Preliminary Study of TexAI: Where Adaptive AI Reimagines Law Enforcement Training

Shreyas Kumar, Charvi Vohra, Manas Rai and Arya Singh
Texas A&M, College Station, USA

shreyas.kumar@tamu.edu
charvivohra@gmail.com
manaskrai10@gmail.com
aryasingh2385@gmail.com

Abstract: Law enforcement agencies today operate at the frontline of data-sensitive decision-making, yet their training systems remain alarmingly analog. This gap has far-reaching consequences: The Police Department unintentionally deleted over eight terabytes of digital evidence, affecting nearly 17,000 criminal cases and causing significant public backlash and judicial delays (NBC 5 Dallas-Fort Worth, 2019). The root of this crisis lies not in technology alone, but in an outdated training paradigm that fails to prepare officers for the ethical, operational, and procedural demands of an AI-driven society. This paper explores how adaptive, explainable AI (XAI) can reframe the relationship between law enforcement and digital governance. We present TEXAI (XAI-powered Knowledge Base for Texas Law Enforcement), an AI-powered prototype built to modernize cybersecurity training in policing. Developed through user interviews and field research, the app combines real-time regulation updates with personalized, scenario-based microlearning-targeting a key challenge: officers forgetting or misunderstanding complex, evolving legal protocols. Our research examines how integrating XAI principles into law enforcement workflows introduces not only technological efficiency but critical epistemological transparency, fostering institutional accountability. We situate this intervention in the broader context of AI's role in public-sector transformation, arguing that ethical deployment of adaptive systems is essential to restoring public trust and preventing catastrophic human error. TEXAI also functions as a case study for how context-aware, role-specific AI tools can evolve through participatory design-responding to both human vulnerability and structural inefficiency. We contrast our solution with existing national systems such as PoliceOne Academy and Axon Academy, highlighting a novel intersection between AI explainability, justice system integrity, and digital literacy. The implications extend beyond law enforcement: in demonstrating how adaptive AI can personalize and democratize professional training in real time, we propose a scalable model for AI's responsible integration into high-stakes, socially critical domains. This work contributes to growing discourse around ethical AI, resilience in digital infrastructure, and the future of labor in AI-mediated institutions.

Keywords: XAI, Law enforcement training, Texas policing, Adaptive AI, Digital governance

1. Introduction

If the financial cost of cybercrime were compared to global economies, it would rank third, behind only the United States and China. The growing scale of these losses reflects how limited defenses and fragmented oversight have encouraged malicious actors to exploit vulnerabilities in increasingly digitized systems. While cybersecurity is often framed as a technical issue, it also exposes a broader challenge: institutions remain unprepared to respond effectively to complex risks in real time. This lack of preparedness is not confined to cyberattacks but extends to many areas where regulation, training, and decision-making intersect.

Recent local incidents illustrate these vulnerabilities. In 2023, the Dallas Police Department lost twenty-two terabytes of critical case data due to a failed transfer, while the Texas Department of Transportation suffered a ransomware attack that disrupted essential operations. These failures revealed more than technological weaknesses; they highlighted a systemic gap in training and oversight that directly impacts public safety. Beyond technology, first responders and officials are often expected to make rapid decisions while navigating intricate legal and procedural frameworks. Under stress, even experienced professionals can make mistakes with severe consequences for civilians.

These realities underscore the need for adaptive systems that go beyond cybersecurity alone. TEXAI builds on lessons from digital vulnerabilities but applies them to a wider spectrum of regulatory, operational, and training challenges. By providing real-time guidance, context-sensitive decision support, and adaptive learning tools, the system addresses the broader issue of how to prepare officers for both routine and high-stakes scenarios in a rapidly evolving environment.

Given these challenges, the goal of this research is to create a modular AI-powered system called TEXAI that enhances resilience, streamlines access to critical protocols, and supports informed decision-making in high-pressure situations. TEXAI is designed to directly address the gaps exposed by the Dallas data loss, offering real-time guidance and secure access to procedures that officers can trust. For instance, during a ransomware lockdown, TEXAI could provide step-by-step instructions to ensure correct responses, minimize operational disruption, and reduce potential harm. This research also aims to evaluate how effective an AI-based solution is compared to existing, traditional methods of protocol access and training. By measuring both efficiency and reliability across these approaches, the study not only develops and tests TEXAI as a practical tool but also gauges the broader role AI can play in improving officer preparedness and bridging the divide between complex systems and their end-users.

In a discussion with a local law enforcement officer, we discovered the legal and cybersecurity training the officers receive every year. The system utilized at the moment is an outdated, powerpoint presentation which is reviewed right before a standardized test held once a year. The officer interviewed mentioned how "...a trend in how officers learn is that we're very hands-on and work better in interactive environments. An app which allows us to discuss our questions and ask any the second a situation arises, would definitely be beneficial… there are so many legal codes we have to memorize for real life situations which makes it easy to make mistakes, so this would remove that risk in our career". We took this feedback in mind when designing our model, especially when it came to the AI chatbot and the way it communicated in different situations with an individual.

TEXAI will feature two tailored interaction modes: one optimized for real-time situations with speed and clarity, and another designed for educational settings with deeper, more detailed explanations. While officers can still toggle between these modes through a simple interface tab, the AI also uses on-device intelligence to automatically detect the nature of a request and respond in the most appropriate mode. The model's tone will be a balanced hybrid, concise and direct during emergencies while remaining professional and polite. Users can request elaboration at any point, ensuring critical information is always accessible without unnecessary delay.

2. Review of Literature

Several platforms provide cybersecurity or compliance-related training for government and enterprise users, including law enforcement, but most fall short in delivering context-specific, adaptive learning experiences suited to the operational realities of police work. Research on adaptive learning systems demonstrates that personalization significantly improves knowledge retention, particularly in professional development settings (Khosravi et al., 2022; Brusilovsky and Millán, 2007). Rafferty et al. (2016) further highlight that adaptive teaching using partially observable Markov decision processes (POMDPs) allows training systems to dynamically adjust to learner performance, though these approaches remain underutilized in policing contexts.

Explainable AI (XAI) frameworks have also been increasingly used in education and compliance systems to support transparency, fairness, and trust (Adadi and Berrada, 2018). Doshi-Velez and Kim (2017) emphasise that interpretability in machine learning is essential for accountability in high-stakes domains, while Lum and Isaac (2016) caution against the risks of opaque predictive policing tools. This suggests strong potential for combining adaptive learning with explainable AI to create training solutions that are both effective and accountable in public safety contexts.

Privacy and security are equally central to technology adoption in law enforcement. Cavoukian (2011) outlines the Privacy by Design framework, which embeds data protection into system architecture from the outset, a principle echoed in the CJIS Security Policy (CJIS Division, 2020). Similarly, national and federal standards such as the Federal Information Security Management Act (FISMA, n.d.), the Federal Risk and Authorization Management Program (FedRAMP, n.d.), and the National Institute of Standards and Technology s Cybersecurity Framework (NIST, n.d.) establish regulatory benchmarks for secure data handling in government systems. Integrating such standards into training technologies ensures compliance while also building trust among officers and the public.

From a systems design perspective, Chen et al. (2019) argue that modular architectures increase adaptability and maintainability in industrial platforms, insights directly applicable to training solutions requiring scalability across diverse law enforcement agencies. Kumar et al. (2025) likewise emphasise the link between cybersecurity maturity and government efficiency, underscoring the importance of resilient, adaptable platforms in the public sector.

Finally, funding and policy incentives shape the adoption of law enforcement technologies. Federal grant programs such as those from the Edward Byrne Justice Assistance Grant Program (n.d.) and the Office of Community Oriented Policing Services (COPS Office, n.d.) illustrate the government's role in supporting training and technology modernization efforts. Embedding adaptive and explainable systems within this policy and funding environment positions new solutions like TEXAI not only as technically innovative but also as aligned with broader institutional frameworks. In sum, existing literature demonstrates that adaptive learning and explainable AI each contribute valuable mechanisms for enhancing training effectiveness and accountability, yet they remain underexplored in law enforcement contexts. Prior research on privacy, compliance, and modularity underscores the need for systems that are secure, transparent, and adaptable, while policy frameworks and funding mechanisms establish pathways for real-world deployment. This convergence highlights the research gap that TEXAI addresses: delivering an adaptive, explainable, and regulation-compliant training platform tailored to the unique demands of police work.

3. Related Work

PoliceOne Academy by Lexipol, widely used among U.S. law enforcement, offers general modules on legal updates, tactical skills, and situational awareness. Some content touches on cybersecurity, but it lacks structured, role-specific modules on data privacy, digital evidence handling, or cyber incident response. It also does not adapt content based on learner performance or evolving digital threats.

KnowBe4 provides comprehensive training on phishing, social engineering, and digital hygiene but is primarily designed for corporate environments. It does not incorporate public sector regulatory frameworks or simulate law enforcement-specific scenarios, such as unauthorized data access during investigations or breaches involving body-worn camera footage.

Similarly, Infosec IQ and Proofpoint Security Awareness provide robust training content but rely on linear, non-adaptive learning structures. These platforms often overlook operational pressures faced by officers and do not provide guidance on state-specific laws or constitutional privacy protections.

TEXAI addresses these gaps by combining adaptive learning technologies with XAI-driven explainability. It emphasizes real-world simulations, compliance with data privacy laws, and role-specific feedback, distinguishing it from existing law enforcement and general cybersecurity education platforms. Unlike traditional training tools, TEXAI is designed to assist officers beyond the classroom, offering real-time tactical guidance during high-pressure situations. For example, when an officer responds to a mugging or another urgent scenario, TEXAI can provide step-by-step procedural guidance, enabling faster, informed, and legally compliant decision-making. This transforms TEXAI into a trusted assistant in live operations, supporting officers with critical insights when it matters most.

4. Methodology

To develop TEXAI as an effective, context-aware AI solution for Texas law enforcement, we adopted a modular and adaptive design approach emphasizing regulatory alignment, real-time usability, and operational relevance. Our methodology integrated technical, educational, and procedural considerations, grounded in both field research and established AI practices (Goodfellow et al., 2016; Russell & Norvig, 2021).

The development process included user-centered research, involving interviews and surveys with law enforcement officers to understand their practical needs, analysis of common cybersecurity and procedural incidents, and iterative prototyping to ensure TEXAI aligned with real-world operational demands (Nielsen, 1993). Each stage was designed to bridge the gap between complex legal and procedural knowledge and the daily decision-making needs of frontline personnel.

A critical component of the methodology is the fine-tuned adapter, developed using Apple's Foundation Models Adapter framework. This adapter enhances the base AI model's domain-specific reasoning without retraining the full model (Houlsby et al., 2019). It is trained on a curated dataset that includes regulation-based queries, real-world law enforcement scenarios, and procedural decision points. To ensure security and compliance, adapters are evaluated on-device, never made public, and operate within a closed privacy-preserving loop, ensuring both responsiveness and cybersecurity integrity.

The methodology further includes the design of three core interaction modes, Informational, Quiz, and Simulation, each serving distinct pedagogical and operational purposes. Informational mode answers officer questions on laws and procedures. Quiz mode delivers multiple-choice and free-response assessments, evaluating responses against a verified dataset. Simulation mode engages users in realistic scenario-based training, assessing actions in accordance with departmental regulations (Salas et al., 2009). A keyword-based control allows concise or detailed responses as needed, improving interaction fluidity.

Conversational memory was selectively implemented for Informational and Simulation modes, allowing TEXAI to maintain context across multi-turn dialogues while conserving resources. A failsafe mechanism clears memory if usage exceeds system thresholds, notifying users without disrupting the interface. Additionally, a long-term memory module is under development to generate session summaries, progress tracking, and training coverage analytics, reinforcing personalization and long-term learning continuity.

5. System Architecture

TEXAI is built with a modular architecture to support rapid updates, legal adaptability, and real-time operational awareness. At the core is the fine-tuned Foundation Model adapter, trained on curated datasets of Texas law enforcement regulations, procedural guidance, and real-world incident case studies. This adapter enables domain-specific learning without retraining the base model and is deployed dynamically through secure application updates.

An onboard controller manages the three primary interaction modes, directing queries to Informational, Quiz, or Simulation based on context and keywords. The keyword system prioritizes concise responses when prompts contain quick while defaulting to a more explanatory tone otherwise.

Conversational memory operates selectively to maintain context in Informational and Simulation modes while conserving system resources. The system discards irrelevant history and triggers a failsafe if memory thresholds are exceeded, notifying the officer while preserving on-screen content. Long-term memory will track learning progress, generating session summaries and checklists to improve training continuity.

Evaluation of TEXAI focuses on three main dimensions: effectiveness, adaptability, and user engagement. Effectiveness will be measured by the model's ability to provide accurate, context-specific responses to regulatory queries and procedural scenarios compared to traditional training methods. Key metrics include accuracy, response time, and comprehension.

Adaptability will be assessed through scenario-based testing in both controlled training environments and simulated field situations. Officers will interact with Informational, Quiz, and Simulation modes, and performance metrics will be collected for scenario completion, decision-making accuracy, and adherence to laws and procedures. The system's ability to integrate and deploy updated adapters will also be evaluated.

User engagement will be analyzed qualitatively through officer feedback and quantitatively through analytics, including mode-switching frequency, requests for expanded explanations, and interactions with long-term memory summaries. Evaluations will also consider the system's support for decision-making under high-pressure conditions and its ability to maintain compliance and procedural accuracy.

6. Design of TEXAI

The design of TEXAI integrates modular architecture, domain-specific adapters, and flexible interaction modes to meet law enforcement needs. The Informational mode allows officers to obtain immediate answers to regulatory queries. Quiz mode reinforces knowledge through structured exercises, while Simulation mode enables real-time scenario practice with context-aware evaluation. Interaction modes can adapt automatically based on the nature of the query using on-device intelligence. Concise or detailed responses are determined by keyword presence or operational context, ensuring that officers receive information appropriate to their current situation. The adapter system supports dynamic updates without retraining the base model, ensuring the system remains legally accurate and contextually relevant. Privacy is ensured through on-device evaluation of adapters and selective conversational memory.

7. Evaluations

Early models tested the AI's language capabilities and responsiveness. Initial results showed that the base model could manipulate language effectively but lacked specialization and context-awareness. Subsequent iterations added regulation-specific datasets, conversational memory, and structured modes, gradually improving accuracy, simulation capability, and adaptive response quality.

8. Testing and Results

To evaluate TEXAI's real-world performance, structured prompts were tested across Informational, Quiz, and Simulation modes. Responses were scored on a 10-point scale with four weighted criteria: relevance, accuracy, and compliance with Texas law (4 points), conciseness or detail depending on context (2 points), correct use of interaction mode (2 points), and effective use of conversational memory when applicable (2 points). There were 2 standalone prompts for informational style questions, 2 conversational prompts to track the bot's conversational memory and accuracy abilities, and 2 activity style prompts intended to activate quiz/simulation mode. The prompts, shown in Figure 1, were as follows:

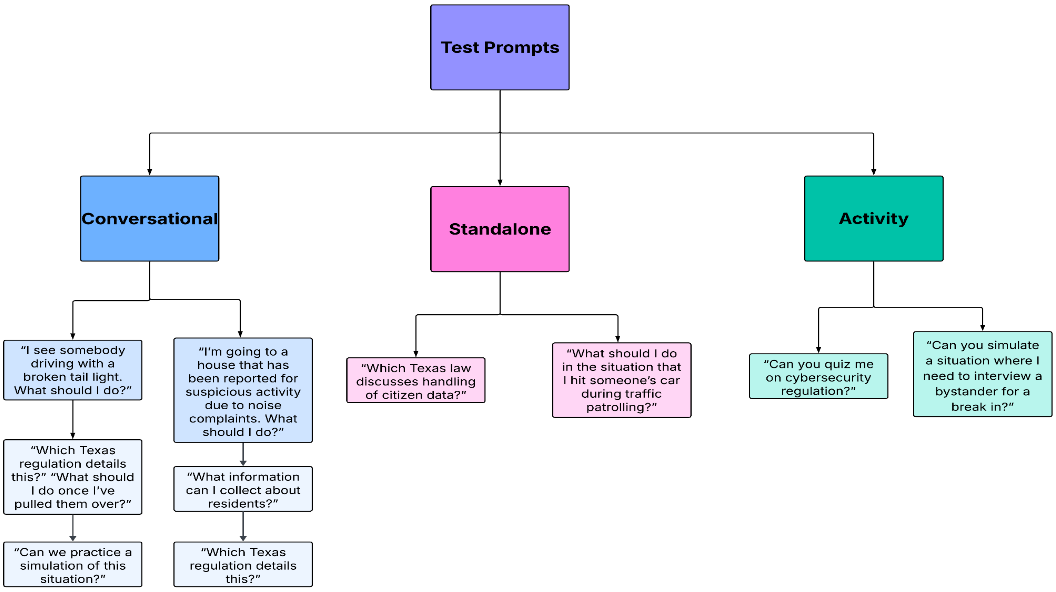

Figure 1: Prompts

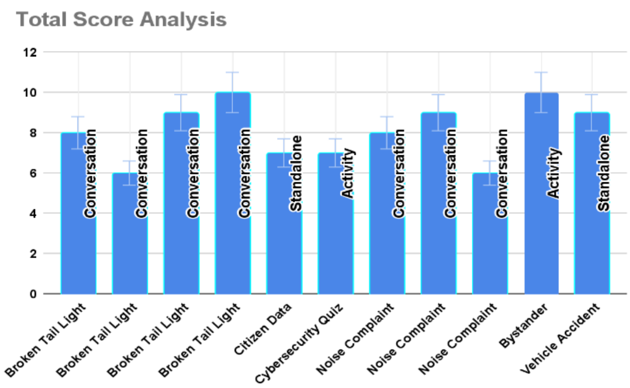

Figure 2: Total score analysis

Mode Score and Memory Score remained consistently strong across all scenarios, reflecting TEXAI's ability to switch modes correctly and maintain continuity during multi-step interactions. Conciseness and Detail varied more noticeably, performing better in activity-oriented and conversational tasks while showing occasional weaknesses in standalone informational prompts.

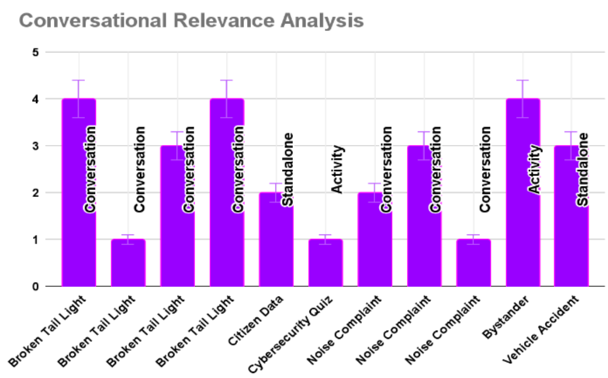

Figure 3: Conversational relevance Analysis

TEXAI generally maintained coherent and contextually appropriate dialogue, though performance was slightly weaker in scenarios requiring precise statutory grounding. This variation likely stems from the current absence of a fully integrated Texas law database, which limits the model's ability to cite legal code directly.

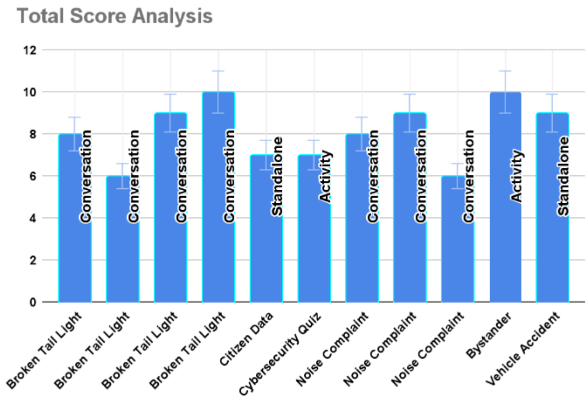

Figure 4: Total score analysis

The chatbot demonstrated strong practical guidance and simulation ability, particularly in roleplay scenarios such as traffic stops, noise complaints, and bystander interviews, where it emphasized officer professionalism, privacy, and safety. For example, in the broken tail light and noise complaint scenarios, the chatbot offered step-by-step

lawful guidance and effective context retention, but its legal referencing was often incomplete or imprecise. In contrast, simulation-based interactions such as interviewing a bystander produced excellent results, earning near-perfect scores due to structured feedback and context-appropriate professionalism. Standalone prompts, such as requests for Texas statutes on citizen data or cybersecurity regulation, highlighted the system's main weakness: statutes were either partially cited, misrepresented, or misapplied, resulting in lower accuracy scores. Overall, TEXAI consistently demonstrated high performance in adaptive coaching and situational practice but revealed clear limitations in legal citation reliability, suggesting the need for integration with more authoritative and structured legal databases.

9. Discussion

The results highlight TEXAI's dual strengths and limitations. Iterative design improved accuracy and context-awareness, while the modular adapter approach allowed legal updates without retraining the full model. Simulation testing confirmed TEXAI's potential as a real-time training and guidance tool: scenarios such as bystander interviews and noise complaints produced high-scoring outputs that emphasized professionalism, privacy, and officer safety. However, repeated testing also revealed a critical gap in legal referencing. In multiple cases, statutes were either incompletely cited or misapplied, reducing confidence in the system's regulatory precision. This weakness was traced to the informal nature of the prototype's training database rather than the model's reasoning process, suggesting that performance can be significantly improved through the integration of more comprehensive, structured legal datasets. Addressing this limitation will be central to TEXAI's future development and deployment in operational contexts.

Beyond this, the theoretical implications of TEXAI’s design and performance point towards a new model of AI-mediated institutional learning. Traditional police training relies on static, uniform instructions that do not reflect the distributed, rapidly evolving nature of modern law enforcement. By contrast, TEXAI operationalizes theories of situated cognition and experimental learning (Lave & Wenger, 1991; Kolb, 1984) by embedding procedural knowledge within real-world contexts and dynamically adjusting feedback based on user input. This adaptive interactivity transforms training from a memorization exercise into a cognitive partnership, where human expertise and AI reasoning co-develop situational fluency.

The system’s explainable AI foundation also aligns with theories of socio-technical accountability, which argue that transparency in algorithmic decision-making reinforces ethical behavior and public trust. From a theoretical standpoint, TEAI can also be understood as a prototype for adaptive governance infrastructure, a theory drawn from systems theory and public administration that emphasizes flexibility, resilience, and responsiveness in managing complex institutions. By allowing continuous adaptation to evolving laws, policies, and cyber threats, modular AI systems can strengthen institutional learning capacity. This aligns with emerging literature on AI-driven public sector transformation, which asserts that intelligent systems can reduce bureaucratic friction and improve compliance through real-time knowledge dissemination.

10. Recommendations and Future Work

The successful performance of TEXAI in conversational memory and simulation modes suggests distinct areas for immediate pilot deployment. We recommend three core actions for law enforcement agencies and training leaders: First, Prioritize Simulation-Based Training by integrating the TEXAI Simulation Mode to replace passive training (e.g., PowerPoint lectures) for complex procedural scenarios, such as traffic stops and bystander interviews. Second, Utilize TEXAI as a Procedural Coach, not a Legal Authority. Given its limitation in precise statutory grounding, TEXAI should be used primarily as a context-aware procedural assistant and ethical guidance coach in the field, not as the sole source for confirming specific Texas legal statutes for court preparation. Third, Support Adaptive Deployment using a phased rollout strategy that leverages the system's adaptive nature to personalize learning and enhance institutional accountability based on identified training gaps.

The findings and limitations revealed during testing necessitate specific directions for future work and research. Development efforts will focus on enhancing long-term memory performance, simulation fidelity, and dynamic response adjustment. Key metrics will include response accuracy, scenario completion, user comprehension, and real-time decision-making support. (We are currently waiting for IRB approval to discuss and begin testing various models with individuals in law enforcement to further develop the app and AI model.) The primary research recommendation is to Integrate Structured Legal Datasets by linking the AI model with an authoritative, structured legal code database to eliminate the critical gap in legal citation reliability, which is essential for operational trust. Researchers should also Conduct Comparative Efficacy Studies (formal comparative studies must be performed to measure TEXAI's long-term impact against traditional training methods and procedural error reduction) and Generalize the Model by exploring adapting the core TEXAI architecture (XAI-powered, adaptive microlearning) to other high-stakes public service domains like emergency medical services (EMS) or fire compliance training, demonstrating its scalability.

11. Conclusion

TEXAI provides a scalable, adaptive AI assistant for Texas law enforcement, combining modular architecture, fine-tuned adapters, multiple interaction modes, and memory systems. The platform bridges the gap between complex legal knowledge and practical operational demands, supporting both training and high-pressure field scenarios. Continued evaluation and iteration will ensure TEXAI remains a reliable, legally compliant, and context-aware tool, enhancing officer preparedness, procedural accuracy, and decision-making in real-world applications.

Ethics Statement: A local police officer was consulted to provide professional insights into training practices. Participation was voluntary, informed consent was obtained, and no personal or identifying information was recorded. The discussion focused solely on professional practices relevant to this study. The anonymous police officer was aware that his answers would be utilized to develop this research paper.

AI Declaration: Portions of this report utilized OpenAI's ChatGPT (GPT-5), which was used to refine wording & suggest formatting.

References

Adadi, A. and Berrada, M., 2018. Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). IEEE Access, 6, pp.52138-52160. Available at: https://doi.org/10.1109/ACCESS.2018.2870052 [Accessed 22 July 2025].

Axon Academy, n.d. What is Axon Academy? A virtual training platform offering product education through online courses, in-person sessions, and certifications for agency users. Available at: https://my.axon.com/s/article/What-is-Axon-Academy?language=en_US [Accessed 23 August 2025].

Brusilovsky, P. and Millán, E., 2007. User models for adaptive hypermedia and adaptive educational systems. In: P. Brusilovsky, A. Kobsa and W. Nejdl, eds. The Adaptive Web. Springer, pp.3-53. Available at: https://doi.org/10.1007/978-3-540-72079-9_1 [Accessed 22 July 2025].

Cavoukian, A., 2011. Privacy by design: The 7 foundational principles. Information and Privacy Commissioner of Ontario. Available at: https://www.ipc.on.ca/wp-content/uploads/resources/7foundationalprinciples.pdf [Accessed 22 July 2025].

Chen, J., He, X., Wen, J. and Xu, Y., 2019. Modular design for smart industrial platforms: A review and future trends. Journal of Industrial Information Integration, 15, pp.1-10. Available at: https://doi.org/10.1016/j.jii.2019.01.001 [Accessed 22 July 2025].

Community Oriented Policing Services (COPS) Office, n.d. Grants. Available at: https://cops.usdoj.gov/grants [Accessed 22 July 2025].

Criminal Justice Information Services (CJIS) Division, 2020. CJIS Security Policy Version 5.9. Federal Bureau of Investigation. Available at: https://www.fbi.gov/services/cjis/cjis-security-policy-resource-center [Accessed 22 July 2025].

Doshi-Velez, F. and Kim, B., 2017. Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608. Available at: https://arxiv.org/abs/1702.08608 [Accessed 22 July 2025].

Edward Byrne Memorial Justice Assistance Grant (JAG) Program, n.d. Program overview. Available at: https://bja.ojp.gov/program/jag/about [Accessed 22 July 2025].

Federal Information Security Management Act (FISMA), n.d. Overview of FISMA. Available at: https://www.dhs.gov/fisma [Accessed 22 July 2025].

Federal Risk and Authorization Management Program (FedRAMP), n.d. About FedRAMP. Available at: https://www.fedramp.gov/about/ [Accessed 22 July 2025].

Infosec, 2025. Infosec IQ Security Awareness Training. Available at: https://www.infosecinstitute.com/security-awareness/ [Accessed 22 July 2025].

KnowBe4, 2025. Security Awareness Training. Available at: https://www.knowbe4.com/ [Accessed 22 July 2025].

Khosravi, H., Kitto, K. and Knight, S., 2022. Personalized learning analytics: A review of the literature. Computers & Education: Artificial Intelligence, 3, p.100056. Available at: https://doi.org/10.1016/j.caeai.2022.100056 [Accessed 22 July 2025].

Kumar, S., Garg, A. and Niranjan, M., 2025. Enhancing government efficiency through cybersecurity hardening. In: Proceedings of the 26th Annual International Conference on Digital Government Research - dg.o. Porto Alegre City, Brazil: ACM. Available at: https://proceedings.open.tudelft.nl/DGO2025/article/view/1047 [Accessed 23 August 2025].

Lexipol, 2025. PoliceOne Academy. Available at: https://www.lexipol.com/solutions/policeone-academy/ [Accessed 22 July 2025].

Lum, K. and Isaac, W., 2016. To predict and serve? Significance, 13(5), pp.14-19. Available at: https://doi.org/10.1111/j.1740-9713.2016.00960.x [Accessed 22 July 2025].

NBC 5 Dallas-Fort Worth, 2019. Dallas Police Department accidentally deletes 8 terabytes of evidence, impacting thousands of cases. Available at: https://www.nbcdfw.com/news/local/dallas-police-department-accidentally-deletes-8-terabytes-of-evidence-impacting-thousands-of-cases/1879909/ [Accessed 22 July 2025].

National Institute of Standards and Technology (NIST), n.d. Cybersecurity Framework. Available at: https://www.nist.gov/cyberframework [Accessed 22 July 2025].

Proofpoint, 2025. Security Awareness Training. Available at: https://www.proofpoint.com/us/security-awareness-training [Accessed 22 July 2025].

Rafferty, A.N., Brunskill, E., Griffiths, T.L. and Shafto, P., 2016. Faster teaching by POMDP planning. Cognitive Science, 40(6), pp.1290-1332. Available at: https://doi.org/10.1111/cogs.12310 [Accessed 22 July 2025].

Texas Tribune, 2023. Texas Department of Transportation hit by ransomware attack, causing service disruptions. Available at: https://www.texastribune.org/2023/05/01/txdot-ransomware-attack-disruption/ [Accessed 22 July 2025].