Determination of Saccade Latency Distributions using Video Recordings from Consumer-grade Devices

Gladynel Saavedra-Peña[1], *Student Member*, Hsin-Yu Lai[1], *Student Member*,
Vivienne Sze[1], *Senior Member, IEEE*, Thomas Heldt[2], *Senior Member, IEEE*

***Abstract*— Quantitative and accurate tracking of neurocognitive decline remains an ongoing challenge. We seek to address this need by focusing on robust and unobtrusive measurement of saccade latency – the time between the presentation of a visual stimulus and the initiation of an eye movement towards the stimulus – which has been shown to be altered in patients with neurocognitive decline or neurodegenerative diseases. Here, we present a novel, deep convolutional-neural-network-based method to measure saccade latency outside of the clinical environment using a smartphone camera without the need for supplemental or special-purpose illumination. We also describe a model-based approach to estimate saccade latency that is less sensitive to noise compared to conventional methods. With this flexible and robust system, we collected over 11,000 saccade-latency measurements from 21 healthy individuals and found distinctive saccade-latency distributions across subjects. When analyzing intra-subject variability across time, we observed noticeable variations in the mean saccade latency and associated standard deviation. We also observed a potential learning effect that should be further characterized and potentially accounted for when interpreting saccade latency measurements.**

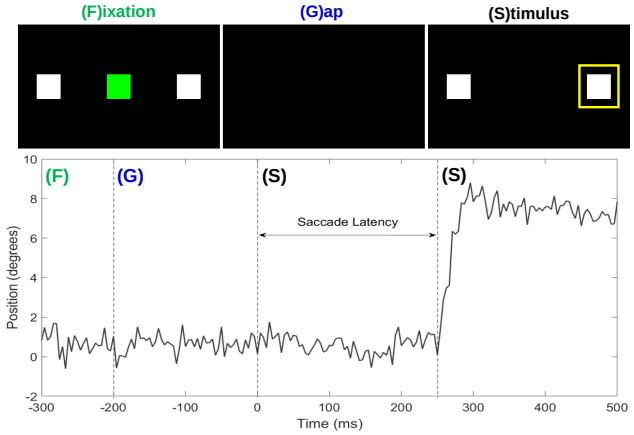

Fig. 1. Top panel: Example of the visual tracking task during a saccade measurement. Only the final 0.1 s of the 1.0 s fixation period is shown. Bottom panel: A corresponding saccadic eye movement.

I. Introduction

Quantitative and accurate tracking of neurocognitive decline remains an ongoing challenge. Repeat observations by care providers are qualitative and suffer from inter-observer variability. Standard neurocognitive and neuropsychological test batteries take significant amount of time to administer, require a trained provider, and suffer from high retest variability. Thus, no objective biomarker currently exists for accurate tracking of neurocognitive decline in the elderly or patients with neurodegenerative diseases. This technology gap is particularly limiting in Alzheimer's Disease (AD), in which expensive and invasive neuroimaging studies are performed to assess a patient's response to candidate treatments.

One way to address the lack of an objective, quantifiable, and accurate metric to track neurodegenerative disease progression is by monitoring changes in a set of digital biomarkers that correlate with disease progression [1]. Digital biomarkers are features of physiological variables obtained through portable platforms, such as laptops and smartphones. In contrast to current diagnostic methods, data from such consumer electronic devices are readily available and can be collected unobtrusively and repeatedly, thus allowing averaging to reduce random variations in observations and thereby enabling longitudinal feature tracking for monitoring of disease progression and quantification of response to treatment. In the context of tracking AD, one such digital biomarker is saccade latency – the time delay between the appearance of a visual stimulus and the eye movement towards the stimulus (Fig. 1). Previous studies have shown significant differences in saccade latency among normal subjects and patients with AD [2] as well as other neurodegenerative diseases [3], [4].

In this work, we propose and test a novel method to enable measurement and tracking of saccade latency outside of the clinical environment using a smartphone camera. Our approach eliminates the need for special-purpose capital equipment or additional sources of light, such as infrared (IR) illumination, to enhance the recording conditions. We also present a new approach to estimate saccade latency by fitting a mathematical model to the raw eye movement tracings. This method is less sensitive to high-frequency noise, compared to the current approach of numerically differentiating raw eye position tracings [5]. Finally, we report on the saccade-latency distributions of individual subjects and track the variation of mean saccade latency in healthy volunteers.

[1]G. Saavedra-Peña, H.-Y. Lai and V. Sze are with the Department of Electrical Engineering & Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA. (sze@mit.edu)

[2]T. Heldt is with the Institute for Medical Engineering & Science and the Department of Electrical Engineering & Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA. (thomas@mit.edu)

The authors would like to thank Prof. Charles Sodini of MIT for helpful discussions and comments on this work. This work was supported in part by a 3M Non-Tenured Faculty Award.

II. Materials and Methods

A. Video recordings

Video recordings on volunteers were approved by MIT's Institutional Review Board, and informed consent was obtained from each participant prior to recording. Subjects were seated centrally in front of a laptop screen, with their chin placed comfortably on a soft chin rest to minimize head movements. An iPhone 6 was also centrally placed, and video recordings were made with the phone's rear-facing camera in slow-motion mode, resulting in a frame rate of 240 fps at a full resolution of 1280×720 pixels. A second monitor was placed behind the subject's head mirroring the laptop's screen. The camera position was chosen to capture the subject's face and the mirrored screen during the task, thus capturing the eye movement and the moment the visual stimulus appeared on the laptop screen.

We used the Psychophysics Toolbox 3 for Matlab [6] to implement and display the visual fixation/stimulus task. A task started with a fixation period in which three squares were presented on the screen, arranged horizontally, against a black background: a green square at the center of the laptop screen and two white squares arranged horizontally (Fig. 1). Subjects were asked to fix their gaze on the green square. After 1.0 s, all three squares disappeared. Following a further 0.2 s, the two lateral squares reappeared in their original position but with one of them bounded by an additional square (the stimulus). Subjects were tasked with moving their eyes to – and subsequently keeping their gaze fixed on – the stimulus (Fig. 1). After the stimulus disappeared, subjects returned their gaze back to the centrally located green square. This task was repeated 40 times per trial, with 20 stimuli each randomized to the right and to the left. Three such trials were conducted in one recording session, resulting in 120 saccade measurements per session.

B. Eye tracking algorithm

To perform eye-tracking, we modified iTracker, a deep convolutional neural network (CNN) for gaze estimation on mobile devices [7]. Details on our modification of iTracker can be found in [8]. Briefly, the original version of iTracker estimates gaze by analyzing separate crops of the right eye, the left eye and the face in each video frame. It also considers a face grid that indicates the location and size of the head within the frame. Our modified version of iTracker only analyzes the face crop and the face grid, resulting in a more computationally efficient algorithm with higher SNR on our data set than the original iTracker algorithm.

C. Modeling horizontal eye movement

To calculate saccade latency, it is necessary to determine the onset of the eye movement toward the target. In prior work, the saccade onset has commonly been defined as an increase in eye velocity above a predefined threshold [2], [9], such as 30 deg/s, where the velocity had been estimated through numerical differentiation and subsequent filtering of the raw eye-position tracing [5]. Such numerical differentiation of experimental data might be permissible for recordings with very high SNR, as obtained from high-end cameras and IR-based eye tracking [10]. Given our use of consumer-grade electronics to record eye movements, point-by-point differentiation would amplify high-frequency noise in the eye-position data, resulting in noisy estimates of velocity and hence saccade onset. To overcome this limitation, we used `lsqcurvefit` in Matlab to fit a hyperbolic tangent model of the form

$$\tilde{x}(t) = A + B \cdot \tanh\left(\frac{t-C}{D}\right)$$

to the eye-position data and differentiate the resulting best-fit solution to obtain a smooth eye-velocity tracing for thresholding (Fig. 2). To quantify the model's goodness of fit, we computed the normalized root-mean-squared error (NRMSE) between the eye-position and model fit. Empirically, we found that a NRMSE<10% is indicative of a good fit.

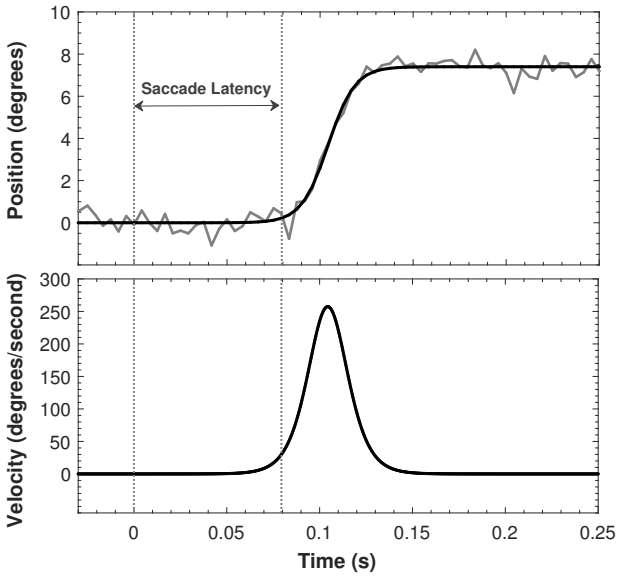

Fig. 2. Top panel: Eye position as estimated by the modified iTracker algorithm (gray) and hyperbolic tangent fit (black). Bottom panel: Eye velocity obtained by differentiating the hyperbolic tangent fit. The dashed line at 0 s indicates the moment of stimulus presentation. The saccade onset is determined by an increase in eye velocity above 30 deg/s (second dotted line). The saccade latency is the time difference between the dotted lines.

D. Data analysis

We censored saccade latencies below 100 ms to guard against anticipatory eye movements [2]. We report individual subject measurements and fit log-normal distributions to each subject's saccade-latency data. The Kolmogorov-Smirnov test (significance level of 0.05) was used to test the null hypothesis that the experimental saccade-latency distributions can be described by a log-normal distribution.

We also report the distribution of the mean saccade latencies across subjects and linear regression analysis on the mean saccade latencies in four subject in whom we had repeat measurements on over ten occasions. The coefficient of determination, R^2, and confidence limits on the slope parameters were calculated.

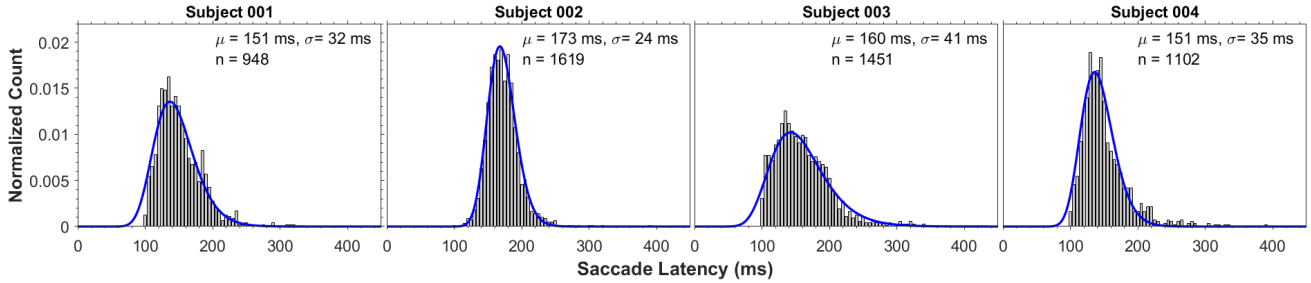

Fig. 3. Saccade latency histograms for four healthy individuals. μ is the sample mean, σ is the associated sample standard deviation, and n is the total number of observations. Saccade latencies below 100 ms were censored. The estimated log-normal probability density functions are shown in blue.

III. Results

A. Saccade-latency distributions in healthy individuals

To date, we have recorded over 11,000 saccade tasks in 21 healthy subjects. The aggregate saccade-latency distributions of four subjects are shown in Fig. 3, where all saccade fits with a NRMSE$<$10% were included. Rejected saccades include those contaminated by blinks or initially directed toward the opposite direction, as well as noisy eye traces. Saccade latencies smaller than 100 ms were censored, as they were considered anticipatory movements [2]. Taking these rejection criteria into account, the average fraction of good saccades per recording session across subjects is 73 $\pm$ 17% (about 88 $\pm$ 20/120). Some subjects had a large tendency to initiate saccades toward the opposite direction or make anticipatory movements, which reduced the fraction of good saccades per session.

The saccade-latency distributions from individual subjects show variable degrees of positive skewness, with saccade latencies larger than 200 or 300 ms not being uncommon. Others have hypothesized that reaction times follow a log-normal distribution, and we tested that hypothesis by fitting censored log-normal distributions to the saccade-latency data from individual recording sessions and also to the saccade latencies from individual subjects aggregated across recording sessions (Fig. 3).

A Kolmogorov-Smirnov test on the individual saccade-latency distributions (one for each recording session) across all subjects indicated that 77 out of 82 (94%) distributions were not significantly different from a log-normal distribution ($p < 0.05$). When the individual data for each subject were aggregated into a single distribution, like in Fig. 3, seven out of ten distributions were not significantly different from a log-normal distribution ($p < 0.05$).

The histogram in Fig. 4 shows the distribution of the mean saccade latencies from the first recording session across all 21 subjects, and indicates that there is a significant range in mean saccade latencies even in the healthy population.

B. Longitudinal analysis of saccade latency

Ten of the 21 subjects have been available for repeat measurements on five or more days, accounting for over 7,500 saccade measurements in our data set. Fig. 5 shows the mean saccade latency across recording sessions from four trained subjects, i.e., they had prior familiarity with the visual tracking task. The error bars represent one sample standard deviation above and below the mean of each session. The dashed lines connecting each subject's data points are the best-fit regression lines. With the exception of one subject, the R^2 values were $\geq$ 0.56. The subject shown in blue had an R^2 value of 0.07. The slope parameters of the linear regression model were statistically different from zero in three of the four subjects studied, indicating a potential learning effect over time.

IV. Discussion

When using saccade latency to track neurodegenerative decline, it is essential to understand the associated intra- and inter-subject variability in healthy subjects in order to put into context the changes seen in patients with neurodegenerative disease. Most clinical studies [2], [11] lack a set-up that allows for unobtrusive and easily repeatable measurements of saccade latency. Often, mean saccade latency measurements from single recording sessions in a number of patients are pooled into a single distribution (Fig. 4). In this approach, intra-subject variability is often not taken into account; very few studies [12] have reported repeat measurements per individual. In this work, we measured more than 900 saccade latencies in each of four subjects (Fig. 3) and demonstrate that the intra-subject variability can be quite substantial. These recordings were enabled by the flexible nature of our measurement system that allows for 120 saccade measurements in less than five minutes. Considering the criteria to reject certain saccade latencies, we are still able to retain, on average, a sizable amount (73%)

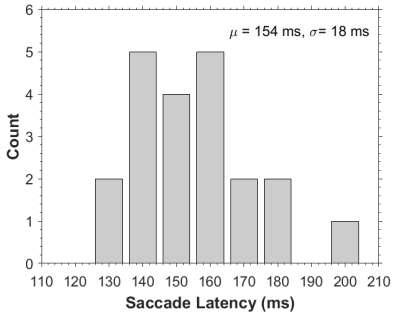

Fig. 4. Distribution of mean saccade latency values from 21 healthy individuals.

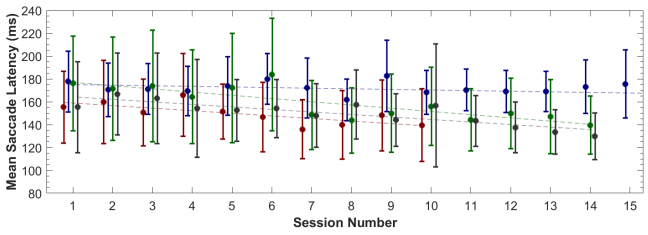

Fig. 5. Changes in the mean saccade latency across recording sessions. An individual subject is represented by a distinct color. The error bars indicate one sample standard deviation above and below the mean for each recording session. The dashed lines indicate linear regression on the mean values across recording sessions.

of data per session across subjects. With this amount of data, we observed that the shape and parameters of the saccade latency distributions differ across normal subjects, as seen in [12]. This suggests that the common approach of pooling data from different subjects may indeed obscure important intra-subject variation that might be important to track.

Moreover, we observed that the saccade latency distribution of the majority of the subjects may be modeled as a log-normal distribution. This observation is consistent with [13], in which neural mechanisms are discussed that might give rise to log-normally distributed reaction times. It might therefore be sufficient to characterize individual saccade-latency distributions using the two parameters of a log-normal distribution ($\log -\mu$ and $\log -\sigma$) and analyze how these parameters change through time. In our longitudinal analysis, however, we evaluated how the mean (μ) and the standard deviation (σ) of the saccacde-latency distributions changed across time. These statistics were chosen over the mode or the median for two major reasons: (1) the mode, mean, and median tracked one another closely; and (2) clinical studies usually report the mean and standard deviation.

One of the goals of our work is elucidating whether repeat measurements of saccade latency can be used to assess neurocognitve decline. Therefore, it is important to understand the behavior and variability of saccade latency in healthy subjects. To aid our understanding, we took multiple measurements of saccade latency per individual across days and analyzed the patterns in these measurements. The recording sessions are sequential but were not necessarily taken on consecutive days. Ultimately, the mean and standard deviation of the saccade-latency distributions were analyzed as a function of the individual recording sessions. Fig. 5 shows that for each subject, the mean saccade latency and the standard deviation change across recording sessions. A linear model on the mean saccade latencies across recording sessions accounted for 7% to 82% of the variability in the measurements. This suggests that although the mean saccade latencies varies with the number of recording sessions, there are other factors (such as tiredness or test-taking strategies) that introduce variability. To quantify the changes in these values for tracking of disease progression, we need to analyze a sufficient amount of recordings to isolate these factors. With our approach to measure saccade latency using laptop and smartphone cameras, it is promising to identify these factors while current clinical methods are too restrictive.

V. Conclusions

We present a method to measure saccade latency outside of the clinical environment using a consumer-grade camera. Furthermore, our implementation of a mathematical model to estimate saccade latency is robust to high-frequency noise, allows for automated rejection of bad saccades and therefore enables efficient large-scale data analysis. We collected over 11,000 saccade latency measurements across 21 healthy volunteers and observed that the saccade-latency distributions of normal subjects have distinctive shapes. A longitudinal analysis also demonstrates that saccade distributions are variable across recording sessions. A deeper understanding of these variations is essential to put into perspective the saccade-latency changes seen in patients with neurocognitive disease.

References

[1] E. Dorsey, S. Papapetropoulos, M. Xiong, and K. Kieburtz, "The first frontier: Digital biomarkers for neurodegenerative diseases," *Digital Biomarkers*, vol. 1, pp. 6–13, 2017.

[2] R. Shafiq-Antonacci, P. Maruff, C. Masters, and J. Currie, "Spectrum of saccade system function in Alzheimer's Disease," *Archives of Neurology*, vol. 60, pp. 1275–1278, 2003.

[3] C. Antoniades, Z. Xu, S. Mason, R. Carpenter, and R. Barker, "Huntington's Disease: changes in sacccades and hand-tapping over 3 years," *Journal of Neurology*, vol. 257, no. 11, pp. 1890–1898, 2010.

[4] R. Perneczky, B. Ghosh, L. Hughes, R. Carpenter, R. Barker, and J. Rowe, "Saccadic latency in Parkinson's Disease correlates with executive function and brain atrophy, but not motor severity," *Neurobiology of Disease*, vol. 43, no. 1, pp. 79–85, 2011.

[5] K. Holmqvist, M. Nyström, R. Andersson, R. Dewhurst, H. Jarodzka, and J. van de Weijer in *Eye Tracking: A Comprehensive Guide to Methods and Measures*, ch. 2, pp. 48–49, Oxford: Oxford University Press, 2011.

[6] "Psychophysics Toolbox 3." `http://psychtoolbox.org/`. Accessed: 2018-02-09.

[7] K. Krafka, A. Khosla, P. Kellnhofer, H. Kannan, S. Bhandarkar, W. Matusik, and A. Torralba, "Eye tracking for everyone," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[8] H.-Y. Lai, G. Saavedra-Peña, C. Sodini, T. Heldt, and V. Sze, "Enabling saccade latency measurements with consumer-grade cameras," *Proceedings of the IEEE International Conference on Image Processing*, 2018.

[9] T. J. Crawford, S. Higham, T. Renvoize, J. Patel, M. Dale, A. Suriya, and S. Tetley, "Inhibitory control of saccadic eye movements and cognitive impairment in Alzheimer's Disease," *Biological Psychiatry*, vol. 57, no. 9, pp. 1052–1060, 2005.

[10] G. Fernández, J. Laubrock, P. Mandolesi, O. Colombo, and O. Agamennoni, "Registering eye movements during reading in Alzheimer's Disease: Difficulties in predicting upcoming words," *Journal of Clinical and Experimental Neuropsychology*, vol. 36, no. 3, pp. 302–316, 2014.

[11] A. L. Boxer, S. Garbutt, W. W. Seeley, A. Jafari, H. W. Heuer, J. Mirsky, J. Hellmuth, J. Q. Trojanowski, E. Huang, S. DeArmond, J. Neuhaus, and B. L. Miller, "Saccade abnormalities in autopsy-confirmed frontotemporal lobar degeneration and Alzheimer's Disease," *Archives of Neurology*, vol. 69, no. 4, pp. 509–517, 2012.

[12] S. Gezeck, B. Fischer, and J. Timmer, "Saccadic reaction times: A statistical analysis of multimodal distributions," *Vision Research*, vol. 37, no. 15, pp. 2119–2131, 1997.

[13] R. Ulrich and J. Miller, "Information processing models generating lognormally distributed reaction times," *Journal of Mathematical Psychology*, vol. 37, no. 4, pp. 513 – 525, 1993.