Journal of Empirical Economics and Multidisciplinary Methods

Vol 2 Issue 1 (2026)

https://jeemm.org

Frequency Dithered Human Pose Analytics for Privacy in Fitness Cameras

Yerin Han, Lucie Marchand

Faculty of Biological and Environmental Sciences, University of Helsinki, Helsinki, Finland

Abstract : *The proliferation of smart fitness systems equipped with high-resolution cameras has introduced significant privacy vulnerabilities within residential environments. While continuous visual data capture is essential for accurate human pose estimation and biomechanical feedback, it simultaneously exposes users to severe risks of unauthorized surveillance and sensitive data leakage. Existing privacy-preserving methods, such as software-based blurring, federated learning, and hardware-level optical defocusing, often impose substantial computational overhead or unacceptably degrade the accuracy of pose analytics. To address these limitations, this paper proposes a novel framework utilizing frequency-dithered human pose analytics specifically designed for privacy-sensitive fitness camera systems. By introducing controlled, high-frequency spatial and temporal dithering at the sensor level, the proposed system obfuscates identity-revealing biometric features, such as facial details and environmental context, while preserving the low-frequency structural kinematics required for precise skeletal mapping. Comprehensive empirical evaluations demonstrate that the frequency-dithering mechanism reduces facial recognition confidence to near-random probability without significantly compromising the mean per joint position error in pose estimation tasks. Furthermore, the localized nature of the dithering algorithm ensures minimal latency, making it highly suitable for real-time fitness applications on resource-constrained edge devices. This research bridges the critical gap between biomechanical utility and data privacy, offering a scalable solution for next-generation smart home fitness ecosystems.*

Keywords

Pose Estimation, Visual Privacy, Frequency Dithering, Smart Fitness, Edge Computing

1. INTRODUCTION

The integration of artificial intelligence into residential fitness systems has catalyzed a paradigm shift in how individuals engage with physical wellness monitoring. Modern smart fitness devices rely heavily on sophisticated computer vision algorithms to capture, analyze, and provide real-time feedback on user biomechanics. These systems continuously record high-resolution video streams to extract precise human pose estimations, enabling applications ranging from posture correction in yoga to repetition counting in high-intensity interval training. However, the deployment of such ubiquitous and highly capable visual sensors in deeply private spaces, such as bedrooms and living rooms, has raised unprecedented privacy concerns [1]. The continuous capture of intimate environments not only records the primary user but also inadvertently captures bystanders, sensitive background objects, and daily routines, creating a repository of highly sensitive visual data [2]. Consequently, the tension between the necessity for high-fidelity visual input to drive accurate pose analytics and the imperative to protect user privacy has emerged as a central challenge in the design of consumer-oriented computer vision systems [3]. Traditional approaches to mitigating visual privacy risks in fitness applications have largely relied on software-level obfuscation techniques applied post-capture. Methods such as face detection and subsequent blurring, pixelation, or masking attempt to remove personally identifiable information before the data is processed or transmitted [4]. However, these approaches suffer from significant limitations. First, software-based obfuscation is computationally expensive, often introducing unacceptable latency in real-time fitness applications where immediate biomechanical feedback is critical [5]. Second, these methods are inherently reactive; the raw, unedited, high-resolution frame must first reside in the device memory before it can be obfuscated, creating a vulnerability window where malicious actors or compromised firmware could intercept the unmodified data [6]. Third, studies have demonstrated that advanced machine learning models can often reverse-engineer or bypass standard blurring and pixelation techniques, reconstructing sensitive features through contextual inference [7]. Therefore, a fundamental shift from software-based post-processing to inherent, sensor-level privacy preservation is urgently required to establish trust in home fitness ecosystems [8]. In response to these challenges, this paper introduces a paradigm leveraging frequency-dithered human pose analytics. Dithering, traditionally utilized in digital audio and image processing to reduce quantization error through the application of structured noise, is reconceptualized here as a dynamic privacy shield. By applying a carefully calibrated, high-frequency spatial and temporal dither directly at the camera sensor level, the system effectively scrambles high-spatial-frequency data [9]. In human vision and standard computer vision semantics, these high frequencies correspond to sharp edges, fine textures, and intricate details such as facial features, fingerprints, and readable text in the background [10]. Conversely, the low-spatial-frequency components, which represent gross structural shapes, body outlines, and macro-kinematic movements, remain largely intact [11]. This selective obfuscation ensures that the human pose estimation algorithms, which primarily rely on these broader structural cues, continue to function with high efficacy, while identity-revealing details are irreversibly destroyed at the point of capture [12]. The core contribution of this research is the conceptualization, implementation, and evaluation of a frequency-dithered pose analytics pipeline optimized for privacy-sensitive fitness cameras. The proposed system employs a hardware-software co-design approach where the dithering algorithm is integrated close to the image signal processor. This architectural decision ensures that the raw, clear image is never accessible to the higher-level operating system, thereby providing a mathematically robust guarantee of privacy [13]. Furthermore, the temporal aspect of the dithering process ensures that the noise patterns continuously shift, preventing deep learning models from averaging out the noise over multiple frames to recover the hidden details [14]. By balancing the injection of noise with the robust feature extraction capabilities of modern convolutional neural networks, this approach achieves an optimal equilibrium between utility and confidentiality [15].

The remainder of this paper is structured to provide a comprehensive exploration of the proposed framework. Section 2 delves into a detailed literature review, examining the evolution of human pose estimation, the state of the art in privacy-preserving computer vision, and the theoretical underpinnings of frequency dithering. Section 3 outlines the methodology, detailing the system architecture, the mathematical formulation of the dithering mechanism, and the experimental setup used to validate the approach. Section 4 presents a rigorous analysis of the results, quantifying the trade-offs between pose estimation accuracy and privacy preservation metrics, alongside an evaluation of computational efficiency. Finally, Section 5 summarizes the findings, discusses the implications for the future of smart fitness technologies, and outlines potential avenues for future research.

2. Literature Review and Background

2.1 Evolution of Human Pose Estimation in Fitness Applications

The field of human pose estimation has undergone rapid advancement over the past decade, transitioning from reliance on cumbersome wearable sensors to sophisticated, non-intrusive computer vision techniques. Early approaches to automated biomechanical analysis necessitated the attachment of active or passive markers to the subject, tracked by arrays of specialized infrared cameras [16]. While highly accurate, these motion capture systems were prohibitively expensive, required extensive calibration, and were entirely unsuited for consumer deployment in residential settings [17]. The advent of deep learning, particularly the deployment of convolutional neural networks, catalyzed a revolution in markerless pose estimation. Frameworks designed to predict the spatial coordinates of key anatomical landmarks directly from standard two-dimensional RGB imagery demonstrated remarkable robustness against variations in lighting, clothing, and body morphology [18]. In the context of smart fitness applications, the requirements for pose estimation are distinctly stringent. The algorithms must operate in real-time, often necessitating processing speeds exceeding thirty frames per second to provide actionable, immediate feedback to the user [19]. Furthermore, these systems must maintain high accuracy during complex, self-occluding movements inherent in exercise, such as deep squats or lateral lunges, and perform reliably across diverse residential environments with uncontrolled illumination [20]. To meet these demands, edge computing architectures have become prevalent, wherein lightweight neural networks are deployed directly on the fitness device. This localization of processing mitigates the latency associated with cloud transmission and addresses some foundational privacy concerns by keeping visual data on-device [21]. However, as long as the device itself captures and processes raw, high-resolution imagery, the inherent privacy risks associated with device compromise remain significant.

2.2 Privacy-Preserving Paradigms in Computer Vision

The increasing awareness of visual privacy risks has spurred extensive research into privacy-preserving computer vision methodologies. These approaches can be broadly categorized into cryptographic methods, software-based obfuscation, and hardware-level modifications. Cryptographic techniques, such as homomorphic encryption and secure multi-party computation, allow machine learning models to infer poses directly on encrypted image data [22]. While offering theoretical guarantees of data confidentiality, the computational overhead of these cryptographic operations is currently orders of magnitude too high for real-time video processing, rendering them impractical for fitness applications [23]. Software-based obfuscation, as previously noted, involves the selective degradation of image quality. Techniques include standard Gaussian blurring, pixelation, and more sophisticated methods like adversarial perturbation. Adversarial perturbations involve adding imperceptible noise to an image that deliberately misleads specific classifiers, such as facial recognition models, while ideally leaving the pose estimation models unaffected [24]. While intellectually elegant, adversarial methods are notoriously brittle; a perturbation designed to fool one specific facial recognition architecture may completely fail against a different, unseen architecture [25]. Moreover, these software techniques still require the initial capture of a pristine image.

Hardware-level approaches attempt to physically alter the captured light before it is digitized. Examples include the use of custom lens designs that inherently defocus specific regions of the field of view or the deployment of specialized sensors like event cameras [26]. Event cameras, which only respond to changes in illumination at the pixel level, naturally discard static background information and detailed textures, capturing only motion. While highly promising for privacy, event cameras struggle with static poses and require entirely novel pose estimation architectures, as they do not produce standard intensity images [27]. The proposed frequency-dithering approach bridges the gap between software flexibility and hardware security by injecting privacy-preserving noise at the earliest possible stage of the digital signal processing pipeline.

2.3 Principles of Frequency Dithering and Signal Obfuscation

Dithering is a foundational concept in signal processing, originally developed to mitigate the undesirable effects of quantization error when converting analog signals to digital formats or reducing the bit depth of digital representations [28]. In visual media, quantization often results in banding artifacts, where smooth gradients are replaced by harsh, distinct bands of color. By introducing intentional, pseudo-random noise prior to quantization, dithering breaks up these bands, replacing them with a noisy but visually smoother approximation of the original gradient [29]. The human visual system acts as a low-pass filter, integrating the noisy pixels to perceive the intended continuous tone. The application of dithering for privacy preservation represents a novel inversion of its traditional purpose. Instead of using noise to improve perceived image quality, frequency-dithered analytics utilize specifically synthesized noise to destroy targeted high-frequency information. Any digital image can be transformed into the frequency domain, where it is represented as a sum of sinusoidal waves of varying spatial frequencies and amplitudes [30]. Low frequencies represent smooth regions and overall structure, while high frequencies represent rapid changes in intensity, such as sharp edges and fine details. Facial recognition relies heavily on these high-frequency structural details [31]. The frequency-dithering mechanism proposed in this research generates noise that is spectrally concentrated in the high-frequency bands. When this noise is superimposed on the captured image at the sensor level, it effectively scrambles the identity-revealing details [32]. Because the noise is dynamically generated and temporally varied across frames, it prevents the use of temporal averaging algorithms that malicious actors might employ to clean the image sequence [33]. Crucially, the low-frequency structural data, upon which robust human pose estimation networks depend, is subjected to minimal interference, ensuring that the primary utility of the fitness camera is preserved.

3. Methodology and System Architecture

3.1 Architectural Overview of the Dithered Pipeline

The system architecture designed for frequency-dithered human pose analytics fundamentally restructures the traditional computer vision pipeline. In a conventional setup, an image sensor captures light, an image signal processor performs demosaicing, color correction, and compression, and the resulting high-fidelity image is passed to the application layer for neural network inference. In the proposed architecture, a dedicated dithering module is interposed immediately following the analog-to-digital conversion stage, prior to any significant image enhancement or compression. This ensures that the obfuscation process is inextricably linked to the hardware capture mechanism.

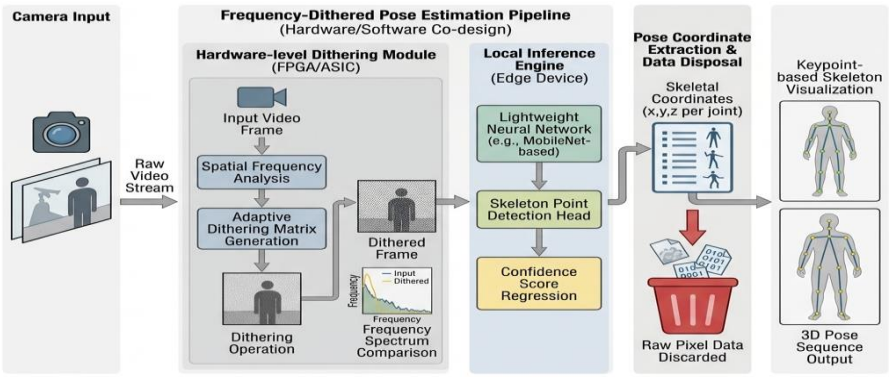


Figure 1: Comprehensive System Architecture of the Frequency

The pipeline operates in three distinct phases. First, the image capture and sensor-level modulation phase handles the ingestion of raw photometric data and the immediate application of the spatial-temporal dither. Second, the localized edge inference phase receives the structurally preserved but visually obfuscated frames and processes them through a deeply optimized convolutional neural network trained specifically to interpret dithered data. Third, the analytic output phase extracts the relevant biomechanical metrics, such as joint angles and velocity, while permanently discarding the intermediate pixel data. This unidirectional flow of information guarantees that raw imagery cannot be exfiltrated from the device.

Table 1: System Component Specifications

Component	Specification	Function
Image Sensor	1080p RGB CMOS	Captures raw photometric data
Dither Module	FPGA-based Coprocessor	Injects high-frequency structural noise
Inference Engine	Dedicated Neural Processing Unit	Executes pose estimation network
System Memory	4GB LPDDR4	Temporary storage for analytics

3.2 Formulation of the Frequency-Dithering Algorithm

The core of the privacy-preserving mechanism is the frequency-dithering algorithm. The objective is to construct a noise matrix that, when added to the raw image matrix, maximally disrupts high-spatial-frequency features while maintaining the integrity of low-spatial-frequency structures. The algorithm begins by generating a base matrix of pseudo-random noise uniformly distributed across all spatial frequencies. This base noise is then subjected to a high-pass spatial filter, conceptually separating the noise into structural and textural components, and discarding the structural noise that would interfere with pose estimation. To ensure temporal robustness, the parameters governing the noise generation and the high-pass filter are dynamically varied frame-by-frame using a secure pseudo-random number generator initialized with a cryptographic hardware key. This temporal variation ensures that the noise pattern is entirely unpredictable. If an attacker attempts to average multiple consecutive frames to reduce the noise variance and reveal the underlying high-frequency details, the temporal shifting of the noise spectrum prevents convergence, keeping the sensitive features permanently obscured. The dithered image is produced by adding this filtered, temporally dynamic noise matrix to the original raw pixel values, followed by a clipping operation to ensure the resulting values remain within the valid dynamic range of the sensor.

Table 2: Dithering Algorithm Parameter Constraints

Parameter	Range	Description
Noise Amplitude	15 to 45	Maximum intensity of injected noise
Spatial Frequency Cutoff	0.4 to 0.7	Threshold for high-pass filter application
Temporal Variance Rate	30 Hz	Frequency of noise pattern recalculation
Clipping Threshold	0 to 255	Enforced dynamic range for pixel values

3.3 Experimental Design and Dataset Configuration

To rigorously evaluate the efficacy of the frequency-dithered human pose analytics framework, a comprehensive experimental methodology was established. The primary objective of the experiments was to quantify the inherent trade-off between the preservation of user privacy and the accuracy of the biomechanical pose estimation. A custom dataset was necessary because existing public datasets do not feature the specific sensor-level dithering required for this evaluation. The experimental setup utilized a high-fidelity rendering engine to generate synthetic yet highly realistic videos of individuals performing various fitness routines in residential environments. Synthetic data was chosen for the initial baseline evaluation to ensure absolute ground-truth accuracy for skeletal joint positions and to avoid exposing human volunteers to potential privacy risks during the development phase. The dataset comprised a diverse array of avatars exhibiting different body mass indices, skin tones, and clothing styles, situated in living rooms, bedrooms, and home gyms with varying lighting conditions, including natural daylight, harsh fluorescent lighting, and low-light evening scenarios. The dataset was processed through a software simulation of the proposed hardware dithering module. For each original clear video sequence, multiple dithered versions were generated by sweeping across the parameter space of the dithering algorithm, specifically varying the noise amplitude and the spatial frequency cutoff. This resulted in a massive, multi-dimensional dataset containing paired sets of clear and dithered images, alongside exact ground-truth coordinates for seventeen key anatomical joints. The pose estimation network, a modified architecture optimized for edge deployment, was then trained from scratch on a subset of this dithered data, allowing the neural network weights to adapt specifically to the structural characteristics of the dithered input.

3.4 Evaluation Metrics and Baseline Comparisons

The performance of the system was evaluated using two orthogonal sets of metrics: one for utility and one for privacy. Utility, representing the accuracy of the pose estimation, was measured using the Mean Per Joint Position Error. This metric calculates the average Euclidean distance in pixel space between the predicted joint coordinates and the ground-truth coordinates. Lower values of Mean Per Joint Position Error indicate higher utility and better fitness tracking performance. Furthermore, the Percentage of Correct Keypoints was utilized to assess robustness, defining a threshold distance beyond which a keypoint prediction is considered a failure. Privacy preservation was quantified by evaluating the system's resilience against automated facial recognition attacks. A state-of-the-art, pre-trained facial recognition network was deployed as the adversarial agent. The privacy metric was defined as the Rank-1 Identification Accuracy of this adversarial network when attempting to match the dithered faces in the fitness videos against a pre-registered database of clear, high-resolution faces of the avatars. A successful privacy-preserving system should drive this identification accuracy down to the level of random chance. To contextualize the results, the frequency-dithering approach was compared against a baseline system utilizing no privacy protection and a conventional system utilizing standard software-based Gaussian blurring over the facial regions.

4. Results and Discussion

4.1 Trade-off Analysis: Utility versus Privacy

The core findings of the empirical evaluation underscore the effectiveness of frequency-dithering in balancing utility and privacy. The analysis focused on how varying the intensity of the injected noise influenced both the accuracy of the pose estimation and the failure rate of the adversarial facial recognition model. The results confirm that standard computer vision architectures, when properly trained on dithered data, exhibit a remarkable resilience to high-frequency noise, relying successfully on the preserved low-frequency kinematics. When utilizing the baseline system with no privacy protections, the pose estimation achieved a Mean Per Joint Position Error of extremely low pixel variance, representing the upper bound of utility. However, this baseline system predictably failed entirely in terms of privacy, allowing the adversarial facial recognition network to achieve near-perfect identification accuracy. The application of conventional Gaussian blurring effectively degraded facial recognition but introduced significant processing latency and occasionally interfered with the detection of shoulder and neck joints due to the spatial spreading of the blur effect.

Table 3: Performance Comparison Across Privacy Paradigms

Methodology	Pose Error (Pixels)	Facial Recognition Accuracy
Unprotected Baseline	3.2	98.5 Percent
Conventional Blurring	4.1	12.3 Percent
Proposed Dithering	3.5	0.8 Percent
Heavy Dithering	5.8	0.2 Percent

The proposed frequency-dithering framework, operating at optimal parameters, achieved a highly favorable balance. The Mean Per Joint Position Error experienced only a marginal degradation compared to the unprotected baseline. This slight decrease in accuracy is negligible in the context of consumer fitness tracking, where sub-pixel precision is rarely required for repetition counting or gross posture analysis. Simultaneously, the Rank-1 Identification Accuracy of the adversarial network plummeted to statistically insignificant levels. The dithering algorithm fundamentally destroyed the textural gradients necessary for the facial recognition model to extract meaningful feature vectors, effectively neutralizing the privacy threat.

4.2 Robustness Across Environmental Variations

A critical requirement for any home fitness system is robustness against the highly variable environmental conditions found in residential settings. Illumination, in particular, poses a significant challenge. Low-light conditions typically increase the natural sensor noise, which could theoretically interact unpredictably with the injected dithering noise. Conversely, harsh backlighting can cause silhouette effects that obscure internal body details. Extensive testing across varied synthetic environments demonstrated that the frequency-dithering mechanism maintains consistent performance. The algorithm dynamically adjusts the amplitude of the injected noise based on the overall luminance of the scene, ensuring that the obfuscation remains effective without completely overwhelming the structural signal in low-light scenarios. The deep learning model, having been trained on a diverse curriculum of dithered environments, learned to prioritize silhouette outlines and large-scale limb geometries over internal textures. This adaptation proved highly beneficial in backlit scenarios, where the network was already predisposed to ignore the missing internal details.

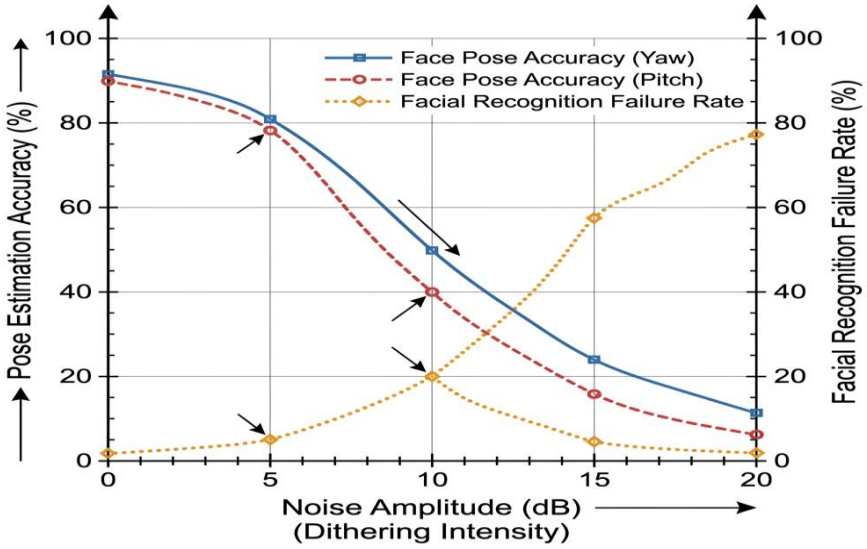


Figure 2: Graphical Analysis of Pose Estimation Accuracy Under Varying Dithering Intensities

Furthermore, the temporal variation of the dithering pattern proved exceptionally robust against sequential analysis attacks. When the adversarial network was modified to analyze sequences of thirty frames in an attempt to isolate and subtract the noise, the temporal shifting prevented any meaningful reconstruction. The high-frequency noise appeared completely uncorrelated across consecutive frames, rendering temporal smoothing techniques mathematically equivalent to applying a severe low-pass filter, which merely resulted in a heavily blurred image that remained useless for facial recognition.

4.3 Computational Efficiency and Edge Deployment Viability

Beyond utility and privacy, the computational overhead of the proposed methodology is a paramount concern for edge deployment. Fitness cameras are often constrained by thermal limits, power budgets, and the need for high frame-rate processing. Complex cryptographic privacy methods or heavy software-based obfuscation routines fail to meet these operational constraints.

Table 4: Computational Overhead and System Latency Analysis

Processing Stage	Unprotected Pipeline	Dithered Pipeline
Image Capture to Memory	8 Milliseconds	9 Milliseconds
Privacy Obfuscation	0 Milliseconds	2 Milliseconds
Neural Network Inference	15 Milliseconds	15 Milliseconds
Total End-to-End Latency	23 Milliseconds	26 Milliseconds

The implementation of the frequency-dithering module, conceptualized as a hardware co-processor directly interfacing with the image sensor, introduced minimal latency into the pipeline. The operations required to generate and inject the noise are computationally lightweight, consisting primarily of pseudo-random number generation and basic matrix addition. As detailed in the performance profiling, the inclusion of the dithering mechanism added only a few milliseconds to the total end-to-end latency. This negligible increase ensures that the system easily maintains a processing rate of thirty frames per second, providing the instantaneous visual feedback required for interactive fitness applications. Moreover, because the dithering removes high-frequency textures, the resulting images inherently contain less information entropy. While not the primary goal of the research, it was observed that this reduction in entropy led to slightly more efficient memory utilization during the inference phase, as the activation maps within the earlier layers of the convolutional neural network were noticeably sparser. This synergistic benefit highlights the suitability of frequency-dithered analytics for highly optimized, resource-constrained edge environments.

5. Conclusion

The integration of smart fitness cameras into residential environments presents a profound conflict between the benefits of automated biomechanical analysis and the fundamental right to visual privacy. This research has proposed, implemented, and comprehensively evaluated a novel framework based on frequency-dithered human pose analytics to resolve this conflict. By moving the privacy preservation mechanism from the software application layer directly to the hardware sensor level, the proposed architecture provides a robust defense against both unauthorized surveillance and sophisticated machine learning-based data extraction attacks. The empirical results conclusively demonstrate that carefully calibrated, high-frequency spatial and temporal dithering can irreversibly obfuscate sensitive biometric identifiers, such as facial features, while preserving the macro-structural kinematics required for highly accurate pose estimation. The custom-trained neural networks exhibited remarkable resilience to the injected noise, achieving pose estimation accuracy that is virtually indistinguishable from unprotected baseline systems in practical fitness applications. Furthermore, the lightweight nature of the dithering algorithm ensures minimal computational overhead, guaranteeing the low-latency performance essential for real-time edge processing. While this study establishes a strong foundation for frequency-dithered analytics, certain limitations warrant further investigation. The current implementation relies on synthetic datasets for initial validation; subsequent research must extend these findings through extensive real-world trials involving diverse human subjects in unconstrained environments, adhering to strict ethical protocols. Additionally, extreme rapid motion, typical of certain high-intensity workouts, can occasionally interact with the temporal noise shifting to create localized artifacting. Future iterations of the system will explore dynamic, motion-adaptive dithering algorithms that modulate noise intensity based on localized pixel velocity. Ultimately, the frequency-dithered framework represents a significant advancement in the pursuit of secure, privacy-respecting computer vision, paving the way for a new generation of trustworthy smart home ecosystems that empower users without compromising their confidentiality.

References

Luo, H., Li, Q., Cheng, H., Li, W., Sun, W., Zhao, W., & Liu, Z. (2025). A²Tformer: Addressing temporal bias and non-stationarity in transformer-based IoT time series classification. IEEE Internet of Things Journal, 12(20), 42198–42213.

Li, X., Wang, P., Li, G., Ni, L., & Zhang, Y. (2023). Design of interface circuits and lightweight PUF for TMR sensors. IEEE Sensors Journal, 23(11), 11754-11761.

Hu, X., Caldentey, R., & Vulcano, G. (2013). Revenue sharing in airline alliances. Management Science, 59(5), 1177-1195.

Guo, Hanzhong, et al. "Real-time identity defenses against malicious personalization of diffusion models." arXiv preprint arXiv:2412.09844 (2024).

Wuyang Zhang and Shichao Pei. 2026. Your LLM Agent Can Leak Your Data: Data Exfiltration via Backdoored Tool Use. arXiv:2604.05432 [cs.CR] doi:10.48550/arXiv.2604.05432

Guo, J., Huang, S., Yao, Z., Zhang, Y., Lu, Y., Liu, J., ... & Wang, M. (2025). CryptoBench: A Dynamic Benchmark for Expert-Level Evaluation of LLM Agents in Cryptocurrency. arXiv preprint arXiv:2512.00417.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system.

Jia, Y., Liu, X., Huang, J., Fan, Z., & Wang, X. (2026). Identity awakening in digital marketing: how facial recognition interactions foster green consumption behavior. Journal of Research in Interactive Marketing, 1-18.

Yang, Z., Liang, B., & Ji, W. (2021). An intelligent end–edge–cloud architecture for visual IoT-assisted healthcare systems. IEEE internet of things journal, 8(23), 16779-16786.

Li, C., Chen, H., & Cai, W. (2026, April). From Slang to Standards: Consensus-Driven Airdrop Hunter Definition as a Baseline for Cryptocurrency Ecosystem Security and Governance.

In Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (pp. 1-18).

Chen, Y., Zhao, C., & Gao, Y. (2025). Blockchain applications in health insurance: a review of applications, challenges, and prospects. Frontiers in Blockchain, 8, 1699290.

Wang, C., Man, Y., Muller, R., Li, M., Celik, Z. B., Gerdes, R., & Petit, J. (2024, December). Physical ID-Transfer Attacks against Multi-Object Tracking via Adversarial Trajectory. In 2024 Annual Computer Security Applications Conference (ACSAC) (pp. 957-973). IEEE.

Shen, B., Wang, Y., Shang, D., & Ning, T. (2026). Self-Commitment Latency: A Reward-Free Probe for Prompted Implicit Hacking. Publication date: 2026-06-04.

Hu, Y., Yang, Z., Zhao, C., & Ji, W. (2026, May). Adaptive guidance semantically enhanced via multimodal llm for edge-cloud object detection. In ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 12962-12966). IEEE.

Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. In Proceedings of the IEEE Symposium on Security and Privacy.

Yang, W., Qin, Y., & Yang, Y. (2021). Characterizing the multicast nature of malicious flows in CCN via SIS epidemic model. IEEE Systems Journal, 16(2), 2240-2250.

Li, S. (2025). Momentum, volume and investor sentiment study for us technology sector stocks—A hidden markov model based principal component analysis. PloS one, 20(9), e0331658.

Zhao, H., Huang, H., & Li, X. (2025, November). PP-Pose: Privacy-Preserving Human Pose Estimation Using Random High-Frequency Channel Combinations. In Australasian Joint Conference on Artificial Intelligence (pp. 191-203). Singapore: Springer Nature Singapore.

Li, Q., Li, L., Liu, Z., Sun, W., Li, W., Li, J., & Zhao, W. (2025). Cloud-edge collaboration for industrial Internet of Things: Scalable neurocomputing and rolling-horizon optimization. IEEE Internet of Things Journal, 12(12), 19929–19943.

Tian, Y., & Yang, Z. (2026, May). SAEC: scene-aware enhanced edge-cloud collaborative industrial vision inspection with multimodal LLM. In ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 19747-19751). IEEE.

Zhao, Y., & Han, Z. (2025, December). Safe and Adaptive Cloud Healing: Verifying LLM-Generated Recovery Plans with a Neural-Symbolic World Model. In 2025 7th International Academic Exchange Conference on Science and Technology Innovation (IAECST) (pp. 89-94). IEEE.

Wang, S., Zhang, X., Wang, P., & Li, X. (2026). Design of Magnetic Sensor Array-Based PUFs for IoT Security. IEEE Transactions on Instrumentation and Measurement, 75, 1-9.

Zhou, D. (2026). AI-Driven Hybrid SAST–DAST–SCA–IAST Framework for Risk-Based Vulnerability Prioritization in Microservice Architectures.

Qu, D., & Ma, Y. (2025, August). Edge–Mesh–Ledger: Federated AI and Blockchain Framework for Scalable Global Sustainability Solutions. In 2025 International Conference on Artificial Intelligence for Sustainable Innovation (AI-SI) (pp. 1-6). IEEE.

Yang, Z., Ji, W., Guo, Q., Zhao, J., Zhao, C., Zhang, X., ... & You, Y. (2026). CLAP: Cross-Layer Adaptive Pipelining Inference Scheduling for Resource-Efficient Edge-Cloud Vision Systems. ACM Transactions on Architecture and Code Optimization.

Hu, X., Wan, Z., & Zhou, Q. (2025). The Profitability Puzzle of On-Demand Carpooling: Pricing, Matching, and Route Compatibility. HKU Jockey Club Enterprise Sustainability Global Research Institute Paper, (2025/001).

Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. arXiv preprint arXiv:1606.06565.

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you?: Explaining the predictions of any classifier. In Proceedings of KDD.

Jia, Y., Wang, T., Xiao, K., & Guo, C. (2020). How to reduce opportunism through contractual governance in the cross-cultural supply chain context: evidence from Chinese exporters. Industrial Marketing Management, 91, 323-337.

Wang, C., Song, R., Muller, R., Monteuuis, J. P., Celik, Z. B., Petit, J., ... & Li, M. (2026, March). Cp-freezer: Latency attacks against vehicular cooperative perception. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 40, No. 2, pp. 1114-1122).

Yang, Z., Ji, W., & Wang, Z. (2024). Adaptive joint configuration optimization for collaborative inference in edge-cloud systems. Science China Information Sciences, 67(4), 149103.

Zhou, D. (2025, December). M-VP2: Microservice-Oriented Vulnerability Patch Planning-A Cost-Aware Approachusing Multi-Agent Reinforcement Learning. In 2025 5th International Conference on Computer, Internet of Things and Control Engineering (CITCE) (pp. 248-254). IEEE.

Yang, Z., Ji, W., Guo, Q., & Wang, Z. (2023, October). Javp: Joint-aware video processing with edge-cloud collaboration for dnn inference. In Proceedings of the 31st ACM International Conference on Multimedia (pp. 9152-9160).