Link to data:

https://atreus.informatik.uni-tuebingen.de/seafile/d/8e2ab8c3fdd444e1a135/?p=%2Fsimulator&mode=list

Eye movement simulation and detector creation to reduce laborious parameter adjustments*

†

Wolfgang Fuhl
University Tuebingen, Perception Engineering
Tuebingen, Baden-Wuerttemberg
wolfgang.fuhl@uni-tuebingen.de

Thiago Santini
University Tuebingen, Perception Engineering
Tuebingen, Baden-Wuerttemberg
thiago.santini@uni-tuebingen.de

Thomas Kuebler
University Tuebingen, Perception Engineering
Tuebingen, Baden-Wuerttemberg
thomas.kuebler@uni-tuebingen.de

Nora Castner
University Tuebingen, Perception Engineering
Tuebingen, Baden-Wuerttemberg
castnern@informatik.uni-tuebingen.de

Wolfgang Rosenstiel
University Tuebingen, Technical Computer Science
Tuebingen, Baden-Wuerttemberg
Wolfgang.Rosenstiel@uni-tuebingen.de

Enkelejda Kasneci
University Tuebingen, Perception Engineering
Tuebingen, Baden-Wuerttemberg
Enkelejda.Kasneci@uni-tuebingen.de

ABSTRACT

Eye movements hold information about human perception, intention and cognitive state. Various algorithms have been proposed to identify and distinguish eye movements, particularly fixations, saccades, and smooth pursuits. A major drawback of existing algorithms is that they rely on accurate and constant sampling rates, impeding straightforward adaptation to new movements such as micro saccades. We propose a novel eye movement simulator that i) probabilistically simulates saccade movements as gamma distributions considering different peak velocities and ii) models smooth pursuit onsets with the sigmoid function. This simulator is combined with a machine learning approach to create detectors for general and specific velocity profiles. Additionally, our approach is capable of using any sampling rate, even with fluctuations. The machine learning approach consists of different binary patterns combined using conditional distributions. The simulation is evaluated against publicly available real data using a squared error, and the detectors are evaluated against state-of-the-art algorithms.

CCS CONCEPTS

• **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability;

KEYWORDS

eye movement, detection, simulation, fixation, saccade, smooth pursuit

---

*Produces the permission block, and copyright information
†

---

1 INTRODUCTION

Eye movements hold valuable information about a subject, her intension and cognitive states [Braunagel et al. 2017; Kübler et al. 2017] and are also important for the diagnosis of defects and diseases of the eyes (many examples can be found in [Leigh and Zee 2015]). Therefore, the detection and differentiation of eye movement types has to be accurate. Most algorithms for eye movement detection apply different dispersion, velocity or acceleration thresholds and validate the detected eye movements based on their duration. This approach seems to be unsatisfactory [Andersson et al. 2017] at its current state. This is partially due to instable or even dynamic sampling rates of eye tracking devices, task specific sources of noise, the interpolation method applied to the data by the eye tracker, and several more [Cornelissen et al. 2002; Duchowski 2002]. Depending on the task at hand, different thresholds are proposed in the literature [Holmqvist et al. 2011]. It is especially difficult to adjust these thresholds for inconsistent sampling rates and noise which is not annotated by the eye tracker. Some commercial eye-tracker differ between tracking the eye and pupil and re-detecting them after a tracking loss, where the latter requires significantly more processing time and thus results in a decreased frame rate. Therefore, the identification of eye movements is still a difficult task; it complicates to confidently generalize research findings across experiments [Andersson et al. 2017].

Classifying eye movements is the process of separating different intervals in the gaze data to certain oculomotor and cognitive processes. For example visual perception during a saccade is severely limited [Kliegl and Olson 1981; Rayner 1998]. In constast to these very fast movements, perception is working during the (much slower) pursuit of a moving object [Rashbass 1961]. Another important eye movement is blinking. While it is not primarily a

movement of the eyeball, the visual intake is limited before the closing and after the opening [Volkmann et al. 1980]. Therefore, this part has to be marked in the eye tracking data [Holmqvist et al. 2011]. Another interesting event measurable in modern high-speed eye trackers are post-saccadic oscillations. The cornea, a crystalline lens, is deformed during a saccade, thus influencing the pupil center estimation through the distortion of the pupil [Hooge et al. 2015; Nyström et al. 2015; Tabernero and Artal 2014]. This event is not grouped to fixations nor saccades [Nyström and Holmqvist 2010]. During this event the subject can perceive but with distortions [Tabernero and Artal 2014].

Up to now the choice of detection algorithm and parameters is up to the researcher, relying on literature values or the unfeasible task to annotate the data manually. The laborious and difficult task when using an algorithm is to adjust its parameters. Unfortunately, theoretically this process has to be repeated multiple times as the quality of the eye-tracking data often varies from subject to subject and between different tasks. If researchers want to analyze novel eye movements for which no gold standard algorithm exists, they have no choice but to annotate the data manually. With the proposed approach it is possible to create detectors theoretically even for yet unknown eye movements. Therefore, we propose to create detectors based on randomly generated binary decisions. We included ten different types of binary decision of which the final detector selects sets. Those sets learn a conditional distribution and can be combined to a single detector. This machine learning approach is random ferns [Ozuysal et al. 2010]. We also propose an eye movement simulator to generate data similar to the data of the eye-tracker and to create a detector based on the simulation. This also enables to create detectors for very specific events such as skewed saccades.

2 RELATED WORK

2.1 Eye movement simulation

While there is well-funded knowledge about the gaze signal itself, its synthesis is still challenging. In the Eyecatch [Yeo et al. 2012] simulator, a Kalman filter is used to produce a gaze signal for saccades and smooth pursuits. While the signal itself was similar to real eye-tracking recordings, the jitter was missing. The first approach for rendering realistic and dynamic eye movements was proposed in [Lee et al. 2002], where the main focus was on saccadic eye movements. It also included smooth pursuits, binocular rotations (vergence) and the combination of eye and head rotations. The first data-driven approaches where proposed in [Ma and Deng 2009] and [Peters and Qureshi 2010]. Both simulate head and eye movements together in order to generate eye-tracking data. The main disadvantage of [Ma and Deng 2009] was that head motion seemed to trigger eye movement. In fact, the head orientation is only changed if the necessary amplitude of the eye is larger than a specific threshold [Murphy and Duchowski 2002] ($\approx 30°$). Another data-driven approach was proposed in [Le et al. 2012], where an automated framework for head motion, gaze and eyelid simulation was developed. The framework generates data based on speech input using trained Gaussian Mixture Models. While this approach is capable of synthesizing non linear data, it only generates unperturbed gaze directions. The approach in [Duchowski and Jörg 2015] models eye rotations using specific eye related quaternions

for oculomotor rotations as proposed in [Tweed et al. 1990]. The main disadvantage of this approach is that the synthetic eyes cannot be rotated automatically. The approach in [Wood et al. 2015] produces gaze vectors and eye images to train machine learning approaches for gaze prediction, but does not synthesize realistic eye movements.

All of the afore mentioned approaches have their origin in computer graphics with the goal to generate visually realistic head movement and gaze data. The main application of those simulators are to produce realistic interacting virtual humans using parametric models [Andrist et al. 2012; Pejsa et al. 2013]. This leads to the disadvantage, that all movements in the generated data are perfect optimal representatives. In reality an important part of natural eye movements is noise introduces either through actual movements such as microsaccades or inaccuracies of the used eye-tracker. The first approach to simulate a realistic scanpath, i.e., a sequence of fixations and saccades, on static images was proposed in [Campbell et al. 2014]. They use a saliency map together with a unified Bayesian model to generate realistic random walks over a stimulus. A pure gaze data simulation approach including noise was proposed in [Duchowski et al. 2015]. Based on this approach, [Duchowski et al. 2016] further improves the noise synthesis by simulating jitter as a normal distribution.

2.2 Detection algorithms

The most prominent fixation and saccade detection algorithm is Identification by Dispersion-Threshold (IDT) [Salvucci and Goldberg 2000]. It uses the data reduction proposed in [Widdel 1984]. The algorithm uses two thresholds, one is for the maximum fixation dispersion and the other for the minimum fixation duration. Another simple to implement algorithm is the Identification by Velocity Threshold (IVT) [Salvucci and Goldberg 2000], where each sample below a chosen velocity threshold is classified as fixation and above as saccade. It is mostly applicable for high speed recordings. Based on the IVT algorithm, a self-adaptive approach was proposed in [Engbert and Kliegl 2003; Engbert and Mergenthaler 2006], where it was developed to detect microsaccades. In that approach, the velocity threshold is automatically adapted to the noise level in the eye-tracking data. An algorithm especially designed to cope with noisy data is the Identification by Kalman Filter (IKF) algorithm [Komogortsev and Khan 2009]. It uses the Kalman filter to predict the next sample value based on previous values. Therefore, it interpolates the data in an online fashion. For classification, two thresholds are used: one for the predicted value (velocity or distance) and one for the minimum fixation duration. Similar to this algorithm an implementation using the ÏǦ2-test instead of the Kalman filter was proposed in [Komogortsev et al. 2010]. In [Veneri et al. 2011], the Covariance Dispersion Algorithm (CDT) was proposed. It is an improvement of the F-tests dispersion algorithm (FDT) [Veneri et al. 2010]. The F-test measures if two data samples belong to the same class and due to the assumption that the data follows normal distributions it is sensitive to noise. The improvement by the covariance matrix is introduced to cope with this problem. The algorithm needs three thresholds, one for the variance, one for the co-variance, and a third threshold for the minimum duration. The identification by a Minimal Spanning Tree

(IMST) [Komogortsev et al. 2010] creates a tree upon the data, where the samples represent the leafs. The goal is to select all samples with a minimum of branches given a connected graph (the data). Hidden Markov Models (HMM) have been proposed in [Kasneci et al. 2014; Komogortsev et al. 2010; Santini et al. 2016; Tafaj et al. 2012] to separate fixations from saccades and even to detect smooth pursuits. The HMM consists of at least two states (fixation and saccade). For each new velocity sample the model decides whether it belongs to the current state (classification) or a state transition has occurred. After each sample, the model is updated to adapt to the data. The first algorithm that detects post saccadic movements was proposed in [Nyström and Holmqvist 2010]. Based on the noise in the data the algorithm also adapts its velocity thresholds. The Binocular-Individual Threshold (BIT) algorithm [van der Lans et al. 2011] was also designed to detect small saccades in noisy data. Therefore, it applies its thresholds to the data of both eyes, following the ideas that both eyes have to perform the same movement. This algorithm also adapts its thresholds automatically. An algorithm detecting fixations, saccades, post saccadic movement and smooth pursuits was proposed in [Larsson et al. 2013]. This algorithms adapts the parameters automatically and is the first method capable of detecting all these eye movements at the same time. For high-speed eye-tracking data an algorithm for fixation, saccade and smooth pursuit detection was proposed in [Larsson et al. 2015]. The algorithm uses three stages to classify the data, starting with a preliminary segmentation and then evaluating each segment again, followed by the final classification.
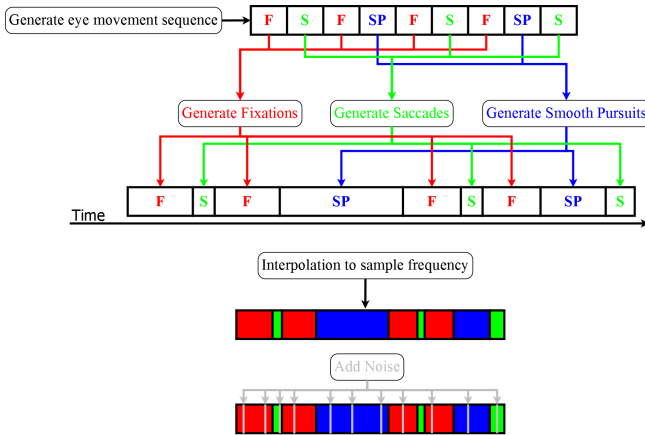
3 SIMULATION

Figure 1: Work-flow of generating eye movement data. First, a sequence of eye movement types is generated. In the second step, a model of each eye movement type is generated (F:Fixation, S:Saccade, SM:Smooth pursuit). This model allows for an almost infinite sampling rate, which is in the next stage interpolated to a target sampling rate (Red:Fixation, Green:Saccade, Blue:Smooth pursuit). Finally, noise is added on top of the signal (gray).

The entire work-flow of the simulator is shown in Figure 1. Generating an eye movement velocity profile is done in four steps. The first step chooses a sequence of eye movement types (Fixation, Saccade, Smooth pursuit) without any time or velocity constrains. Afterwards, each movement type in this sequence is assigned a velocity profile generated by preliminary set parameters. The mathematical model behind these profiles allows sampling at an extremely high, almost arbitrary rate. The target sampling rate is obtained by interpolating this frequency, which also allows for dynamically adjusting the target sampling rate. In the last step, noise is added which represents measurement errors. Each step of this eye movement simulator is described in the following subsections in more detail. The simulator also includes a random walker generator to model fixation direction [Engbert et al. 2011]; saccade and smooth pursuit directions are generated randomly (but consistently within a movement) since this are stimuli- and task-dependent.

3.1 Eye movement sequence

Generating a sequence of eye movement types can be done either by sampling from a uniform distribution, setting it manually, or by following construction constrains. In case of the uniform distributed eye movements, the generator script randomly selects between three types of eye movements. If the amount of each type is specified a priori, the probability is automatically adjusted. This means that after each insertion the probabilities are computed based on the remaining quantity of each type to favor higher quantities. This process can also be constrained, e.g., by forcing the algorithm to insert a saccade after each fixation or before a smooth pursuit.

3.2 Fixation

Fixations are generated based on two probability distributions which can be specified and parametrized. The first distribution determines its duration, the second the consistency of the fixation. For the duration and consistency the minimum and maximum can be set. As distributions, the simulator provides Normal and Uniform random number generation. For the Normal distribution, the standard deviation can be specified. consistency describes the fluctuations in the velocity profile and is used as such in the entire document.
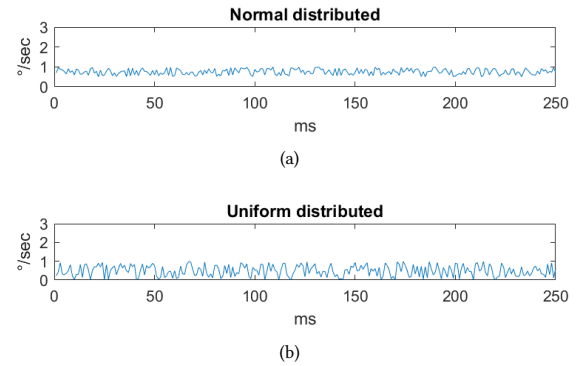

Figure 2: Generated fixation based on a Normal (a) and Uniform (b) distribution.

In Figure 2, two artificially generated fixations are shown. The consistency was set to one degree per second and the standard deviation for the normal distribution to two (Figure 2 (a)). As can be seen in the figure, the Uniform distribution looks more similar compared to real data although we have set the consistency very high with one degree per second.

3.3 Saccade

The most complex part of the eye movement generator are the saccades. For the length, we follow the same approach as for the fixations, in which a minimum and maximum length has to be set. The selectable distributions are Normal and Uniform. The result of the length also influences the maximum speed of the saccade. Therefore, the two random numbers are multiplied (both in the range between zero and one). This means that shorter saccades are limited to lower maximal velocities. To generate the velocity profile, minimum, maximum and the distribution tyoe have to be set.

The most characteristic property of a saccade is its velocity profile. In our simulator this is generated as a Gamma distribution. Therefore, the minimum and maximum skewness has to be specified. In [Van Opstal and Van Gisbergen 1987] it was found that the Gamma function can be considered suitable to approximate saccade profiles (yet not perfect). To achieve more realistic data, a consistency minimum, maximum and distribution can be specified. This generates the jitter along the velocity profile.

Figure 3 shows some generated saccades of fixed length length. We simulated two large and two slightly left skewed saccades. The maximum velocity was selected from a range between 300 and 500 degrees per second. As can be seen from the Figure, the profile contains on- and offset of a saccade. The profile itself is smooth and follows the Gamma distribution. Post-saccadic movement is as of now missing in the simulator. In Figure 3(b) and (d), a small amount of jitter was added to simulate measurement inaccuracy. This usually occurs through the approximation on image pixels or ellipse fit inaccuracy in pupil detection.

3.4 Smooth pursuit

For generating smooth pursuits we also simulate the onset following the findings in [Ogawa and Fujita 1998]. The authors did not provide a final function for the description of the velocity profile but visualized and described it precisely. The shape of the onset of a smooth pursuit follows a non linear growing function similar to the sigmoid function. While this equation is not scientifically proven, our framework allows to simply replace it once a better model is available. The most complex part of the pursuit model is the onset, followed by a regular movement.

The parameters that can be specified are the minimum and maximum length together with their distribution type. For the velocity and the length of the onset the same parameters can be adjusted. To include the measuring error, the consistency parameters are also configurable. For the pursuit itself we included linear growing, decreasing and constant profiles. In case of the growing, again the minimum, maximum and consistency function can be specified. Figure 4 shows simulated smooth pursuits. For the visualization of the linear decreasing and increasing function, extreme values were used. The first column shows a smooth pursuit for a constantly
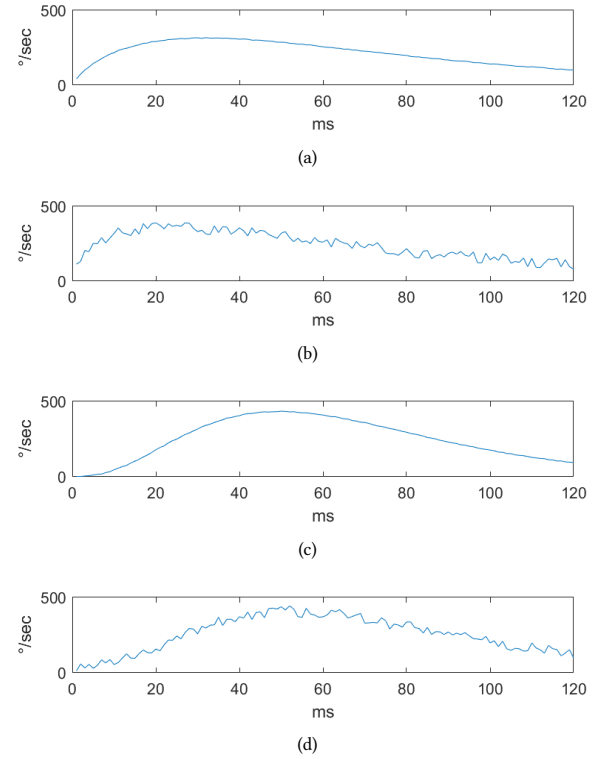
(a)

(b)

(c)

(d)

Figure 3: Generated saccades with jitter (b,d) and without (a,c). For (a) and (b), the distribution was skewed to the left. In (c) and (d), Gamma distribution was only slightly skewed.

moving object, which is often observed in laboratory experiments. The increasing and decreasing profiles are for objects which move further away or come closer to the subject with a constant speed. Other profiles may occur in real settings too, where the object has a slightly varying speed but these are future extensions of the generator and not part of this paper.

3.5 Sampling

After generating and linking the eye movements, they have to be interpolated to a sampling rate. This is necessary to simulate different recording frequencies. Here it is important to mention that not all modern eye trackers record at a constant frequency. On the one hand image acquisition rates can vary depending on illumination changes that affect the aperture time of the camera and timestamps generated by the eye-tracker can vary in accuracy. On the other hand, image processing time, e.g. for eye and pupil detection, are not necessarily constant and might change depending on how easy the pupil can be identified. For example, detection of the pupil is usually more time-consuming than keeping track of a previously detected pupil. Some systems, especially when running on mobile devices, may run into a state where frames are dropped in order to maintain real time performance. We found systems where the timestamps are generated by the CPU time (which may be inaccurate for fast sampling rates) and even timestamps that
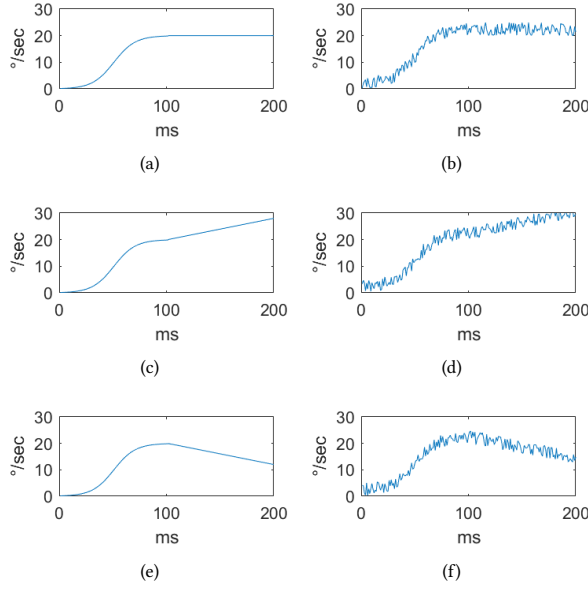
Figure 4: Generated smooth pursuits with jitter (b,d,f) and without (a,c,e). For (a) and (b), the pursuit movement was constant. In (c,d) and (e,f) it was linear increasing and decreasing.
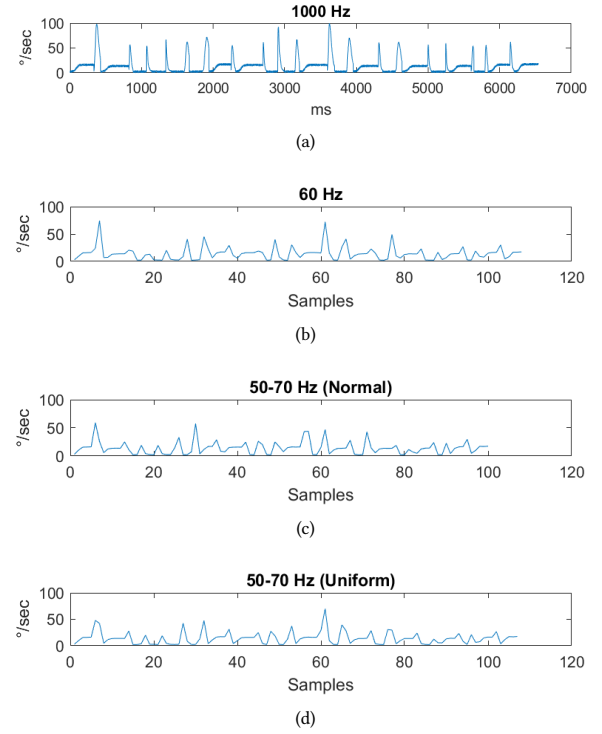

Figure 5: Generated velocity profile of an eye movement sequence (a). In (b), the data is sampled at 60Hz without variations. (b) and (c) vary between 50 and 70 Hz with the Normal and the Uniform distribution.

are generated after image processing. Therefore, our simulator is capable of simulating varying sampling rates. The parameters for this step are the minimum and maximum sampling rate and also the consistency function. The interpolation itself computes the mean of all values from the last sampling position to the new sampling position.

In Figure 5(a) a generated velocity profile is shown. The initial sampling frequency was set to 1000 Hz but any other sampling rate is possible. For (b), a constant sampling frequency of 60 Hz was used. In (c), the sampling frequency varies between 50 and 70 Hz (with a mean of 60 Hz), wherein the Normal distribution was used as random number generator. It differs significantly from the constant sampling rate in (a) and also has a different length. For (d), the sampling frequency also varied between 50 and 70 Hz with the difference that the Uniform distribution was used as random number generator. The length is therefore similar to the constant sampling rate but it still differs especially for the saccadic peeks.

3.6 Noise

For generating noise, two distributions are used: one for the location where to place the noise in the data and the second for the velocity change to apply. Therefore, the user has to specify the types for both distributions and the minimum and maximum velocity of noise. The amount of noise is specified as a percentage of the samples that should be influenced.

Figure 6 shows two types of Noise added to the velocity profile shown in (a). The amount of noise added was 10%. For the Normal distributed noise in (b) it can be seen that the peaks are mostly high.

In comparison to it, the Uniform distributed noise in (c) produces more peaks of different heights.

4 DETECTOR CREATION

In state-of-the-art algorithms there are thresholds for upper and lower limits as well as for ranges which have to be fulfilled. The main disadvantage is that those thresholds are difficult to adjust to new data, where the sampling rate is not constant or no time information is given [Andersson et al. 2017]. Another issue with those thresholds is that for some data they work very well while for more noisy data they do not work at all or need intensive preprocessing (such as smoothing filters and outliers detection). Our idea is to use the traditional thresholding approach but to adapt the algorithm to the data. The first step in our algorithm is to randomly generate different types of thresholding approaches and thresholds. The following binary decisions are generated:

$$
\begin{array}{ll}
|P1 - P2| < TH1 & |P1 - P2| > TH1 \\
P1 < TH1 \; and \; P2 < TH2 & P1 < TH1 \\
P1 > TH1 \; and \; P2 > TH2 & P1 > TH1 \\
P1 < TH1 \; and \; P2 > TH2 & P2 < TH2 \\
P1 > TH1 \; and \; P2 < TH2 & P2 > TH2,
\end{array}
$$

where $P1$ and $P2$ are two samples of the generated sequence, e.g., two velocities. These points are not required to be sequential, in
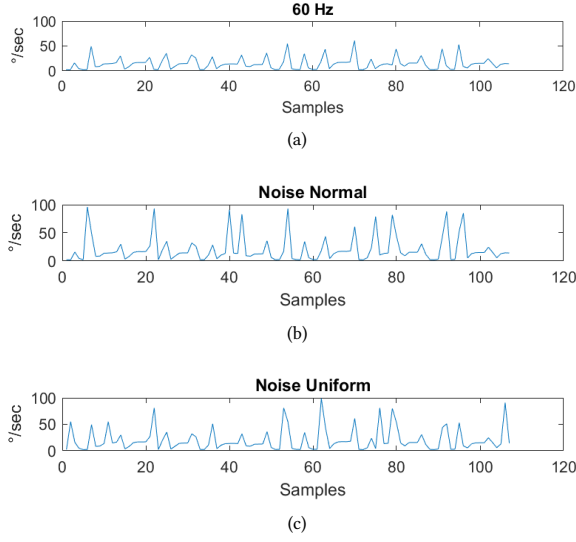
Figure 6: Generated velocity profile of an eye movement sequence (a). In (b), noise is added based on a Normal distribution and in (c) a Uniform distribution was used.

fact $P1$ can be earlier or later than $P2$ in the sequence of samples. Their relative offsets to the sample that is currently in consideration for being classified is generated randomly. Therefore, the binary decisions consist of up to two distances to the current classification position and up to two thresholds ($TH1$, $TH2$). Based on the two distances, the sample positions $P1$ and $P2$ are computed. For those distances there is also the option to restrict them to samples preceding the current position, so that classification can be performed online, without the knowledge of future samples. An example of
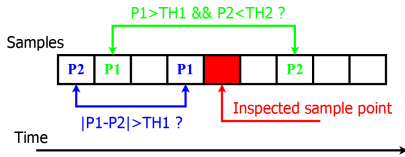

Figure 7: Two binary decisions (green and blue) for the inspected sample position (red).

such binary decisions is shown in Figure 7. Red is the inspected sample for which a decision is to be made. The two binary decisions are colored blue and green. Blue is an example for an online applicable decision and green would be applied to an already complete recording or delayed to the recording. As can be seen, for each generated binary decision five parameters have to be selected (the fifth being the equation to use). A scan of all possible combinations of each parameter value would be far too extensive. Therefore, we generate hundreds of thousands randomly. Key to a good decision making is the selection of those parameter combinations that are relevant for determining an eye movement type. This means that we can assign a higher feature quality value to decisions that result in a true binary condition for a certain detection task, and vice

versa. Afterwards, only the top ten percent with the highest quality values assigned to them are further used.

The next step is to combine these binary decisions to a detector. This is done by computing a conditional distribution for a randomly selected set of ten binary decisions (B).

$$D(X, B) = p(X|B) \tag{1}$$

Equation 1 describes this conditional distribution. X is the current sample point, B are the binary decisions and p is the probability of the sample to be an eye movement type. Figure 8 shows the
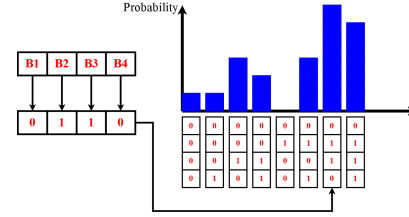

Figure 8: The binary decisions (red) represent an index in the distribution (blue) which holds the probability.

conditional distribution p in blue. Each binary decision represents a digit, that is used as index in the conditional distribution. This combines multiple binary decisions to one detector. For training we use two distributions, one for valid examples and one for wrong examples. This allows to compute the final distribution without negative probabilities or to stop at zero. The difference between both distributions is the final score for a sample ($p_{positive}(X|B) - p_{negative}(X|B)$). The computation of each distribution is a simple lookup by the binary decisions index number and increasing the respective histogram index. The increment has to be normed to equalize the amount of positive and negative samples (the ratio between both occurrences in the training set).

After training of the conditional distributions for the so-called ferns, we have thousands of randomly selected weak detectors available. To create one strong detector we combine multiple of these, again randomly. Therefore, we compute a quality score for each fern similar as it was done for the binary decisions, and again consider only the top ten percent. Afterwards, we randomly select ten ferns and combine them under the independence assumption to one strong detector.

$$C(X, D_1, ..., D_n) = D_1(X, B_1) * ... * D_n(X, B_n) \tag{2}$$

Equation 2 describes this computation where $D_1, ..., D_n$ are the ferns and $B_1, ..., B_n$ are their binary decisions. For each eye movement type, we randomly generate hundreds of such detectors and score them as described for the binary decisions. After scoring, we select the top ten percent and evaluate them in combination. For combining the classifiers of different eye movement types, we consider only the combinations of equally ranked classifiers, e.g. the best fixation classifier with the best saccade classifier.

$$Type(X) = \begin{cases} Fixation & C_1 > C_2, C_3, C_4 \\ Saccade & C_2 > C_1, C_3, C_4 \\ Smooth\ Pursuit & C_3 > C_1, C_2, C_4 \\ Noise & C_4 > C_1, C_2, C_3 \end{cases} \tag{3}$$

The final result is obtained using Equation 3, where the highest probability decides which type is detected. The best combination of individual movement classifiers for the evaluation on the training set is selected as final detector. This also shows that it is easy to extend the detector for novel eye movement types or to train it only for specific events like the beginning of a movement or movement combinations such as regressions during reading.

5 EVALUATION

For the entire evaluation we transformed the input signal to a sample per sample velocity signal. This was done to have an extremely challenging signal to simulate and for the detection. Most of the state-of-the-art algorithms apply smoothing or compute the velocity from multiple sequential samples. While such methods can always be applied the purpose of this evaluation is to show that our approach can adapt to the noise level even for challenging conditions. This sample per sample velocity signal is used in the entire evaluation. The evaluated data sets of annotated eye movements are chosen from [Dorr et al. 2010; Larsson et al. 2013; Santini et al. 2016] and contain multiple annotators. We evaluated each annotation separately, meaning that each annotator was evaluated as ground truth independent of the others.

The first subsection presents the evaluation of the proposed simulator, where examples are given showing real velocity profiles of recorded saccades from publicly available data sets. In the second part of this section, the trained detectors are evaluated and compared to state-of-the-art algorithms.

5.1 Evaluation of the simulation

Figure 9 shows the per sample point squared error as whisker plots of our simulator in comparison to the publicly available data sets. The error was computed based on the squared difference between each sample. Therefore, we simulated each fixation, saccade, and smooth pursuit ten times with the same length as in the available data sets. For a fixation, the simulator got the information of the mean velocity and the standard deviation to generate a profile. The information of a saccade was the peak velocity and the position of this peak. For smooth pursuits, the simulator got the information of the mean velocity and the standard deviation.

As can be seen in Figure 9(b), the error for saccades was the largest. This is due to the noisy signal which comes from the sample per sample velocity. Figure 10 shows some saccades which produced high squared errors. The red line corresponds to the simulation result, whereas the blue line corresponds to the real data. As can be seen, the course of the velocity profile is well simulated, which is well in line with previous findings in [Van Opstal and Van Gisbergen 1987]. The high errors originate mainly from measurement inaccuracies in the real data. This also highlights the difficulty in detecting eye movements in such a signal. For the data set from [Santini et al. 2016](I-BDT), the error for saccades was lowest. This is due to the low sampling rate of the used eye tracker (30 Hz), for which large fluctuations do not occur. This is similar to smoothing or using multiple samples for the velocity computation. In contrast, the smooth pursuits error was the largest in the I-BDT data set. This is because in such low sampling rates the onset of a smooth pursuit is hardly represented. Our simulator is capable

of simulating this (sampling 3.5) but for the evaluation it was not used. We only used the generators to simulate the eye movements.

5.2 Evaluation of the detectors

For comparison, we adapted the algorithms from [Santini et al. 2016] (I-BDT), [Nyström and Holmqvist 2010] (EV), and [Larsson et al. 2013] (LS) to work with the velocity profile instead of x, y coordinates. We chose those algorithms because all come together with a self-adaptation procedure. I-BDT initializes itself on a part of the data and continues updating its probability distributions during runtime. For initialization, in the evaluation we provided the algorithm with the computed mean and standard deviation of the data it was evaluated on instead of the initial 15 seconds as done in the original implementation. The algorithm EV automatically adjusts its thresholds. Therefore, we provided it with the appropriate computed values for minimum velocity, maximum velocity, minimum duration, maximum duration, mean noise velocity etc. from the data it was evaluated on. LS is the representative for a segmentation-based self-adapting algorithm. We provided the statistical data similar to EV. As for I-BDT and EV, this means that for each evaluation we computed the statistics for the algorithms as initialization of their parameters based on the data they are evaluated on. This was done to simulate a handcrafted initialization. For the proposed approach we used simulated data to train and select a detector which was afterwards evaluated on the annotated data set. We did not use any post- or preprocessing of the data, nor segmentation or similar.

Data	Alg.	Detection Rate (%)			
		Fixation	*Saccade*	*Pursuit*	*Noise*
Larsson	EV	92.47	4.37	0.00	35.11
	IBDT	0.10	0.01	96.32	0.00
	LS	23.73	72.69	56.54	3.98
	Proposed	61.86	76.56	61.88	23.43
GazeCom	EV	94.86	0.78	0.00	43.19
	IBDT	0.03	0.02	99.85	0.00
	LS	49.18	70.79	43.18	5.86
	Proposed	70.55	70.40	51.60	20.27
I-BDT	EV	92.47	0.61	0.00	71.46
	IBDT	22.48	24.22	97.29	0.00
	LS	29.50	13.04	93.14	10.44
	Proposed	72.50	69.44	64.83	20.88

Table 1: Percentage of correctly detected samples per data set.

Table 1 shows the correctly detected results per data set for all evaluated algorithms. As can be seen, our approach results in constantly balanced accuracies for fixations, saccades, and smooth pursuits for all data sets without having seen any of the real recordings. The other algorithms tend to prefer different types of eye movement. I-BDT for example cannot handle high speed recordings because the probability for smooth pursuits dominates. This is additionally supported by noise in the data which the algorithm is not designed to detect. The probably best performing competitor is LS [Larsson et al. 2013]. For EV [Nyström and Holmqvist 2010] it
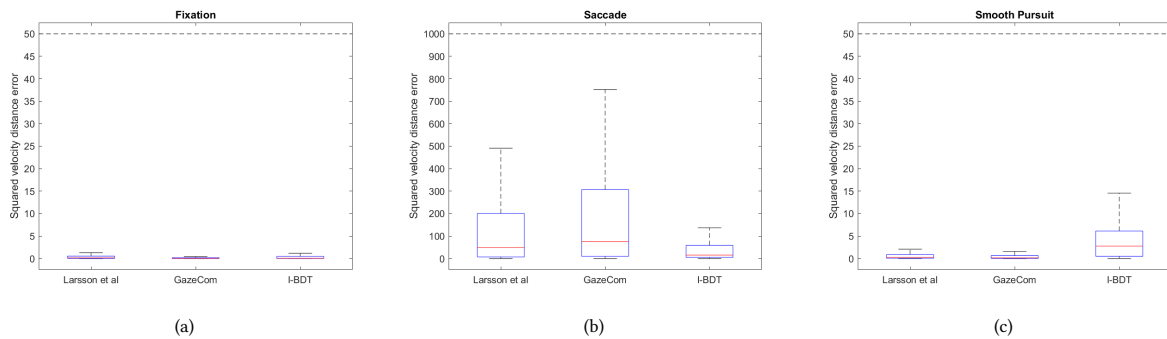
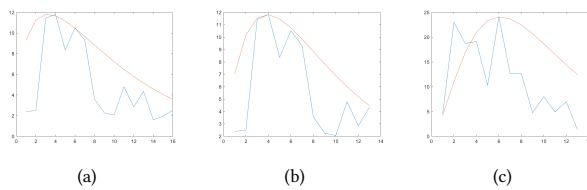Figure 9: Squared velocity error for the simulation per data set.

Figure 10: Saccades with high squared error. Red is the simulation and blue the real data.

has to be mentioned, that it does not detect smooth pursuits at all, significantly simplifying the classification problem as the probably hardest to classify class [Komogortsev and Karpov 2013] (where velocities are *somewhere in the middle*) is left out. Still noise and saccades are often confused.

The results of our detector could simply be improved by applying post processing to avoid too short or too long durations. The detectors could also be selected by evaluating combinations on the training set not only the ones which are ranked equally. Another improvement is to detect only the start and ending points of eye movements and set the data in between accordingly (similar to the segmentation in [Larsson et al. 2013]) but this is out of scope of this paper and will be part of further research.

6 CONCLUSION

We proposed a novel eye movement detection approach which is based on machine learning. It is capable of training detectors for specific eye movements and expendable to new findings. The detectors are capable of outperforming the state-of-the-art and adaptable to new challenges from new eye trackers. In addition the detectors can be trained for offline and online analysis enabling a second validation and refinement stage for eye movement detection. The underlying simulator, which generates the training data is based on scientific findings to generate the velocity profiles of eye movements. It is capable of simulating any static or dynamic sampling rate and allows to select different distributions for noise, sampling shift, eye tracker accuracy etc. Further research will be the extension of the simulator to be also capable of generating post saccadic, optokinetic and vestibulo-ocular movement. Additionally,

a validation and correction extension will be developed to refine eye movement data based on known velocity profiles and validate gaze positions.

REFERENCES

Richard Andersson, Linnea Larsson, Kenneth Holmqvist, Martin Stridh, and Marcus Nyström. 2017. One algorithm to rule them all? An evaluation and discussion of ten eye movement event-detection algorithms. *Behavior Research Methods* 49, 2 (2017), 616–637.

Sean Andrist, Tomislav Pejsa, Bilge Mutlu, and Michael Gleicher. 2012. Designing effective gaze mechanisms for virtual agents. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. ACM, 705–714.

Christian Braunagel, David Geisler, Wolfgang Rosenstiel, and Enkelejda Kasneci. 2017. Online Recognition of Driver-Activity Based on Visual Scanpath Classification. *IEEE Intelligent Transportation Systems Magazine* 9, 4 (2017), 23–36.

Daniel J Campbell, Joseph Chang, Katarzyna Chawarska, and Frederick Shic. 2014. Saliency-based bayesian modeling of dynamic viewing of static scenes. In *Proceedings of the Symposium on Eye Tracking Research and Applications*. ACM, 51–58.

Frans W Cornelissen, Enno M Peters, and John Palmer. 2002. The Eyelink Toolbox: eye tracking with MATLAB and the Psychophysics Toolbox. *Behavior Research Methods, Instruments, & Computers* 34, 4 (2002), 613–617.

Michael Dorr, Thomas Martinetz, Karl R Gegenfurtner, and Erhardt Barth. 2010. Variability of eye movements when viewing dynamic natural scenes. *Journal of vision* 10, 10 (2010), 28–28.

AT Duchowski and S Jörg. 2015. Modeling Physiologically Plausible Eye Rotations. *Proceedings of Computer Graphics International* (2015).

Andrew Duchowski, Sophie Jörg, Aubrey Lawson, Takumi Bolte, Lech Świrski, and Krzysztof Krejtz. 2015. Eye movement synthesis with 1/f pink noise. In *Proceedings of the 8th ACM SIGGRAPH Conference on Motion in Games*. ACM, 47–56.

Andrew T Duchowski. 2002. A breadth-first survey of eye-tracking applications. *Behavior Research Methods, Instruments, & Computers* 34, 4 (2002), 455–470.

Andrew T Duchowski, Sophie Jörg, Tyler N Allen, Ioannis Giannopoulos, and Krzysztof Krejtz. 2016. Eye movement synthesis. In *Proceedings of the Symposium on Eye Tracking Research and Applications*. ACM, 147–154.

Ralf Engbert and Reinhold Kliegl. 2003. Microsaccades uncover the orientation of covert attention. *Vision research* 43, 9 (2003), 1035–1045.

Ralf Engbert and Konstantin Mergenthaler. 2006. Microsaccades are triggered by low retinal image slip. *Proceedings of the National Academy of Sciences* 103, 18 (2006), 7192–7197.

Ralf Engbert, Konstantin Mergenthaler, Petra Sinn, and Arkady Pikovsky. 2011. An integrated model of fixational eye movements and microsaccades. *Proceedings of the National Academy of Sciences* 108, 39 (2011), E765–E770.

Kenneth Holmqvist, Marcus Nyström, Richard Andersson, Richard Dewhurst, Halszka Jarodzka, and Joost Van de Weijer. 2011. *Eye tracking: A comprehensive guide to methods and measures*. OUP Oxford.

Ignace Hooge, Marcus Nyström, Tim Cornelissen, and Kenneth Holmqvist. 2015. The art of braking: Post saccadic oscillations in the eye tracker signal decrease with increasing saccade size. *Vision research* 112 (2015), 55–67.

Enkelejda Kasneci, Gjergji Kasneci, Thomas C Kübler, and Wolfgang Rosenstiel. 2014. The applicability of probabilistic methods to the online recognition of fixations and saccades in dynamic scenes. In *Proceedings of the Symposium on Eye Tracking Research and Applications*. ACM, 323–326.

Reinhold Kliegl and Richard K Olson. 1981. Reduction and calibration of eye monitor data. *Behavior Research Methods* 13, 2 (1981), 107–111.

Oleg V Komogortsev, Denise V Gobert, Sampath Jayarathna, Do Hyong Koh, and Sandeep M Gowda. 2010. Standardization of automated analyses of oculomotor fixation and saccadic behaviors. *IEEE Transactions on Biomedical Engineering* 57, 11 (2010), 2635–2645.

Oleg V Komogortsev and Alex Karpov. 2013. Automated classification and scoring of smooth pursuit eye movements in the presence of fixations and saccades. *Behavior research methods* 45, 1 (2013), 203–215.

Oleg V Komogortsev and Javed I Khan. 2009. Eye movement prediction by oculomotor plant Kalman filter with brainstem control. *Journal of Control Theory and Applications* 7, 1 (2009), 14–22.

Thomas C Kübler, Colleen Rothe, Ulrich Schiefer, Wolfgang Rosenstiel, and Enkelejda Kasneci. 2017. SubsMatch 2.0: Scanpath comparison and classification based on subsequence frequencies. *Behavior Research Methods* 49, 3 (2017), 1048–1064.

Linnéa Larsson, Marcus Nyström, Richard Andersson, and Martin Stridh. 2015. Detection of fixations and smooth pursuit movements in high-speed eye-tracking data. *Biomedical Signal Processing and Control* 18 (2015), 145–152.

Linnéa Larsson, Marcus Nyström, and Martin Stridh. 2013. Detection of saccades and postsaccadic oscillations in the presence of smooth pursuit. *IEEE Transactions on Biomedical Engineering* 60, 9 (2013), 2484–2493.

Binh H Le, Xiaohan Ma, and Zhigang Deng. 2012. Live speech driven head-and-eye motion generators. *IEEE transactions on Visualization and Computer Graphics* 18, 11 (2012), 1902–1914.

Sooha Park Lee, Jeremy B Badler, and Norman I Badler. 2002. Eyes alive. In *ACM Transactions on Graphics (TOG)*, Vol. 21. ACM, 637–644.

R John Leigh and David S Zee. 2015. *The neurology of eye movements*. Vol. 90. Oxford University Press, USA.

Xiaohan Ma and Zhigang Deng. 2009. Natural eye motion synthesis by modeling gaze-head coupling. In *IEEE Virtual Reality Conference*. IEEE, 143–150.

Hunter Murphy and Andrew T Duchowski. 2002. Perceptual gaze extent & level of detail in VR: looking outside the box. In *ACM SIGGRAPH conference Abstracts and Applications*. ACM, 228–228.

Marcus Nyström, Richard Andersson, Måns Magnusson, Tony Pansell, and Ignace Hooge. 2015. The influence of crystalline lens accommodation on post-saccadic oscillations in pupil-based eye trackers. *Vision research* 107 (2015), 1–14.

Marcus Nyström and Kenneth Holmqvist. 2010. An adaptive algorithm for fixation, saccade, and glissade detection in eyetracking data. *Behavior research methods* 42, 1 (2010), 188–204.

Tadashi Ogawa and Masahiko Fujita. 1998. Velocity profile of smooth pursuit eye movements in humans: pursuit velocity increase linked with the initial saccade occurrence. *Neuroscience research* 31, 3 (1998), 201–209.

Mustafa Ozuysal, Michael Calonder, Vincent Lepetit, and Pascal Fua. 2010. Fast keypoint recognition using random ferns. *IEEE transactions on Pattern Analysis and Machine Intelligence* 32, 3 (2010), 448–461.

Tomislav Pejsa, Bilge Mutlu, and Michael Gleicher. 2013. Stylized and performative gaze for character animation. In *Computer Graphics Forum*, Vol. 32. Wiley Online Library, 143–152.

Christopher Peters and Adam Qureshi. 2010. A head movement propensity model for animating gaze shifts and blinks of virtual characters. *Computers and Graphics* 34, 6 (2010), 677–687.

C1 Rashbass. 1961. The relationship between saccadic and smooth tracking eye movements. *The Journal of Physiology* 159, 2 (1961), 326–338.

Keith Rayner. 1998. Eye movements in reading and information processing: 20 years of research. *Psychological bulletin* 124, 3 (1998), 372.

Dario D Salvucci and Joseph H Goldberg. 2000. Identifying fixations and saccades in eye-tracking protocols. In *Proceedings of the Symposium on Eye Tracking Research and Applications*. ACM, 71–78.

Thiago Santini, Wolfgang Fuhl, Thomas Kübler, and Enkelejda Kasneci. 2016. Bayesian identification of fixations, saccades, and smooth pursuits. In *Proceedings of the Symposium on Eye Tracking Research and Applications*. ACM, 163–170.

Juan Tabernero and Pablo Artal. 2014. Lens oscillations in the human eye. Implications for post-saccadic suppression of vision. *PloS one* 9, 4 (2014), e95764.

Enkelejda Tafaj, Gjergji Kasneci, Wolfgang Rosenstiel, and Martin Bogdan. 2012. Bayesian online clustering of eye movement data. In *Proceedings of the Symposium on Eye Tracking Research and Applications*. ACM, 285–288.

Douglas Tweed, Werner Cadera, and Tutis Vilis. 1990. Computing three-dimensional eye position quaternions and eye velocity from search coil signals. *Vision research* 30, 1 (1990), 97–110.

Ralf van der Lans, Michel Wedel, and Rik Pieters. 2011. Defining eye-fixation sequences across individuals and tasks: the Binocular-Individual Threshold (BIT) algorithm. *Behavior Research Methods* 43, 1 (2011), 239–257.

AJ Van Opstal and JAM Van Gisbergen. 1987. Skewness of saccadic velocity profiles: a unifying parameter for normal and slow saccades. *Vision research* 27, 5 (1987), 731–745.

Giacomo Veneri, Pietro Piu, Pamela Federighi, Francesca Rosini, Antonio Federico, and Alessandra Rufa. 2010. Eye fixations identification based on statistical analysis-case study. In *Cognitive Information Processing (CIP), 2010 2nd International Workshop on*. IEEE, 446–451.

Giacomo Veneri, Pietro Piu, Francesca Rosini, Pamela Federighi, Antonio Federico, and Alessandra Rufa. 2011. Automatic eye fixations identification based on analysis of variance and covariance. *Pattern Recognition Letters* 32, 13 (2011), 1588–1593.

Frances C Volkmann, Lorrin A Riggs, and Robert K Moore. 1980. Eyeblinks and visual suppression. *Science* 207, 4433 (1980), 900–902.

Heino Widdel. 1984. Operational problems in analysing eye movements. *Advances in psychology* 22 (1984), 21–29.

Erroll Wood, Tadas Baltrusaitis, Xucong Zhang, Yusuke Sugano, Peter Robinson, and Andreas Bulling. 2015. Rendering of eyes for eye-shape registration and gaze estimation. In *Proceedings of the IEEE International Conference on Computer Vision*. 3756–3764.

Sang Hoon Yeo, Martin Lesmana, Debanga R Neog, and Dinesh K Pai. 2012. Eyecatch: simulating visuomotor coordination for object interception. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 42.