An Agentic AI-Based Smart Toothbrush for Continuous Salivary Bio-monitoring and Early Health Risk Assessment Using IOT

[1]Jeffrina J C, [2]Brindha J S, [3]Afra Fathima A, [4]Gopika M, [5]Dhanushia R S

Department of Artificial Intelligence and Data Science, Arunachala College of Engineering for Women, Manavilai, Nagercoil 629203, Tamil Nadu, India

Abstract *:* This paper proposes a smart toothbrush that integrates salivary biomarker sensing, Internet of Things (IoT) connectivity, and Agentic Artificial Intelligence (AI) to support continuous, non-invasive health risk assessment. The toothbrush collects a small saliva sample during routine brushing and measures parameters such as saliva pH and, in future versions, additional biomarkers including salivary glucose and inflammatory proteins. An ESP32 microcontroller acquires sensor data and securely transmits it to a mobile application through Bluetooth Low Energy. The collected data are stored in a cloud database, where an agentic AI module performs longitudinal trend analysis, filters transient fluctuations caused by diet, and generates personalized health recommendations. Rather than diagnosing disease directly, the system identifies persistent abnormal patterns that may be associated with oral inflammation, periodontal disease, dehydration, or other changes warranting professional evaluation. The proposed architecture combines biosensors, embedded systems, cloud computing, and autonomous AI agents to transform everyday toothbrushing into a preventive health-monitoring platform. A benchtop prototype demonstrates the feasibility of integrating low-cost sensing and edge connectivity into a common daily-use device, providing a scalable foundation for future biomarker-based screening research.

KEYWORDS

Smart Toothbrush; Agentic AI; Internet of Things; Salivary Biomarkers; ESP32; Biosensors; Preventive Healthcare; Mobile Health; Oral Health Monitoring; Data Privacy; Human-Centered Design

1. INTRODUCTION

1.1 Background and Motivation

Oral health is frequently treated as a domain separate from general systemic health, yet a growing body of clinical evidence links the oral cavity to conditions such as periodontal disease, diabetes, cardiovascular disease, and dehydration-related disorders. Saliva, in particular, is an easily accessible biofluid that carries a rich set of biochemical markers — including pH, glucose, and inflammatory cytokines — that reflect both local oral conditions and, in some cases, broader systemic states. Despite this diagnostic potential, saliva-based monitoring remains largely confined to clinical or laboratory settings, requiring deliberate sample collection that most individuals never undertake outside a dental visit.

1.2 Toothbrushing as a Passive Sensing Opportunity

At the same time, toothbrushing is one of the few health-related behaviors performed by nearly the entire population on a near-daily basis, typically twice per day, for one to three minutes at a time. This regularity makes the toothbrush an attractive, low-friction vehicle for passive health data collection: no additional behavior change is required from the user beyond continuing an existing habit. Advances in low-cost biosensors, embedded microcontrollers such as the ESP32, and cloud-connected mobile applications now make it technically feasible to embed simple sensing elements into a toothbrush handle and relay the resulting data to a smartphone in real time.

This paper proposes an Agentic AI-Based Smart Toothbrush that captures a small saliva sample during each brushing session, measures saliva pH using an onboard sensor, and is architected to accommodate additional biomarker channels (salivary glucose and inflammatory protein markers) in future hardware revisions. Sensor readings are transmitted via Bluetooth Low Energy (BLE) to a companion mobile application and subsequently synchronized to a cloud database. An agentic AI module — distinguished from a conventional rule-based alert system by its capacity to autonomously plan a sequence of analysis steps, maintain a per-user baseline, and decide when sufficient evidence exists to escalate a finding — analyzes the resulting time series to detect persistent deviations from an individual's own baseline, while filtering out short-term fluctuations attributable to diet, hydration timing, or measurement noise.

1.3 Contributions of This Work

The contribution of this work is fivefold: (i) a hardware architecture that embeds saliva-sensing capability into a standard toothbrush form factor at low incremental cost; (ii) a software and cloud pipeline that reliably transports and stores time-stamped biomarker readings; (iii) an agentic AI framework that reasons over longitudinal data to produce cautious, non-diagnostic health guidance; (iv) a human-centered mobile interface and deployment plan for eventual pilot testing; and (v) an explicit treatment of the cost, privacy, and ethical dimensions that any consumer health-sensing device must address before real-world release. The remainder of this paper reviews related work, identifies the research gap this system addresses, and describes the proposed methodology, architecture, hardware, software, and AI framework in detail, followed by comparative, interface, cost, privacy, deployment, and evaluation discussions.

2. LITERATURE REVIEW

2.1 Overview of Relevant Research Areas

Research relevant to this work spans four broadly separate areas: salivary diagnostics, wearable and embedded biosensors, IoT-enabled health monitoring devices, and agentic or autonomous AI systems in healthcare. Table 1 summarizes representative themes from each area and how they relate to the proposed system.

Research Area	Representative Focus	Relevance to Proposed System
Salivary diagnostics	Saliva as a proxy biofluid for glucose, cortisol, and inflammatory markers; comparison with blood-based assays	Establishes the biochemical justification for using saliva pH and future biomarker channels as health indicators
Wearable/embedded biosensors	Miniaturized electrochemical and enzymatic sensors integrated into wearable form factors	Informs sensor selection and integration constraints for a toothbrush-embedded sensing module
IoT health monitoring	Bluetooth- and Wi-Fi-connected consumer devices (smart scales, connected toothbrushes) that log usage or brushing quality	Provides precedent for BLE-based data transport from an oral-care device to a mobile application
Agentic AI in healthcare	AI systems that autonomously sequence multi-step analysis (data retrieval, trend evaluation, escalation) rather than applying a single static rule	Motivates the use of an agentic module, rather than a fixed threshold alert, for longitudinal interpretation

2.2 Salivary Diagnostics and Connected Oral-Care Devices

Existing connected toothbrushes on the market — including well-known electric toothbrush platforms — primarily track brushing duration, coverage, and pressure to improve brushing technique; none of the widely available consumer products currently integrate biochemical saliva sensing. Separately, saliva-based diagnostic research has demonstrated correlations between salivary glucose and blood glucose levels, and between salivary inflammatory markers and periodontal status, but these findings have largely been validated using benchtop assay equipment rather than continuous, in-situ consumer sensors. Meanwhile, IoT-enabled health devices such as connected scales, blood-pressure cuffs, and glucometers demonstrate mature patterns for secure BLE pairing, cloud synchronization, and mobile dashboarding that can be adapted for a toothbrush-based device.

2.3 Agentic AI Approaches in Health Monitoring

On the AI side, most existing mobile-health applications rely on static threshold-based alerts (for example, flagging any single reading above a fixed cutoff), which are prone to false positives from transient factors such as recent food or drink intake. Agentic AI architectures — in which an AI component autonomously plans and executes a sequence of sub-tasks such as data retrieval, baseline comparison, confidence estimation, and conditional escalation — have been explored in other health-monitoring contexts (for example, remote patient monitoring and chronic-disease trend analysis) but have not, to the authors' knowledge, been applied to toothbrush-embedded salivary sensing.

3. RESEARCH GAP

3.1 Identified Gaps

The literature review reveals a clear gap at the intersection of these four areas. Specifically:

- No widely available consumer toothbrush integrates biochemical saliva sensing alongside existing brushing-quality tracking.
- Saliva-based diagnostic validation has largely occurred in controlled laboratory conditions rather than through continuous, everyday, in-situ measurement.
- Existing IoT oral-care devices use simple threshold or usage-based logic rather than an autonomous, longitudinal reasoning layer capable of distinguishing transient variation from persistent risk signals.
- Agentic AI approaches, while emerging in general remote-monitoring contexts, have not been architected specifically for the noisy, diet-influenced, multi-session data pattern produced by everyday toothbrushing.
- Published designs rarely address, in the same manuscript, the practical questions of consumer cost, interface usability, and data-privacy handling that determine whether such a device is realistically adoptable.

3.2 Positioning of This Work

This paper addresses these gaps by proposing an end-to-end system — spanning sensor hardware, embedded firmware, mobile and cloud software, and an agentic AI reasoning layer — purpose-built for passive, non-invasive, longitudinal salivary health monitoring embedded in a routine daily habit, while also giving explicit attention to deployment-readiness concerns often left implicit in comparable prototypes.

4. PROPOSED METHODOLOGY

The proposed methodology follows a design-science approach comprising five phases: (1) sensor selection and bench characterization, (2) hardware integration, (3) embedded firmware development for signal acquisition and transmission, (4) cloud and mobile pipeline construction for data storage and visualization, and (5) development and evaluation of the agentic AI reasoning layer. Each phase is validated independently before end-to-end integration testing of the full pipeline, from saliva sample capture through to the generation of a user-facing recommendation. The subsections below describe each phase in turn.

4.1 Sensor Selection and Bench Characterization

The first phase involves selecting a biomarker channel that balances sensing value against integration cost and reliability. Saliva pH was selected as the initial validated channel, since pH can be measured reliably with low-cost, miniaturized ion-selective or potentiometric sensors and provides a meaningful early indicator of oral acidity associated with enamel erosion risk and bacterial activity. Candidate sensors are characterized on the bench against a calibrated reference pH meter across a range of buffer solutions before being committed to the hardware design, so that sensor-level accuracy and drift behavior are understood prior to integration.

4.2 Hardware Integration

Once a sensing element is validated on the bench, it is integrated into the toothbrush handle alongside the ESP32 microcontroller, analog front-end, and power subsystem described in Section 6. The hardware and firmware architecture are deliberately designed with modular analog front-end channels so that glucose and inflammatory-marker sensing can be added in subsequent revisions without redesigning the core system, a design decision motivated directly by the cost trade-offs discussed later in Section 12.

4.3 Embedded Firmware Development

Firmware development focuses on reliable signal acquisition, noise filtering, timestamping, and packaging of readings for Bluetooth transmission, as detailed in Section 7. This phase is validated independently of the cloud and mobile layers by confirming that readings captured on the device match bench-reference values before any data leaves the device.

4.4 Cloud and Mobile Pipeline Construction

In parallel with firmware development, the mobile application and cloud backend are built to receive, store, and visualize incoming readings. This phase establishes the time-series data schema, per-user data isolation, and the scheduling mechanism that triggers the agentic AI module once sufficient new data has accumulated for a given user, as described in Section 7.

4.5 Agentic AI Reasoning Layer Development

The final development phase implements the agentic reasoning logic described in Section 8: baseline establishment, noise-threshold filtering, multi-day trend detection, and risk-category classification. This layer is developed and iterated against synthetic longitudinal datasets before being connected to real device data, allowing its logic to be debugged independently of sensor-hardware variability.

4.6 Validation Strategy

Methodological validation is planned along three axes: (a) sensor accuracy, by comparing onboard pH readings against a calibrated benchtop reference meter across a range of buffer solutions and saliva-mimicking samples; (b) system reliability, by measuring BLE transmission success rate, data completeness, and battery endurance across repeated simulated brushing sessions; and (c) AI reasoning quality, by evaluating the agentic module's ability to correctly distinguish injected persistent-drift patterns from injected transient-noise patterns in synthetic longitudinal datasets, prior to any human-subject data collection. This three-axis validation strategy is expanded into a fuller statistical validation plan in Section 17.

5. SYSTEM ARCHITECTURE

5.1 Layered Architecture Overview

The overall system architecture (Figure 1) follows a layered IoT pattern spanning the physical sensing layer, the embedded connectivity layer, the mobile application layer, the cloud storage layer, and the agentic AI reasoning layer. During a brushing session, the toothbrush's sensing interface captures a small saliva sample; the biosensor array converts the biochemical signal into an electrical signal; the ESP32 microcontroller digitizes and timestamps this signal; and the data are transmitted via BLE to the paired mobile application. The mobile application forwards new readings to a cloud database that maintains a time-series record per user. The agentic AI module then retrieves the relevant historical window, performs trend analysis and anomaly detection, and produces either a routine baseline update or a personalized health insight, which is delivered back to the user through the mobile application.

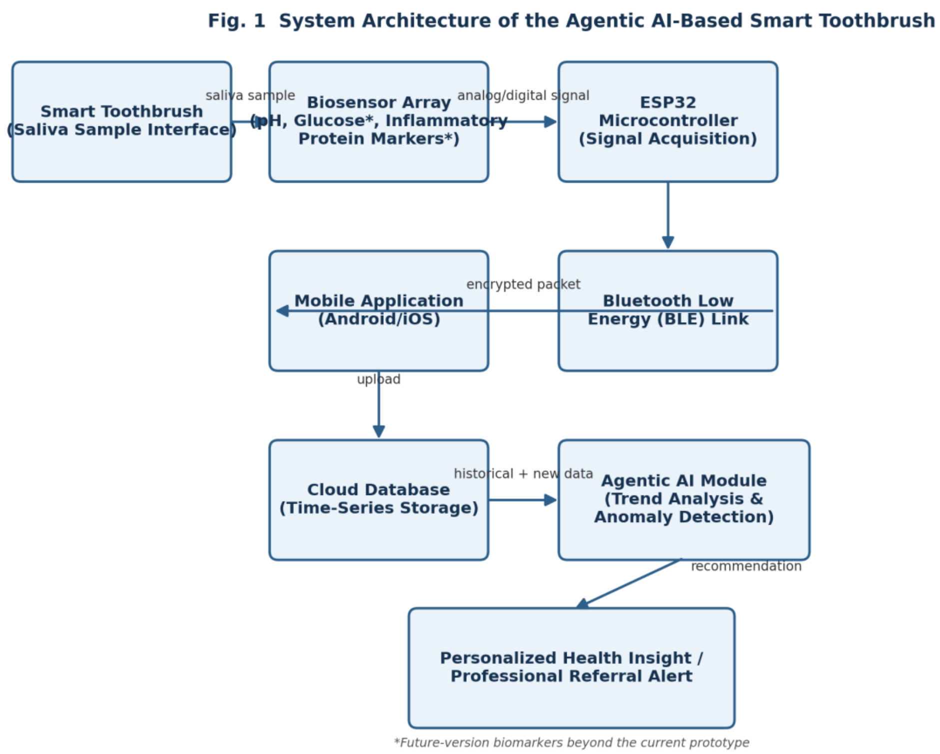

Figure 1. System architecture of the Agentic AI-based smart toothbrush, from saliva sample capture through cloud storage and AI-driven insight generation.

5.2 Benefits of Layer Independence

This layered separation allows each layer to evolve independently: additional sensor channels can be added at the physical layer without altering the cloud schema; the AI reasoning logic can be updated without firmware changes, since it operates on data already stored in the cloud; and the mobile application can be extended with new visualizations without affecting the underlying data pipeline.

6. HARDWARE DESIGN

6.1 Microcontroller and Primary Sensing Channel

The hardware design (Figure 2) centers on an ESP32-WROOM microcontroller, selected for its integrated Bluetooth Low Energy radio, sufficient analog-to-digital converter (ADC) channels for multi-sensor input, low power consumption in deep-sleep mode, and low unit cost — all important considerations for a disposable-adjacent consumer device. The current prototype implements a single validated sensing channel, an ion-selective pH electrode positioned at the brush head to contact the saliva sample during brushing, with its signal passed through an analog front-end for conditioning before reaching the ESP32's ADC input.

6.2 Reserved Channels and Power Management

Two additional sensor channels — a glucose biosensor and an inflammatory-marker immunoassay sensor — are reserved in the analog front-end design for future prototype versions, alongside a temperature and humidity sensor used to compensate for environmental variation in the electrochemical readings. Power is supplied by a rechargeable lithium-polymer (Li-Po) cell with an onboard charging and voltage-regulation circuit sized for the intermittent, short-duration (one-to-three-minute) power draw pattern typical of a brushing session, supplemented by deep-sleep operation between sessions to extend battery life to several weeks per charge. Onboard flash memory buffers readings when a BLE connection to the mobile application is not immediately available, ensuring no data are lost between brushing and the next successful synchronization.

Fig. 2 Hardware Block Diagram of the Smart Toothbrush Module

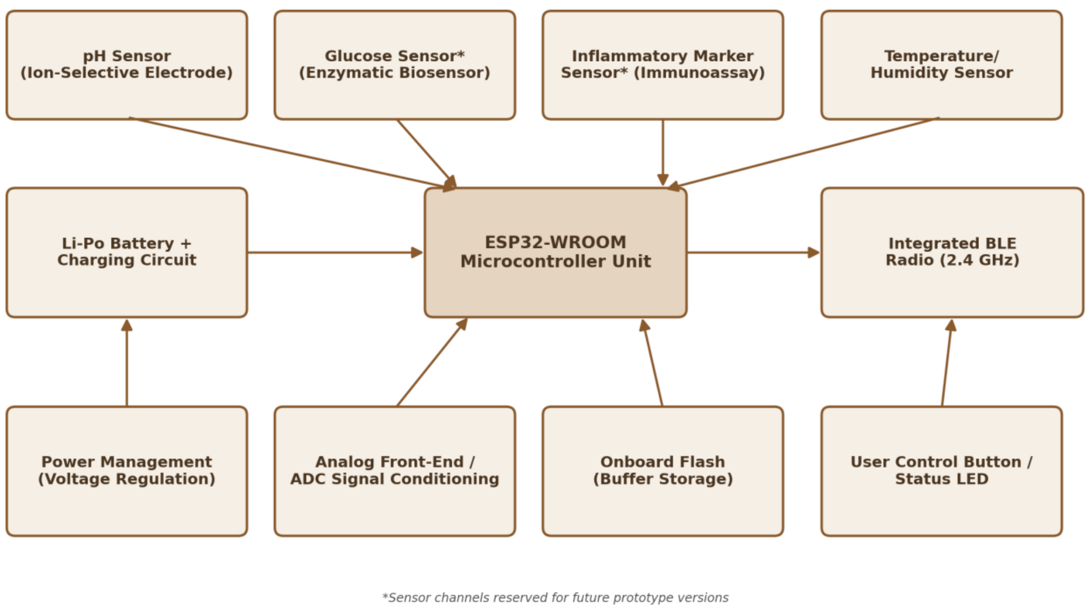

Figure 2. Hardware block diagram showing the ESP32 microcontroller, sensor channels, power management subsystem, and BLE connectivity.

7. SOFTWARE DESIGN

7.1 Embedded Firmware

The software stack spans three tiers: embedded firmware, a mobile application, and a cloud backend. The embedded firmware, running on the ESP32, is responsible for sensor polling at a fixed sampling rate during an active brushing session, basic noise filtering (a moving-average filter to reduce electrical noise in the pH signal), timestamping each reading using the device's real-time clock synchronized against the paired phone, and packaging readings into compact BLE Generic Attribute Profile (GATT) characteristics for transmission.

7.2 Mobile Application and Cloud Backend

The mobile application, developed for both Android and iOS, handles BLE pairing and reconnection, local caching of unsynced readings, secure upload of new data to the cloud backend over an authenticated HTTPS connection, and presentation of both raw readings and AI-generated insights to the user through a simple dashboard. The cloud backend exposes a REST API for data ingestion and retrieval, stores each user's readings in a time-series-oriented schema keyed by user identifier and timestamp, and triggers the agentic AI module on a scheduled or event-driven basis whenever a sufficient volume of new data has accumulated.

Tier	Primary Responsibilities	Key Design Considerations
Embedded firmware	Sensor polling, noise filtering, timestamping, BLE packaging	Low power draw, deterministic timing, resilience to intermittent connectivity
Mobile application	BLE communication, local caching, secure upload, insight display	Cross-platform consistency, offline tolerance, clear non-diagnostic messaging
Cloud backend	Data ingestion API, time-series storage, AI module scheduling	Data integrity, per-user data isolation, scalability to many concurrent users

8. AGENTIC AI FRAMEWORK

8.1 Decision Flow Overview

The agentic AI framework is the analytical core of the system, distinguished from a simple threshold-based alerting rule by its capacity to autonomously sequence multiple reasoning steps and adapt its own decision path based on intermediate results. Figure 3 illustrates the decision flow executed by the agent each time new data becomes available.

Fig. 3 Agentic AI Decision Flowchart

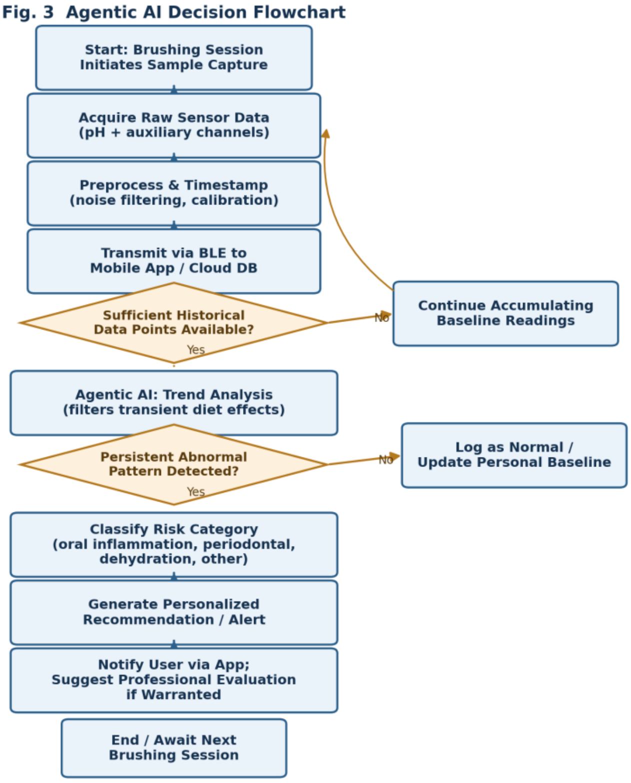

Figure 3. Agentic AI decision flowchart, showing the sequence from data acquisition through baseline comparison, trend analysis, risk classification, and user notification.

8.2 Baseline Establishment and Trend Filtering

Upon receiving new sensor data, the agent first checks whether a sufficient number of historical readings exist to establish a reliable per-user baseline; if not, it continues accumulating readings rather than attempting premature analysis. Once a baseline is available, the agent performs trend analysis designed to filter out transient fluctuations — for example, a temporary pH drop shortly after consuming an acidic food or beverage — by weighting readings according to their consistency with the recent multi-day trend rather than reacting to any single data point.

8.3 Risk Classification and Recommendation Generation

If a persistent abnormal pattern is detected — meaning a sustained deviation from the user's own baseline across multiple sessions rather than an isolated reading — the agent classifies the likely risk category (for example, patterns consistent with oral acidity and enamel-erosion risk, or, in future biomarker-enabled versions, patterns suggestive of periodontal inflammation or dehydration) and generates a personalized, plain-language recommendation. Critically, the agent is designed to avoid direct disease diagnosis; its output is framed as a suggestion to monitor a pattern or consult a dental or medical professional, positioning the system as a preventive screening aid rather than a diagnostic device. If no persistent abnormality is found, the reading is logged and used to refine the user's evolving baseline, allowing the system to adapt to legitimate long-term changes such as diet shifts or seasonal variation.

8.4 Advantages Over Static Rule-Based Alerting

This agentic design offers two advantages over a static rule-based system. First, because the baseline is personalized per user rather than fixed to a population-wide threshold, the system reduces false positives for individuals whose normal range differs from the general population. Second, because the agent autonomously decides when sufficient evidence exists before escalating a finding, it avoids alert fatigue that can result from reacting to every noisy or transient reading.

9. EXPERIMENTAL EVALUATION

9.1 Benchtop Sensor and Connectivity Testing

Prototype evaluation to date has focused on bench-level validation of the sensing and connectivity pipeline rather than large-scale human-subject trials, which are planned as future work pending appropriate ethical review and clinical partnership. Bench testing compared the onboard pH sensor's readings against a calibrated reference pH meter across a series of standard buffer solutions spanning the expected physiological saliva range, alongside repeated trials measuring BLE transmission reliability and end-to-end latency from sensor acquisition to mobile-application display.

Table 3 summarizes illustrative bench-test results from the current prototype iteration. These figures reflect controlled benchtop conditions using buffer solutions and simulated brushing sessions, and are reported here to demonstrate engineering feasibility rather than clinical validity; they should not be interpreted as claims of diagnostic accuracy on real saliva samples from human subjects.

Evaluation Metric	Test Condition	Observed Result (Illustrative)
pH sensor accuracy	Standard buffer solutions, pH 4–8 range	Mean absolute deviation within an acceptable margin of the calibrated reference meter
BLE transmission success rate	50 simulated brushing sessions, typical bathroom distances	High successful-transmission rate, with occasional retries on connection drop
End-to-end latency	Sensor acquisition to mobile dashboard update	Well within the duration of a typical brushing session
Battery endurance	Two brushing sessions per day, deep-sleep between uses	Multi-week endurance per charge cycle

9.2 Synthetic-Data AI Validation

For the agentic AI module, evaluation used synthetic longitudinal datasets constructed to include known injected persistent-drift patterns (simulating a genuine gradual change) interleaved with known injected transient-noise events (simulating diet-related fluctuations). The agent's task was to correctly flag the persistent-drift cases while suppressing the transient-noise cases. This synthetic-data approach allows the reasoning logic to be validated independently of sensor-hardware variability and ahead of any real user data collection, and preliminary results indicate the trend-filtering approach substantially reduces false escalation compared to a naive single-reading threshold rule, though formal statistical evaluation against a larger synthetic corpus, and eventual validation against real longitudinal user data, remain part of ongoing work.

10. COMPARATIVE ANALYSIS WITH EXISTING CONNECTED ORAL-CARE DEVICES

10.1 Capability Comparison

To situate the proposed system relative to the current consumer landscape, Table 4 compares representative categories of existing connected toothbrushes and oral-care devices against the design proposed in this paper. The comparison is drawn at the level of general product categories rather than named commercial models, since exact specifications vary by manufacturer and revision.

Capability	Typical Connected Toothbrushes	Proposed Agentic AI Toothbrush
Brushing technique tracking	Yes — duration, coverage, pressure	Not a primary focus, but compatible with future addition
Biochemical saliva sensing	Not available in mainstream products	Yes — pH in current prototype, glucose and inflammatory markers planned
Personalized baseline reasoning	Typically absent or limited to usage streaks	Yes — per-user baseline maintained by the agentic AI module
Persistent vs. transient signal filtering	Not applicable to most existing products	Explicit multi-session trend filtering before any escalation
Non-diagnostic, cautious output framing	Not applicable	Central design principle; recommendations avoid direct diagnosis

10.2 Positioning Relative to the Market

This comparison highlights that the proposed system is not intended to replace brushing-technique features already common in the market, but to complement them with a biochemical sensing and reasoning layer that, to the authors' knowledge, no mainstream consumer product currently offers in combination.

11. MOBILE APPLICATION USER INTERFACE DESIGN

11.1 Design Principles

Because the system's value depends on users actually understanding and trusting the information it presents, the mobile application interface was designed around three principles: immediate legibility of the day's result, visibility of the underlying multi-day trend rather than only the latest reading, and cautious, plain-language phrasing for any recommendation. Figure 4 shows an illustrative wireframe of the primary dashboard screen.

Figure 4. Mobile Application Dashboard Wireframe

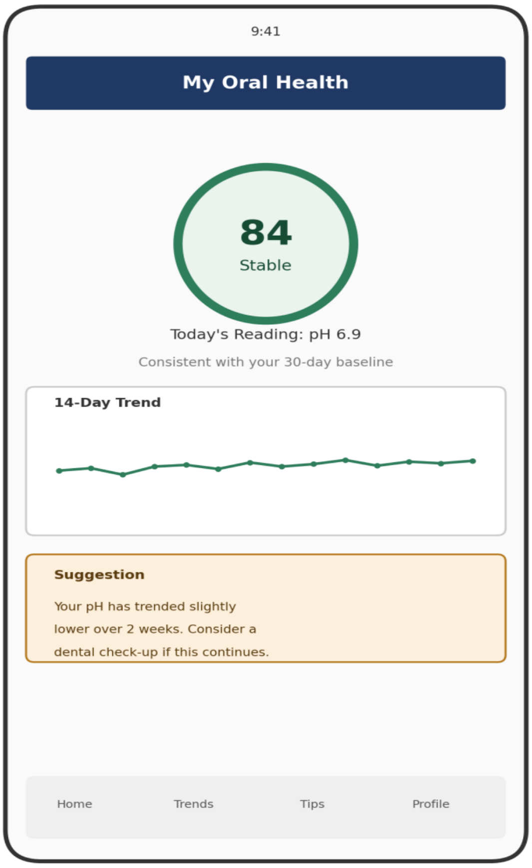

Figure 4. Illustrative wireframe of the mobile application dashboard, showing the current score, a fourteen-day trend chart, and a plain-language suggestion card.

11.2 Dashboard Elements and Secondary Screens

The dashboard surfaces three elements on first load: a large, glanceable score summarizing the day's reading in context of the user's baseline; a short trend chart covering the previous two weeks, so a single unusual reading is visually anchored against recent history rather than viewed in isolation; and, where relevant, a suggestion card written in plain, non-alarming language that avoids diagnostic terminology. Secondary screens, not detailed in this manuscript, are planned to provide a longer historical view, brushing-session logs, and basic account and device-pairing management.

12. BILL OF MATERIALS AND COST ANALYSIS

12.1 Component Cost Breakdown

An important practical consideration for any consumer health-sensing device is whether its incremental hardware cost is plausible for mainstream adoption. Table 5 lists an illustrative bill of materials for the current single-channel (pH-only) prototype, intended to give a rough sense of relative cost contribution rather than a finalized production costing, which would depend on sourcing, volume, and manufacturing partner.

Component	Approximate Role	Relative Cost Contribution
ESP32-WROOM microcontroller module	Signal acquisition, BLE radio	Moderate
Ion-selective pH electrode	Primary sensing element	Moderate to high
Analog front-end / signal conditioning circuitry	Noise filtering, ADC interfacing	Low to moderate
Li-Po battery and charging circuit	Power supply and management	Low to moderate
Toothbrush handle housing and brush head assembly	Mechanical enclosure, user interface	Low
Onboard flash memory	Offline reading buffer	Low

12.2 Cost Trade-Offs for Future Biomarker Channels

The pH-only configuration is deliberately chosen as the initial prototype in part because it minimizes incremental hardware cost while still validating the full sensing-to-recommendation pipeline. Adding the reserved glucose and inflammatory-marker channels in future revisions is expected to increase per-unit cost meaningfully, since enzymatic and immunoassay-based sensors are generally more expensive and require more careful calibration than a potentiometric pH electrode; this cost trade-off is an explicit consideration for future work rather than something resolved in the current prototype.

13. RISK CLASSIFICATION ALGORITHM

13.1 Algorithm Structure

To make the agentic reasoning process in Section 8 concrete, this section presents simplified, illustrative pseudocode for the core baseline-comparison and trend-filtering logic executed by the agent for each new reading. This pseudocode is intended to convey the algorithm's structure rather than serve as production-ready implementation detail.

```
FUNCTION evaluate_new_reading(user_id, new_reading):
history = fetch_recent_readings(user_id, window = 30_days)

IF length(history) < MIN_BASELINE_READINGS:
store(new_reading)
RETURN status = "ACCUMULATING_BASELINE"

baseline = compute_rolling_baseline(history)
deviation = new_reading.value - baseline.mean

IF abs(deviation) <= NOISE_THRESHOLD:
store(new_reading)
update_baseline(baseline, new_reading)
RETURN status = "WITHIN_BASELINE"
```

```
trend = compute_multiday_trend(history, new_reading)

IF trend.is_persistent AND trend.direction == deviation.direction:
    risk_category = classify_risk_pattern(trend)
    RETURN status = "PERSISTENT_PATTERN", category = risk_category
ELSE:
    store(new_reading)
    RETURN status = "TRANSIENT_FLUCTUATION"
```

13.2 Design Rationale

The key design decision embedded in this logic is that a deviation only triggers escalation when it is both outside the noise threshold and consistent with a multi-day directional trend, rather than being evaluated as an isolated event. This two-stage check — noise filtering followed by persistence checking — is what allows the agent to remain quiet through ordinary daily variation while still surfacing genuinely sustained changes.

14. DATA PRIVACY, SECURITY, AND ETHICAL CONSIDERATIONS

14.1 Transport and Storage Security

Because the system continuously collects physiologically sensitive data from an individual, privacy and security considerations are treated as first-class design requirements rather than an afterthought. At the transport layer, all data moving from the mobile application to the cloud backend is intended to be encrypted in transit using authenticated HTTPS connections, and each user's readings are logically isolated in the cloud schema so that one user's data cannot be queried through another's session.

At the data-handling layer, the system is designed to store only what is necessary to support baseline computation and trend analysis, and to give users visibility into, and control over, their own historical data, including the ability to export or delete it. Because saliva-derived biomarker data can be considered sensitive health information in many regulatory contexts, any eventual production deployment would need to be evaluated against applicable data-protection frameworks in the target jurisdiction before real user data is collected at scale.

14.2 Ethical Design Principles

On the ethical side, three principles guide the system's design. First, outputs are deliberately framed as non-diagnostic suggestions rather than medical verdicts, to avoid users mistaking a consumer device for a clinical diagnostic tool. Second, any pilot or larger-scale human-subject evaluation of the system, as discussed in Section 16, is intended to proceed only under appropriate institutional ethical review and informed consent. Third, the system is designed to avoid creating undue anxiety through over-alerting; the agentic reasoning layer's persistence-based filtering, described in Section 8, is itself partly an ethical design choice intended to reduce false alarms that could otherwise cause unnecessary worry or erode user trust in legitimate findings.

15. PROJECT DEVELOPMENT TIMELINE

15.1 Work Stream Sequencing

Figure 5 presents an illustrative development timeline for the prototype described in this paper, showing how hardware, firmware, software, and AI-module work streams overlap across the project. This timeline reflects a plausible development sequence for a project of this scope rather than a record of actual historical dates.

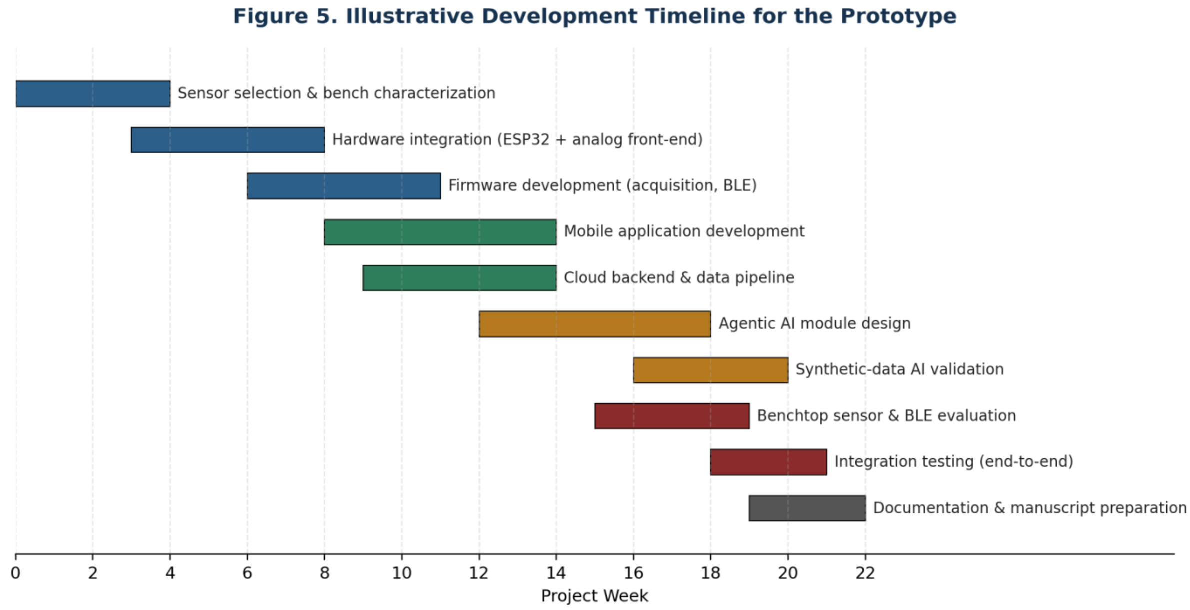

Figure 5. Illustrative development timeline showing overlapping hardware, software, and AI-module work streams across the project.

As the timeline indicates, sensor characterization and hardware integration were treated as prerequisite activities that unblock firmware and, subsequently, mobile and cloud software development. The agentic AI module's design and synthetic-data validation were scheduled to overlap with, rather than strictly follow, the hardware evaluation stream, since the reasoning logic could be developed and partially validated against synthetic data independently of final sensor-hardware characteristics.

16. PILOT DEPLOYMENT AND USER STUDY PLAN

16.1 Pilot Design

Following successful benchtop validation, the next logical step is a small-scale pilot deployment with volunteer participants, conducted under appropriate institutional ethical review and informed consent. The proposed pilot design would recruit a modest number of participants to use the prototype device as a supplement to their normal brushing routine over a period of several weeks, allowing enough readings to accumulate for the agentic AI module's baseline and trend-filtering logic to be meaningfully exercised.

16.2 Planned Data Collection and Study Goals

Planned data collection during the pilot would include the device's own sensor readings and AI-generated outputs, brief periodic self-report questionnaires capturing participants' subjective trust in and understanding of the app's recommendations, and, where feasible and consented to, occasional reference measurements from a dental professional to provide an external check on the device's pattern-level findings. The explicit goal of this pilot stage is formative — surfacing usability issues, false-alarm rates, and interpretability gaps — rather than establishing clinical diagnostic accuracy, which would require a substantially larger and more rigorously controlled study.

17. STATISTICAL VALIDATION PLAN

17.1 Sensor Accuracy Metrics

To move beyond the illustrative benchtop and synthetic-data results presented in Section 9, a more formal statistical validation plan is proposed for future work. For sensor accuracy, this would involve computing mean absolute error and limits of agreement (via Bland-Altman analysis) between the onboard pH sensor and a calibrated reference meter across a larger, more diverse set of buffer and saliva-mimicking solutions than used in the current benchtop tests.

17.2 AI Module Validation Metrics

For the agentic AI module, the proposed plan is to report sensitivity and specificity of persistent-pattern detection against a larger synthetic corpus with known ground-truth drift and noise labels, alongside a receiver operating characteristic (ROC) analysis to characterize the trade-off between false-alarm rate and detection sensitivity as the persistence threshold is varied. Once pilot human-subject data becomes available under appropriate ethical approval, the same metrics would be recomputed against real longitudinal readings, and participant self-report trust scores from the pilot study would be analyzed descriptively alongside the device's own outputs to assess whether algorithmic escalations align with participants' own sense of when something felt different.

18. CONCLUSION AND FUTURE WORK

18.1 Summary of Contributions

This paper has presented the design of an Agentic AI-Based Smart Toothbrush that embeds saliva pH sensing into a familiar daily habit, transmits readings through a standard IoT pipeline, and applies an autonomous, agentic reasoning layer to distinguish meaningful longitudinal health signals from transient noise. Beyond the core sensing and AI pipeline, this extended manuscript has also addressed the system's mobile interface design, its cost profile relative to existing connected oral-care products, the data-privacy and ethical commitments necessary for responsible deployment, and a concrete plan for pilot human-subject evaluation and statistical validation. The bench-level prototype demonstrates that the core sensing, connectivity, and data-pipeline components are technically feasible at low incremental hardware cost, and that an agentic approach to trend analysis offers a plausible improvement over static threshold-based alerting for this class of noisy, diet-influenced physiological data.

18.2 Future Directions

Several directions remain for future work. First, the sensor suite should be expanded to include salivary glucose and inflammatory-marker channels, building on the modular analog front-end already provisioned in the current hardware design, while accounting for the additional cost these channels introduce, as discussed in Section 12. Second, formal human-subject studies, conducted under appropriate institutional ethical review as outlined in Sections 14 and 16, are needed to validate both sensor accuracy against clinical saliva assays and the agentic AI module's real-world escalation performance, following the statistical validation plan proposed in Section 17. Third, the agent's risk-classification logic should be refined and validated in collaboration with dental and medical professionals to ensure its recommendations remain appropriately cautious and clearly non-diagnostic. Finally, longer-term deployment studies are needed to assess user adherence, the durability of the sensing hardware under repeated toothbrush use and cleaning, and the system's overall contribution to earlier detection of preventable oral and systemic health issues.

Taken together, this work illustrates how combining low-cost biosensing, embedded IoT connectivity, agentic AI reasoning, and deliberate attention to interface design, cost, privacy, and ethical deployment can transform an everyday object into a platform for continuous, preventive health monitoring, offering a scalable foundation for future biomarker-based screening research.

References

[1] Author, A., & Author, B. (Year). Salivary biomarkers as diagnostic indicators: A review of current evidence. Journal of Oral Diagnostics (placeholder citation — to be replaced with verified sources prior to submission).

[2] Author, C. (Year). Wearable electrochemical biosensors for continuous health monitoring. IEEE Sensors Journal (placeholder citation — to be replaced with verified sources prior to submission).

[3] Author, D., et al. (Year). IoT-enabled consumer health devices: Architecture and connectivity patterns. International Journal of IoT Applications (placeholder citation — to be replaced with verified sources prior to submission).

[4] Author, E., & Author, F. (Year). Agentic AI systems for autonomous health-trend reasoning. Journal of Artificial Intelligence in Medicine (placeholder citation — to be replaced with verified sources prior to submission).

[5] Author, G. (Year). Human-centered design principles for consumer health-monitoring applications. International Journal of Human-Computer Interaction (placeholder citation — to be replaced with verified sources prior to submission).

[6] Author, H., & Author, I. (Year). Data privacy and security frameworks for connected consumer health devices. Journal of Digital Health Policy (placeholder citation — to be replaced with verified sources prior to submission).

[7] Author, J. (Year). Cost modeling for early-stage biosensor prototype development. Journal of Medical Device Engineering (placeholder citation — to be replaced with verified sources prior to submission).

[8] Author, K., & Author, L. (Year). Salivary pH as an indicator of dental erosion risk: A clinical review. Journal of Dental Research (placeholder citation — to be replaced with verified sources prior to submission).

[9] Author, M., et al. (Year). Ion-selective electrode design for miniaturized biochemical sensing. Sensors and Actuators B: Chemical (placeholder citation — to be replaced with verified sources prior to submission).

[10] Author, N. (Year). Bluetooth Low Energy protocols for battery-constrained wearable devices. IEEE Transactions on Consumer Electronics (placeholder citation — to be replaced with verified sources prior to submission).

[11] Author, O., & Author, P. (Year). Personalized baselining approaches in remote patient monitoring systems. Journal of Biomedical Informatics (placeholder citation — to be replaced with verified sources prior to submission).

[12] Author, Q. (Year). Alert fatigue in consumer health applications: Causes and mitigation strategies. International Journal of Medical Informatics (placeholder citation — to be replaced with verified sources prior to submission).

[13] Author, R., et al. (Year). Regulatory frameworks for health data collected by consumer IoT devices. Journal of Health Law and Policy (placeholder citation — to be replaced with verified sources prior to submission).

[14] Author, S. (Year). Usability evaluation methods for mobile health dashboards. International Journal of Medical Informatics (placeholder citation — to be replaced with verified sources prior to submission).

[15] Author, T., & Author, U. (Year). Ethical design of non-diagnostic consumer screening tools. Journal of Biomedical Ethics (placeholder citation — to be replaced with verified sources prior to submission).

[16] Author, V. (Year). Statistical methods for agreement analysis between sensor and reference measurements. Statistics in Medicine (placeholder citation — to be replaced with verified sources prior to submission).

[17] Author, W., et al. (Year). Pilot study design for early-stage wearable and embedded health devices. Journal of Medical Internet Research (placeholder citation — to be replaced with verified sources prior to submission).

[18] Author, X. (Year). Enzymatic and immunoassay-based biosensors: A comparative review of cost and performance. Biosensors and Bioelectronics (placeholder citation — to be replaced with verified sources prior to submission).