A Real-Time Camera-Based Eye-Tracking System for Nystagmus Subtype Identification and Motion Correction

Design Team 4: iCrutch

Archis Shankaran, Dennis Ngo, Simren Shah, Kenzi Griffith, Immanuel Etoh, Evan Batten, Travis Tran, Esha Venkat, Jay Tailor, Edward Kuwera, Michael Repka, Michelle Zwernemann, Kemar Green

NPJ Digital Medicine

https://www.nature.com/npjdigitalmed/

Description of Requirements: Title - 15 words, avoiding punctuation, idioms, and puns. Abstract - 150 words, no subheadings. Introduction - No subheadings allowed. Results & Methods -Subheadings should be used. Discussion - No subheadings, limitations, or conclusions sections permitted. Other required sections - Data availability, code availability (if applicable), author contributions, competing interests, and references (limited to ~60). In most NPJ articles, the Results and Discussion sections precede the Methods, and the Results briefly summarize the relevant methods for each subsection.

Abstract

Acquired nystagmus, characterized by repetitive, involuntary eye movements, remains challenging to treat with current pharmacological and surgical therapies, which are often ineffective, non-specific, and unable to adapt to progressive symptom changes. Adaptive stimulation therapies, which depend on accurate, real-time detection and motion correction, may be necessary for effective treatment, yet these capabilities have not been fully achieved. To address this, we present a real-time system integrating camera-based tracking and frequency analysis to distinguish pathological oscillations from voluntary gaze shifts, enabling nystagmus detection and corrective motion vector computation. A computational model predicts the necessary counteracting motion to stabilize gaze, mapping corrective movements to extraocular muscles. Preliminary results demonstrate high accuracy (82.0%) and low response latency (0.01-0.022 ms) in nystagmus detection and motion compensation, establishing a foundation for future stimulation therapies. This system represents a step toward adaptive, closed-loop interventions for nystagmus management.

Introduction

Acquired nystagmus is a condition characterized by involuntary, repetitive eye movements that develops later in life and disrupts visual fixation. These uncontrolled oscillations result in oscillopsia, a perception of a continuously shifting visual field, making it difficult for patients to read, recognize faces, maintain balance, and perform everyday tasks requiring stable gaze. The condition can arise from various causes, including multiple sclerosis, stroke, head trauma, and neurodegenerative diseases, often leading to progressive visual impairment [1-2]. Unlike normal eye movements, which are precisely regulated by neural circuits in the brainstem and cerebellum, nystagmus reflects a dysfunction in ocular motor control, making it particularly challenging to manage.



Fig 1: Extraocular muscle anatomy and characteristic nystagmus waveforms. (a) Schematic depiction of the six extraocular muscles responsible for precise eye movement. Arrows indicate the directions each muscle pulls the eye when contracting. (Extraocular muscles. Contributed by OpenStax, License: CC BY 4.0). (b) Eye movement waveform morphologies illustrating distinct nystagmus subtypes, plotted as position versus time. (Wagle et al., *Frontiers in Neurology*, 2022).

Current therapeutic approaches for acquired nystagmus focus primarily on symptomatic relief rather than addressing the underlying neurophysiological mechanisms. Pharmacological interventions, including gabapentin and memantine, can reduce nystagmus amplitude but often come with systemic side effects and variable efficacy across patients [3-4]. Botulinum toxin injections into the extraocular muscles can transiently dampen oscillations but frequently lead to ptosis and diplopia, limiting their long-term viability [5]. Surgical procedures, such as tenotomy and muscle repositioning, attempt to realign the ocular motor system but require careful patient selection and do not prevent disease progression [6-7]. Furthermore, these treatments lack adaptability to the fluctuating nature of nystagmus, offering only partial or inconsistent improvement. The inability of current approaches to provide a dynamic and personalized solution

underscores the need for novel strategies that leverage recent advances in neurotechnology and computational modeling.

Emerging innovations in wearable eye-tracking technology have created new possibilities for real-time characterization and adaptive intervention in nystagmus. High-speed cameras now enable precise, continuous tracking of eye position and velocity, allowing for the differentiation of pathological oscillations from normal saccadic and pursuit movements [8]. Innovative signal processing techniques can further refine this analysis by classifying distinct nystagmus subtypes and predicting compensatory motor responses [9]. In this study, we present a novel eye-tracking and classification system, "iCrutch," that integrates real-time velocity computation to detect pathological eye movements, compute counteracting motion vectors, and map corrective actions to extraocular muscle groups (Fig. 2). This system lays the groundwork for a future closed-loop neuromodulation therapy capable of dynamically stabilizing gaze, offering a transformative approach to managing acquired nystagmus.



Fig 2: Real-time Nystagmus Detection and Correction Schematic for iCrutch. System overview depicting video frame processing, pupil tracking with Kalman filtering, velocity

calculation, frequency analysis, nystagmus detection based on velocity magnitude and zero-crossing rate, and real-time corrective anti-vector generation for eventual gaze stabilization.

Results

Real-time estimation of eye movement velocities

To enable rapid detection of nystagmus eye movements, we developed a real-time tracking system that estimates horizontal and vertical eye velocities from video frames. Using Mediapipe's FaceMesh model to detect pupil centers, pixel-to-degree normalization was achieved. The system's velocity outputs were compared against those produced by Otosuite, a clinically validated standard for eye movement analysis, to validate the accuracy of the algorithm's predictions [10]. All Otosuite outputs are referred to as "ground truth." A 50-second clinician-provided nystagmus video was processed simultaneously through both systems. Frame-by-frame comparisons were performed across all four tracking channels (Left Eye X/Y, Right Eye X/Y). Predicted velocities from iCrutch showed strong agreement with ground truth, achieving R² scores of 0.825–0.900 (Fig. 3a–d). Mean absolute errors were low, with values of 1.25 degrees/second (Left Eye X), 1.59 degrees/second (Left Eye Y), 1.09 degrees/second (Right Eye X), and 1.15 degrees/second (Right Eye Y). The distribution of absolute errors conveyed that 88–92% of velocity estimates deviated by less than 3 degrees/second from ground truth (Fig. 4a–d). Velocity detection performance was slightly higher for the right eye compared to the left eye. A full summary of velocity estimation performance, including mean errors, accuracies (<3 degrees/second), and R² scores across all tracking channels, is provided in Fig. 5. These results confirm that the iCrutch system can deliver clinical-grade velocity measurements with high accuracy.



Fig 3: Validation of real-time eye velocity estimation by iCrutch against clinical

standard. Scatterplots show predicted eye velocities from the iCrutch system compared with ground-truth velocities from Otosuite across left and right eyes and horizontal and vertical axes, achieving R² values of 0.825 (a) left eye vertical axis, 0.893 (b) right eye vertical axis, 0.900 (c) right eye horizontal axis, and 0.886 (d) left eye horizontal axis.



Fig 4: Distribution of absolute velocity estimation errors across eyes and axes. Histograms of absolute errors between iCrutch-estimated and Otosuite ground-truth velocities for **(a)** left eye vertical, **(b)** right eye vertical, **(c)** left eye horizontal, and **(d)** right eye horizontal movements. In all cases, 88–92% of estimates deviated by less than 3 degrees/second from ground truth.

	Left Eye X	Left Eye Y	Right Eye X	Right Eye Y
Mean Error				
(degrees/second)	1.251969421	1.591656422	1.086720172	1.14779202
Accuracy (<3				
degrees/second)	90.0%	88.0%	90.0%	92.0%
R ² Score	0.886	0.825	0.900	0.893

Fig 5: Summary table of velocity estimation performance across eyes and axes. Mean error (degrees/second), accuracy rates (percentage of frames with velocity deviations <3 degrees/second from ground truth), and R² scores for iCrutch-predicted eye velocities compared to Otosuite ground-truth measurements. Results are reported separately for horizontal (X) and vertical (Y) axes of both eyes.

Differentiating nystagmus from voluntary saccades

Accurate classification of eye movement type is critical to prevent unnecessary corrective feedback during normal eye movements. We evaluated classification performance using three clinician-annotated videos: one containing only saccadic movements, one containing only nystagmus movements, and one mixed video with 50 labeled time points conveying ground-truth classification. For the saccades-only video, iCrutch correctly rejected pathological detection 82.6% of the time (specificity), while on the nystagmus-only video, 81.5% of events were correctly classified as pathological (sensitivity). Overall classification accuracy across the mixed dataset was 82.0% (Fig. 6).



Fig 6: Confusion matrix evaluating nystagmus detection performance. The 50 frames classified by the iCrutch system are shown. True positives (nystagmus correctly detected) = 22; false negatives (nystagmus missed) = 5; false positives (saccades incorrectly classified as nystagmus) = 4; and true negatives (saccades correctly classified) = 19. Rows indicate the actual label, and columns indicate the predicted label.

Muscle targeting based on corrective anti-vectors

In addition to nystagmus classification, iCrutch was designed to identify which extraocular muscles should be engaged to counteract detected pathological movement. We validated muscle targeting performance by comparing iCrutch predictions to evaluations by a clinical expert across 50 annotated trials. Videos were analyzed using the full iCrutch detection pipeline, and predicted muscle activations were cross-checked against clinician-identified targets. The overall targeting accuracy was 94%, with medial rectus movements classified with perfect accuracy (23/23 trials correct) and lateral rectus movements classified correctly in 80% of cases (Fig. 7).



Fig 7: Confusion matrix for predicted versus actual muscle activation classifications. The 50 trials classified by the iCrutch system across four muscle activation categories are shown: Lateral, Lateral/Inferior, Medial, and Medial/Superior. True positives included 12 lateral, 4 lateral/inferior, 23 medial, and 8 medial/superior classifications. Minor misclassifications occurred between lateral and lateral/inferior activations. Rows indicate the actual muscle activation label; columns indicate the predicted label.

Latency of velocity detection and corrective vector generation

To ensure corrective feedback occurred before eye reversal during a nystagmus oscillation, we targeted a maximum system latency of 83 ms. This threshold was based on the upper bound of typical horizontal nystagmus frequencies (~6 Hz), where the available time for corrective action corresponds to half the oscillation period [11]. To evaluate system latency, timestamps were recorded at the moment of frame acquisition and again at the moment when the corresponding corrective anti-vector was rendered on screen. The difference between these timestamps represented the total computational delay.

Across multiple validation sessions, the iCrutch system exhibited an average corrective vector latency of 0.01 ms for right eye tracking and 0.022 ms for left eye tracking (Fig. 8). Both values were well below the 83 ms threshold, ensuring that corrective vectors could be calculated and displayed within a fraction of a nystagmus cycle. These findings confirm that iCrutch achieves the low-latency requirements necessary for real-time closed-loop nystagmus stabilization.





Discussion

The results demonstrate that iCrutch is capable of real-time nystagmus detection and mapping corrective intervention. iCrutch achieved high spatial velocity accuracy, with mean errors ranging from 1.09–1.59 degrees/second across all tracking channels and accuracies of velocity estimates ranging from 88%–92%. R² scores between 0.825–0.900 demonstrate strong correlations with a clinical standard, Otosuite, validating the reliability of the velocity measurements for clinical and research use. Accurate velocity detection is critical for determining the appropriate magnitude of stimulation required for extraocular muscle activation. Tracking accuracy was slightly higher for the right eye, likely due to better lighting and camera alignment during video acquisition, which improved landmark visibility.

Beyond velocity estimation, iCrutch demonstrated the ability to reliably differentiate between pathological nystagmus and voluntary saccades, achieving an overall classification accuracy of 82.0%. Specificity for correctly rejecting saccades was 82.6%, and sensitivity for detecting nystagmus was 81.5%, despite challenges posed by low-amplitude oscillations and slow rhythmic saccades. The system's ability to distinguish pathological from voluntary movement patterns is essential for enabling selective, real-time corrective interventions for nystagmus.

In addition to robust velocity detection, our system demonstrated the ability to map corrective vectors to specific muscle groups, a critical step toward closed-loop control. Muscle targeting predictions achieved 94% accuracy when compared to clinical assessments, with perfect classification of medial rectus movements and slightly lower performance for lateral rectus movements, which was likely due to increased variability in lateral tracking caused by smaller observable displacements during lateral gaze and increased sensitivity to head rotations.

12

Nevertheless, the system's ability to decompose eye movements into discrete muscle activations lays the groundwork for future closed-loop neuromuscular stimulation interventions.

Latency testing confirmed that iCrutch operates well within the required timeframe for effective correction. Fast corrective vector computation times (0.01 ms right eye, 0.022 ms left eye) ensure that interventions can occur well within the critical time window (<83 ms). These low latencies ensure that corrective vectors are generated quickly enough to meaningfully reduce oscillation amplitude before reversal, a critical requirement for real-time symptom management.

Despite these promising results, several limitations remain. Tracking performance could regress under poor or dark lighting conditions, off-center head positions, or in cases of atypical facial geometry, suggesting areas for further algorithmic improvements. Muscle targeting accuracy for compound movements remains to be tested, particularly for patients exhibiting nystagmus on multiple axes. Although iCrutch was validated using clinician-recorded datasets, larger-scale validation across diverse patient populations and nystagmus subtypes will be necessary to fully confirm the system's clinical utility.

Future development efforts will focus on expanding eye tracking robustness across varied environmental conditions (poor lighting, off-angle faces, eye closure, or partial occlusion), integrating real-time extraocular muscle targeting, and implementing wearable glasses for continuous patient use. Longer-term clinical studies will be essential to assess the durability, tolerability, and therapeutic effectiveness of iCrutch in real-world settings. With further testing and clinical validation, iCrutch has the potential to substantially improve visual stability and quality of life for affected patients. Its ability to correct pathological eye movements in real time offers a unique opportunity for therapeutic intervention, potentially transforming current

13

treatment strategies. By combining real-time tracking, classification, and muscle-specific corrective outputs, iCrutch establishes a foundation for an adaptive closed-loop therapy capable of dynamically managing acquired nystagmus.

Methods

Prototype design

The goal of iCrutch was to enable real-time detection and correction of the visual instability caused by acquired nystagmus. Upon detection, the system calculates a corrective anti-vector by inverting and normalizing the instantaneous velocity vector. Anti-vector generation occurs within 83 ms of movement onset, ensuring that corrective feedback aligns with natural visual-motor latencies. Although closed-loop neuromuscular stimulation is a future objective, the current prototype focuses on real-time detection and corrective vector output. Subsequent sections detail the algorithms responsible for pupil tracking, pathological classification, and anti-vector computation.

Eve tracking software: detection and analysis algorithms

From video data inputs, eye position and movement were recorded using the FaceMesh model of the Mediapipe library (version 0.8.10) to estimate facial and iris landmarks [13]. Tracking was performed at one face per frame, with iris landmark indices 469–472 (left eye) and 474–477 (right eye) used to localize pupil positions. Pupil center positions were computed as the mean coordinates of the four defined iris landmarks and filtered using a custom 1D Kalman filter to reduce noise. The Kalman filter was initialized with a state vector containing position and

velocity, a state transition matrix assuming constant velocity, and process and observation noise covariances set to diag([1, 3]) and 10, respectively. The Kalman filter parameters were iteratively tested and updated to minimize tracking noise. Velocity estimates were obtained by differencing the filtered pupil positions across frames and dividing by the inter-frame interval. High-pass filtering was applied to the velocity signals using a Butterworth filter with a 0.2 Hz cutoff frequency, preserving rapid oscillations characteristic of nystagmus while dampening lower frequencies. To standardize measurements across variations in subject distance and scaling, instantaneous pupil velocities, were normalized by the estimated eye width and converted from pixels/second to degrees/second using the following formula:

$$Velocity_{degrees/second} = Velocity_{pixels/second} \times \frac{30}{EyeWidth_{pixels}}$$

The eye width was estimated per frame by calculating the pixel distance between medial and lateral canthi, the corner of the eye where the upper and lower eyelids meet (landmarks 263–362 for the left eye, and 33–133 for the right eye). Pixel-to-degree conversion was calibrated by taking images of a calibration target placed at a known distance and measuring the angle between the medial and lateral canthi. Across sample videos, the average horizontal visual angle was about 30°, which was used for normalization.

Once velocities were identified, candidate nystagmus events were detected based on two criteria calculated over a sliding 1-second window: (1) a peak pupil velocity magnitude between 15–60 degrees/second, estimated as the 90th percentile of the filtered velocity distribution, and (2) a zero-crossing rate (ZCR) of the velocity signal exceeding 3 Hz. Zero-crossings were defined as sign changes in the high-pass filtered velocity trace, and ZCR was computed by counting crossings divided by twice the analysis window duration [14]. The ZCR threshold was

empirically optimized using clinician-labeled datasets to maximize differentiation between continuous pathological oscillations and isolated voluntary saccades. Nystagmus detection required sustained threshold satisfaction for at least 0.5 seconds to minimize false positives. These parameters reliably differentiate continuous nystagmus oscillations from voluntary saccades, which exhibit isolated high-velocity bursts and lower ZCR [5]. For frames with detected eye movements, anti-movement vectors were computed by inverting and normalizing the instantaneous pupil movement vectors and scaling them proportionally to pupil velocity. Muscle identification was based on a vector decomposition: positive X components above a threshold of 0.7 degrees/second indicated activation of the medial rectus; negative X components indicated the lateral rectus; positive Y components indicated the superior rectus; and negative Y components indicated the inferior rectus.

Real-time visualizations of pupil tracking, nystagmus detection, anti-movement vectors, and muscle activation overlays were implemented using OpenCV (version 4.5.5). A graphical overlay depicting the right eye's extraocular muscles was superimposed onto the video stream with consistently updated highlights based on inferred muscle activation. For offline analysis, position traces, velocity profiles, and anti-movement vector magnitudes were stored and visualized using custom scripts written in Python (version 3.8.8) and Matplotlib (version 3.7.2).

Evaluation of eye velocity measurements

To confirm the accuracy of the nystagmus velocity values calculated by the iCrutch system, outputs were compared to Otosuite, a widely used clinical standard. A clinician-provided video of nystagmus was loaded into the iCrutch system. The algorithm processed the video and outputted velocity values for the left and right eyes in the X and Y axes. The same dataset was

16

analyzed by both systems, and resulting velocity values for both eyes and axes were compared frame-by-frame across a 50-second interval. The test assumed that (1) Otosuite's velocity output represents a clinically validated gold standard for nystagmus assessment and (2) the input dataset is raw and unsmoothed, to ensure accurate frame-level comparison. The goal was for the iCrutch algorithm to match Otosuite's velocity measurements with \geq 90% accuracy. Accuracy was defined as the proportion of data points where the iCrutch output falls within ±10 degrees/second of the Otosuite value. Additionally, the mean error and R² score was analyzed for each axis of both eyes, calculated using the following equations:

Mean error:
$$\frac{1}{n}\sum_{i=1}^{n}$$
 (ViCrutch1_i - Votosuite_i)

Accuracy: (Number of Points Within $\pm 10 \text{ px}$) / (Number of Trials) × 100%

Where *ViCrutch*1_i *and Votosuite*_i represent velocities outputted by iCrutch system and Otosuite respectively.

Evaluation of nystagmus differentiation

The nystagmus differentiation capabilities of the iCrutch system were evaluated through a clinician guided comparison test. Three clinician-labeled videos were obtained from Dr. Kemar Green (Johns Hopkins Medical Institutions): one containing only saccades, one containing only nystagmus, and one containing a mix of both movement types with 50 annotated time points indicating the expected classification. Each video was loaded into the system and analyzed using the complete algorithm pipeline, which included pupil tracking, velocity detection, and nystagmus classification. For the saccades-only video, any detection of nystagmus was counted

as a false positive. For the nystagmus-only video, any missed detection was counted as a false negative. The mixed video was used to evaluate overall classification accuracy at the 50 annotated time points. To support validation, a script that automatically logged the detection outputs into a structured CSV file was developed. For each detected event, the script recorded the classification label, the corresponding timestamp, and an extracted image frame of the eye position. This enabled efficient post-processing and cross-referencing against ground truth labels. The system's classification accuracy was evaluated based on the proportion of correctly identified movements, with a target of \geq 90% accuracy for distinguishing nystagmus from voluntary saccades.

Evaluation of extraocular muscle targeting

Beyond detecting nystagmus, the system needs to produce corrective anti-vectors mapped to extraocular muscle activations. The muscle targeting test was aimed to verify the accuracy of the system in identifying which extraocular muscles should be stimulated to counteract abnormal eye movements from nystagmus. To test this, a two-pronged methodology was used. Initially, the algorithm analyzed eye movement data from videos, a total of 50 trials, to predict which extraocular muscles required stimulation. These predictions were then compared against assessments by a neuro-ophthalmologist, who independently reviewed the same videos to determine the necessary muscle stimulations. A prediction was considered correct if it matched the clinician's assessment. Following this, a digital eye model was utilized to confirm the accuracy of muscle stimulation. By inputting videos for specific counteracting vectors into the model, the stimulation of opposing muscles was verified based on whether the correct muscles were activated.

Evaluation of latency

The latency test measured the time between frame acquisition and the visual rendering of the corresponding eye velocity and corrective anti-vector overlay. Using Python's time.time() function, both the moment each video frame was captured and the moment the anti-vector was drawn on-screen were captured. The difference between these two timestamps was recorded for each frame for each eye, providing a running estimate of system latency. Latency values were stored in a buffer and visualized both in real time (on-screen) and retrospectively (plotted over time) to evaluate consistency and detect possible delays. This test quantified the real-time performance of the velocity detection and feedback rendering pipeline, ensuring that corrective cues are displayed with minimal delay after movement onset.

The test was considered successful if the time to vector calculation was under 83 ms. This threshold was determined by considering the upper bound of typical horizontal nystagmus frequencies (6 Hz). To make sure that corrective feedback could be delivered before the eye changed directions, the time required for a single movement during a nystagmus oscillation was calculated. This was calculated by taking half the period of the oscillation, given by $\frac{1}{2 \times frequency}$. At a maximum frequency of 6 Hz, this yielded a minimum latency requirement of approximately 83 ms. Ensuring corrective detection and feedback well within this window was critical to provide proper correction prior to the onset of the next oscillation.

Data Availability

The datasets generated and analyzed during the current study are available from the corresponding author upon reasonable request. Due to privacy considerations and the sensitive

nature of patient-derived eye-tracking data, access will be granted following appropriate ethical approvals and data-sharing agreements.

Author Contributions

A.S., D.N., S.S., K.G., I.E., E.B., T.T., and E.V. conceived the study. A.S., D.N., S.S., K.G., E.B., and T.T. designed the study, collected data, and conducted data analyses. A.S. and E.B. developed the nystagmus detection and counteraction pipeline. D.N., S.S., K.G., T.T, E.B., and E.V. drafted the manuscript. K.G., E. K., and M.R. provided clinical expertise and supervised the study. All authors – A.S., D.N., S.S., K.G., I.E., E.B., T.T., E.V., M.Z., J.T., E.K., and M.R. – have read and approved the manuscript.

Competing Interests

The authors declare no competing interests.

References

- Tarnutzer, A. A., & Straumann, D. (2018). Nystagmus. Current opinion in neurology, 31(1), 74–80. <u>https://doi.org/10.1097/WCO.00000000000517</u>
- Sekhon RK, Rocha Cabrero F, Deibel JP. Nystagmus Types. In: StatPearls. StatPearls Publishing; 2024. Accessed September 16, 2024.
- 3. Abadi RV. Mechanisms underlying nystagmus. J R Soc Med. 2002;95(5):231-234.
- Yasaei R, Katta S, Patel P, et al. Gabapentin. [Updated 2024 Feb 21]. In: StatPearls. Treasure Island (FL): StatPearls Publishing; 2024 Jan-. Available from: <u>https://www.ncbi.nlm.nih.gov/books/NBK493228/</u>

- Ospina LH. Dealing with Nystagmus. J Binocul Vis Ocul Motil. 2018;68(4):99-109. doi:10.1080/2576117X.2018.1493311.
- 6. Lee J. Surgical management of nystagmus. J R Soc Med. 2002;95(5):238-241.
- Stahl JS, France TD. The long-term outcomes of the Anderson-Kestenbaum procedure. Front Ophthalmol. 2023;3:1247385. doi:10.3389/fopht.2023.1247385.
- Sakazaki, H., Noda, M., Dobashi, Y., Kuroda, T., Tsunoda, R., & Fushiki, H. (2025). Monitoring Nystagmus in a Patient With Vertigo Using a Commercial Mini-Infrared Camera and 3D Printer: Cost-Effectiveness Evaluation and Case Report. JMIR formative research, 9, e70015. <u>https://doi.org/10.2196/70015</u>
- Kong, S., Huang, Z., Deng, W., Zhan, Y., Lv, J., & Cui, Y. (2023). Nystagmus patterns classification framework based on deep learning and optical flow. Computers in biology and medicine, 153, 106473. <u>https://doi.org/10.1016/j.compbiomed.2022.106473</u>
- Natus Medical Incorporated. (2024). *Otosuite* (Version 5.1) [Computer software]. Natus Medical Incorporated. <u>https://www.natus.com/products-services/otosuite/</u>
- Abadi, R. V., & Scallan, C. J. (2000). Waveform characteristics in congenital nystagmus. *Documenta Ophthalmologica*, 100(2–3), 287–304. https://doi.org/10.1023/A:1002734422320
- Robinson, D. A. (1964). The mechanics of human saccadic eye movement. *The Journal of Physiology*, *174*(2), 245–264. <u>https://doi.org/10.1113/jphysiol.1964.sp007011</u>
- 13. Lugaresi, C., et al. (2019). "MediaPipe: A Framework for Building Perception Pipelines." *arXiv preprint arXiv:1906.08172.*
- 14. Oppenheim, A.V., Schafer, R.W. (2010). Discrete-Time Signal Processing. Pearson.