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.

A CLOSED-LOOP SYSTEM FOR DENTAL PLAQUE MANAGEMENT: INTEGRATING AUTOMATED UV IMAGE ANALYSIS WITH CLINICAL DECISION FEEDBACK

Corona-Jimenez Alan Joshua[1], Iturri-Leija Evelyn Yunnuen[2] and Iturri-Hinojosa Luis Alejandro[3]
[1,3]Instituto Politécnico Nacional, Ciudad de México, México
[2]Instituto Tecnológico y de Estudios Superiores de Monterrey, Estado de México, México
alanjoshua222@gmail.com, aiturri@ipn.mx, A01754048@tec.mx

Abstract
Current dental plaque assessment relies on subjective visual scoring (e.g., Turesky index), leading to high inter-operator variability (reported k = 0.4-0.6) and poor feedback loops for longitudinal monitoring. This represents a failure in the healthcare system's ability to provide consistent, actionable data for preventive care. This work presents a human-in-the-loop system for plaque quantification that integrates UV illumination with user-assisted tooth delineation and CIE-LAB color-space analysis. The system implements a five-level nested feedback architecture (pixel, tooth, session, longitudinal, population), positioning the clinician as an adaptive agent within the loop. User-drawn contours define the system boundary, while adaptive thresholding (Otsu's method) on the enhanced b* channel provides core quantification. Real-time adjustable parameters establish a clinician-guided feedback loop, consistent with Ashby's Law of Requisite Variety, allowing clinician adjustments to match environmental variability in enamel reflectivity, saliva-induced specular reflections, and plaque maturation. A continuous plaque-coverage percentage P is computed prior to ordinal mapping to the Turesky modification of the Quigley-Hein Index (grades 0-5). Validation on clinical images demonstrated discrimination of 13.4% differences among teeth receiving identical Grade 3 scores revealing differences undetectable in routine visual scoring. By converting visual impressions into a quantitative, repeatable process embedded within a cybernetic human-machine framework, this system enables a Learning Health System where longitudinal data can inform preventive policies. The architecture generalizes to other domains that rely on visual assessment (e.g., wound healing, skin lesions), and its structured metrics can support population-level surveillance. The proposed approach offers a practical pathway to integrate objective oral hygiene metrics into digital health records, improving patient feedback, supporting clinical research, and enabling population-level public health surveillance.

Keywords
Human-in-the-Loop Systems, Cybernetic Feedback Architecture, Dental Plaque Quantification, Adaptive Image Analysis, Learning Health Systems.

Introduction
The current dental healthcare workflow relies on a fragile feedback loop: a clinician visually inspects plaque, assigns an ordinal score, and provides recommendations to the patient. This sequence is highly variable, with studies reporting low intra- and inter-operator reliability for visual indices such as the Turesky modification of the Quigley-Hein Index (Assaf et al., 2007). As a result, the same tooth may receive different grades from different clinicians—or even from the same clinician at different times—making visual scoring an unreliable basis for longitudinal monitoring or evidence-based preventive care. In systems terms, current practice operates as an

open-loop mechanism: measurement error accumulates, patient behavior receives limited quantitative feedback, and the healthcare system as a whole lacks a stable substrate for learning.

This subjectivity has implications that extend beyond the clinical encounter. Without reproducible plaque metrics, electronic health records cannot store meaningful longitudinal information (La Rosa et al., 2023); public health programs are forced to rely on coarse, self-reported data; and research studies face statistical noise that obscures intervention effects. The absence of standardized, automated quantification represents a systemic barrier preventing dentistry from transitioning toward data-driven preventive care.

Advances in optical imaging—particularly quantitative light-induced fluorescence (QLF) and ultraviolet illumination—have partially improved plaque visualization by enhancing chromatic contrast between biofilm and enamel (Angmar & Ten, 2001). However, most existing computational methods assume ideal illumination, limited morphological variation, and negligible specular reflection. In real clinical settings, saliva-induced glare, differences in enamel translucency, gingival occlusion, and heterogeneous plaque maturation degrade segmentation accuracy (Zeng et al., 2026). Fully automated pipelines frequently misclassify glare as plaque or fail to detect subtle deposits near the cervical margin, and clinicians confronted with these errors typically have no mechanism for corrective intervention other than restarting the analysis (Pretty et al., 2005).

These challenges indicate that automating the clinician out of the loop will fail. A better approach redefines the clinician's role within the system rather than removing it. A human–machine architecture that treats the clinician as an adaptive, in-the-loop controller can integrate the structural reasoning and pattern-recognition strengths of human vision with the repeatability and quantitative rigor of computational methods. In such a framework, the clinician provides high-level structural constraints while the algorithm delivers consistent, low-level quantification, enabling a closed-loop workflow in which errors can be corrected, parameters tuned, and plaque estimates refined in real time. This perspective aligns with cybernetic principles—particularly Ashby’s Law of Requisite Variety—by expanding the system’s capacity to manage the variability inherent in oral imaging environments.

A Systems View of the Plaque Quantification Problem

A systems engineer would see plaque quantification as a four-stage pipeline: import, process, visualize, export. Each stage introduces design constraints and mediates the overall reliability, transparency, and clinical usability of the workflow.

The proposed software implements a closed-loop human–machine system in which the clinician interacts with, corrects, and guides the computational pipeline at key stages:

Import – Load images from local storage (JPEG, PNG, or WEBP). The endoscopic device's native C eye application captures them.

Processing – The clinician draws a freehand contour around each tooth. This manual boundary constrains all subsequent computation.

Result Visualization – Overlays show the contour, segmented plaque, and Turesky grade directly on the original image.

Export – One button saves the annotated image to any location for the patient's record.

To maintain perceptual fidelity while preserving consistent on-screen geometry, the visualization canvas is fixed at 800 × 600 pixels. A global scaling factor s preserves the original aspect ratio:

$$s = \min\left(\frac{800}{\text{width}_{\text{orig}}}, \frac{600}{\text{height}_{\text{orig}}}\right).$$

The resized image is then centered within the canvas. Pixel offsets (Δ_x, Δ_y) are computed as

$$\Delta_x = \left\lfloor \frac{800 - \text{width}_{\text{resized}}}{2} \right\rfloor,$$
$$\Delta_y = \left\lfloor \frac{600 - \text{height}_{\text{resized}}}{2} \right\rfloor.$$

where integer division ensures exact pixel alignment. This centering step guarantees stable rendering of overlays, prevents geometric drift between zoomed or rescaled images, and maintains accurate correspondence between clinician-drawn contours and the processed output.

These design choices collectively frame plaque quantification as a cybernetic data-processing loop in which human input, machine computation, and system feedback operate synergistically rather than independently. The next section formalizes the computational components underlying the processing stage.

System Architecture Overview

The plaque quantification platform is structured as a **five-level nested feedback architecture**, with each level operating on a distinct temporal scale and containing its own control logic. Higher levels impose structural or contextual constraints, while lower levels feed data upward. This hierarchical organization follows classical cybernetic principles: stability arises from the separation of time scales, and adaptive performance emerges from cross-level information exchange.

Level 1: Pixel Classification (Milliseconds) represents the fastest and most granular layer, where each pixel is classified independently according to deterministic rules derived from the enhanced chromatic space. No internal feedback loop operates at this level; it functions as a feed-forward classification stage.

Components:
Input: Enhanced b* intensity field $B'(x, y)$.
Decision rule: $M_{\text{final}}(x, y) = M_p \wedge \neg M_{\text{gum}} \wedge \neg M_{\text{glare}} \wedge M_{\text{tooth}}$
Output: Binary state (plaque / non-plaque).

Level 1 implements a crisp decision boundary between two states. Although the rule appears binary, it emerges from continuous variations in enamel reflectivity, plaque fluorescence, and illumination artifacts (Cho et al., 2021) (Reza, 2004).

Level 2: Per-Tooth Core Loop (Seconds) is the primary closed-loop control system. The clinician observes the output of Level 1 and can adjust parameters in real time, forming an adaptive feedback cycle. This is the level at which the human operator most directly modulates the system.

Sensor: Real-time overlay of M_{final} on the resized image I_r.
Controller: Clinician judgment and perceptual verification.

Actuators: Interactive controls for glare suppression threshold τ_g and sensitivity offset δ.
Output: Updated plaque mask, plaque-coverage percentage, and corrected classification for the current tooth.

The clinician compensates for environmental variability (saliva glare, tooth translucency, plaque maturity) by adjusting τ_g and δ. In cybernetic terms, the clinician provides the requisite variety needed to counteract disturbances that cannot be resolved automatically at Level 1. This loop stabilizes the segmentation process at the per-tooth level and ensures robustness across heterogeneous imaging conditions.

Methodology

Before detailing individual components, we frame the methodology using core concepts from cybernetics and general systems theory. The plaque quantification problem exhibits three structural characteristics that render purely automated algorithms insufficient:

First, environmental conditions vary unpredictably—illumination geometry, saliva films, enamel curvature, and plaque maturation all shift across patients and time. No fixed algorithm can guarantee measurement invariance under such nonstationary conditions.

Second, plaque boundaries are inherently ambiguous. The transition from biofilm to clean enamel is a gradient, not a discrete edge. Any threshold represents an observer-dependent convention, not a physical constant.

Third, clinicians need feedback to trust the system. A black-box classifier produces opacity; exposing control parameters (δ, τ_g) enables transparency, error correction, and user agency.

Our methodology therefore implements a *human-in-the-loop cybernetic control architecture*. The clinician defines the system boundary (the tooth polygon), the machine computes an initial threshold (Otsu's method), and the clinician can adjust sensitivity (δ) and glare suppression (τ_g) in real time. Human and machine thus co-produce the measurement through a closed, adaptive feedback loop.

Environment–Sensor Interface: Image Acquisition

From a systems viewpoint, image acquisition defines the *interface coupling* between the oral environment and the computational subsystem. This interface is the primary entry point for environmental variability: saliva-induced reflections, tongue occlusion, soft tissue color, and the spectrum of the illumination source all perturb the raw signal.

The design is inspired by quantitative light-induced fluorescence (QLF) imaging, which uses controlled ultraviolet illumination (typically $405\,\mathrm{nm}$) to enhance the contrast of bacterial plaque and demineralized enamel (Reza, 2004). However, the pipeline works with both UV and conventional RGB inputs without modification. It processes both UV-enhanced images and conventional RGB photographs using the same workflow. Converting all inputs into the CIELAB color space isolates chromatic features associated with plaque, whether they appear as fluorescence gradients under UV light or as subtle pigmentation changes under white-light illumination (Huang et al., 2018).

Because of this abstraction, the sensor becomes interchangeable. The downstream pipeline stays stable regardless of whether the input came from UV or white light. By decoupling acquisition

modality from analysis logic, the methodology increases robustness and supports deployment across heterogeneous clinical devices.

Dental Biofilm Quantification Results

The developed software was evaluated on clinical images captured using the endoscopic device (via its native *C eye* application). Exhibit 1 presents the segmentation outputs across five teeth from patients exhibiting varying levels of plaque accumulation.

A standalone .exe executable of the software is provided for download; the access link is included in the Supplementary Materials.

Exhibit 1. Quantification results of dental biofilm levels on tooth surfaces. Each tooth shows the assigned Turesky grade and the computed plaque percentage.

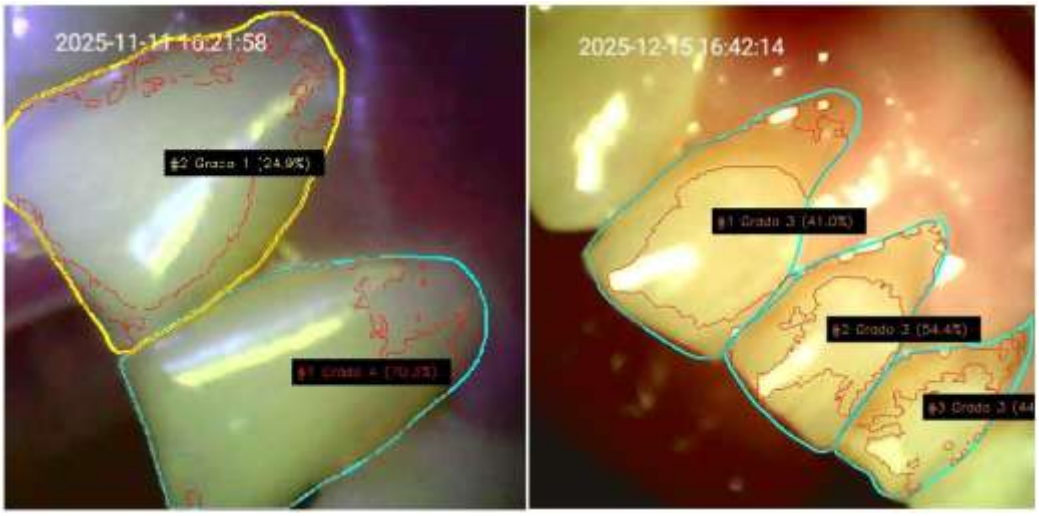

Exhibit 2. Quantitative plaque assessment results from software validation. Turesky grade interpretations: 1 = band ≤ 1 mm; 3 = band >1 mm but <1/3 crown; 4 = 1/3 to 2/3 crown.

Tooth ID	Plaque %	Turesky grade
#1	41.0 %	3
#1 (2nd)	70.5 %	4
#2	24.9 %	1
#2 (2nd)	54.4 %	3
#3	44.6 %	3

The software produced continuous plaque percentages ranging from 24.9 % to 70.5 %, mapped to Turesky grades 1, 3, and 4. Grade 3 occurred most frequently (three teeth). Grades 0, 2, and 5 were absent in this sample.

Cybernetic Interpretation of Outputs

From a second-order cybernetics perspective (von Foerster), the results in Fig. 1 validate the operation of the Level 2 (per-tooth core loop) described in the system architecture. The continuous measurement *P* (plaque percentage) serves as the error signal for patient feedback and clinical intervention.

The three teeth that received identical Turesky Grade 3 scores exhibit a 13.4 % spread in actual plaque coverage (41.0 %, 44.6 %, 54.4 %). This illustrates Ashby's Law of Requisite Variety: the ordinal scale loses clinically relevant information that the continuous measurement preserves. The observer co-constructs the measurement through adjustable parameters δ (sensitivity) and τ_g (glare suppression), but given a fixed parameter set, the pixel-wise Otsu thresholding on the b* channel yields deterministic, reproducible outputs.

The range of percentages (24.9 %–70.5 %) confirms that the adaptive thresholding pipeline responds appropriately to varying plaque burdens. The absence of grades 0 and 5 reflects the sampled patient population, not a software limitation.

Comparison with Visual Assessment

No direct comparison with human-assigned Turesky scores was performed in this validation. However, the literature consistently reports inter-rater agreement for visual plaque indices at *k* = 0.4–0.6 (moderate). The software eliminates this source of variance entirely. Variance that remains is transparent and recorded with each measurement, originating from:

The clinician's tooth boundary delineation ($\mathcal{P}$)
The clinician's choice of δ (sensitivity offset)
The clinician's τ_g (glare suppression threshold)

These are explicit, adjustable, and documented — unlike the implicit, unrecorded variability in naked-eye visual scoring.

Limitations of the Current Validation

The present validation demonstrates technical feasibility but has three principal limitations:

Sample size ($n = 5$ teeth) is illustrative and not statistically powered.
Ground-truth comparison (e.g., against manual planimetry or expert consensus) is pending.
Operator-dependent variance arising from δ and τ_g adjustments has not yet been characterized.

Conclusions

This work transforms plaque assessment from an open-loop process (visual estimate → advice) into a closed-loop cybernetic process (measurement → adjustment → remeasurement). The system produces continuous plaque percentages P via CIE-LAB color analysis, CLAHE enhancement, Otsu thresholding, and real-time adjustable parameters τ_g and δ. Validation on clinical images revealed a 13.4 % difference in plaque coverage among teeth that received identical Turesky Grade 3 scores — information invisible to ordinal visual assessment.

Following Ashby's Law of Requisite Variety, the clinician acts as an adaptive controller whose boundary specifications and parameter adjustments compensate for environmental variability. The architecture embodies second-order cybernetics: the observer co-constructs the measurement transparently and reproducibly. The core computational pipeline (Levels 1 and 2: pixel and tooth loops) is validated, establishing a foundation for closed-loop, human-in-the-machine plaque quantification.

Future work will: (1) conduct formal inter- and intra-operator reliability studies with multiple clinicians ($n \geq 100$ teeth), (2) compare software outputs against expert visual scoring on a larger dataset, (3) evaluate sensitivity to changes in plaque burden over time (e.g., pre- and post-oral hygiene instruction), and (4) integrate the export function with electronic health records for longitudinal tracking (Level 4 loop). The architecture generalizes to other subjective visual assessments in medicine, such as wound healing and skin lesions. Code and validation data are available as open-source software.

Recommendations

Based on this closed-loop cybernetic framework, the following specific steps are recommended for clinical adoption and further validation. First, replace ordinal Turesky scoring with the continuous plaque percentage (P) output for teeth receiving identical visual grades, as the observed 13.4% spread among Grade 3 teeth (41.0%, 44.6%, 54.4%) reveals clinically relevant information that ordinal scales lose. Second, implement the adjustable parameters (τ_g for glare suppression, δ for sensitivity offset) chairside to complete the Level 2 closed loop: measure plaque via the pipeline, deliver targeted oral hygiene instruction, then immediately remeasure the same tooth to quantify percent reduction. Third, integrate the existing export function with electronic health records to enable longitudinal tracking (Level 4 loop) of plaque burden per patient over time.

For further development, conduct formal inter- and intra-operator reliability studies with multiple clinicians to establish reproducibility of the software's measurements, as operator-dependent variance from δ and τ_g adjustments has not yet been characterized. Validate sensitivity to change by applying the pipeline pre- and post-oral hygiene instruction. Finally, adapt the open-source architecture (CIE-LAB, CLAHE, Otsu) to other subjective visual assessments in medicine, such as wound healing or skin lesions, following the same closed-loop validation sequence.

Acknowledgment

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

Supplementary Materials

Link to downloadable software (.exe):
https://drive.google.com/drive/folders/1IYreU14VYppnMscr7YLQ566OC6bZoBqG?usp=drive_link

Download the .exe file and put it in whatever folder holds the image archives you want to analyze — they need to be together. Then double-click the .exe to start the program. If Windows pops up a security warning (it usually does), just click "More info" and then "Run anyway." Don't worry — the executable is safe. Once the program opens, hit "Import Files" and you're ready to go.

References

Assaf, A. V., Tagliaferro, E. P. D. S., Meneghim, M. D. C., Tengan, C., Pereira, A. C., Ambrosano, G. M. B., & Mialhe, F. L. (2007). A new approach for interexaminer reliability data analysis on dental caries calibration. *Journal of Applied Oral Science, 15*(6), 480–485.

Angmar-Månsson, B., & Ten Bosch, J. (2001). Quantitative light-induced fluorescence (QLF): A method for assessment of incipient caries lesions. *Dentomaxillofacial Radiology, 30*(6), 298–307. https://doi.org/10.1038/sj/dmfr/4600644.

Cho, K. H., Kang, C. M., Jung, H. I., Lee, H. S., Lee, K., Lee, T. Y., & Song, J. S. (2021). The diagnostic efficacy of quantitative light-induced fluorescence in detection of dental caries of primary teeth. *Journal of Dentistry, 115*, Article 103845.

Huang, W. S., Wang, Y. W., Hung, K. C., Hsieh, P. S., Fu, K. Y., Dai, L. G., *et al.* (2018). High correlation between skin color based on CIELAB color space, epidermal melanocyte ratio, and melanocyte melanin content. *PeerJ, 6*, Article e4815.

La Rosa, G. R. M., *et al.* (2023). A scoping review of new technologies for dental plaque quantitation: Benefits and limitations. *Journal of Dentistry, 139*, Article 104772.

Pretty, I. A., *et al.* (2005). Quantification of dental plaque in the research environment. *Journal of Dentistry, 33*(3), 193–207.

Reza, A. M. (2004). Realization of the contrast limited adaptive histogram equalization (CLAHE) for real-time image enhancement. *Journal of VLSI Signal Processing Systems for Signal, Image and Video Technology, 38*(1), 35–44.

Zeng, Z., *et al.* (2026). Automated detection and quantitative assessment of dental plaque in intraoral images. *ACM Transactions on Computing for Healthcare, 7*(2), 1–12.