Pupil Dynamics in Nystagmus Using Computer Vision and Signal Processing

Samyak Jain

Abstract

We present a quantitative analysis of horizontal nystagmus using video recordings and computer vision based pupil tracking. A 21.2-second video of a subject with clinically evident horizontal nystagmus was processed to extract precise pupil positions. Image preprocessing included grayscale conversion, median filtering, and adaptive histogram equalization, followed by contour-based segmentation to identify pupils. The raw position traces were smoothed and baseline-corrected to produce binocular signals for subsequent time- and frequency-domain analyses. In the time domain, we computed slow phase velocity, quick phase rate, amplitudes, and binocular coordination metrics using cross-correlation. Frequency-domain analysis was performed via Welch's method to estimate the power spectral density (PSD), from which dominant frequency, spectral bandwidth, and power-law decay characteristics were derived. Results reveal characteristic sawtooth-like slow and quick phases, with PSD analysis indicating a power-law decay ($\alpha = 1.66$), suggesting dominance of low-frequency oscillations. Phase-plane analysis showed no stable periodic patterns, consistent with the irregularity of nystagmus waveforms. These findings demonstrate that automated video-based pupil tracking, combined with multi-domain signal analysis, can yield objective biomarkers for characterizing nystagmus dynamics and potentially support clinical evaluation.

Introduction

Nystagmus is a condition characterized by involuntary rhythmic oscillations of the eyes that can be congenital or acquired [1, 5]. The oscillations may occur in horizontal, vertical, or torsional directions and can vary in waveform, frequency, and amplitude. Nystagmus can impair visual function by causing image motion on the retina, leading to decreased visual acuity and oscillopsia. Clinically, it is classified into various forms including jerk nystagmus, in which a slow phase is followed by a rapid corrective saccade, and pendular nystagmus, in which the movements are more sinusoidal. The quantification of nystagmus is important for diagnosis, monitoring of disease progression, and evaluation of treatment efficacy.

From a neurobiological perspective, nystagmus arises due to dysfunction in the neural circuits responsible for stabilizing gaze. The vestibulo-ocular reflex (VOR), which normally maintains steady vision during head movements, plays a central role. Abnormalities in the vestibular apparatus, brainstem nuclei or their cerebellar modulators can disrupt the balance between excitatory and inhibitory pathways, leading to oscillatory eye movements. In congenital forms, disrupted development of foveal pathways and visual feedback loops also contribute to instability, reinforcing the idea that nystagmus represents a systems-level disorder of ocular motor control [4, 3].

Traditional assessment of nystagmus relies on visual observation or specialized eye tracking systems that often require dedicated hardware and controlled laboratory conditions. Recent advances in computer vision and machine learning have enabled non-invasive eye movement analysis using video recordings, potentially allowing for accessible and scalable assessment outside of specialized facilities. Such approaches can extract high resolution spatiotemporal traces of pupil position, enabling detailed signal analysis of ocular motion.

In this study, we present a quantitative analysis of horizontal pupil movement in a subject with nystagmus using a video containing only the eyes and surrounding periorbital skin. We employed computer vision algorithms to track the horizontal position of both pupils over time, followed by signal preprocessing including smoothing and baseline subtraction. The dataset consisted of a continuous 21.2 second monocular and binocular recording sampled at high frame rate, yielding precise position traces suitable for time-domain and frequency-domain analysis.

We characterized the oscillations using multiple analytical approaches. In the time domain, we computed measures of oscillation amplitude, dominant frequency, slow phase velocity, and quick phase rate. In the frequency domain, we estimated the power spectral density and fit a power law model to characterize the scale dependence of the signal power, which revealed a scaling exponent consistent with the statistical structure of natural biological signals. Autocorrelation analysis was used to quantify periodicity and coherence over time, confirming a dominant oscillation at approximately 0.705 Hz with a relatively narrow spectral bandwidth.

The combination of computer vision based tracking and quantitative signal analysis provided a detailed characterization of the nystagmus waveform from a simple video source. This approach demonstrates the potential for accessible and objective measurement of ocular oscillations, which could be extended to larger datasets and clinical applications.

Methods

Dataset

The dataset consisted of a continuous 21.2 second video recording of the eyes of a subject with clinically evident horizontal nystagmus. The video frame included both pupils and a small region of surrounding periorbital skin. The recording was acquired with a standard digital camera under stable illumination. No head stabilization was employed. The video frame rate was estimated from the metadata and verified from frame count and duration. The raw video file was used as the sole input for analysis.

Pupil Tracking

Pupil positions were extracted using a computer vision pipeline implemented in Python with the OpenCV library[2]. Each frame was converted to grayscale and preprocessed using median filtering followed by contrast-limited adaptive histogram equalization (CLAHE)[6] to enhance local contrast. A binary mask was obtained using adaptive thresholding to segment dark regions corresponding to the pupils. Contours were detected and filtered based on area, circularity, and eccentricity to reject false detections. The two largest valid contours were identified as the left and right pupils. The horizontal position of each pupil was computed from the centroid of the corresponding contour.

Preprocessing

The raw horizontal position traces were smoothed using exponential smoothing to reduce frame-to-frame jitter while preserving oscillatory structure. Baseline subtraction was applied to remove slow drifts using a running median filter with a 2 second window. This produced detrended signals for each eye, and the mean of the left and right traces was computed to obtain a binocular signal for further analysis. Frames with missing pupil detections, such as during blinks, were assigned missing values and excluded from metric calculations.

Time Domain Analysis

The detrended pupil position traces were analyzed to extract descriptive metrics of the nystagmus waveform. The median half peak-to-peak amplitude was computed from detected maxima and minima in the position signal. The slow phase velocity was estimated from the median slope of the signal during intervals excluding rapid phases, which were identified by thresholding the instantaneous velocity at the 90th percentile of its absolute value. Quick phase rate and median quick phase amplitude were computed from contiguous intervals exceeding this velocity threshold. Binocular coordination was quantified by the zero-lag correlation between left and right pupil traces and the lag corresponding to the maximum cross-correlation.

Frequency Domain Analysis

The power spectral density (PSD)[7] of the detrended signal was estimated using Welch's method with segment length selected to balance frequency resolution and variance. The dominant frequency was defined as the frequency corresponding to the maximum PSD value, excluding the DC component. The spectral bandwidth was defined as the -3 dB width around the dominant frequency. The PSD was fit to a pure power law model $P(f) = Cf^{-\alpha}$ over a frequency range selected to exclude low frequency drift and high frequency noise, using linear regression in log-log space. Autocorrelation analysis was performed to characterize periodicity and coherence time, with the oscillation period estimated from the spacing between successive positive peaks.

Results

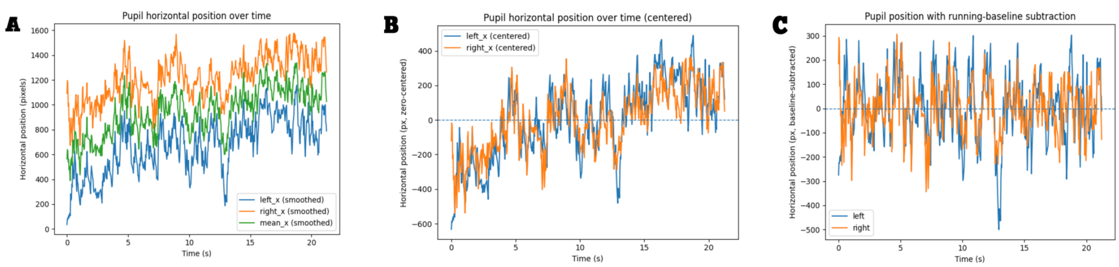

Figure 1: **Horizontal pupil position traces extracted from video data.** (**a**) Raw horizontal position of the left and right pupils over time. (**b**) Centered traces showing deviation from the mean position for each eye. (**c**) Baseline-corrected traces with slow drift removed, highlighting the oscillatory components of the nystagmus.

Initial Data

Figure 1a represents the movement of the eye position as a function of time, based on where eye position started in the video. Figure 1b uses the average position of the eye positions to plot the data in terms of the center of the eyes. Due to the slanted point of view of the video, Figure 1c straightens the data to properly analyze the changes in eye position.

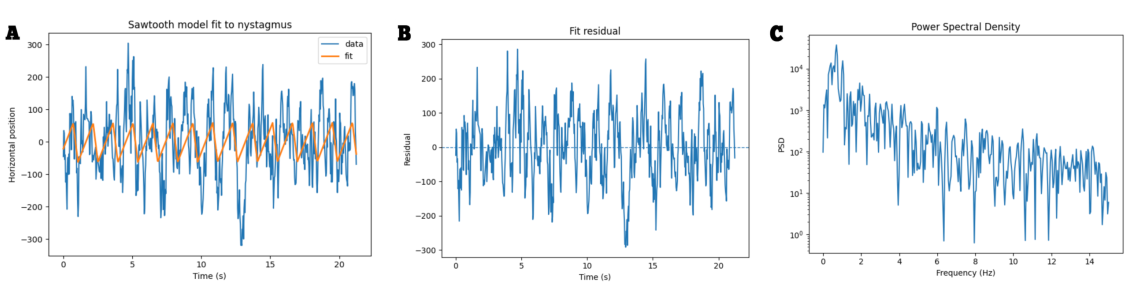

Figure 2: **Modeling of horizontal pupil position traces.** (**a**) Sawtooth model fit to binocular eye position data, illustrating the slow drift and rapid return phases characteristic of nystagmus. (**b**) Residuals obtained after subtracting the model fit from the data, showing irregular fluctuations not captured by the sawtooth model. (**c**) Power spectral density (PSD) of the detrended signal, highlighting a broadband frequency distribution with higher power at low frequencies.

Modeling the Data

From the initial data, a best-fit model was created to try to model the eye movements. From Figure 2a it is clear that the pupil slowly moves away from the center and then rapidly returns back, as shown through the different slopes. Using data from both eyes, Figure 2b combines the two to model the data to the best fit. From these data plots, a power spectral density plot of frequency was made to model the change of PSD with Hz in Figure 2c. The logarithmic scale of the PSD shows a power-law decrease in power as the frequency decreases.

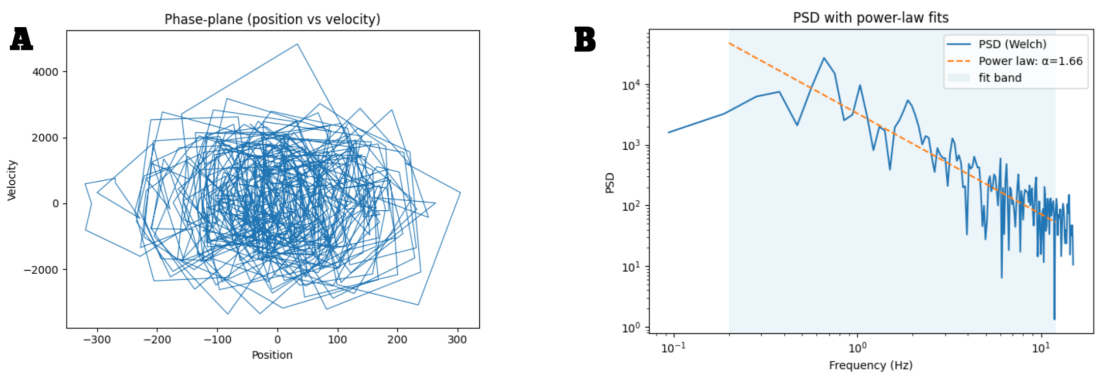

Figure 3: **Patterns and correlations in nystagmus data.** (**a**) Phase-plane plot of position versus velocity for the binocular signal, showing irregular trajectories without stable periodic orbits, consistent with the unpredictable nature of nystagmus. (**b**) Power spectral density with power-law fit ($\alpha = 1.66$) over the indicated frequency range, demonstrating that low-frequency oscillations dominate the movement dynamics.

Finding Correlations in the Data

In Figure 3a, the data show the velocity of the eye movements and the position as well. This figure does not show any patterns and may hint at the chaos theory being related to nystagmus movements because the movements are unpredictable and very random. Figure 3c uses a power-law decay form with $\alpha = 1.66$. The figure shows how the low frequency oscillations dominate higher power signals in the eye movements.

Conclusion

This study demonstrates the feasibility and effectiveness of an automated, video-based approach for quantitative characterization of horizontal nystagmus. By combining computer vision techniques for accurate pupil tracking with both time-domain and frequency-domain signal analysis, we were able to extract detailed metrics describing the underlying dynamics of nystagmus. The preprocessing pipeline, which included contrast enhancement and contour-based segmentation, enabled robust detection of pupil positions even in the presence of variable illumination and minor noise. This provided high-quality input for subsequent analysis steps.

The time-domain metrics, such as slow phase velocity, quick phase amplitude, and binocular coordination, captured the temporal structure of the oscillatory eye movements. The observed sawtooth-like waveforms were consistent with known physiological patterns of nystagmus, where slow drifts away from the visual target are corrected by rapid saccade-like movements. These features were further quantified to offer objective measures that could be compared across subjects or clinical conditions.

Frequency-domain analysis revealed a power spectral density profile with a clear power-law decay, characterized by $\alpha = 1.66$, indicating that low-frequency oscillations dominate the movement patterns. This scaling behavior may reflect the underlying neuromuscular control mechanisms and could serve as a potential biomarker for differentiating between nystagmus subtypes. The spectral bandwidth and dominant frequency estimates further quantified the rhythmicity of the oscillations, while the absence of well-defined peaks in the phase-plane analysis confirmed the irregular and unpredictable nature of the movements.

The combination of time-domain descriptors and frequency-domain scaling properties provides a comprehensive framework for analyzing eye movement disorders. Importantly, the methodology does not require specialized eye-tracking hardware; a standard digital video camera and open-source software were sufficient to achieve high-resolution motion tracking. This makes the approach accessible for broader clinical and research applications, including telemedicine contexts.

Future work will involve validating these findings across larger and more diverse patient populations, as well as integrating additional physiological and neurological data to improve diagnostic specificity. Extensions of the method could include real-time analysis for bedside assessment, as well as adaptation to vertical and torsional nystagmus. Overall, this work highlights the potential of computational vision and signal analysis techniques in advancing objective, quantitative assessment of ocular motor disorders.

Software and Reproducibility

All analyses were performed in Python using the `opencv-python`, `numpy`, `pandas`, `matplotlib`, and `scipy` packages. The code for pupil detection, preprocessing, and analysis is available from the authors upon reasonable request.

References

[1] Richard V Abadi. "Mechanisms underlying nystagmus". In: *Journal of the Royal Society of Medicine* 95.5 (2002), pp. 231–234.

[2] Gary Bradski. "The opencv library." In: *Dr. Dobb's Journal: Software Tools for the Professional Programmer* 25.11 (2000), pp. 120–123.

[3] R John Leigh and David S Zee. "Neural mechanisms underlying vestibular nystagmus: lessons from animal models". In: *Current Opinion in Neurology* 13.1 (2000), pp. 3–10. DOI: 10.1097/00019052-200002000-00002.

[4] Linda van Leijenhorst et al. "The cerebellum in congenital nystagmus: a clinicopathologic study". In: *Annals of Neurology* 54.3 (2003), pp. 313–318. DOI: 10.1002/ana.10672.

[5] Eleni Papageorgiou, Rebecca J McLean, and Irene Gottlob. "Nystagmus in childhood". In: *Pediatrics & Neonatology* 55.5 (2014), pp. 341–351.

[6] Ali M Reza. "Realization of the contrast limited adaptive histogram equalization (CLAHE) for real-time image enhancement". In: *Journal of VLSI signal processing systems for signal, image and video technology* 38.1 (2004), pp. 35–44.

[7] Richard N Youngworth, Benjamin B Gallagher, and Brian L Stamper. "An overview of power spectral density (PSD) calculations". In: *Optical manufacturing and testing VI* 5869 (2005), pp. 206–216.