Automated gaze direction scoring from videos collected online through conventional webcam.

Contributors:

Fraser, A., Gattas, S. U., Hurman, K., Robinson, M., Duta, M., & Scerif, G.

Abstract

As online research has become more prevalent, researchers have been investigating the possibility of replicating techniques that go beyond measuring simple behaviour alone. One such method could leverage the webcam of the participants' device to collect information about eye gaze direction. Several packages have been developed for collecting such data, but they can lead to high attrition, and require extensive and potentially frustrating calibration procedures. This hinders all research, and in particular, data collection including children and participants with neurodevelopmental difficulties.

To overcome this issue, we developed GazeScorer, a package that uses basic image processing techniques and a simple one-point calibration to score horizontal gaze orientation. We based our work on experience gained from infant research, in which a researcher manually scores horizontal gaze orientation from individual video frames. Using videos collected in two browser-based remote studies, one including adults, and one including children, we achieved low participant attrition, single-frame point calibration, and demonstrated a good level of inter-rater reliability between GazeScorer and a manual scorer.

Our package provides a new participant-friendly but reliable resource for researchers working on studies in which information about static horizontal gaze orientation is sufficient.

Credit Statement

Alex Fraser Conceptualization, Methodology, Software, Formal analysis, Investigation, Resources, Data curation, Writing - Original draft, Writing - Review & Editing, Visualization, Funding acquisition;

Sylvia Gattas Conceptualization, Methodology, Investigation, Resources, Writing - Review & Editing, Writing -Review & Editing;

Katie Hurman Data Curation, Writing - Review & Editing;

Martin Robinson Conceptualization, Writing - Review & Editing, Supervision;

Mihaela Duta Conceptualization, Methodology, Software, Formal analysis, Investigation, Resources, Data curation, Writing - Review & Editing, Visualization, Supervision;

Gaia Scerif Conceptualization, Methodology, Formal analysis, Investigation, Writing - Review & Editing, Supervision, Project administration, Funding acquisition;

For many years, psychology researchers have attempted to realise the potential of collecting data remotely over the internet. Early data collection methods required a high level of technical expertise to allow building of experiments (A. Anwyl-Irvine et al., 2021), but new JavaScript packages have made online cognitive research more accessible (jsPsych: de Leeuw, 2015; Lab.js: Henninger et al., 2019; Psychojs: Peirce et al., 2019), and platforms such as Gorilla (A. L. Anwyl-Irvine et al., 2019) and Pavlovia (Peirce et al., 2019) have made hosting and running experiments even more accessible. These platforms have demonstrated high timing precision, comparable to lab-based alternatives, with timing suitable for experiments which need less than 100ms reaction time accuracy (see A. Anwyl-Irvine et al., 2021 for a more extensive history and comparison with lab-based software). Online experimentation has allowed researchers to collect larger and more diverse data sets (Buhrmester et al., 2011; Peer et al., 2017; Sheehan, 2018), helping them to address issues raised as part of the Reproducibility in Psychology crisis (Open Science Collaboration, 2015) and to reach larger audiences with greater diversity (Arnett, 2008; Rad et al., 2018; Henrich et al., 2010). It has also proved invaluable during the COVID-19 pandemic as researchers were forced online due to lockdowns and social distancing measures. Thus far, online experiments have proven powerful in collecting questionnaire data, as well as the accuracy and reaction time of manual responses. However, many researchers who rely on complementary cognitive measures that are traditionally collected either in the lab or with costly portable equipment have been limited in their ability to collect data online.

One such complementary method is eye tracking. Traditional lab-based eye trackers project infrared ("IR") light, which reflects from the eye and is detected by IR-cameras. In controlled situations this can result in highly accurate and precise measures of a participant's gaze position when calibrated in relation to a screen. Eye tracking provides a rich data source for studying overt attention, with a huge potential for different measures (Holmqvist et al., 2011) – including 120 individual measures and analysis for eye tracking data categorised by Holmqvist et al (2011). As personal computing has become cheaper and more powerful, access to devices with built in webcams has become a relative norm across countries, and researchers have begun to investigate whether it is possible to leverage them to extract gaze position information. There are currently several software packages that have attempted to replicate the abilities of an eye tracker using conventional webcams. Many of these packages are in the early stages of development and are yet to be widely used in active research for evaluation.

Here we compare three approaches attempting to replicate eye tracking data remotely and discuss the pros and cons of each. To do this we have identified three factors that we felt were functionally important for researchers who are attempting to use eye tracking software remotely. The first factor is the spatial resolution: lab-based eye trackers are often able to achieve high levels of precision and accuracy (Dalrymple et al., 2018), and

typically report pixel level estimates of gaze location. Many researchers are reliant on the high spatial resolution offered by lab-based eye trackers and therefore are looking to replicate these abilities with remote eye tracking methodologies. The second factor is the ease of participant interaction and the consequences challenging interactions have on data collection. In the lab eye-tracking environments and calibration procedures can be optimised and controlled by the researcher to minimise the impact on the participants. However, with remote data collection, calibration procedures can be more difficult to manage, meaning that they are either excessively involved (requiring significant input from the participant and the researcher) or may result in lower quality data. In addition, whereas in the lab it may be possible to make informed judgments on the quality of the calibration and run recalibration if required, this may not be an option when collecting data remotely. The third factor is whether the technique explicitly measures overt visual attention. Lab-based eye tracking techniques measure discrete eye movements, allowing direct measurement of overt attention. Thus, it is possible to draw conclusions related to the relationships between visual stimuli and different types of attentional processes.

Lab-based eye trackers are considered the gold-standard, and while the factors discussed here are not an exhaustive list of considerations when assessing eye tracking techniques, they provide a framework for comparing current remote eye tracking options available to researchers. In the next section we consider three remote techniques and discuss how well each one is able to perform against these three factors. I will then discuss the implications for remote eye tracking and present new methodologies that will contribute to the growing field.

Comparison of current remote eye tracking techniques

WebGazer.js

The first package that has gained significant traction is WebGazer.js (Papoutsaki et al., 2016). WebGazer is a JavaScript package that accesses the webcam through the participant's browser. It then detects the pupil and uses a regression analysis model informed by interactions with the participant to estimate gaze location as they interact with a website. Papoutsaki and colleagues found that WebGazer was comparable to a consumer grade eye tracker, and they suggested that WebGazer is suitable for research where only approximate gaze location is required. However, the core methodological foundation of WebGazer is its ability to consistently update its estimated gaze position based on user interaction through mouse clicks (Papoutsaki et al., 2016). Working on the assumption that participants' eye gaze closely correlates with mouse position when they produce a mouse click (Hauger et al., 2011), the software regularly updates the calibration to measure gaze between click events. This is ideal when the participant is tasked with viewing a website and

therefore regularly clicks to update the calibration, but it is not practical for many current psychological paradigms where either tight experimental control is required or mouse clicking is not possible because of participants' characteristics.

As WebGazer's code is open source, researchers have been able to adapt its pupil detection and regression model to use a more traditional 'point-style' calibration to be used with paradigms where mouse movements were not possible (Madsen et al., 2021; Semmelmann & Weigelt, 2018; Yang & Krajbich, 2020). However, difficulties have arisen when using this style of calibration in active experiments. Several environmental factors such as brightness and lighting direction can heavily impact WebGazer's ability to detect the participant's face, leading to a potentially frustrating situation for participants, which may explain high and prohibitive levels of attrition, even with compliant adult participants (around 60% for both Semmelmann et al., 2017 and Yang & Krajbich, 2020). In addition, by not using the native mouse style calibration, it is not possible to update the calibration required for the model to work effectively, resulting in interruption of the task to recalibrate regularly to maintain accuracy (Semmelmann et al., 2017; Yang & Krajbich, 2020). Together, these calibration issues may result in poor data quality and lead to significant frustration for the participant, which in turn may require either a high level of researcher engagement (to walk participants through the calibration), or very large recruitment samples to compensate for poor data quality and high levels of attrition. In addition, participants for whom long and language dependent calibration routines are not an option (young children, individuals with attentional difficulties and / or poor language comprehension) are out of reach with this approach.

Against the key factors we identified earlier, WebGazer has demonstrated reasonable spatial resolution using discrete measurements of the eyes, therefore providing direct measures of overt attention. While not as precise as lab-based eye trackers, it is sufficient for researchers to investigate multiple regions of interest throughout a task. However, there may be issues for how participants engage with the software, resulting in difficulties maintaining the calibration required, and therefore frustration and high levels of attrition. This will impact data quality, as well as limiting usability for participants with specific requirements.

MouseView.js

An alternative approach to calibration-free measures of eye tracking is mouse tracking. A recent JavaScript package, Mouseview.js (A. L. Anwyl-Irvine et al., n.d.), has provided an alternative to webcam-based eye tracking for online experimentation. MouseView.js works by projecting a distorted layer over the image with foveal viewpoint, replacing the on-screen mouse location so that the participant can investigate a scene in a way that attempts to replicate visual attention. In their experiments, the authors demonstrated good correlation with their previously collected eye tracking data (A. L. Anwyl-Irvine et al., n.d.), replicating the findings of previous mouse tracking experiments (Lingnau et al., 2008; McConkie & Rayner, 1975; Rayner, 2014). This approach provides the opportunity for researchers to measure attention in visual search paradigms in a way that can replicate aspects of visual attention, but it does not measure overt visual attention and therefore may lose aspects of bottom-up, more diffuse elements of involuntary attention that are also a desirable factor in eye-tracking. However, it does not require extensive calibration and does provide very high spatial resolution.

Against the key factors we defined earlier, Mouseview.js provides high spatial resolution and offers easy participant interactions. However, it does not directly measure overt attention in the way that eye-tracking does. Therefore, it can provide an effective method for researchers who are interested in deliberate, voluntary attention shifts, but may not be suitable for researchers who are interested in more reflexive, automatic attention (A. L. Anwyl-Irvine et al., n.d.). Although this may limit the overall usability of mouse tracking, these techniques may offer the unique opportunity of measuring voluntary shifts in attention, while limiting the potential impact of more reflexive mechanisms attention, offering a valuable additional use case to measure differences between these attentional processes.

Manual scoring from conventional webcams

An alternative technique that has been used previously is manual scoring. Manual scoring of gaze location has been used for research on infants and young children for many years (Fantz, 1965; Teller, 1979). In a modern manual scoring setup, a child sits in-front of a screen while a series of stimuli are presented. Either in combination with eye-tracking, or in isolation, a video recording is taken of the infant as they view stimuli. A manual scorer then scores the amount of time the infant spent looking at visual stimuli, allowing the comparison between the time spent attending to one stimulus over another, or even if they are attending to the screen at all. Using this technique, it is possible to gauge the interest of an infant who is too young to respond directly, with now classical and pervasively used paradigms including preferential looking procedures (Houston-Price et al., 2007; Leo & Simion, 2009;

Teller, 1979) and looking-while-listening tasks (Fernald et al., 2008; Swingley, 2012). In addition to use with infants and young children, gaze data have been highly informative for both young and old individuals whose ability to comply with complex experimental instructions and mouse clicks, or other behavioural responses is more limited. Venker & Kover (2015) have provided a discussion of the advantages and disadvantages of using gaze data from participants who have neuro developmental conditions (NDCs) via manual scoring or automated eye trackers. They list several factors to be considered when deciding on the type of eye tracking method to use, and these are also applicable to online data collection with neurotypical infants, children and adults.

A number of researchers have been investigating the possibility of collecting video data for the purpose of infant looking paradigms using webcams in the participants' own homes. This method of data collection could potentially be problematic as the quality of the video may not be sufficient and the smaller screen of most home computers may not be large enough to elicit sufficient eye movement for accurate gaze scoring (C. E. ; Venker & Kover, 2015). However, several research groups have been able to replicate lab-based experimental findings using manual gaze scoring of home-collected webcam video-data and have found the quality to be adequate for manual scoring (Hermes, n.d.; Scott et al., 2017; Scott & Schulz, 2017; Semmelmann et al., 2017). Specifically, there are now multiple platforms that provide researchers with the ability to create experiments designed for babies and infants, while simultaneously collecting video through the webcam; namely, E-babylab (Hermes, n.d.) and Lookit (Scott et al., 2017; Scott & Schulz, 2017). Both platforms have demonstrated that they were able to collect video data that provided adequate quality to replicate the findings of studies that had previously been collected in the lab, highlighting the potential of manual video coding methods for remote data collection going forward. Against our key factors, manual scoring offers the ability to measure overt visual attention without the need for complex or invasive calibration procedures. However, it has very limited spatial resolution, with typically only three observable locations — left, right or middle.

It is clear that, while all of the packages discussed here offer the ability to remotely collect data that replicate data from an eye tracker, none of the options are able to adequately fulfil all three of the key factors we had identified. This does not mean that these techniques are not useful for those considering remote eye tracking research, but it does highlight the need to seriously consider which factors are important for specific research questions and the importance of selecting the techniques that best match desired hypotheses. Furthermore, the variety of methods now available suggests that this is an emerging field that will continue to grow as technology improves and demand increases. In the following section we discuss our own future developments in this arena, specifically in relation to automating the manual scoring method.

Manual coding – The (painful) alternative approach to identifying gaze position

There are considerable limitations of manual scoring that may be off putting for researchers. Firstly, there is the limited spatial resolution as manual scorers are only able to offer limited gaze information, such as left or right directional choices or whether the participant is attending the screen at all (Fernald et al., 2008; C. E. Venker et al., 2020). This contrasts with a lab-based eye tracker which typically offers a much finer level of spatial resolution, providing wider applications for research (Dalrymple et al., 2018). Secondly, when a manual scorer is tasked with coding the gaze data of a participant, it is not possible to entirely remove the risk of bias coming from external contextual cues (Oakes, 2012). Finally, the time and resources required for manual coding make it a very labour-intensive method. A trained manual scorer may take approximately 1-2 hours to score 4 minutes of video, plus any training and validation required (C. E. ; Venker & Kover, 2015). This raises the question: Is it possible for an analysis pipeline to estimate the gaze orientation of a participant with comparable accuracy to a manual scorer from videos collected using conventional webcams?

Chouinard et al., (2019) utilised this approach with Amazon Web Services (AWS), Rekognition video analysis (Mishra, 2019), a deep learning platform that is able to detect features of a face including the pupil, which can then be used to create a custom classifier to estimate gaze location. The authors found that they were able to classify gaze in a preferential looking paradigm at above chance levels of accuracy (Chouinard et al., 2019). However, there are a number of limitations to using Amazon Web Services. Firstly, there is a cost associated with using the service, although it may be more cost effective than lab-based eye tracking, depending on the number of participants assessed over time. Secondly, although the classifier itself is open source, the Amazon Rekognition software is not, so it is not possible to identify how the software is working at its core. Finally, the use of web-based services to analyse video may have privacy concerns that need to be considered, particularly if Amazon is retaining videos to improve their current algorithms.

The current approach: Automated gaze direction detection

We propose a new approach to overcoming the issues that arise when using conventional webcams to measure eye gaze data. Instead of attempting to replicate the high-resolution positional information that lab-based eye-trackers provide, our approach instead focuses on providing detected gaze orientation. At its core our approach tackles gaze data as a “simple” image processing problem: Can an automated algorithm reliably classify eye gaze orientation based on images encompassing face landmarks and the eyes? At this stage we aim to produce such an algorithm with two goals in mind: 1) that it has the ability to accurately distinguish gaze orientation while subjects follow a target stimulus, and

2) that it removes the need for excessively long and involved calibration procedures prior to data collection. Reliable software without complex calibration is important to minimise rates of participant attrition, even for adult samples, and it is critical for hard-to-calibrate populations. Automatic gaze scoring will not be able to fulfil all three of our key factors as it will only have limited spatial resolution, but we believe it will provide an additional resource that will expand on available remote eye tracking solutions.

Here we present an analysis pipeline that investigates the possibility of automatic gaze scoring of video data. This replicates the ability of manual scorers when making gaze direction judgements using only basic video processing methods. Building on a similar approach to that of Chouinard et al., (2019), it does however not rely on AWS to operate and therefore enhances opportunity for open protocol sharing. To test the software, video data was collected from participants following a target stimulus as it moved to and from static locations across the screen. The videos were then analysed by both the GazeScorer software (referred to as the automatic scorer) and a manual scorer. It was predicted that the automated scorer (“AS” henceforth) and the manual scorer (“MS” henceforth) would demonstrate a good level of agreement when making judgments on participants’ gaze direction.

Study 1 – Adult sample

Methods

Data Collection

Materials

Visual stimuli

Videos that presented a target dot moving on the screen on a light grey background were generated using PsychoPy (Peirce et al., 2019) and recorded as screen captures at a resolution of 1920x1080. Following a 5-second countdown, the dot was presented in the centre of the screen for 3000 ms. The dot then travelled along the horizontal plane to the left side of the screen and paused there for 1000 ms. Following the pause, the dot returned to the centre and paused there for 1000 ms before moving to the right-hand side of the screen and pausing there for 1000 ms. Finally, the dot returned to the centre of the screen and paused there for 1000 ms before disappearing. A blank screen was then presented for 2000 ms before the end of the video.

Browser-based remote study

The study comprised of the presentation of the videos described above while simultaneously collecting video data from the webcam of the participants' device. The study was administered using the Gorilla platform (A. L. Anwyl-Irvine et al., 2019) and participants were instructed to access it using the Google Chrome browser.

Participants

For Study 1, 20 adult participants were recruited through direct contact by the researchers. Only 16 participants reported demographic information. Of these participants there were 9 who identified as female and 7 who identified as male. Their ages ranged from 21 to 32 ($M = 27.8$). All participants were made aware about how their video footage would be collected and stored and fully informed consent was received before data collection began.

Procedure

Participants were sent a direct link to the remote study where they first had to read an online information sheet and complete the consent form. Participants were then presented with a set up procedure in which they were shown a live webcam feed alongside a demonstration of a number of positional and lighting requirements to optimise data collection. Once the participant had confirmed that they had met these requirements they proceeded to record an example video which was then played back to them to ensure the study could run. If the recording was successful and the participant was happy to continue, they were presented with the study stimuli.

Participants were first displayed short instructions about how to proceed. When they were ready, they pressed a 'Record' button that loaded the video stimuli and started the webcam recording. The video opened and played in full screen. At the end of the stimuli video the study would automatically stop the video playback and webcam recording.

Gaze Scoring

Each video recording was first trimmed to the length of the stimuli presentation and then split into individual frame images using ffmpeg and saved as png images. The number of individual frame images depended on the frame rate of the participant's webcam: for three participants the webcam had a frame rate of approximately 15 fps, the remaining participants webcam had a frame rate of approximately 29 fps. The resulting frame images were used to score the horizontal orientation of the participant's gaze on a frame-by-frame basis by a human scorer and an algorithmic scorer.

Manual Gaze scoring

Manual scoring of the video footage was conducted by a researcher who was completely naïve to the study and video stimuli. An online task was built in Gorilla to allow the MS to view images of the individual frames from the video footage. The individual frame images were presented in chronological order, but the MS could also navigate backwards and forwards between the frames. The MS marked the frame where there was a change in gaze orientation, or when the participant blinked.

Algorithmic Gaze Scoring

The individual frame images were processed in python, using the OpenCV (Bradski, 2000) and dlib (King, 2009) packages, and custom routines. Each individual frame image was first converted to greyscale. Then, the bounding box of the face and individual eyes were extracted using dlib frontal face detector and 68-point face landmarks detector, respectively (Álvarez Casado & Bordallo López, 2021). Images of the individual eyes were then extracted and processed individually to detect the iris area and infer from that the gaze orientation. Figure 1 shows the distribution of the width and height of the eye bounding box across all the individual frames from all participants. The figure shows that the eye width varied from 21 to 57 pixels, around an average of 35 pixels.

Figure 1. Width and Height (in pixels) for the detected eye region across all frames.

The detection of the iris assumes that the iris is the largest dark and symmetrical patch within the eye. Each individual eye image was first smoothed with a median blur filter to reduce noise. A threshold at 50% of the image grey scale was then derived and used to create a binary contrast image in which all the points corresponding to original grey values above the threshold are set to 1 (white) and the rest to 0 (black). In this contrast image, the iris would be inside the largest area of black. However, depending on the individual features of the eye and the relative brightness of the iris, sclera and eye lids, the largest patch of black might also contain parts of the upper eye and/or relatively darker parts of the sclera. These extraneous areas would also make the black patch asymmetrical, and this insight was used in an iterative process that incrementally removed the extraneous pixels such that the patch converged to a more symmetrical shape. The centre of the resulting patch was then considered as the centre of the iris. The horizontal orientation of the gaze was then judged

relative to a reference point represented by the average position of the centre of the iris in a small number of frames at the start of the recording when it was assumed the participants looked at the centre of the screen. Frames were marked as blinks if the aspect ratio of the eye bounding box dropped below one standard deviation from the mean. The resulting gaze orientation scoring was median filtered to remove noise.

Gaze scoring comparison

Both AS and MS data were processed to assign the last known value to all frames that were marked as a blink and a record was kept of the number of frames for which the previously known value was recorded. Once all frames were processed, a Cohen-Kappa score was calculated to compare Inter-Rater Reliability (IRR) between the automatic and the manual scorers. The two scorers were entered as the judges and individual frames were entered as subjects. Initial comparisons were made for an aggregate of frames across participants to assess overall IRR, but subsequent analyses were run at a participant level to investigate differences in quality between individual participants. Following the convention from (McHugh, 2012) comparisons with a Cohen's k score below 0.2 were considered to have no agreement, scores between 0.21 and 0.39 had minimal agreement, scores between 0.4 and 0.59 had weak agreement, scores between 0.6 and 0.79 had moderate agreement, scores between 0.8 and 0.9 were had good agreement, and a score above 0.9 was considered near perfect agreement. Based on these conventions, we proposed a minimum of moderate agreement ($k = 0.6$) for the automatic gaze detection to be successful.

The comparison between manual and algorithmic scoring was performed separately for:

1. Frames of the video footage corresponding to the dot moving (dynamic phase)
2. Frames of the video footage corresponding to the dot in a static position (static phase)

A gaze orientation score average was calculated for each frame across all participants for both MS and AS score. This measure allows a visual evaluation of the gaze orientation score distribution along the video timeline.

Results

The average orientation score across all participants for both manual and algorithmic scorers across the entire duration of the stimulus video were plotted and can be found in Figure 2. The time intervals for which the dot was static were also marked. The figure shows that there is good agreement between the two scorers. However, there was a higher level of agreement between the AS and the MS when the dot was static, compared to the frames in which the target was moving.

Figure 2. Grand average plots for the sample $N = 20$. The purple line represents the automatic scorer, and the yellow lines represent the manual scorer. The SE around the M is plotted for each point. The green line marks the location of the target stimuli during the static phase.

Analysis of static phase

There was a near perfect level of agreement between the two judges for the frames in the static phase (of $N = 4400$): left eye ($k = .960$, 95% CI [.951, .969]), right eye ($k = .957$,

95% CI [.948, .967]), and the overall judgments ($k = .964$, 95% CI [.955, .973]). Inspection of the confusion matrix found that the scorers agreed on 4335 frames of the total 4400 frames. Of the frames where there was disagreement, 33 occurred for left orientation judgments, marginally more disagreements for right orientation judgments with a total of 34. There were no frames where an opposite directional judgment occurred. To examine how the Cohen's k scores are distributed across participants, the tests were conducted for each participant individually. The results of the Cohen's k tests can be found in Figure 3. There was disparity between the left and right eyes for a small number of participants. There was only one case of agreement falling below the established Cohen's kappa threshold of 0.61. However, this does not appear to have had a critical impact as the alternate eye achieved very good agreement with the MS. In a number of cases the overall judgment does not select the eye that has the highest level of agreement, which will be considered in more detail in the discussion.

Analysis of dynamic phase

As with the static target frames, Cohen's k tests were conducted on frames where the target dot was in motion. Individual tests were conducted for the left and right eye, as well as the overall judgments. The MS and AS scored each frame as individual subjects for a total of $N = 3576$. There was a weak level of agreement between the two scorers for the left eye ($k = .514$, 95% CI [.490, .538]), right eye ($k = .501$, 95% CI [.477, .525]), and the overall judgments ($k = .472$, 95% CI [.448, .497]). Looking at the data from the overall judgments, it was found that the scorers agreed on 2383 of the total frames. Scorers disagreed on direction judgments at a similar level, with 590 times for a left direction and 703 times for a right gaze. All disagreements for a direction judgment were made between a directional choice and the centre. That is, at no point did either the MS or AS score the gaze in the opposite direction to the other. To investigate inter-rater reliability for individual participants, a series of Cohen's k tests were run for each participant across both eyes and the overall judgments (Figure 4). Again, it appears that there is disparity between the quality of judgments based on each eye when the dot was moving. Emphasis on the potential for improving these disparities will be considered in more detail in the discussion. The results of these analyses suggest that although there are trends towards agreement, the quality of the AS was lower compared to the MS while the target stimulus was moving.

Figure 3. Cohen's k scores for individual participants during the dynamic phase, whiskers represent SE around the M . Only two participants (Participants 001 and 020) found agreement above the minimum threshold for all points. Five other participants demonstrated at least one point above the 0.6 threshold. Although the threshold fell within the confidence interval for each point for four of the six participants.

These overall results do not provide insight into why agreement is lower between the AS and the MS during dynamic phase. A likely explanation is that the AS was able to classify gaze direction better at wider excursions of the target stimulus. To investigate whether excursion could explain larger discrepancies, a normalised measure of excursion for the target stimulus was calculated for each frame, with zero assigned to the centre and one for the maximum excursion from the centre. The overall data were then divided into two new datasets based on excursion of the target stimulus. The first subset included frames in the first 50% of the dynamic phases – frames with a normalised excursion greater than zero and less than 0.5 ($N = 1746$). The second subset included all frames in the second 50% of the dynamic phases, with a normalised excursion greater than 0.5 and less than 1 ($N = 1573$). The analysis of the aggregate Cohen's k scores for the first half of the movement phases found that there was no agreement between the AS and MS ($k = .153$, 95% CI [.098, .207]). There was a weak level of agreement for the second half of the movement frames ($k = .485$, 95% CI [.450, .521]), but the results were still below the $k = 0.6$ cut established earlier. To investigate whether agreement increases at more extreme levels of excursion, the analysis was run on a subset of the last 75% of frames – with a normalised excursion value which was greater than 0.75 and below 1 ($N = 741$). At this level of excursion there was moderate agreement between the two scorers ($k = .624$, 95% CI [.575, .674]). As the target moves further from the centre of the screen the level of agreement between the scorers increases from poor to good, which would follow the logic that the more a participant shifts their attention away from the centre, the more discrete eye movements are made and the better

Figure 4. Cohen's k scores for the three subsets of excursion. Frames that correspond to the first half of the dynamic phase are Red, the second half of the dynamic phase are green, and the final 25% of the phase is in blue. In most cases it appears that the further the target stimulus was from the centre, the higher the level of agreement. Although in two cases (participants 010 and 015) it appeared that agreement was higher in the first portion of the dynamic phase.

these are detected, in agreement, by both the MS and AS. Therefore, both the MS and the AS are more likely to detect gaze direction changes. The inter-rater reliability scores are plotted by participant for all three subsets of the data in (Figure 5).

While the moving target frames were only achieving a fair level of agreement between the MS and AS, the static frames were able to achieve near perfect agreement across all participants. Overall, these results demonstrate that GazeScorer can effectively score gaze direction across three locations when the target stimulus is static. However, when the target stimulus was moving, the level of agreement fell below the threshold for good agreement. Agreement between the two scorers improved the further the target stimulus was away from the centre, but with the current data it is not possible to establish performance at different levels of excursion as the impact of eye movements on image quality may reduce agreement.

Study 2 – Child sample

Methods

Participants

Parents were sent a separate link to the information and consent prior to receiving the full experimental link to make sure they were informed about how their identifying data (i.e., video footage) would be collected and turned into research data (numerical data points) before storage. Once fully informed consent was received, an appointment was made with the researcher for a data collection session and an experiment link was sent to the participant. Ethical approval for the study was granted by the Interdivisional Research Ethics Committees of the University of Oxford (REF: R73342/RE001).

Data were collected from a sample of 20 children (six female). Their ages ranged from 3 to 12 years old, with a mean age of 6.15 years. Of the 20 participants, data collected for four participants were rejected due to poor video footage. In these cases, three of the participants' faces were either partially or entirely out of the camera frame during the initial calibration frames, interfering with the ability to accurately detect the face and eye movements. One of the participants had excessive reflections in their glasses that impacted the ability to detect the eye regions for a substantial number of frames. The footage from the remaining $N = 16$ (five female; mean age = 6.38; age range 3-6 years) was adequate for both manual and automatic scoring.

Procedure

In the experimental session, parents and guardians were instructed to access Google Chrome (Chrome) and to ensure their privacy settings allowed them to share their screen with researchers. Parents were contacted through a video call to monitor the participant's progress. This method also allowed the researcher to control the participants' computer remotely if required, which allowed researchers more independent control of the study, although the parent or guardian were asked to stay in the same room for child safeguarding reasons. Once the researcher, participant and guardian were ready, they accessed the Gorilla hosted experiment via a link to open on the participant's computer. The researcher walked the participant through child assent, reading to them even if they could read along as well. The participant was then walked through a setup stage with instructions, including picture examples, showing how the participant should sit within the frame and minimise inconsistent lighting conditions. Once the participant was positioned correctly, they continued

to the task presentation. The video stimuli were the same as in the adult study. However, the video was presented twice to the children to ensure a good recording was achieved. The first video was used to minimise training effects unless excessive head movements during the initial recording limited analysis, which occurred in four of the 17 participants. While the majority of participants were able to comply with instructions, participants younger than 5 years were sat with their parent to encourage them to stay still and focus on the task.

Results

The average orientation score across all the participants was plotted for both the automatic and the algorithmic scorers alongside the estimated location of the target stimuli. As with the adult participants, both the AS and MS demonstrated general agreement in the direction of the gaze during the static phase. As with the adult participants, the data were split into two subsets based on the phase of the target stimulus (static or dynamic). However, unlike in the adult participants, it appears that during the static phase the agreement was attenuated during the initial gaze orientation movement towards the left, but not during the subsequent movement to the right. During the dynamic phase, a similar pattern is seen as in the adult participants. The MS appears to make more defined gaze orientation judgments resulting in a steeper slope when compared to the AS, which demonstrates a more gradual slope, resulting in a lower level of agreement in the Dynamic phase.

Figure 5. Grand average plots for the sample $N = 16$. The purple line represents the automatic scorer, and the yellow lines represent the manual scorer. The SE around the M is plotted for each point. The green line marks the location of the target stimuli during the static phase.

Analysis of static phase

A series of Cohen's k tests were conducted individually for the left, right, and overall scores, entering the aggregate data across all the participants. Each frame was treated as an individual subject ($N = 3723$). There was a moderate agreement between the two scorers for the left eye ($k = .830$, 95% CI [.811, .850]) and strong level of agreement for the right eye ($k = .811$, 95% CI [.791, .831]) and the overall judgments ($k = .850$, 95% CI [.833, .867]). Inspection of the confusion matrix of the overall judgments found that, in total, the MS and AS agreed on 3480 of 3723 frames. Of the frames where there were disagreements between a centre and directional gaze judgment, 132 occurred on left gaze judgments and 111 occurred in right gaze judgments. Unlike in the adult group, there were three frames where the MS and the AS detected opposing gaze orientations, with the MS detecting a left gaze and the AS detected a right gaze. Cohen's k tests were run on individual participants and were plotted in Figure 7. While all participants in the adult group demonstrated at least a single point approaching near perfect agreement, it is not the case with four of the participants in the children's sample. Of these four, one achieved above the $k = 0.6$ cut off and would be considered as having acceptable agreement, but the remaining three all fell below the acceptable minimum level. In addition, for five participants the overall score did not correspond with the optimal eye when compared to the manual scorer. Of these five: Two participants showed relatively comparable scores between the left and right eye and were above the acceptable minimum agreement score. Two participants' alternative eye was approaching near perfect agreement with the MS. Only one participant demonstrated a mismatch between the overall score and the optimal eye. Possible reasons for these differences will be considered in more detail in the discussion.

Figure 6. Cohen's k scores for the child sample ($N = 16$). The whiskers represent the SE around the M . AS with the adult sample, all participants have at least one point that is above the minimum agreement ($k = 0.6$, represented by the dashed line). While many participants show near perfect agreement (greater than 0.9) with at least one point, there are a number of participants who do not (PID: 005, 009, 013, 016).

Analysis of dynamic phase

Analyses of the dynamic phase frames were conducted for the left, right and overall scores for the aggregate data ($N = 3025$). There was a weak level of agreement for the overall judgments ($k = .488$ 95% CI [.462, .514]), left judgments ($k = .469$, 95% CI [.442, .496]), and right judgments ($k = .489$, 95% CI [.444, .516]; see Figure 8). Inspection of the confusion matrix of the agreement found that 2045 of the 3025 frames were in agreement. There was disagreement between a centre and directional judgment for 974 frames (532 for the left eye and 442 for the right eye). For judgments of opposing direction, there were 5 frames in which the MS detected a left gaze, and the AS detected a right gaze, and 1 frame in the opposite direction.

Figure 7. Cohen's k scores for the child sample during the dynamic phase. The whiskers represent SE around the M . As in the adult sample performance was significantly reduced when the target stimulus was moving.

To investigate the impact of excursion of the target stimulus on these results the Overall data were split into three subsets to match the subsets in the adult sample. The first consisted of frames where the excursion of the target was less than 50%. The second consisted of frames where the target was above 50% excursion. Finally, a third subset of frames involved a target excursion above 75%. Analyses of these subsets found that agreement between the MS and the AS was minimal in the subset of frames with the target excursion below 50% ($N = 1479$; $k = .307$, 95% CI [.256, .359]). For the subset of excursion above 50% the agreement between the scorers was greater, but still minimal ($N = 1330$; $k = .460$, 95% CI [.421, .499]). Finally, when frames with target excursion over 75% were analysed, the level of agreement rose but still only demonstrated weak agreement ($N = 625$; $k = .530$, 95% CI [.473, .587]). The Cohen's k scores for individual participants for all subsets can be found in Figure 9. Unlike with the adult participants, there does not appear to be a consistent increase in agreement as the target moves further from the central location. This

may be reflective of differences between adults and children when making overt attentional shifts, which will be considered in more detail in the discussion.

Figure 8. Cohen's k scores for the three subsets of excursion. Frames that correspond to the first half of the dynamic phase are red, the second half of the dynamic phase are green, and the final 25% of the phase is in blue. The majority of cases demonstrated progressively better agreement as the target stimulus moved further from the centre.

Although there is no direct comparison between the two samples, it appears that the automatic gaze scoring of child participants is similar to that of the adult participants, albeit with slightly attenuated agreement between scorers. Good agreement during the static subset of frames is promising for experiments that would require static stimuli. There are more stark differences during the movement phases that highlight issues that may impact younger participants more than adult participants. These results are promising for the use of GazeScorer with children as well as adults, although more data is required from children of specific age groups to establish what impact age has on the differences between children and adults.

Discussion

Online gaze orientation detection: Automatic gaze scoring via image processing

Our approach attempts to move away from current browser-based methods of eye tracking, that aim to replicate the abilities of specialist in-lab equipment, and instead focus on classifying gaze orientation in a reliable and user-friendly way. At the same time, our approach expands on an established method of manual scoring in a way that reduces the labour required for the technique. This automated approach differs from current alternatives because it focuses on algorithmic image processing, rather than on identifying gaze position co-ordinates through complex regression analysis and interactions with mouse movements. We do not believe that the approaches available to leverage webcam data are to be compared in absolute terms, but rather in terms of how they are best suited for different purposes. In particular, we believe that our distinct approach could apply to multiple experimental questions and target participant samples, and it will be particularly powerful for researchers who aim to collect and analyse out-of-lab eye-direction information from young or neurodivergent participants who are least likely to travel to the lab and / or comply to long calibrations or extended verbal instructions and mouse-tracking demands. Our key goals, related to the three key desirable features of webcam based gaze data collection we identified in our introduction, were: 1) to test the ability of GazeScorer to accurately distinguish the directional gaze of participants while they followed a target stimulus, with reasonable spatial resolution; 2) to remove the need for excessively long and involved calibration procedures prior to data collection; 3) to capture overt attention direction in participants of different ages. Having identified the central gaze direction, we conducted a series of comparisons between the AS and the MS, naïve to the position of the stimulus on screen. We collected data from two samples of participants: The first was an adult sample with participants between the ages of 21 and 32. The second was a preliminary sample of children aged between 3.5 and 12.

The two age groups demonstrated comparable results during the static phase and GazeScorer displayed good to near perfect levels of agreement in gaze direction judgements with our MS in this phase, achieving good spatial resolution for simple left-middle-right judgements. Indeed, many researchers measuring gaze in experimental conditions focus their work on fixations on static stimuli, so we investigated the accuracy of GazeScorer for frames in which the target stimulus was static, separately from those when the stimulus was moving. For the static stimuli, it was found that there was near perfect agreement between GazeScorer and the MS across frames. There was variation in agreement between GazeScorer and the MS for the individual participants' scores, but no participant fell below the established minimum for moderate agreement in the adult

participants, and only three of the 16 children were below the minimum. In addition, for moving frames with the largest excursion, agreement between the MS and GazeScorer was moderate for adults. As a whole, these findings are promising and positive for researchers who are interested in capturing fixations and for target regions of interest presented on the horizontal dimension on participants' home screen with a reasonable degree of spatial resolution. With regards to our second key feature, i.e., the need to avoid long, complicated and verbally demanding calibrations, GazeScorer performed well for both children and adults, with a brief and friendly set of initial positioning / lighting instructions that could also be monitored remotely by the experimenter during testing sessions. Finally, we captured participants' overt attention to target stimuli with ease, across age groups.

In summary, gaze direction was very well identified by GazeScorer for static frames and the most peripheral moving frames acquired through a browser and with a variety of webcams, converging with judgements by a MS, but with a much-reduced time commitment. This suggests great promise for using GazeScorer as a scoring package for static and peripherally presented stimuli online, collected with minimal calibration requirements and no requirement to use a mouse for calibration. This approach is by no means the only approach for remote online gaze data collection, but our promising findings suggest that it offers an alternative method. Here we will consider some of the limitations of the automatic scoring algorithm and possible future solutions, as well as limitations associated with browser-based video data collection. We will then discuss our results in relation to a number of these alternative approaches to understand what GazeScorer adds to the current resources available to researchers. Finally, we will conclude by discussing the future directions for GazeScorer.

Current Limitations of GazeScorer

We discovered a number of limitations to GazeScorer in its current form. First, for the dynamic target phase, the level of agreement was moderate across all frames. Again, there was variation across participants, but for most participants at least one of the eyes approached a good level of agreement, but this was influenced by the degree of excursion and by the age of the participants. For adults, the more extreme positional subset of frames (beyond 75% of moving frames) yielded an acceptable level of agreement. However, in the children's sample, there was not such a clear-cut example of excursion resulting in improved agreement between the two scorers. In addition, there were some disparities between data collected from adults and children, especially during the dynamic phase. The overall gaze direction judgment measure was generally a robust indicator but failed under certain circumstances. Finally, collecting participants' data as videos introduced ethical implications

that must be considered very carefully. We tackle each of these limitations and suggestions for improvement in turn.

Poor agreement during the dynamic phase

As a participant shifts their attention to follow a stimulus, they make overt eye movements to maintain foveal attention on the target. With lab-based eye trackers utilising IR technology, fine movements can be detected from the retinal reflection. Unlike eye trackers, GazeScorer estimates gaze location by identifying the position of the iris within the eye. This is complicated when the eye is moving, as video capture may then result in motion-blur as images are saved from the webcam feed. That is, because the webcam will capture frames during an eye movement, the method and time taken to encode each image may result in a blurring effect if the eye is actively moving when the frame is captured. This small amount of blur is unlikely to affect a manual scorer when making their decisions, but at a pixel level it may result in poor performance. This could be improved by adapting the iris detection algorithm in future iterations. However, this raises the important question of what a user would want from the GazeScorer software. Many researchers who would utilise the software are primarily interested in fixations and therefore would not necessarily use frames during periods of excessive eye movements. As such, strong agreement found in all participants during the static phase shows promise for research where the primary focus is during fixation periods.

Requirement of wide excursion of target stimuli

Another factor that impacted the quality of the agreement between GazeScorer and the manual scorer was the excursion of the target from the centre. In our adult data, acceptable agreement was only found when the target was over a 75% excursion. This was because the further the target moved from the initial central point, the more the eye had to physically move. As the eye moves away from the centre, the middle point of the detected iris also moves further from the established centre gaze location and therefore, it is more likely that the gaze would be detected as moving outside of the central fixation boundary. Our current research can help guide where users should place their stimuli in relation to the centre in any future experiments. There are other factors that may also impact the level of excursion, such as the size of the monitor, the distance sat from the monitor, or how eye motion may impact our current findings. In our study, participants used their own home computers and were instructed to sit at a comfortable distance, and we achieved good performance in all participants. In future, more research will be needed to establish how screen size and distance impact GazeScorer's ability to detect gaze orientation. In addition, our results on excursion are based on the dynamic phase and therefore may be impacted by

motion blur discussed above. Research that establishes performance at different excursion points while the target remains static will also help establish thresholds for stimuli placement.

In the children's sample, unlike the adult sample, the relationship between excursion and agreement during dynamic phases was not as clear cut. Our results suggest that the disparity is unlikely to simply be the result of structural, facial differences between children and adults. That is, the face detection libraries and the custom iris detection algorithms were able to operate with children as young as 3 years, and therefore any differences between the groups are likely to be the result of behaviour during video recording - for example, more head movement or poorer positioning within the webcam frame. This is evident in comparisons between levels of excursion during the Dynamic phase of the target stimuli in the children's data. In the adult group there is improvement in agreement at wider levels of target excursion, but such clear improvement is not seen in the children's data. As it is unlikely that it is physical differences, it is still advisable to place targets at a wider excursion to amplify eye movements. A more systematic approach is necessary to identify any behaviours in children that may impact the agreement between the scorers, and propose strategies on how to mitigate or minimise their impact on GazeScorer in future.

Accuracy of the unifying gaze measure

In the current experiment, the automatic gaze detection software produced a unifying gaze measure, referred to as the overall judgment. This unifying measure attempts to give a classification utilising information from both eyes to provide a single measure. Currently it works in a relatively simplistic way, by selecting the eye that makes an abduction movement (movement away from the midline, e.g., a left eye makes a left look). During development it was found that this was the most reliable eye movement for detection. However, in some instances this failed if an abduction orientation was not identified. Arguably this option is not required as both eye outputs are available to researchers to make their own decisions on how to utilise each eye, but as GazeScorer develops it may be necessary for future applications where post-hoc decisions cannot be used to decide the best eye to use. Therefore, it will be necessary to refine the unifying measure using more sophisticated inputs. For example, utilising luminosity measures, dlib's face detection confidence outputs, and temporal vs nasal eye moments to weight which eye is selected more accurately.

Collection and storage of online video data

In its current form, GazeScorer works with videos recorded during an experimental session. While this is commonly used within research, it has its own ethical considerations that can require significant ethical and security implications. Specifically, a key principle of modern psychological research is the right for participants to remain anonymous, which is

simply not possible when recording videos. Therefore, it is important that researchers take significant precautions to protect the participant and make sure they are comfortable with the experimental set up. In addition to the ethical implications, there are also legal concerns regarding storage of identifiable data that require more considerations than anonymised behavioural response data. Recent changes in data protection law, specifically the implementation of General Data Protection Regulation (GDPR) in the UK and Europe, put an extra layer of data protection responsibility on researchers who are working with personally identifiable data. Of note, developmental psychologists have used video recordings as part of psychological research for decades, and as such have well established pipelines to maintain the integrity and security of such data. Furthermore, retaining video data has additional benefits that may not be immediately obvious. For example, it provides the opportunity for confirmatory manual checks after any automatic scoring occurs; video recordings do not require any specialist kit or a dedicated internet connection, so research can be conducted in any location; and it is possible to record videos in relation to non-computer displayed stimuli as no involved calibrations are required. These options make recorded video data ideal for developmental research where retaining all trials is important but may also provide advantages to other populations.

Comparison to current online eye tracking methods

This approach is by no means the only approach for remote online gaze data collection, but our promising findings suggest that it offers an alternative approach. Previously we discussed a number of methods that attempt to replicate eye tracking data in relation to three desirable factors that would be important for remote eye tracking data: spatial resolution, participant usability, and whether or not the data reflect a true measure of overt attention. The two software solutions discussed, Webgazer.js and MouseView.js both offered the opportunity for collecting high-spatial resolution gaze data but were limited by usability issues among some participant samples in psychological paradigms and the inability to measure overt eye movement, respectively. In addition, we considered evidence that videos that were collected over the internet using conventional webcams could be used for manual coding of gaze orientation. These findings provide an alternative solution to gaze scoring that, while still only meeting two of the three established factors (participant usability and overt attention indices, while limited in spatial resolution), offered wider accessibility for researcher where the current software solutions are not appropriate - laying the groundwork for our software approach.

Here, we have demonstrated an automated analysis pipeline that was able to replicate the abilities of a manual scorer. This is comparable to work by Chouinard et al.

(2019), which used Amazon Web Services Rekognition software alongside a custom classifier to detect gaze orientation. The use of Amazon Web Services may raise some data privacy concerns among researchers and participants, as it would require processing and storing the data remotely. In addition, not all elements of this pipeline are open source elements. Instead, our approach offers the ability to analyse data locally using Python, and using open source and Open Science principles. However, based on the current data, it is difficult to compare the quality of two approaches directly. While we used data collected with a task designed to collect ideal data for scoring gaze orientation, Chouinard and colleagues used video recordings from an active infant preferential looking experiment that was not optimised a priori for automated scoring. Furthermore, Chouinard's sample consisted of children between 24 and 36 months, whereas among our sample of children the ages ranged from three to 12 years. More work will need to be done to make a direct comparison, but the relative success of these approaches, with both adults and children, confirms that it is possible to get reliable gaze orientation data without involved calibrations. And while the loss of spatial resolution is not suitable for some, it may represent a functional reality of collecting eye gaze data remotely.

Together, these three approaches (WebGazer, MouseView.js and GazeScorer) to measuring attention can provide an array of choices for researchers. Each method comes with their individual pros and cons, and it is important that researchers bear this in mind when designing and building their experiments. If a researcher requires a high precision measure, then WebGazer or MouseView.js will offer that, but it is important that they work within the limitations of these methods. GazeScorer offers an option for those who do not need a high level of spatial precision but require minimal calibration to ensure participant comfort and to maintain engagement. This would help to reduce participant attrition seen in current studies and make these experiments more accessible for research working with children or neurodivergent participants.

Conclusions and future directions

GazeScorer has demonstrated very good agreement with a manual scorer in optimal conditions (i.e., a static target stimulus with a wide excursion from the centre) - the question now remains whether this approach is a viable alternative for other methods of webcam eye tracking. Ultimately this depends on the requirements of the researcher. In research with young children where preferential looking paradigms are used regularly, the established constraints on spatial resolution, static targets, and wide excursion, are an appropriate trade off to avoid difficult calibration procedures. In the short term we will work to develop GazeScorer as a hybrid automatic-manual analysis that can offer a superior approach to either one or the other independently, by offering the ability for unbiased analysis of all frames, with only minimal confirmatory manual intervention. Moving beyond the hybrid

system, future development will focus on improving the gaze detection algorithm to optimise it for use with live-webcam inputs. This will allow us to transition GazeScorer to fully automatic gaze orientation detection, which will hopefully be made available to integrate alongside current browser-based experimental packages.

When working with neurotypical adults, researchers who are used to the fine precision associated with lab-based eye trackers may need to consider their options more carefully. However, it is important to remember that aiming for very precise spatial resolution is only viable if the data collection method can demonstrate good validity and reliability in active research environments. In the “wild”, i.e., when collecting data out of the lab and online, GazeScorer provides proof-of-concept for a viable alternative to mouse tracking-based methods when high spatial precision recording is not required. Future development and testing are required to optimise the algorithm, the usability of the software, and establish its limitations. In addition, we are committed to providing the software open source to allow researchers to build on and adapt the software as is necessary for their requirements, in full alignment with Open Science practices.

Acknowledgments

We are grateful to all children, young people and families who took part in our study.

GazeScorer development was supported by a Health Research Bridging Salary Scheme. - Bridging Salary Award to AF, a COVID-19 Rebuilding Research Momentum Fund Award to AF, a John Fell Oxford University Press Fund Award and Leverhulme Trust Project Grant (RPG-2022-072) to GS.

Open Science Statement

As requested by Behavior Research Methods, we have made the associated code available through an Open Science Framework (OSF) Project (<https://osf.io/kr8d7/>). The code provided demonstrates the analysis pipeline for extracting gaze orientation estimates from three participants who have provided consent for the use of their videos. In addition, we provide the fully anonymized data for our comparisons between the automatic and manual scorers.

References

- Álvarez Casado, C., & Bordallo López, M. (2021). Real-time face alignment: Evaluation methods, training strategies and implementation optimization. *Journal of Real-Time Image Processing*, 0123456789. <https://doi.org/10.1007/s11554-021-01107-w>
- Anwyl-Irvine, A., Dalmaijer, E. S., Hodges, N., & Evershed, J. K. (2021). Realistic precision and accuracy of online experiment platforms, web browsers, and devices. *Behavior Research Methods*, 53(4), 1407–1425. <https://doi.org/10.3758/s13428-020-01501-5>
- Anwyl-Irvine, A. L., Armstrong, T., & Dalmaijer, E. S. (n.d.). *MouseView.js: Reliable and valid attention tracking in web-based experiments using a cursor-directed aperture*. <https://www.mouseview.org/>
- Anwyl-Irvine, A. L., Massonnié, J., Flitton, A., Kirkham, N., & Evershed, J. K. (2019). Gorilla in our midst: An online behavioral experiment builder. *Behavior Research Methods*. <https://doi.org/10.3758/s13428-019-01237-x>
- Bradski, G. (2000). The opencv library. *Dr Dobb's J. Software Tools*, 25, 120–125.
- Chouinard, B., Scott, K., & Cusack, R. (2019). Using automatic face analysis to score infant behaviour from video collected online. *Infant Behavior and Development*, 54(September 2018), 1–12. <https://doi.org/10.1016/j.infbeh.2018.11.004>
- Dalrymple, K. A., Manner, M. D., Harmelink, K. A., Teska, E. P., & Elison, J. T. (2018). An Examination of Recording Accuracy and Precision From Eye Tracking Data From Toddlerhood to Adulthood. *Frontiers in Psychology*, 9, 803. <https://doi.org/10.3389/fpsyg.2018.00803>
- de Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, 47(1), 1–12. <https://doi.org/10.3758/s13428-014-0458-y>
- Hauger, D., Paramythis, A., & Weibelzahl, S. (2011). Using Browser Interaction Data to Determine Page Reading Behavior. In J. A. Konstan, R. Conejo, J. L. Marzo, & N.

Oliver (Eds.), *User Modeling, Adaption and Personalization* (Vol. 6787, pp. 147–158).

Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-22362-4_13

Henninger, F., Shevchenko, Y., Mertens, U. K., Kieslich, P., & Hilbig, B. (2019). *lab.js: A free, open, online study builder*. 6. <https://doi.org/10.31234/osf.io/fqr49>

King, D. E. (2009). Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10, 1755–1758.

McHugh, M. L. (2012). Lessons in biostatistics interrater reliability: The kappa statistic. *Biochemica Medica*, 22(3), 276–282.

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, 51(1), 195–203. <https://doi.org/10.3758/s13428-018-01193-y>