Toward Always-On PSOG: Wearable Gaze Tracking System via a PPG-Derived Optical Analog Front-End

Luca Francesco Raduzzi
Politecnico di Milano
Milano, Italy
lucafrancesco.raduzzi@polimi.it

Carlo Pezzoli
Politecnico di Milano
Milano, Italy
carlo.pezzoli@polimi.it

Marco Paracchini
Politecnico di Milano
Milano, Italy
marcobrando.paracchini@polimi.it

Alessandro Perricone
Politecnico di Milano
Milano, Italy
alessandro.perricone@polimi.it

Daniele Bani
EssilorLuxottica
Milano, Italy
daniele.bani@luxottica.com

Luca Merigo
EssilorLuxottica
Milano, Italy
luca.merigo@luxottica.com

Filippo Melloni
EssilorLuxottica
Milano, Italy
filippo.melloni@luxottica.com

Marco Marcon
Politecnico di Milano
Milano, Italy
marco.marcon@polimi.it

Abstract

This work presents a low-power wearable Photosensor Oculography (PSOG) gaze-tracking system for smart glasses, repurposing the MAX86181 optical analog front-end, originally designed for photoplethysmography (PPG) to measure periocular reflectance patterns. Infrared LEDs and photodiodes are arranged around the eyes, using time-multiplexed illumination to generate 32-element feature vectors (4 PDs × 4 LEDs × 2 eyes). The system's integrated analog/digital ambient-light cancellation preserves signal integrity across indoor to outdoor conditions, as validated on a humanoid robot platform. A lightweight Mahalanobis distance-based classifier achieves ~97% gaze-region classification accuracy for both 5-zone and 9-zone partitions in the controlled robotic evaluation at 58.6 μs integration time. Power consumption ranges from 11 mW at 50 fps to 18.5 mW at 150 fps, corresponding to an estimated operating time from approximately 34 hours to 20 hours, respectively, on a 100 mAh, 3.7 V LiPo battery under ideal conversion efficiency. These results demonstrate the feasibility of repurposing a PPG-derived optical front-end for low-power PSOG in smart-glasses form factors. While the present validation is limited to controlled experiments on a humanoid robot and coarse gaze-region classification, the approach represents a promising step toward always-on wearable eye tracking.

CCS Concepts

• **Hardware** → **Functional verification**; • **Human-centered computing** → *Ubiquitous and mobile computing systems and tools*; • **Computer systems organization** → Embedded hardware; Embedded software.

ETRA '26, Marrakesh, Morocco

ACM ISBN 979-8-4007-2519-7/26/06
https://doi.org/10.1145/3797246.3806215

Keywords

Eye tracking, Photosensor Oculography, Embedded systems, Hardware, Gaze estimation, Classification

ACM Reference Format:
Luca Francesco Raduzzi, Carlo Pezzoli, Marco Paracchini, Alessandro Perricone, Daniele Bani, Luca Merigo, Filippo Melloni, and Marco Marcon. 2026. Toward Always-On PSOG: Wearable Gaze Tracking System via a PPG-Derived Optical Analog Front-End. In *2026 Symposium on Eye Tracking Research and Applications (ETRA '26), June 01–04, 2026, Marrakesh, Morocco.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3797246.3806215

1 Introduction and State of the Art

Eye tracking measures gaze direction and eye-movement dynamics, providing a direct window into visual attention, perception, and visuomotor control. By quantifying where and for how long a person fixates within a scene, it has become a core tool in psychology and neuroscience, as well as in human-computer interaction, marketing research, and medical diagnostics [Duchowski 2017; Mele and Federici 2012; Wade and Tatler 2005]. In the last decade, the push toward pervasive sensing (AR/VR headsets and smart glasses) has shifted the emphasis from laboratory accuracy alone to long-term wearability, low power consumption, mechanical integration, and robustness to uncontrolled lighting and motion.

Most commercial and research-grade systems still rely on video-based oculography (VOG), in which one or more cameras image the eye and gaze is inferred from features such as the pupil center and corneal reflections. Classical high-precision solutions exploit Purkinje images [Cornsweet and Crane 1973; Crane and Steele 1985], while many wearable implementations adopt pupil-corneal-reflection (PCCR) geometries coupled with calibration models. More recently, appearance-based methods infer gaze directly from eye images using machine learning, improving robustness under challenging imaging conditions [Larrazabal et al. 2019; Zhang et al. 2015]. Despite their maturity and accuracy, camera-based pipelines require continuous image acquisition, illumination control, and non-trivial processing, and can be sensitive to environmental illumination, occlusions, and eye-camera slippage, factors that translate into increased power draw and tighter opto-mechanical constraints in lightweight eyewear.

This has motivated growing interest in camera-free sensing modalities. Electro-oculography (EOG) can offer low-power eye-movement signals but is affected by drift and user-dependent electrode interfaces [Estrany and Fuster-Parra 2022], while scleral search coils achieve excellent precision at the cost of invasiveness and limited practicality for everyday use [Sprenger et al. 2008]. Event-based cameras and other emerging sensors have demonstrated very high-speed gaze tracking, yet they still rely on dedicated optics and relatively complex processing pipelines [Angelopoulos et al. 2021]. In parallel, non-imaging optical approaches have re-emerged as a promising compromise between robustness, integration, and energy efficiency.

Among these, Photosensor Oculography (PSOG) is a compact non-imaging technique that uses infrared (IR) emitters and photodiodes to measure reflected light from the periocular region [Kumar and Krol 1992; Van Den Berg and Spekreijse 1997]. As the eye rotates, the relative contribution of sclera, iris, and pupil within each emitter-detector geometry changes, producing systematic reflectance modulations that can be mapped to gaze with minimal computation. PSOG can achieve high sampling rates, low latency, and very low power consumption, making it attractive for always-on wearable platforms. At the same time, PSOG performance is strongly influenced by sensor placement, inter-subject anatomy, frame slippage, and nonlinear reflectance-to-gaze relationships, and it must contend with ambient light interference and motion artifacts, especially outdoors [Palmero et al. 2023; Pezzoli et al. 2026].

Although examples of prototypes embedding a PSOG tracker in a pair of smart glasses were already attempted by [Pezzoli et al. 2025] and [Pettenella et al. 2025], a central barrier to widespread PSOG adoption in consumer eyewear are not only the gaze-inference model, but the acquisition chain: achieving a high signal-to-noise ratio (SNR) while maintaining a compact footprint and a tight power budget under intense and time-varying ambient illumination. Interestingly, PSOG shares key hardware requirements with photoplethysmography (PPG): both measure small AC variations riding on a dominant DC component, typically within a bandwidth below ∼100 Hz, and both are susceptible to displacement artifacts and ambient light pickup. This observation opens an alternative design space: repurposing mature, low-power optical AFEs originally developed for PPG as the front-end for reflectance-based eye tracking.

In this work, we explore this direction by leveraging the MAX86181 optical analog front-end (AFE) as the core of a wearable PSOG acquisition platform. The MAX86181 integrates multi-channel, high-resolution optical readout with programmable LED driving and on-chip ambient light cancellation, offering an unusually compact and low-power solution for embedded reflectance sensing [Analog Devices 2026]. Building on this capability, we implement a time-multiplexed illumination strategy in which individual LEDs are pulsed sequentially while multiple photodiodes are sampled in parallel, yielding a richer set of spatial reflectance measurements per time step without resorting to imaging optics. In addition, we exploit the device's layered ambient-rejection mechanisms, including analog and digital cancellation to preserve headroom and mitigate saturation under harsh illumination. Finally, we evaluate the resulting system in controlled and repeatable conditions, including standard indoor lighting and stress-test illumination representative of outdoor environments, to assess whether a PPG-derived AFE can enable robust, low-power PSOG suitable for smart-glasses integration.

2 Optical Signal Characteristics for Eye Motion Detection

Although photoplethysmography (PPG) and optical eye tracking target fundamentally different physiological phenomena, both rely on controlled optical emission and photodetection of light reflected by biological tissues. PPG is a technique used to measure variations in blood volume within microvascular tissue: it operates by illuminating tissue with one or more LEDs and measuring the intensity of reflected (or transmitted) light with a photodiode. The resulting signal encodes the periodic absorption changes driven by the cardiac cycle, and its informative content is confined to a relatively narrow and low-frequency band, typically below a few tens of Hz. In contrast, optical eye tracking based on reflectance does not measure volumetric absorption changes but rather the geometric modulation of reflected light due to ocular rotation. When the eye rotates, the relative position of sclera, iris, and pupil changes with respect to a fixed emitter–detector configuration, producing a variation in the intensity and distribution of the reflected optical signal that is directly coupled to gaze angle. Despite these differences in physiological origin, the two modalities share a number of fundamental hardware requirements. In both cases, the detected signal consists of a dominant DC component with a superimposed AC variation that carries the information of interest: in fact the relevant frequency content of this AC component typically lies below approximately 100 Hz. Furthermore, both modalities are susceptible to similar sources of disturbance which are mainly mechanical displacement artifacts and ambient light interference: those phenomena must be adequately mitigated at both the hardware and signal-processing levels to keep the acquisition chain at a good level of sensing in terms of SNR. These shared characteristics suggested that integrated circuits developed for PPG acquisition could represent viable candidates for optical eye tracking as well. This observation significantly broadened the scope of the component search: the market for PPG-dedicated integrated solutions is considerably larger and more mature than that for eye tracking-specific hardware, offering a wider range of devices optimized for low power consumption, miniaturization, and high-sensitivity optical front-ends. Among the solutions surveyed, the MAX86181 emerged as particularly suitable, combining a programmable multi-LED driver, a high-resolution photodetector acquisition chain, and an exceptionally compact and low-power profile, properties that aligned well with the constraints imposed by a flexible, wearable implementation.

3 System Design and Integration

The prototype developed consists of a wearable smart-glasses platform integrating the MAX86181 optical analog front-end (AFE) together with multiple infrared LEDs and photodiodes positioned around the periocular region. The MAX86181 is a quad-channel, high-resolution optical acquisition system featuring two programmable LED drivers and four independent 20-bit current-mode ADCs: although originally designed for photoplethysmography (PPG), its architecture makes it suitable for reflectance-based eye motion sensing.

The distribution of LEDs and photodiodes around the periocular region was determined through preliminary optical simulations and analysis of prior reflectance-based eye tracking configurations, with the objective of maximizing sensitivity to sclera–iris reflectance transitions during ocular rotation, while minimizing direct LED-to-photodiode coupling and inter-channel optical crosstalk.

The physical quantity of interest is the surface reflectance ρ, defined as the ratio between reflected and incident optical flux: this factor, combined with the detector responsivity, yields the following.

$$\rho = \frac{\Phi_r}{\Phi_i} \qquad I_{ph} = \mathcal{R}\,\Phi_r \tag{1}$$

where Φ_i is the incident optical flux emitted by the LED, Φ_r is the reflected flux collected by the photodiode, and $\mathcal{R}$ is the photodiode responsivity (A/W). Substituting the reflectance definition gives:

$$I_{ph} = \mathcal{R}\,\rho\,\Phi_i \qquad I_{ph} \propto \rho \tag{2}$$

As the eye rotates, ρ varies with gaze angle due to the different near-infrared optical properties of sclera, iris, and pupil: this angular modulation of ρ constitutes the informative signal extracted by the system.

In practice, the photocurrent I_{ph} is further influenced by emitter–detector geometry, tissue absorption, and angular emission/collection profiles: due to such many factors its magnitude is several orders of magnitude smaller than the LED drive current, imposing strict requirements on front-end sensitivity and noise performance.

To achieve a measurable signal without violating optical safety limits, the LED drive current was set to 9 mA in compliance with IEC 62471, moreover preliminary tests indicated that the expected photocurrent at the detector input is on the order of a few nanoamperes. With the MAX86181 configured for a 8 μA full-scale input range, the useful signal occupies only a small fraction of the converter dynamic range : for an N-bit ADC with full-scale current I_{FS}, the least significant bit (LSB) and the corresponding quantization noise power are given by:

$$\mathrm{LSB} = \frac{I_{FS}}{2^N} \qquad \sigma_q^2 = \frac{\mathrm{LSB}^2}{12} \tag{3}$$

and the RMS quantization noise is:

$$\sigma_q = \frac{\mathrm{LSB}}{\sqrt{12}} \qquad \sigma_q \ll I_{ph} \tag{4}$$

Under the selected configuration ($I_{FS} = 8\,\mu$A, $N = 19$), the LSB is on the order of a few picoamperes, making quantization noise negligible compared to shot noise and ambient light fluctuations.

To maximize the collected signal while maintaining energy efficiency, the ADC integration time was set to 58.6 μs, since LED activation is confined to the integration window, illumination and signal acquisition are temporally overlapped and so the energy dissipated per measurement cycle can be approximated as:

$$E_{meas} = V_{LED}\,I_{LED}\,T_{int} \qquad P_{avg} = E_{meas}\,f_{frame} \tag{5}$$

where T_{int} is the integration time and f_{frame} is the frame rate. The integration time therefore represents a primary design parameter for trading signal-to-noise ratio against power consumption.

3.1 Illumination Scheme and Signal Conditioning

The acquisition strategy adopted is based on time-multiplexed illumination: at each measurement cycle, a single LED is activated while all four photodiodes are sampled simultaneously through the device's independent acquisition chains, after which the next LED in the sequence is switched on and the process repeats. This approach allows the eye to be illuminated from multiple spatial positions in rapid succession, extracting richer geometrical information about its orientation than would be possible with a single fixed emitter.

Once the illumination scheme had been defined and simulated, preliminary tests were performed to better assess the expected signal magnitude and the possible impact of ambient light disturbances involving the same LED and photodiode models that were selected for the final setup. LEDs and photodiodes with peak spectral response in the near-infrared region were selected, where the spectral irradiance of sunlight is naturally attenuated relative to the visible range. This choice provides an intrinsic degree of ambient-light rejection, which is particularly valuable in outdoor operating conditions, where broadband illumination would otherwise degrade signal quality or drive the acquisition chain into saturation. A series of ambient-light acquisitions was carried out under different lighting conditions by using a luxmeter to measure the ambient light intensity; the photocurrent generated by the photodiode was recorded, with the device oriented approximately normal to the sun: the resulting data are reported in Figure 1. In parallel, additional experiments were conducted to estimate the expected photogenerated signal by evaluating the coupling efficiency between the light emitted by the LEDs and the current produced by the photodiodes placed in front of them. The data shown in Figure 2 indicate that, for a 9 mA LED driving current, chosen to comply with the eye-safety constraints IEC62471, it is on the order of tens of microamperes, in an environment where full-sunlight conditions can generate more than 100 μA of baseline current. Starting from these preliminary data, the system architecture was based on an STM Nucleo board, used to control and interface with the MAX86181 which was positioned centrally between the two lenses in order to minimize the analog routing length between the photodiodes and the AFE inputs along the flexible-PCB traces; this layout reduced parasitic capacitance and electromagnetic interference, yielding a measurable improvement in signal integrity.

A built-in photocurrent background cancellation function was integrated to accommodate DC levels both analogically and digitally: at the analog level, this stage is capable of cancelling up to 200 μA of DC photocurrent, providing the main contribution to ambient-light immunity. At the digital level, each illumination event (exposed sample) was bracketed by dark samples acquired with the LED turned off, and on-board subtraction was applied to remove

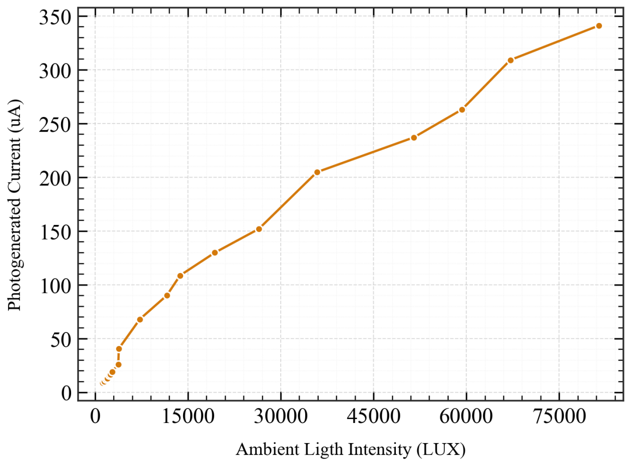

Figure 1: Photogenerated Current in high ambient light conditions

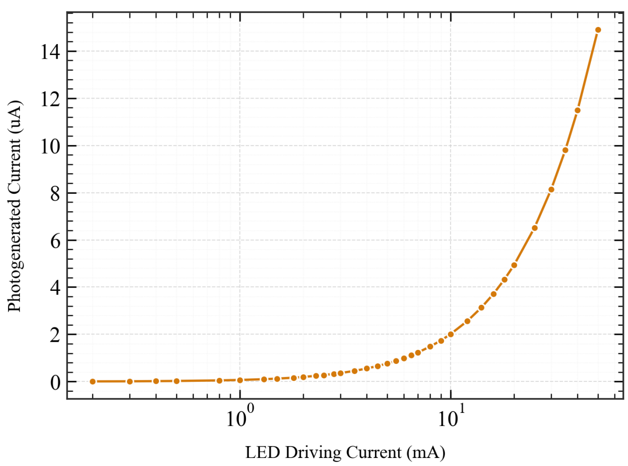

Figure 2: Photogenerated current with infrared LED emitted light

residual ambient contributions persisting after the analog stage. In addition, real-time averaging of repeated exposed samples was supported prior to writing the result to the internal FIFO, thereby lowering the effective noise floor of each output sample without increasing the data rate transferred to the host microcontroller. Together, these mechanisms give the system a robust and layered rejection of both stochastic noise and structured ambient interference, which is essential for reliable operation in an uncontrolled wearable environment.

3.2 Region Classifier

Once the raw data were acquired and digitised, they were transmitted via UART to a host PC: there, the contribution of each LED was extracted from the raw photodiode signals, yielding a feature vector of 32 elements (4 photodiodes × 4 LEDs × 2 eyes), on which the classification stage was executed. The purpose of this stage was to determine the region where the gaze direction was pointing by analysing the pattern of signals generated by the different LED–photodiode combinations. The algorithm used was a lightweight distance-based classifier based on the Mahalanobis distance[Duda et al. 2000], consistent with our previous work. The main advantage of this approach is that it accounts not only for the distance between feature vectors but also for the statistical structure of the data, incorporating both the variance of each feature and the correlations between them [Pettenella et al. 2025]. During the training phase, the classifier estimates the mean feature vector and the covariance matrix associated with each class. The covariance matrix captures the statistical dispersion of the signals and their mutual correlations, while its inverse is later used to weight the contribution of each feature during classification. During inference, each newly acquired feature vector is compared against the statistical model of every class. The algorithm evaluates the squared Mahalanobis distance between the test sample and each class distribution, and the sample is assigned to the class that minimises this distance.

The squared Mahalanobis distance is defined as:

$$d_c^2 = (x - \bar{x_c})^T \Sigma_c^{-1} (x - \bar{x_c}) \tag{6}$$

where d_c is the distance between the test point x and the class c, $\bar{x}_c$ is the mean of the training features for class c, and Σ_c^{-1} is the inverted covariance matrix for the class c.

4 Experimental Acquisition and Model verification

Once the prototype was assembled, a series of hardware verification and validation tests were conducted to verify the correct functioning of the acquisition chain; the experiments were designed to assess three core requirements that the platform must satisfy: first, as a baseline, the system must be capable of reliably detecting eye position under standard indoor lighting conditions. Second, it must maintain the same capability when exposed to demanding outdoor illumination, where ambient light represents a significant source of disturbance. Third, both tasks must be accomplished within a power consumption budget compatible with a wearable, battery-powered device intended for extended use. To ensure repeatability, all tests were conducted on the Ami Desktop humanoid robot, whose eyes were programmed to execute a controlled sequence of movements covering a predefined set of gaze positions. The field of view was partitioned in two different ways, first into 5 zones and then into 9, in order to evaluate two classifiers of different spatial resolution and assess how classification accuracy scales with the number of target regions. All experiments were performed under two distinct illumination conditions: the indoor scenario was characterized by an ambient light level of approximately 1000 lux, representative of a typical well-lit indoor environment. The *stress test* was instead conducted by emulating harsh outdoor illumination through the use of halogen lamps capable of producing up to 100,000 lux, pushing the system's ambient light rejection to its operational limits.

As shown in Figure 3, the acquired signals exhibit a characteristic step-like pattern that directly reflects the discrete gaze positions assumed by the eye during the experiment. Each trace corresponds to

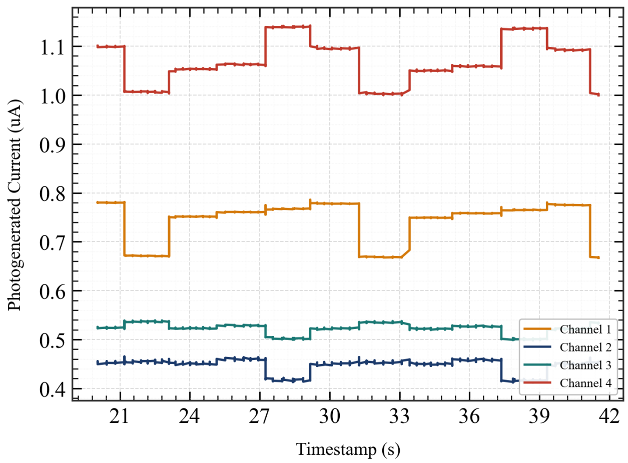

Figure 3: Photogenerated Current signal

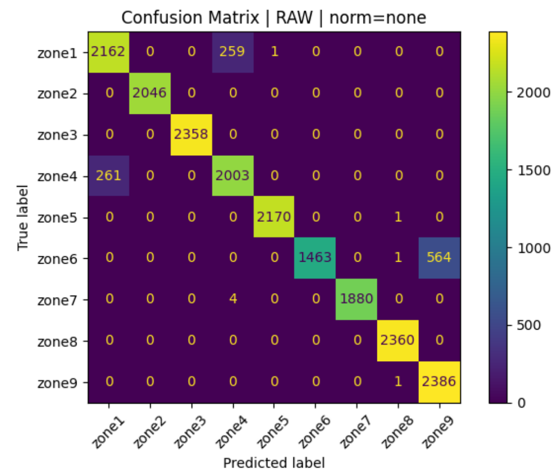

Figure 4: Confusion matrix for the 9-zone classifier under high-illumination conditions.

a single photodiode channel, and the contribution of each LED can be distinguished as it is activated in sequence. The step transitions in the signals correspond to the eye moving from one predefined zone to the next, with the five-zone partitioning producing step patterns of different granularity. During the experiments illustrated in Figure 3, the ambient light level was gradually increased, starting from a dark condition and reaching a fully illuminated condition toward the end of the acquisition. A small DC fluctuation is observable due to imperfect DC-level cancellation; however, by referring to Figure 1, one would expect a much larger photocurrent that could drive the system into saturation. In practice, the setup continues to operate correctly and acquires the signal reliably despite the increasing ambient illumination, demonstrating its robustness to ambient-light variations. In the setup, the field of view was limited between -30 and +30 degrees horizontally and -20 and +20 vertically, then subdivided into 9-zone or 5-zone depending on the granularity we aimed to achieve. The robot was programmed to fixate each point for 2 seconds, spaced by 5 degrees on the horizontal and 4 degrees on the vertical, that covered the whole field of view available. A uniformly distributed random subsample comprising 20% of the observation points was designated as the test set and excluded from the training data prior to model fitting. This procedure ensured that distinct instances were used during the training and evaluation phases of the classifier, thereby mitigating the risk of overfitting. Lastly, power-consumption tests were performed with the MAX86181 operated in two conditions: deep-sleep mode, which represents the lowest-power state with the internal oscillator disabled, and the measurement state, during which the LEDs are turned on and data are acquired and streamed to the microcontroller. In the measurement state, tests were repeated while varying the frame rate which is strictly related to the LED on-time and therefore expected to be a key contributor to power consumption; moreover measurements were conducted both in darkness and under ambient illumination to verify how the ambient-light cancellation scheme affects the overall power usage and classification accuracy.

5 Results

The acquisitions described in section 4 were analyzed for both the 9-zone and 5-zone classifiers: with an integration time of 58.6 μs and on-board averaging over four samples, and with both analog and digital ambient-light cancellation enabled, the algorithm achieved an accuracy of approximately 97 % in both dark and illuminated conditions. Figure 4 reports the confusion matrix for the 9-zone classifier under the final operating configuration. Most predictions lie on the main diagonal, confirming that classification errors are limited and mainly occur between neighboring gaze regions.

Power-consumption results are summarized in Table 1: among the parameters investigated, the frame rate emerged as the primary factor affecting power usage since it is directly correlated to the LED's on time duration and to the numbers of data present on the data bus. In deep-sleep mode, the device consumes only 7.16 mW, while enabling ambient-light cancellation in a worst-case condition of 100 klux increases the power by a fixed amount of about 2 mW. The reported power values were obtained by measuring the voltage drop across a shunt resistor placed on the IC supply line. Therefore, the reported power refers to the MAX86181 during LED driving, photodiode acquisition, on-chip processing, and I2C communication with the microcontroller: it does not include the power consumed by the STM Nucleo development board, the host PC, or battery-conversion losses. A frame is defined as a complete acquisition cycle composed of eight sequential LED–photodiode measurements, which can be regarded as a single snapshot of the periocular reflectance pattern. The number of frames acquired per second therefore has a direct impact on the overall energy budget. The required frame rate depends on the target application: high-speed saccadic eye movements may require sampling rates of around 150 frames per second, whereas tracking slower voluntary gaze shifts can be reliably achieved with frame rates in the 50–100 frames-per-second range, enabling a substantially more favorable power profile, consuming $\sim\frac{220\mu J}{\text{frame}}$

Table 1: Power consumption as a function of frame rate

Power Consumed	Frame Rate
11 mW	50 fps
14.9 mW	100 fps
18.5 mW	150 fps

6 Conclusions

The work presented in this paper demonstrates that repurposing a highly integrated PPG analog front-end for reflectance-based eye tracking is not only technically viable but represents a genuinely attractive direction for the development of next-generation wearable gaze sensing systems. The MAX86181, through its combination of high-resolution optical acquisition, on-board ambient light cancellation, and flexible configuration of power-critical parameters, proved capable of supporting reliable coarse gaze-region discrimination across indoor and emulated harsh outdoor illumination conditions in a controlled robotic setup. From a power consumption standpoint, the architecture already exhibits a favorable profile, and the margin for further optimization remains significant. Features such as dynamic drive current adaptation and frame rate scaling, which were not fully exploited in the current prototype, offer concrete levers to reduce energy dissipation in real operating scenarios, pushing the system closer to the autonomy requirements of a practical all-day wearable device. Looking ahead, perhaps the most compelling perspective opened by this work is the possibility of embedding the classification and gaze estimation logic directly on-board, eliminating the dependence on an external host for signal processing and moving toward a fully self-contained eye tracking module. Combined with the inherent advantages of flexible PCB integration in terms of form factor and wearability, this trajectory points to a system that could help bridge the gap between laboratory-grade eye tracking performance and the constraints of everyday wearable use, a challenge that remains largely open in the field, and one that this architecture appears well positioned to address. The present evaluation was designed as a controlled hardware-verification study rather than a human-subject validation: the use of a robotic platform ensured repeatable eye positions and illumination conditions, but it does not capture human-specific variability such as blinking, eyelid and eyelash occlusion, anatomical differences, sensor slippage, or motion artifacts. Future work will extend the evaluation to human participants and finer-grained gaze estimation.

Privacy and ethics statement

The methods and findings presented in this study, based on data acquired in a controlled environment using a robotic system without involving human participants or personal data, pose no risks to privacy, fairness, safety, or human rights beyond those already addressed by institutional review board protocols and adheres to ethical and safety standards. We foresee minimal societal risks associated with this work and believe it contributes positively to eye-tracking research.

References

Analog Devices. accessed 2026. MAX86181: High SNR Quad-Channel AFE for Optical Spectroscopy. Product page / datasheet. Two programmable LED drivers, four 20-bit receive channels, ambient light cancellation (ALC).

Anastasios N. Angelopoulos, Julien N.P. Martel, Amit P. Kohli, Jörg Conradt, and Gordon Wetzstein. 2021. Event-Based Near-Eye Gaze Tracking Beyond 10,000 Hz. *IEEE Transactions on Visualization and Computer Graphics* 27, 5 (2021), 2577–2586. doi:10.1109/TVCG.2021.3067784

Tom N Cornsweet and Hewitt D Crane. 1973. Accurate two-dimensional eye tracker using first and fourth Purkinje images. *Journal of the Optical Society of America* 63, 8 (1973), 921–928.

Hewitt D Crane and Carroll M Steele. 1985. Generation-V dual-Purkinje-image eyetracker. *Applied optics* 24, 4 (1985), 527–537.

Andrew T. Duchowski. 2017. *Eye Tracking Methodology: Theory and Practice* (3rd ed.). Springer, Cham, Switzerland. doi:10.1007/978-3-319-57883-5

Richard O. Duda, Peter E. Hart, and David G. Stork. 2000. *Pattern Classification (2nd Edition)* (2 ed.). Wiley-Interscience.

Bartomeu Estrany and Pilar Fuster-Parra. 2022. Human Eye Tracking Through Electro-Oculography (EOG): A Review. In *International Conference on Cooperative Design, Visualization and Engineering*. Springer, 75–85.

Arvind Kumar and Glenn Krol. 1992. Binocular infrared oculography. *The Laryngoscope* 102, 4 (1992), 367–378.

Agostina J Larrazabal, CE García Cena, and César Ernesto Martínez. 2019. Video-oculography eye tracking towards clinical applications: A review. *Computers in biology and medicine* 108 (2019), 57–66.

Maria Laura Mele and Stefano Federici. 2012. Gaze and eye-tracking solutions for psychological research. *Cognitive Processing* 13, 1 (01 Aug 2012), 261–265. doi:10.1007/s10339-012-0499-z

Cristina Palmero, Oleg V Komogortsev, Sergio Escalera, and Sachin S Talathi. 2023. Multi-rate sensor fusion for unconstrained near-eye gaze estimation. In *Proceedings of the 2023 Symposium on Eye Tracking Research and Applications*. ACM, 1–8.

Alberto Pettenella, Daniele Maria Crafa, Jacopo Spagnoli, Carlo Pezzoli, Marco Paracchini, Susanna Di Giacomo, Carlo Fiorini, Sean Byrne, Tommaso Ongarello, Luca Merigo, and Marco Carminati. 2025. Development of a Low-Power Wearable Eye Tracker based on Hidden Photodetectors. In *Proceedings of the 2025 Symposium on Eye Tracking Research and Applications (ETRA '25)*. Association for Computing Machinery, New York, NY, USA, Article 117, 7 pages. doi:10.1145/3715669.3727347

Carlo Pezzoli, Marco Brando Mario Paracchini, Daniele Maria Crafa, Marco Carminati, Luca Merigo, Tommaso Ongarello, and Marco Marcon. 2026. Design and Implementation of an Eyewear-Integrated Infrared Eye-Tracking System. *Sensors* 26, 7 (2026). doi:10.3390/s26072065

Carlo Pezzoli, Emanuele Santoro, Marco Brando Mario Paracchini, Giulio Marano, Daniele Bani, Luca Francesco Raduzzi, Daniele Maria Crafa, Marco Carminati, Luca Merigo, Tommaso Ongarello, Marco Marcon, and Stefano Tubaro. 2025. Low-Power Hierarchical Network: Pervasive Eye-Tracking on Smart Eyewear. In *Proceedings of the 2025 Symposium on Eye Tracking Research and Applications (ETRA '25)*. Association for Computing Machinery, New York, NY, USA, Article 121, 7 pages. doi:10.1145/3715669.3727346

Andreas Sprenger, Birte Neppert, Sabine Köster, Steffen Gais, Detlef Kömpf, Christoph Helmchen, and Hubert Kimmig. 2008. Long-term eye movement recordings with a scleral search coil-eyelid protection device allows new applications. *Journal of neuroscience methods* 170, 2 (2008), 305–309.

Thomas JTP Van Den Berg and Henk Spekreijse. 1997. Near infrared light absorption in the human eye media. *Vision research* 37, 2 (1997), 249–253.

Nicholas Wade and Benjamin Tatler. 2005. *The Moving Tablet of the Eye: The origins of modern eye movement research*. Oxford University Press. doi:10.1093/acprof:oso/9780198566175.001.0001

Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. 2015. Appearance-based gaze estimation in the wild. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 4511–4520. doi:10.1109/CVPR.2015.7299081