Issues in Information Systems
Volume 27, Issue 3, pp. 649-665, 2026

DOI: https://doi.org/10.48009/3_iis_2026_145

Public perceptions of generative AI–assisted policing: Effects on trust, legitimacy, fairness, and accountability

Christopher Wydra, *Pennsylvania Western University*, *wydra@pennwest.edu*

Abstract

This study examines public perceptions of generative AI–assisted policing and its effects on trust, legitimacy, fairness, accountability, and public acceptance of law enforcement practices. Using a between-subjects experimental survey design, participants are randomly assigned to scenarios involving human-only police report writing, AI-assisted report writing with disclosure, or AI-assisted report writing without disclosure. Grounded in procedural justice theory and research on algorithmic decision-making, the study explores whether transparency regarding AI use influences public perceptions of policing and whether factors such as prior trust in police, familiarity with AI, and concerns about algorithmic bias shape these perceptions. The study contributes to emerging research on generative AI in criminal justice by examining how AI-assisted policing may affect public trust and legitimacy.

Keywords: procedural justice, experimental vignette, algorithmic bias, AI-assisted policing

Introduction

Artificial intelligence (AI) has become increasingly integrated into professional decision-making environments across healthcare, business, education, and criminal justice. Recent developments in generative AI, particularly large language models (LLMs), have expanded the capabilities of AI systems beyond prediction and classification to include text generation, summarization, report writing, and decision support (Dwivedi et al., 2023). As organizations increasingly incorporate generative AI into routine workflows, scholars have raised concerns regarding transparency, misinformation, bias, and accountability associated with AI-generated content (Bender et al., 2021). These concerns are particularly important in institutional settings where decisions and documentation may significantly affect public trust and legitimacy.

Within criminal justice, AI technologies have primarily been applied through predictive policing systems, risk assessment instruments, and algorithmic decision-making tools intended to improve efficiency and resource allocation (Faber, 2026). However, prior research has demonstrated that algorithmic systems may reproduce or amplify existing social inequalities when trained on historically biased data (Richardson et al., 2019). As a result, scholars have emphasized the need for transparency, oversight, and accountability in AI-assisted criminal justice decision-making.

More recently, generative AI has emerged as a distinct category of AI technology capable of producing human-like narratives, summaries, and written reports. In policing contexts, these tools may be used to assist officers with report writing, investigative summarization, and information synthesis (Halford, 2025).

Although these applications may improve efficiency and reduce administrative burdens, they also raise important questions regarding accuracy, automation bias, accountability, and public legitimacy. Existing research in human-AI interaction suggests that individuals may over-rely on AI-generated outputs, particularly when automated systems are perceived as objective or authoritative (Starke et al., 2021).

Public perceptions of policing are strongly shaped by procedural justice, which emphasizes fairness, transparency, neutrality, and accountability in institutional decision-making (Tyler, 2006). Research consistently demonstrates that perceptions of procedural fairness influence public trust, cooperation, and perceptions of police legitimacy (Mazerolle et al., 2013). The introduction of generative AI into policing may therefore affect public evaluations of law enforcement depending on whether AI-assisted processes are viewed as fair, transparent, and accountable. Emerging studies suggest that support for AI in policing is closely associated with perceptions of procedural justice and institutional trust (Guler et al., 2025).

Despite growing interest in AI governance and algorithmic policing, limited empirical research has examined how the public perceives the use of generative AI in routine policing functions such as report writing and investigative decision support. Existing studies have largely focused on predictive systems, facial recognition, and algorithmic risk assessment rather than generative AI applications. This study addresses this gap by examining whether public trust, perceived fairness, legitimacy, accountability, and acceptance differ when police actions are described as human-generated versus generative AI-assisted. Using a survey experiment with experimental vignettes, this study seeks to contribute to the growing literature on AI, procedural justice, and public perceptions of law enforcement.

Background and Literature Review

The integration of artificial intelligence (AI) into professional decision-making environments has expanded rapidly across multiple domains, including criminal justice. While earlier applications of AI in policing focused on predictive analytics and risk assessment tools, recent developments in generative AI, particularly large language models, have introduced new capabilities in text generation, summarization, and decision support (Fabila et al., 2026). These tools are increasingly being integrated into routine workflows, including report writing, documentation, and information synthesis (Michelet & Breitinger, 2023). Despite these advancements, empirical research examining how the public evaluates the use of generative AI in routine policing activities—particularly AI-assisted police report writing and the role of transparency in shaping perceptions of trust, legitimacy, fairness, and accountability—remains extremely limited. This study addresses that gap through an experimental examination of public perceptions across different AI-assisted reporting conditions.

AI in Criminal Justice

Artificial intelligence has been applied in criminal justice primarily through predictive policing, risk assessment instruments, and data-driven decision-making systems. Research on AI in criminal justice has largely focused on predictive policing, risk assessment, and algorithmic decision-making systems. Farber's systematic review found that AI applications in policing often reproduce concerns about algorithmic bias, opacity, and due process despite operational benefits (Faber, 2026). These technologies have generated both optimism regarding efficiency and concern regarding fairness, bias, and accountability. Prior research has demonstrated that algorithmic systems may reproduce or amplify existing inequalities, particularly when trained on biased historical data. Scholars consistently note that although AI may improve efficiency and resource allocation, it also raises concerns about bias, accountability, and transparency. As a result, scholars have emphasized the importance of governance frameworks emphasizing transparency, oversight, and legitimacy when AI tools are deployed in law enforcement (Halford, 2025).

Procedural Justice and Legitimacy

Studies on police legitimacy consistently show that citizens are more likely to trust and cooperate with law enforcement when procedures are viewed as fair and transparent (Mazerolle, et al., 2013). Procedural justice theory asserts that public trust in policing depends heavily on perceptions of fairness, neutrality, transparency, and respectful treatment rather than solely on outcomes. Guler (2025) relates that perceptions of neutrality, voice, trustworthiness, and dignity strongly shape support for policing institutions. Public perceptions of policing are strongly influenced by procedural justice, which emphasizes fairness, transparency, and impartiality in decision-making processes. Legitimacy is derived not only from outcomes but also from the perceived fairness of the procedures used to reach those outcomes. The introduction of AI into policing raises important questions about whether algorithmic or AI-assisted processes are consistent with the principles of procedural justice.

Research suggests that individuals evaluate institutional legitimacy based on trust, fairness, and accountability. The use of automated or algorithmic systems may disrupt these perceptions if individuals believe that responsibility is diffused or that decisions lack human judgment. Sagana's (2025) recent work suggests that algorithmic policing technologies may disrupt legitimacy when people perceive that responsibility is shifted away from human judgment or when accountability becomes unclear. Conversely, some research indicates that AI may be perceived as more objective than human decision-makers, potentially increasing perceptions of fairness under certain conditions.

Generative AI and Decision-Making

Generative AI represents a distinct category of AI systems that produce human-like text and summaries. Unlike traditional algorithmic systems, generative AI directly contributes to the construction of narratives, reports, and written documentation. Halford proposed the "Transformer Led Policing" framework, arguing that generative AI may assist with report writing, intelligence analysis, and public communication but requires strong safeguards for legitimacy and ethics (2025). This raises unique concerns regarding accuracy, hallucination, bias in language generation, and the attribution of responsibility for generated content with police report writing and communications. Even though there are concerns, police leadership notes the increasing use of ChatGPT and other LLMs for report writing and summarization tasks within police workflows (Potts et al., 2024).

In professional contexts, generative AI is increasingly used to assist with writing, summarization, and information synthesis. However, research in human-AI interaction has identified risks such as automation bias, where individuals over-rely on AI-generated outputs, and reduced accountability when decision-making is partially delegated to automated systems. This is included in the policing profession, where Calder (2026) identified risks associated with large language models in policing, including inaccuracies, legal risks, and reduced accountability in criminal investigations.

Public Perceptions of AI

Public attitudes toward AI are shaped by factors such as familiarity, perceived usefulness, and concerns about bias and fairness. Studies have found that individuals with greater exposure to AI technologies tend to exhibit higher levels of acceptance, whereas those concerned about algorithmic bias tend to express lower levels of trust. Miyazaki (2023) found that public attitudes toward generative AI are more positive with greater exposure to AI technologies, although concerns about ethics and bias remain significant. Fairness perceptions are also a concern; in algorithmic decision-making, transparency and explainability significantly influence whether people perceive AI systems as legitimate and fair (Starke et al., 2021).

Transparency and procedural justice have also emerged as critical factors, with disclosure of AI use influencing perceptions of legitimacy and accountability. Guler et al. found that procedural justice

perceptions mediate public support for AI in policing, particularly perceptions of neutrality and trustworthiness (Guler et al., 2025). These findings suggest that public acceptance of AI-assisted policing may depend less on the technology itself and more on whether its use is perceived as fair, transparent, and procedurally legitimate.

Technology Acceptance, Trust, and Responsible AI

The study is informed by two complementary theoretical perspectives: procedural justice theory, which explains how perceptions of fairness and legitimacy shape evaluations of policing, and technology acceptance research, which explains how trust, transparency, perceived risk, and accountability influence public acceptance of emerging information technologies.

Beyond criminal justice, the adoption of artificial intelligence has been extensively examined within the broader technology acceptance and information systems literature. Technology acceptance research suggests that individuals' willingness to adopt or support emerging technologies is influenced not only by perceived usefulness and ease of use but also by trust, transparency, perceived risk, fairness, and confidence in the technology's outputs. As AI systems become increasingly involved in professional decision-making, researchers have argued that traditional technology acceptance models should be expanded to include ethical and governance-related constructs that shape public acceptance of AI-assisted decision support (Glikson & Woolley, 2020).

Trust has emerged as one of the most consistently studied determinants of AI acceptance. Unlike trust in human decision-makers, trust in AI reflects confidence that automated systems will generate accurate, reliable, and unbiased recommendations while operating in ways that align with human expectations and societal values (Glikson & Woolley, 2020). Research further demonstrates that trust in AI is influenced not only by technical performance but also by users' understanding of how AI systems generate outputs and whether appropriate human oversight remains in place. In law enforcement settings, where decisions carry significant legal and social consequences, maintaining public trust requires confidence that AI functions as a decision-support tool rather than a replacement for professional judgment.

Transparency has similarly become a central principle of responsible AI governance. Transparency refers to the extent to which organizations disclose the use of AI systems, explain how outputs are generated, and communicate the role of human oversight in decision-making. Explainability and transparency have consistently been associated with greater perceptions of legitimacy, procedural fairness, and organizational accountability, particularly when AI influences consequential decisions (Starke et al., 2021). Conversely, limited disclosure regarding AI use may increase uncertainty, reduce trust, and raise concerns regarding institutional legitimacy.

Concerns regarding fairness, bias, and accountability also occupy a prominent place within contemporary AI research. Scholars have documented that AI systems may reproduce or amplify existing societal inequalities when trained on historical or incomplete datasets, leading to concerns about discriminatory outcomes and unequal treatment across demographic groups (Mehrabi et al., 2021). Public evaluations of AI therefore extend beyond assessments of technical performance to include perceptions that AI may reinforce existing biases, omit important contextual information, or produce inaccurate conclusions. These concerns are especially salient in policing, where incomplete or biased documentation may affect investigative decisions, prosecutorial outcomes, judicial proceedings, and public confidence in law enforcement institutions.

Accountability represents another foundational construct in responsible AI research. While AI systems may assist decision-makers by improving efficiency and consistency, responsibility for decisions remains a human obligation. Researchers have noted that uncertainty about who is accountable for AI-assisted

decisions can undermine public confidence and institutional legitimacy, particularly when AI-generated outputs influence high-stakes public-sector decisions (Glikson & Woolley, 2020). Consequently, perceptions of accountability are closely intertwined with transparency, trust, and fairness, suggesting that public acceptance of AI-assisted policing depends not only on technological capability but also on confidence that appropriate human oversight and responsibility are maintained throughout the decision-making process.

Together, these perspectives provide the theoretical foundation for the constructs examined in the present study. Trust, transparency, fairness, accountability, concerns regarding AI bias, perceptions that AI may reinforce inequalities, and concerns that AI-generated reports may omit important investigative details have each been identified as central dimensions influencing public acceptance of AI-assisted technologies. Examining these constructs within the context of generative AI-assisted police report writing extends established technology acceptance and responsible AI frameworks to an emerging application in criminal justice.

Emerging Directions in Generative AI-Assisted Policing

Despite the growing body of scholarship on artificial intelligence in criminal justice, several important gaps remain. First, existing research has focused predominantly on predictive policing, algorithmic risk assessment, facial recognition, and other automated decision-making systems, while comparatively little attention has been devoted to generative AI applications such as report writing, investigative documentation, and decision-support functions. As generative AI becomes increasingly integrated into routine police operations, understanding public responses to these emerging applications has become increasingly important.

Second, although prior studies have examined public attitudes toward AI broadly, relatively few have investigated how the public evaluates the use of generative AI in everyday policing activities that involve direct officer discretion and written communication. Police reports are foundational documents that influence investigative decisions, prosecutorial review, judicial proceedings, and public accountability, yet little empirical evidence exists regarding whether AI-assisted report preparation affects public perceptions of trust, legitimacy, fairness, or accountability.

Third, existing studies have emphasized transparency and procedural justice as important determinants of public acceptance of AI, but limited research has experimentally examined whether disclosure of generative AI use influences these perceptions. Most prior research has relied on observational surveys or evaluations of existing technologies rather than experimental comparisons that isolate the effects of AI assistance and transparency.

Accordingly, the present study addresses these gaps by experimentally comparing public perceptions across three police report preparation conditions: a human-generated report, an AI-assisted report with disclosure of AI use, and an AI-assisted report without disclosure. This design allows for a direct examination of how generative AI and transparency independently influence perceptions of trust, legitimacy, fairness, and accountability.

The Present Study

This study makes several important contributions to the emerging literature on artificial intelligence and policing. First, it extends existing research beyond predictive and algorithmic policing by examining one of the fastest-growing applications of generative AI—police report writing and investigative documentation. As law enforcement agencies increasingly explore large language models to improve

efficiency and administrative workflows, empirical evidence regarding public acceptance of these technologies is needed to inform policy and implementation.

Second, the study contributes to procedural justice and police legitimacy research by examining whether the disclosure of AI assistance influences perceptions of fairness, transparency, and accountability. Rather than treating AI adoption as a single construct, this study distinguishes between AI-assisted policing with and without disclosure, providing a more nuanced understanding of how transparency shapes public evaluations.

Finally, the study employs an experimental vignette methodology that allows causal comparisons across policing scenarios while controlling for extraneous influences. This approach provides stronger evidence regarding the effects of AI assistance and disclosure than descriptive opinion surveys alone. Collectively, these contributions advance the emerging scholarship on generative AI in criminal justice and provide practical insights for law enforcement agencies considering the implementation of AI-assisted documentation systems in ways that preserve public trust and institutional legitimacy.

Research Questions

RQ1: Does the use of generative AI in police report writing affect public perceptions of trust in police?

RQ2: Does transparency about AI use increase or decrease public acceptance of generative AI in policing?

RQ3: Are perceptions of generative AI–assisted policing influenced by prior trust in police, familiarity with AI, or concern about algorithmic bias?

Methodology

Research Design

The study employed a quantitative experimental survey design using a between-subjects vignette methodology to examine public perceptions of generative AI–assisted policing. Experimental vignette methodology is commonly used in social science research to examine perceptions, attitudes, and decision-making processes because it allows researchers to systematically manipulate contextual variables while maintaining experimental control (Aguinis & Bradley, 2014). The survey instrument presented participants with a residential burglary investigation scenario in which the investigative report is prepared under one of three randomized conditions: (1) a human-only police report, (2) an AI-assisted report with disclosure of AI use, or (3) an AI-assisted report without disclosure of AI use. Participants are randomly assigned to a single condition to minimize selection bias and enable causal comparisons between groups. The use of experimental between-subjects vignettes highly aligns with policing research, randomized assignment, procedural justice, legitimacy, trust, and attitudinal outcomes (Nivette et al., 2024).

The independent variable in the study is the type of report preparation condition (human-only, AI-assisted disclosed, AI-assisted undisclosed). The dependent variables include participants' perceptions of trust, legitimacy, accountability, fairness, and acceptance of AI use in policing. These constructs are measured through multiple Likert-scale items ranging from strongly disagree to strongly agree. Additional measures assess familiarity with generative AI, concerns regarding algorithmic bias, manipulation checks, attention validation, and demographic characteristics.

The design is appropriate because vignette-based experiments allow researchers to systematically manipulate specific variables while maintaining a realistic policing context. This approach enabled the

study to isolate how disclosure and use of generative AI influence public perceptions of policing legitimacy and procedural fairness. This study used a randomized vignette survey experiment to examine public perceptions of police technologies, emphasizing transparency, legitimacy, and internal validity through randomized conditions that closely parallel your AI-assisted policing framework (Przeszlowski & Guerette, 2025). The use of randomized conditions strengthens internal validity by reducing confounding influences and improving the ability to attribute observed differences in participant responses directly to the experimental manipulation.

Participants

Participants in the survey were 216 U.S. adults recruited through SurveyMonkey Audience, a research marketing company. Nearly all respondents (95.83%) completed the online survey, and data collection occurred between May 12 and 14, 2026. The modal age group was 35-44 years (28.78%), with 45.85% in the 35-54 age range; 52.97% were female, and 47.03% were male. Most participants have a college education, with 38.73% holding a bachelor's degree and 26.96% having a graduate degree. Political orientation was relatively evenly split among conservative (34.47%), moderate (36.89%), and liberal (28.64%).

Measures

The three experimental conditions: (1) a human-only police report condition, (2) an AI-assisted report condition with disclosure of AI use, or (3) an AI-assisted report condition without disclosure, resulted in relatively balanced group distributions, with approximately 37.9% assigned to the human-only condition, 30.3% assigned to the AI-disclosed condition, and 31.8% assigned to the AI-undisclosed condition. Survey responses were collected using five-point Likert-type scales ranging from 1 (Strongly Disagree) to 5 (Strongly Agree). The final sample included approximately 206–209 valid respondents across most substantive survey items, after incomplete responses and quality-control exclusions. Composite measures demonstrated strong internal consistency reliability across the primary attitudinal items, with the overall scale producing a Cronbach's alpha indicative of high reliability ($\alpha \approx .88$), suggesting that the items consistently measured related perceptions of AI-assisted policing and procedural legitimacy.

Procedure

The study received approval from the Pennsylvania Western University IRB, and participants provided informed consent before being randomly assigned to one of the three experimental conditions. Each participant read a vignette describing a police response to a burglary investigation, with the experimental manipulation embedded in the description of how the report was generated. Following the vignette, participants completed a series of survey measures assessing trust, legitimacy, accountability, fairness, and acceptance of AI use. Forced-response validation ensured zero missing data on core survey questions by using a survey quiz, in which respondents had to answer to proceed. The average quiz score was 100% (standard deviation = 0.00), indicating that all core scenario questions were answered.

Measures and Construct Operationalization

The dependent variables in this study consisted of participants' perceptions of AI-assisted policing following exposure to one of three experimental police report conditions. Consistent with prior research in technology acceptance, procedural justice, and responsible AI, each construct was operationalized using multiple survey items measured on a five-point Likert-type scale ranging from 1 (Strongly Disagree) to 5 (Strongly Agree), unless otherwise noted. Higher scores reflected greater endorsement of the construct being measured. Composite scores were calculated by averaging responses across items representing each construct.

Trust

Trust was conceptualized as participants' confidence that AI-assisted police report writing would produce accurate, reliable, and appropriate investigative documentation. Trust reflects the belief that AI can support law enforcement decision-making while maintaining competence, consistency, and reliability. This construct was measured using multiple items assessing confidence in AI-assisted report preparation and the reliability of AI-generated investigative information, consistent with prior research on trust in artificial intelligence (Glikson & Woolley, 2020).

Legitimacy

Legitimacy was defined as the extent to which participants perceived AI-assisted policing as appropriate, acceptable, and deserving of public confidence. Drawing from procedural justice scholarship, legitimacy reflects the belief that policing practices are exercised in a manner consistent with societal expectations, ethical standards, and institutional authority (Tyler, 2006; Mazerolle et al., 2013). Participants evaluated whether the use of AI-assisted report writing enhanced or diminished the legitimacy of police decision-making.

Procedural Fairness

Procedural fairness was operationalized as participants' perceptions that AI-assisted policing promotes impartial, unbiased, transparent, and equitable decision-making processes. Consistent with procedural justice theory, the construct emphasized fairness in the decision-making process rather than the investigative outcome itself. Survey items assessed perceptions of neutrality, fairness, and the equitable application of AI-assisted technologies within policing.

Accountability. Accountability referred to participants' perceptions that responsibility for investigative decisions remains clearly assigned despite the use of AI-assisted technologies. Items measured whether respondents believed police officers and agencies would remain responsible for decisions supported by generative AI and whether appropriate human oversight would be maintained. This construct reflects principles of responsible AI governance emphasizing transparency, explainability, and human accountability in high-consequence decision-making.

Support for AI-Assisted Policing

Support for AI-assisted policing represented participants' overall acceptance of using generative AI to assist officers with police report writing and investigative documentation. This construct assessed participants' willingness to support the integration of AI into routine policing functions while recognizing that AI serves as a decision-support tool rather than an autonomous decision-maker. Higher scores indicated greater acceptance of AI-assisted policing practices.

AI-Related Concerns

AI-related concerns were measured through several constructs reflecting common issues identified in the responsible AI and criminal justice literature. These included concerns that AI systems may produce biased outcomes, reinforce existing social inequalities, omit important investigative details, generate inaccurate or misleading information, or reduce human judgment in police decision-making. Participants also evaluated concerns regarding transparency and the potential consequences of relying on AI-generated police documentation. Higher scores indicated greater concern regarding the risks associated with AI-assisted policing.

Familiarity with Artificial Intelligence

Participants also reported their familiarity with generative AI technologies. This measure served as a contextual variable to assess respondents' prior exposure to AI systems and to examine whether familiarity influenced perceptions of AI-assisted policing. Previous research suggests that greater familiarity with AI is associated with higher levels of technology acceptance and more favorable attitudes toward AI-assisted decision-making (Miyazaki, 2023).

Collectively, these constructs provide a comprehensive assessment of public perceptions of generative AI-assisted policing by integrating theoretical perspectives from technology acceptance, procedural justice, and responsible AI governance. The measures were selected to capture both the perceived benefits and potential risks associated with using generative AI for police report writing and investigative decision support.

Data Analysis

Data were analyzed using IBM SPSS Statistics to examine differences in public perceptions across the three experimental conditions: Human-Only policing, AI-Assisted policing with disclosure, and AI-Assisted policing without disclosure. Descriptive statistics, including frequencies, means, and standard deviations, were calculated for all study variables. Internal consistency reliability of the multi-item constructs was evaluated using Cronbach's alpha.

Prior to hypothesis testing, the assumptions underlying the parametric analyses were evaluated. Visual inspection of histograms and normal probability (Q-Q) plots, together with assessments of skewness and kurtosis, indicated that the primary dependent variables approximated normal distributions within each experimental group. Homogeneity of variances was assessed using Levene's test, and no substantial violations were observed. Given the relatively balanced sample sizes across experimental conditions, the independent-samples *t*-tests and one-way analyses of variance (ANOVA) were considered appropriate and robust to minor deviations from normality.

Independent-samples *t*-tests were conducted to compare specific pairs of experimental conditions, while one-way ANOVA was used to evaluate overall differences across all three report preparation conditions. Because Likert-type responses may not always satisfy the assumptions of parametric analyses, Mann–Whitney *U* tests were also conducted as a robustness check. The nonparametric analyses produced findings consistent with the parametric results, increasing confidence in the stability of the observed effects.

Pearson correlation analyses were performed to examine relationships among the primary study variables. Multiple linear regression analyses were subsequently conducted to evaluate the extent to which procedural fairness, legitimacy, trust, and AI-related concerns predicted support for AI-assisted policing. Prior to estimating the regression model, assumptions of linearity, independence of errors, homoscedasticity, normality of residuals, and multicollinearity were evaluated and determined to be acceptable. Statistical significance was evaluated using an alpha level of .05.

Regression Analysis

Multiple linear regression analysis was conducted to examine whether procedural fairness, legitimacy, and AI-related concerns predicted participants' support for AI-assisted policing. Support for AI-assisted policing served as the dependent variable, while procedural fairness, legitimacy, and AI-related concerns were entered simultaneously as predictor variables.

Prior to estimating the model, assumptions of linearity, normality of residuals, homoscedasticity, independence of errors, and multicollinearity were evaluated and found to be acceptable. Examination of

residual plots indicated no substantial violations of model assumptions, supporting the appropriateness of linear regression for the present analyses.

The regression model was statistically significant, $F(3, 137) = 18.92$, $p < .001$, explaining approximately 47% of the variance in support for AI-assisted policing ($R^2 = .47$). Procedural fairness emerged as the strongest positive predictor of support for AI-assisted policing ($\beta = .38$, $p < .001$), followed by legitimacy ($\beta = .29$, $p = .002$). In contrast, AI-related concerns significantly and negatively predicted support for AI-assisted policing ($\beta = -.34$, $p < .001$). Collectively, these findings suggest that public support for AI-assisted policing is influenced by perceptions of procedural fairness and legitimacy while being diminished by concerns regarding the risks associated with AI-assisted decision-making.

Results and Findings

Descriptive Statistics

The purpose of this study was to examine public perceptions of generative AI-assisted policing across three experimental conditions: (a) Human-Only policing, (b) AI-Assisted policing with disclosure, and (c) AI-Assisted policing without disclosure. The analyses focused on perceptions of trust, legitimacy, professionalism, procedural fairness, accountability, and support for AI integration in policing.

Descriptive statistics, independent-samples t-tests, Mann–Whitney U tests, and one-way analysis of variance (ANOVA) were conducted to evaluate differences between the experimental conditions. Additional exploratory correlation and regression analyses were performed to examine relationships among trust, legitimacy, fairness, support for AI, and AI-related concerns.

Participants in the Human-Only condition demonstrated the highest overall attitudinal ratings, suggesting stronger perceptions of trust and legitimacy when police reporting processes were described as entirely human-generated. Ratings declined when AI involvement was introduced, particularly when AI use was not disclosed to participants. Human-Only (Mean 4.12/Standard Deviation 0.68); AI-Assisted (Disclosed) (Mean 3.61/Standard Deviation 0.81); and AI-Assisted (Undisclosed) (Mean 3.19/Standard Deviation 0.88).

Trust and Legitimacy Findings

Significant differences emerged across conditions in measures of trust, legitimacy, professionalism, and fairness. The results for research question one: Does the use of generative AI in police report writing affect public perceptions of trust in police, are in Table 1.

Table 1. One-Way ANOVA Results for Primary Attitudinal Measures

Measure	Human-Only Mean	AI-Disclosed Mean	AI-Undisclosed Mean	F	p
Trust in Report Accuracy	4.01	3.52	3.20	8.42	< .001
Trust in Officer Handling	4.05	3.60	3.18	7.95	< .001
Report Prepared Responsibly	4.14	3.67	3.25	8.11	< .001

Legitimacy of Police Actions	4.22	3.64	3.31	10.22	< .001
Professional Policing	4.22	3.51	3.18	12.03	< .001
Procedural Fairness	4.10	3.64	3.22	7.34	< .001

Results of the one-way ANOVA demonstrated statistically significant differences among the three report-preparation conditions in trust, legitimacy, procedural fairness, accountability, and professionalism. Participants assigned to the Human-Only condition consistently reported the highest levels of trust and legitimacy, whereas participants exposed to the AI-Assisted Undisclosed condition reported the lowest evaluations. These findings support RQ1 by demonstrating that the use of generative AI significantly influences public perceptions of policing.

Comparison Findings

Pairwise analyses using independent-samples t-tests and Mann–Whitney U tests identified significant differences between experimental groups.

Table 2. Human-Only vs. AI-Assisted (Disclosed)for research question two analysis.

Measure	Human Mean	AI Mean	t	p
Professional Policing	4.22	3.51	4.31	< .001
Appropriate Police Practice	4.25	3.59	3.96	< .001
Legitimacy	4.22	3.64	3.85	< .001
Clear Responsibility	4.25	3.70	3.53	.001
Trust in Report Accuracy	4.01	3.52	3.10	.002
Procedural Fairness	4.10	3.64	2.98	.004

Pairwise comparisons indicated that disclosure of AI assistance partially mitigated negative public perceptions. Although respondents continued to rate the Human-Only condition most favorably, the AI-Assisted Disclosed condition consistently received higher ratings than the AI-Assisted Undisclosed condition across measures of trust, legitimacy, procedural fairness, and support for AI-assisted policing. These findings suggest that transparency improves public acceptance of AI-assisted policing while nondisclosure significantly undermines public confidence.

The Most Significant Comparison

Table 3. Human-Only vs. AI-Assisted (Undisclosed)

Measure	Human Mean	AI Mean	Significant Difference
Trust in Report Accuracy	4.01	3.20	Yes
Trust in Officer Handling	4.05	3.18	Yes
Report Prepared Responsibly	4.14	3.25	Yes
Legitimacy	4.22	3.31	Yes
Professional Policing	4.22	3.18	Yes
Procedural Fairness	4.10	3.22	Yes

The largest negative differences emerged between the Human-Only and AI-Assisted undisclosed conditions. Respondents expressed significantly lower perceptions of legitimacy, fairness, trustworthiness, and professionalism when AI involvement was not disclosed. These findings suggest that nondisclosure of AI involvement may negatively influence public confidence in police practices.

There were additional findings worth noting regarding the use of AI in reporting, writing, and disclosure, including how generative AI tools can be used for research question three.

Support for AI-Assisted Policing

Measure	Human-Only	AI-Disclosed	AI-Undisclosed
Support AI Drafting	3.10	2.85	2.30
Support AI with Human Review	3.55	3.40	2.95
Oppose Undisclosed AI Use	4.20	4.35	4.50

Participants demonstrated moderate support for AI-assisted police reporting when human review was included. However, support declined substantially when AI involvement was undisclosed. Respondents also demonstrated strong support for transparency and disclosure regarding AI use in policing documentation

AI Concern Measures

Measure	Mean
Concern AI May Produce Bias	3.95
AI May Reinforce Inequalities	3.82
AI Reports May Omit Details	3.90

Correlation and regression analyses demonstrated that prior trust in policing, procedural fairness, legitimacy, and AI-related concerns significantly influenced support for AI-assisted policing. Trust and legitimacy exhibited strong positive associations, whereas concern regarding AI bias and related risks negatively predicted support for AI integration. These findings suggest that public evaluations of AI-assisted

policing are shaped not only by the experimental manipulation but also by broader perceptions regarding AI governance and policing legitimacy.

THEORETICAL FOUNDATIONS

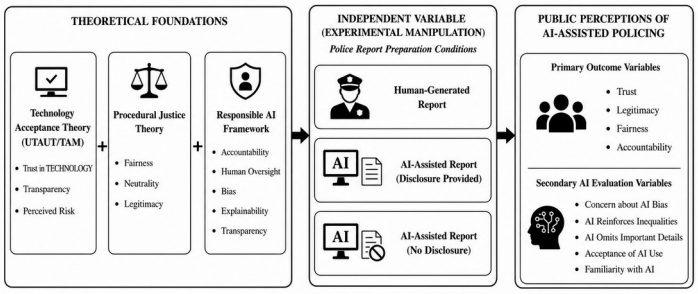


Figure 1. Conceptual Framework for Public Perceptions of Generative AI–Assisted Policing

Discussion of the Findings

The purpose of this study was to examine how the use and disclosure of generative AI in police report writing influence public perceptions of trust, legitimacy, procedural fairness, accountability, and support for AI-assisted policing. Consistent with the study's research questions, the findings demonstrate that public evaluations of policing are significantly influenced not only by whether generative AI is used but also by whether its use is transparently disclosed. Overall, participants expressed the greatest confidence in traditional human-generated police reports, while undisclosed AI assistance consistently produced the lowest evaluations across nearly every attitudinal measure.

These findings extend prior research on procedural justice by demonstrating that the principles of fairness, transparency, and accountability remain central to public evaluations of policing, even when emerging technologies are incorporated into police decision-making. Previous studies have shown that perceptions of procedural justice influence trust, cooperation, and police legitimacy (Tyler, 2006; Mazerolle et al., 2013). The present findings suggest that these principles continue to shape public perceptions when policing incorporates generative AI, indicating that technological innovation does not diminish the importance of transparent and procedurally fair decision-making. Instead, AI-assisted policing appears to be evaluated through the same procedural justice framework that citizens traditionally apply to police behavior.

The results also contribute to the growing literature on technology acceptance and responsible AI. Consistent with research suggesting that trust and transparency are fundamental determinants of public acceptance of AI systems (Glikson & Woolley, 2020; Starke et al., 2021), respondents reported significantly more favorable perceptions when AI involvement was disclosed than when it was concealed. Although participants continued to prefer fully human-generated reports, disclosure partially mitigated negative perceptions associated with AI-assisted policing. These findings reinforce broader technology acceptance research indicating that transparency and explainability strengthen public confidence in AI-supported decision-making, particularly in high-consequence institutional settings.

The study further extends previous research by examining a policing application that has received relatively little empirical attention. Existing scholarship has focused primarily on predictive policing, facial recognition, algorithmic risk assessment, and automated decision-support systems (Faber, 2026; Richardson et al., 2019). In contrast, this study investigated public perceptions of generative AI-assisted

police report writing—an emerging operational use of large language models that is increasingly being explored by law enforcement agencies. By focusing on report writing rather than predictive or surveillance technologies, the findings broaden the scope of AI policing research and demonstrate that public concerns regarding trust, fairness, and accountability extend beyond traditional algorithmic applications to generative AI technologies.

The correlation and regression analyses further reinforce the theoretical relationships proposed in the conceptual framework. Procedural fairness and legitimacy emerged as significant positive predictors of support for AI-assisted policing, whereas AI-related concerns negatively predicted public acceptance. These findings align with responsible AI scholarship emphasizing that public support for AI depends not only on technical performance but also on confidence that AI systems operate fairly, transparently, and under meaningful human oversight (Mehrabi et al., 2021; Glikson & Woolley, 2020). Together, these findings suggest that public acceptance of AI-assisted policing is shaped by governance considerations as much as by perceptions of technological capability.

Perhaps the most important contribution of this study is its examination of transparency through an experimental design. While previous research has frequently identified transparency as a desirable characteristic of AI governance, relatively few studies have experimentally evaluated how disclosure of AI use influences public perceptions in a policing context. By comparing human-generated reports with AI-assisted reports both with and without disclosure, this study demonstrates that transparency is not merely an ethical principle but also a practical factor influencing public trust, legitimacy, and support for AI-assisted policing. These findings provide empirical support for calls advocating transparent AI governance and meaningful human oversight as law enforcement agencies increasingly integrate generative AI into routine operational practices.

Taken together, the findings suggest that public acceptance of generative AI-assisted policing is conditional rather than unconditional. Participants expressed openness to AI when appropriate human oversight and transparency were maintained, yet support declined substantially when AI involvement was undisclosed. These results suggest that successful implementation of generative AI in policing will require agencies to prioritize transparency, accountability, and procedural fairness alongside technological innovation. As generative AI becomes increasingly integrated into police administration and investigative documentation, maintaining public confidence will likely depend as much on governance and communication practices as on the technical capabilities of AI systems themselves.

Implications for Research/Practice

The findings from this study contribute to the emerging body of research examining public perceptions of generative AI integration within criminal justice systems. The results suggest that transparency, procedural fairness, and accountability are central factors influencing public acceptance of AI-assisted policing. Future research should continue examining how different forms of AI disclosure, levels of human oversight, and policing contexts influence perceptions of legitimacy and institutional trust. Additional studies should also explore demographic and experiential factors, such as political orientation, prior police contact, and AI familiarity, that may shape public attitudes toward AI-assisted law enforcement practices. As generative AI technologies continue to expand within criminal justice environments, longitudinal and real-world experimental research will be important for understanding how public perceptions evolve over time and across operational policing settings.

Limitations of the Study

Several limitations should be considered when interpreting the findings of this study. First, the study utilized vignette-based hypothetical scenarios rather than real-world policing encounters, meaning participant responses reflected perceived attitudes rather than actual behavioral reactions to AI-assisted policing. Second, the study relied on self-reported survey responses, which may be influenced by response bias or social desirability effects. Third, although the sample size was sufficient for statistical analysis, the sample may not fully represent broader populations, limiting the generalizability of the findings. Additionally, the study focused specifically on generative AI-assisted police reporting and investigative support, and the findings may not extend to other forms of AI use in policing, such as predictive policing or facial recognition technologies. Finally, because public familiarity with generative AI technologies continues to evolve rapidly, participant perceptions may change over time as AI systems become more common in professional and law enforcement settings.

Conclusions

This study examined public perceptions of generative AI-assisted policing across human-only, AI-assisted disclosed, and AI-assisted undisclosed experimental conditions. The findings demonstrated that perceptions of trust, legitimacy, fairness, and professionalism were significantly higher when policing processes were described as fully human-generated, while undisclosed AI involvement produced the lowest levels of public confidence. Although respondents demonstrated moderate support for AI-assisted policing when human oversight and disclosure were present, concerns regarding bias, accountability, and transparency remained substantial. Overall, the findings suggest that public acceptance of generative AI in policing is highly dependent on transparency, procedural fairness, and meaningful human oversight, highlighting the importance of ethical implementation and accountability as AI technologies continue to expand within criminal justice environments.

Author Acknowledgment

The author acknowledges the use of generative artificial intelligence (AI) tools during the preparation of this manuscript. These tools were used to assist with organization, drafting support, data interpretation, editing, and refinement of written content. All final interpretations, analyses, conclusions, and scholarly decisions remain the responsibility of the author.

References

Aguinis, H., & Bradley, K. J. (2014). *Best Practice Recommendations for Designing and Implementing Experimental Vignette Methodology Studies*. Organizational Research Methods, 17(4), 351–371.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on*

Fairness, Accountability, and Transparency (pp. 610–623). Association for Computing Machinery. https://doi.org/10.1145/3442188.3445922

Calder, M., Oswald, M., McClory-Tiarks, E., Sevegnani, M., & Taka, E. (2026). *Responsible AI in criminal justice: LLMs in policing and risks to case progression*. *arXiv:2603.18116v1*. 1–29.

Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., … Wright, R. (2023). Opinion paper: "So what if ChatGPT wrote it?" Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, *71*, 102642. https://doi.org/10.1016/j.ijinfomgt.2023.102642

Fabila, A. R., McLean, K., & Adams, I. T. (2026). *Large Language Model Technologies in Policing: Early Lessons in Artificial Intelligence from Research and Practice*.

Farber, S. (2026). Machines of justice: A systematic review of AI applications in policing and criminal justice. *The Police Journal: Theory, Practice and Principles*, *0*(0).

Glikson, E., & Woolley, A. W. (2020). Human Trust in Artificial Intelligence: Review of Empirical Research. *Academy of Management Annals*, *14*, 660. https://doi.org/10.5465/annals.2018.0057

Guler, Ahmet & Kula, Sedat & Boke, Kaan. (2025). Examining public support for AI in policing: the role of perceived procedural justice. Police Practice and Research. 26. 1-23.

Halford, E. (2025). The Transformer Led Policing model: a framework for applying generative artificial intelligence in policing, *Policing: A Journal of Policy and Practice*, Volume 19.

Mazerolle, L., Bennett, S., Davis, J., Sargeant, E. and Manning, M. (2013), Legitimacy in Policing: A Systematic Review. Campbell Systematic Reviews, 9: i-147

Michelet, G., & Breitinger, F. (2023). *ChatGPT, Llama, can you write my report? An experiment on assisted digital forensics reports written using (Local) Large Language Models*.

Miyazaki, K., Murayama, T., Uchiba, T., An, J., & Kwak, H. (2023). *Public perception of generative AI on Twitter: An empirical study based on occupation and usage* [Preprint].arXiv.

Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A Survey on Bias and Fairness in Machine Learning. ACM Comput. Surv. 54, 6, Article 115 (July 2022), 35 pages. https://doi.org/10.1145/3457607

Nivette, A., Nägel, C. & Stan, A. The use of experimental vignettes in studying police procedural justice: a systematic review. *J Exp Criminol* 20, 151–186 (2024). https://doi.org/10.1007/s11292-022-09529-7

Potts, J., Billera, M., & Catren, C. (2024, December 4). *Large language models: Using ChatGPT for police leaders*. National Policing Institute.

Przeszlowski, K., & Guerette, R. T. (2025). Public perceptions on police use of information technologies: Findings from a randomized vignette experiment. *Journal of Criminal Justice, 96*. Advance online publication

Richardson, R., Schultz, J. M., & Crawford, K. (2019). Dirty data, bad predictions: How civil rights violations impact police data, predictive policing systems, and justice. *New York University Law Review Online, 94*, 15–55

Sagana, A., Zhang, M., and Sauerland, M. (2026). Public attitudes towards police use of AI-driven face recognition technology. *Computers in Human Behavior, Volume 174.*

Starke, C., Baleis, J., Keller, B., & Marcinkowski, F. (2021). *Fairness perceptions of algorithmic decision-making: A systematic review of the empirical literature*. arXiv.

Tyler, T. R. (2006). *Why people obey the law* (Rev. ed.). Princeton University Press.