7. Comparative International Legal and Governance Analysis

While terminology and statutory frameworks differ internationally, democratic legal systems demonstrate a clear convergence around the requirements of identifiable authorship and procedural transparency.

7.1 Canada

Canadian evidence law strictly distinguishes between refreshed memory and substituted narrative, imposing broad constitutional disclosure duties under *R. v. Stinchcombe* ([1991] 3 S.C.R. 326). In an AI-mediated environment,

Stinchcombe obligations will inevitably expand to encompass the technical infrastructure of the report itself: original prompt histories, confidence scores, underlying model versions, and system interaction logs. This raises complex litigation hurdles regarding whether proprietary vendor algorithms constitute missing first-party evidence or third-party records.

7.2 United States

In the United States, criminal litigation will focus heavily on *Brady v. Maryland* (373 U.S. 83) and *Giglio v. United States* (405 U.S. 150). If an LLM generates multiple distinct draft variations before the officer accepts the final version, those discarded iterations, prompt instructions, and system parameters represent textbook impeachment material that must be disclosed to the defense to protect constitutional due process.

7.3 United Kingdom

The UK disclosure environment, governed by the Criminal Procedure and Investigations Act 1996 (CPIA), already faces institutional strain over digital evidence management. The National Police Chiefs' Council (NPCC) AI Strategy emphasizes explainability and human accountability. Machine-generated narratives introduce severe friction into this framework, as the probabilistic nature of LLMs makes it impossible to provide an auditable, step-by-step account of why a specific phrase or causal inference was generated.

7.4 European Union

Under the European Union AI Act (Regulation (EU) 2024/1689), AI systems utilized in law enforcement are classified as high-risk technologies. Article 14 mandates strict human oversight capabilities to prevent automation bias. However, the EU framework will face significant friction when confronting narrative reporting, as standard technical oversight metrics cannot adequately detect or regulate the subtle cognitive contamination occurring within an officer's mind during draft review.

8. Evidentiary Risk, Opaque Systems, and Adversarial Asymmetry

Because police narratives directly shape charging decisions, bail hearings, and judicial interpretations, narrative framing carries immense institutional power. The introduction of generative AI into this loop, however, deepens a structural crisis within adversarial litigation: **adversarial asymmetry**. Under standard criminal discovery models, the defense possesses a constitutional right to confront state witnesses and challenge the methods by which evidence was compiled. Opaque, machine-mediated interpretation effectively truncates this right, rendering parts of the data lineage opaque to traditional adversarial testing.

When an interpretive narrative is generated by a Large Language Model, its text is the product of multi-dimensional vector math, weight matrices, and localized probabilistic scoring. Neither the reviewing officer nor the third-party commercial vendor can provide a traceable, step-by-step rationalization of why a specific verb was chosen, why a specific conversational exchange was omitted, or how the chronology was stabilized. This introduces an unchallengeable, closed-system intermediary directly into the evidentiary chain, creating an unmeasurable hurdle for classic cross-examination mechanisms.

Furthermore, this opacity creates severe friction across three operational digital layers:

- **Prompt Engineering as Testimony:** System prompts, hidden pre-prompts, and RAG (Retrieval-Augmented Generation) instructions act as invisible guardrails that shape how an AI filters BWC footage. If a prompt instructs an engine to "prioritize signs of civilian resistance," the resulting draft is inherently skewed. System instructions are no longer merely code; they are an influential form of pre-authored, algorithmic interaction that must be fully disclosable.
- **Opaque Metadata and Hyper-Parameters:** Fluctuations in model temperatures, system-level top-p configurations, and localized probability distributions directly determine text variance. A shift in these parameters can alter a narrative from a conservative log into a speculative, highly contextual reconstruction, yet this metadata remains hidden behind proprietary vendor firewalls.
- **The Proprietary Secret Barrier:** When defense counsel requests access to base training data sets, prompt iteration histories, or neural weights to test for systematic confirmation bias, vendors routinely invoke trade-secret protections. This structural imbalance violates the principles of fundamental justice, forcing the accused to confront a proprietary commercial algorithm insulated from judicial scrutiny whose internal interpretive mechanics are legally shielded from adversarial testing.

9. Artificial Narrative Certainty and the Erosion of Reasonable Doubt

Criminal conviction standards are explicitly designed to tolerate and evaluate human uncertainty. Proof beyond a reasonable doubt depends heavily on the trier of fact analyzing the natural imperfections of human recollection: hesitation, shifting chronology, fragmented sequencing, and linguistic awkwardness. These organic gaps are not flaws; they are indicators of authenticity that provide the necessary landscape through which reasonable doubt is assessed.

Generative AI suppresses this landscape by substituting organic human imperfection with artificial narrative consistency. By optimizing for fluency and coherence, the machine masks the true level of situational uncertainty present during the event. This introduces an unprecedented systemic risk: it shifts the operational balance of criminal adjudication by artificially reducing the visible indicia of doubt within state evidence, undermining cross-examination before adversarial testing even begins.

10. Architectural Separation: Preservation vs. Interpretation

To insulate democratic criminal proceedings from the liabilities of adversarial asymmetry and cognitive overwrite, the operational architecture of public safety technologies must enforce an absolute, infrastructural separation between **Evidence Preservation** and **Narrative Interpretation**.

SYSTEM ARCHITECTURE GOVERNANCE FRAMEWORK

Sovereign Preservation Stack

- **Verbatim Automation:** Word-for-word transcript mapping with immutable audio sync tracking.
- **Linear Metadata:** Strict hashing and auditable timeline markers.
- **Zero Cognitive Drift:** No generative text framing, summarization, or synthesis.
- **Court Status:** Transparent, provable, and fully defensible baseline evidence.

Generative Interpretation Stack

- **Probabilistic Synthetics:** Automated summaries and chronologies built on predictive sequence math.
- **Opaque Vector Weights:** Black-box semantic filtering masking text generation logs.
- **Automation Bias Risk:** High cognitive overwriting and memory alignment vulnerabilities.
- **Court Status:** High disclosure liabilities and structural trial asymmetry.

By enforcing this clear technical boundary, public safety infrastructures ensure that data processing remains restricted to objective recording and translation. The human investigator remains the sole interpretive author, keeping the cognitive ownership of investigative synthesis uncompromised and fully challengeable in open court.

11. Provenance, Human Authorship, and the Future of AI-Assisted Policing

Preserving the institutional legitimacy of AI-assisted policing requires a definitive shift from passive oversight to active provenance tracking and technological isolation.

11.1 Technical Safeguards for Future Governance

To ensure that machine-mediated investigative environments can withstand constitutional and adversarial challenges, governance models must implement the following structural mandates:

- **Mandatory Redline/Diff Tracking:** Any police narrative produced with AI assistance must be submitted to the court with an immutable, automated "track changes" view, explicitly highlighting what text was generated by the model versus what was added, deleted, or modified by the human author.
- **System-Level Prompt Logging:** All prompt templates, system instructions, retrieval-augmented generation (RAG) contexts, and model parameters must be treated as formal components of the investigative file, preserved immutably for disclosure.
- **Version History Preservation:** Systems must log and preserve every iterative draft generated by the model during the report-writing process, preventing the erasure of alternative narrative interpretations.

- **Mandatory Cognitive Attestation:** Officers must provide a specific, reviewable attestation confirming which sections of the final report reflect independent, refreshed recollection versus sections adopted directly from machine-synthesized text.

12. Conclusion

The integration of generative AI into police reporting represents a profound technological transformation that challenges the core asset of criminal adjudication: human accountability. When a machine assumes responsibility for organizing chronology, prioritizing relevance, and structuring narrative flow, the resulting report ceases to function as an authentic contemporaneous record of direct human cognition. It shifts definitively into a post-event synthesized reconstruction, mediated by an unauditably algorithm operating beyond judicial discovery.

This structural evolution fundamentally alters the operational dynamics of democratic justice. A polished, internally consistent machine narrative masks the baseline situational uncertainty of a chaotic event, processing away the organic fractures and cognitive gaps through which courts evaluate reasonable doubt. By substituting human imperfection with artificial narrative clarity, generative systems subtly alter the adjudicative balance, shielding state evidence from effective adversarial testing before cross-examination even begins.

The future legitimacy of public safety technologies must therefore rest upon an absolute baseline rule: verbatim capture systems preserve source evidence; generative narrative frameworks simulate interpretation. To safeguard procedural fairness, protect judicial discovery transparency, and preserve the deep democratic principles underlying criminal adjudication, justice systems must enforce a strict institutional boundary against the cognitive outsourcing of narrative formation. Efficiency must never be bought at the expense of systemic legitimacy. The integrity of the rule of law demands that the deployment of coercive state power remains traceably, transparently, and unalterably human.

References and Legal Authorities

Jurisprudential Authorities

Brady v. Maryland, 373 U.S. 83 (1963).

Giglio v. United States, 405 U.S. 150 (1972).

R. v. Fliss, 2002 SCC 16, [2002] 1 S.C.R. 535.

R. v. Lifchus, [1997] 3 S.C.R. 320.

R. v. Meddoui, 1990 CanLII 2592 (ABCA), 61 C.C.C. (3d) 345.

R. v. McCarroll, 2008 ONCA 715, 238 C.C.C. (3d) 404.

R. v. Starr, 2000 SCC 40, [2000] 2 S.C.R. 144.

R. v. Stinchcombe, [1991] 3 S.C.R. 326.

R. v. Wilks, 2005 MBCA 99, 201 C.C.C. (3d) 11.

Woolmington v. DPP, [1935] A.C. 462 (H.L.).

United States Federal Rule of Evidence 612 (Writing Used to Refresh a Witness's Memory).

United States Federal Rule of Evidence 803(5) (Recorded Recollection).

Academic Literature & Institutional Frameworks

Blaskovits, Brittany, and Craig Bennell. "Exploring the Potential Impact of Body Worn Cameras on Memory in Officer-Involved Critical Incidents: A Literature Review." *Journal of Police and Criminal Psychology* 35 (2020): 251–262.

Dawes, D., W. Heegaard, M. Brave, G. Paetow, B. Weston, and J. Ho. "Body-Worn Cameras Improve Law Enforcement Officer Report Writing Accuracy." *Journal of Law Enforcement* 4, no. 6 (2015): 1–21.

European Union. *Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)*, Article 14 (Human Oversight).

Ferguson, Andrew Guthrie. "Generative Suspicion and the Risks of AI-Assisted Police Reports." *Northwestern University Law Review* 120, no. 2 (2025).

Loftus, Elizabeth F. *Eyewitness Testimony*. Cambridge, MA: Harvard University Press, 1996.

Loftus, Elizabeth F., and John C. Palmer. "Reconstruction of Automobile Destruction: An Example of the Interaction Between Language and Memory." *Journal of Verbal Learning and Verbal Behavior* 13 (1974): 585–589.

National Police Chiefs' Council. *Artificial Intelligence (AI) Strategy*. London: NPCC, 2025.

U.S. Department of Justice, Office of Community Oriented Policing Services. "Using AI to Write Police Reports." *Community Policing Dispatch* 18, no. 1 (January 2025).

White, Michael D. *Police Officer Body-Worn Cameras: Assessing the Evidence*. Washington, DC: U.S. Department of Justice, Office of Justice Programs, 2014.

Wigmore, John Henry. *Evidence in Trials at Common Law*. Revised by James H. Chadbourn. Boston: Little, Brown & Co.