PupilMetrics Neuro — Pupil Deformation Analysis

Boundary Morphology, Autonomic Innervation Theory, and Machine Learning Detection - Building the quantitative foundation for segmental autonomic assessment.

Overview

The pupil is not a perfect circle. Under high-resolution imaging, the boundary of the pupil sphincter reveals subtle regional irregularities — localized flattenings where the border is pressed inward, and protrusions where it pushes outward beyond the average radius. These deformations are not random. They occur at specific clock positions, persist across imaging sessions, and correspond to anatomically plausible zones of differential autonomic innervation.

PupilMetrics Neuro measures these deformations using two independent methods: a classical computer vision pipeline that maps boundary deviation per clock hour, and a trained machine learning model that predicts a 12-element angular deformation vector from the full iris image. Together, they offer a quantitative map of the pupil's shape — one that may carry clinically meaningful information about the regional state of the autonomic nervous system.

The Anatomy Behind the Shape

The Sphincter Is Not a Uniform Ring

The pupillary sphincter is a circular smooth muscle approximately 0.75 mm wide that runs around the inner margin of the iris. While it is often described as a single unit, its innervation and mechanical properties are not uniform around the circumference. The sphincter receives parasympathetic input via short ciliary nerves derived from the ciliary ganglion — but those nerves arrive in discrete bundles at specific meridional positions, and the density and functional state of those fibers can vary regionally.

The pupillary dilator is a thin radial smooth muscle that runs outward from the sphincter toward the iris root. It receives sympathetic input via long ciliary nerves from the superior cervical ganglion. Like the sphincter, dilator innervation is not perfectly uniform — regional differences in fiber density, neurotransmitter release, and receptor sensitivity are well-established in comparative anatomy studies.

At any given clock position on the pupil boundary, the observed radius represents the **local mechanical equilibrium** between:
- The inward pull of the sphincter at that sector (parasympathetic drive)
- The outward pull of the dilator at that meridian (sympathetic drive)
- The elastic tension of the surrounding iris stroma

When this local balance is disturbed — by a regional change in either innervation or mechanical properties — the boundary deforms.

The Theoretical Framework

Pupil Flattenings: A Parasympathetic Tone Hypothesis

A **flattening** is a sector of the pupil boundary where the measured radius falls below the median pupil radius. The boundary is pushed inward at that clock position — the sphincter is winning the tug-of-war at that meridian.

The theoretical interpretation that PupilMetrics Neuro reports is grounded in the following reasoning:

Excess parasympathetic tone at a sector → regional sphincter hypercontraction → inward displacement → flattening.

However, the more clinically interesting case is the inverse: if a flattening persists without obvious cholinergic stimulation, it may represent a **structural or functional change in the dilator at that meridian** — the sympathetic side is weaker, allowing the parasympathetic sphincter to dominate and draw the boundary inward. In this reading, a persistent flattening is a marker of reduced sympathetic tone or dilator weakness at that sector.

This interpretation parallels how we think about the ANW (autonomic nerve wreath) in the iris context: a "drawn-in" collarette, pulled toward the pupil, is associated in historical research with parasympathetic dominance or structural tightening of the iris stroma. A sector-specific flattening of the pupil margin is a more localized, higher-resolution version of the same phenomenon.

Important qualifier: The interpretation of flattenings as markers of parasympathetic dominance or sympathetic weakness is a theoretical framework derived from the anatomy of differential innervation and from observational research traditions. It is not a validated clinical diagnostic standard. PupilMetrics Neuro presents the morphological finding and the associated research context — the clinical significance is for the physician to determine.

Pupil Protrusions: A Sympathetic Innervation Hypothesis

A **protrusion** is a sector of the pupil boundary where the measured radius exceeds the median pupil radius. The boundary is pushed outward at that clock position — the dilator is dominant at that meridian.

The theoretical interpretation:

Elevated sympathetic tone at a sector → regional dilator hyperactivation → outward displacement → protrusion.

Conversely, a sector where the sphincter has reduced tone or fiber density would allow the dilator to expand the boundary at that meridian, producing the same finding through a different mechanism.

What makes protrusions particularly interesting from an autonomic research perspective is their potential relationship to **segmental sympathetic outflow**. The superior cervical ganglion that drives pupillary dilation receives preganglionic fibers from the thoracic spinal cord. Different visceral territories — cardiac, pulmonary, hepatic, renal, urogenital — have topographic representation in the sympathetic chain. If that topographic specificity is reflected, even partially, in the pattern of regional dilator activation, then the clock position of a pupil protrusion might carry information about *which sympathetic territory* is showing elevated tone.

This is the theoretical basis for the zone associations that PupilMetrics Neuro displays with each protrusion finding:

Zone (Right Eye)	Associated Sympathetic Territory
Upper-central (11:30–1:30)	Cerebral and psychomotor sympathetic patterns
Middle-temporal (8:30–10:30)	Cardiac and respiratory sympathetic zones
Lower-temporal (6:30–8:30)	Hepatobiliary and metabolic sympathetic zones
Lower-basal (4:30–6:30)	Renal and pelvic sympathetic hyperactivity
Lower-nasal (3:30–5:30)	Urogenital sympathetic activation

Note: OD and OS zone mappings are mirror images of each other, reflecting the anatomical crossing of sympathetic pathways.

Again: these associations are research-derived theoretical frameworks, not validated diagnostic criteria. They represent the most scientifically coherent interpretation of historical observational data within the context of known sympathetic neuroanatomy.

The Classical Computer Vision Detection Pipeline

Boundary Extraction and Filtering

The classical deformation analysis begins with the pupil boundary points already identified by the core iris detection algorithm. For deformation analysis to be valid, a minimum of 24 boundary points are required — fewer points cannot reliably characterize a 12-sector angular profile.

RANSAC-based outlier rejection is applied before any deformation metric is computed. The algorithm:
1. Sorts boundary radii to compute Q1, Q3, and the interquartile range (IQR)
2. Removes points beyond Q1 − 1.5×IQR and Q3 + 1.5×IQR
3. From the surviving set, computes median radius and MAD (median absolute deviation)
4. Applies a Gaussian-consistent multiplier (1.4826) to scale MAD to an equivalent standard deviation
5. Retains only points within 2.5 MAD of the median radius

This two-pass rejection eliminates eyelid intrusions, specular reflections, and motion artifacts that would otherwise produce false deformation signals.

Clock-Hour Zone Analysis

After outlier rejection, each boundary point is mapped to one of 12 clock hours based on its angular position. For each clock hour, the algorithm computes the median fractional deviation from the overall median radius:

deviation_pct = (radius - median_radius) / median_radius × 100

A 3-point smoothing operation (weighted 1-2-1) is applied across adjacent clock hours to reduce single-sector noise from sparse boundary sampling.

Zones are then mapped from clock hours to anatomical zones (upper-central, middle-temporal, etc.) with separate mappings for OD and OS to account for the left-right anatomical mirror.

Significance Thresholding

Raw deformation percentages are calibrated by a boundary factor and clamped to a maximum reportable value. A minimum threshold (`ClinicalThresholds.minReportable`) is applied — deformations below this threshold are not reported, preventing noise from masquerading as clinical findings.

Surviving findings are graded:
- **Mild** — detectable deviation, warrants documentation
- **Moderate** — clinically meaningful deviation, warrants monitoring
- **Significant** — marked regional deformation, warrants clinical evaluation

Up to 4 flattenings and 4 protrusions are reported per eye, sorted by magnitude — the most prominent deformations lead the list.

The Machine Learning Deformation Model

Why a Second Model?

The classical pipeline operates on boundary points extracted from detected pupil contours. Its accuracy depends on the quality of the contour detection — if the pupil boundary is partially occluded, low-contrast, or detected with sparse point coverage, the angular profile will be unreliable.

The ML deformation model addresses this by operating directly on the full iris image, learning to predict the deformation vector from visual patterns in the image rather than from an explicitly extracted boundary. It provides an independent, complementary measurement that is particularly useful when classical boundary detection is suboptimal.

Architecture and Input

The model is a convolutional neural network exported to ONNX format (opset 17), taking a 224×224 RGB input tensor with **ImageNet normalization**:

pixel_normalized = (pixel_value / 255.0 - mean[c]) / std[c]

where `mean = [0.485, 0.456, 0.406]` and `std = [0.229, 0.224, 0.225]` — the standard ImageNet statistics used by all major pretrained vision architectures.

Before inference, the iris center and radius (detected by the classical pipeline) are used to crop a **1.4× radius region around the iris center**, isolating the iris and inner pupil boundary from the surrounding sclera and periocular tissue. This guided crop significantly improves deformation prediction accuracy compared to a naive center crop.

The crop-and-resize pipeline runs on a **background isolate** (Dart's concurrency primitive), keeping the main UI thread fully responsive during the CPU-intensive image processing step.

Output: The 12-Element Angular Deformation Vector

The model produces **12 continuous outputs**, one per 2-hour clock window:

Index	Clock Window	Anatomical Position
0	11:30–1:30	12 o'clock (upper-central)
1	12:30–2:30	Upper right
2	1:30–3:30	Upper-nasal (OD) / Upper-temporal (OS)
3	2:30–4:30	Mid-nasal (OD) / Mid-temporal (OS)
4	3:30–5:30	Lower-nasal
5	4:30–6:30	Basal
6	5:30–7:30	6 o'clock (lower-basal)
7	6:30–8:30	Lower-temporal (OD)
8	7:30–9:30	Middle-temporal
9	8:30–10:30	Upper-temporal (OD)
10	9:30–11:30	Upper-right
11	10:30–12:30	Near 12 o'clock

Positive values indicate protrusion (outward boundary expansion).
Negative values indicate flattening (inward boundary compression).

Raw model outputs are in normalized z-score space. The model was trained with per-output standardization — each of the 12 outputs was normalized to zero mean and unit variance during training. At inference time, the app denormalizes each output using the stored training statistics:

deformation_pct = normalized_output × (label_std[i] + ε) + label_mean[i]

The training statistics (`deformation_model_stats.json`) capture the population-level mean and standard deviation of each clock-hour deformation across the training dataset. The mean values near zero (range approximately −0.87 to +1.91 percentage points) reflect that the average eye has near-symmetric deformation, with slight asymmetry in certain clock positions. The standard deviations (range approximately 6.95–10.07 percentage points) reflect the true biological variability in deformation magnitude across individuals.

Integrity Verification

Before any inference is performed, the app computes the SHA-256 hash of the loaded ONNX model bytes and compares it against a hardcoded expected hash embedded in the application binary. If the hashes do not match — indicating the model file was corrupted, accidentally replaced, or tampered with — the model is marked unavailable and no deformation results are produced.

This integrity check ensures that clinical outputs are always derived from the exact trained model, not from an accidentally substituted file.

Display

The ML deformation vector is displayed in the analysis report as a **12-bar mini chart**, with each bar representing the signed deformation at that clock position. The chart is color-coded:
- Positive bars (protrusions) extend in one direction
- Negative bars (flattenings) extend in the opposite direction

The **peak deformation** — the single clock position with the largest positive deviation — is reported numerically alongside the chart, identifying the most prominent protrusion in the angular profile.

What This May Reveal About the Autonomic Nervous System

A New Window Into Regional Autonomic Tone

Standard autonomic assessment tools — heart rate variability, galvanic skin response, blood pressure variability, thermoregulatory testing — measure the *aggregate* output of the autonomic nervous system. They integrate across all sympathetic and parasympathetic pathways and produce a global, undifferentiated measure of ANS tone.

Pupil deformation mapping offers something fundamentally different: **a spatially resolved readout of autonomic balance at specific iris meridians**, acquired non-invasively in under a second.

If the theoretical framework holds — that regional protrusions reflect sympathetic dominance at that meridian, and regional flattenings reflect parasympathetic dominance or sympathetic weakness — then a high-resolution deformation map could provide the first clinically accessible picture of *how* autonomic tone is distributed across the iris innervation network, not just what the aggregate level is.

Potential Research Applications

1. Chronic autonomic dysfunction

Conditions like postural orthostatic tachycardia syndrome (POTS), dysautonomia, and diabetic autonomic neuropathy produce heterogeneous patterns of sympathetic and parasympathetic dysfunction that current tools measure only crudely. Serial pupil deformation mapping might reveal whether these conditions produce characteristic spatial patterns of sphincter-dilator imbalance — and whether those patterns change with treatment.

2. Lateralized autonomic injury

The sympathetic innervation of the pupil is ipsilateral — the right superior cervical ganglion drives the right pupil's dilator, and vice versa. Comparing OD and OS deformation profiles could reveal asymmetric autonomic tone that accompanies lateralized neurological injury, stroke, or Horner syndrome — findings that bilateral aggregate measures would miss entirely.

3. Pharmacological profiling

Different drug classes produce different patterns of regional autonomic effect. An opioid that suppresses locus coeruleus sympathetic outflow may produce a different angular deformation profile than a peripherally-acting alpha-1 agonist that directly stimulates the dilator. Systematic characterization of drug-class-specific deformation patterns could add a spatial dimension to the pharmacodynamic profiling that the Drug Effect Monitor currently performs.

4. TBI recovery trajectory

Traumatic brain injury disrupts autonomic pathways in a pattern-specific way depending on the injury location. A frontal contusion may differentially affect the cortical modulation of certain iris meridians — a finding that would be invisible to global PLR metrics but potentially detectable in the deformation profile across serial visits.

5. Longitudinal personal baseline

Each individual has a characteristic deformation profile — a personal "pupil fingerprint" that reflects their underlying iris anatomy and baseline autonomic tone. Serial monitoring could detect departures from personal baseline that global metrics would classify as within-population normal, improving sensitivity for individual-level change detection.

An Honest Assessment of What We Know

The theoretical framework connecting pupil boundary deformation to segmental autonomic tone is scientifically coherent and anatomically grounded — but it is not yet validated against an established gold standard. The zone associations displayed in PupilMetrics Neuro are derived from the historical observational literature, which was collected without the imaging resolution and computational analysis that modern pupillometry provides.

What we can say with confidence:
- Pupil boundary deformations are real, measurable, and reproducible under controlled imaging conditions
- They occur at specific clock positions that differ systematically across individuals
- They can be detected and quantified by both classical CV and trained ML methods
- The theoretical link to differential autonomic innervation is anatomically plausible

What requires further investigation:
- Whether specific deformation patterns correlate with independently measured sympathetic or parasympathetic activity at those meridians
- Whether the zone-organ associations observed historically hold up under prospective quantitative analysis
- Whether deformation patterns change systematically with autonomic interventions (pharmacological, physiological, surgical)
- What the normal intraindividual variability is across imaging sessions and lighting conditions

PupilMetrics Neuro is positioned to generate exactly the kind of structured, high-resolution, longitudinally consistent data that could answer these questions. Every scan that documents both the classical deformation profile and the ML angular vector, timestamped, patient-linked, and exportable, is a data point in what could become the first systematic quantitative study of segmental pupil morphology as a clinical endpoint.

The Machine Learning Model at a Glance

Property	Value
Format	ONNX (opset 17)
Input	224×224×3 RGB, ImageNet normalization
Preprocessing	Iris-guided crop (1.4× radius), background isolate
Outputs	12 continuous values (one per 2-hour clock sector)
Output sign	Positive = protrusion; Negative = flattening
Denormalization	Per-output z-score, trained mean and std
Integrity check	SHA-256 hash comparison on load
Inference location	On-device, no network required
Availability	Gracefully absent — null returned if model unavailable
Classical complement	RANSAC boundary pipeline runs in parallel; both are displayed