

The Ethics of Spatial Computing: Integrating Digital Provenance into Human-Centric UX to Mitigate Synthetic Misinformation

Madhvesh Kumar *Independent Researcher*

Atlanta, USA

madhveshkumar@hotmail.com

Deepika Singh *Independent Researcher*

Atlanta, USA deepikapals@yahoo.co.in

Abstract—The proliferation of generative AI models capable of producing photorealistic synthetic media has introduced an acute epistemological challenge for spatial computing platforms. Augmented and extended reality (AR/XR) headsets overlay digital artifacts onto the physical world, creating an attack surface where users can no longer rely on perceptual intuition to distinguish authentic objects from AI-generated fabrications rendered in three-dimensional space. This paper proposes a control-systems framework for natively embedding cryptographic digital provenance—specifically, markers conforming to the Coalition for Content Provenance and Authenticity (C2PA) standard—within the rendering pipelines of spatial computing interfaces. We formalize the provenance verification loop as a multi-input, multi-output (MIMO) feedback control problem, where sensor inputs, content metadata, and user-facing trust indicators form an integrated signal chain subject to latency, error-rate, and privacy constraints. Closed-loop stability is proven analytically for all positive parameter values of the assumed first-order human-response model, and a numerical simulation illustrates the transient response of the adaptive trust-indicator controller under literature-calibrated conditions, including a simulated adversarial content burst and variable user responsiveness. An ethical analysis grounded in Value-Sensitive Design (VSD)

weighs informational transparency against biometric and spatial data privacy, and proposes a multi-level governance model spanning device, platform, and regulatory layers. A proof-of-concept implementation targeting the Unity rendering engine with a custom C2PA manifest parser provides evidence that the Ingest-Verify-Annotate-Present pipeline is computationally feasible within the 11.1 ms frame budget of a 90 Hz display. The contribution is a theoretically grounded framework with proof-of-concept validation rather than a deployment-ready system: the human-response model has not yet been identified with human participants, detection-accuracy and adversarial-robustness benchmarks remain to be conducted, and the proposed 3D provenance bindings await standardization. We argue that provenance-aware spatial interfaces are a promising, though by themselves insufficient, safeguard against synthetic misinformation, and that moving them toward deployment demands coordinated action across hardware manufacturers, standards bodies, content platforms, and regulatory authorities.

Index Terms—*Spatial computing, digital provenance, C2PA, deepfake detection, AR/XR, MIMO control systems, AI ethics, synthetic misinformation, human-computer interaction, value-sensitive design.*

I. INTRODUCTION

A. The Convergence of Spatial Computing and Generative AI

Spatial computing has moved from speculative prototype to consumer product within a compressed timeline. Apple's Vision Pro, Meta's Quest 3, and Microsoft's HoloLens 2 have collectively normalized the idea of persistent digital overlays anchored to physical environments [6]. These platforms merge camera-fed reconstructions of the real world with rendered digital content, producing a unified perceptual field in which the provenance of any given visual element is, by default, opaque to the user.

Simultaneously, generative adversarial networks (GANs), diffusion models, and neural radiance fields (NeRFs) have reached a point of output fidelity where synthetic three-dimensional objects and human avatars can be introduced into mixed-reality scenes with minimal perceptual discontinuity [9]. The threat model is no longer hypothetical. A spatially rendered deepfake—an AI-generated person, product, or environmental feature placed within a user's augmented

field of view—could influence purchasing decisions, distort situational awareness in safety-critical settings, or serve as a vector for targeted disinformation campaigns. Recent surveys of deepfake generation and detection methods have catalogued more than forty distinct synthesis architectures capable of producing outputs that defeat human perceptual judgment under controlled conditions [10].

The convergence of these two technological trajectories—immersive spatial interfaces and high-fidelity generative models—creates a compound risk that neither community has adequately addressed in isolation. AR/XR platform designers have focused on rendering fidelity, spatial registration accuracy, and interaction paradigms; generative AI researchers have focused on synthesis quality, controllability, and detection. The gap between these communities is the provenance layer: a systematic mechanism for certifying the origin, transformation history, and authenticity of every digital asset that enters the user's perceptual field.

B. The Provenance Gap in Immersive Interfaces

Traditional web and mobile platforms have begun to adopt content provenance standards. The C2PA specification, jointly developed by Adobe, Microsoft, the BBC, Intel, and Arm under the umbrella of the Content Authenticity Initiative, provides a cryptographic chain of custody for digital assets: each transformation applied to an image, video, or audio file is recorded in a tamper-evident manifest bound to the content via hashing and digital signatures [13]. Browser extensions and social-media platform integrations can surface these manifests to end users as trust indicators. The W3C Verifiable Credentials standard provides complementary infrastructure for asserting and validating the identity of content signers [14].

Spatial computing interfaces, however, lack an analogous mechanism. The rendering pipeline in AR/XR headsets ingests 3D models, textures, lighting maps, and spatial anchors from heterogeneous sources—developer SDKs, cloud-streamed assets, user-generated content, and real-time sensor data—without a unified provenance layer. No mainstream headset operating system currently exposes content-origin metadata to the wearer. The result is a provenance gap: the more seamlessly digital and physical realities are blended, the less capacity the user retains to interrogate the authenticity of what they perceive.

This gap is not merely an oversight; it reflects a structural asymmetry in incentives. Platform vendors benefit commercially from seamless blending, as the perceived quality of an AR experience improves when the boundary between real and rendered is imperceptible. Provenance indicators, by contrast, introduce visual friction. Resolving this tension requires a framework that treats provenance not as an intrusion upon experience quality, but as a first-class design parameter within the rendering control loop.

C. Research Objectives and Contributions

This paper addresses three interrelated questions:

- Technical feasibility. How can C2PA-compliant provenance manifests be integrated into the asset-ingestion and rendering pipelines of spatial computing platforms without introducing perceptible latency or degrading user experience?
- Control-theoretic modeling. Can the provenance verification and trust-indicator presentation process be formalized as a closed-loop MIMO control system, and if so, what are the stability and performance bounds imposed by real-time operating constraints?
- Ethical trade-offs. What privacy risks emerge when headset-mounted sensors, provenance verification subsystems, and user-facing trust indicators interact, and how should these risks be governed under established ethical frameworks?

The principal contributions are:

- A four-stage pipeline architecture (Ingest-Verify-Annotate-Present) for provenance-aware spatial rendering, with implementation-level detail sufficient for reproduction.
- A MIMO feedback control formalization of the adaptive trust-indicator system, with analytical stability conditions that hold for all positive parameter values of the assumed plant model, and numerical simulation results supported by time-domain, frequency-domain, and parameter-sensitivity figures.
- A proof-of-concept implementation in Unity demonstrating sub-frame-budget provenance verification for representative 3D assets.

- A qualitative design-property comparison positioning the framework against invisible watermarking, soft-binding, blockchain, and 2D C2PA approaches, together with a security argument for the manifest-stripping threat, and an ethical analysis grounded in the Value-Sensitive Design methodology of Friedman and Hendry [4], producing a multi-level governance model with prospective policy recommendations.

D. Scope of the Contribution

We state the contribution type explicitly. This paper proposes a theoretically grounded control-systems framework for provenance-aware spatial rendering, supported by proof-of-concept validation; it does not present a validated, deployment-ready solution to synthetic misinformation. The evidence offered is of three distinct kinds and should be read as such: (i) measured evidence—end-to-end pipeline latency on commodity XR hardware (Section III-B); (ii) analytical evidence—closed-loop stability, which we show holds for every positive value of the human-response parameters within the first-order model class (Section IV-D); and (iii) simulated and literature-grounded evidence—transient performance under assumed user and attacker models (Section IV-E) and a qualitative comparison with, and security argument against, existing provenance methods (Section IV-F).

Three things are not established here. The human-response plant has not been identified with human participants; detection accuracy, false-positive rates, computation overhead relative to competing methods, and robustness to real manifest-stripping attacks have not been benchmarked; and the 3D C2PA bindings are forward-compatible proposals rather than ratified standards. We regard these as prerequisites for deployment rather than optional refinements. Accordingly, the governance recommendations of Section V are prospective and conditional, and Section VI-C specifies the empirical program required before deployment claims can be made.

II. RELATED WORK

A. Content Provenance and Authenticity Standards

The C2PA specification (version 1.3, released 2023, and subsequently advanced through the 2.x series) defines a data model for content credentials: signed, tamper-evident manifests that record the identity of the creator, the tool used, the actions performed, and the cryptographic binding between the manifest and the content payload [13]. The specification builds on prior work in the JPEG Trust initiative and the International Press Telecommunications Council (IPTC) photo metadata standard. Several implementations exist for two-dimensional media, including Adobe's Content Credentials toolchain and the open-source c2patool library. However, no published work has addressed the extension of C2PA manifests to three-dimensional asset formats (glTF, USDZ, FBX), nor has any prior study examined the integration of provenance verification into a real-time 3D rendering pipeline.

B. Deepfake Detection in Two and Three Dimensions

The deepfake detection literature has grown rapidly since 2018, with methods spanning frequency-domain analysis, physiological signal detection, and learned classifiers [10]. More recent work has extended detection to 3D-reconstructed content, particularly NeRF-generated scenes [9]. Zhao et al. [15] demonstrated that invisible

image watermarks—one class of provenance signal—can be provably removed using generative AI, highlighting the fragility of watermark-only strategies and motivating the cryptographic-manifest approach adopted in this paper.

A distinct line of research has examined user perceptions of deepfakes and the effectiveness of labeling interventions. Groh et al. [5] conducted a large-scale study ($n = 15,016$) comparing human and algorithmic deepfake detection accuracy, finding that brief contextual warnings improved human detection performance by 6–10 percentage points. This finding supports the hypothesis that machine-mediated provenance indicators can meaningfully augment human judgment in immersive contexts.

C. Trust and Perception in AR/XR Environments

Research on trust in augmented reality has established that users exhibit a strong tendency to accept spatially registered digital content as veridical, a phenomenon Dey et al. [1] term “AR authority bias.” In their systematic review of 291 AR user studies, they found that trust in AR-presented information exceeded trust in equivalent screen-based information by a statistically significant margin across multiple task domains. This bias amplifies the risk posed by unprovenanced synthetic content and underscores the need for countervailing trust indicators.

Grubert et al. [6] proposed a taxonomy of context-awareness in AR systems, identifying content provenance as an underexplored dimension of contextual metadata. Their framework stops short of proposing specific provenance mechanisms but establishes the conceptual vocabulary adopted in this paper.

D. Control-Theoretic Approaches to Human–Computer Interaction

The application of control theory to HCI problems has a lineage extending from manual control theory in aviation [8] to contemporary work on adaptive user interfaces. Müller et al. [11] formalized gaze-contingent rendering as a feedback control problem, demonstrating that foveated rendering quality can be modulated in closed loop with saccade-driven feedback while maintaining perceptual constancy. Their work establishes the precedent that human visual perception can be treated as a plant within a control loop, provided the modeling assumptions are empirically grounded and the ethical dimensions are addressed. Our work extends this paradigm from rendering quality to content authenticity.

E. Value-Sensitive Design and AI Ethics

Value-Sensitive Design (VSD) is a theoretically grounded methodology for incorporating human values into technology design [4]. VSD proceeds through three interlinked investigations—conceptual, empirical, and technical—to identify stakeholders, enumerate relevant values (e.g., autonomy, privacy, trust, fairness), and translate those values into design requirements. Floridi’s [3] ethical framework for information provides complementary foundations, particularly his account of informational self-determination as a prerequisite for autonomous agency.

In the AI ethics domain, the EU’s AI Act (Regulation 2024/1689) has introduced mandatory transparency requirements for AI-generated content, including labeling obligations for deepfakes. However, the Act’s provisions are media-format-agnostic and do not address the specific challenges posed by spatially registered synthetic content in AR/XR environments [2]. This regulatory gap motivates the governance recommendations in Section V of this paper.

F. Hybrid AI Models and Integrated System-Level Modeling

Our framework sits within a broader movement toward hybrid AI and integrated system-level modeling, in which learned components are coupled to explicit analytical or control-theoretic structure rather than deployed as monolithic black boxes. Several recent contributions illustrate this trajectory across disparate domains. Lal et al. [16] combine deep compressive sensing with hardware-efficient models for efficient ECG compression and reconstruction, exemplifying the co-design of learned representations with resource-constrained execution targets—precisely the regime in which our on-device verification pipeline must operate. Fazzinga et al. [17] revisit the notions of extension and acceptance over incomplete abstract argumentation frameworks, offering formal machinery for reasoning under partial or contested information that parallels the discrete trust-state semantics introduced in Section III. Das et al. [18] survey computer-vision approaches to wildfire detection and emission quantification, a domain that, like ours, demands real-time inference fused with downstream decision logic under safety-critical constraints. Finally, Das et al. [19] present UrSL-CNN for Urdu sign-language translation, underscoring how human-centric AI systems must integrate perceptual modeling with accessibility and trust considerations. Collectively, these works situate our contribution within an emerging design philosophy that treats AI components as elements of a larger, analyzable system—the same philosophy that motivates our control-theoretic treatment of trust-indicator dynamics.

III. PROPOSED FRAMEWORK: PROVENANCE-AWARE SPATIAL RENDERING

A. Architecture Overview

We propose a four-stage pipeline—Ingest, Verify, Annotate, Present—that operates within the rendering loop of a spatial computing operating system. The pipeline is designed as a modular middleware layer that interposes between the asset-loading subsystem and the scene-graph renderer, requiring no modifications to the underlying graphics API (Vulkan, Metal, or OpenGL ES). Fig. 1 illustrates the data flow and the feedback path that couples the user-facing presentation stage back to the MIMO controller.




Fig. 1. Data-flow architecture of the Ingest-Verify-Annotate-Present (IVAP) pipeline. Assets enter at the Ingest stage; the trust signal σ propagates through Verify and Annotate to the Present stage, while gaze and query telemetry close the loop through the MIMO PI controller that modulates indicator saliency (u_1) and verification depth (u_2).

Stage 1—Ingest. Every digital asset entering the rendering pipeline—whether loaded from local storage, streamed from a

Signal	Visual Indicator	Meaning
V	Solid green ring with check glyph; high salience	Verified — all three checks (integrity, signer, temporal) pass
P	Dashed amber ring with caution glyph; medium salience	Partially verified — signer authenticity confirmed; integrity or temporal check inconclusive
U	Hollow grey outline; low, non-intrusive salience	Unverified — no manifest present (unattested asset)
F	Solid red ring with strike glyph; maximum salience	Failed — one or more checks explicitly failed (e.g., tamper detected)

remote server, or generated on-device by an AI model—is intercepted by a provenance extraction module. This module inspects the asset container for an embedded C2PA manifest. For glTF assets, the manifest is stored as a custom extension (KHR_c2pa_provenance); for USDZ assets, as an auxiliary payload referenced in the asset's root layer; for proprietary formats, as an appended JUMBF (JPEG Universal Metadata Box Format) container, consistent with the C2PA specification's extensibility provisions. Assets lacking a manifest are flagged as unattested and assigned a null provenance record.

The Ingest module also computes a content hash (SHA-256) of the asset's geometric and texture data, independent of any manifest-provided hash. This dual-hash strategy enables detection of post-signing modifications: if the Ingest-computed hash diverges from the manifest-declared hash, the asset is flagged as tampered.

Stage 2—Verify. Extracted manifests are validated against a certificate authority (CA) trust chain. The verification subsystem checks three properties:

- **Manifest integrity:** the cryptographic hash binding the manifest to the asset content has not been broken, using the dual-hash comparison from Stage 1.
- **Signer authenticity:** the manifest was issued by a recognized and unrevoked certificate. The subsystem maintains a local trust-anchor store populated with root certificates from established CAs (e.g., DigiCert, Let's Encrypt) and domain-specific CAs. Certificate revocation is checked against a locally cached Certificate Revocation List (CRL), updated every 24 hours.
- **Temporal validity:** the manifest's timestamps fall within acceptable bounds, preventing replay attacks.

Verification results are encoded as a discrete trust signal $\sigma \in \{V, P, U, F\}$: *Verified* (all three checks pass), *Partially verified* (signer authenticity confirmed but integrity or temporal checks inconclusive), *Unverified* (no manifest present), or *Failed* (one or more checks explicitly failed).

Stage 3—Annotate. The trust signal σ is attached to the asset's scene-graph node as persistent metadata. This annotation survives frame updates, spatial re-anchoring events, and level-of-detail transitions. For dynamically updated assets (e.g., streamed volumetric video), re-verification is triggered at a configurable interval (default: every 30 frames, or 333 ms at 90 Hz).

Stage 4—Present. A UX layer translates trust signals into human-perceptible indicators. Drawing on research in ambient information display [12] and the empirical finding that contextual warnings improve deepfake detection by 6–10 percentage points [5], we

propose a graduated indicator set, summarized in Table I and illustrated in Fig. 2.

TABLE I TRUST SIGNAL INDICATOR DESIGN



Fig. 2. Graduated trust-indicator set anchored to a spatial asset. Each badge couples a distinct color, glyph, and ring style to one of the four trust states; salience and modality are modulated at run time by the controller of Section IV.

The intensity and modality of these indicators are governed by two parameters—salience and verification depth—dynamically modulated by the feedback controller described in Section IV. Users may also manually override the indicator intensity, consistent with the autonomy-preserving principle of VSD [4].

B. Proof-of-Concept Implementation

To assess the computational feasibility of the IVAP pipeline, we developed a proof-of-concept implementation as a Unity 2023 LTS plugin targeting the Meta Quest 3 via the OpenXR runtime. The implementation comprises four C# modules corresponding to the four pipeline stages, totaling approximately 2,400 lines of code.

Ingest module. A custom AssetPostprocessor intercepts assets during import and at runtime via the Addressables loading API. The module parses C2PA manifests embedded in glTF files using a lightweight JUMBF parser (548 lines). SHA-256 hashing of mesh vertex buffers and texture pixel data is performed via Unity's NativeArray API to minimize managed-memory allocations.

Verify module. Certificate chain validation is implemented using the .NET X509Chain class with a custom trust store. CRL caching uses a write-ahead log stored in the application's persistent data directory, updated via background coroutine every 24 hours. The full verification path executes in 1.8 ms mean, 3.1 ms 99th percentile on the Quest 3's Snapdragon XR2 Gen 2 processor.

Annotate module. Trust metadata is stored as a custom component (ProvenanceAnnotation) attached to the asset's GameObject, exposing the trust signal, manifest summary, and verification timestamp.

Present module. Trust indicators are rendered as post-processing shader effects applied to the asset's bounding geometry, with indicator parameters exposed as uniform variables driven by the controller.

Latency benchmarks. End-to-end pipeline latency was measured across a corpus of 200 glTF assets ranging from 500 to 500,000 triangles (Table II). Fig. 3 plots the mean and 99th-percentile latency for each complexity tier against the 11.1 ms frame budget.

TABLE II IVAP PIPELINE LATENCY BENCHMARKS (QUEST 3)

Complexity	Mean	P99	Budget (%)
Low (< 5K tris)	2.1 ms	3.4 ms	19%
Medium (5K–50K)	3.8 ms	5.9 ms	34%
High (50K–500K)	6.2 ms	8.7 ms	56%




Fig. 3. End-to-end IVAP latency by asset complexity, mean and 99th percentile, relative to the 11.1 ms frame budget of a 90 Hz display. Even the high-complexity P99 case retains roughly 2.4 ms of headroom.

All measurements fall within the 11.1 ms frame budget, indicating that provenance verification is computationally feasible on this platform without exceeding the per-frame budget. These benchmarks establish latency feasibility only; they do not measure detection accuracy, false-positive rates, or robustness to adversarial manipulation, which are scoped as future work in Section VI-C.

IV. MIMO CONTROL-THEORETIC FORMALIZATION

A. System Definition

The adaptive trust-indicator system is formalized as a two-input, two-output (2×2) MIMO feedback control system. Two distinct actuator variables (indicator salience and verification depth) are jointly modulated in response to two distinct measured outputs (user trust-state and system verification load). Fig. 4 shows the decentralized two-loop structure that this section develops.




Fig. 4. Decentralized 2×2 MIMO feedback structure. Two PI controllers $C_1(s)$ and $C_2(s)$ regulate the trust-state (y_1) and verification-load (y_2) outputs; the triangular plant couples verification depth u_2 into the trust channel through the cross term G_{12} , while u_1 does not load the compute channel.

State variables. Let $y_1(t) \in [0, 1]$ denote the estimated user trust-state, computed as a weighted exponential moving average:

$$y_1(t) = \alpha \cdot g(t) + \beta \cdot q(t) + (1 - \alpha - \beta) \cdot (1 - d(t))$$

where $g(t)$ is normalized gaze dwell-time on provenance indicators, $q(t)$ is the query rate (queries per minute, normalized), $d(t)$ is the indicator dismissal rate, and $\alpha = 0.5$, $\beta = 0.3$. Intuitively, (1) says that estimated trust rises when the user attends to and queries provenance cues and falls when the user habitually dismisses them; the weights α and β set the relative influence of attention versus active querying.

Let $y_2(t) \in [0, 1]$ denote the system verification load: the fraction of the per-frame computational budget consumed by the Verify stage. Reference signals. $r_1 = 0.7$ (default trust-state target) and $r_2 = 0.4$ (default verification-load target). Error signals. $e_1(t) = r_1 - y_1(t)$ and $e_2(t) = r_2 - y_2(t)$. Control inputs. $u_1(t) \in [0, 1]$ is indicator salience, and $u_2(t) \in [0, 1]$ is verification depth, where 0 denotes heuristic-only checking and 1 denotes full cryptographic verification.

B. Plant Model

The plant captures the coupled dynamics of user perception and system computation. Because indicator salience drives user trust through a first-order cognitive lag while verification depth maps almost instantaneously to compute load, the plant is conveniently written entry-by-entry rather than as a dense matrix:

$$G_{11}(s) = K_1 / (\tau_1 s + 1), \quad G_{12} = K_{12}, \quad G_{21} = 0, \quad G_{22} = K_2$$

where $K_1 = 0.8$ (scaled from [5]), $\tau_1 = 2.0$ s (consistent with [1]), $K_{12} = 0.15$ (weak cross-coupling), and $K_2 = 0.9$ (measured from the proof-of-concept). The triangular structure ($G_{21} = 0$) encodes a physical decoupling: visual indicators do not consume significant compute, so salience never loads the verification channel, whereas verification depth affects both compute and—weakly, through G_{12} —user confidence. This triangularity is what permits the decentralized design of Section IV-C.

C. Controller Design

We employ a decentralized PI controller with anti-windup, exploiting the triangular plant structure to design two independent single-input, single-output (SISO) loops. Each controller has the standard proportional-plus-integral form, in which the proportional term reacts to the instantaneous error and the integral term removes steady-state offset:

$$C_1(s) = K_{p1} + K_{i1} / s$$

with $K_{p1} = 1.2$, $K_{i1} = 0.3$, selected for a gain-crossover (open-loop unity-gain) frequency of ≈ 0.33 rad/s, which trades response speed for the large stability margins quantified below.

$$C_2(s) = K_{p2} + K_{i2} / s$$

with $K_{p2} = 0.8$, $K_{i2} = 0.5$. Both controllers implement clamping anti-windup, bounding the integrator output to $[0, 1]$ so that actuator saturation does not cause integral wind-up during sustained load bursts.

D. Stability Analysis

Loop 2. With a PI controller acting on the pure-gain plant $G_{22} = K_2$, the closed-loop characteristic equation is $(1 + K_{p2}K_2)s + K_{i2}K_2 = 0$, i.e. $1.72s + 0.45 = 0$, giving a single real pole at $s = -0.262$ and a time constant of 3.8 s. Because the characteristic equation is first order with positive coefficients for any $K_2 > 0$, the loop is stable for every positive plant gain.

Loop 1. The open-loop transfer function is the product of the PI controller and the first-order plant:

$$L_1(s) = C_1(s)G_{11}(s) = (K_{p1} + K_{i1}/s) \cdot K_1 / (\tau_1 s + 1)$$

Substituting the parameter values and clearing denominators yields the closed-loop characteristic polynomial:

$$2s^2 + 1.96s + 0.24 = 0$$

whose roots are $s_1 = -0.1434$ and $s_2 = -0.8366$. Both poles are real and negative, confirming asymptotic stability with no oscillatory transient. The dominant pole at $s_1 = -0.1434$ corresponds to a time

constant of $1/0.1434 \approx 7.0$ s; the conservative 4τ rule would predict a 2% settling time near 28 s, whereas the full nonlinear simulation (Section IV-E) settles in ≈ 23 s because the dominant mode carries only part of the response residue.

Structural stability under plant uncertainty. Because the human-response parameters are literature-grounded estimates rather than identified values, it matters that stability does not hinge on them. For general parameters, the Loop 1 characteristic polynomial is $\tau_1 s^2 + (1 + K_{p1}K_1)s + K_{i1}K_1 = 0$. For any $\tau_1 > 0$ and $K_1 > 0$, and with the design gains K_{p1} , $K_{i1} > 0$, all three coefficients are strictly positive, which for a second-order polynomial is necessary and sufficient for both roots to lie in the open left-half plane. Since the plant is triangular and the controller diagonal, the closed-loop poles of the full 2×2 system are exactly those of the two loops, and Loop 2 is stable for any $K_2 > 0$. Stability is therefore guaranteed across the entire first-order model class, independent of the specific values of K_1 and τ_1 ; the unidentified parameters affect performance—settling time and the onset of actuator saturation examined in Section IV-E—but not stability. This guarantee is conditional on the model structure itself: if empirical identification reveals higher-order dynamics or appreciable perceptual delay, the analysis must be repeated for the identified plant.

A frequency-domain evaluation of $L_1(s)$ gives a gain-crossover frequency $\omega_c \approx 0.33$ rad/s and a phase margin of $\approx 109^\circ$. Because $L_1(s)$ is a minimum-phase type-1 system whose phase never reaches -180° , the gain margin is effectively unbounded (no phase-crossover frequency exists). These margins—far above the conventional 6 dB / 45° thresholds—indicate ample robustness, deliberately purchased at the cost of response speed. These margins are depicted in Fig. 5.




Fig. 5. Left: closed-loop pole map. All poles of both loops lie in the open left-half plane, so the system is asymptotically stable. Right: open-loop magnitude response of Loop 1, annotated with the ≈ 0.33 rad/s gain-crossover frequency and $\approx 109^\circ$ phase margin; the absence of a -180° phase crossing yields an unbounded gain margin.

Cross-coupling effect. The G_{12} term introduces a disturbance pathway from u_2 to y_1 . Because Loop 2 settles in ≈ 3.8 s relative to Loop 1's ≈ 7.0 s, this disturbance acts as a slowly varying bias that Loop 1's integral action eliminates in steady state, as the simulation in Section IV-E confirms.

E. Numerical Simulation

We simulated the closed-loop system as a discrete-time model of the two coupled loops (time step 0.01 s) with clamping anti-windup, under three scenarios summarized in Table III. The same simulator produces Figs. 6 and 7, so the plotted curves and the tabulated numbers are mutually consistent.

TABLE III SUMMARY OF CLOSED-LOOP SIMULATION RESULTS

Scenario	Settling / recovery	Overshoot	dev	Stable
A: Nominal step	≈ 23 s	0.0%	0.0%	Yes
B: Adversarial burst	≈ 6 s recov.	0.0%	0.01	Yes
C: Variable K_1	steady track	—	0.043	Yes

Column 2 reports the 2% settling time for Scenario A, the disturbance-recovery time for Scenario B, and steady tracking for Scenario C; |dev| is steady-state error (A), peak trust-state dip (B), or maximum deviation from r_1 (C).

Scenario A—Nominal operation. A step change in r_1 from 0.5 to 0.7 at $t = 5$ s, with r_2 held at 0.4. As Fig. 6 shows, $y_1(t)$ reaches the reference within ≈ 23 s with no overshoot and zero steady-state error, while $y_2(t)$ remains within 2% of its reference throughout—confirming that a step in the trust channel does not perturb the verification channel, as the triangular plant predicts.




Fig. 6. Scenario A—nominal step response. Top: outputs y_1 (trust-state) and y_2 (verification load) track their references $r_1 = 0.7$ and $r_2 = 0.4$ with no overshoot; the shaded band marks the $\pm 2\%$ settling region. Bottom: the corresponding control inputs u_1 (salience) and u_2 (verification depth).

Scenario B—Adversarial content burst. At $t = 10$ s, a burst of 50 unverified assets is injected, spiking the verification load. Loop 2 responds by reducing $u_2(t)$ from 0.44 to 0.23 to protect the frame budget. The cross-coupling through G_{12} induces only a small transient dip in $y_1(t)$ of ≈ 0.01 , which Loop 1 erases within ≈ 6 s (Fig. 7, left). The small magnitude of the dip is itself evidence of effective decoupling: the verification channel can shed load under attack with negligible disturbance to the user-facing trust estimate. Scenario B models the computational effect of an attack—a surge in verification demand under an assumed attacker model—and therefore tests the controller's load-shedding behavior, not the cryptographic resistance of the manifests; the latter is analyzed in Section IV-F and remains to be evaluated empirically.

Scenario C—Variable user responsiveness. Holding r_1 fixed, we vary the cognitive gain K_1 sinusoidally within $\pm 19\%$ of nominal (0.65–0.95,

period 60 s). The controller holds $y_1(t)$ to a maximum deviation of 0.043—comfortably inside the ± 0.08 design band (Fig. 7, center)—indicating robustness to moderate parameter variation within the assumed model. A complementary static sensitivity sweep (Fig. 7, right) reveals an actuator-saturation knee near $K_1 \approx 0.62$: below this value the salience actuator saturates ($u_1 = 1$) and steady-state trust tracking begins to degrade. This boundary quantifies the validity envelope of the fixed-gain design and directly motivates the adaptive and gain-scheduled extensions discussed in Section VI.




Fig. 7. Robustness analysis. Left: Scenario B adversarial load burst— y_1 dips by only ≈ 0.01 as u_2 sheds load, recovering in ≈ 6 s. Center: Scenario C sinusoidal variation of K_1 within $\pm 19\%$, with y_1 held inside the ± 0.08 band. Right: static sensitivity sweep over K_1 , exposing the salience-saturation knee near $K_1 \approx 0.62$.

F. Comparative Positioning Against Existing Provenance Methods

A full empirical comparison against deployed watermarking and provenance systems would require a shared 3D benchmark corpus and human-subject evaluation, which we scope as future work in Section VI. To situate the contribution in the interim, Table IV offers a qualitative comparison across the properties that matter most for spatially rendered content: resistance to tampering, native support for 3D asset formats, compatibility with a real-time rendering pipeline, resistance to forgery of a valid provenance claim (including by AI-based removal and re-embedding), and privacy posture. The assessment is grounded in the cited literature—most directly the provable-removability result of Zhao et al. [15] for invisible watermarks—rather than in new measurements. The ratings are analytical assessments of the properties each approach is designed to provide; they are not measured performance and should not be read as demonstrating superior detection accuracy, lower computation overhead, or lower false-positive rates, none of which are benchmarked here.

TABLE IV QUALITATIVE DESIGN-PROPERTY COMPARISON OF PROVENANCE APPROACHES FOR SPATIAL CONTENT

Method	Tamper -evidence	Nativ e 3D	Real -time	Forger y resist.	Privac y
Invisible watermarking	Low	No	Med	Low	Med.
Soft binding (perceptual)	Med.	Partial	Med	Med.	Med.
Blockchain ledger	High	No	Low	High	Low
C2PA (2D media)	High	No	Med	High	Med.
IVAP (this work)	High	Yes	High	High	High

Two points follow. First, the cryptographic-manifest approaches (blockchain, 2D C2PA, and IVAP) are, by construction, stronger than the perceptual schemes on tamper-evidence and forgery resistance, because they bind a signature to a content hash rather than to a fragile perceptual signal. Second, IVAP differs from prior cryptographic schemes in design scope—native 3D-format handling and operation inside a hard real-time rendering loop, the latter supported by the measurements of Table II—while its device-local verification design (Section V) is intended to provide the strongest privacy posture. Of the IVAP ratings, only real-time compatibility is backed by measurement in this paper.

Security argument against manifest stripping. Zhao et al. [15] show that AI-based regeneration removes invisible watermarks; the analogous attack on a manifest scheme is to strip, alter, or re-sign the manifest. We have not evaluated such attacks empirically, but the security property on which IVAP relies can be stated precisely. Under the dual-hash and trust-list design of Section III-A, each attack maps to a determinate trust state: (i) stripping the manifest, or regenerating the asset so that it no longer carries one, yields $\sigma = U$; (ii) modifying geometry or textures while retaining the manifest breaks the hash binding and yields $\sigma = F$; (iii) re-signing with a key that does not chain to the trust list fails the signer-authenticity check and yields $\sigma = F$; and (iv) replaying an outdated manifest is caught by the temporal check. An adversary can therefore downgrade an asset to Unverified but cannot, without compromising a trusted signing key, cause a synthetic asset to be presented as Verified.

This is a guarantee of unforgeability, not of persistence: cryptographic manifests, like watermarks, can be removed. Protection against stripping therefore depends on how U-state content is treated by default and on the governance measures of Section V that disincentivize stripping by intermediaries, rather than on an algorithm-level defense; key compromise remains a residual risk bounded only by revocation latency. Quantifying these properties—detection accuracy, false-positive rate, and computation overhead against a benchmark of stripped, modified, and re-signed 3D assets, alongside watermarking and soft-binding baselines—is the adversarial study specified in Section VI-C.

V. ETHICAL ANALYSIS: PRIVACY, TRANSPARENCY, AND AUTONOMY

A. Methodological Framework: Value-Sensitive Design

We structure the ethical analysis using the tripartite VSD methodology of Friedman and Hendry [4]: conceptual investigation (identifying stakeholders and values), empirical investigation (examining stakeholder experience), and technical investigation (assessing how design supports or undermines values). While a full empirical investigation is beyond the scope of this paper, we conduct the conceptual and technical investigations in depth.

Stakeholders. Direct stakeholders include AR/XR headset users, content creators (both human and AI-assisted), and platform operators. Indirect stakeholders include bystanders whose likeness or environment may be captured by headset sensors, regulatory bodies, and civil society organizations.

Values. The central values in tension are informational transparency (the right to accurate provenance information), privacy (control over personal data generated by the verification system), autonomy (the right to independent authenticity judgments), and fairness (equitable access to attestation tools).

B. The Transparency Imperative

The case for provenance transparency rests on Floridi's [3] account of informational self-determination: users have a right to know whether perceived objects are physical-world artifacts, verified augmentations, or synthetic fabrications. The "AR authority bias" documented by Dey et al. [1]—users afford greater credibility to AR-presented information than to screen-based equivalents—compounds the risk. Provenance markers reduce this cognitive burden by externalizing the verification function, providing the informational basis for autonomous judgment without automating the judgment itself.

C. Privacy Risks of Provenance Infrastructure

Three categories of privacy risk emerge:

Spatial telemetry. The MIMO feedback controller requires gaze direction, dwell-time, and interaction events. Kröger et al. [7] demonstrated that eye-tracking data from VR headsets can infer cognitive load, emotional state, personality traits, and clinical conditions. In this context, gaze telemetry reveals how users evaluate what they see—a metacognitive signal of exceptional sensitivity.

Verification metadata. Each provenance check potentially involves querying external CAs and content-origin databases. Without safeguards, these queries leak information about encountered content.

Indicator interaction logs. We term this category epistemic telemetry: a record of how users respond to uncertainty, with potential to reveal political orientation, susceptibility to misinformation, and technological literacy.

D. Toward a Value-Sensitive Governance Model

Resolving the tension between transparency and privacy requires governance at three levels. These recommendations are prospective: they describe the governance that a deployment would require and are conditional on (i) ratification of 3D provenance bindings (Section VI-B) and (ii) empirical validation of the indicator design and human-response model (Section VI-C). They are offered to inform standards and policy development, not as guidance for integrating the present prototype into production platforms.

Device level. On-device verification (demonstrated in our proof-of-concept) eliminates cloud dependencies. Controller state variables should be stored in volatile memory and discarded at session termination.

Platform level. Data minimization by default; platform-level APIs designed as zero-knowledge protocols where feasible.

Regulatory level. The EU AI Act [2] mandates content labeling but lacks spatial-computing-specific provisions. We recommend domain-specific annexes addressing (1) mandatory provenance labeling for AI-generated 3D assets; (2) prohibitions on provenance-stripping by intermediaries; (3) audit requirements for adaptive trust-indicator systems; and (4) equitable access provisions for content-signing tools.

E. Epistemic Justice and Access Equity

A provenance system risks reproducing existing inequalities if content creators must obtain digital certificates and integrate signing tools that require resources unavailable to independent artists, citizen journalists, and communities in the Global South. If unattested content is systematically treated as less trustworthy, the framework could inadvertently marginalize under-resourced voices. Mitigation requires free, open-source signing tools; simplified workflows in

consumer-grade creation apps (e.g., Blender, Unity Personal); and recognition of alternative attestation mechanisms such as community vouching and decentralized identity networks.

VI. DISCUSSION

A. Generalization Beyond AR/XR

Although we have developed the framework for AR/XR, neither the IVAP pipeline nor its control-theoretic core is intrinsically tied to head-mounted displays. The essential pattern—ingest heterogeneous content, verify cryptographic provenance, annotate with a discrete trust state, and present a salience-modulated indicator under feedback—applies to any domain in which a stream of artifacts of uncertain origin must be surfaced to a human decision-maker in real time, and in which over- or under-trust both carry cost.

Three domains illustrate the transfer. In automated journalism and social-media feeds, an analogous controller could modulate how prominently a content-credential badge is shown as a function of reader engagement and verification cost, calibrating trust without overwhelming the interface. In clinical decision-support, where AI-generated or AI-enhanced imagery is increasingly interleaved with primary acquisitions, the same trust-state machinery could govern how aggressively a system flags synthetically reconstructed regions to a radiologist. In autonomous-vehicle and robotics perception, provenance of map tiles, sensor overlays, and V2X messages must be verified within a hard real-time budget, exactly the regime our latency analysis targets. In each case the two-loop decomposition—one loop regulating a human-facing trust signal, the other regulating the verification effort against a compute budget—carries over, with only the plant identification (the human or operator response model) requiring domain-specific calibration.

B. Standardization Status and Multi-User Scalability

Our adaptation of C2PA to 3D assets relies on custom extensions, and their standardization status must be stated plainly. The C2PA specification has matured rapidly: the 2.x series consolidates the trust model around a Conformance Program and an official Trust List, and the work is being advanced toward international standardization under ISO/IEC. Regulatory pull is concrete—Article 50 of the EU AI Act requires machine-readable marking of AI-generated content, with obligations phased in during 2026—so a robust provenance substrate is becoming a compliance requirement rather than an option. C2PA's manifests support both hard binding (a cryptographic hash of the asset, strongly tamper-evident but broken by re-encoding) and soft binding (a durable perceptual identifier or watermark that survives transformations but offers weaker integrity guarantees); a production 3D deployment would combine the two.

For 3D specifically, no ratified provenance extension yet exists. glTF 2.0 (Khronos) is the dominant real-time interchange format, and OpenUSD/USDZ (Pixar–Apple) the dominant authoring and archival format, but neither currently defines a standardized content-credential channel. Our *KHR_c2pa_provenance* extension and USDZ auxiliary-payload binding should therefore be read as a forward-compatible proposal, not a deployed standard. Encouragingly, cross-organization coordination—through the Metaverse Standards Forum and an active Khronos–Alliance for OpenUSD liaison—is precisely the venue in which such an extension would be ratified, and we frame our binding to align with that process. Until such an extension is ratified, the framework cannot be integrated interoperably into production

AR/XR platforms, and the deployment-oriented discussion in this paper should be read as prospective.

On scalability, the proof-of-concept verifies assets for a single wearer, whereas shared and large-scale XR sessions raise two further problems. First, multi-user consistency: when many headsets render the same scene, provenance verdicts and indicator states should be consistent across participants, which argues for verifying once at a trusted edge or session host and distributing signed verdicts rather than re-verifying per device. Second, throughput: a populated scene may stream thousands of assets, exceeding per-frame verification capacity; here the verification-load loop of Section IV becomes the natural admission-control mechanism, prioritizing salient or newly introduced assets and amortizing re-verification across frames. A networked-control extension that treats inter-headset communication delay and verdict caching as part of the plant is the subject of ongoing work.

C. Limitations and Future Work

Model calibration. The MIMO plant assumes a first-order lag for user cognitive response, calibrated from aggregate statistics rather than from a controlled study. Individual variation in processing speed, AR expertise, and task context will produce plant-model mismatch, and the parameters K_1 , τ_1 , α , and β remain literature-grounded estimates. The sensitivity analysis of Section IV-E quantifies the resulting validity envelope (stable tracking for $\pm 19\%$ variation in K_1 , with a saturation knee near $K_1 \approx 0.62$), but it is not a substitute for empirical identification. The analytical result of Section IV-D guarantees stability for all positive K_1 and τ_1 within the first-order model class, but not the validity of that model class itself; the human-in-the-loop model, which is the conceptual core of the framework, therefore remains unvalidated against real human dynamics.

Empirical evaluation. This paper provides latency benchmarks, a stability and robustness analysis, and a literature-grounded qualitative comparison, but it deliberately does not report a user study, a deepfake-detection accuracy benchmark, or an adversarial-robustness evaluation against a generative attacker, because we did not collect such data and decline to present simulated figures as if they were measured. We instead scope a concrete program of future work: (i) a within-subjects study measuring trust calibration and task performance with and without indicators, which would directly identify the plant and the weights α , β ; (ii) a quantitative benchmark on a labeled corpus of provenanced, stripped, and synthetic 3D assets, reporting detection accuracy, false-positive rate, and computation overhead for IVAP alongside watermarking and soft-binding baselines, converting the qualitative ratings of Table IV into measured comparisons; and (iii) an adversarial study in which an attacker attempts manifest stripping and re-signing, evaluating the dual-hash and trust-list defenses end to end. We regard these three studies as prerequisites for any deployment claim rather than as optional extensions, and the conclusions of this paper are limited accordingly.

Engineering scope. The proof-of-concept targets a single engine (Unity) and headset (Quest 3); generalization requires API-specific adaptation, though the IVAP architecture is engine-agnostic. As noted in Section VI-B, the 3D C2PA extensions are not yet ratified. Simulation code and the proof-of-concept are available as supplementary material.

D. Implications for Control Systems Research

This work situates trust calibration within a formal MIMO control framework, extending the precedent of Müller *et al.* [11] and opening several avenues:

- **Adaptive control:** replacing fixed-gain PI with model-reference adaptive control (MRAC) or gain scheduling to learn individual cognitive-response characteristics online—directly motivated by the saturation knee identified in Section IV-E.
- **Robust control:** formulating plant uncertainty as a structured perturbation set and applying μ -synthesis for formal robustness guarantees.
- **Networked control:** extending to multi-user AR scenarios with communication constraints and verdict-consistency requirements, as outlined in Section VI-B.

The ethical dimension of treating human cognition as a “controlled variable” deserves sustained attention. The feedback loop modulates the informational environment in which judgment operates; the governance mechanisms in Section V are intended to maintain the boundary between informing and manipulating.

VII. CONCLUSION

Spatial computing is rewriting the contract between perception and reality. As generative AI makes synthetic three-dimensional content indistinguishable from physical-world inputs, the burden of authenticity verification can no longer rest on unaided human cognition. This paper has argued that cryptographic digital provenance—C2PA-conformant manifests integrated into the rendering pipeline—is a promising, technically feasible, and ethically navigable mechanism for restoring informational transparency in AR/XR environments.

We have presented the IVAP pipeline architecture, established its computational feasibility through a proof-of-concept achieving sub-frame-budget latency on current-generation hardware, and formalized the adaptive trust-indicator system as a MIMO feedback control problem whose stability is guaranteed across the assumed first-order model class and whose performance is characterized in simulation through time-domain, frequency-domain, and sensitivity analyses. The ethical analysis, grounded in Value-Sensitive Design, has identified privacy risks and proposed a multi-level governance model with prospective recommendations. The contribution is thus a theoretically grounded framework with proof-of-concept validation; user studies, quantitative detection and adversarial-robustness benchmarks, and standardization of the 3D bindings remain prerequisites for deployment.

The path forward demands coordinated action: dedicated silicon for on-device verification; ratification of C2PA extensions for 3D content through the Khronos and OpenUSD ecosystems; spatial-computing-specific regulatory provisions; freely accessible content-signing tools; empirical user studies to calibrate plant models and validate indicator designs; and adversarial and detection-accuracy benchmarks against manifest-stripping attacks. The alternative—an immersive information environment in which synthetic and authentic content are permanently indistinguishable—is not a future that responsible technologists, policymakers, or citizens can accept.

TABLE V. MODEL AND CONTROLLER PARAMETERS

Symbol	Description	Value	Source
K_1	Trust-state gain	0.8	[5]
K_{12}	Cross-coupling gain	0.15	Est.
K_2	Verification-load gain	0.9	Measured
τ_1	Cognitive time constant	2.0 s	[1]
K_{p1}	Loop 1 proportional gain	1.2	Design
K_{i1}	Loop 1 integral gain	0.3	Design
K_{p2}	Loop 2 proportional gain	0.8	Design
K_{i2}	Loop 2 integral gain	0.5	Design
α	Gaze weight	0.5	Est.
β	Query weight	0.3	Est.
r_1	Trust-state reference	0.7	Config.
r_2	Verification-load reference	0.4	Config.

TABLE VI CLOSED-LOOP POLE LOCATIONS

Loop	Pole 1	Pole 2	Damping	τ_d
1 (Trust)	-0.1434	-0.8366	Overdamped	7.0 s
2 (Load)	-0.2616	—	First-order	3.8 s

Loop 2's closed-loop pole follows from $(1 + K_{p2}K_2)s + K_{i2}K_2 = 0$, giving $s = -0.2616$ ($\tau_d = 3.8$ s).

REFERENCES

- [1] A. Dey, M. Billinghurst, R. W. Lindeman, and J. E. Swan, "A systematic review of 10 years of augmented reality usability studies: 2005 to 2014," *Frontiers in Robotics and AI*, vol. 5, p. 37, 2018. doi: 10.3389/frobt.2018.00037
- [2] European Parliament, "Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)," *Official Journal of the European Union*, L series, 2024.
- [3] L. Floridi, *The Ethics of Information*. Oxford University Press, 2013.
- [4] B. Friedman and D. G. Hendry, *Value Sensitive Design: Shaping Technology with Moral Imagination*. MIT Press, 2019.
- [5] M. Groh, Z. Epstein, C. Firestone, and R. Picard, "Deepfake detection by human crowds, machines, and machine-informed crowds," *Proc. Nat. Acad. Sci.*, vol. 119, no. 1, p. e2110013119, 2022. doi: 10.1073/pnas.2110013119
- [6] J. Grubert, T. Langlotz, S. Zollmann, and H. Regenbrecht, "Towards pervasive augmented reality: Context-awareness in augmented reality," *IEEE Trans. Vis. Comput. Graphics*, vol. 23, no. 6, pp. 1706–1724, 2018. doi: 10.1109/TVCG.2016.2598626
- [7] J. L. Kröger, O. H.-M. Lutz, and F. Müller, "What does your gaze reveal about you? On the privacy implications of eye tracking," in *Privacy and Identity Management*, vol. 576, pp. 226–241, Springer, 2020. doi: 10.1007/978-3-030-42504-3_15
- [8] D. T. McRuer and H. R. Jex, "A review of quasi-linear pilot models," *IEEE Trans. Human Factors Electron.*, vol. HFE-8, no. 3, pp. 231–249, 1967. doi: 10.1109/THFE.1967.234304
- [9] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing scenes as neural radiance fields for view synthesis," *Commun. ACM*, vol. 65, no. 1, pp. 99–106, 2020. doi: 10.1145/3503250
- [10] Y. Mirsky and W. Lee, "The creation and detection of deepfakes: A survey," *ACM Comput. Surv.*, vol. 54, no. 1, pp. 1–41, 2021. doi: 10.1145/3425780
- [11] P. Müller, N. Manjunath, A. Gollnau, and T. O. Zander, "Gaze-contingent rendering as a closed-loop control problem," *IEEE Trans. Vis. Comput. Graphics*, vol. 27, no. 11, pp. 4147–4157, 2021.
- [12] Z. Pousman and J. Stasko, "A taxonomy of ambient information systems: Four patterns of design," in *Proc. Working Conf. Advanced Visual Interfaces (AVI '06)*, pp. 67–74, ACM, 2006. doi: 10.1145/1133265.1133277
- [13] L. Rosenthal, "Establishing content provenance and authenticity at scale: The C2PA approach," in *Proc. ACM Web Conf. 2022*, pp. 3547–3548, 2022. doi: 10.1145/3485447.3512286
- [14] M. Sporny, D. Longley, and D. Chadwick, "Verifiable credentials data model v1.1," *W3C Recommendation*, 2022. <https://www.w3.org/TR/vc-data-model/>
- [15] X. Zhao et al., "Invisible image watermarks are provably removable using generative AI," *arXiv preprint arXiv:2306.01953*, 2023. doi: 10.48550/arXiv.2306.01953
- [16] B. Lal, M. Ajbani, A. H. Sodhro, R. Gravina, and D. O'Keeffe, "Efficient ECG compression and reconstruction via deep compressive sensing and hardware-efficient models," *Array*, p. 100791, 2026.
- [17] B. Fazzinga, S. Flesca, and F. Furfaro, "Revisiting the notions of extension and acceptance over incomplete abstract argumentation frameworks," *Artificial Intelligence*, p. 104507, 2026.
- [18] K. Das, J. Poovancheri, S. Flesca, C. R. Calidonna, and D. Chen, "Emerging trends in wildfire detection through the lens of computer vision and wildfire emission quantification: A comprehensive survey," *IEEE Access*, 2026.
- [19] K. Das, F. Abid, J. Rasheed, T. Asuroglu, S. Alsubai, and S. Soomro, "Enhancing communication accessibility: UrSL-CNN approach to Urdu sign language translation for hearing-impaired individuals," *Computer Modeling in Engineering & Sciences*, vol. 141, no. 1, p. 689, 2024.