Verification as the Standard: Ethical Use of Artificial Intelligence in Journalism

Atlanta Press Club
Innovation in Journalism Award Submission

Matthew Terrell
Freelance Journalist and Assistant Professor of Communication & Media at Kennesaw State University

Ethical AI use in journalism centers on human accountability.

Journalists remain responsible for verifying every factual claim before publication.

PURPOSE AND SCOPE

Newsrooms are shrinking while the amount of information reporters are expected to evaluate keeps expanding. At the same time, journalists are being asked to work faster, publish more frequently, and still maintain the same standards of accuracy and fairness. Because of that pressure, many reporters are beginning to experiment with new tools—including generative artificial intelligence—and trying to figure out where those tools fit inside real reporting work.

This paper comes from a practical question, not a theoretical one: Can a working journalist use AI without lowering reporting standards?

To test that, I used AI while producing a feature article for the *Atlanta Journal-Constitution* about the history and modern role of mulberry trees in Georgia. The article required interviews, archival research, historical synthesis, and narrative storytelling. Rather than speculate about AI in the abstract, I documented exactly how I used it, what it helped with, where it failed, and what safeguards were necessary to keep the reporting accurate.

The purpose of this paper is not to claim AI is better than traditional reporting or that it should replace journalists. Instead, it shows what actually happened when AI was introduced into a real reporting workflow under human control. Every stage of the process—research, drafting, verification, editing, and post-publication review—was examined to determine whether the final work remained factually sound and ethically defensible.

By explaining the workflow step-by-step, this paper offers transparency about how the article was produced, identifies the risks the technology introduced, and outlines the practical rules that made its use responsible.

DOCUMENTATION RECORD

"Atlanta's mulberry trees went from colonial experiment to forager fave", as published in the *Atlanta Journal-Constitution*, Is reproduced in Appendix A

EXECUTIVE SUMMARY

Journalism Is Verification, Not Just Writing—
and AI Changes How We Account for Intellectual Labor

Journalism is often described as writing. That's understandable—writing is the part the public sees. But writing is not the profession. Writing is the output.

Journalism is the process behind the output.

A news story requires a reporter to determine what is true, what is uncertain, what is credible, what is misleading, and what information is fair to publish. That involves interviews, document review, contextual judgment, and ethical decisions about harm and accuracy. The article is the final step of that process, not the process itself.

This distinction matters when discussing artificial intelligence. AI systems produce language. They do not verify reality.

AI can help organize notes, summarize research, suggest structure, or assist with drafting. What it cannot do is determine whether a claim is accurate, whether a source is reliable, whether a quote is fairly represented, or whether publication is responsible. Those are journalistic judgments, and they require a human reporter who is accountable for the result.

The central finding of this project is straightforward: AI can be used ethically in reporting only when the journalist retains full control over verification. The moment a reporter relies on the system to decide what is true, the work stops being journalism and becomes unverified content.

CORE POSITION

This paper does not argue that AI writes journalism. It argues the opposite. Writing and journalism are related but different kinds of work.

- **Writing is production work.** It involves composing sentences, organizing paragraphs, and shaping presentation. Tools can assist this.
- **Journalism is knowledge work.** It asks: How do we know this? Who says so? What evidence supports it? What context is missing? Who could be harmed by publication? That responsibility cannot be delegated to a language model.

Professional reporting standards emphasize accountability for accuracy before publication. The long-standing understanding inside the profession is that journalism is a discipline of verification. This project applies that principle to AI: the system may assist with drafting, but it is never treated as a source and never trusted to establish facts.

WHAT THIS PROJECT DID

Using the *Atlanta Journal-Constitution* mulberry feature as a case study, this paper documents an AI-assisted reporting workflow built around human control:

1. **Reporting first.** Interviews, field observations, and historical research were completed before AI was used for drafting.
2. **Ethical constraints.** The process followed the Ethical Wheel of Prompting: read inputs, read outputs, evaluate for accuracy and relevance, and edit before use.
3. **Traceable sourcing.** Every factual claim was mapped back to an interview, document, or historical reference. The audit trail — not the AI text — is the journalism.
4. **Failure awareness.** The paper identifies predictable error patterns (hallucination, misattribution, confident falsehoods, and citation mimicry) and treats them as normal behavior the reporter must guard against.

KEY FINDINGS

1. **AI helps after reporting exists**
   AI works best once a journalist already understands the material. It can organize notes, structure sections, draft transitions, and compare revisions. Used this way, it reduces repetitive labor and frees time for interviews, research, and editing.

2. **AI is not a truth system**
   Large language models produce plausible language, not verified facts. They generate confident errors, invented details, and false citations. These are not rare glitches — they are predictable behaviors.

3. **Verification is the bright line**
   The workable model is simple: AI assists thinking. Humans verify. Humans remain accountable.

4. **Journalism labor is actually two different kinds of labor**
   AI forces a distinction the profession has long blurred—

Journalism (knowledge work):
sourcing, interviewing, verification, credibility judgment, fairness decisions, and accountability.

Writing (composition work):
outlining, drafting, editing, restructuring, and stylistic refinement.
AI can assist the second category.

PROJECT OVERVIEW

This project was undertaken independently by the author and was not commissioned or directed by Kennesaw State University, the *Atlanta Journal-Constitution*, or any other organization. It grew out of a practical question encountered during ordinary reporting: how should a working journalist responsibly use generative AI tools that are quickly becoming part of everyday knowledge work?

I have worked as a writer and journalist across local, national, and international publications, including the *Atlanta Journal-Constitution, Huffington Post, San Francisco Weekly, Slate, Mic,* and *VICE*. My academic training—an MFA and BFA in writing from the Savannah College of Art and Design and an MA in Communication from Georgia State University—provided grounding in research methods, narrative construction, and journalistic ethics. This project applies those skills to a new reporting environment rather than treating AI as an abstract technological debate.

Instead of theorizing, the project documents a real reporting case—a feature article on the history and modern significance of mulberries in Georgia. The story required far more than a short lifestyle piece. Producing it involved historical research, interviews, scientific context, archival material, and narrative structure similar to long-form magazine reporting.

The reporting included:

- historical research into 18th- and 19th-century agricultural development in Georgia
- botanical and horticultural background on mulberry species
- interviews with local residents, foragers, and nonprofit organizers
- synthesis of archival and contemporary sources into a single narrative

AI was not used to replace any of this reporting. Interviews were conducted by the author, sources were located and evaluated by the author, and all factual claims were verified by the author. AI was used only after research existed, primarily to help organize notes, compare drafts, and assist with drafting and revision.

This project therefore treats AI not as an automated reporting system, but as a support tool operating inside a traditional reporting workflow.

Journalism has repeatedly adapted to new tools. Reporters moved from typewriters to word processors, from physical archives to searchable databases, and from film photography to digital editing. Each change raised ethical concerns at the time, but the core responsibility of the journalist—verify information and present it fairly—remained unchanged.

Generative AI presents a similar moment. The technology does not determine what is true. It predicts language patterns. The reporter still decides what sources are credible, what facts are supported, and what belongs in print. Used carefully, AI can reduce repetitive labor such as organizing research or restructuring drafts, allowing more time for interviewing, verification, and editorial judgment.

This white paper documents that process in detail: how the workflow was structured, what worked, what failed, and where ethical boundaries must remain firm.

ETHICAL FRAMEWORK:
USING THE ETHICAL WHEEL OF PROMPTING

Using generative AI responsibly requires more than technical skill. It requires a set of habits that keep the reporter, not the software, in control of what is published. For this project, my workflow was guided by the **Ethical Wheel of Prompting (EWP)**, a practical decision-making framework developed by Jeanne Beatrix Law, Ph.D., which I first encountered at a Kennesaw State University research summit.

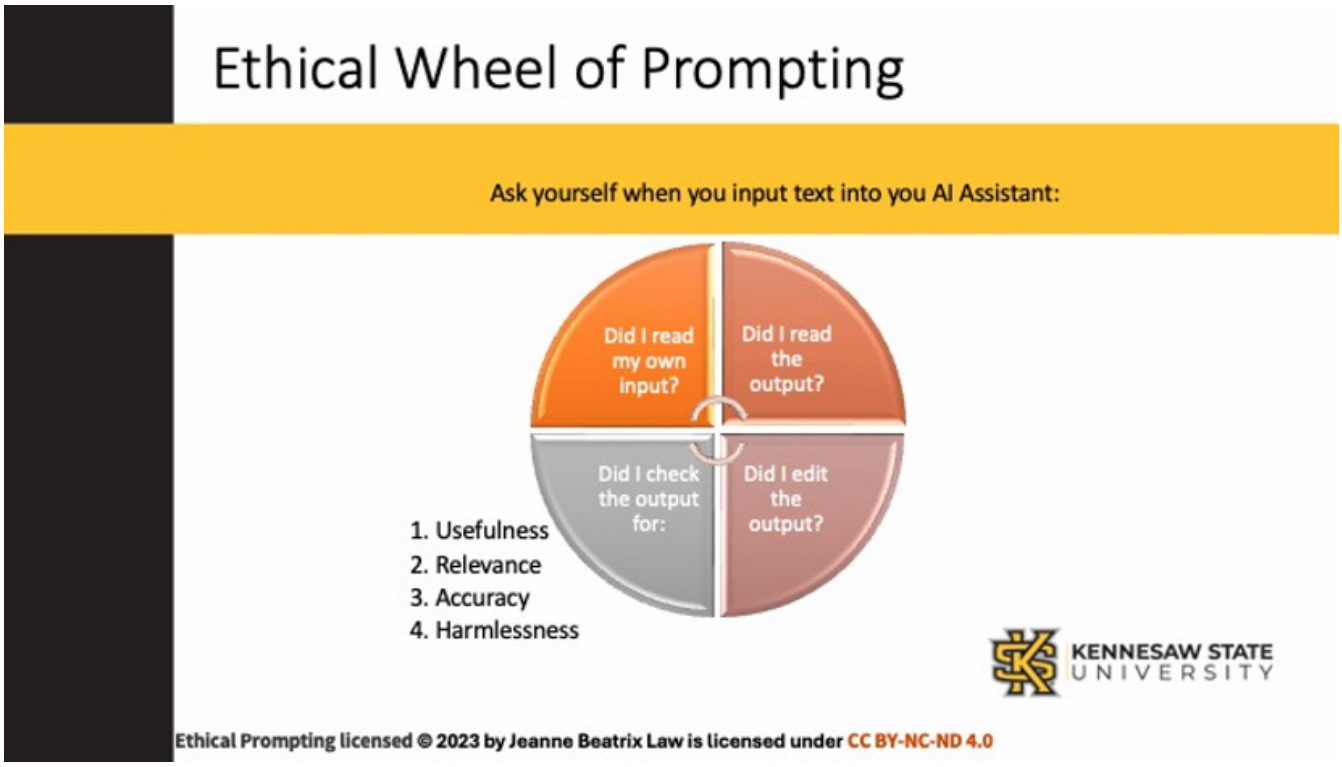

The value of the framework is not theoretical. It functions as a checklist for how a person should behave when working with AI. Instead of treating AI output as information, the model forces the user to treat it as a draft that must be evaluated before it enters a journalistic workflow.

The EWP centers on four rules:

1. **Know your material before prompting.**
   I read and understood interviews, research, and notes before asking the system to process them. AI was not used to interpret sources I had not already examined myself.
2. **Read every output carefully.**
   Fluency is not accuracy. AI text often sounds authoritative even when it is incomplete or wrong. Every response was reviewed line by line before it could influence the article.
3. **Evaluate for accuracy, relevance, and harm.**
   Output was checked against original sources and judged the same way any draft would be judged in an editorial setting: Is it correct? Is it fair? Does it mislead? Could it cause harm if published?
4. **Rewrite before use.**
   No AI-generated text was treated as final copy. Any usable material was edited, rewritten, or discarded so that the final article reflected verified information and my own editorial judgment.

In practice, the framework prevented a common mistake in AI use: confusing polished language with verified fact. The system could help organize notes, compare drafts, or suggest structure, but it could not determine what was true. Verification, attribution, and fairness remained human decisions.

The result is important. AI did not function as a source, a reporter, or an editor. It functioned as a drafting and organizational aid inside a traditional reporting process. Interviews were conducted by me, sources were evaluated by me, and every factual claim was verified by me before publication.

Put simply: the Ethical Wheel of Prompting does not make the reporting ethical—it ensures that the journalist remains responsible for it.

WORKFLOW METHODOLOGY

This section explains how AI was actually used in producing the mulberry feature—and what it was not allowed to do. The key boundary was simple: reporting and verification were human tasks; AI was used only for organization and drafting support.

1. **Reporting Came First**
   All reporting was completed before AI was introduced into the writing process. This included—

   Interviews
   - A Decatur family actively foraging mulberries
   - The executive director of the nonprofit Concrete Jungle, which maps and harvests fruit from trees across Atlanta

   Research
   - Colonial Georgia agricultural records (Trustee Garden)
   - Botanical and horticultural sources on mulberry cultivars
   - Historical accounts of the 19th-century silk and mulberry speculation boom

   This material formed the factual record of the article. None of it was generated by AI.

2. **Source Material Was Reviewed Before Prompting**
   AI was not used to discover facts. It was used only after I had already read and understood the source material.

   Before prompting the system, I:
   - read the sources,
   - took notes,
   - selected relevant excerpts when necessary,
   - and confirmed their relevance.

   No text was pasted into prompts without first being reviewed. This prevented the system from interpreting material I did not understand myself.

3. **AI Was Used for Organization, Not Authority**
   AI was used as a drafting and structural assistant. Typical uses included:
   - outlining sections,
   - organizing historical sequences,
   - suggesting transitions,
   - compressing notes,
   - and identifying gaps in coverage.

   It was never treated as a reporting source. Any factual statement suggested by the system was treated as unverified until checked independently.

4. **Narrative Decisions Remained Human**
   Feature writing requires interpretation: deciding chronology, emphasis, and meaning. Those decisions were not delegated.

   AI could suggest structure, but I determined:
   - story arc,
   - thematic connections,
   - tone,
   - and significance.

   The system assisted with wording. It did not decide what the story meant.

5. **Verification Was Entirely Human**
   Every factual claim was confirmed against original sources.
   - Interview material was checked against recordings and notes.
   - Historical claims were verified through archives and reference works.
   - Contextual claims were checked against primary or reputable secondary sources.

   AI could propose language, but it could not confirm accuracy. Verification remained my responsibility.

6. **What AI Actually Replaced**
   AI did not replace reporting. It replaced mechanical writing labor.

   Specifically, it reduced time spent on:
   - reorganizing drafts,
   - rewriting transitions,
   - summarizing notes,
   - and comparing alternate versions.

   The practical result was more time spent reporting and fact-checking rather than formatting and restructuring text.

WORKING RULE

The workflow followed a consistent principle:
The human gathers and verifies the information.
The AI helps arrange the words.
The final article remained a reported piece because every fact was sourced and verified independently of the system.

WHAT ARTIFICIAL INTELLIGENCE ACTUALLY DOES WELL IN JOURNALISM

When Used Inside a Verified Human Workflow

Artificial intelligence is not a reporter. It cannot witness events, judge credibility, or verify facts.

What it can do is assist the work that happens **after reporting.**

The simplest way to understand its value: **AI helps handle language. Journalists handle reality.**

When used this way, it becomes a practical newsroom tool rather than a substitute for reporting.

Below are the areas where AI proved genuinely useful in producing the mulberry feature—

1. **Organizing Complex Material**
   AI is especially strong at structuring large amounts of information.
   It can:
   - group related facts,
   - outline sections,
   - sequence historical timelines,
   - and suggest narrative order.

   In the mulberry story, it helped arrange a 200+-year arc—from pandemic backyard foraging to colonial silk experiments to modern urban ecology—into a readable structure. However, deciding why those events mattered remained a human task.

   AI could organize the material, but it could not interpret significance.

2. **Summarizing and Navigating Sources**
   AI worked effectively as a navigation tool for research I had already gathered.
   It helped:
   - compress notes,
   - identify key themes,
   - and point me back to relevant passages in long documents.

   But every referenced claim was checked in the original source.

   Used properly, AI does not replace research—it helps you move through it faster.

   The rule was simple: If AI mentioned a fact, I located it myself before using it.

3. **Interview Preparation**
   Before interviews, AI was useful for preparation.
   It helped:
   - generate question lists,
   - identify knowledge gaps,
   - and clarify unfamiliar terminology.

   This did not replace interviews. It improved them by allowing me to arrive better informed and ask more precise follow-up questions.

4. **Editing and Drafting Support**
   AI functioned most reliably as a drafting assistant.
   It helped:
   - smooth transitions,
   - clarify unclear paragraphs,
   - compare alternate phrasings,
   - and identify structural gaps.

   In the mulberry article, it was particularly helpful when moving between contemporary reporting and historical sections, where readers needed clear signals of time and context. The reporting still came from my sources.

   AI only helped refine how the material was expressed.

5. **Research Direction and Story Planning**
   AI was useful in identifying where information might exist.
   It suggested:
   - search terms,
   - related topics,
   - and potential archives or experts.

   However, it was never treated as an archive itself. It pointed toward information; it did not provide evidence.

CENTRAL PRINCIPLE

AI accelerates thinking work, but it does not perform journalism.

Journalism requires verification and accountability.

The person who verifies the facts is the author.

In this project, that responsibility remained entirely human.

HOW ARTIFICIAL INTELLIGENCE MISLEADS JOURNALISTS

A Field Guide to Hallucination, Error, and False Authority

Large language models don't lie. They simulate.

They generate sentences that look right—not sentences that are right.

This core reality introduces serious factual risks to journalism, a field grounded in verification, not linguistic plausibility. When journalists mistake fluency for fact, errors enter the record. Worse still, AI can short-circuit the journalist's own critical thought process—offering complete-seeming answers that halt further investigation.

This section outlines 10 distinct error patterns—each grounded in a different mechanism—that journalists must learn to recognize.

1. **Fabricated Specifics (Classic Hallucination)**

   Definition:
   AI invents concrete facts: fake people, studies, quotes, statistics, dates.

   Why it happens:
   The model is rewarded for completion and plausibility. It fills in gaps with statistically likely-seeming answers, not verified ones.

   Mulberry Risk Example:
   Artificial Intelligence might claim that "a 1913 USDA study found mulberries were the most drought-resistant fruit in Georgia," when no such study exists. Or it might fabricate a quote from Amelia Trace Lerner about child hunger rates, based on similar-sounding interviews.

2. **Implied Authorship / Misattributed Sources**

   Definition:
   Real information, wrong attribution.

   Why it happens:
   The model remembers what was said, but not who said it. It predicts attribution based on probability, not accuracy.

   Mulberry Risk Example:
   AI could claim a passage on mulberries came from NPR when it was actually a quote from *The Georgia Review.* Or it might assign a botanical fact stated by Concrete Jungle staff to a completely unrelated historian.

3. **Composite Truth (The Plausible Article Problem)**

Definition:
Multiple real facts are stitched into a single narrative that never existed as such.

Why it happens:
AI compresses similar source material into a smooth, synthesized summary—but the resulting "article" may not exist anywhere.

Mulberry Risk Example:
You might get a section that claims: "During the Dust Bowl, Georgia farmers relied on mulberries for soil restoration and children's nutrition, as documented in state archives and WPA reports."

Each part sounds possible—but that narrative is entirely AI-generated.

4. **Authority Tone Bias**

Definition:
Uncertain or unverified information is presented in confident, structured prose.

Why it happens:
Language models are trained to imitate confident rhetorical style. They can't signal doubt the way human writers can.

Mulberry Risk Example:
A paragraph might say:
"Concrete Jungle has distributed over 2 million pounds of mulberries since 2010, proving their central role in Atlanta's food equity movement."

Sounds official—but is unverifiable and inflated. The authoritative cadence tricks the journalist into lowering their guard.

5. **The Verification Collapse Effect**

Definition:
AI shortcuts the journalistic process by providing a full answer—removing the friction that usually triggers deeper fact-checking.

Why it's dangerous:
When research yields contradictions or gaps, journalists investigate. But if AI fills those gaps too cleanly, no verification is triggered.

Mulberry Risk Example:
Instead of researching the economic lifespan of a mulberry tree, a journalist asks AI—which confidently outputs "Mulberry trees produce peak fruit for 45 years," short-circuiting the need to verify through agricultural extension offices or botany sources.

6. **Contextual Drift**

Definition:
Correct facts placed in the wrong context.

Why it happens:
AI merges patterns of similar topics. It knows the fact, but not its timeline, actor, or application.

Mulberry Risk Example:
It might state that mulberries were a staple food in the Civil Rights-era Black Panther food programs—confusing their use with more recent nonprofit work like Concrete Jungle, which harvests mulberries in Atlanta today.

7. **Temporal Errors (The Silent Hallucination)**

Definition:
Time collapses. Events are sequenced incorrectly or causality is implied falsely.

Why it happens:
LLMs have weak models of chronology. They assume "typical" time patterns.

Mulberry Risk Example:
AI could claim that Thomas Jefferson planted mulberries in response to a silk embargo, when in fact his experiments with mulberry cultivation predated the Embargo of 1807—an error that shifts historical causation.

8. **Citation Mimicry**

Definition:
AI generates sentence structures that look sourced—but no traceable citation exists.

Why it happens:
LLMs are trained on academic and journalistic tone. They learn what citations sound like—not how to retrieve or verify them.

Mulberry Risk Example:
A phrase like "According to a 2004 CDC study on fruit access in Atlanta…" may appear—but no such CDC study exists. You search and find nothing. The sentence simply mimicked citation form.

9. **Memory Anchoring**

Definition:
Even after you verify that an AI claim is false, your brain still retains it as "plausible."

Why it matters:
You may unconsciously shape interview questions, angles, or tone around falsehoods you've mentally "stickied."

Mulberry Risk Example:
Let's say AI claims that mulberry leaves were used in Georgia moonshine production. You verify that it's false. But now the detail feels colorful—and you still ask your source: "Did you ever hear about mulberry leaves being used that way?" You've already shaped your frame.

10. Confirmation Engine Effect

Definition:
When a journalist prompts with a leading question, AI supports the implied belief.

Why it happens:
The model interprets user belief as a signal and tries to complete the pattern.

Mulberry Risk Example:
Prompt: "Are mulberries the most overlooked superfood in Georgia history?"

AI response: "Yes, mulberries are a forgotten powerhouse… widely used in Cherokee medicine and early Southern diets."

It supports your frame without interrogating it—reinforcing bad hypotheses.

CONCLUSION AND PROFESSIONAL RECOMMENDATIONS

Journalism, Artificial Intelligence, and the Future of Verification Work

This project began with a practical question: Can a journalist use generative artificial intelligence without weakening the integrity of the published work? The answer is yes—but only under one condition: AI may assist writing. It may not replace journalism.

Journalism is not the production of text. It is a process for verifying reality, interpreting what that verified information means, and presenting it responsibly to the public. Writing is the visible product of that process, but underneath it are reporting, sourcing, verification, editorial judgment, and accountability. Readers do not trust journalists because they produce sentences. Readers trust journalists because they produce defensible claims about the world.

Generative AI operates differently. It produces language, not knowledge. It predicts plausible phrasing, not truth. It does not witness events, evaluate credibility, or bear responsibility for publication. Any factual accuracy in an AI output exists only because a human reporter verifies it. Without verification, an AI-assisted article is not journalism—it is merely writing.

When the journalist retains control over sourcing, verification, and interpretation—and uses AI only to assist organization or drafting—the integrity of the work can remain intact. When those responsibilities are transferred to the machine, journalism stops occurring.

Journalistic credibility depends on traceability: every claim must be connectable to a human source, document, observation, or interview. The evidence audit in this project appendix demonstrates that an AI-assisted article can meet that standard when verification work is performed and documented by the reporter.

The boundary is clear: AI participation in writing is acceptable. AI participation in verification is not.

REFRAMING INTELLECTUAL LABOR IN JOURNALISM

For generations, authorship in journalism has been associated with writing. This study suggests the profession has misidentified its core activity. Writing is not the defining labor of journalism. Verification is.

A reporter's expertise is demonstrated by the ability to:

- identify credible sources
- recognize unreliable claims
- understand context
- interpret meaning responsibly
- and accept accountability for publication

AI can help generate prose. It cannot assume responsibility.

Once this distinction is recognized, the conversation shifts. The relevant question is no longer "AI or no AI." The relevant question becomes whether verification remains human-controlled.

RECOMMENDATIONS FOR NEWSROOMS AND JOURNALISTS

1. **Build Policies Around Verification, Not Software**
   Avoid both blanket bans and unrestricted use.
   A workable standard:
   - AI may assist drafting and organization.
   - AI may not be treated as a reporting source.
   - No factual claim may appear in publication without independent verification.

   Policies should regulate responsibility, not tools.

2. **Require Traceable Sources**
   Every factual statement in an AI-assisted article should connect to:
   - a named human source,
   - a documented record,
   - a primary historical source,
   - or direct observation.

   AI output is never evidence. AI is a tool, not a witness.

3. **Adopt Internal Verification Audits**
   Reporters should maintain a simple source audit documenting where claims originate. This need not be public disclosure. It functions like a fact-checking file long used in investigative and magazine journalism.

   This preserves trust while allowing responsible technological adoption.

4. **Provide AI Literacy Training**
   Journalists do not need to become engineers, but they must understand predictable failure modes:
   - hallucinated facts
   - false citations
   - overconfident tone
   - probabilistic output

   Without training, reporters may overtrust the tool. With training, they can use it critically.

5. **Editors Should Evaluate Process**
   Editors should assess whether reporting occurred, not whether a drafting assistant was used.
   The central editorial question should be:

   "Can this article be defended?"

   Not

   "Was AI involved?"

6. **Encourage Open Discussion in Newsrooms**

The greatest risk is not AI use. The greatest risk is hidden AI use.

Many journalists currently fear that discussing AI openly could result in discipline or termination. That environment discourages responsible practice and removes editorial oversight.

Newsrooms should encourage:

- non-punitive disclosure
- shared experimentation
- collaborative policy formation
- transparent discussion

Ethical adoption requires psychological safety.

7. **Replace Conflict With Professional Dialogue**

Public debate often frames AI as either a replacement for journalists or an existential threat. Neither framing is useful.

AI is not a journalist.

But journalists will increasingly work alongside AI-assisted tools.

Journalism has survived photography, radio, television, databases, and the internet by strengthening its core principles rather than rejecting new tools. The same path exists here.

APPENDIX MATERIALS

The following appendices provide the full documentation supporting the methods, claims, and conclusions discussed in this white paper. Because the central argument of this project concerns verification and accountability, the supporting materials are included so readers may independently examine the reporting process, source base, and evidence-based traceability of the case study article. Rather than summarize sources secondhand, this document makes them available for direct review.

APPENDIX A—VERIFICATION RECORD AND SOURCE DOCUMENTATION

Link: https://tinyurl.com/Terrell-Appendix-A

Appendix A contains the complete evidence record for the *Atlanta Journal-Constitution* mulberry feature used as the case study throughout this paper. Together, these materials allow readers to follow the reporting process from published article back to original sources.

Included Materials:

A1. Published Article
The original article as it appeared in the *Atlanta Journal-Constitution* online edition.

A2. Evidence Audit (Line-by-Line Claim Traceability)
A sentence-level audit of the article. Each assertion in the published piece is identified, categorized by type (e.g., interview statement, historical fact, contextual claim, descriptive observation), and matched to its originating source.

This audit demonstrates how every factual claim can be traced to a human interview, documented reference, or direct observation.

A3. Primary Interviews

- Transcript of interview with Amelia Trace Lerner
- Transcript of interview with Katherine Kennedy, Executive Director of Concrete Jungle

A4. Historical and Reference Sources

- Trustees Garden (Georgia Encyclopedia)
- White Mulberry Tree (Georgia Encyclopedia)
- The National Mulberry Craze (JSTOR Daily)
- Amy Chambliss, 1960 article on the history of mulberries in Georgia (The Georgia Review)

These materials allow independent verification of the reporting and provide the documentation underlying the evidence audit.

Appendix B—Extended Analysis and Supplemental Discussions

Link: https://tinyurl.com/Terrell-Appendix-B

Appendix B includes supporting discussions and operational considerations that extend beyond the core argument of the white paper but may be useful for journalists, editors, and organizations considering practical implementation.

B1. Practicalities of Using AI in a Reporting Workflow
A practical overview of how AI can be incorporated into reporting and writing processes, including preparation, drafting, revision, and verification practices.

B2. When AI Fails at Simple Tasks: Lessons from a Workflow Breakdown
A case study examining failures encountered during document formatting and pagination. This section discusses operational reliability and why generative AI should not be treated as a production or layout system.

B3. Limitations of AI in Journalism: Ethical Filtering and Reporting Gaps
An examination of how safety filtering and content moderation systems may affect reporting, particularly when journalists cover sensitive, controversial, or traumatic subjects.

B4. AI in Journalism: Practical Limitations and Emerging Questions
A collection of remaining open questions, unresolved issues, and professional considerations surrounding the long-term integration of AI tools into journalistic practice.