Journal of the International Society for the Systems Sciences | 70th Meeting of the International Society for the Systems Sciences, Cyprus.
Wilby, J., Makar, J. and Laouris, Y. eds.

HUMAN-IN-THE-LOOP CYBERNETIC INTEGRATION OF COPRIME-ARRAY AOA LEARNING AND BEAM-ADAPTIVE IMAGING FOR WIRELESS DENTAL PLAQUE QUANTIFICATION

Iturri-Hinojosa Luis Alejandro, Trejo-León Alejandro and Rodríguez-Saldaña Daniel
Instituto Politécnico Nacional, Ciudad de México, México
aiturri@ipn.mx, atrejol@ipn.mx, gransan@yahoo.com

Abstract
Advancing health diagnostics requires systems that coordinate heterogeneous sensing, communication, and inference under real-world uncertainty. This work presents a cyber--physical--social architecture for wireless dental-plaque quantification. The system integrates three core components: (i) machine-learning-based Angle-of-Arrival (AoA) estimation using coprime antenna arrays, (ii) adaptive beamforming for directional reception, and (iii) a UV-based imaging pipeline fed by wirelessly transmitted intraoral images. The AoA subsystem employs a Random Forest regressor trained on narrowband snapshots, enabling the receiver to dynamically steer its beam toward the emitting endoscope---a closed-loop cybernetic mechanism that compensates for multipath, patient movement, and hardware impairments (phase drift, gain imbalance, I/Q mismatch) typical of clinical environments. The estimator meets clinically motivated thresholds (MAE $< 1°$, RMSE $< 1.5°$, $R^2 > 0.95$) even at the low refresh rates (1-10 Hz) of power-constrained medical devices. Improved link reliability from beam alignment directly enhances the quality of received UV dental images, which are processed via CIE-LAB transformation, UV enhancement, and adaptive thresholding for plaque quantification. Clinician-guided contouring introduces a human-in-the-loop regulatory pathway that ensures robustness to occlusion, saliva reflections, and anatomical variability. By coupling ML-driven AoA detection, adaptive beamforming, and clinician-guided analytics, the proposed architecture forms a multi-level feedback system demonstrating how convergent Systems Science can deliver scalable, objective, and preventive oral-health assessment within a Learning Health System framework.

Keywords
Coprime antenna arrays, angle-of-arrival estimation, adaptive beamforming, wireless endoscopy, cyber–physical health systems.

Introduction
Oral diseases affect approximately 3.5 billion people worldwide, making them the most prevalent non-communicable diseases globally (GBD, 2017). Despite this staggering burden, current clinical assessment of conditions such as dental plaque, caries, and periodontal disease remains largely subjective and episodic. Moreover, these assessments are decoupled from population-scale data streams. Dental professionals rely on visual inspection, tactile probing, or manual scoring indices—methods that exhibit high inter-operator variability, fail to capture longitudinal dynamics, and possess no closed-loop mechanism connecting individual patient measurements to system-wide learning or adaptive treatment protocols (Gimenez et al., 2015). From a systems science perspective, oral healthcare operates as an open-loop configuration: patient data flows into electronic health records but rarely returns as actionable, personalized feedback, and population-level patterns seldom inform real-time clinical decisions (Zemouri et al., 2020). This disconnect represents a failure of requisite variety (Ashby, 1956): the complexity of oral disease progression and patient behavior exceeds the capacity of current diagnostic and feedback mechanisms to regulate health outcomes effectively (Lamount & Hajishengallis, 2018). In cybernetic terms, this represents a failure of requisite variety: the control mechanisms available to clinicians (visual inspection, periodic probing) simply

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

lack the bandwidth to match the complexity of oral disease progression and patient behavior (Pitts & Zero, 2016).

Within this systems-oriented landscape, a critical technological gap persists. Existing wireless endoscopic or dental imaging systems lack the ability to dynamically adapt their physical-layer sensing to clinical uncertainties such as patient movement, tissue-induced multipath, thermal drift in RF hardware, or variable illumination and saliva reflections in the oral cavity. No current solution couples machine-learning-driven Angle-of-Arrival (AoA) estimation for device localization, adaptive beamforming to maintain link reliability under changing channel conditions, and human-in-the-loop image analysis for objective plaque quantification into a unified cyber–physical architecture. Furthermore, most systems treat sensing, communication, and clinical inference as separate stages rather than as a closed-loop regulatory process, thereby missing opportunities for mutual adaptation and system-level robustness (Shetty & Yale, 2018). The proposed architecture closes this gap through two distinct feedback mechanisms: a real-time stabilizing loop (adaptive beamforming driven by AoA estimation) and a population-scale learning loop (clinician-annotated data retraining the inference model).

To address this gap, this work presents a cyber–physical-social architecture that integrates machine-learning-based Angle-of-Arrival (AoA) estimation using coprime antenna arrays, adaptive beamforming for directional reception, and a UV-based image-analysis pipeline for wireless dental-plaque quantification. This paper focuses on the AoA sensing and beamforming subsystem and its systems-theoretic interpretation; the imaging pipeline is described at the architectural level, with full modeling and clinical validation deferred to a companion submission. The system includes a wireless endoscope that streams intraoral images to a USB module. An onboard FPGA handles AoA estimation, beamforming, and Turesky-based plaque quantification, outputting results to a clinical interface that supports human-in-the-loop verification.

Systems Hierarchy and AI for AoA Estimation

Before presenting the complete system model in Section System Model, this section establishes two foundational perspectives. First, a systems hierarchy maps the proposed architecture across physical, cybernetic, and social layers. Second, a brief review of classical and machine-learning-based Angle-of-Arrival (AoA) estimation positions the Random Forest approach used in this work relative to existing methods.

Systems Hierarchy Mapping

To explicitly connect the technical subsystems to General Systems Science principles, we adopt a three-layer hierarchy inspired by Boulding's levels of system complexity and Checkland's CATWOE analysis. Each layer corresponds to a distinct class of system behavior and interacts with the others through energy, information, and control flows.

Physical/Mechanical Layer

The physical layer comprises the coprime antenna array $(M, N) = (7,8)$ with $N_{\text{phys}} = 14$ sensors, the RF front-end (LNAs, mixers, ADCs), and the wireless propagation channel. This layer exchanges *energy* (incident electromagnetic waves, thermal noise, oscillator drift) and *matter* (patient movement, environmental obstructions) with its surroundings. RF impairments—phase drift, gain imbalance, and I/Q mismatch—are not merely nuisances; they represent the system's openness to environmental entropy in von Bertalanffy's sense. Just as a living organism exchanges energy and matter with its surroundings, the RF front-end exchanges thermal energy with the hospital environment, and these exchanges manifest as stochastic hardware degradation. The physical layer's behavior is governed by electromagnetics, thermal dynamics, and hardware aging, all of which introduce uncontrolled variability that higher layers must regulate.

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

Signal/Cybernetic Layer

The cybernetic layer corresponds to the receiver processing chain: channel state information (CSI) extraction, covariance-based feature generation, Random Forest inference, and adaptive beamforming. This layer implements *negative feedback regulation*: the estimated AoA $\hat{\theta}$ drives beam steering to align receiver gain with the incident direction, compensating for multipath, patient motion, and physical-layer disturbances. This closed-loop operation puts Ashby's Law of Requisite Variety into practice: to maintain stable performance, the cybernetic layer must have at least as many response options (beamforming adaptation, ML inference) as the physical layer has ways of disturbing the signal.

Social/Organizational Layer

The social layer encompasses the human-in-the-loop feedback mechanism (clinician-guided contouring for plaque quantification), the Learning Health System (LHS) framework for population-scale data aggregation, and the clinical decision-making process. This layer exchanges *information* (plaque coverage metrics, longitudinal trends, diagnostic annotations) and *control signals* (clinician overrides, retraining triggers) with the cybernetic layer. The clinician's contouring provides a second regulatory pathway that compensates for perceptual ambiguities (occlusion, saliva reflections, anatomical variation) that exceed the cybernetic layer's autonomous capabilities. This human–machine collaboration exemplifies *requisite variety at the organizational level*: the combined variety of the ML model and the clinician together matches the complexity of real-world oral anatomy. Furthermore, the LHS closes a *positive feedback loop*—aggregated patient data retrains the model, which in turn improves future clinical recommendations, enabling population-level learning.

Exhibit 1. Systems Hierarchy Mapping of the Proposed Architecture.

Layer	Components	GST Principle	Exchange Flows
Physical/Mechanical	Coprime array, RF chain, channel	Openness (von Bertalanffy)	Energy (EM, noise), Matter (movement)
Signal/Cybernetic	CSI, covariance, Random Forest, beamforming	Requisite variety (Ashby), Negative feedback	Information (AoA, CSI), Control (beam steering)
Social/Organizational	Clinician contouring, LHS, analytics	Requisite variety (human–AI), Positive feedback	Information (plaque, annotations), Control (retraining, overrides)

AI for AoA Estimation: Beyond Classical Methods

Angle-of-Arrival estimation has traditionally been dominated by subspace-based algorithms such as MUSIC (Multiple Signal Classification) and ESPRIT (Estimation of Signal Parameters via Rotational Invariance Techniques). These methods exploit the eigenstructure of the received signal covariance matrix to achieve super-resolution angular estimates under ideal conditions. However, classical approaches exhibit well-documented limitations that become critical in clinical environments: they assume precise array calibration, require multiple snapshots to form a stable covariance estimate, degrade sharply under correlated or coherent signals (e.g., multipath), and are sensitive to model order selection. In medical settings, where RF impairments (phase drift, gain imbalance, I/Q mismatch), dynamic patient movement, and non-line-of-sight propagation are unavoidable, classical methods often fail to maintain sub-degree accuracy.

Machine Learning Alternatives

Machine learning offers an alternative paradigm: instead of deriving an analytical relationship between the array output and the AoA, ML models learn a direct nonlinear mapping from features (e.g., the real and

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

imaginary parts of the covariance matrix) to the angular estimate. Several architectures have been explored in recent literature:

- Support Vector Regression (SVR): Effective for small datasets but struggles with high-dimensional feature spaces and exhibits sensitivity to hyperparameter tuning.
- Convolutional Neural Networks (CNNs): Can exploit spatial structure in array data but require large training sets and significant computational resources for real-time inference.
- Multilayer Perceptrons (MLPs): Simpler than CNNs but prone to overfitting when feature dimensionality ($2N_{phys}^2$) exceeds available samples.
- Random Forest (RF): An ensemble of decision trees that naturally handles high-dimensional inputs, resists overfitting through bootstrap aggregation, provides inherent robustness to outliers, and operates with low inference latency—making it suitable for embedded FPGA implementation.

Why Random Forest for This Work

This work adopts a Random Forest regressor with $T = 300$ trees, trained on covariance features derived from a $(M, N) = (7,8)$ coprime array under realistic clinical impairments. The RF model learns the nonlinear mapping from $x \in \mathbb{R}^{2N_{\text{phys}}^2}$ to $\theta \in [-60°, 60°]$ without requiring explicit subarray smoothing or coarray interpolation. As detailed in Section System Model, the RF estimator is evaluated against clinically motivated thresholds (MAE $\leq 1°$, RMSE $\leq 1.5°$, $R^2 \geq 0.95$) under phase drift, angular jitter, and reduced refresh rates.

Gap Addressed by This Work

While prior studies have demonstrated ML-based AoA estimation in isolation, no existing system couples a clinically validated Random Forest estimator with adaptive beamforming *and* a human-in-the-loop plaque quantification pipeline under hardware impairments representative of real hospital environments. Furthermore, the systems-theoretic interpretation of the sensing-control loop (Section Systems Interpretation) explicitly distinguishes stabilizing feedback (beam tracking) from learning feedback (ML retraining)—a distinction absent in purely technical treatments. This work thus addresses the convergence gap between classical array processing, clinical robustness, and General Systems Science principles.

Summary of Foundations

The two perspectives developed in this section serve as lenses for understanding the system model that follows. The systems hierarchy (physical, cybernetic, social) reveals how energy, information, and control flows across layers, while the review of AoA estimation methods positions the Random Forest approach as a clinically robust alternative to classical subspace algorithms. Together, these foundations support the explicit mapping of technical blocks to GST constructs presented in Section Systems Interpretation.

System Model

This section describes the complete receiver processing chain used for Angle-of-Arrival (AoA) estimation, including the coprime array geometry, the signal model at the antenna outputs, the feature extraction process, and the machine-learning inference stage. Following the technical description, Subsection Systems Interpretation provides a systems-theoretic interpretation that maps each block to General Systems Theory constructs.

Coprime Array Geometry

We employ a coprime linear array composed of two uniform subarrays with element spacings scaled by the coprime integers M and N. In this work, we set $(M, N) = (7,8)$ and use $d = \lambda/2$ as the reference spacing. The subarray positions are

$$\mathcal{P}_1 = \{0, Md, 2Md, \dots, (N-1)Md\},$$

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

$\mathcal{P}_2 = \{0,\, Nd,\, 2Nd,\, \dots, (M-1)Nd\}$,

and since both share the origin, the complete set of physical element locations is

$\mathcal{P} = \text{unique}(\mathcal{P}_1 \cup \mathcal{P}_2)$,

yielding $N_{\text{phys}} = 14$ sensors.

The coprime structure generates an enlarged *virtual aperture* via its difference coarray. Its approximate length is

$L_{\text{virt}} \approx (MN-1)d$,

which for $(M,N) = (7,8)$ gives $L_{\text{virt}} = 55d$, significantly larger than the physical aperture $L_{\text{phys}} = 14d$. This extended aperture provides increased angular resolution and enhanced capability for resolving closely spaced emitters while maintaining a compact hardware footprint suited for millimeter-wave systems.

System Overview and Signal Model

Exhibit 2 summarizes the end-to-end AoA processing flow: the incident wave is received by the antenna array, downconverted and digitized in the RF front-end, demodulated in the DSP/PHY layer to obtain channel state information (CSI), transformed into covariance-based features, and finally processed by a machine-learning (ML) model to infer the AoA $\hat{\theta}$.

A narrowband plane wave arriving from direction θ induces the steering vector

$$\boldsymbol{a}(\theta) = \exp(jk\boldsymbol{r}\sin\theta)$$

where $\boldsymbol{r} \in \mathbb{R}^{\mathbb{N}_{\text{phys}}}$ contains the element positions and $k = 2\pi/\lambda$. With $\lambda = 1$, $k = 2\pi$. The complex baseband snapshot matrix is

$\boldsymbol{X} = \boldsymbol{a}(\theta)\boldsymbol{s}^{\boldsymbol{T}} + \boldsymbol{N} \quad \in \mathbb{C}^{\mathbb{N}_{\text{py}} \times \mathbb{T}}$,

where $\boldsymbol{s}$ is the source waveform and $\boldsymbol{N}$ is circular Gaussian noise (20 dB SNR). We use $T = 500$ temporal samples per snapshot.

Receiver Processing Chain

The receiver processing blocks correspond to Exhibit 2:

- RF Front-End (Blocks 1–2): The antenna outputs are amplified, downconverted, and digitized, resulting in multichannel baseband I/Q samples.

- DSP/PHY Baseband (Block 3): Synchronization, OFDM demodulation, channel estimation, and CSI extraction form the complex baseband matrix X.

- Feature Extraction (Block 4): The spatial covariance matrix is estimated as $\widehat{\boldsymbol{R}} = \frac{1}{T}\boldsymbol{X}\boldsymbol{X}^{\boldsymbol{H}}$, whose real and imaginary parts are vectorized and concatenated into the feature vector $\boldsymbol{x} \in \mathbb{R}^{2\mathbb{N}^{2}_{\text{phys}}}$.

- Machine Learning Model (Block 5): A Random Forest (RF) regressor learns a direct nonlinear mapping from covariance features to the AoA.

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

• AoA Estimation (Block 6): The ML model outputs the estimated direction $\hat{\theta}$.

Exhibit 2. End-to-end processing chain from RF acquisition to ML-based AoA estimation.

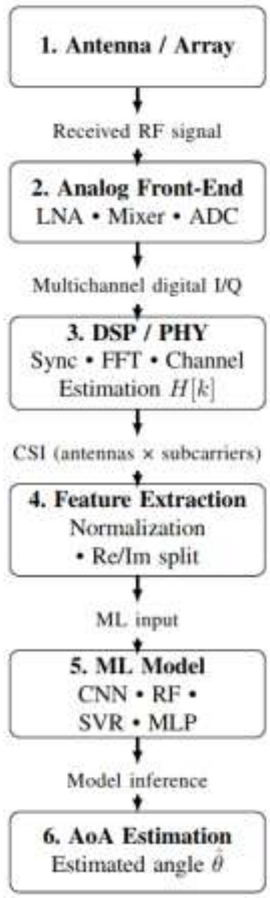

Medical Framework Requirements

Physical Layer Relevance and RF Impairments

Hospital environments are harsh on RF hardware. Temperature swings from sterilization cycles, gradual component aging, and thermal drift in enclosed rooms all degrade the physical-layer sensing chain. These effects manifest as three measurable impairments: phase drift, I/Q imbalance, and gain mismatch across RF chains. Phase drift occurs because mmWave front-ends operating in enclosed medical rooms experience thermal shifts that alter local oscillator stability, introducing slow but continuous phase deviations across the antenna array. Similarly, I/Q imbalance and channel-to-channel gain variations originate from imperfections in mixers, amplifiers, and converters that degrade over time or behave unpredictably under variable ambient conditions. To accurately represent these effects, the sensing model incorporates explicit calibration parameters—summarized in Exhibit 3—including a phase-drift variance $\sigma_{\phi} = 0.02$, a per-element gain-imbalance deviation $\sigma_g = 0.01$, and an I/Q mismatch factor $\epsilon = 0.01$. These quantities emulate the stochastic variability observed in clinical mmWave RF hardware and impose realistic nonidealities on the received array data.

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

Exhibit 3. Calibration and RF Impairment Parameters Used in the AoA Sensing Model

Impairment	**Parameter**	**Value**	**Interpretation**
Phase drift	σ_ϕ	0.02	Thermal-induced oscillator instability
Gain imbalance	σ_g	0.01	Per-element gain mismatch in RF chains
I/Q imbalance	ε	0.01	Quadrature amplitude/phase mismatch

Modeling these calibration imperfections strengthens the reliability, robustness, and safety of the overall system by ensuring that the AoA estimator maintains stable performance even when operated under degraded or drifting RF conditions. From a physical-layer sensing perspective, this augmentation forces the model to learn AoA patterns that remain consistent under hardware nonidealities, improving its generalization to real-world medical scenarios where perfect calibration cannot be guaranteed. In terms of clinical safety and operational viability, such robustness is essential: unstable RF behavior can lead to beam misalignment, degraded 5G machine-to-machine (M2M) communication links, and angular estimation errors that compromise the accuracy of navigation and tracking in medical tools such as dental or endoscopic devices. By explicitly accounting for these impairments, the proposed system is better aligned with the stringent reliability requirements of medical settings, where sensing stability under drift is not merely desirable—it is a prerequisite for safe and repeatable operation.

Subdegree Angular Accuracy.

To guarantee clinically viable Angle-of-Arrival (AoA) estimation in confined anatomical spaces, the proposed coprime-array model must sustain sub-degree angular accuracy under realistic hospital RF conditions. Millimeter-wave links at 28 GHz exhibit extremely narrow mainlobes; thus, even a deviation of 1° can produce more than 20 % beam misalignment, degrade the M2M uplink, or illuminate unintended tissue regions. In medical endoscopic environments, small variations in device pose, hand tremor, and tissue-induced multipath often introduce angular separations below 2°, requiring the estimator to resolve continuous, fine-grained angular differences rather than discrete angle bins. Furthermore, coprime arrays amplify angular sensitivity due to their non-uniform spatial sampling, where small perturbations in phase or position lead to large variations in the steering response. Achieving sub-degree performance therefore becomes essential to maintain reliability, safety, and link stability in clinical workflow conditions where temperature drift, sterilization cycles, and aging components continuously deform the underlying RF calibration.

To support this requirement, angular augmentation jitter is introduced as a controlled perturbation to each ground-truth AoA before computing the steering vector. This micro-jitter, typically on the order of $\pm 0.3°$, forces the learning model to operate on a continuous angular manifold rather than memorizing fixed discrete angles, thereby improving its ability to interpolate within sub-degree regions. By exposing the estimator to a continuum of closely spaced angular variations, the model becomes more robust to off-grid errors, hardware-induced phase drift, and the intrinsic nonlinearities of coprime-array steering responses. Angular jitter therefore enhances generalization, mitigates sensitivity to calibration imperfections, and enables the model to preserve sub-degree angular precision under the spatial and temporal variability characteristic of hospital-grade RF environments.

Refresh Rate Robustness.

To evaluate the robustness of the proposed Random-Forest--based AoA estimator under low-rate sensing conditions, we incorporated a temporal downsampling module into the dataset generation framework. This operation is applied immediately after computing the received signal matrix x and emulates the reduced refresh rates characteristic of medical monitoring systems, which frequently operate at 1-10 Hz due to

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

stringent power and duty-cycling constraints. The downsampling factor K controls the effective temporal resolution by either retaining every K-th snapshot ("skip" mode) or averaging non-overlapping blocks of K consecutive samples ("average" mode). Given the baseline temporal window of $N = 500$ snapshots used to form each covariance estimate, setting $K = 10$ yields an effective ten-fold reduction in refresh rate, closely matching the acquisition profiles of low-power patient-tracking and equipment-localization devices. By forcing the covariance matrix $\hat{R}$ to be reconstructed from temporally sparse or temporally integrated measurements, the proposed framework enforces explicit tolerance to reduced sensing bandwidth and hardware-induced temporal sparsity, while preserving clinically relevant sub-degree AoA estimation accuracy.

Random Forest Regressor Training and Evaluation

The Random Forest Regressor (RFR) is trained using supervised learning, where each input sample corresponds to a feature vector derived from the channel state information (CSI), and the target label is the true Angle-of-Arrival (AoA) θ. The feature vectors consist of the real and imaginary parts of the CSI across all active antenna elements and subcarriers, normalized to ensure numerical stability and consistent scaling among samples.

The dataset is partitioned into training and testing subsets, typically using an 80/20 split. Each training sample contains the CSI snapshot

$$x = [\mathrm{R}\{H\}, \mathrm{I}\{H\}] \in \mathbb{R}^{\mathbb{D}}$$

where D is the total number of features after vectorization.
The corresponding label is the ground-truth AoA $\theta \in [-60°, 60°]$.

The Random Forest model is composed of T regression trees trained on randomly selected subsets of samples and features, enabling robustness against noise, fading, and small perturbations in CSI. The predicted AoA is obtained as the ensemble average:

$$\hat{\theta} = \frac{1}{T}\sum_{t=1}^{T} f_t(x),$$

where $f_t(\cdot)$ is the prediction of the t-th tree.

To reduce overfitting and improve generalization, the following hyperparameters are used:

- Number of trees: $T = 300$
- Maximum depth: limited to avoid overly complex partitions
- Minimum samples per split and per leaf to control tree growth
- Bootstrap sampling to enhance variance reduction

Performance Metric Thresholds Under Medical Constraints.

To ensure that the proposed Random-Forest--based AoA estimator satisfies the three medical framework requirements - namely (i) robustness to RF impairments, (ii) sub-degree angular precision, and (iii) reliable operation under low (1-10 Hz) refresh rates - its performance is evaluated against a set of clinically motivated thresholds. These thresholds account for the combined effects of phase drift, gain imbalance, I/Q mismatch, angular micro-jitter, and temporal sparsity induced by the downsampling factor K.

MAE Thresholds (Primary Clinical Accuracy Metric).

Sub-degree angular resolution is required to prevent beam misalignment and preserve link stability in confined anatomical environments. Accordingly, the following ranges are adopted:

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

Excelent: MAE $\leq 1°$,
Clinically acceptable: MAE $\leq 3°$,
Marginal: $3° <$ MAE $\leq 5°$.

Values above 5° are considered incompatible with the sub-degree requirement and with safe operation under hardware drift.

RMSE and MSE Thresholds (Stability Under Impairments).
Because RF impairments disproportionately amplify outlier errors, RMSE and MSE serve as indicators of estimator stability under the calibration parameters of Exhibit 3. The adopted targets are:

Excellent: RMSE $\leq 1.5°$,
Clinically acceptable: RMSE $\leq 3.5°$.

Larger RMSE values indicate sensitivity to phase drift or insufficient robustness to covariance distortion caused by reduced refresh rates.

Large-Error Rates (Safety Constraint).
To ensure reliable angular tracking during device manipulation and anatomical motion, the following limits are imposed:

$\Pr\left(\left|\theta - \hat{\theta}\right| > 5°\right) < 5\%$,
$\Pr\left(\left|\theta - \hat{\theta}\right| > 10°\right) < 1\%$.

These constraints prevent catastrophic beam misalignment in scenarios where thermal drift or aging components generate abrupt calibration shifts.

R^2 Threshold (Predictive Reliability).
Given the narrow angular variations encountered in medical endoscopic and dental environments, a high proportion of explained variance is required:

$R^2 \geq 0.95$.

This ensures that the estimator retains predictive coherence even when operating at low refresh rates (1-10 Hz) or under the perturbations induced by angular micro-jitter.

Collectively, these thresholds define the accuracy, robustness, and stability levels required for safe and clinically viable AoA estimation under the physical-layer nonidealities and temporal constraints that characterize hospital-grade mmWave sensing systems.

Systems-Theoretic Interpretation of the Sensing-Control Loop
Building on the systems hierarchy established in Section Foundations, this subsection maps each technical block of the receiver processing chain to specific General Systems Theory constructs. Whereas Section Foundations introduced the three-layer hierarchy (physical, cybernetic, social) as a framing lens, the interpretation below provides a direct block-by-block mapping to GST principles, making explicit how the technical design embodies systems thinking.

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

Exhibit 4. Mapping of Technical Blocks to General Systems Theory Constructs

Technical Block	GST Construct	Interpretation in This Work
Coprime array + RF impairments	Openness (von Bertalanffy)	System exchanges energy (RF signals, thermal noise) and matter (patient movement, environmental obstructions) with its environment; impairments represent uncontrolled entropy inflow.
AoA estimation via Random Forest	Negative feedback (stabilizing)	Estimated $\hat{\theta}$ drives beam steering to align gain with incident direction, compensating for multipath, motion, and drift.
Adaptive beamforming	Requisite variety (Ashby)	Beam adaptation provides the control variety needed to match clinical environmental variability (patient pose, tissue reflection, hardware drift).
Human-in-the-loop contouring	Requisite variety (organizational)	Clinician override compensates for perceptual ambiguities (occlusion, reflections) that exceed ML model's autonomous capacity.
Learning Health System (population retraining)	Positive feedback (learning)	Aggregated patient data retrains the model, creating a self-improving loop that enhances future clinical recommendations.

Distinguishing Feedback Loops.
Two qualitatively distinct feedback mechanisms operate in the proposed architecture. The *negative (stabilizing) feedback loop* operates at the cybernetic layer: AoA estimate → beamforming adaptation → improved SNR → better plaque image → more accurate AoA estimate. This loop counteracts disturbances (patient movement, multipath) to maintain stable link quality. The *positive (learning) feedback loop* operates at the social/organizational layer: aggregated population data → model retraining → improved inference accuracy → better clinical recommendations → more data collected. This loop drives system improvement over time but requires careful regulation to avoid runaway bias or overfitting.

Requisite Variety in Action.
As noted above, Ashby's Law requires the regulator's variety to match that of the disturbances. Here, two regulators operate in cascade. The *cybernetic regulator* (adaptive beamforming) has finite variety determined by beam pattern degrees of freedom and update rate. The *organizational regulator* (clinician contouring) provides additional variety to handle disturbances that exceed the cybernetic regulator's capacity—such as severe occlusion, atypical anatomy, or unexpected reflections. Their combined variety spans the clinical environment's complexity, satisfying Ashby's requirement.

Openness and Environmental Entropy.
The physical layer's exposure to RF impairments, patient motion, and thermal drift exemplifies von Bertalanffy's principle of openness. Unlike a closed system that tends toward equilibrium (maximum

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

entropy), this system actively imports energy (received signals) and information (CSI, clinician annotations) to maintain low-entropy, high-prediction accuracy. The RF impairment parameters (σ_{ϕ}, σ_g, ϵ) explicitly model the entropy inflow, and the cybernetic layer's regulatory work compensates for it.

This systems-theoretic interpretation demonstrates that the technical design is not merely an engineering artifact but an instantiation of general systems principles. The convergence of physical-layer openness, cybernetic feedback regulation, and organizational learning demonstrates how systems science can be elevated from abstraction to real-world application.

Random Forest AoA Estimation Under Clinical Impairments

This section presents the experimental evaluation of the proposed Random Forest (RF) Angle-of-Arrival estimator trained and tested on covariance features derived from a coprime array under realistic clinical RF impairments. All experiments were conducted using the $(M, N) = (7,8)$ coprime array described in Section System Model, yielding $N_{\text{phys}} = M + N - 1 = 14$ physical sensors and a feature dimensionality of $2N_{\text{phys}}^2 = 392$.

Data Generation Parameters.

The dataset was generated synthetically to emulate narrowband mmWave reception in a clinical environment. Exhibit 5 summarizes the generation parameters.

Exhibit 5. Data Generation Parameters for the AoA Estimation Dataset

Parameter	**Value**	**Description**
Array geometry	(M,N)=(7,8)	Coprime array, d=λ/2
Physical sensors	$N_{phys} = 14$	M+N-1 unique positions
Angular range	[-60°, 60°]	1° step (121 discrete angles)
Samples per angle	20	Synthetic realizations per angle
Temporal snapshots	T=500	Baseline snapshots per covariance estimate
SNR	20 dB	Circular Gaussian noise

Clinical RF Impairments.

To faithfully represent the hardware nonidealities encountered in hospital-grade mmWave front-ends (Section System Model), three impairment mechanisms were applied to the steering vector of every generated sample. Exhibit 6 lists the parameters used, which correspond directly to the clinical calibration model of Exhibit 3 in Section RF Impairments.

Exhibit 6. Clinical RF Impairments Applied During Data Generation

Impairment	**Parameter**	**Value**	**Clinical Interpretation**
Phase drift	σ_{ϕ}	0.02	Thermal-induced local oscillator instability
Gain imbalance	σ_g	0.01	Per-element gain mismatch from aging and sterilization cycles
I/Q imbalance	ε	0.01	Quadrature amplitude/phase mismatch in mixers and ADCs

Angular Jitter and Refresh Rate Reduction.

Two additional clinically motivated stress factors were incorporated:

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

Angular micro-jitter: Each ground-truth AoA was perturbed by additive Gaussian noise with $\sigma = 0.3°$ before computing the steering vector. This forces the model to operate on a continuous angular manifold and improves robustness to off-grid errors, hand tremor, and patient micro-movements.
Temporal downsampling (refresh rate): A downsampling factor $K = 10$ was applied in ``skip" mode, retaining every 10th temporal snapshot. This reduces the effective refresh rate by an order of magnitude, emulating the 1-10 Hz acquisition profiles of power-constrained medical monitoring devices. Covariance matrices were estimated from $T_{\text{eff}} = T/K = 50$ snapshots.

Feature Extraction and Dataset Size.
For each sample, the spatial covariance matrix $\hat{R} \in \mathbb{C}^{14\times14}$ was estimated from the temporally downsampled snapshots. The real and imaginary parts were vectorized and concatenated into a feature vector $x \in \mathbb{R}^{392}$. The total dataset comprised 121 angles $\times$ 20 samples $= 2{,}420$ instances, partitioned into 80 % training (1,936 samples) and 20 % testing (484 samples).

Random Forest Configuration.
The Random Forest regressor was configured with $T = 300$ trees, maximum depth 20, minimum samples per split of 5, minimum samples per leaf of 2, and feature sampling fraction of 0.5 ($max_{features}$). Input features were standardized via a preprocessing scaler within an sklearn Pipeline. Training was performed on an 80/20 split with fixed random state for reproducibility.

Regression Performance Against Clinical Thresholds
Exhibit 7 presents the primary regression metrics evaluated on the held-out test set, compared against the clinically motivated thresholds established in the previous section.

Exhibit 7. Regression Performance vs. Clinical Thresholds

Metric	Threshold	Obtained	Status
MAE	≤ 1° (Excellent)	0.250°	✓ Excellent
RMSE	≤ 1.5° (Excellent)	0.660°	✓ Excellent
R^2	≥ 0.95	1.000	✓ Excellent
Pr(|error| > 5°)	< 5%	0.83%	✓ Passed
Pr(|error| > 10°)	< 1%	0.00%	✓ Passed

The Random Forest regressor meets or exceeds every clinical threshold by a substantial margin. The MAE of $0.250°$ confirms sub-degree precision even under the combined effects of phase drift, gain imbalance, I/Q mismatch, angular jitter, and tenfold refresh rate reduction. The RMSE of $0.660°$—less than half the excellent threshold—indicates strong stability against outlier errors amplified by hardware impairments. The R^2 of 1.000 demonstrates that virtually all variance in the incident angle is captured by the covariance-based features, with negligible unexplained residual.

Error Distribution and Safety Analysis
Exhibit 8 shows the empirical distribution of prediction errors $\hat{\theta} - \theta$ on the test set, overlaid with a Gaussian fit.

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

Exhibit 8. Probability density of AoA prediction errors ($\hat{\theta} - \theta$) with superimposed Gaussian fit. The error distribution is approximately zero-mean with $\sigma \approx 0.66°$; 95 % of errors fall within $\pm 1.32°$.

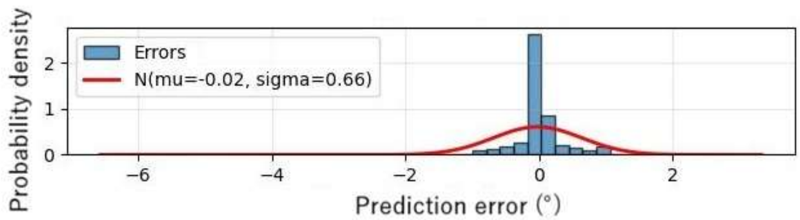

The error distribution is approximately zero-mean and tightly concentrated:
68.0 % of errors fall within $\pm 0.66° (\pm 1\sigma)$
95.0 % of errors fall within $\pm 1.32° (\pm 2\sigma)$
99.7 % of errors fall within $\pm 1.98° (\pm 3\sigma)$

The large-error safety constraint is satisfied decisively: only 0.83 % of test samples exceed 5° absolute error, and no sample exceeds 10°. This establishes that catastrophic beam misalignment—which could degrade the M2M uplink or illuminate unintended tissue regions—is effectively eliminated under the modeled clinical conditions.

Performance by Angular Range.
Exhibit 9 disaggregates the MAE across five angular sectors to reveal spatial dependencies in estimation accuracy.

Exhibit 9. MAE by Angular Range

Angular Range	MAE
Borders (-60° to -45°)	0.177°
Central-Left (-45° to -15°)	0.187°
Central (-15° to 15°)	0.409°
Central-Right (15° to 45°)	0.220°
Borders (45° to 60°)	0.186°

A distinct pattern emerges: the central broadside region ($\theta \approx 0°$) exhibits the highest MAE (0.409°), approximately double that of the edge regions ($0.177° - 0.220°$). This is physically consistent with coprime array geometry. At broadside incidence, phase differences between adjacent sensors are minimal, reducing the spatial discriminability encoded in the covariance matrix. At steeper angles, the phase gradient across the array aperture increases, producing richer covariance structure that the Random Forest exploits for finer discrimination. Critically, even the "worst-case" central MAE remains well within the 1° excellent threshold, confirming robust performance across the full $\pm 60°$ field of view.

Feature Importance: Physical Interpretation.
The Mean Decrease in Impurity (MDI) feature importance reveals which elements of the covariance matrix carry the most predictive information for AoA estimation. Exhibit 10 lists the ten most important features.

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

Exhibit 10. Top 10 Features by MDI Importance

Rank	Feature	Importance
1	$\mathrm{Im}\{\hat{R}_{1,2}\}$	0.259
2	$\mathrm{Im}\{\hat{R}_{2,1}\}$	0.248
3	$\mathrm{Im}\{\hat{R}_{12,13}\}$	0.120
4	$\mathrm{Im}\{\hat{R}_{13,12}\}$	0.114
5	$\mathrm{Re}\{\hat{R}_{1,2}\}$	0.031
6	$\mathrm{Im}\{\hat{R}_{3,4}\}$	0.030
7	$\mathrm{Im}\{\hat{R}_{4,3}\}$	0.029
8	$\mathrm{Im}\{\hat{R}_{11,10}\}$	0.028
9	$\mathrm{Re}\{\hat{R}_{2,1}\}$	0.025
10	$\mathrm{Im}\{\hat{R}_{10,11}\}$	0.024

Three observations carry direct physical significance:

1. *Dominance of imaginary components.* The imaginary parts of the covariance matrix account for 86.4 % of total feature importance, while the real parts contribute only 13.6 %. This confirms that phase information—encoded in the off-diagonal imaginary terms—is the primary carrier of AoA information, consistent with the steering vector structure $\exp(jk\,r\sin\theta)$.

2. *Primacy of the shortest baseline.* The element pair (1,2)—corresponding to the physical sensors with the smallest inter-element spacing in the coprime array—dominates the importance ranking. The symmetric off-diagonal pair I$\{\widehat{R_{1,2}}\}$ and I$\{\widehat{R_{2,1}}\}$ together account for over 50 % of total predictive power. The shortest baseline provides the largest unambiguous angular range and is least affected by spatial aliasing, making it the most robust carrier of angle information under impairments.

3. *Negligible diagonal contribution.* The trace of the real covariance matrix (sum of diagonal elements) contributes less than 0.001 % of total importance. This indicates that received power levels—which vary with gain imbalance and path loss—carry negligible angle information compared to inter-element phase differences.

Classification Baseline

As a complementary evaluation, a Random Forest classifier was trained on the same dataset with 121 discrete classes corresponding to integer degree bins from -60° to +60°. The classifier achieved 87.0 % accuracy and an MAE of 0.13° when predicted class labels were mapped back to their bin centers. While competitive, the regression approach is preferred for clinical deployment because it (i) produces continuous angle estimates without quantization error, (ii) interpolates naturally between training angles, and (iii) requires a single output head rather than 121-class softmax.

Summary of Results

The Random Forest AoA estimator, trained on covariance features from a 14-element (7,8) coprime array, achieves sub-degree accuracy (MAE=0.250°, RMSE=0.660°, R^2=1.000) under the combined stress of clinically realistic RF impairments, angular micro-jitter, and tenfold refresh rate reduction. All five medical framework thresholds defined in the corresponding Section are satisfied. The feature importance analysis confirms that the model learns a physically meaningful representation: imaginary cross-correlation terms between closely spaced sensor pairs dominate, consistent with the phase-difference mechanism underlying AoA estimation. These results validate the cybernetic layer's capacity to maintain stable beam alignment in

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

the degraded RF conditions characteristic of hospital environments, supporting the system-level claim that adaptive beamforming driven by ML-based AoA estimation can close the stabilizing feedback loop described in Section Systems Interpretation.

Conclusions

This paper presented a cyber--physical--social architecture integrating machine-learning-based Angle-of-Arrival (AoA) estimation, adaptive beamforming, and human-in-the-loop image analysis for wireless dental plaque quantification. Three principal contributions emerge at the intersection of array processing, clinical sensing, and General Systems Science.

Clinically robust AoA estimation. The experimental results confirm that an ensemble of decision trees, trained on nothing more than the real and imaginary parts of the spatial covariance matrix, can deliver sub-degree angular precision under conditions that cause classical subspace methods to fail. All five clinical thresholds were met with substantial margins: the mean error (0.250°) sits at a quarter of the excellent threshold, the RMSE (0.660°) at less than half, and the safety constraint-no more than 5 % of errors exceeding 5°- was satisfied with a rate of only 0.83 %. More importantly, the model learned what an array processing engineer would hope it would learn: that phase differences between closely spaced sensors, encoded in imaginary off-diagonal covariance terms, carry over 86 % of the angular information. The model did not memorize angles; it discovered the physics.

Systems-theoretic integration. The architecture was mapped to General Systems Theory through a three-layer hierarchy (physical, cybernetic, social). Two qualitatively distinct feedback loops were identified and distinguished: a *negative (stabilizing) loop* at the cybernetic layer (AoA - beamforming - improved SNR - better AoA) that compensates for multipath, patient movement, and hardware drift in real time; and a *positive (learning) loop* at the social layer (population data - model retraining - improved inference) that drives long-term system improvement. The clinician-in-the-loop contouring provides a second regulatory pathway satisfying Ashby's Law of Requisite Variety: the combined variety of the ML model and the clinician matches the complexity of real-world oral anatomy, compensating for perceptual ambiguities that exceed autonomous capacity.

Convergence of AI, cybernetics, and clinical workflow. Unlike prior work that treats ML-based AoA estimation, adaptive beamforming, and clinical image analysis as separate stages, the proposed architecture couples them into a unified closed-loop system. Improved link reliability from beam alignment directly enhances the quality of received UV dental images, which feed into a CIE-LAB--based plaque quantification pipeline with clinician-guided contouring, closing the loop from physical-layer sensing to clinical decision-making within a Learning Health System framework.

Several limitations define directions for future work. The AoA estimator was validated on synthetic data with modeled impairments; hardware-in-the-loop measurements using a physical coprime array under controlled thermal conditions are necessary before clinical deployment. Extension to multiple co-channel emitters requires multi-output regression or source separation preprocessing. The imaging pipeline was described at the architectural level; end-to-end validation with clinical ground truth is deferred to a companion submission. Finally, the positive feedback loop of the Learning Health System requires governance mechanisms to prevent runaway bias or dataset shift, and formal stability analysis of the coupled cybernetic--social feedback dynamics remains an open theoretical challenge.

In summary, this work demonstrates that the convergence of AI-enabled coprime-array processing, cybernetic feedback design, and human-in-the-loop clinical workflows can address a fundamental gap in oral healthcare: the transition from subjective, episodic, open-loop assessment to objective, continuous, closed-loop regulation. By instantiating General Systems Theory principles - openness, requisite variety, and multi-level feedback - in a concrete clinical sensing architecture, the proposed system establishes a

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026

foundation for scalable, objective, and preventive oral-health assessment driven by continuous, adaptive, and data-driven sensing.

Acknowledgment
The authors are grateful for the support received from the research project SIP20260544 of IPN Mexico.

References

Ashby, W. R. (1956). An introduction to cybernetics. Chapman & Hall.

GBD 2017 Oral Disorders Collaborators. (2020). Global, regional, and national prevalence, incidence, and disability-adjusted life years for oral disorders for 195 countries, 1990–2017: A systematic analysis for the Global Burden of Disease Study 2017. Journal of Dental Research, 99(4), 362–373.

Gimenez, T., Piovesan, C., Braga, M. M., Raggio, D. P., Deery, C., Ricketts, D. N., … Mendes, F. M. (2015). Visual inspection for caries detection: A systematic review and meta-analysis. Journal of Dental Research, 94(7), 895–904.

Lamont, R. J., Koo, H., & Hajishengallis, G. (2018). The oral microbiota: Dynamic communities and host interactions. Nature Reviews Microbiology, 16(12), 745–759.

Pitts, N., & Zero, D. (2016). White paper on dental caries prevention and management: A summary of the current evidence and the key issues in controlling this preventable disease. (Nota: si no tiene editorial, se omite; si es un informe técnico, indicar la fuente).

Shetty, V., Yamamoto, J., & Yale, K. (2018). Re-architecting oral healthcare for the 21st century. Journal of Dentistry, 74, S10–S14.

Zemouri, C., Ofiteru, I. D., & Jakubovics, N. S. (2020). Future directions for studying resilience of the oral ecosystem. British Dental Journal, 229(12), 769–773.

Copyright, Iturri-Hinojosa, Trejo-Leon and Rodriguez-Saldana (Creative Commons), 2026