

OcuGuard: A Privacy-Preserving Mobile Multimodal Framework for Real-Time Digital Eye Strain Mitigation

Nila Jayakumar¹, Vijay K. Madiseti²

¹Alpharetta High School, Alpharetta, USA

²College of Computing, Georgia Institute of Technology, Atlanta, USA

Email: nilarjk@gmail.com, vkm@gatech.edu

How to cite this paper: Jayakumar, N. and Madiseti, V.K. (2026) OcuGuard: A Privacy-Preserving Mobile Multimodal Framework for Real-Time Digital Eye Strain Mitigation. *Journal of Information Security*, 17, 498-513.

<https://doi.org/10.4236/jis.2026.174022>

Received: September 7, 2026

Accepted: October 17, 2026

Published: October 20, 2026

Copyright © 2026 by author(s) and

Scientific Research Publishing Inc.

This work is licensed under the Creative

Commons Attribution-NonCommercial

International License (CC BY-NC 4.0).

<http://creativecommons.org/licenses/by-nc/4.0/>



Open Access

Abstract

Prolonged use of digital displays, combined with widespread shifts to remote work and education, has coincided with a marked rise in myopia across children and adults. Existing screen-monitoring tools mainly provide retrospective, time-based logs and lack real-time physiological data or vision-specific diagnostics. This paper introduces OcuGuard, a privacy-centric, cross-platform mobile framework for iOS and Android that continuously monitors and mitigates digital eye strain using multimodal sensor fusion. With real-time computer vision processed on a secure backend, OcuGuard tracks gaze, blink frequency and dynamics, viewing distance, and ambient light. These data streams feed a machine learning regression model, calibrated with isotonic regression, to compute an individualized daily eye strain index. Users see this index through an interpretable *Eye Nutrition Label*, visually modeled on the format of the U.S. Food and Drug Administration (FDA) Nutrition Facts panel (not an FDA-endorsed or regulatory label). We also design and implement closed-loop, automated interventions—such as context-aware display brightness control, micro-blink reminders, and gamified eye exercises—while transmitting only short-lived, encrypted video frames to a backend for feature extraction (frames are discarded immediately after processing and never persisted) and keeping derived scores and long-term storage on the user's device by default, sharing only aggregated summary reports with an explicitly configured guardian account (Section 3.4) to protect privacy. The pipeline is architected to support alignment with the principles of the Children's Online Privacy Protection Act (COPPA) and the Health Insurance Portability and Accountability Act (HIPAA).

Keywords

Digital Eye Strain, Computer Vision, Multimodal Fusion, Myopia Risk

1. Introduction

1.1. Background and Motivation

By 2050, the American Academy of Ophthalmology projects that nearly 50% of the global population will meet clinical criteria for myopia [1], marking a major shift in ophthalmic public health. At the same time, daily use of electronic visual display devices has risen sharply, driven by remote and hybrid education and reduced time outdoors. CDC data show that from July 2021 to December 2023, adolescents averaged at least 4 hours of screen time per day [2]. The parallel rise in myopia and screen use suggests an association worth investigating and highlights the need for automated, mobile technologies that can continuously measure real-time visual load. It is important to note that this paper targets short-term, reversible asthenopic eye strain rather than myopia itself: myopia is a long-term structural change in refractive error that requires longitudinal clinical evidence and follow-up measures beyond the scope of the framework presented here. Elevated screen use and digital eye strain are treated as behaviorally related but mechanistically distinct outcomes. Throughout this paper, *digital eye strain* (used interchangeably with *eye strain* and *ocular strain*) denotes the short-term, screen-related visual and ocular discomfort that is the outcome of interest, while *visual load* denotes the distinct, sensor-derived input construct (e.g., blink patterns, screen distance, ambient light, screen time) hypothesized to contribute to that outcome; visual load is treated as a predictor, not a symptom.

1.2. Limitations of Existing Approaches

Current mobile operating systems provide restricted utility frameworks that lack real-time, domain-specific visual feedback. Most commercial tools only offer retrospective, low-resolution screen-time summaries (typically weekly), giving generic per-app usage data with little physiologic or biophysical meaning. System features focus mainly on temporal self-regulation, like time-management dashboards and user-defined limits. Common mobile health tools—such as Android’s ambient “20-20-20” reminder notifications—depend on voluntary compliance [3], cannot measure instantaneous ocular strain or accommodative load, and do not provide an immediate, continuous, clinically informative visual health dashboard.

1.3. Research Questions

To address these technical boundaries, this investigation formulates and motivates three core research questions, which the architecture presented in Section 3 is designed to address and which future empirical work (Section 5) is intended to evaluate:

- **RQ1:** To what quantitative extent does relative smartphone-to-ocular viewing distance modulate short-term asthenopic eye strain?
- **RQ2:** What operational data streams can be safely harvested from active mobile device ecosystems, and how can they be structurally ingested by real-time tracking models?
- **RQ3:** How can persistent, positive visual behaviors be continuously reinforced through localized real-time application outputs?

1.4. Summary of Contributions

To address the limits of open-loop system architectures, we present OcuGuard, a proactive multi-modal software framework for quantitatively assessing and mitigating digital eye strain. OcuGuard combines continuous client-side monitoring of user-interface interactions with a secure remote computational backend. The framework includes: 1) an adaptive scoring engine producing a fixed-range (1 to 999) score calibrated to each user's own history, 2) automated closed-loop structural interventions, 3) an interpretable diagnostic ledger modeled on the U.S. Food and Drug Administration (FDA) Nutrition Facts panel format, and 4) a privacy-oriented data processing pipeline—combining ephemeral, encrypted backend feature extraction with on-device storage of derived data by default (Section 3.5)—designed to support institutional user privacy and data protection requirements, with a road map toward on-device (edge) feature extraction, while the scoring model remains server-side, described in Section 3.1.

2. Related Work

2.1. Timer-Based Applications

Previous interventions have mainly targeted screen-viewing behavior using software-based time limits. *BlinkMate* is an environment-level behavioral intervention platform for media-streaming ecosystems, specifically integrated into Android TV architectures [4]. Optimized for pediatric users, it applies noninvasive display-dimming sequences or programmatic playback interruptions to implicitly trigger normal blink reflexes. Caregivers can configure pause frequency, intervention duration, and daily exposure limits through an authenticated management interface. However, this architecture operates as an open-loop control system and lacks computer vision modules to continuously sense, infer, and adapt to the user's real-time physiological and behavioral state.

2.2. Hardware-Dependent Tracking Systems

To quantify physiological variability directly, modern diagnostic paradigms use dedicated external sensor subsystems. *EyeStrainNet* is a representative multi-modal monitoring platform that uses long-wave infrared (LWIR) imaging to model the spatiotemporal evolution of ocular surface temperature [5]. It is optimized to characterize tear-film evaporation during prolonged, close-range interaction with digital displays. Although it uses deep learning—most notably a one-

dimensional convolutional neural network (1D-CNN) on temporally ordered thermal image sequences—the system still depends on tethered, externally mounted cameras, which severely limits its suitability for everyday consumer mobile use.

2.3. Desktop Computer Vision Frameworks

To remove the need for dedicated external hardware, software-only computer vision frameworks use built-in webcams on desktop systems. *Lumina AI* uses a 68-point facial landmark model with Eye Aspect Ratio (EAR) vectors to measure blink frequency on Windows and macOS [6]. Written in Python and packaged with Electron, it enforces the “20-20-20” rule by scheduling and displaying micro-exercise prompts [7]. Its advanced visualization dashboard and compliance logging are fully available only with an active localized Meibomian Gland Dysfunction (MGD) profile, and the framework is not optimized for mobile or other resource-constrained platforms.

2.4. Distributed Background Applications

Recent advances in mobile optimization frameworks increasingly use model pruning to enable efficient inference in background processing. The *EyeNet* architecture—specifically the EyeNetV3 variant—is a highly optimized convolutional neural network with about one million parameters, trained via Knowledge Distillation (KD) from a deeper ResNet-18 teacher network [8]. Deployed as a pruned TensorFlow Lite (.tflite) model, it attains 87.16% diagnostic accuracy while running as a background task on modern mobile operating systems, without noticeable impact on frame rate or CPU-induced display latency.

3. The OcuGuard Framework

3.1. System Architecture

The proposed OcuGuard framework comprises three concurrently executing runtime layers that jointly optimize computational efficiency and user privacy:

1) **Client Presentation Layer:** Built with React Native and Expo, this layer renders real-time composite strain metrics, delivers interactive animated ocular exercises, presents a daily Eye Nutrition Label modeled on the FDA Nutrition Facts panel format, and manages on-device visual alerts triggered by backend events.

2) **Backend Processing Layer:** A server-side environment running an optimized C++ build of the OpenFace computer vision engine [9] [10]. To reduce overhead and protect data, a single camera frame is sent from the mobile device over an encrypted WebSocket to this backend once per second, forming a continuous unidirectional stream. OpenFace extracts only numerical feature descriptors from volatile RAM and immediately discards the raw image buffer; frames are never written to persistent storage. The resulting numeric coordinates are returned to the device via the persistent WebSocket loop, enabling low-latency, high-precision inference. As a planned future direction, we intend to export the feature-extraction model via ONNX to CoreML (iOS) and TensorFlow Lite (An-

droid), moving this step fully on-device and eliminating network transmission of the real-time video stream entirely, while the heavier regression/scoring model described in Section 3.3 would remain server-side.

3) **Encrypted Storage Layer:** An on-device cryptographic storage subsystem that securely stores derived statistical features, mitigating risks of global telemetry and cross-device tracking.

3.2. Multimodal Sensor Data Fusion

To estimate users' visual load and infer resulting digital eye strain without specialized clinical tools, OcuGuard simultaneously records, timestamps, and co-registers five environmental and biometric data streams, then integrates them into a unified regression model. Each stream below is treated as a candidate input signal contributing to the visual-load estimate that feeds the eye-strain score; none of these signals is presented as an independently, clinically validated measure of digital eye strain itself, and their individual and joint predictive value is a target of future validation work (see Section 5):


- **Front-Facing Camera Core:** Feeds the OpenFace algorithm [9] [10] to detect blinks via Action Unit 45 (AU45) presence and intensity, OpenFace's Eye-Aspect-Ratio-based measure of eyelid closure (approximately equivalent to an EAR threshold of 0.20 sustained for two or more consecutive frames), distinguishing full from partial eyelid closures. It computes real-time gaze vectors to localize and quantify fixation on screen regions and fixation duration, flagging sustained, stable fixations as candidate markers of visual load. Screen-to-head distance is estimated from OpenFace's head-pose translation output (pose_Tz), which gives monocular depth in millimeters from the fitted 3D head-pose model across all 68 landmarks, rather than interpupillary-distance triangulation; distances below approximately 300 - 400 mm are flagged as too close. These distance and gaze signals, together with continuous engagement duration, are combined into a proposed *ocular-demand index*—an increasing composite in which shorter screen distance and more extreme gaze deviation contribute more heavily—reported as a normalized, unitless proxy rather than a measurement in physical units (see Section 5). Expected measurement error for these signals under common conditions such as eyeglasses, variable illumination, and head rotation has not yet been empirically characterized (see Section 5).
- **Ambient Light Sensing Array:** Uses the device's integrated light sensor to measure absolute ambient illuminance and capture high-contrast differences between display backlight output and low-luminance environments.
- **OS Brightness Percentage:** Continuously sampled from operating system settings. Combined with the time-of-day weighting described below and cumulative session duration, brightness feeds a proposed *blue-light exposure risk* proxy (defined in the next item); it does not, by itself, correspond to a measured spectral retinal dose.

- **Temporal Location Offsets:** Uses geolocation and time-of-day data to compute a time-weighting coefficient, *Time_Weight*, that is 1.0 during daylight hours and scales up to approximately 2.5 after 10:00 PM, reflecting reduced dark adaptation and pupil dilation at night. This is combined with OS brightness percentage and cumulative session duration into a working blue-light exposure risk proxy:

$$\text{Blue Light Risk} = \text{Brightness\%} \times \text{Time_Weight} \times \text{Session_Minutes}.$$

This formula is informed by published research on mesopic vision and blue-light sensitivity, but it is not a single, universally accepted clinical standard; we report it as a proxy metric pending validation against self-reported strain data (see Section 5), not as a clinically validated measurement.

- **Operating System Application Usage Statistics:** Uses Apple's Screen Time API and Android's UsageStatsManager to link real-time ocular metrics to specific apps, avoiding nonspecific, device-level aggregation. The resulting per-application strain ranking is shown in **Figure 1**.




Source: Screenshot by the authors.

Figure 1. Eye strain by application.

3.3. Eye Strain Score and Eye Nutrition Label

The model converts fused inputs into an index on a fixed 1 to 999 scale, calibrated to each user's own historical distribution (see below). Low values indicate optimal eye conditions (green), mid-range values indicate moderate digital eye strain (yellow), and values above 750 indicate severe digital eye strain requiring real-time system interventions (red). **Figure 2** shows the home screen displaying the current score alongside its underlying ocular metrics.


The score is computed in the backend using a machine learning regression model. Input features include blink frequency relative to a healthy baseline of 15 - 20 blinks/min, incomplete blink ratios over 60-second windows, gaze deviation from screen center, screen-to-head distance, OpenFace Action Unit intensities used as candidate indicators of eyelid fatigue [9] [10], and uninterrupted screen time. To account for individual physiological differences, the raw regression outputs undergo **Isotonic Calibration**, normalizing metrics to each user's historical distribution rather than fixed population templates [11] [12]. This calibration is intended so that a given percentile of a user's own strain distribution maps to a comparable position on the 1 - 999 scale across users. However, because no common labeled reference or cross-user validation currently exists (see Section 5), we do not claim that a threshold of 750 corresponds to an equivalent physiological stress state across users; we treat 750 as a provisional, per-user-relative alert threshold pending such validation.



Source: Screenshot by the authors.

Figure 2. Eye strain score/home.

At the end of each day, local session logs are aggregated to generate an *Eye Nutrition Label*, delivered as an in-app notification card. Modeled on the FDA Nutrition Facts panel to improve clinical interpretability and standardize visual presentation [13], the card summarizes peak and mean eye-strain score, per-application usage durations, cumulative blink deficits (magnitude and time below healthy rates), an estimated blue-light exposure risk score ($\text{Brightness\%} \times \text{time-of-day weighting} \times \text{session duration}$, as defined in Section 3.2—not a measured spectral retinal dose), and relative per-application strain severity rankings to flag high-risk software. Historical records are archived longitudinally to visualize weekly and monthly trends. An example label is shown in Figure 3.



Source: Screenshot by the authors.


Figure 3. Eye nutrition label.

3.4. Parental Controls and Closed-Loop Interventions


OcuGuard employs state-dependent intervention parameters to support behavioral modification and to manage pediatric screen exposure. Some of these parameters trace to established clinical guidance—the 15 - 20 blinks/min healthy baseline and the 20-minute interval underlying the 20-20-20 rule [7]—while others, such as the 750 composite-index threshold and the specific AU05/AU07 fatigue combination below, are provisional design choices pending the validation

discussed in Section 5:

- **Micro-Blink Alerts:** Activated when real-time blink frequency decreases below 10 blinks per minute for a continuous duration of 45 seconds. Under these conditions, the system renders an animated graphical overlay designed to re-engage the child user, which persists until OpenFace [9] [10] verifies a compliant, full corneal occlusion event.
- **Contextual Display Dimming:** Initiated when the composite eye-strain index exceeds the predefined high-risk threshold (750+). The mechanism differs by platform. On Android, the app sets `WindowManager.LayoutParams.screenBrightness` on its own window (via expo-brightness or a bare native module) to a low value (e.g., 0.01), dimming the screen to near-black; this is a standard, permission-free public API. On iOS, direct brightness control below the system-defined minimum is a private API and is not App Store-compliant, so the same visual effect is approximated with a full-screen, high-opacity `UIWindow` overlay (`UIWindowLevel.alert + 1`)—implemented in React Native/Expo as a full-screen `Modal` with a dark, semi-transparent background—combined with reducing brightness to the OS-permitted minimum via expo-brightness. On both platforms, this dimming is scoped to OcuGuard’s own foreground window: it does not persist if the user switches to another app. System-wide dimming that follows the user across apps would require a Device Admin profile (Android) or Apple’s Screen Time entitlement (iOS); this is a substantial additional architectural undertaking and is not part of the current design (see Section 5).
- **Animated Lottie Exercises:** Automatically triggered following 20 minutes of uninterrupted application foreground activity. The intervention presents an animated, interactive protocol that operationalizes the conventional “20-20-20” rule [7], enforces full blink targets, and guides deliberate accommodative and vergence shifts. Task adherence is continuously monitored by OpenFace [9] [10] to objectively confirm completion.
- **Ophthalmic Action Units (AUs):** Fatigue states are inferred by cross-referencing fine-grained facial musculature dynamics via OpenFace Action Units [9] [10]. A critical fatigue threshold is reached when the system detects concurrent elevation of AU05 (upper eyelid raising, indexing muscular strain) and AU07 (eyelid tightening, indexing involuntary squinting), which triggers a dedicated fatigue alert workflow.
- **Guardian Management Controls:** Caregivers provision supervisory devices using Apple’s Family Controls entitlement or Android’s Family Link core frameworks. The guardian interface supports remote configuration of application-specific visual-strain limits, device lockout policies, and the automated weekly dissemination of eye-health summary reports (PDF) via encrypted email channels, aggregating adherence metrics and temporal trajectories of daily score patterns. The exercise interface and guardian controls are shown in **Figure 4(a)** and **Figure 4(b)**, respectively.



(a)



(b)

Source: Screenshot by the authors.

Figure 4. (a) Eye exercises; (b) Parental controls.

3.5. Privacy and Compliance Architecture

Because the system processes pediatric facial data, privacy-protective design principles are built into its architecture. In the current production implementation, camera frames are transmitted from the client to OcuGuard's backend over an encrypted WebSocket connection (Section 3.1); frames are held only in volatile memory on the backend, are never written to persistent storage, and are irreversibly discarded immediately after facial landmarks are extracted, leaving only derived numerical parameters, which are stored on the user's device via the Encrypted Storage Layer. Camera frames themselves are never shared with third parties; the one exception to on-device-only storage of derived data is the aggregated guardian summary report described in Section 3.4, which is generated on-device and transmitted via encrypted email only to a caregiver the account holder has explicitly configured as a guardian. As a planned future direction, we intend to move real-time feature extraction fully on-device via an ONNX export path to CoreML (iOS) and TensorFlow Lite (Android), eliminating network transmission of the raw video stream entirely while keeping the heavier scoring/regression model server-side; we expect this migration to further strengthen the system's privacy posture, particularly for pediatric users. This data-minimization design is intended to support alignment with the Children's Online Privacy Protection Act (COPPA) child safety requirements [14]—including mechanisms for verified parental consent and user-initiated data deletion—and with Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule standards [15] for managing personal health tracking records. We note that regulatory compliance is a legal and organizational determination that depends on factors beyond system architecture (e.g., policies, audits, and operational practices), and is not established by architectural design alone.

4. Architectural Comparison

To determine the architectural placement of OcuGuard within the contemporary design landscape, the framework is compared against prevailing technical design paradigms along a set of qualitative architectural features. This comparison is feature-level and does not constitute an empirical evaluation with human participants, deployed devices, or measured outcomes; see Section 5 for a discussion of what remains to be empirically validated.

Table 1 criteria are defined as follows: **Continuous Video Eye Tracking** indicates real-time, frame-level eye/gaze tracking from a camera, as opposed to periodic screenshots or manual logging; **Live Metrics/Adherence Dashboard** indicates an in-app, continuously updating view of measured metrics or exercise adherence, as opposed to a static or retrospective-only report; **Embedded Parental Dashboard** indicates caregiver controls integrated directly into the same application rather than a separate product; **Active Intervention (Dimming/Pauses)** indicates the system autonomously changes device behavior (e.g., dimming, pausing playback) rather than only issuing a notification; **Mobile Environment Optimi-**

zation indicates the system is designed and packaged for iOS/Android rather than desktop or television platforms; **Edge AI Constraint/Compression** indicates the on-device model is optimized (e.g., pruned, quantized) for mobile hardware constraints; and **COPPA/HIPAA Alignment Profile** indicates that the system’s design choices (e.g., data locality, consent flows) are intended to support alignment with these regulations, not that formal regulatory compliance has been certified (see Section 5).

As shown in **Table 1**, existing system configurations have major functional limitations when compared architecturally along these dimensions. Purely timer-based approaches, such as BlinkMate, lack active computer vision and cannot reliably infer user interaction states. Hardware-dependent platforms like EyeStrainNet offer robust physiological signal measurements but are impractical to deploy because they require tethered external infrared sensors. Desktop-focused computer vision solutions, such as Lumina AI, achieve high tracking accuracy with Eye Aspect Ratio-based metrics but are limited to static desktop environments and are not optimized for mobile. In contrast, OcuGuard overcomes these constraints with a mobile-optimized, privacy-preserving, closed-loop framework that integrates real-time visual telemetry with immediate, context-aware behavioral interventions.

Table 1. Comprehensive architectural feature matrix and baseline comparison.

Architectural Feature	Timer-Based App (e.g., BlinkMate)	Hardware Tracking (e.g., EyeStrainNet)	Desktop CV App (e.g., Lumina AI)	OcuGuard (Our System)
Continuous Video Eye Tracking	No	Yes	Yes	Yes
Infrared/Thermal Hardware Req.	No	Yes	No	No
Live Metrics/Adherence Dashboard	No	No	Yes (MGD Only)	Yes
Embedded Parental Dashboard	Yes	No	No	Yes
Guided Interactive/Animated Exercises	No	No	Yes	Yes
Active Intervention (Dimming/Pauses)	Yes	No	No	Yes
Mobile Environment Optimization	No	No	No	Yes
Edge AI Constraint/Compression	No	No	No	No
COPPA/HIPAA Alignment Profile	No	No	No	Yes

5. Limitations

OcuGuard is presented in this paper as an architectural proposal. The design and its components have not yet undergone empirical evaluation with human participants or deployed devices, and no model described here has yet been trained or fitted on real data. We detail these limitations explicitly below, both to scope the paper’s current contribution accurately and to outline the validation work required before the claims in Section 3 can be treated as established rather than proposed.

Absence of empirical validation. The comparison in Section 4 evaluates Ocu-

Guard against existing systems along qualitative architectural features; it is not a user study, and no participant, device, or session data has yet been collected to answer RQ1 - RQ3 quantitatively. A planned next step is a pilot study with a defined sample size, data-collection protocol, and consent process, reporting numerical results for each research question.

Unvalidated scoring model. The regression model, the 1 - 999 scale, the isotonic-calibration procedure, and the 750 “severe strain” threshold described in Section 3.3 are proposed design targets rather than a model that has been trained, fitted, or validated on data. No regression target, training/validation/test split, or calibration dataset currently exists. Accordingly, the claim that a threshold of 750 corresponds to an equivalent physiological stress state across users is a design goal, not an empirically demonstrated result, and should be read as such until tested in future work.

Sensor signals as candidate proxies. Blink rate, incomplete blink ratio, gaze deviation, screen-to-head distance, and OpenFace Action Unit intensities are treated throughout this paper as candidate proxies for visual load and digital eye strain (Section 3.2). None of these signals has yet been independently validated against a clinically defined digital-eye-strain outcome, including under common real-world conditions such as eyeglasses, variable illumination, or head rotation, and their individual or combined predictive value remains to be established. The blue-light exposure risk and ocular-demand formulas described in Section 3.2 are similarly working proxy metrics with defined input signals but no fixed, clinically standardized formula; validating them against self-reported strain data is part of the research contribution we intend to pursue.

Feasibility and performance unknowns. This paper does not report empirical measurements of end-to-end latency, CPU/GPU or memory utilization, battery consumption, supported OS versions, or behavior under camera-permission or background-execution restrictions. These are particularly relevant given that current mobile operating systems place substantial limits on background camera access and on third-party control of system-level display parameters (Section 3.4); whether the contextual-dimming and continuous-monitoring mechanisms described here can operate reliably under these constraints has not yet been demonstrated. In particular, the contextual display dimming described in Section 3.4 is currently scoped to OcuGuard’s own foreground window on both platforms and does not persist if the user switches to another app; achieving system-wide dimming would require a Device Admin profile (Android) or Screen Time entitlement (iOS), which is a substantial additional undertaking outside the current design.

Network dependency and biometric data transmission. The current production architecture requires an active network connection to a backend server for real-time feature extraction (Section 3.1), introducing latency and connectivity dependence not present in a fully on-device design, and meaning that raw camera frames leave the user’s device—encrypted in transit and never per-

sisted—rather than remaining strictly local. We view eliminating this network hop, via the on-device ONNX/CoreML/TensorFlow Lite migration described in Section 3.1 and Section 3.5, as a priority for future work, particularly given the sensitivity of pediatric biometric data; no empirical results for this on-device variant exist yet.

Reproducibility. Because no model has yet been trained, we do not report hyperparameters, training procedures, or dataset composition. We intend to include these, together with the pilot evaluation described above, in a future revision once a trained model exists.

Scope of health claims. OcuGuard is intended as a wellness and behavioral-feedback tool rather than a diagnostic medical device; the Eye Nutrition Label and eye-strain score are not clinical diagnoses of any ophthalmic condition. The framework targets short-term, reversible digital eye strain; any relationship between its use and longer-term outcomes such as myopia progression would require separate, longitudinal clinical study and is outside the scope of the present work.

Statements regarding alignment with COPPA and HIPAA describe architectural design choices intended to support such alignment; they do not constitute a legal certification of compliance, which depends on organizational and operational factors beyond system architecture.

6. Conclusion

We introduce OcuGuard, a privacy-conscious, multimodal mobile intervention framework aimed at addressing digital eye strain, a short-term behavioral factor of interest amid the worldwide rise in myopia. The system features a multi-sensor fusion layer embedded in a React Native, cross-platform mobile client, enabling real-time extraction of physiological and behavioral indicators via an encrypted backend processing pipeline, with derived indicators stored on-device by default and shared only as aggregated guardian summary reports where explicitly configured (Section 3.4). OcuGuard combines active computer vision outputs, ambient light data, and application usage logs within an optimized scoring engine. This synthesized user profile is translated into a practical *Eye Nutrition Label*, modeled on the visual format of the FDA Nutrition Facts panel, and paired with automated, closed-loop interventions, including context-aware screen dimming and guided, animated visual exercises. Its current architecture processes raw camera frames on an encrypted backend with strict ephemeral-data handling—no persistence, no third-party sharing—while storing derived scores on-device by default; a planned migration to on-device feature extraction via ONNX export to CoreML and TensorFlow Lite would eliminate network transmission of the video stream entirely, while the heavier scoring model would remain server-side. This design supports design-level alignment with COPPA guidelines and HIPAA Privacy Rule standards, while maintaining practicality for deployment at large scale across consumer mobile devices.

Declaration

CCS Concepts: Computing methodologies → Computer vision; Human-centered computing → Mobile systems and services.

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

- [1] Ang, M. (2024) Sharp Rise in Myopia around the World. American Academy of Ophthalmology (AAO) Young Ophthalmologists Committee.
<https://www.aao.org/young-ophthalmologists/yo-info/article/sharp-rise-in-myopia-around-world>
- [2] Centers for Disease Control and Prevention (CDC) (2024) Daily Screen Time among Teenagers: United States, July 2021-December 2023.
<https://www.cdc.gov/nchs/products/databriefs/db513.htm>
- [3] Eye Center of Northern Colorado (2023) 7 Ways to Protect Your Eyes from Screen Time. Patient Education Portal.
<https://www.eyecenternoco.com/blog/7-ways-to-protect-your-eyes-from-screen-time/>
- [4] Dubey, T., Dubey, A., Kothari, D. and Gupta, M. (2025) Blink Mate: A Smart Screen Control Application to Mitigate Digital Eye Strain in Children. *International Journal on Science and Technology*, **16**, 1-10. <https://doi.org/10.71097/ijst.v16.i2.6398>
- [5] EyeStrainNet Research Group (2025) Evaluation of Ocular Surface Microclimate Dynamics and Thermal Responses to Screen Displacement Tasks. *Journal of Thermal Biology*, **133**, Article ID: 104280.
<https://www.sciencedirect.com/science/article/abs/pii/S0306456525002372>
- [6] Dhankar, U., Kumar, A., Pradhan, N. and Gupta, N. (2025) Lumina AI: Real-Time Eye Wellness Monitoring Using Advanced Computer Vision and Machine Learning. *International Journal for Research in Applied Science and Engineering Technology*, **13**, 1636-1641. <https://doi.org/10.22214/ijraset.2025.74809>
- [7] American Optometric Association (AOA) (2023) Protecting Your Vision with the 20-20-20 Rule. Patient Health and Wellness Guidelines.
<https://www.aoa.org/AOA/Images/Patients/Eye%20Conditions/20-20-20-rule.pdf>
- [8] Quang Thao, L., Duc Cuong, D., Manh Hung, V., Thanh Vinh, L., Trong Nghia, D., Ha Hai, D., *et al.* (2023) Eye Strain Detection during Online Learning. *Intelligent Automation & Soft Computing*, **35**, 3517-3530.
<https://doi.org/10.32604/iasc.2023.031026>
- [9] Baltrusaitis, T., Robinson, P. and Morency, L. (2016) OpenFace: An Open Source Facial Behavior Analysis Toolkit. 2016 *IEEE Winter Conference on Applications of Computer Vision (WACV)*, Lake Placid, 7-10 March 2016, 1-10.
<https://doi.org/10.1109/wacv.2016.7477553>
- [10] Baltrusaitis, T., Zadeh, A., Lim, Y.C. and Morency, L. (2018) OpenFace 2.0: Facial Behavior Analysis Toolkit. 2018 *13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, Xi'an, 15-19 May 2018, 59-66.
<https://doi.org/10.1109/fg.2018.00019>
- [11] Barlow, R.E., Bartholomew, D.J., Bremner, J.M. and Brunk, H.D. (1972) Statistical Inference under Order Restrictions: The Theory and Application of Isotonic Regres-

sion. John Wiley & Sons.

- [12] Zadrozny, B. and Elkan, C. (2002) Transforming Classifier Scores into Accurate Multiclass Probability Estimates. *Proceedings of the 8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Edmonton, 23-26 July 2002, 694-699. <https://doi.org/10.1145/775047.775151>
- [13] U.S. Food and Drug Administration (2020) How to Understand and Use the Nutrition Facts Label. U.S. Department of Health and Human Services. <https://www.fda.gov/food/new-nutrition-facts-label/how-understand-and-use-nutrition-facts-label>
- [14] Federal Trade Commission (2013) Children's Online Privacy Protection Rule, 16 C.F.R. Part 312. *Federal Register*, **78**, 3972-4014.
- [15] Department of Health and Human Services (2002) Standards for Privacy of Individually Identifiable Health Information (HIPAA Privacy Rule), 45 C.F.R. Parts 160 and 164. *Federal Register*, **67**, 53182-53273.