Brief communication: Three errors and two problems in a recent paper: *gazeNet*: End-to-end eye-movement event detection with deep neural networks (Zemblys, Niehorster, and Holmqvist, 2019)

Lee Friedman¹

Published online: 13 April 2020
© The Psychonomic Society, Inc. 2020

Abstract

Zemblys et al. (*Behavior Research Methods*, 51(2), 840–864, 2019) reported on a method for the classification of eye-movements (“*gazeNet*”). I have found three errors and two problems with that paper that are explained herein. **Error 1:** The *gazeNet* classification method was built assuming that a hand-scored dataset from Lund University was all collected at 500 Hz, but in fact, six of the 34 recording files were actually collected at 200 Hz. Of the six datasets that were used as the training set for the *gazeNet* algorithm, two were actually collected at 200 Hz. **Problem 1** has to do with the fact that even among the 500 Hz data, the inter-timestamp intervals varied widely. **Problem 2** is that there are many unusual discontinuities in the saccade trajectories from the Lund University dataset that make it a very poor choice for the construction of an automatic classification method. **Error 2** The *gazeNet* algorithm was trained on the Lund dataset, and then compared to other methods, not trained on this dataset, in terms of performance on this dataset. This is an inherently unfair comparison, and yet nowhere in the *gazeNet* paper is this unfairness mentioned. **Error 3** arises out of the novel event-related agreement analysis employed by the *gazeNet* authors. Although the authors intended to classify unmatched events as either false positives or false negatives, many are actually being classified as true negatives. True negatives are not errors, and any unmatched event misclassified as a true negative is actually driving kappa higher, whereas unmatched events should be driving kappa lower.

Keywords Eye-movement classification · Deep neural networks · Article commentary

Error number 1

The *gazeNet* paper depends on the hand scoring of the eye-tracker data from Lund University (hereafter called the *Lund2013*.¹) The first error has to do with this dataset, on which the *gazeNet* algorithm was trained and developed. Here is how the *gazeNet* authors describe the *Lund2013*:

“It consists of monocular eye-movement data of participants viewing images, videos, and moving

dots. The eye-movements of all participants were recorded with the SMI Hi-Speed 1250 eye-tracker, running at a sampling frequency of 500 Hz. Two domain experts then manually segmented data into fixations, saccades, post-saccadic oscillations (PSO), smooth pursuit, blinks, and undefined events. A comprehensive description of the *Lund2013* dataset and the coding process can be found in Larsson et al. (2013).” (page 3, left column, last paragraph)

The 34 recordings where subjects were observing images were included in their analysis. Unfortunately, the *gazeNet* authors were unaware of the fact that in six² of the 34 files the median inter-timestamp intervals were 5000μs, for a nominal sample rate of 200 Hz rather than the expected 500 Hz. Two of these 200-Hz datasets were included in the *gazeNet* training set (of a total of six), one each was in the validation sets for rater RA and rater MN, and one each was

¹Available for download at <http://www.humlab.lu.se/en/person/MarcusNyström/>

✉ Lee Friedman
lfriedman10@gmail.com

¹ Derrick M5, Department of Computer Science, Texas State University, 601 University Drive, San Marcos, TX 78640, USA

²Highlighted in blue in Appendix A

in the test sets for rater RA and rater MN. The gazeNet paper states that “Trials were 4 or 10 s long.” (column 2, page 3). This statement is false. If one subtracts the first timestamp from the last timestamp, one will see that all of the files are almost exactly 10 s. Since the gazeNet algorithm was trained and validated on improper data, it would seem to the present author that, at the least, the algorithm is not optimally trained. Since the gazeNet model was trained on improper data, there is basis for concern about the value of the reported gazeNet values in the gazeNet paper.

Problem 1

The gazeNet authors apply the MNH (Friedman et al., 2018) to the *Lund2013*. The MNH (Friedman et al., 2018) was developed using the EyeLink 1000, at a 1000-Hz sampling rate. (I created the MNH.) It is worth considering whether the application of the MNH to the *Lund2013* is appropriate. Most of the *Lund2013* signals were sampled near 500 Hz, but as noted above, some of the data were sampled at 200 Hz. Although the gazeNet authors correctly adjusted the sampling frequency parameter of the MNH to 500 Hz prior to application, they obviously did not adjust either the MNH or the data for the recordings at 200 Hz. Furthermore, even for studies sampled at 500 Hz, there is significant variability in the inter-timestamp intervals in the *Lund2013* (Fig. 1). According to technical staff at SR-Research, the inter-timestamp intervals for the EyeLink 1000 are all exactly 1 ms. So, prior to applying the MNH to this data, this instability in the sampling rate would need to be addressed.³

Problem 2

There is a further problem with the *Lund2013*. A simple inspection of the recordings in the *Lund2013* reveals many discontinuities in saccades. (See Fig. 2 for examples.) In a large number of samples during saccades (for this analysis only files scored by MN were used), there is a large gap in horizontal pixel position followed by a very short gap in pixel position. The blue dots in Fig. 2 point to the sample

Fig. 1 Inter-sample intervals (ISI) for the *Lund2013*. This histogram reflects all the data from all the time stamps for the fourteen 500 Hz (nominal) recordings that were coded by rater MN. At that sample rate, samples should be exactly 2000 μ s apart. Note the bimodal nature of this histogram and the considerable variability around the nominal rate

after a large horizontal position move and followed by a small move. To automatically detect these events, saccade samples were searched to find samples where the horizontal distance from the previous sample was greater than three times larger than the distance to the next sample. Also, the distance to the next sample had to be less than 5.0 pixels. Since such events can occur naturally at the end of saccades, the search was stopped six samples before the end of each saccade. There were 292 saccades in the entire dataset, and 109 (37%) had a least one of these unusual events detected (see Fig. 3). The detection was not perfect, but the vast majority were indeed these discontinuities. Three saccades had four such events. These discontinuities tend to occur more often in large saccades, and also at a time when eye velocity is high.

These abnormalities are easily spotted when viewing the signals. I certainly think it is highly relevant for the reader to be informed about these unusual saccade artifacts, and their presence certainly should have been reported. Given all of these issues with the *Lund2013*, I would not recommend its use, without correction, to develop and test a machine learning algorithm designed to classify eye movements generally.

³As the creator of the MNH, I cannot endorse the application of the MNH to data collected at 500 Hz. I do not have extensive experience assessing its performance on such data. It was written for a reading task recorded with the EyeLink 1000 at 1000 Hz. However, I would suggest that those who intend to use it on such data anyway, interpolate their data to 1000 Hz using the “Piecewise Cubic Hermite Interpolating Polynomial” (PCHIP) method available in MATLAB.

Fig. 2 Eight examples, of eight different saccades from eight different files from the *Lund2013*. File 1 = TH34-img-Europe-labeled-MN, File 2 = TH34-img-vy-labeled-MN, File 3 = TL20-img-konijntjes-labeled-MN, File 4 = TL28-img-konijntjes-labeled-MN, File 5 = UH27-img-vy-labeled-MN, File 6 = UH29-img-Europe-labeled-MN, File 7 = UH33-img-vy-labeled-MN, File 8 = UL43-img-Rome-labeled-MN.

These kinds of discontinuities would obviously lead to an mis-estimation of saccade velocity, which the MNH very much depends on, as does the algorithm of (Nyström & Holmqvist, 2010). It is not possible to accurately assess the

The title for each plot contains the file number and the saccade number. The hand scoring contained in the dataset was used to determine where saccades occur (shown in *red*). Non-saccade data are shown in *green*, and the *blue dots* indicate the occurrence of the unusual discontinuities in this data (see text for explanation). These plots all employed the hand classification performed by rater “MN”

classification accuracy of the MNH on the *Lund2013* until this problem is solved, assuming it can be solved.⁴

Although the gazeNet paper was developed using eye movement data that had several unusual properties, it is a general method that could, in theory, be applied to any similar eye-movement dataset.

Fig. 3 Frequency histogram of number of discontinuous events per saccade

Error number 2

On page 5 of the gazeNet paper, the authors state:

“...gazeNet was evaluated together with four other event detection algorithms on the Lund2013-image-test dataset, genSet and two other eye movement datasets.”

⁴There are classification issues with respect to the GazeCom and humanFixationEvaluation datasets that are also evaluated in the gazeNet paper. Some of the issues with GazeCom are outlined in Djanian (2019), (see section 3.1.4, with six illustrations of misclassified saccades). This evaluation will be amplified in a future manuscript. For our analysis of the approach gazeNet applies to the human-FixationEvaluation dataset, see “GazeNetCritiqueAddendum.pdf” at <https://digital.library.txstate.edu/handle/10877/8649>

On page 17, the authors state:

“The main exception is in the performance of MNH in detecting fixations: at the event level it outperforms IRF (but not gazeNet) in the lund2013-image-test...”

The data on this comparison is prominently displayed in Table 8.

This makes it clear that the MNH and other algorithms were compared to gazeNet on the Lund2013 dataset.

It should be noted that machine learning algorithms are trained to minimize classification error for the provided ground-truth labels. If a great many labeled saccades in the training set have the discontinuities illustrated above, then the network is expected to detect such unconventional events during inference. As MNH was not designed to identify discontinuous saccades, there is no question that the gazeNet algorithm has an unfair advantage when classifying the Lund2013 dataset, where such discontinuous trajectories are labeled as standard saccades, but the problem is even more extensive than this. In addition to saccade discontinuities, the machine learning approach can “memorize” dataset specific patterns (saccade amplitude, frequency of saccades, general motion patterns, etc.). Any model comparison is unfair if only one algorithm was trained on the data used to compare methods and the other methods were not trained on this dataset. Nowhere in the gazeNet paper is this inherent unfairness acknowledged.⁵

Error number 3

The third error has to do with the calculation of the event-level agreement analysis in the gazeNet paper.⁶ As part of that paper, the authors presented a new method to

assess event-level agreement between two classifiers.⁷ One classifier was referred to as “ground truth” (“gt”), and in the gazeNet paper, was a human rater. The other classifier was referred to as an algorithm which makes “predictions” (“pr”). The final result was a Cohen’s kappa for each comparison.

For ground truth (“gt”), I also used the *Lund2013*, described above. Specifically, the 34 recordings when subjects were looking at images are the focus of this report as was the case for the gazeNet paper. Like the gazeNet paper, I treated all of the *Lund2013* as if it was collected at the nominal sampling rate of 500 Hz. Issues of data quality are not relevant to the point I am trying to make in this section. The names of the original datasets and the names I used are listed in Appendix A. Eye-movement data were classified into fixation, saccades, post-saccadic oscillations, smooth pursuit, blinks, and undefined.

For my comparison, “predictive” algorithm, I employed the MNH (Friedman et al., 2018). As the creator of the MNH, it was easy to get the classifications. I reiterate that data quality is not an issue for the point I want to make regarding event-level agreement. I am not trying to accurately estimate the classification accuracy of the MNH on the *Lund2013*. The MNH classifies data into fixations, saccades, post-saccadic oscillations, artifact, small RIONEPS noise (Friedman et al., 2018; Abdulin et al., 2017) and large RIONEPS noise. All sections of each dataset that were not labeled as fixation, saccade, or PSO were removed from both the gt scored and the pr scored data. Any dataset that contained less than 1000 remaining samples was not analyzed further. This resulted in 28 datasets.⁸

The description of the algorithm in the gazeNet paper does not include pseudocode, nor does it precisely describe the actual steps in the analysis. Also, this description fails to introduce the concept of true negatives, which as we will see below, are critical to consider. Therefore, I will take some time to introduce this method using my own figures and tables. The input to the analysis is a two-column dataset, with one column listing the original classes (fixations, saccades, and PSOs), on a sample by sample basis, from the gt method (*Lund2013* hand scoring), and one column listing the original classes from the pr (MNH) algorithm. Fixations were coded as 1, saccades as 2, and PSOs as 3. Herein, I will be focused exclusively on obtaining kappas for fixation events, saccade events, and PSO events. For illustrative purposes, I will be focused on the event-related kappa for PSO events for my 12th dataset (“Set_12.csv”). To

⁵In the paper introducing the MNH (Friedman et al., 2018), three other algorithms were compared to the MNH on EyeLink-1000 data collected at 1000 Hz. Two were developed or designed for such EyeLink data. The third algorithm (Nyström & Holmqvist, 2010) was developed on an SMI HiSpeed device. In this case, I received personal communication from Dr. Nyström indicating that that there should be no problem as long as the sampling rate parameter was modified. Also, Dr. Holmqvist, who self-identified as one of the reviewers, raised no issues about the appropriateness of applying their algorithm to EyeLink data.

⁶I received substantial initial assistance related to this error from a graduate student (Vladyslav Prokopenko). He is my colleague at Texas State University. Before I got involved, he had looked carefully at the event-level analysis code and told me that he was confident that it did not do what the gazeNet authors intended, but he was not able to describe the exact nature of the error, nor was he able to point to the lines in the code which caused the error. Mr. Prokopenko also provided four Python code assignment statements that corrected the error in the original code.

⁷The Python code which performs the event-level agreement analysis for the gazeNet paper is available at <https://github.com/r-zemblyus/ETeval>.

⁸Datasets are at <https://doi.org/10.18738/T8/LKSTW4>

produce an evaluation of say, PSO events, all events that are not PSOs are combined into non-PSO events (Appendix B, Fig. 1)

The next step was to match gt binary events to pr binary events. This was done by determining the degree of overlap between events. The amount of overlap between each gt event and each pr event was calculated. Then each overlap was considered for matching from the longest to shortest and events were matched based on the overlap, if, and only if, both events were not already matched to other events. A single pr event can only be matched to a single gt event.

The data format for input to the kappa function used in this code was two vectors, one representing the class (1 or 0) of the gt event, and the other representing the class (1 or 0) of the pr event. Unmatched gt events and unmatched

pr events are also given similar codes. The final codes for the PSO evaluation for set number 12 are shown in Table 1. Gt events that are matched to pr events are labeled as “Matched”. Although the codes for these matched events are typically either 1-1 or 0-0, this was not always the case. In the gazeNet analysis, two events can be matched, but one could be a PSO and the other could be a “not-PSO”. For example, event # 14 in this list (Table 1) shows a match, where the gt event was a PSO and the pr event was a not-PSO. Also, there are unmatched gt events and unmatched pr events. Although the two vectors that are labeled in Table 1 as gt_label and pr_label are all that are required to compute kappa using the particular function employed in the original Python code, statisticians typically convert such information into a contingency table like that shown in Table 2. Once

Table 1 Codes for events to be input to the kappa function. PSO analysis, Set_12.csv

Count	Status	gt_label	pr_label	Original classification ^a	Corrected gt label	Corrected pr label	Correct Classification
1	Matched	0	0				
2	Matched	1	1				
3	Matched	0	0				
4	Matched	1	1				
5	Matched	0	0				
6	Matched	1	1				
7	Matched	0	0				
8	Matched	1	1				
9	Matched	0	0				
10	Matched	1	1				
11	Matched	0	0				
12	Matched	1	1				
13	Matched	0	0				
14	Matched	1	0				
15	Matched	0	0				
16	Matched	1	1				
17	Matched	0	0				
18	UnMatched_GT	0	0	TN	0	1	FP
19	UnMatched_GT	1	0	FN			FN
20	UnMatched_GT	0	0	TN	0	1	FP
21	UnMatched_GT	1	0	FN			FN
22	UnMatched_GT	1	0	FN			FN
23	UnMatched_GT	0	0	TN	0	1	FP
24	UnMatched_GT	0	0	TN	0	1	FP
25	UnMatched_GT	1	0	FN			FN
26	UnMatched_PR	0	1	FP			FP
27	UnMatched_PR	0	0	TN	1	0	FN
28	UnMatched_PR	0	1	FP			FP
29	UnMatched_PR	0	1	FP			FP

^aTN = True Negative, FP = False Positive, FN = False Negative

Table 2 Contingency table

	PR=1	PR=0
GT=1	TP(1,1)	FN(1,0)
GT=0	FP(0,1)	TN(0,0)

the data are in this form, Cohen's kappa can be easily calculated.

To calculate kappa for this binary form of the data, all events in each input file need to be classified as either true positive (TP), true negative (TN), false positive (FP), or false negative (FN).

The results of the matching process and code labeling is illustrated in Table 1 (see also Appendix B Fig. 2). In Table 1, note that some unmatched gt events are labeled 0-0. In terms of the contingency table presented in Table 2, these would be interpreted as true negatives. However, all unmatched events must be considered as errors, and coded either 1-0 or 0-1. True-negatives (0-0) are not errors. I quote from the gazeNet paper:

“The remaining unmatched events are then labeled as false positives or false negatives, depending on whether they, respectively, occur in the ground truth or the algorithm event stream.” (page 6, right column, 2nd line)⁹

True negatives should only occur when non-PSO gt events are being correctly matched to non-PSO pr events. True negatives are an indication of good performance. That unmatched events are being labeled as true negatives is the indication that the original code is not performing correctly. If one looks at Table 1, one will see that all unmatched gt events have pr codes of 0 (blue highlight). You will also see that for all unmatched pr events, the corresponding gt code was 0 (green highlight). Therefore, although all of the unmatched events are errors and should be coded as either 1-0 or 0-1 (false negatives or false positives), many of them (5 of 12 or 42%) are actually coded as true negatives (red highlight). Treating unmatched events as true negatives has the effect of increasing kappa toward 1.0, and therefore must be an error. There are two lines in the code that are responsible for the error (see Table 3).

Prior to the first line of code (1), `evt_pr` is a list of the codes (0 or 1) for matched events. The variable “`set_gt`” contains a list of all the unmatched gt events. The part of assignment statement (1) that is:

```
pd.DataFrame(np.zeros(len(set_gt)))
```

⁹There are four potential types of mismatches, unmatched gt-true events, unmatched gt-false events, unmatched pr-true events and unmatched pr-false events. Unfortunately, the gazeNet authors do not specify exactly how each of these should be coded.

is appending a series of 0 codes to the `evt_pr` list that is the same length as `set_gt`. Appending a list of 0 values is simply not correct. I propose that what is needed is to append a list of codes that are opposite of the paired gt codes. So, if an event is unmatched, it must be coded either 0-1 or 1-0. I have prepared a corrected version of the event-level agreement Python code¹⁰. In assignment statements (3) and (4) (Table 3), the codes that will be appended for unmatched events are created as `inverted_gt` and `inverted_pr`. If the unmatched gt event has a code of 0, the pr code for the same event will be 1. If the unmatched pr event has a code of 0, the gt code for the same event will be 1. The original code overestimates the number of true negatives and underestimates the number of false negatives and false positives. Therefore, the original code produces kappa estimates that are too high. In Fig. 4, I present distributions of incorrect Kappa_gN (kappa_gazeNet) values as well as the corrected Kappa_Cor values for fixations, saccades and PSOs.

The gazeNet authors make a strong case for using kappa over an F1 score for sample-level agreement statistics. With regard to the event-level they state:

“Note that when evaluating the detector's performance at the event level instead of the sample level, the distribution of data across the classes is less unbalanced. Nevertheless, Cohen's kappa is also a robust measure to use in this case, and will thus be used throughout the paper.”

There is an alternative way to describe the situation. With binary classifications, every true event is followed by a false event. Thus, the distribution of classes is either perfectly balanced or only trivially unbalanced. An F1 score would be appropriate. Other researchers use the F1 score to describe their event-level agreement results (Hooge et al., 2017; Startsev et al., 2019). There is no concept of a true negative in the context of the F1 score, and there is no ambiguity regarding the classification of false negatives (Misses) and false positives (false alarms). The gazeNet authors could have preferred one matching criteria over another, (e.g., “largest overlap” versus “overlapping test-events occurring earliest in time”), and still analyzed the performance using an F1 score.

Also, in the section in the gazeNet paper describing this event-level evaluation there were two additional errors. The statement:

¹⁰See <https://digital.library.txstate.edu/handle/10877/8649>, where you can download ImprovedEventLevelKappa.zip. This code reproduces the original erroneous kappas and the new, corrected kappas. It does not depend on the original ETEval code. MATLAB code is also available to graphically illustrate the event-level agreement results, as in Appendix B, Figs. 1 and 2.

Table 3 Python assignment statements

Original code snippets:

```
evt_pr = pd.concat((evt_pr, pd.DataFrame(np.zeros(len(set_gt))))) (1)
```

```
evt_gt = pd.concat((evt_gt, pd.DataFrame(np.zeros(len(set_pr))))) (2)
```

Corrected code snippets:

```
inverted_gt = (~_etdata_gt.evt.loc[set_gt, 'evt'].values.astype(bool)).astype(int) (3)
```

```
inverted_pr = (~_etdata_pr.evt.loc[set_pr, 'evt'].values.astype(bool)).astype(int) (4)
```

```
evt_pr = pd.concat((evt_pr, pd.DataFrame(inverted_gt))) (5)
```

```
evt_gt = pd.concat((evt_gt, pd.DataFrame(inverted_pr))) (6)
```

“Hoppe and Bulling (2016) used a majority vote approach and calculated the F1 score.”

is false. Hoppe and Bulling (2016) do not express their event-level agreement results in terms of F1 scores.

Also, in this section, the gazeNet authors make the following statement:

“Hooge et al. (2017) developed another event level F1 score and used it to compare human coders for assessing fixation coding performance. Hooge et al. looked for overlapping test-events occurring earliest in time and labeled them as hits. The remaining

unmatched ground truth fixations were labeled as misses, while the remaining unmatched test fixations were labeled as false alarms. This works when there is a single event class, but when there is more than one event type to match, the Hooge et al. approach does not work well.” (page 6, left column, near top)

This notion does not make sense to me. The Hooge et al. approach (Hooge et al., 2017) is designed to compute an F1 score, and an F1 score is an average of precision and recall. Precision and recall are only defined in the two-class case. They are not defined in the multi-class case.

Fig. 4 Violin plots of Cohen kappa values before correction (“gN”, for gazeNet) and after correction (“Cor”, for correct). Mean values are shown in green. The interquartile range is shown as a black vertical bar. For fixation, there was a decrease of 18.27% after correction in mean kappa ($t = 9.18$, $df = 19$, p value = 0.00000002, paired t test, two-tailed). For saccades, there was a decrease of 16.20% after correction ($t = 6.55$, $df = 22$, p value = 0.000001). For PSOs, there was a decrease of 33.78% ($t = 8.61$, $df = 26$, p value = 0.000000004). Cases

where kappa.gN was equal to 1.0 were not included in this analysis (fixation $N = 8$, saccade $N = 5$, PSO $N = 1$), since there was nothing to correct in these cases. In three cases, all involving PSOs, the percent decrease was greater than 70% (0.42 to 0.1, 0.35 to 0.08 and 0.42 to 0.03). If cases in which kappa.gN = 1.0 are included, the mean percent decrease across all events was 21.1%, but in eight cases the percent decrease was greater than 50%

A note on nomenclature

In the machine learning literature, including the gazeNet paper, criteria which emerge from human observation, intelligence, insight, and intuition are pejoratively referred to as “hand-crafted”. The hand has nothing to do with such criteria. Jewelry is handcrafted. Clay sculpture sculpture is handcrafted. I propose that such criteria be properly referred to as criteria which emerge from human observation, intelligence, insight and intuition (*OI*³). Alternatively, criteria emerging from traditional approaches might be referred to as “mindful” whereas machine learning approaches might be referred to as “mindless”.

An earlier version of this manuscript is posted on arXiv (Friedman, 2019)

Acknowledgements I would like to express my gratitude to Mr. Vladyslav Prokopenko for all of his assistance with this work. Another graduate student colleague, Mr. Dmytro Katrychuk, made changes to my GitHub repository that facilitates setting up the working environment to use the code. He also significantly improved the Python script to increase its comprehensibility and to remove the dependency on the original ETeval code base. I very much appreciate his assistance.

Compliance with Ethical Standards

Conflict of interests I have no conflict of interest.

Appendix: A

Lund2013 dataset names.

Original file names	Renamed as	Original file names	Renamed as
TL28_img_konijntjes_labelled_MN.mat	Set_01.csv	UL23_img_Europe_labelled_RA.mat	Set_18.csv
TH34_img_Europe_labelled_MN.mat	Set_02.csv	UL31_img_konijntjes_labelled_RA.mat	Set_19.csv
TH34_img_vy_labelled_MN.mat	Set_03.csv	UL39_img_konijntjes_labelled_RA.mat	Set_20.csv
TL20_img_konijntjes_labelled_MN.mat	Set_04.csv	UL43_img_Rome_labelled_RA.mat	Set_21.csv
UH21_img_Rome_labelled_MN.mat	Set_05.csv	UL47_img_konijntjes_labelled_RA.mat	Set_22.csv
UH33_img_vy_labelled_MN.mat	Set_06.csv	TH38_img_Europe_labelled_RA.mat	Set_23.csv
UL23_img_Europe_labelled_MN.mat	Set_07.csv	TH46_img_Rome_labelled_RA.mat	Set_24.csv
UL31_img_konijntjes_labelled_MN.mat	Set_08.csv	TH50_img_vy_labelled_RA.mat	Set_25.csv
UL39_img_konijntjes_labelled_MN.mat	Set_09.csv	TL44_img_konijntjes_labelled_RA.mat	Set_26.csv
UL43_img_Rome_labelled_MN.mat	Set_10.csv	TL48_img_Europe_labelled_RA.mat	Set_27.csv
UL47_img_konijntjes_labelled_MN.mat	Set_11.csv	TL48_img_Rome_labelled_RA.mat	Set_28.csv
TL28_img_konijntjes_labelled_RA.mat	Set_12.csv	UH27_img_vy_labelled_RA.mat	Set_29.csv
TH34_img_Europe_labelled_RA.mat	Set_13.csv	UH29_img_Europe_labelled_RA.mat	Set_30.csv
TH34_img_vy_labelled_RA.mat	Set_14.csv	UH47_img_Europe_labelled_RA.mat	Set_31.csv
TL20_img_konijntjes_labelled_RA.mat	Set_15.csv	UH27_img_vy_labelled_MN.mat	Set_32.csv
UH21_img_Rome_labelled_RA.mat	Set_16.csv	UH29_img_Europe_labelled_MN_Corrected.mat	Set_33.csv
UH33_img_vy_labelled_RA.mat	Set_17.csv	UH47_img_Europe_labelled_MN.mat	Set_34.csv

Appendix: B

Additional Figures

Appendix B. Figure 1. Original three-level classification reduced to a binary classification: PSO vs. Non-PSO. In the original data from gt, I have fixations (*red*), saccades (*blue*), and PSOs (*green*). After conversion to binary labels, Non-PSOs are shown in black and PSO events are shown in *grey*. The first event, a fixation, and the

second event, a saccade, are grouped together as a single non-PSO event. The PSO labeled in the original series as event number 3 becomes the second binary event, PSO = true. A similar conversion to binary labels was made to the pr events

Appendix B. Figure 2. Binary event matching and relevant codes. The first gt event (#01) was a not-PSO and it was matched to the first pr event. The second gt event was a PSO event, and it was matched to the second pr event. GT events #03, #04, #05, and #06 are not matched but gt event #07 was a non-PSO and was matched to pr event #03. Gt event #08 was matched to pr event #04. The green codes (“0-0”, “1-1” and “0-1”) refer to the final labels that are to be input to the Python kappa function. For matched events, the codes appear above

the gt line, and for unmatched events the codes appear above the word “unmatched”. The lower code refers to the code assigned to pr, and the upper code refer to the code assigned to gt. Unmatched pr events do occur in this data but are not shown in this figure. The *red arrows* indicate unmatched events that are coded as true negatives (“0-0”). For illustrative purposes, the first event in both the gt and pr were truncated

References

- Zemblys, R., Niehorster, D. C., & Holmqvist, K. (2019). gazeNet: End-to-end eye-movement event detection with deep neural networks. *Behavior Research Methods*, 51(2), 840–864. <https://doi.org/10.3758/s13428-018-1133-5>
- Larsson, L., Nystrom, M., & Stridh, M. (2013). Detection of saccades and postsaccadic oscillations in the presence of smooth pursuit. *IEEE Transactions on Biomedical Engineering*, 60(9), 2484–2493. <https://doi.org/10.1109/TBME.2013.2258918>
- Friedman, L., Rigas, I., Abdulin, E., & Komogortsev, O. V. (2018). A novel evaluation of two related and two independent algorithms for eye movement classification during reading. *Behavior Research Methods*. <https://doi.org/10.3758/s13428-018-1050-7>.
- Djanian, S. (2019). Eye Movement Classification Using Deep Learning, Master Thesis. Department of Electronic Systems, Aalborg University, Fredrik Bajers Vej 7B,9220 Aalborg East, Denmark, [https://projekter.aau.dk/projekter/da/studentthesis/eye-movement-classification-using-deep-learning\(d61882dd-d297-48c5-bcc5-500f81dd78c3\).html](https://projekter.aau.dk/projekter/da/studentthesis/eye-movement-classification-using-deep-learning(d61882dd-d297-48c5-bcc5-500f81dd78c3).html).
- Nyström, M., & Holmqvist, K. (2010). An adaptive algorithm for fixation, saccade, and glissade detection in eyetracking data. *Behavior Research Methods*, 42(1), 188–204. <https://doi.org/10.3758/brm.42.1.188>
- Abdulin, E., Friedman, L., & Komogortsev, O. V. (2017). Method to Detect Eye Position Noise from Video-Oculography when Detection of Pupil or Corneal Reflection Position Fails. arXiv e-prints arXiv:1709.02700.
- Hooge, I. T. C., Niehorster, D. C., Nystrom, M., Andersson, R., & Hessels, R. S. (2017). Is human classification by experienced untrained observers a gold standard in fixation detection? *Behavior Research Methods*, 50(5), 1864–1881. <https://doi.org/10.3758/s13428-017-0955-x>
- Hoppe, S., & Bulling, A. (2016). End-to-End Eye Movement Detection Using Convolutional Neural Networks. arXiv e-prints. Retrieved from <https://ui.adsabs.harvard.edu/abs/2016arXiv160902452H>.
- Startsev, M., Agtzidis, I., & Dorr, M. (2019). 1D CNN with BLSTM for automated classification of fixations, saccades, and smooth pursuits. *Behavior Research Methods*, 51(2), 556–572. <https://doi.org/10.3758/s13428-018-1144-2>
- Friedman, L. (2019). Three errors and two problems in a recent paper: gazeNet: End-to-end eye-movement event detection with deep neural networks (Zemblys, Niehorster, and Holmqvist, 2019), arXiv e-prints, arXiv:1908.11436.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.