FAIR AND JUST PROSECUTION

AI-GENERATED POLICE REPORTS:

HIGH-TECH, LOW ACCURACY, BIG RISKS

JUNE 2025

ABOUT FJP & ACKNOWLEDGEMENTS

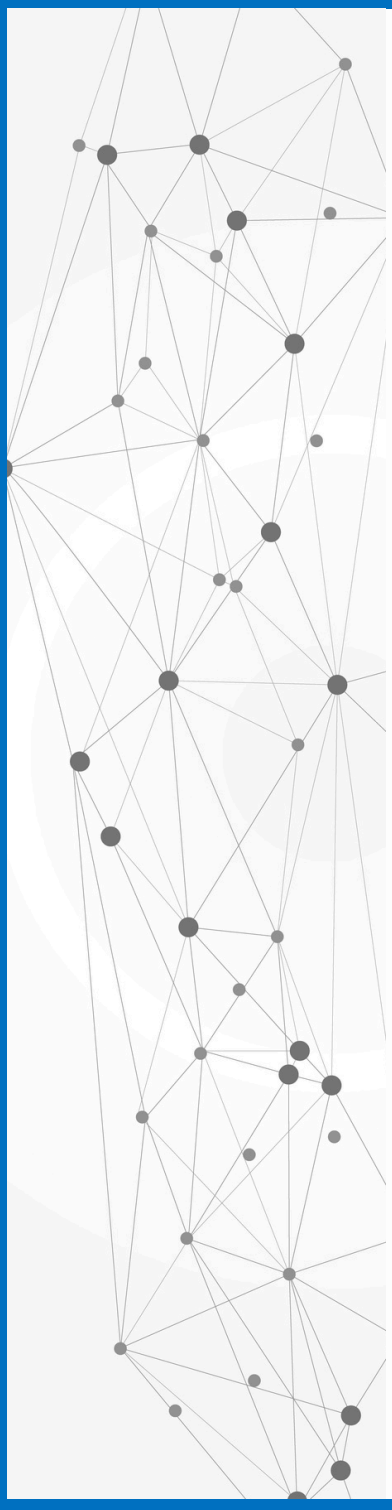

Fair and Just Prosecution (FJP) brings together elected district attorneys as part of a network of like-minded leaders committed to change and innovation. FJP hopes to enable a new generation of prosecutive leaders to learn from best practices, respected experts, and innovative approaches aimed at promoting a justice system grounded in fairness, equity, compassion, and fiscal responsibility. In furtherance of those efforts, FJP's "Issues at a Glance" briefs provide district attorneys with information and insights about a variety of critical and timely topics. These papers give an overview of the issue, key background information, ideas on where and how this issue arises, and specific recommendations to consider. They are intended to be succinct and to provide district attorneys with enough information to evaluate whether they want to pursue further action within their office. For each topic, Fair and Just Prosecution has additional supporting materials, including model policies and guidelines, key academic papers, and other research. If your office wants to learn more about this topic, we encourage you to contact us.

Thanks to the many people who contributed to this "Issues at a Glance" brief, including members of FJP's team who provided guidance, input, research, and assistance: Rebecca Blair, Amy Fettig, Monica Fuhrmann, Jennifer Krantz, and lead author Christina Green. FJP is also grateful to Emily Bender, Andrew Ferguson, Kyle Fox, Brandon Garrett, and Katie Kinsey for their invaluable comments, insights, and input.

TABLE OF CONTENTS

SUMMARY

This FJP “Issues at a Glance” brief provides an overview of the concerns related to the use of artificial intelligence (AI) tools to generate police reports.

While AI-generated police report technology is marketed as a solution to reduce administrative burdens and improve report accuracy, these tools pose significant risks to both individual rights and public safety systems outcomes that must be considered before law enforcement agencies adopt them.

This brief will explore many of the emerging issues related to the use of AI-generated police reports, including their susceptibility to inaccuracy and bias, privacy and data ownership concerns, impacts on police accountability and public trust in law enforcement, and the serious legal issues these tools raise. District attorneys are encouraged to consider these factors in determining their own policies and practices around the use of these reports, and to prepare their offices to navigate the novel legal issues such reports may evoke.

BACKGROUND

AI-generated police reporting tools rely on natural language processing and machine learning technology to transform recorded audio into draft narratives. Typically, audio from body-worn cameras or other recording devices is uploaded to a system where speech recognition software converts it into text. Then, pre-trained language models use this transcription as the foundation to generate a narrative resembling a traditional police report. The draft report is organized using algorithms that identify and categorize key details on subjects such as the individuals involved, the actions described, and the incident's context. Notably, these algorithms are proprietary to the vendor and remain inaccessible to law enforcement and other stakeholders in the criminal legal system. Officers are then tasked with reviewing these drafts, making necessary edits, and verifying accuracy before final submission.[2]

Proponents of AI-generated reports argue this technology can revolutionize policing by automating one of its most labor-intensive tasks.[3] Officers often spend hours drafting reports after each incident, which takes them away from fieldwork and can contribute to burnout. Additionally, AI could, in theory, help reduce inconsistencies in reports by mitigating human errors, like accidental omissions of minor details or delays in reporting after incidents. However, while these potential benefits are appealing, integrating AI into police reporting triggers a host of new challenges and complexities.

The effectiveness of these systems depends on the quality of the input data, the robustness of the algorithms, and the quality of oversight during the review process. Errors in transcription, misinterpretation of context, or the inability to distinguish overlapping voices can lead to inaccuracies that may compromise the integrity of police documentation, especially in the complex and high-stakes context of criminal investigations. Moreover, reliance on AI raises critical concerns about accountability and the erosion of human judgment in crafting reports that play a pivotal role in legal proceedings.

Despite the claims of companies like Axon, who assert their generative AI police reporting products are rigorously tested, the absence of independent evaluations raises serious concerns about the performance of these systems. Without external validation, it is impossible to determine how consistently these tools produce accurate outputs. Furthermore, without public disclosure of testing results, there is no way for stakeholders — whether they be defense attorneys, the general public, or even law enforcement officers — to assess the tool's effectiveness or understand the limitations of its output. Independent, transparent testing is crucial to ensure that AI systems used in police work do not undermine prosecutorial integrity, particularly in ways that may disproportionately affect vulnerable groups or lead to wrongful convictions.

In response to these and other concerns, in September of 2024, the King County Prosecuting Attorney's Office in Seattle became the first prosecutorial office in the country to publicly announce that it would not accept AI-generated reports from police agencies. Given the lack of oversight or regulation, there is currently no way to know how many police departments are actively using AI tools to assist in report generation or how other prosecutorial offices have responded to the introduction of such tools. However, as AI continues to proliferate rapidly across industries, more and more prosecutorial offices will undoubtedly have to confront complicated questions around the role of AI in their local justice systems.

CONCERNS REGARDING THE USE OF AI-GENERATED POLICE REPORTS

AI TOOLS ARE PRONE TO INACCURACY

The inability of generative AI tools to produce reliably accurate outputs is a primary reason these tools are currently unsuited for use in situations with such high stakes as criminal investigations. Generative AI tools have a demonstrated propensity to "hallucinate" made-up information, producing misleading or demonstrably false output.[4] This makes the tools susceptible to a range of errors, such as misidentifying the presence of a weapon, misrepresenting the actions or intentions of a suspect, or even introducing fictitious dialogue into a report. For instance, there have been cases in which AI-generated police reports referenced officers who were never present at the scene.[5]

Emily M. Bender, a linguistics professor and director of the University of Washington's Computational Linguistics Laboratory, argues this is not a fixable glitch but rather "inherent in the mismatch between the technology and the proposed use cases."[6] AI language models are designed to recognize patterns and mimic human speech; in its current form, the technology is not equipped to discern truth from fiction. Much like an autocomplete feature, it can fill in the blanks with plausible-sounding language but cannot guarantee accuracy.[7]

The quality of the output provided by AI police report generation systems is also highly contingent on the quality of the input provided; even minor problems with body camera footage, witness statements, and other inputted materials can lead to significant errors in the final report. Problems with input data that lead to inaccuracies include:

- **Poor-quality audio:** Among the most pressing challenges are background noise, muffled audio, and distorted speech, which can drastically reduce the model's effectiveness in transcribing and summarizing conversations accurately.[8] For example, the model might mis-transcribe key phrases or miss entire sections of dialogue in a recording taken in a noisy environment with sirens, crowds, or other environmental sounds. Since the model tends to fill gaps in the recording with plausible-sounding language, inaccurate or fabricated dialogue can be difficult to spot.

- **Overlapping voices:** AI technologies are often ineffective when the input data features multiple speakers, particularly when there are rapid exchanges, interruptions, or voices with similar tonal qualities. This limitation can lead to inaccuracies in attributing statements, potentially resulting in misrepresentations of key details or loss of conversational context.[9] For instance, if two officers or witnesses speak simultaneously, the model may combine their statements into one, creating confusion in the generated report.

- **Speech Bias:** Speech recognition models, while improving, still perform less accurately for certain speaker groups—including those with non-native accents, those with certain regional accents, children, and older speakers—making their outputs susceptible to bias for those with less common speech patterns.[10] This could lead to the omission or misinterpretation of critical testimony, affecting the report's reliability.[11]

- **Technical issues:** Poor recording quality, low microphone sensitivity, and data compression artifacts can compromise audio fidelity, causing even state-of-the-art models to fail in capturing nuances or specific terminology critical for legal or investigative accuracy.

- **Environmental factors:** Wind, heavy rain, or other environmental sounds captured on recordings can confuse the model, further reducing transcription accuracy. In emergency scenarios where clear communication is essential, these challenges can have profound implications for the completeness and reliability of the final police report.[12]

> *The increased use of AI presents completely new challenges for prosecutors to navigate. Prosecutors rely on police reports that are precise and accurate to make charging decisions. We have communicated our concerns about AI-generated reports to our law enforcement partners and have received assurances that AI will not be used to generate these critical materials.*
>
> **— Hennepin County Attorney Mary Moriarty**

CASE STUDY

FROM KING COUNTY PROSECUTING ATTORNEY'S OFFICE

The King County Prosecuting Attorney's Office (KCPAO) does not accept any police report narratives that have been produced with the assistance of artificial intelligence (AI) – police reports must be produced entirely by the authoring officer.

The KCPAO understands that staffing levels are extremely short in some departments, and there is a real need to free up as much time as possible for officers to be on the street. We do not fear advances in technology – but we do have legitimate concerns about some of the products on the market now.

In one example we have seen, an otherwise excellent report that used AI included a reference to an officer who was not even at the scene. This is one type of error that could easily go unnoticed by a reviewing officer given the volume of material required to be reviewed on deadline. And when an officer on the stand alleges that their report is accurate – they will be proven wrong.

The consequences would be devastating for that case, the community and the officer. Furthermore, it will subject them to potential impeachment ramifications and leave them without a way to establish that theirs was an error of oversight, and not an intentional falsehood.

AI continues to develop, and we are hopeful that we will reach a point in the near future where these reports can be relied on. For now, the King County Prosecuting Attorney's Office has made the decision not to accept any police narratives that were produced with the assistance of artificial intelligence.

THE USE OF AI RISKS DILUTING ACCOUNTABILITY IN LAW ENFORCEMENT

Police reports are not just bureaucratic formalities; they are official records that rely on first-hand observations and require accuracy in every detail. Writing these reports forces officers to reconstruct incidents step by step, engaging in critical thinking, attention to detail, and memory recall.[13] This active process can reinforce their accountability and sense of ownership over the final narrative.

Outsourcing this task to AI fundamentally changes the officer's role, reducing their involvement to reviewing and approving pre-generated text. This leads to cognitive offloading, a psychological phenomenon that occurs when individuals rely on technology to perform mental tasks, thereby conserving effort but reducing engagement and retention.[14] Research shows that people who offload tasks to technology engage less deeply with the material, which can lead to complacency in the face of overlooked or incomplete information.[15] While proponents argue that officers remain responsible for ensuring the accuracy of AI-generated reports during editing, replacing the active narrative-building process with the more passive task of reviewing AI-generated text may make officers less likely to fully process incidents or retain key details, weakening their capacity to ensure the overall integrity of the report.

AI also introduces new risks of shielding police misconduct. Systems designed to filter bias or refine language may inadvertently obscure details critical to identifying issues, such as profiling or excessive force. For example, an AI tool might sanitize descriptions of an incident or remove problematic dialogue that would otherwise draw the attention of officers or prosecutors, making it harder for oversight bodies to detect and address misconduct. This unintended shielding of wrongdoing could further undermine the accountability traditional report writing is meant to uphold.

AI TOOLS CAN REINFORCE BIASES

Although AI systems are often hailed as neutral and unbiased, AI models are fundamentally shaped by the data sets they are trained on, and many are trained on language and information that reflect deeply entrenched systemic biases. If the language models are trained to incorporate patterns of over-policing, racial profiling, or discriminatory enforcement practices—issues endemic in law enforcement—they will inevitably reproduce and even amplify these injustices.

Moreover, attempts to mitigate these biases through debiasing techniques introduce their own challenges. Debiasing often leads to more "hallucination" errors, whereby the AI generates outputs that are factually inaccurate or misleading. For example, Google's Gemini chatbot produced images of Black people in Nazi uniforms when asked to generate images of German soldiers from World War II.[16] By treating all racial groups interchangeably, the AI ignored historical context, leading to offensive and misleading outputs.

AI-generated police reports can also create feedback loops that further entrench bias. Once an AI-generated report is entered into the system, it can serve as training data for future reports, effectively institutionalizing and spreading biases. A report that unfairly portrays someone as suspicious based on biased data might be used as evidence that similar individuals are suspicious in future encounters. This loop can disproportionately affect marginalized communities, who are already more likely to be subject to biased policing practices.[17]

Adding to these concerns is the opaque nature of AI tools, often described as "black boxes." These systems operate without transparency, making it difficult to understand how conclusions are reached or to identify potential errors and biases embedded in their processes.[18] This lack of visibility means law enforcement agencies risk unknowingly basing critical decisions on flawed or biased information.

AI-GENERATED POLICE REPORTS POSE HIGH-STAKES RISKS IN LEGAL PROCEEDINGS

Police reports are the backbone of criminal prosecutions and are relied upon for accurate documentation of arrests, police interactions, and probable cause determinations. These reports provide essential details that inform investigations, court proceedings, and legal strategies. When cases take months or years to reach trial, officers often depend on their reports to refresh their memory about key details of an incident.

If an AI tool generates a police report, even a minor error can have profound legal and constitutional implications. Legal challenges that may arise as a result of AI-generated police reports include those related to:

Fourth Amendment concerns: Fourth Amendment analyses often turn on minute details concerning the precise timing, sequence, location, and nature of discrete actions. Even small factual errors in police reports can lead to incorrect determinations about the admissibility of evidence and the constitutionality of searches. For example, if an officer conducts a legal traffic stop and sees contraband in plain view on top of the glove compartment, but an AI-generated report mistakenly says the contraband was found inside the glove compartment, it could undermine the finding of probable cause and lead the court to rule the search unlawful. If a car owner refuses to consent to a search but a passenger advises the owner to allow it, and an AI-generated report mistakenly attributes the consent to the owner rather than the passenger, it could mislead the court into finding the search was voluntary. If the police search multiple rooms in a shared living space and the AI-generated report incorrectly places the discovery of contraband in the wrong room, it could confuse the issue of who has standing to challenge the search. These types of factual errors — common in AI-generated outputs — can have serious constitutional implications, affecting motions to suppress evidence, altering the course of post-conviction proceedings, and ultimately undermining the fairness and accuracy of judicial outcomes.

***Brady* challenges:** Defendants whose police reports are generated by AI tools could argue they are entitled to information on the underlying technology used to create the report—including its accuracy record, the algorithms it uses, and the data it was trained on—as part of their right to exculpatory evidence under *Brady v. Maryland*.[19] Much of this information would likely be inaccessible to police or prosecutors, as it is considered proprietary by the vendors. This could lead to complex legal challenges around the admissibility of AI-generated reports absent such disclosures, potentially endangering prosecutions and consuming prosecutorial office resources on litigation.

The reliability of officer testimony: As discussed above, AI police reporting tools put the onus on individual officers to catch factual inaccuracies in generated reports, and the process of cognitive offloading may lead to complacency by the reviewing officer, allowing small mistakes to slip through unnoticed. A prosecutorial office made aware of an inaccuracy in an AI-generated police report may be obligated to flag the officer who approved the report as an unreliable witness and to disclose the inaccuracy to defendants in future cases in which the officer is involved, even if the inaccuracy was unintentional.

The risks of inaccuracies are magnified during cross-examination. Defense attorneys could argue that the AI, not the officer, authored the report, thereby challenging the officer's ability to confirm details with first-hand knowledge. This line of questioning could become particularly damaging if errors in the report come to light. Even a seemingly minor error could damage an officer's credibility in court, undermining the prosecution's case and damaging the officer's reputation.[20]

In an internal memo to law enforcement partners, Chief Deputy Daniel Clark from the King County Prosecuting Attorney's Office warned officers that if they certify a report with factual errors, the consequences could be "devastating" for the case and subject them to inclusion on the office's Brady list, which tracks officers who have been found to make untrue statements or whose testimony cannot be relied upon in court.[21]

These issues will likely strain prosecutor offices, many of which lack the resources to litigate such technical challenges effectively. Moreover, the added resource burden could detract from other prosecutorial work, harming the office's ability to effectively and expeditiously manage its caseload.

AI TOOLS CARRY PRIVACY CONCERNS

AI-generated reports raise significant privacy concerns, particularly regarding how personal information is captured and stored.[22] When AI systems process body camera footage, they may not be able to differentiate between relevant evidence and unrelated private details.[23] For instance, if a bystander's personal conversation is captured in the background of a recording, that information could inadvertently become part of the police report — even if it has no bearing on the case. This data is then shared with the external vendors who control the technology, where it could be stored indefinitely.

This lack of control introduces significant risks, including data misuse, unauthorized access, and the erosion of public trust in law enforcement practices.

THE TIME-SAVING BENEFITS OF AI TOOLS MAY BE MINIMAL

One of the most compelling arguments favoring AI-generated reports is that they save time. However, the first independent study of Axon's Draft One found no evidence to support this claim.[24] The report suggests that, because many police departments already have templates to streamline report writing and the AI tool still requires that officers perform extensive data entry, outsourcing the narrative writing process may not substantially reduce the overall time necessary to complete a report.

Moreover, the time investment required to write a report traditionally discourages frivolous over-filing, helping officers and departments focus on significant cases and avoid creating a backlog. If AI made over-filing easier, it could overwhelm district attorney offices with excessive reports, bogging down their resources and stretching prosecutorial attention thin.

RELIANCE ON AI COULD ERODE PUBLIC TRUST IN LAW ENFORCEMENT

A damaging possible consequence of using AI-generated police reports is their potential to erode public trust, particularly in communities already burdened by historical over-policing, racial profiling, and excessive use of force by law enforcement. Many of these communities have longstanding skepticism toward police, and introducing AI into a process as critical as report-writing risks amplifying existing tensions, including by:

- **Demonstrating a propensity for error**: AI-generated reports are prone to inaccuracies stemming from flawed algorithms, poor data inputs, or misinterpretations of context.[25] Even minor errors in these reports can lead to serious consequences, such as wrongful arrests or convictions. These incidents could degrade public trust in the capacity of police to reliably investigate and address crimes.

- **Shifting officer behavior**: AI-generated reports may require officers to narrate incidents in ways that prioritize the needs of the transcript, rather than the natural flow of conversation. This could fundamentally alter the dynamics of police interactions. Officers may feel pressured to speak in ways that emphasize specific actions or responses to ensure the transcript is comprehensive for the AI model's purposes. This approach would disrupt otherwise straightforward communication, creating forced and unnatural exchanges.

- **Weakening perceptions of officer accountability**: Relying on AI may make law enforcement agencies appear less transparent and accountable for their actions. This reliance can create the perception that officers are detached from the consequences of their decisions, leading communities to feel that law enforcement is increasingly impersonal and inaccessible.

- **Highlighting the legal system's lack of safeguards for untested tools**: The rapid adoption of untested AI tools risks creating lasting negative perceptions of the technology before its full potential can be realized. Although advancements in AI could eventually lead to reliable applications that genuinely support public safety, early missteps associated with errors and injustices may cement skepticism in the public consciousness. If AI tools, like those used in policing, become synonymous with unjust outcomes, communities will likely remain resistant to future efforts to integrate AI into law enforcement practices.

LEARNING FROM THE PAST: HISTORICAL PRECEDENTS HIGHLIGHT RISKS

Privacy concerns associated with AI-generated reporting tools echo challenges seen in other untested surveillance technologies. Baltimore's aerial surveillance program serves as a particularly relevant example. Designed to solve crimes by deploying planes to capture extensive footage of the city, the program operated without the public's knowledge or consent, sparking significant backlash. Critics argued the initiative amounted to mass surveillance, violating Fourth Amendment protections against unreasonable searches.[26] The courts ultimately agreed, ruling that such widespread, unregulated surveillance infringed on citizens' reasonable expectations of privacy.[27] This case illustrates the dangers of deploying advanced surveillance tools without clear limitations, safeguards, or accountability mechanisms.

Although AI-generated police reports and Baltimore's aerial surveillance program differ in function, they share fundamental privacy and civil liberties risks. Both involve outsourcing sensitive data — whether body camera footage or aerial surveillance records — to systems that diminish departmental control and transparency. In both cases, the lack of oversight creates opportunities for misuse and raises constitutional concerns, while also highlighting the critical need for transparency, oversight, and well-defined limitations to protect civilian privacy and uphold trust in law enforcement.

LOOKING FORWARD

Although AI may streamline specific report-writing tasks, it is significantly susceptible to error, diminishes direct officer responsibility for reports, risks reinforcing entrenched biases, threatens to trigger complex and resource-intensive litigation, infringes on individuals' privacy, and holds the potential to damage the public's already strained trust in the criminal legal system. While some challenges associated with AI in policing could likely be mitigated over time, the field has yet to initiate much of the work necessary to make these tools potentially suitable for use in criminal investigations.

Developments necessary to lay the groundwork for the use of this technology include:

- **Increased accuracy**: Before police departments can outsource the critical work of report drafting to AI, these tools must undergo external validation studies that show them to be reliably accurate in interpreting input data from complex and varied situations.

- **Development of transparency and accountability mechanisms**: Officers must review AI-generated reports and be trained to understand the technology and its limits. Departments and vendors must work together to create a feedback loop for correcting systematic errors. Departments must develop processes for catching inaccuracies and designating individual accountability for end products.

- **Eliminating biases in training data**: AI systems must be trained on unbiased datasets to avoid reinforcing historical prejudices. Unfortunately, this is not yet possible with current technology: current AI systems are not equipped to remove or reliably mitigate bias, and the historical data sets AI models are trained on inevitably reflect entrenched societal biases. Until more sophisticated training techniques are developed, implemented, and rigorously tested, we cannot trust AI-generated police reports will deliver unbiased outcomes.

- **Establishing clear legal frameworks**: Courts need to develop guidelines on the admissibility of AI-generated reports and determine the extent to which defendants' Brady rights extend to information on the underlying technology behind AI tools.

- **Transparent third-party evaluations**: It is critical to have independent evaluations that consider the context of AI's proposed use. Accuracy must be demonstrated not only in general but also specifically for the populations impacted by these tools.

CONCLUSION

AI-generated police reporting technology presents significant risks that far outweigh its purported benefits. While automation offers the potential to streamline aspects of law enforcement, even minor errors can have serious legal and societal consequences. The high stakes of criminal investigations and prosecutions demand the utmost standards of accuracy, fairness, and transparency – unfortunately, these are all areas in which AI falls short.

Adopting AI-generated reporting without robust safeguards, independent audits, and measurable improvements in accuracy and bias mitigation risks undermining the trust essential for effective policing. For this reason, law enforcement agencies must proceed with extreme caution when considering these tools. Similarly, prosecutors should monitor whether these tools are being used by their local law enforcement agencies and implement internal policies declining to accept reports generated by these tools or implementing safeguards around their use.

Until AI technology can reliably meet the rigorous demands required of policing, law enforcement must prioritize transparent and ethical practices that protect constitutional rights and maintain public confidence. Rushing to implement flawed systems risks undermining the credibility of law enforcement, eroding trust, and perpetuating systemic biases that compromise the principles of equity and justice.

ENDNOTES

1. The terms “district attorney,” “DA,” or “elected prosecutor” are used generally to refer to any chief local prosecutor, including State’s Attorneys, Prosecuting Attorneys, and Attorneys General with local jurisdiction.
2. Axon, “Axon Reimagines Report Writing with Draft One, a First-of-Its-Kind AI-Powered Force Multiplier for Public Safety,” April 23, 2024, https://investor.axon.com/2024-04-23-Axon-reimagines-report-writing-with-Draft-One,-a-first-of-its-kind-AI-powered-force-multiplier-for-public-safety.
3. Ian T. Adams, “Large Language Models and Artificial Intelligence for Police Report Writing,” *CrimRxiv*, February 28, 2024, https://www.crimrxiv.com/pub/c5lj2rmy/release/1.
4. Adam Tauman Kalai and Santosh S. Vempala, “Calibrated Language Models Must Hallucinate,” *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*, March 20, 2024, https://doi.org/10.48550/arXiv.2311.14648.
5. Todd Bishop, “Internal Memo: Don’t Use AI for Police Reports, Prosecutor Tells Seattle-Area Law Enforcement,” *GeekWire*, September 26, 2024, https://www.geekwire.com/2024/internal-memo-dont-use-ai-for-police-reports-prosecutor-tells-seattle-area-law-enforcement/.
6. Matt O’Brien, “Chatbots Sometimes Make Things Up. Is AI’s Hallucination Problem Fixable?,” *AP News*, August 1, 2023, https://apnews.com/article/artificial-intelligence-hallucination-chatbots-chatgpt-falsehoods-ac4672c5b06e6f91050aa46ee731bcf4.
7. *Ibid*.
8. Andrew Guthrie Ferguson, “Generative Suspicion and the Risks of AI-Assisted Police Reports,” *Northwestern Law Review*, forthcoming, July 17, 2024, https://ssrn.com/abstract=4897632.
9. *Ibid*.
10. Siyuan Feng, Bence Mark Halpern, Olya Kudina, and Odette Scharenborg, “Towards inclusive automatic speech recognition,” *Computer Speech & Language* 84 (2024), https://www.sciencedirect.com/science/article/pii/S0885230823000864.
11. *Ibid*.
12. *Ibid*.
13. Han Yu and Natalie Monas, “Recreating the Scene: An Investigation of Police Report Writing,” *Journal of Technical Writing and Communication* 50, no.1 (2018): 35–55, https://journals.sagepub.com/doi/10.1177/0047281618812441,
14. Sandra Grinschgl, Frank Papenmeier, and Hauke Meyerhoff, “Consequences of Cognitive Offloading: Boosting Performance but Diminishing Memory,” *Journal of Experimental Psychology* 74, no. 9 (2021):1477-96, https://pmc.ncbi.nlm.nih.gov/articles/PMC8358584/.
15. *Ibid. See also* Michael Gerlich, “AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking,” *Societies* 15, no. 1 (2025), https://www.mdpi.com/2075-4698/15/1/6.
16. Nico Grant, “Google Chatbot’s AI Images Put People of Color in Nazi-Era Uniforms,” *The New York Times*, February 22, 2024, https://www.nytimes.com/2024/02/22/technology/google-gemini-german-uniforms.html.
17. Clare Veal, Monty Raper, and Peter Waters, “The Perils of Feedback Loops in Machine Learning: Predictive Policing,” *Gilbert + Tobin*, February 20, 2023, https://www.gtlaw.com.au/knowledge/perils-feedback-loops-machine-learning-predictive-policing.
18. Lou Blouin, “AI’s Mysterious ‘Black Box’ Problem, Explained,” *University of Michigan-Dearborn News*, March 6, 2023, https://umdearborn.edu/news/ais-mysterious-black-box-problem-explained.
19. Ferguson, “Generative Suspicion.”
20. Ferguson, “Generative Suspicion.”
21. Bishop, “Internal Memo.”
22. Morgan Sullivan, “Examining Privacy Risks in AI Systems,” *Transcend*, December 1, 2023, https://transcend.io/blog/ai-and-privacy.
23. Ferguson, “Generative Suspicion.”

24. Ian T. Adams, Matt Barter, Kyle McLean, et al., "No Man's Hand: Artificial Intelligence Does Not Improve Police Report Writing Speed," *Journal of Experimental Criminology* (2024), https://doi.org/10.1007/s11292-024-09644-7.
25. Grant, "Google Chatbot."
26. Janus Rose, "Baltimore Has Officially Ended Its Creepy Spy Plane Program," *VICE*, February 4, 2021, https://www.vice.com/en/article/baltimore-has-officially-ended-its-creepy-spy-plane-program/.
27. The Fourth Circuit Court of Appeals ruling found that the BPD's surveillance system violates the Fourth Amendment because persistent surveillance of outdoor movements invades people's reasonable expectation of privacy. It explained that "allowing the police to wield this power unchecked is anathema to the values enshrined in our Fourth Amendment." *Leaders of a Beautiful Struggle v. Baltimore Police Department*, 20-1495 (2021), https://www.ca4.uscourts.gov/opinions/201495A.P.pdf.

FOR MORE INFORMATION:
Contact FJP at info@fairandjustprosecution.org