PupillaRX: On-Device Pupillary Light Reflex Analysis from Smartphone Video Recording

Author: Alex Marinica

This work presents PupillaRX, a smartphone-first pipeline for pupillary light reflex recording, quality control, and metric extraction. The system combines a self-collected, manually annotated eye dataset with a lightweight MobileNet-style multi-head CNN that jointly predicts pupil and iris geometry, blink state, and eye side. Integrated into a mobile app, PupillaRX provides guided video capture, real-time quality feedback, on-device analysis, and instant metric visualization within seconds. Results obtained show the potential for low-cost, portable, and easy-to-operate pupillometry in real-world settings, as a practical complement to dedicated clinical pupillometers.

1 Smartphone Application

The interface is intentionally simple for emergency staff and nurses: an eye cutout centers the pupil, live prompts coach steady phone holding, and a visible countdown guides baseline, torch-on, and recovery timing. During capture, system provides real-time quality feedback and supports immediate retry when quality drops, improving reproducibility without compromising capture quality. After analysis, users can hover to inspect the pupil size in time, compare left and right eye metrics, and save recordings locally for offline lookup later.

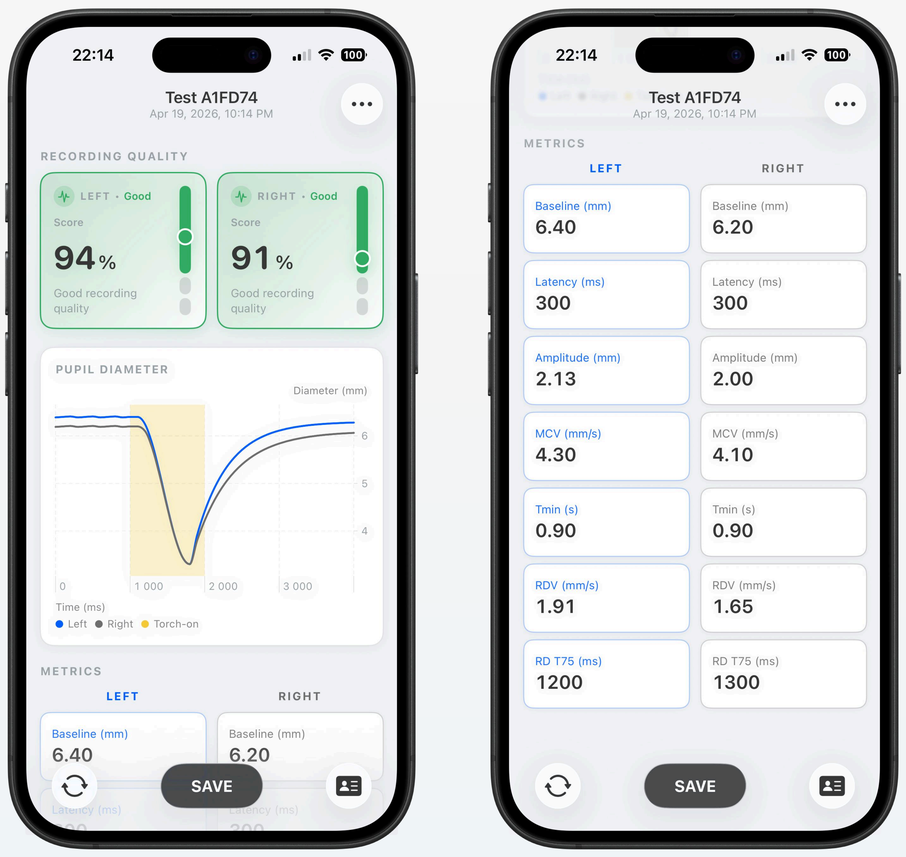

Fig. 1. User interface after analysis of both eyes, showing recording quality and derived metrics.

2 Multi-head CNN

Our model is a lightweight MobileNet-style multi-head CNN optimized for fast on-device inference. It predicts pupil and iris centers, pupil/iris radii, blink state, and eye side from each frame. In our best training run (grouped 5-fold validation), with about 71k trainable parameters.

$$\mathcal{L}_{\text{total}} = 1.0\,\mathcal{L}_{\text{heatmap}} + 6.0\,\mathcal{L}_{\text{radii}} + 0.5\,\mathcal{L}_{\text{blink}} + 1.0\,\mathcal{L}_{\text{ratio}} + 0.2\,\mathcal{L}_{\text{eye_side}}$$

Fig. 2. Shared eye-CNN takes a 128×128 frame and outputs heatmaps, ellipse radii, blink state, and eye side.
Fig. 3. Multi-task training loss function used.

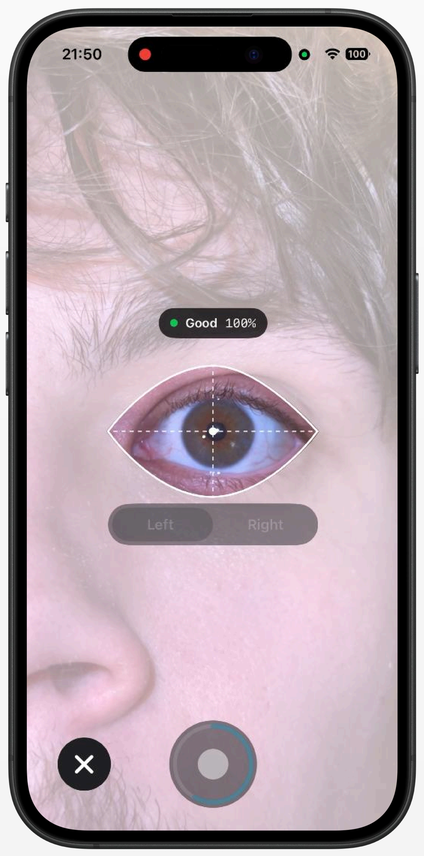

Fig. 4. Photo displaying a view during illumination of the eye using a torch while recording the eye from ~20cm distance

3 Recording Quality

We score each recording by combining two reliability terms:

- fraction of usable frames: u_{frame}
- stability of baseline phase: u_{baseline}

and aafter that we subtract artifact penalties. At frame level, we flag issues such as low model confidence, blink-like frames, missing pupil/iris detections, excessive pupil–iris center offset, unstable pupil-to-iris ratio, blur, and saturation. The final quality score is:

$$\text{score} = 0.45\,u_{\text{frame}} + 0.55\,u_{\text{baseline}} - \text{penalties}$$

PLR metrics are reported only when score ≥ 0.70, so low-quality captures do not produce misleading outputs.

4 Dataset

We built a self-annotated dataset from scratch. It contains around 2,000 close-up eye frames. Recordings were captured directly in the app on devices ranging from older to newer/flagship iPhone models. From each ~5-second, 60 fps video, we kept every 13th frame to reduce near-duplicate temporal samples and overfitting risk. The annotation schema tracks pupil and iris ellipses, confidence labels (high/medium/low), quality flags (occlusion, glint, blink), and eye side (left/right), with data collected indoors and outdoors across bright and dark conditions. The age distribution is 80% (12-25) and 20% (31-72).

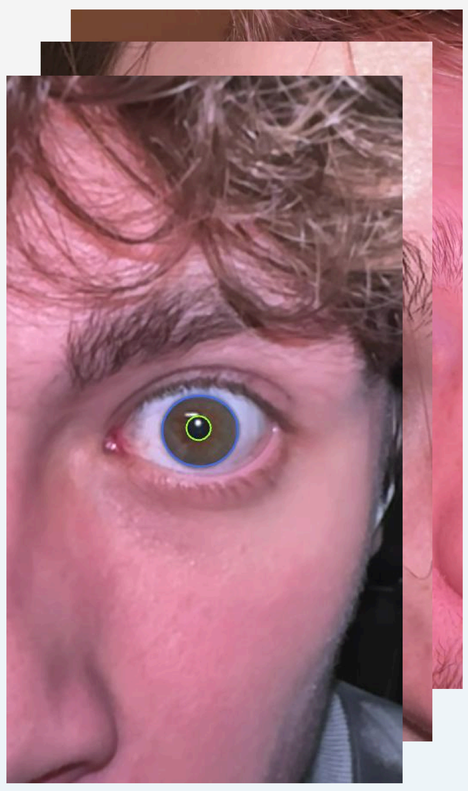

Fig. 5. Annotated images showing pupil (green) and iris (blue).

5 Results

PupillaRX suggests that neurologically relevant pupillary reflex analysis can be performed on a smartphone. In our best run, the model reached a ratio error of 0.048, pupil center error of 2.87 px, iris center error of 3.38 px, and eye-side accuracy of 0.908. The app can guide users toward higher-quality recordings and report clinically relevant metrics, including baseline diameter, latency, amplitude, constriction/redilation velocity, time to minimum, and T75 redilation. However, whether this approach could replace clinical pupillometers still requires careful validation through dataset enrichment, paired clinical comparisons, and prospective studies.

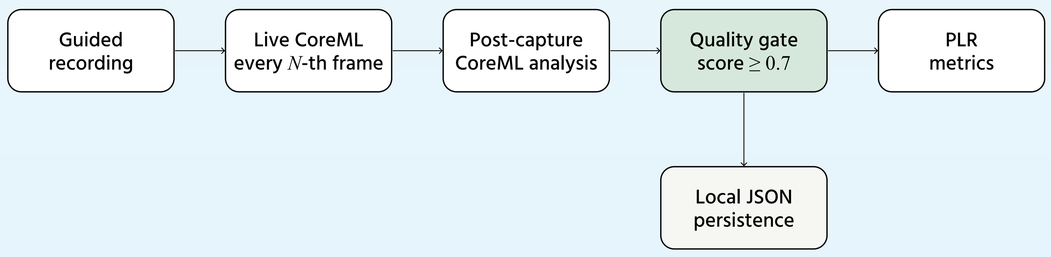

Fig. 6. Illustrates how the whole end-to-end workflow works. Starting with recording of an eye, proceeds with pre-processing, inference of video frames and finishes with data for interpretation.