

Journal of Visualized Experiments

A Methodology for Capturing Joint Visual Attention Using Mobile Eye-trackers

--Manuscript Draft--

Article Type:	Invited Methods Article - Author Produced Video
Manuscript Number:	JoVE60670R2
Full Title:	A Methodology for Capturing Joint Visual Attention Using Mobile Eye-trackers
Section/Category:	JoVE Behavior
Keywords:	Educational Research, Collaborative Learning, Mobile Eye-tracking, Joint Visual Attention, Multimodal Learning Analytics.
Corresponding Author:	Bertrand Schneider Harvard University Cambridge, UNITED STATES
Corresponding Author's Institution:	Harvard University
Corresponding Author E-Mail:	bertrand_schneider@g.harvard.edu
Order of Authors:	Bertrand Schneider
Additional Information:	
Question	Response
Please indicate whether this article will be Standard Access or Open Access.	Standard Access (US\$1200)

TITLE:**A Methodology for Capturing Joint Visual Attention Using Mobile Eye-Trackers****AUTHORS AND AFFILIATIONS:**

Bertrand Schneider¹

¹Learning, Innovation, and Technology Lab, Graduate School of Education, Harvard University

Corresponding Author:

Bertrand Schneider (bertrand_schneider@gse.harvard.edu)

KEYWORDS:

educational research, collaborative learning, mobile eye-tracking, joint visual attention, multimodal learning analytics

SHORT ABSTRACT:

Using multimodal sensors is a promising way to understand the role of social interactions in educational settings. This paper describes a methodology for capturing joint visual attention from colocated dyads using mobile eye-trackers.

LONG ABSTRACT:

With the advent of new technological advances, it is possible to study social interactions at a microlevel with unprecedented accuracy. High frequency sensors, such as eye-trackers, electrodermal activity wristbands, EEG bands, and motion sensors provide observations at the millisecond level. This level of precision allows researchers to collect large datasets on social interactions. In this paper, I discuss how multiple eye-trackers can capture a fundamental construct in social interactions, joint visual attention (JVA). JVA has been studied by developmental psychologists to understand how children acquire language, learning scientists to understand how small groups of learners work together, and social scientists to understand interactions in small teams. This paper describes a methodology for capturing JVA in colocated settings using mobile eye-trackers. It presents some empirical results and discusses implications of capturing microobservations to understand social interactions.

INTRODUCTION:

JVA has been extensively studied over the last century, especially by developmental psychologists studying language acquisition. It was quickly established that joint attention is more than just a way to learn words but rather a precursor to children's theories of mind¹. Thus, it plays a significant role in many social processes, such as communicating with others, collaborating, and developing empathy. Autistic children, for instance, lack the ability to coordinate their visual attention with their caregivers, which is associated with significant social impairments². Humans need joint attention to become functional members of society, to coordinate their actions, and to learn from others. From children acquiring their first words, teenagers learning from schoolteachers, students collaborating on projects, and to groups of adults working toward

common goals, joint attention is a fundamental mechanism to establish common ground between individuals³. In this paper, I focus on the study of JVA in educational research. Understanding how joint attention unfolds over time is of primary importance for the study of collaborative learning processes. As such, it plays a predominant role in socioconstructivist settings.

The exact definition of joint attention is still debated⁴. This paper is concerned with a subconstruct of joint attention (JA), namely JVA. JVA happens when two subjects are looking at the same place at the same time. It should be noted that JVA does not provide any information about other important constructs of interest in the study of JA, such as monitoring common, mutual, and shared attention, or more generally, awareness of the cognition of another group member. This paper operationalizes and simplifies JVA by combining the eye-tracking data from two participants and analyzing the frequency in which they align their gazes. For a more comprehensive discussion, the interested reader can learn more about the study of the JA construct in Siposova et al.⁴.

Over the past decade, technological advances have radically transformed research on JVA. The main paradigm shift was to use multiple eye-trackers to obtain quantitative measures of attentional alignments, as opposed to qualitatively analyzing video recordings in a laboratory or ecological setting. This development has allowed researchers to collect precise, detailed information about dyads' visual coordination. Additionally, eye-trackers are becoming more affordable: until recently, their use was reserved to academic settings or large corporations. It is now possible to purchase inexpensive eye-trackers that generate reliable datasets. Finally, the progressive inclusion of gaze-tracking capabilities into existing devices like high-end laptops and virtual and augmented reality headsets suggests that eye-tracking will soon become ubiquitous.

Because of the popularization of eye-tracking devices, it is important to understand what they can and cannot tell us about social interactions. The methodology presented in this paper marks a first step in this direction. I address two challenges in capturing JVA from multiple eye-trackers: synchronizing the data on 1) the temporal scale, and 2) on the spatial scale. More specifically, this protocol makes use of fiducial markers placed in real-world environments to inform computer vision algorithms where participants are orienting their gaze. This new kind of methodology paves the way to rigorous analysis of human behavior in small groups.

This research protocol complies with the guidelines of Harvard University's human research ethics committee.

PROTOCOL:

1. Participant Screening

1.1. Ensure that participants with normal or corrected-to-normal vision are recruited. Because participants will be asked to wear a mobile eye-tracker, they can wear contact lenses but not regular eyeglasses.

2. Preparation for the Experiment

2.1. Eye-tracking devices

2.1.1. Use any mobile eye-tracker capable of capturing eye movement in real world environments.

NOTE: The mobile eye-trackers used here were two Tobii Pro Glasses 2 (see **Table of Materials**). In addition to specialized cameras that can track eye movements, the glasses are also equipped with an HD scene camera and a microphone so that the gaze can be visualized in the context of the user's visual field. These glasses capture gaze data 50 times per second. Other researchers have used ASL Mobile Eye⁵, SMI⁶, or Pupil-labs⁷, all of which provide video streams from the scene camera and eye-tracking coordinates at varying sampling rates (30–120 Hz). The procedure below might vary slightly with other eye-tracking devices.

2.2. Fiducial Markers

2.2.1. The two steps below (i.e., temporal and spatial alignments) require the use of fiducial markers. There are several computer vision libraries that provide researchers with these markers and algorithms to detect them on an image or video feed. The protocol described uses the Chilitag library⁸.

2.3. Temporal alignment

2.3.1. Because the eye-tracking data are recorded on two separate units, ensure that the data are properly synchronized (**Figure 1**). Two main methods can be used. This manuscript only covers the first method, because server synchronization works differently with each brand of mobile eye-tracker.

2.3.1.1. Briefly show a fiducial marker on a computer screen to mark the beginning and the end of a session. This is similar to a visual "hand clap" (**Figure 2**).

2.3.1.2. Alternatively, use a server to synchronize the clocks of the two data collection units. This method is slightly more accurate and recommended if a higher temporal accuracy is required.

2.4. Spatial alignment

2.4.1. To find if two participants are looking at the same place at the same time, map their gazes to a common plane. This plane can be a picture of the experimental setting (see the left side of **Figure 3**). Carefully design this image before the experiment.

2.4.2. Size of the fiducial markers: The general size of the fiducial markers depends on the algorithm used to detect them from the eye-tracking video. Surfaces close to the participants can

have smaller fiducial markers, while surfaces further away from them need to be larger, so that they look similar from the participants' perspective. Try different sizes beforehand to make sure that they can be detected from the eye-tracking video.

2.4.3. Number of fiducial markers: To make the process of mapping gaze points into a common plane successful, make sure to have several fiducial markers visible from the participants' point of view at any given time.

2.4.4. Location of the fiducial markers: Frame relevant areas of interest with strips of fiducial markers (e.g., see the laptop screen on **Figure 3**).

2.5. Finally, run pilots to test the synchronization procedure and determine the optimal location, size, and number of fiducial markers. Eye-tracking videos can be processed through a computer vision algorithm to see if the fiducial markers are reliably detected.

3. Running the experiment

3.1. Instructions

3.1.1. Instruct participants to put on the eye-tracking glasses as they would a normal pair of glasses. Based on the participants' distinct facial features, nose pieces of different heights may need to be used to preserve data quality.

3.1.2. After turning on the eye-tracker, have the participants clip the recording unit to themselves to allow for natural body movement.

3.2. Calibration

3.2.1. Instruct the participants to look at the center of the calibration marker provided by Tobii while the calibration function of the software is enabled. Once the calibration is complete, recording can be started from within the software.

3.2.2. Instruct participants to not move the mobile eye-trackers after calibration. If they do, the data are likely to be inaccurate and the calibration procedure will need to be performed again.

3.3. Data monitoring

3.3.1. Monitor the data collection process during the study and ensure that the eye-tracking data are being collected properly. Most mobile eye-trackers can provide a live stream on a separate device (e.g., a tablet) for this purpose.

3.4. Data export

3.4.1. After the recording session is complete, instruct the participant to remove the eye-tracking

glasses and the data collection unit. Turn off the unit.

3.4.2. Extract data using another software, Tobii Pro Lab, by removing the SD card from the data collection unit importing the session data. Tobii Pro Lab can be used to replay the video, create visualizations, and export the eye-tracking data as comma-separated (.csv) or tab-separated (.tsv) files.

4. Preprocessing the dual eye-tracking data

4.1. Sanity checking eye-tracking data

4.1.1. Check the eye-tracking data visually after data collection. It is not uncommon for some participants to have missing data. For example, some particular eye physiology can pose problems to eye-tracking algorithms, the glasses might shift during the experiment, the data collection software might crash, etc.

4.1.2. Use descriptive statistics to check how much data were lost during each session and exclude sessions that have significant amounts of missing or noisy data.

4.2. Temporal alignment

4.2.1. Trim the data from each mobile eye-tracker to only include interactions between the participants. This can be achieved by using the method described above (i.e., presenting two special fiducial markers to participants at the start and the end of the session). These fiducial markers can then be detected from the eye-tracking video to trim the datasets.

4.3. Spatial alignment

NOTE: To detect whether two participants are looking at the same place at the same time, it is necessary to remap the participants' gaze onto a common plane (i.e., an image of the experimental setting). A computational method for achieving this goal is a homography (i.e., a perspective transformation of a plane). From a technical perspective, two images of the same planar surface in space are related by a homography matrix. Based on a common set of points, this matrix can be used to infer the location of additional points between two planes. In **Figure 3**, for example, if a computer vision algorithm knows where the fiducial markers are on the handout, it can remap the gaze of the participant onto the common plane on the left side. The white lines connect the two sets of points shared by the video feed of each participant and the scene, which are then used for building the homography to remap the green and blue dots on the left side.

4.3.1. Use the Python version of OpenCV, e.g., to compute the homography matrix from the fiducial markers and then to remap the eye-tracking data to the scene of the experimental setting (or any other suitable library in your language of choice). OpenCV provides two useful functions: **findHomography()** to get the homography matrix, and **perspectiveTransform()** to transform the

point from one perspective to the other.

4.3.2. To use **findHomography()**, run with two arguments: the X,Y coordinates of the source points (i.e., the fiducial markers detected from the participants' scene video, shown on the right in **Figure 3**) and the corresponding destination points (i.e., the same fiducials markers detected on the scene image, shown on the left in **Figure 3**).

4.3.3. Feed the resulting homography matrix into the **perspectiveTransform()** function, along with a new point that needs to be mapped from the source image to the destination image (e.g., the eye-tracking data shown as a blue/green dot on the right side of **Figure 3**). The **perspectiveTransform** function returns the new coordinate of the same point on the scene image (i.e., the blue/green dots shown on the left side of **Figure 3**).

NOTE: For more information, the OpenCV official documentation provides sample code and examples to implement the homography: docs.opencv.org/master/d1/de0/tutorial_py_feature_homography.html.

4.4. Sanity checking the homography

4.4.1. Complete section 4.3 for the entire session, and perform a homography on each frame of the mobile eye-tracking video to check the quality of the homography. While there are no automated ways to estimate the accuracy of the resulting eye-tracking data, videos like the one shown in **Figure 4** should be used to manually sanity check each session.

4.4.2. If the quality is lower than expected, consider additional parameters to improve the results of the homography:

4.4.2.1. Number of fiducial markers detected: Only perform the homography if enough fiducial markers can be detected from the video stream. This number can be determined by examining the video produced above.

4.4.2.2. Location of the fiducial markers: If different markers are at different depths and orientations, the quality of the homography usually increases when the markers closest to the gaze coordinates are selected, given that there are enough markers to build a robust homography.

4.4.2.3. Orientation of the fiducial markers: Combining fiducial markers that have different orientations (e.g., horizontal and vertical) will produce inaccurate homographies. It is recommended to first detect which plane or areas of interests (AOIs) the participant is looking at (e.g., the computer screen, the cheat sheet, the table, see **Figure 3**) and then use the fiducial markers on this plane for the homography.

4.4.2.4. Quality of the video stream: Sudden head movements can blur video frames and make the data unusable, because fiducial markers cannot be reliably detected (**Figure 4**). The

methodology of this paper is not appropriate for experiments that involve a lot of sudden head movements.

5. Analyzing the dual eye-tracking data

5.1. Missing data

5.1.1. In order to make sure that the data were properly remapped onto the reference image, produce visualization graphs (e.g., **Figure 5**, **Figure 6**) and descriptive statistics to check how much data are missing.

5.2. Cross-recurrence graphs

5.2.1. Use cross-recurrence graphs⁹ to represent visual synchronization between two participants (**Figure 6**), where the X-axis represents time for the first participant, and the Y-axis represents time for the second participant. Black squares indicate that participants are looking at the same area, a black diagonal line describes two subjects looking at the same thing at exactly the same time, and black squares off the diagonal line describes when two subjects looking at the same thing with a time lag. Finally, differentiating between missing data (white square) and existing data with no JVA (gray squares) helps identify problematic sessions. This provides researchers with a visual sanity check.

5.3. Computing JVA

5.3.1. After filtering for missing data, compute a metric for JVA by counting the number of times that participants' gazes are in the same radius in the scene (defined below) in a -2/+2 s time window. Divide this number by the number of valid data points that can be used to compute JVA. The result of the division represents the percentage of time that two subjects were jointly looking at the same place. This last step is necessary to avoid inflating the scores of groups with more data after the homography.

NOTE: Two parameters need to be set before JVA can be computed, the minimal distance between two gaze points, and the time window between them (**Figure 7**): 1) Time window: An early foundational study¹⁰ used a single eye-tracker to measure JVA between a listener and a speaker. The researchers asked a first set of participants ("speakers") to talk about a television show whose characters were displayed in front of them. A second set of participants ("listeners") then watched the same show while listening to audio recording of the speakers. The eye movements of the speakers and listeners were compared, and it was found that a listener's eye movements closely matched a speaker's eye movement with a delay of 2 s. In subsequent work¹¹, researchers analyzed live dialogues and found that a delay of 3 s best captured moments of JVA. Because each task is unique and might exhibit different time lags, it is also suggested to explore how different time lags affect the results of a given experiment (see Schneider et al.¹², for example). Overall, it is common to look for JVA in a $\pm 2/3$ s time window depending on the experimental task and then explore how different time lags might change the results. 2) Distance

between gazes: there is no empirically defined distance between two gazes for them to count as JVA. This distance is dependent on the research questions defined by the researchers. The research questions should inform the size of the targets of interest. In the example seen in **Figure 7**, a radius of 100 pixels on the scene image (blue/green circles) was chosen for the analysis because it is sufficient to capture when participants are looking at the robot in the maze, as well as at similar user interface elements on the computer screen, which are the two main areas of interest for this experimental task.

REPRESENTATIVE RESULTS:

The methodology presented above was used to study students who were following a vocational training program in logistics ($n = 54$)¹². In this experiment, pairs of students interacted with a Tangible User Interface (TUI) that simulated a small-scale warehouse. The fiducial markers placed on the TUI allowed the research team to remap students' gazes onto a common plane and compute levels of JVA. Findings indicated that groups who had higher levels of JVA tended to do better at the task given to them, learned more, and had a better quality of collaboration¹³ (**Figure 8**, left side). Dual eye-tracking datasets also allowed us to capture particular group dynamics like the free-rider effect. We estimated this effect by identifying who was likely to have initiated each moment of JVA (i.e., whose gaze was there first) and who responded to it (i.e., whose gaze was there second). We found a significant correlation between learning gains and the students' tendency to equally share the responsibility of initiating and responding to offers of JVA. In other words, groups in which the same person always initiated moments of JVA were less likely to learn (**Figure 8**, right side) and groups where this responsibility was equally shared were more likely to learn. This finding shows that we can go beyond merely quantifying JV, and actually identify group dynamics and productivity using dual eye-tracking data.

FIGURE AND TABLE LEGENDS:

Figure 1. Each participant generates two video feeds with the X,Y coordinates of their gaze on each video frame: This methodology addresses synchronizing the data temporally and spatially between the participants.

Figure 2. A methodology for synchronizing the two datasets: Briefly showing unique fiducial marker on a computer screen to tag the start and the end of the activity.

Figure 3. Using fiducial markers disseminated in the environment to remap participants' gazes onto a common plan (left side): White lines indicate fiducial markers that have been detected in both images.

Figure 4. Examples of poor data quality: Left: A blurred frame from the eye-tracking video caused by a sudden head movement. Fiducial markers could not be detected in this image. Right: A failed homography where the fiducial marker data were not properly synchronized with the video feed.

Figure 5. Heatmaps: Left: A heatmap of the eye-tracking data remapped onto the experimental scene. This visualization was used as a sanity check for the homography. Right: A group that had too much missing data and had to be discarded.

Figure 6. Cross recurrence graph generated from three dyads to visualize JVA: P1 represents time for the first participant, P2 represents time for the second participant. Black squares show JVA; gray squares show moments where participants are looking at different places; white squares show missing data. Square along the main diagonal indicates moments where participants looked at the same place at the same time. This visualization was used as a sanity check for measures of JVA from the combined eye-tracking data.

Figure 7. A video frame where JVA was detected between two participants (red dots): Richardson et al.¹¹ recommend looking at a time window of ± 2 s. when computing JVA. Additionally, researchers need to define the minimal distance between two gaze points to count as JVA. A radius of 100 pixel was chosen on the middle image above.

Figure 8. Examples of results: Data from Schneider et al.¹² where the percentage of time looking at the same place at the same time was correlated with participants' quality of collaboration: $r(24) = 0.460$, $P = 0.018$ (left side) and imbalances in initiating/responding to offers of JVA was correlated with their learning gains: $r(24) = -0.47$, $P = 0.02$ (right side).

DISCUSSION:

The methodology described in this paper provides a rigorous way to capture JVA in colocated dyads. With the emergence of affordable sensing technology and improved computer vision algorithms, it is now possible to study collaborative interactions with an accuracy that was previously unavailable. This methodology leverages fiducial markers disseminated in the environment and uses homographies as a way to remap participants' gazes onto a common plane. This allows researchers to rigorously study JVA in colocated groups.

This method includes multiple sanity checks that need to be performed at various point of the experiment. Because this is a complex procedure, researchers need to make sure that the resulting datasets are complete and valid. Finally, it is recommended to conduct pilot studies before the actual experiment, and to reconstruct participants' interactions though a video after data collection is completed (**Figure 3, Figure 4, Figure 5, Figure 6**).

There are several limitations associated with this method:

Number of participants. While this methodology works well for two participants, analysis become more complicated with larger groups. Fiducial markers can still be used to remap gazes onto a ground truth but knowing how to identify JVA becomes a more nuanced process. Should JVA be defined as the times when everyone is looking at the same place at the same time, or when two participants are gazing at the same place? Additionally, visualizations like the cross-recurrence graph become impractical with more than 2–3 people.

Settings. The method described in this paper is appropriate for small, controlled settings (e.g., laboratory studies). Open-ended settings, such as outdoors or large spaces, are usually too complicated to instrument with fiducial markers and thus can limit the usefulness of the eye-

tracking data. Additionally, the fiducial markers can be distracting and clutter the environment. In the future, better computer vision algorithms will be able to automatically extract common features between two perspectives. There are already algorithms that exist for this purpose, but we found that the level of accuracy was not yet acceptable for the type of experiment described above.

AOIs. Related to the point above, computing homography and the cross-recurrence graph work well with a stable number of areas of interest, but corrections have to be made when comparing different tasks with different numbers of areas of interest.

Use of equipment. Mobile eye-trackers can be obtrusive, affecting participants' behavior or failing to work with particular eye physiology.

In conclusion, the methodology described in this paper is a promising way to study colocated interactions. It allows researchers to capture a precise metric for JVA, which is a critical construct in the social sciences¹. Additionally, it is possible to detect more fine-grained indicators of collaborative learning through this methodology¹² compared to traditional qualitative analyses. In short, it is a more efficient and accurate way to study social interactions.

Potential application of this method includes designing interventions to support collaboration through real-time eye-tracking data. Some pioneering work has produced shared gaze visualizations using remote eye-trackers, which has been shown to benefit collaborative learning from a distance¹⁴. Dyads who could see the gaze of their partner in real time exhibited more JVA, collaborated better and achieved higher learning gains compared to a control group. Future work will examine whether this kind of intervention can support collaborative processes in colocated settings (e.g., through virtual or augmented reality headsets).

DISCLOSURES:

The authors declare that they have no competing financial interests.

ACKNOWLEDGEMENTS:

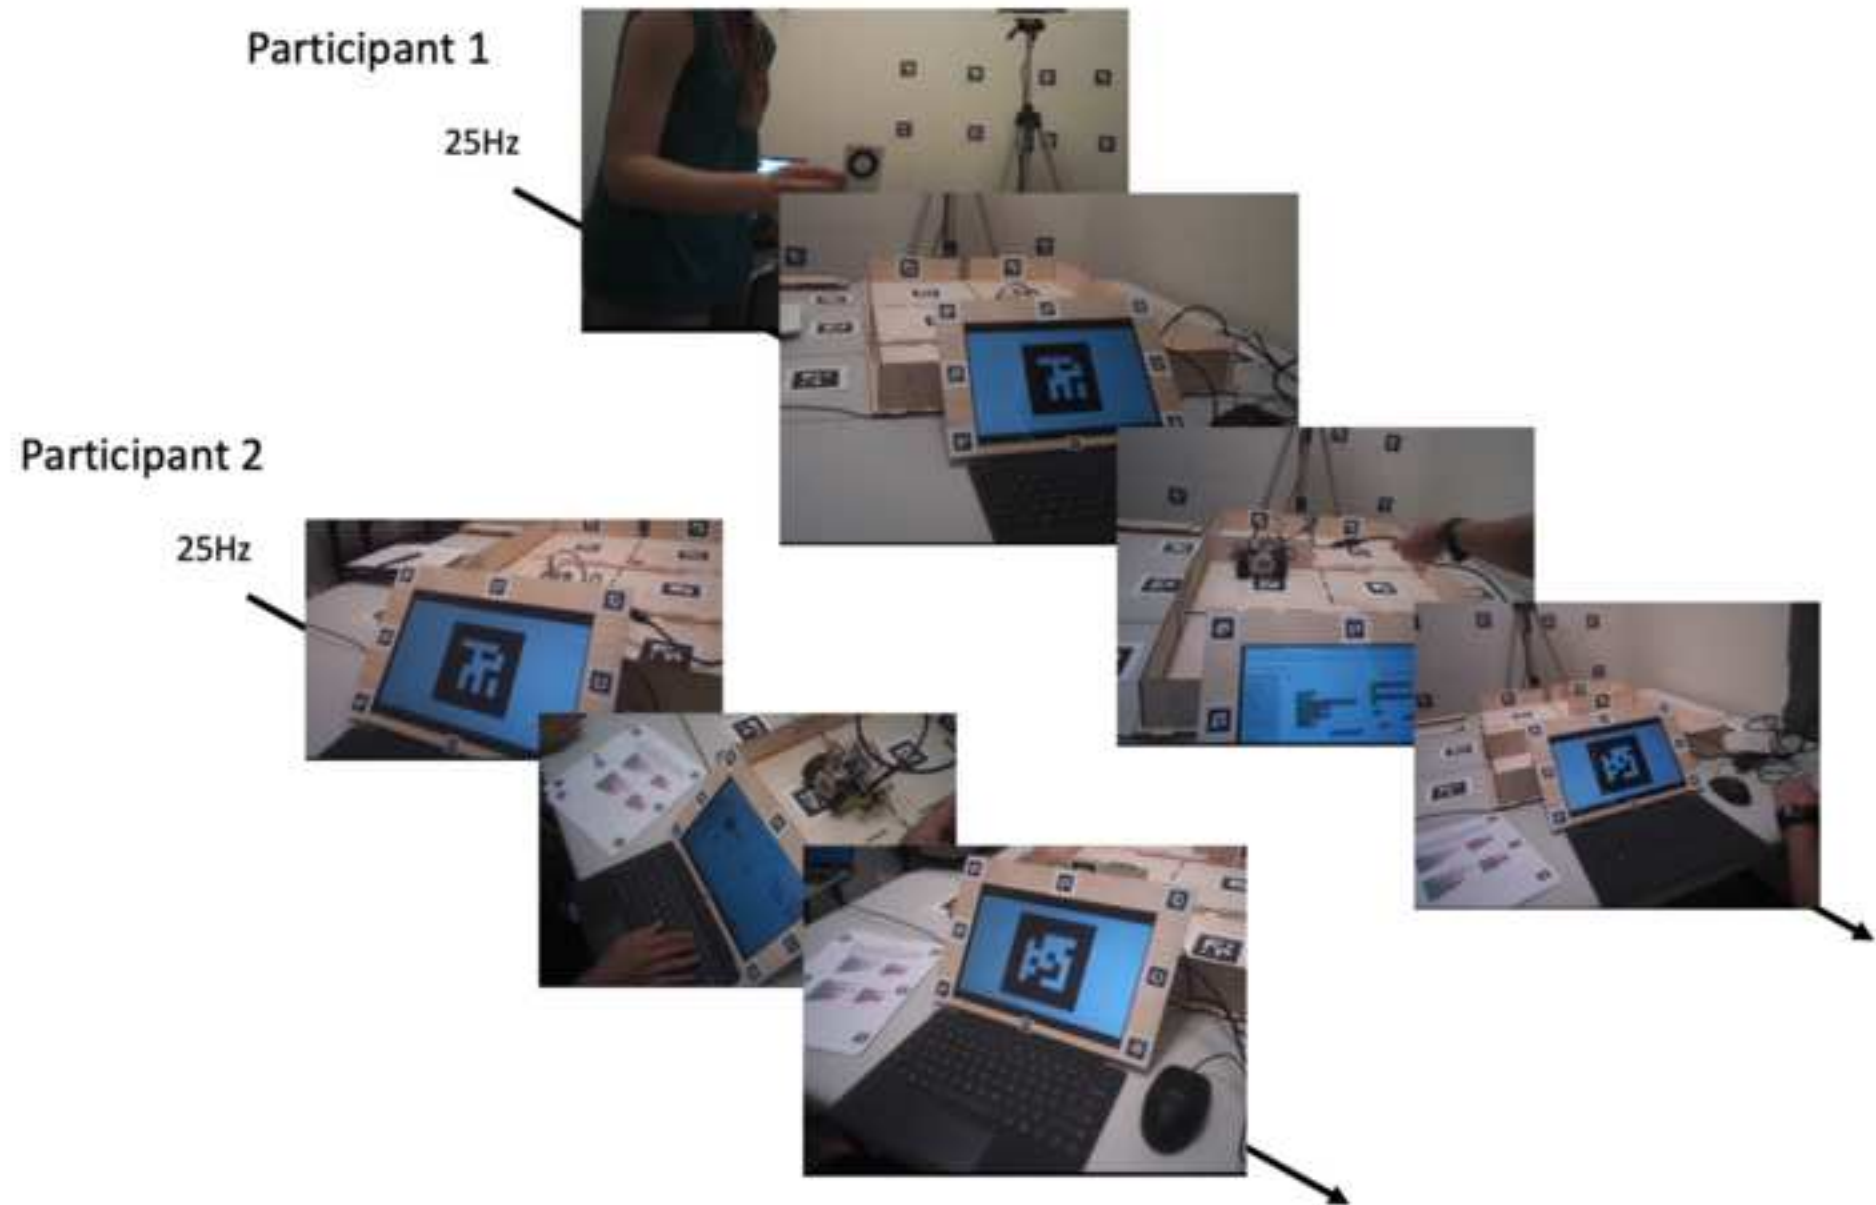
The development of this methodology was supported by the National Science Foundation (NSF #0835854), the Leading House Technologies for Vocation Education, funded by the Swiss State Secretariat for Education, Research and Innovation, and the Harvard School of Education's Dean Venture Fund.

REFERENCES:

1. Tomasello, M. Joint attention as social cognition. In: Moore, C., Dunham, P. J. (eds.) Joint attention: Its origins and role in development. pp. 103–130. Lawrence Erlbaum Associates, Inc, Hillsdale, NJ, England (1995).
2. Mundy, P., Sigman, M., Kasari, C. A longitudinal study of joint attention and language development in autistic children. *Journal of Autism and Developmental Disorders*. **20**, 115–128 (1990).
3. Clark, H. H., Brennan, S. E. Grounding in communication. In: Resnick, L. B., Levine, J. M.,

- Teasley, S. D. (eds.) Perspectives on socially shared cognition. pp. 127–149. American Psychological Association, Washington, DC, US (1991).
4. Siposova, B., Carpenter, M. A new look at joint attention and common knowledge. *Cognition*. **189**, 260–274 (2019).
 5. Gergle, D., Clark, A. T. See What I’m Saying?: Using Dyadic Mobile Eye Tracking to Study Collaborative Reference. In: Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work. pp. 435–444. ACM, New York, NY, USA (2011).
 6. Renner, P., Pfeiffer, T., Wachsmuth, I. Spatial References with Gaze and Pointing in Shared Space of Humans and Robots. In: Freksa, C., Nebel, B., Hegarty, M., Barkowsky, T. (eds.) Spatial Cognition IX. pp. 121–136. Springer International Publishing (2014).
 7. Shvarts, A. Y. Automatic detection of gaze convergence in multimodal collaboration: a dual eye-tracking technology. *The Russian Journal of Cognitive Science*. **5**, 4 (2018).
 8. Bonnard, Q. et al. Chilitags: Robust Fiducial Markers for Augmented Reality [software]. <https://github.com/chili-epfl/qml-chilitags> (2013).
 9. Jermann, P., Mullins, D., Nüssli, M.-A., Dillenbourg, P. Collaborative Gaze Footprints: Correlates of Interaction Quality. Connecting Computer-Supported Collaborative Learning to Policy and Practice: CSCL2011 Conference Proceedings., Volume I - Long Papers, 184–191 (2011).
 10. Richardson, D. C., Dale, R. Looking To Understand: The Coupling Between Speakers’ and Listeners’ Eye Movements and Its Relationship to Discourse Comprehension. *Trends in Cognitive Sciences*. **29**, 1045–1060 (2005).
 11. Richardson, D. C., Dale, R., Kirkham, N. Z. The Art of Conversation Is Coordination Common Ground and the Coupling of Eye Movements During Dialogue. *Psychological Science*. **18**, 407–413 (2007).
 12. Schneider, B. et al. Using Mobile Eye-Trackers to Unpack the Perceptual Benefits of a Tangible User Interface for Collaborative Learning. *ACM Transactions on Computer-Human Interaction*. **23**, 1–23 (2016).
 13. Meier, A., Spada, H., Rummel, N. A rating scheme for assessing the quality of computer-supported collaboration processes. *Int. J. Comput.-Support. Collab. Learn.* **2**, 63–86 (2007).
 14. Schneider, B., Pea, R. Real-time mutual gaze perception enhances collaborative learning and collaboration quality. *Journal of Computer-Supported Collaborative Learning*. **8**, 375–397 (2013).








Participant 1

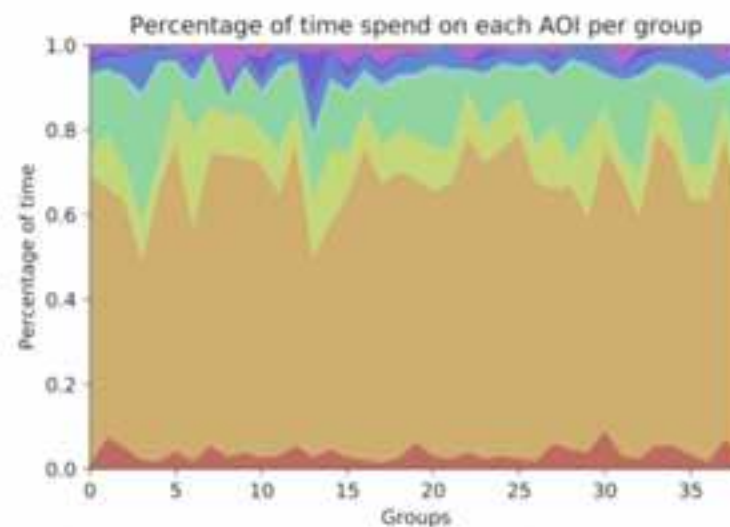



Participant 2

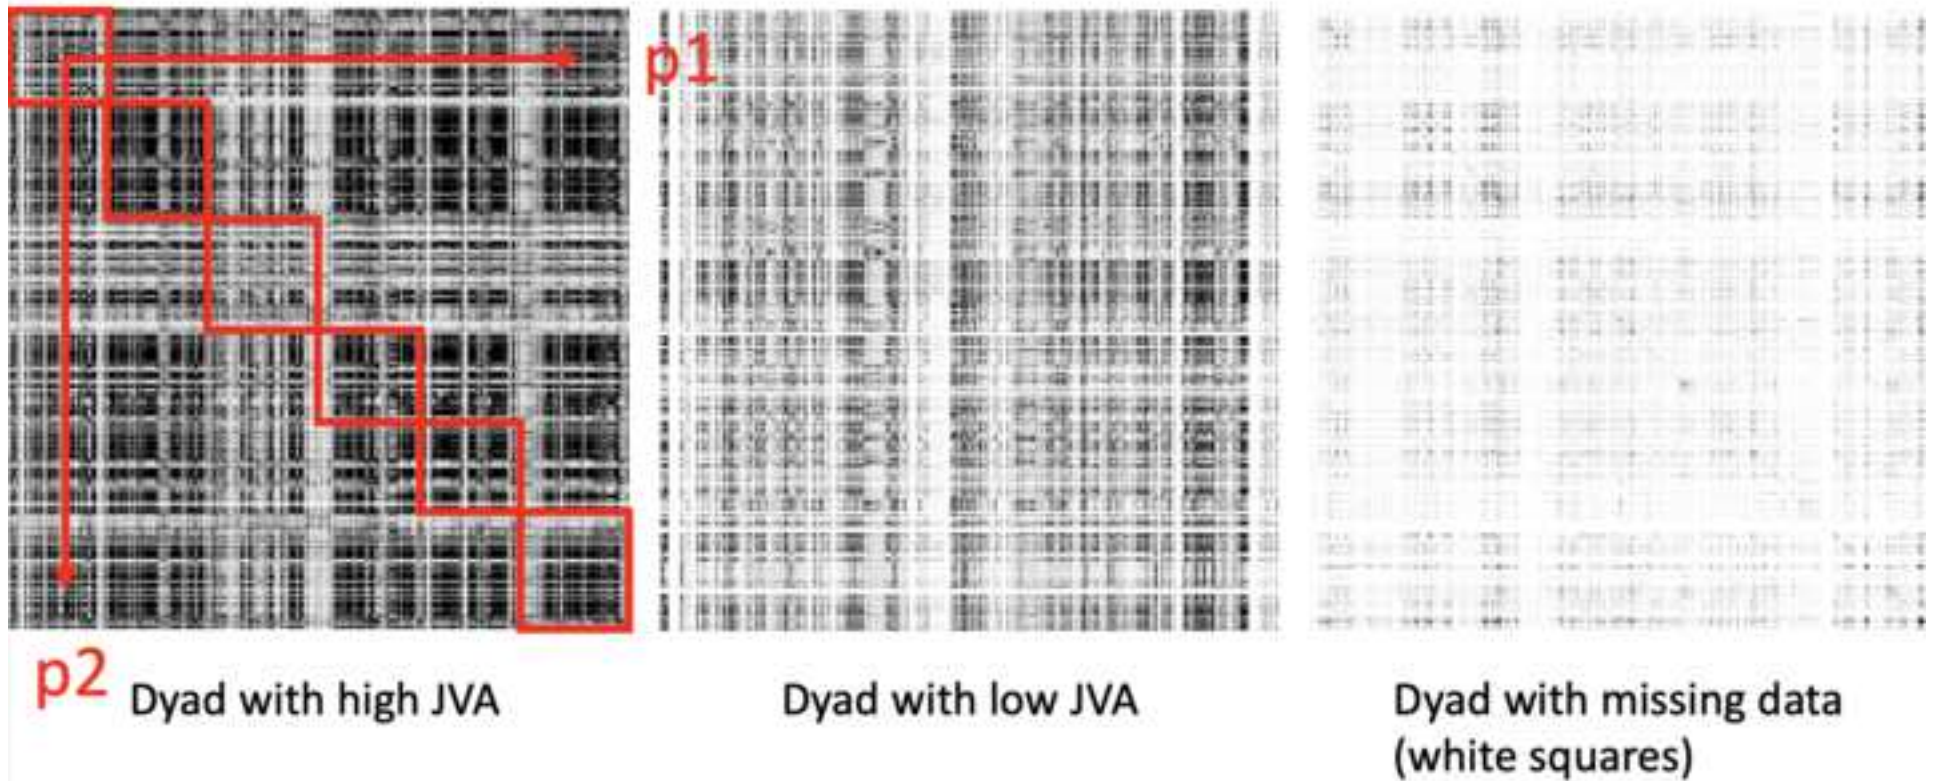




Sanity check #1: Heatmaps



Sanity check #2: Cross Recurrence Graphs



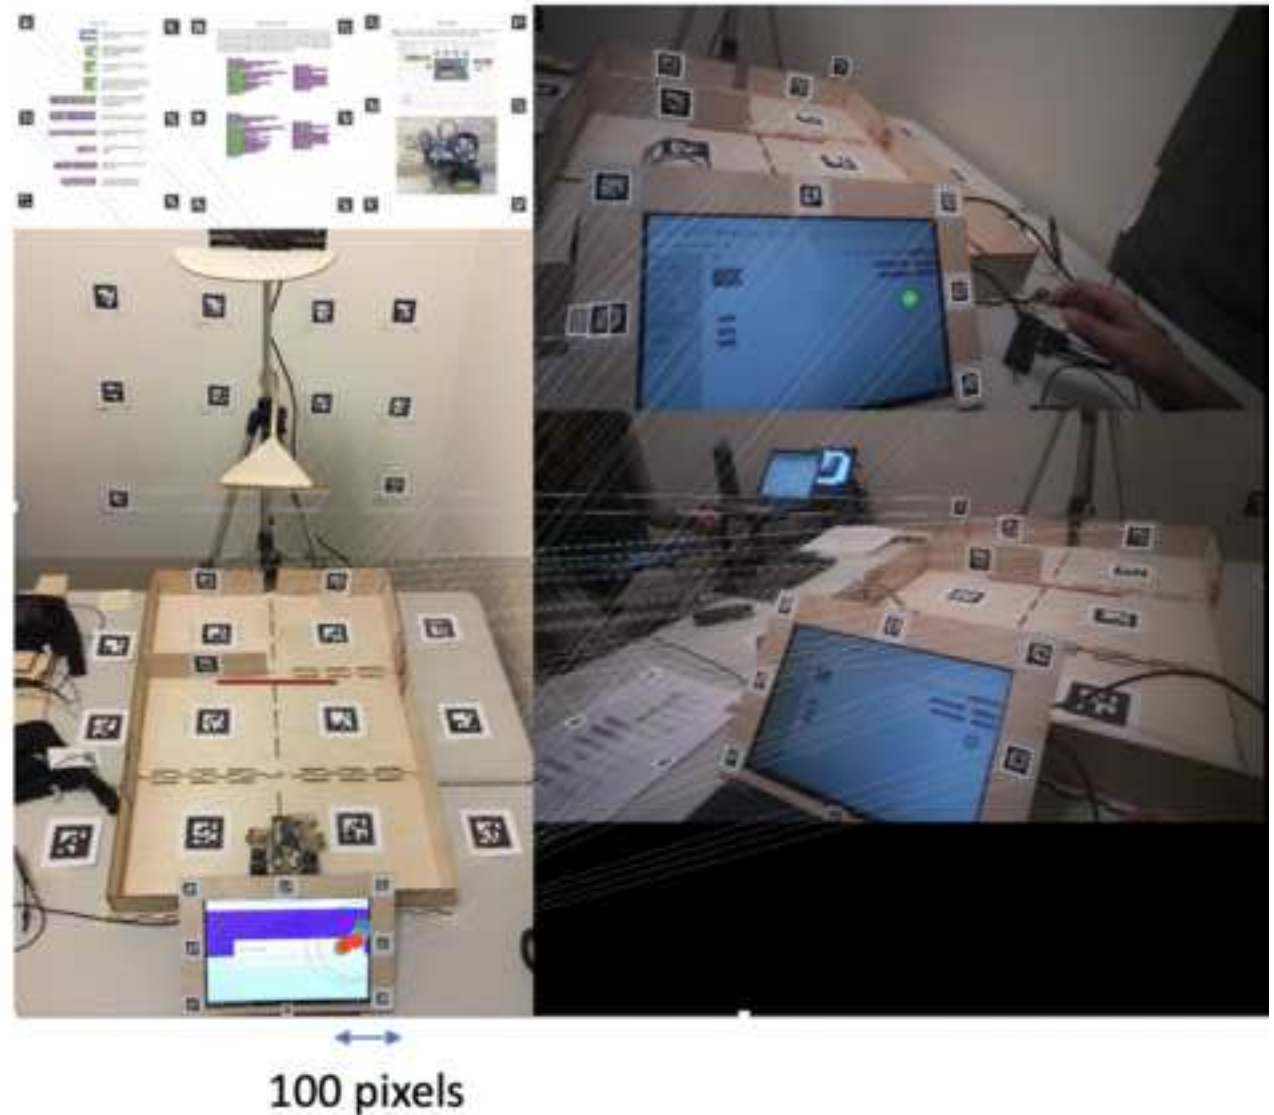
Two parameters for Computing Joint Visual Attention: time and distance

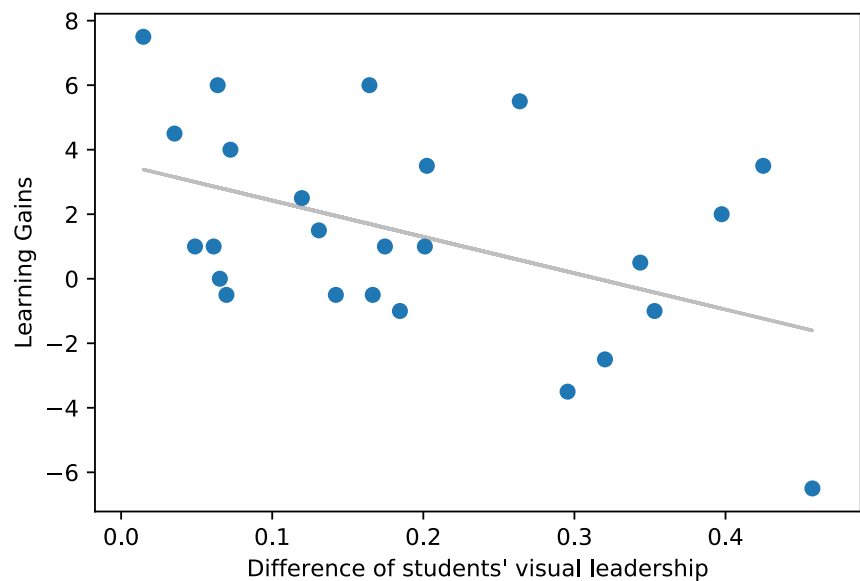

Time Window

-2sec / + 2sec
(Richardson and Dale, 2005)

Distance (radius)

To be determined
based on the task
and setting





Name of Material/ Equipment	Company	Catalog Number
Tobii Glasses 2	Tobii	N/A
Fiducial markers	Chili lab – EPFL, Switzerland	N/A

Comments/Description

<https://www.tobiipro.com/product-listing/tobii-pro-glasses-2/>

<https://github.com/chili-epfl/chilitags>



Harvard
Graduate School
of Education

Professor Bertrand Schneider
Longfellow Hall, office 333
School of Education
Harvard University
13 Appian Way Cambridge, MA 02138
bertrand_schneider@gse.harvard.edu

August 12th, 2019

Dear JoVE editors,

I wish to submit a revised version of the manuscript: “A Methodology for Capturing Joint Visual Attention Using Mobile Eye-trackers” for consideration by the Journal of Visualized Experiments. I have made the following changes to the manuscript and video:

1. Section 0:46-3:00 of the video now appears in the manuscript.
2. The written protocol is in the imperative voice (when applicable).
3. Protocol step 4c: There is more detail on how the python functions are implemented (which includes additional references and/or links to example code).
4. Figure 8: the two panels are combined into one image file and both are described in the caption.
5. Showing Tobii logos in the video: as agreed by email, we haven't made any changes to citing Tobii in the video.
6. 3:01, 3:50, 6:38, 7:13, 8:23 - the white background has been extended to eliminate the pillarboxing. The background is now in a consistent color.
7. The title card is placed at the very beginning of the video as well as the end.
8. There are chapter title cards before each section of the video.
9. The article ID number (60670) was included in the video file name and uploaded to dropbox.

Please don't hesitate to let me know if other changes are necessary.

Sincerely,

Bertrand Schneider, PhD
Assistant Professor
Graduate School of Education
Harvard University

SPRINGER NATURE LICENSE
TERMS AND CONDITIONS

Oct 04, 2019

This Agreement between Bertrand Schneider ("You") and Springer Nature ("Springer Nature") consists of your license details and the terms and conditions provided by Springer Nature and Copyright Clearance Center.

License Number	4682160372499
License date	Oct 04, 2019
Licensed Content Publisher	Springer Nature
Licensed Content Publication	International Journal of Computer-Supported Collaborative Learning
Licensed Content Title	Leveraging mobile eye-trackers to capture joint visual attention in co-located collaborative learning groups
Licensed Content Author	Bertrand Schneider, Kshitij Sharma, Sebastien Cuendet et al
Licensed Content Date	Jan 1, 2018
Licensed Content Volume	13
Licensed Content Issue	3
Type of Use	Journal/Magazine
Requestor type	academic/university or research institute
Is this reuse sponsored by or associated with a pharmaceutical or a medical products company?	no
Format	electronic
Portion	figures/tables/illustrations
Number of figures/tables/illustrations	1
Will you be translating?	no
Circulation/distribution	1 - 29
Author of this Springer Nature content	yes
Title of new article	A Methodology for Capturing Joint Visual Attention Using Mobile Eye-trackers

Lead author	Bertrand Schneider
Title of targeted journal	Journal of Visualized Experiments
Publisher	MyJove Corp.
Expected publication date	Jan 2020
Portions	Figure 10 - two scatter plots describing correlations between eye-tracking measures and outcomes of interest.
Requestor Location	Bertrand Schneider 3432 20th St SAN FRANCISCO, CA 94110 United States Attn: Bertrand Schneider
Total	0.00 USD
Terms and Conditions	

Springer Nature Customer Service Centre GmbH **Terms and Conditions**

This agreement sets out the terms and conditions of the licence (the **Licence**) between you and **Springer Nature Customer Service Centre GmbH** (the **Licensor**). By clicking 'accept' and completing the transaction for the material (**Licensed Material**), you also confirm your acceptance of these terms and conditions.

1. Grant of License

- 1. 1.** The Licensor grants you a personal, non-exclusive, non-transferable, world-wide licence to reproduce the Licensed Material for the purpose specified in your order only. Licences are granted for the specific use requested in the order and for no other use, subject to the conditions below.
- 1. 2.** The Licensor warrants that it has, to the best of its knowledge, the rights to license reuse of the Licensed Material. However, you should ensure that the material you are requesting is original to the Licensor and does not carry the copyright of another entity (as credited in the published version).
- 1. 3.** If the credit line on any part of the material you have requested indicates that it was reprinted or adapted with permission from another source, then you should also seek permission from that source to reuse the material.

2. Scope of Licence

- 2. 1.** You may only use the Licensed Content in the manner and to the extent permitted by these Ts&Cs and any applicable laws.
- 2. 2.** A separate licence may be required for any additional use of the Licensed Material, e.g. where a licence has been purchased for print only use, separate permission must be obtained for electronic re-use. Similarly, a licence is only valid in the language selected and does not apply for editions in other languages unless additional translation rights have been granted separately in the licence. Any content owned by third parties are expressly excluded from the licence.

2. 3. Similarly, rights for additional components such as custom editions and derivatives require additional permission and may be subject to an additional fee. Please apply to Journalpermissions@springernature.com/bookpermissions@springernature.com for these rights.

2. 4. Where permission has been granted **free of charge** for material in print, permission may also be granted for any electronic version of that work, provided that the material is incidental to your work as a whole and that the electronic version is essentially equivalent to, or substitutes for, the print version.

2. 5. An alternative scope of licence may apply to signatories of the [STM Permissions Guidelines](#), as amended from time to time.

3. Duration of Licence

3. 1. A licence for is valid from the date of purchase ('Licence Date') at the end of the relevant period in the below table:

Scope of Licence	Duration of Licence
Post on a website	12 months
Presentations	12 months
Books and journals	Lifetime of the edition in the language purchased

4. Acknowledgement

4. 1. The Licensor's permission must be acknowledged next to the Licenced Material in print. In electronic form, this acknowledgement must be visible at the same time as the figures/tables/illustrations or abstract, and must be hyperlinked to the journal/book's homepage. Our required acknowledgement format is in the Appendix below.

5. Restrictions on use

5. 1. Use of the Licensed Material may be permitted for incidental promotional use and minor editing privileges e.g. minor adaptations of single figures, changes of format, colour and/or style where the adaptation is credited as set out in Appendix 1 below. Any other changes including but not limited to, cropping, adapting, omitting material that affect the meaning, intention or moral rights of the author are strictly prohibited.

5. 2. You must not use any Licensed Material as part of any design or trademark.

5. 3. Licensed Material may be used in Open Access Publications (OAP) before publication by Springer Nature, but any Licensed Material must be removed from OAP sites prior to final publication.

6. Ownership of Rights

6. 1. Licensed Material remains the property of either Licensor or the relevant third party and any rights not explicitly granted herein are expressly reserved.

7. Warranty

IN NO EVENT SHALL LICENSOR BE LIABLE TO YOU OR ANY OTHER PARTY OR ANY OTHER PERSON OR FOR ANY SPECIAL, CONSEQUENTIAL, INCIDENTAL OR INDIRECT DAMAGES, HOWEVER CAUSED, ARISING OUT OF OR IN CONNECTION WITH THE DOWNLOADING, VIEWING OR USE OF THE MATERIALS REGARDLESS OF THE FORM OF ACTION, WHETHER FOR BREACH OF CONTRACT, BREACH OF WARRANTY, TORT, NEGLIGENCE, INFRINGEMENT OR OTHERWISE (INCLUDING, WITHOUT LIMITATION, DAMAGES BASED ON LOSS OF PROFITS, DATA, FILES, USE, BUSINESS OPPORTUNITY OR CLAIMS OF THIRD PARTIES), AND WHETHER OR NOT THE PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION SHALL APPLY NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY PROVIDED HEREIN.

8. Limitations

8. 1. *BOOKS ONLY*:Where '**reuse in a dissertation/thesis**' has been selected the following terms apply: Print rights of the final author's accepted manuscript (for clarity, NOT the published version) for up to 100 copies, electronic rights for use only on a personal website or institutional repository as defined by the Sherpa guideline (www.sherpa.ac.uk/romeo/).

9. Termination and Cancellation

9. 1. Licences will expire after the period shown in Clause 3 (above).

9. 2. Licensee reserves the right to terminate the Licence in the event that payment is not received in full or if there has been a breach of this agreement by you.

Appendix 1 — Acknowledgements:

For Journal Content:

Reprinted by permission from [**the Licensor**]: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**]
[**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication)]

For Advance Online Publication papers:

Reprinted by permission from [**the Licensor**]: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**]

[**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication), advance online publication, day month year (doi: 10.1038/sj.[JOURNAL ACRONYM].)]

For Adaptations/Translations:

Adapted/Translated by permission from [**the Licensor**]: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication)]

Note: For any republication from the British Journal of Cancer, the following credit line style applies:

Reprinted/adapted/translated by permission from [**the Licensor**]: on behalf of Cancer Research UK: : [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication)]

For Advance Online Publication papers:

Reprinted by permission from The [**the Licensor**]: on behalf of Cancer Research UK: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication), advance online publication, day month year (doi: 10.1038/sj.[JOURNAL ACRONYM].)]

For Book content:

Reprinted/adapted by permission from [**the Licensor**]: [**Book Publisher** (e.g. Palgrave Macmillan, Springer etc)] [**Book Title**] by [**Book author(s)**] [**COPYRIGHT**] (year of publication)

Other Conditions:

Version 1.2

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.
