An adaptive algorithm for fixation, saccade, and glissade detection in eyetracking data

MARCUS NYSTRÖM AND KENNETH HOLMQVIST
Lund University, Lund, Sweden

Event detection is used to classify recorded gaze points into periods of fixation, saccade, smooth pursuit, blink, and noise. Although there is an overall consensus that current algorithms for event detection have serious flaws and that a de facto standard for event detection does not exist, surprisingly little work has been done to remedy this problem. We suggest a new velocity-based algorithm that takes several of the previously known limitations into account. Most important, the new algorithm identifies so-called *glissades*, a wobbling movement at the end of many saccades, as a separate class of eye movements. Part of the solution involves designing an adaptive velocity threshold that makes the event detection less sensitive to variations in noise level and the algorithm settings-free for the user. We demonstrate the performance of the new algorithm on eye movements recorded during reading and scene perception and compare it with two of the most commonly used algorithms today. Results show that, unlike the currently used algorithms, fixations, saccades, and glissades are robustly identified by the new algorithm. Using this algorithm, we found that glissades occur in about half of the saccades, during both reading and scene perception, and that they have an average duration close to 24 msec. Due to the high prevalence and long durations of glissades, we argue that researchers must actively choose whether to assign the glissades to saccades or fixations; the choice affects dependent variables such as fixation and saccade duration significantly. Current algorithms do not offer this choice, and their assignments of each glissade are largely arbitrary.

The goal of *event detection* is to parse gaze coordinates recorded from an eyetracker into periods of fixation, saccade, smooth pursuit, and blink. What does not fit into these categories can be considered noise. Since the duration, prevalence, and direction of such events reflect ongoing cognitive processes in humans, the conversion to events is an essential part in the vast majority of eye movement studies (see, e.g., Irwin, 2004, and the review by Rayner, 1998, pp. 374–375).

Event detection may sound easy, but several works have shown that the most common dependent variables, such as fixation duration, saccadic amplitude, and number of fixations, depend critically on the recording software with its internal filters and on the choice of algorithm, as well as on the settings of the selected algorithm (Blignaut, 2009; Inchingolo & Spanio, 1985; Karsh & Breitenbach, 1983; Salvucci & Goldberg, 2000; Shic, Scassellati, & Chawarska, 2008). This situation makes the output of most algorithms less reliable and many studies less comparable than would be desired.

Another largely unexplored reason that can explain some of the variation in event detection results is the behavior of the eye at the end of many saccades, where the eye sometimes does not stop directly but wobbles a while before coming to a stop. This type of wobbling movement is called a dynamic or *glissadic* overshoot (or undershoot) and, although frequently reported in the literature, is not

explicitly taken into account by event detection algorithms. Glissades are therefore treated unsystematically and differently across algorithms, and even within the same algorithm; one glissade may be assigned to the saccade, whereas the next one, instead, is merged with the fixation.

The aim of this article is to design a velocity-based saccade detection algorithm that addresses many of the limitations of current algorithms. The goal is to develop a data-driven algorithm that leaves as few subjective settings to the end user as possible. The improvements include *data-driven thresholds* for peak and saccade onset/offset detection, *adaptive threshold adjustment* based on local noise levels, physical constraints on eye movements to *exclude noise and jitter*, and new recommendations for *minimum allowed fixation and saccade duration*. Since it is unclear how attention and visual intake are correlated with eye movement in the general case, the proposed algorithm is designed to identify *oculomotor* events in eye movement data and leave the important question of related cognitive processes to the individual researcher.

The main novelty of this algorithm, however, is the option to *define the glissadic movement as a separate event*; we demonstrate this by investigating the prevalence and nature of glissades as they occur during reading and during scene perception. Moreover, we will in-

M. Nyström, marcus.nystrom@humlab.lu.se

vestigate how glissades influence computed saccade and fixation measures. To demonstrate the performance of the new algorithm, it is compared against two of the most commonly used event detection algorithms: the identification by dispersion threshold (I-DT, as defined in Salvucci & Goldberg, 2000) and a velocity-based algorithm (commercial implementation of the algorithm in Smeets & Hooge, 2003).[1]

GLISSADES

Since glissade detection is one of the main novelties in the new algorithm, let us start by taking a look at what they are and how they might influence event detection. Figure 1 plots the *x*-, *y*-coordinates and angular velocity of a portion of reading data. At the end of the two leftmost saccades, one can clearly see how the eye overshoots its intended target and immediately corrects this by a smaller, saccade-like movement. This movement is known as a *dynamic overshoot* (rapid postsaccadic movement; Kapoula, Robinson, & Hain, 1986) or a *glissadic overshoot* (slower postsaccadic movement; Weber & Daroff, 1972). For simplicity, in the remainder of this article, we use the term *glissade* for all types of movements that happen due to over-undershoots.

Whether data exhibit glissades or not depends on a variety of factors, of which one is the recording method. For instance, with video-based pupil and corneal reflex systems, glissades are common, but not in data from scleral search coils. Frens and Van der Geest (2002) found that coils alter the saccade dynamics and smooth out glissadic eye movements, possibly because the coils affect oculomotor command signals. With dual Purkinje image eyetrackers, glissades appear to be amplified as an effect of slippage of the lens, relative to the rest of the eye, at the end of

the saccade (Deubel & Bridgeman, 1995). When recorded data are then filtered, filter properties influence saccade and glissade dynamics (see Inchingolo & Spanio, 1985), and low-pass filtering may heavily reduce the detection of recorded glissades. Stampe (1993) even designs the filters with the explicit intention to remove glissades from data.

Kapoula et al. (1986) found glissades to be largely idiosyncratic and overrepresented in the abducting eye and in small saccades. They also observed that glissades rarely occur simultaneously in both eyes. It also has been reported that saccades accompanied by blinks increase the tendency to make a glissade (Rottach, Das, Wohlgemuth, Zivotofsky, & Leigh, 1998). The general belief has been that glissades are due to "mistakes" in saccadic programming and that they therefore serve no useful purpose (Kapoula et al., 1986).

Whatever their cause or function, it is a fact that glissades affect the result of a fixation and saccade detection algorithm. The best way to deal with this influence is to acknowledge them and prepare the algorithm for their existence in the data.

EVENT DETECTION ALGORITHMS

Given gaze coordinates produced by the eyetracker, an event detector should robustly classify these raw data into events such as fixations, saccades, and smooth pursuit periods, as well as blinks and noise. This may sound fairly straightforward, but the task presents users with several difficulties, among which the selection of threshold is the best known. It has been shown that event detection algorithms are very sensitive to the choice of thresholds that they need to distinguish—for example, a saccade sample from a fixation sample (see Blignaut, 2009; Shic, Chawarska, & Scassellati, 2008).
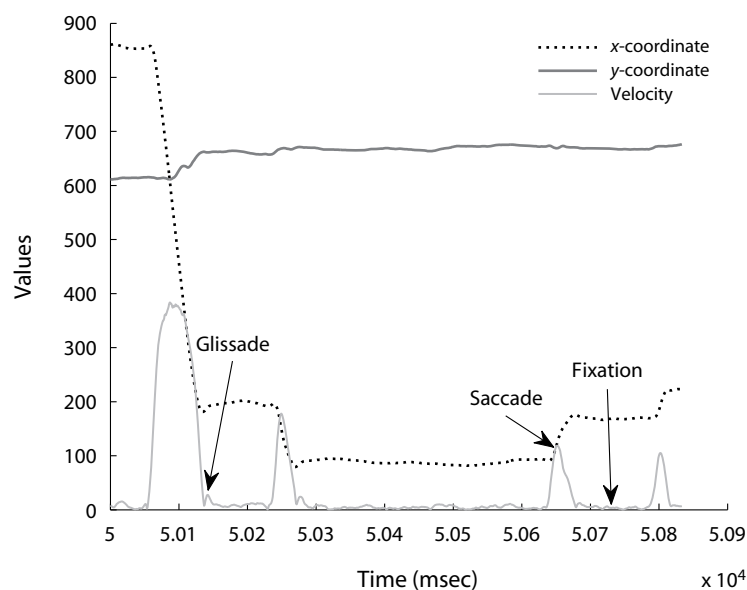

Figure 1. Example of reading saccades ending with glissadic overshoot (first two saccades).

Current Algorithms and Their Shortcomings

Event detection algorithms classify gaze data on the basis of *dispersion*, *velocity*, and *acceleration* criteria (or combinations thereof; see Duchowski, 2003; Salvucci & Goldberg, 2000).

Dispersion-based algorithms typically identify gaze samples as belonging to a fixation if the samples are located within a spatially limited region (about 0.5°) for a minimum period of time: the minimum allowed fixation duration (usually, in the range 80–150 msec). Saccades are then detected implicitly as "everything else" (often excluding blinks and jitter).

The most common dispersion-based algorithm, the I-DT, first initiates a duration window over the gaze data with a length equal to the minimum fixation duration threshold (see Figure 2). This temporal window is then expanded sample by sample until the dispersion of the included (x, y) points exceeds the dispersion threshold. Dispersion is calculated as the average of the largest horizontal and vertical distances (Δx-max and Δy-max) between any two samples within the window. Finally, the fixation location is found as the centroid of all window samples, and at the end of the just-finished fixation, a new window is started where the search for new fixations continues. Dispersion-based algorithms exist in a number of varieties, which differ in how dispersion is calculated and the condition for ending a search.

Algorithms belonging to the I-DT family can be found in commercial software provided by, for example, Applied Science Laboratories (ASL), SensoMotoric Instruments (SMI), and Tobii Technology, and in the analysis software Gazetracker. Although the I-DT algorithm is well established and an integrated part of many commercial softwares, it is very sensitive to noise and drifts in the data and is poor at providing accurate temporal estimates of event onsets and offsets. Recent work has also shown that the output from I-DT is affected not only by the choice of minimum fixation duration and dispersion thresholds, but also on how dispersion is calculated; keeping other parameters fixed, this factor alone can yield significant differences in the number and durations of fixations (Blignaut, 2009).

The simplest type of *velocity-based algorithm* uses one threshold, which is typically set by the experimenter after an initial check on the noise levels and velocities of saccades in data. The velocity threshold decides whether a velocity sample belongs to a fixation (everything below the threshold) or a saccade (everything above the threshold) (Bahill, Brockenbrough, & Troost, 1981; Salvucci & Goldberg, 2000). Although it is very simple, it is rarely, if at all, used in any real implementations, since it is very sensitive to noise around the value of the threshold. Practical implementations of velocity-based algorithms instead use a more sophisticated, multipass strategy in which saccade peaks are *detected* in the first pass and the second pass is used to find the onset and offset of the saccades. For example, Smeets and Hooge (2003) presented such an approach to saccade detection, where they complemented the detection threshold with a simple shape estimator that checks whether the saccade peak is located in the central portion of the detected saccade; otherwise, it is discarded. Smeets and Hooge used a peak saccade detection threshold of 75°/sec. Although velocity-based algorithms are more transparent than dispersion-based algorithms, since they use the instantaneous velocity to find onsets and offsets of saccades, algorithms such as those in Smeets and Hooge can output unphysiologically short fixations (<20 msec), due to noise and glissades present in the data. Another potential source of error is the setting of a peak velocity threshold, which critically decides whether small saccades will be detected or not.

Velocity-based algorithms are sometimes combined with *acceleration criteria* to find saccade onsets and offsets—in particular, if online saccade detection is required. For example, the EyeLink software in its *cognitive configuration* uses velocity, acceleration, and motion thresholds of 30°/sec, 8,000°/sec^2, and 0.15°, respectively (SR Research, 2007). The last threshold is used to ensure that the eye has moved sufficiently before a saccade is
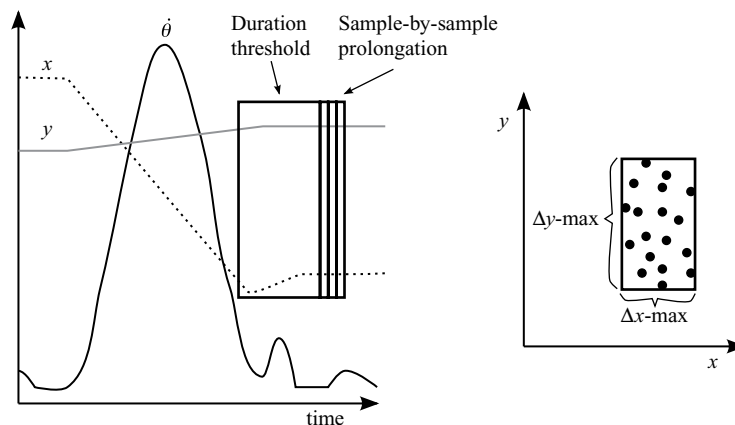

Figure 2. Principle for duration and dispersion calculation in the identification by dispersion algorithm. The initial duration window is expanded over time sample by sample (left-hand graph) until the average of Δx-max and Δy-max exceeds the dispersion threshold.

detected. Such thresholds are typically based on "rules of thumb" coming from research labs using EyeLink eyetrackers, and it is unclear how well such an algorithm would perform on gaze data collected from other labs using another type of eyetracker. Moreover, using acceleration thresholds carries with it the need for additional filtering, since the level of noise is magnified twice through each numerical differentiation.

In sum, the results of event detection depend heavily on how raw gaze coordinates are filtered to produce velocity and acceleration profiles, on the detection algorithm, and on the thresholds the algorithm uses to classify the data. Moreover, the same thresholds are used without regard to individual differences in data quality across participants and trials, which are known to vary. Clearly, future event detection algorithms should take these factors carefully into account.

A New Algorithm for Fixation, Saccade, and Glissade Detection

On the basis of the previously described shortcomings of current algorithms for event detection—in particular, their unsystematic treatment of glissadic eye movements—we propose in this section a new velocity-based algorithm that is developed to robustly detect fixations, saccades, and glissades; blinks are included in the noise category. The algorithm is based on the saccade detection *velocity* algorithm by Smeets and Hooge (2003) and does *not* use positional data directly. Using velocity, instead of position, data to identify saccades was, in fact, recommended already by Bahill et al. (1981), since it provides more accurate and intuitive information about the precise onset and offset of a saccade.

The proposed algorithm contains five main steps: filtering and denoising, peak saccade detection, saccade onset/offset detection, glissade detection, and, finally, fixation detection. These steps are summarized as pseudocode in Table 1. The novelties of the algorithm are, besides explicit glissade detection, an adaptive, data-driven peak saccade detection threshold, a new method to define saccade onsets and offsets, and the use of thresholds motivated by physiological limitations of eye movements.

Filtering and denoising. The output of an eyetracker contains a list of triplets (x_i, y_i, t_i), where (x_i, y_i) is the gaze position acquired at time t_i. These coordinates contain not only information about how the eye moved, but also various types of noise deriving from both the eyetracker and the participant. The role of filtering and denoising is to remove this unwanted variation in the eye movement data.

The first step is to calculate velocity and acceleration profiles from the gaze coordinates. There is a range of methods available for such calculation (see, e.g., Bahill et al., 1981; Inchingolo & Spanio, 1985). Duchowski (2003) described a protocol that first identifies the instantaneous sample-to-sample movement, θ_i, between two consecutive gaze coordinates in degrees of visual angle. If the sampling frequency f_s Hz is known, the angular velocity can be calculated as $\dot{\theta}_i^{raw} = f_s \theta_i$ for the ith sample. However, if a simple sample-to-sample subtraction is used to calculate instantaneous velocity, the output will be very noisy. To reduce the velocity and acceleration noise, two finite impulse response (FIR) filters $\mathbf{h}_k$ and $\mathbf{g}_k$ can be used as follows:

$$\dot{\theta}_i = f_s \sum_{j=0}^{k} \theta_{i+j} \mathbf{h}_j, \, i \in [0, n-k) \tag{1}$$

and

$$\ddot{\theta}_i = f_s \sum_{j=0}^{k} \dot{\theta}_{i+j} \mathbf{g}_j, \, i \in [0, n-k), \tag{2}$$

where k is the filter length and n is the size of the sample vector. In his book, Duchowski describes a five-tap velocity filter and a seven-tap acceleration filter, both with their shapes adapted after the typical appearance of velocity (one large velocity peak) and acceleration (two peaks pointing in opposite directions) during a saccade. The length of the filters should be designed to cover a typical saccade length. Such pattern-matching filters amplify parts of the signals with a similar appearance to the filters while attenuating other portions of the signal.

Saccades with glissades do not have the idealized dynamic behavior assumed by the filters Duchowski (2003) presented. In fact, glissades could be significantly attenuated during filtering. In this article, we use the Savitzky–

Table 1
Pseudocode of the Proposed Algorithm for Fixation, Saccade, and Glissade Detection

For each trial
1. Calculate angular velocities and accelerations. Remove noise.
2. Iteratively find velocity peaks (samples larger than a threshold)
3. Saccade detection: For each velocity peak
   (a) Go back until velocity <= saccade onset threshold
   (b) Go forward until velocity <= saccade offset threshold (adaptive)
   (c) Make sure the saccade duration > minimum saccade duration
4. Glissade detection: Glissades are detected if
   (a) Low-velocity detection: velocity > saccade offset threshold
      within a fixed time window after saccade offset.
   (b) High-velocity detection: velocity > peak saccade threshold
      within a fixed time window after saccade offset.
5. Fixation detection: Fixations are defined by samples that are
   (a) neither saccades, glissades or noise AND
   (b) longer than the minimum fixation duration

Golay (SG) FIR smoothing filter, which makes no strong assumption on the overall shape of the velocity curve and is reported to have a good performance in terms of preserving high-frequency detail in the signal while maintaining both temporal and spatial information about local maxima and minima (Savitzky & Golay, 1964). In principle, this filter finds the polynomial function that best describes the raw data and differentiates the polynomial analytically to avoid adding noise. It then resamples the differentiated polynomial to the original sampling frequency. Applying the `sgolay` function in MATLAB, we can find a set of SG differentiation filters, which, given the raw gaze coordinates (x_i, y_i), output smoothed velocity $(\dot{x}_i, \dot{y}_i)$ and acceleration $(\ddot{x}_i, \ddot{y}_i)$ data for the x- and y-dimensions separately. Then the total angular velocity and acceleration are calculated as the Euclidean distance of the x- and y-components:

$$\dot{\theta}_i = f_s \phi \sqrt{\dot{x}_i^2 + \dot{y}_i^2} \tag{3}$$

and

$$\ddot{\theta}_i = f_s \phi \sqrt{\ddot{x}_i^2 + \ddot{y}_i^2}, \tag{4}$$

where ϕ denotes a factor for converting from pixels to visual degrees. Velocity and acceleration data were appropriately adjusted to compensate for the time shift introduced by the filters. Figure 3 shows the dramatic effect such a second-order polynomial filter with a length of 24 samples (20 msec) has, as compared with simple sample-by-sample differences.

Another type of filtering and denoising targets samples with a position outside of the display or with unphysiological eye movements. For example, when the eye is closed, the eyetracker reports (0, 0) coordinates and/or
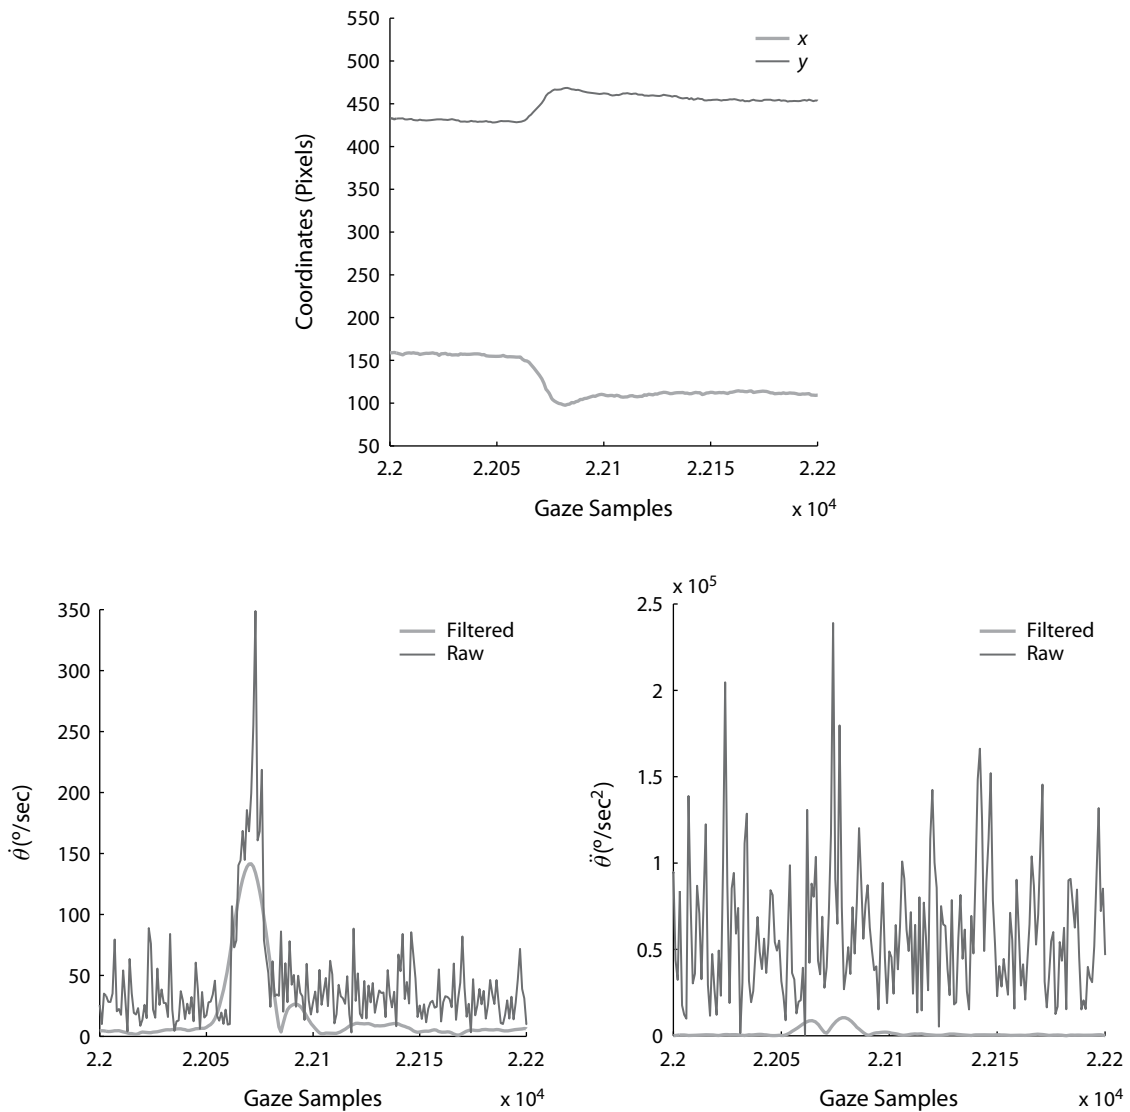

Figure 3. Effect of filtering on velocity and acceleration (1250-Hz data). The *raw* data represent numerical differentiation using sample-by-sample differences. Data were filtered using the Savitzky–Golay filter, with second-order polynomials and 20-msec filter length. All data represent absolute values.
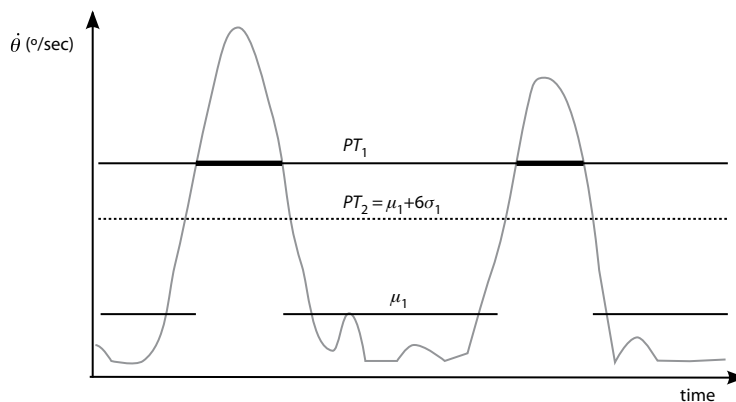
Figure 4. Principle for iterative estimation of saccadic velocity threshold. An initial threshold PT_1 is selected. For all samples with velocities below PT_1, the average velocity μ_1 and standard deviation σ_1 are calculated. The next threshold, PT_2, is then set as $\mu_1 + 6\sigma_1$. The iterative process is continued until $|PT_n - PT_{n-1}|$ is smaller than 1°/sec, and PT_n is selected as the final velocity threshold θ_{PT}. This iteration can be done separately for each participant or even for each trial.

that the pupil diameter equals zero. All such data are removed. Furthermore, when velocity $\dot\theta > 1,000$°/sec or acceleration $\ddot\theta > 100,000$°/sec², data are taken to represent a movement of the eye that is not physiologically possible, and such samples are also omitted from further analysis (see Duchowski, 2003, for an overview on velocity and acceleration thresholds). To make sure that samples before the start and after the end of a noisy period are not left to contaminate the data, the algorithm then searches for onset and offset of noise. Noise on-offset is detected when the velocity reaches the median value of the velocities over the whole trial (which is close to the average noise inside fixations).

Velocity threshold estimation. After filtering and denoising, the algorithm searches for velocity peaks by identifying samples where the velocity is larger than a velocity threshold, $\dot\theta > \dot\theta_{PT}$. In current implementations of saccade detection algorithms (that explicitly detect saccades and implicitly detect fixations), the choice of $\dot\theta_{PT}$

is very important for the outcome. If it is too large, short saccades will be overlooked and if it is too small, samples belonging to a fixation may erroneously be detected as saccade samples. Today, this threshold is set according to an experimenter's individual judgment or, maybe even worse, by the default settings in the analysis software coming with the eyetracker.

We instead propose an iterative, data-driven approach to finding a suitable threshold, illustrated in Figure 4. The iterative algorithm is given an initial peak velocity detection threshold PT_1, which could be in the range 100°–300°/sec, but the choice is not critical as long as there are saccades with peak velocities reaching this threshold. For all samples with velocities lower than PT_1, the average (μ_z) and standard deviation (σ_z) are calculated. The threshold is updated as $PT_n = \mu_{n-1} + 6\sigma_{n-1}$ for each iteration. This sequence of thresholds will converge at a level that is low enough to detect as many saccades as the noise level in data allows and, at the same time, high enough to avoid
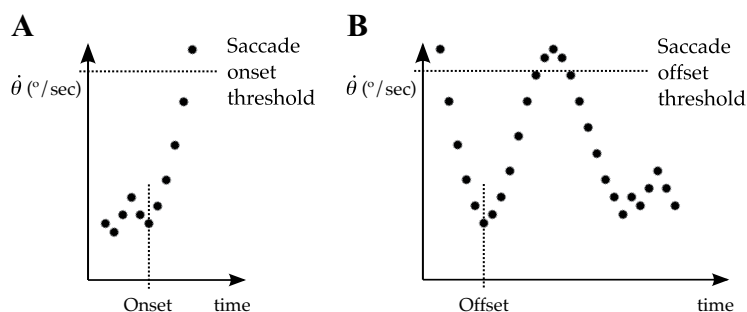

Figure 5. Onset and offset. (A) Saccade onset is found by searching backward from the saccade onset threshold $\dot\theta_{ST}^{onset} = \mu_z + 3\sigma_z$ (not to be confused with the velocity threshold $\dot\theta_{PT}$) until the first local minimum is found. This sample is defined as the onset of the saccade. (B) Saccade offset is detected using a weighted combination of the global noise of the trial and the local noise just before the saccade.

Table 2
Summary of Parameter Settings
Used in the Event Detection Algorithm

Setting	Value
Filter	Savitzky–Golay (`sgolay` in MATLAB)
Filter order	2
Filter length	2× min saccade duration
Peak velocity detection threshold	$\dot\theta_{PT} = \mu_z + 6\sigma_z$
Saccade onset threshold	$\dot\theta_{ST}^{\text{onset}} = \mu_z + 3\sigma_z$
Saccade offset threshold	$\dot\theta_{ST}^{\text{offset}} = \alpha\dot\theta_{ST}{}^{\text{onset}} + \beta\dot\theta_t, \dot\theta_t = \mu_t + 3\sigma_t$
Max saccade velocity	1,000°/sec
Max saccade acceleration	100,000°/sec²
Min saccade duration	10 msec
Min fixation duration	40 msec
α, β	0.7, 0.3
Window length for glissade search	saccade end + min fixation duration

false detection of saccades when intrafixational noise is high. The safety margin of 6σ is considered a good robust level and is also used in microsaccade detection algorithms (Engbert & Kliegl, 2003). Our algorithm continues to iterate until $|PT_n - PT_{n-1}|$ is smaller than 1°/sec between two iterations, and then the final velocity threshold, $\dot\theta_{PT}$, is set as PT_n. For the reading data used to evaluate the algorithm, average values for fixation velocity were 5.44 ± 4.55°/sec ($\mu_n \pm \sigma_n$), giving peak velocity thresholds around 33°/sec (but the individual variation was large across participants). In scene perception data, fixation velocity values were 5.40 ± 3.97°/sec.

Two large advantages exist in this method. First, the choice of a peak detection threshold is data driven and not influenced by the user. Second, the velocity threshold can be calculated separately for each user, or even each trial of data, to adapt to varying noise levels over a number of recordings.

Saccade detection. For each detected saccade peak (those detected after the last iteration), the algorithm searches backward (from the leftmost peak saccade sample) and forward (from the rightmost peak saccade sample) in time for the saccade onset and offset.[2] Saccade onset, which typically is simple to detect, due to its well-behaved velocity profile, is defined as the first sample that goes below the *saccade onset threshold* $\dot\theta_{ST}^{\text{onset}} = \mu_z + 3\sigma_z$ and where $(\dot\theta_i - \dot\theta_{i+1}) \geq 0$. Figure 5A illustrates the principle.

Saccade offset is less obvious, because this is where we find the glissades. In order to make the algorithm better at detecting glissades, we defined $\dot\theta_{ST}^{\text{offset}}$ as a weighted combination of the velocity noise during the whole trial, $\dot\theta_{ST}^{\text{onset}}$, and a locally adaptive noise factor, $\dot\theta_t = \mu_t + 3\sigma_t$. To avoid contamination from glissadic movements, the local noise factor is calculated over the velocity samples within a window with size $\tau_{\min}$ msec (40 msec, the minimum allowed fixation duration) and *preceding* the saccade currently being processed. The local noise factor $\dot\theta_t$ is used to adapt to local variation in noise, for each saccade, which is important since noise levels may vary, for instance, between data at the center of the stimulus and peripheral saccades or due to participant-dependent factors, such as downward lashes, lenses, or glasses (Duchowski, 2003). The idea of adaptive, noise-contingent thresholding is not new but has been described, for example, by Duchowski (pp. 117–119; on the basis of acceleration thresholds). To

ills. Many marketers have believed that they can build a brand in such a way as to gain and sustain competitive advantage, only to see it falter and fail. Equally, organizations have failed to harness the possibilities that marketing may be able to deliver, to in turn falter and fail.

Therefore marketing must be viewed objectively, and a course in marketing does not make a marketer. It is experience, and the knowledge gained from that experience, that makes for a potentially more objective and thus 'successful' marketer.

Main ways of approaching a marketing strategy

Business, just like life, for that matter, is rarely quite so rational. Decisions can be taken from a variety of perspectives and logic and rationality may provide only one perspective. One of the main

Figure 6. Example of a text stimulus.

the authors' knowledge, no systematic testing of this type of algorithm has been undertaken so far. Saccadic offset is then defined as the first sample that goes below the *saccade offset threshold* $\dot{\theta}_{ST}^{offset} = \alpha \dot{\theta}_{ST}^{onset} + \beta \dot{\theta}_t$ and where $(\dot{\theta}_i - \dot{\theta}_{i+1}) \leq 0$. Furthermore, we propose a minimum saccade duration threshold of 10 msec (12 samples), which is large enough to avoid noise being falsely categorized as saccades but small enough to include short saccades (~1°). Finally, we exclude saccades that are preceded by a period where $\mu_t > \dot{\theta}_{PT}$, since this indicates that there was no period of stillness prior to the saccade onset (most often, indicating recording imperfections).

Glissade detection. Since there is no widespread or accepted definition of the oculomotor movement we call the glissade, we provide two alternative definitions, dubbed *low-velocity* and *high-velocity* glissade detection. Low-velocity detection targets slower glissadic movements with small amplitudes, whereas high-velocity finds only long high-velocity glissades. In this article, we define the two groups of glissades as mutually exclusive; that is, low-velocity glissades are not a subset of high-velocity glissades.

The high-velocity criterion requires that the velocity curve within a τ_{min} (40) msec window after the saccadic offset raises above the peak saccade threshold, $\dot{\theta}_{PT}$, and down below it, at least once. In other words, a high-velocity glissade has a velocity peak that would qualify it for saccadic status. The low-velocity criterion is identical, except that it requires only that the velocity curve raises above the saccade offset threshold $\dot{\theta}_{ST}^{offset}$, as in Figure 5B.

The *onset of the glissade* is defined as the offset of the preceding saccade. As with saccade detection, the *glissade*

Table 3
Summary (Means ± Standard Deviations When Applicable)
of Average Results for Reading and Scene Perception

Measure	Reading	Scene Perception
	Proposed Algorithm	
Fixation duration (msec)	193.7 ± 100.0 (N = 24,936)	263.6 ± 185.4 (N = 7,963)
Saccade duration (msec)	42.5 ± 18.0 (N = 26,203)	47.2 ± 16.8 (N = 8,597)
Glissade duration (msec)	22.2 ± 9.8 (N = 12,525)	25.0 ± 9.8 (N = 5,081)
Fixation velocity (°/sec)	5.44 ± 4.55	5.40 ± 3.97
Saccade peak velocity (°/sec)	167.5 ± 112.7	230.6 ± 139.2298
Saccade peak acceleration (°/sec^2)	9,916.5 ± 6,114.2	12,779.1 ± 6,390.9
Max peak velocity (°/sec)	933.9	845.5
Max peak acceleration (°/sec^2)	99,343	65,752
% glissadic saccades	47.8	59.1
	I-DT (40 msec, 0.5°)	
Fixation duration (msec)	128.1 ± 80.3 (N = 47,467)	150.6 ± 112.4 (N = 24,845)
Saccade duration (msec)	11.4 ± 19.4 (N = 26,866)	14.8 ± 20.8 (N = 12,996)
Glissade duration (msec)	n.a.	n.a.
Fixation velocity (°/sec)	n.a.	n.a.
Saccade peak velocity (°/sec)	119.5 ± 146.1	133.6 ± 166.5
Saccade peak acceleration (°/sec^2)	85,985.6 ± 135,472.3	112,678.9 ± 169,443.1
Max peak velocity (°/sec)	999.86	998.86
Max peak acceleration (°/sec^2)	993,401	995,058
% glissadic saccades	n.a.	n.a.
	I-DT (100 msec, 1°)	
Fixation duration (msec)	213.5 ± 97.6 (N = 28,703)	244.9 ± 133.2 (N = 15,079)
Saccade duration (msec)	15.8 ± 27.3 (N = 18,552)	20.8 ± 27.4 (N = 10,823)
Glissade duration (msec)	n.a.	n.a.
Fixation velocity (°/sec)	n.a.	n.a.
Saccade peak velocity (°/sec)	178.7 ± 156.2	211.5 ± 173.7
Saccade peak acceleration (°/sec^2)	116,237.2 ± 155,978.8	169,331.7 ± 186,858.7
Max peak velocity (°/sec)	1,000	1,000
Max peak acceleration (°/sec^2)	1,000,000	1,000,000
% glissadic saccades	n.a.	n.a.
	Velocity Algorithm	
Fixation duration (msec)	171.6 ± 108.2 (N = 32,438)	210.4 ± 159.9 (N = 16,116)
Saccade duration (msec)	32.4 ± 17.2 (N = 31,061)	36.7 ± 18.4 (N = 14,786)
Glissade duration (msec)	n.a.	n.a.
Fixation velocity (°/sec)	n.a.	n.a.
Saccade peak velocity (°/sec)	162.1 ± 117.4	205.0 ± 145.3
Saccade peak acceleration (°/sec^2)	160,557.2 ± 130,699.4	205,099.0 ± 174,269.6
Max peak velocity (°/sec)	994.5	993.9
Max peak acceleration (°/sec^2)	993,697	977,468
% glissadic saccades	n.a.	n.a.

Note—The identification by dispersion (I-DT) and velocity algorithms were run in the SMI BeGaze 2.2 implementation. No postprocessing, such as removal of absurdly short fixations, was applied to the output from these two algorithms.
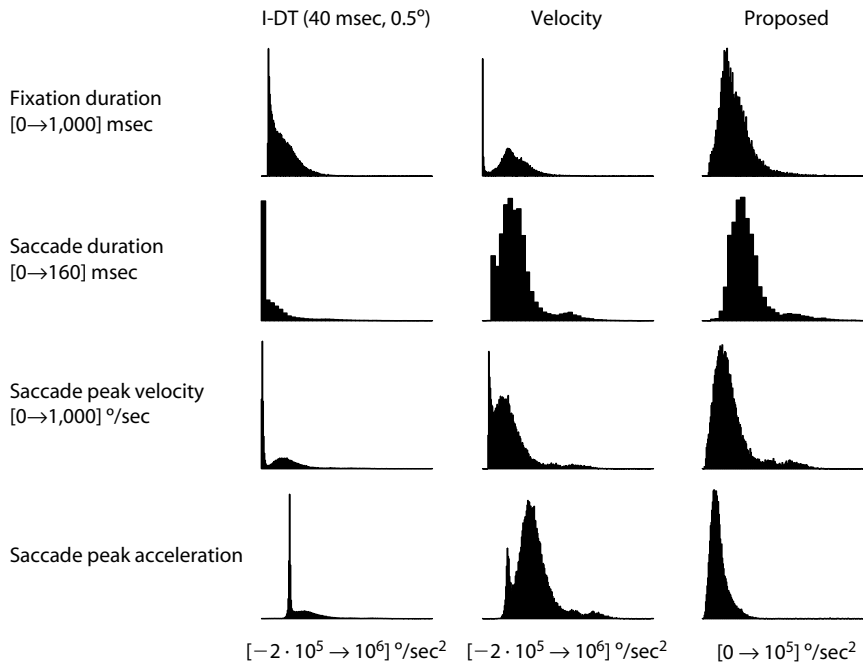
Figure 7. (Relative) histogram comparisons across algorithms for reading data. Bin size is 4 msec for duration data. Values for *x*-axes are given to the left in the figure, except for acceleration, where values are given at the bottom. I-DT, identification by dispersion threshold.

offset is defined when $(\dot{\theta}_i - \dot{\theta}_{i+1}) \leq 0$ after the last velocity peak sample in the glissade. Glissades with an amplitude larger than their preceeding saccades were omitted.

Fixation detection. Since our evaluation stimuli are static, fixations are everything that is not noise, saccades, or glissades. However, in order to exclude extremely short fixations, we use an additional minimum duration criterion, τ_{min}. In the literature, there is no consensus on how to choose the minimum duration, but there are reports of durations as short as 50 msec (Rayner, 1998,
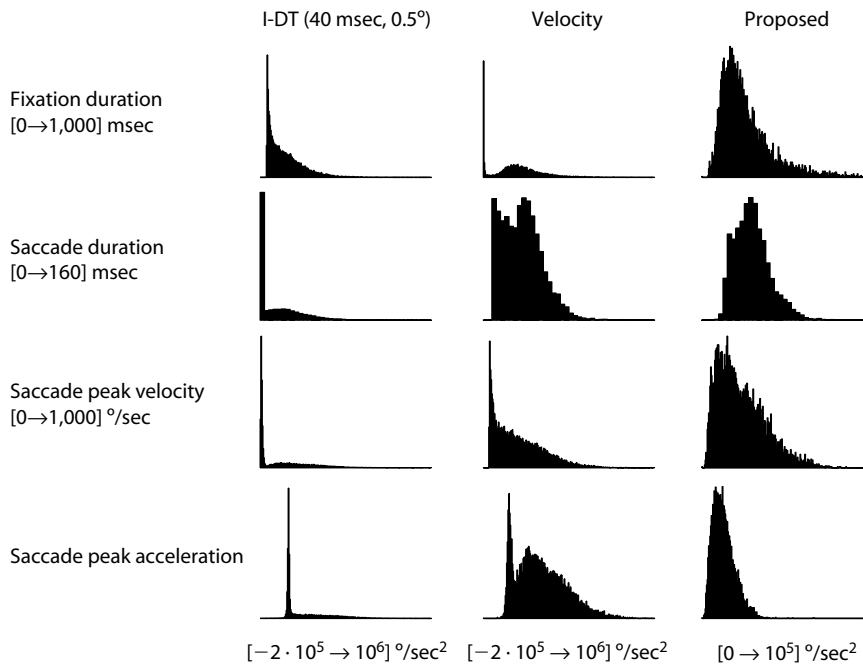

Figure 8. Histogram comparisons across algorithms for scene perception data. I-DT, identification by dispersion threshold.
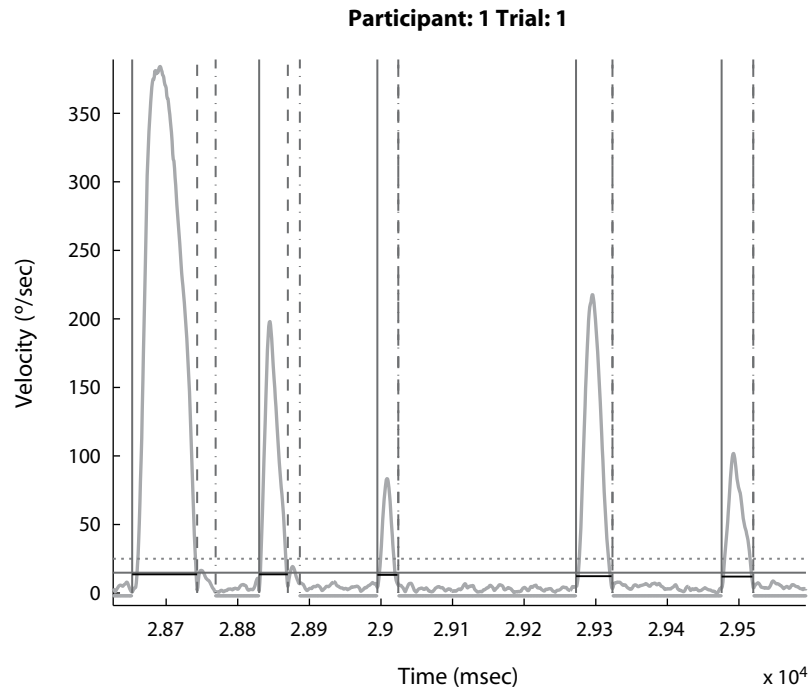
Participant: 1 Trial: 1

Figure 9. Results for fixation, saccade, and glissade detection of 1 person during reading. Solid, dashed, and dashed-dotted lines mark the saccade onset, saccade offset, and the glissade offset, respectively. Fixations are indicated by the thick lines at the bottom. The black horizontal lines within the saccades represent the saccade onset threshold after it has been adjusted by an estimation of the local noise level. As can been seen, the fixation velocity shown in the figure is lower than the global fixation velocity, since the adaptive threshold is lower than the global, trial-based one.

p. 376). In most studies, however, the minimum fixation duration thresholds are around 100–200 msec (Salvucci & Goldberg, 2000). There are probably three reasons why the duration span is large. First, a fixation could be defined as an *oculomotor event*—that is, the time that the eye is still—but also as the period of *visual intake* during this period. Due to saccadic suppression, these two definitions are not equal, but the visual system begins to shut off before the oculomotor fixation has ended and does not fully recover until a bit into the next fixation (Matin, 1974). Thus, some oculomotor fixations may have very little visual intake, and some researchers therefore select a duration threshold based on an estimation of the shortest period for visual intake (Rötting, 2001). The second reason is due to current fixation detection algorithms and their settings; several studies have shown that typical dependent variables, such as fixation duration and saccade length, critically depend on the settings (see, e.g., Shic, Chawarska, & Scassellati, 2008). Finally, as will be shown in this article, short fixations can vary significantly in duration depending on whether you assign glissades to saccades or fixations. We have manually identified several oculomotor fixations in the data, especially during reading, with durations below 50 msec, and we therefore use $\tau_{\min} = 40$ msec as the minimum allowed duration.

An overview of all the algorithmic settings is shown in Table 2.

DATA USED FOR EVALUATION AND COMPARISONS

Reading

The reading data originate from a study in which over 300 students from the Lund University School of Economics and Management read an English text on a computer screen for 10–20 min, followed by a comprehension test. The texts described a case scenario, and the participants were instructed to read the text well enough to be able to answer some questions about it afterward. The text was divided into 16 images with a resolution of $1{,}024 \times 768$ pixels, containing black letters on a midgray background. Each letter spanned approximately 15×15 pixels (0.5° of visual angle). An example of one text image is shown in Figure 6.

The experimental setup consisted of one control computer running SMI's recording software iViewX with default settings and communicating with another computer running MATLAB to display the texts in full screen. Text images were shown on a Samsung 19-in. flat screen with a size of 377×300 mm and viewed from a distance of 670 mm.

After a 13-point calibration, the participants were left to read the text at their own pace. They could only move forward through the text (by clicking the mouse). Gaze positions were recorded monocularly (although viewing was binocular) with the SMI HiSpeed eyetracker at 1250 Hz throughout the experiment.
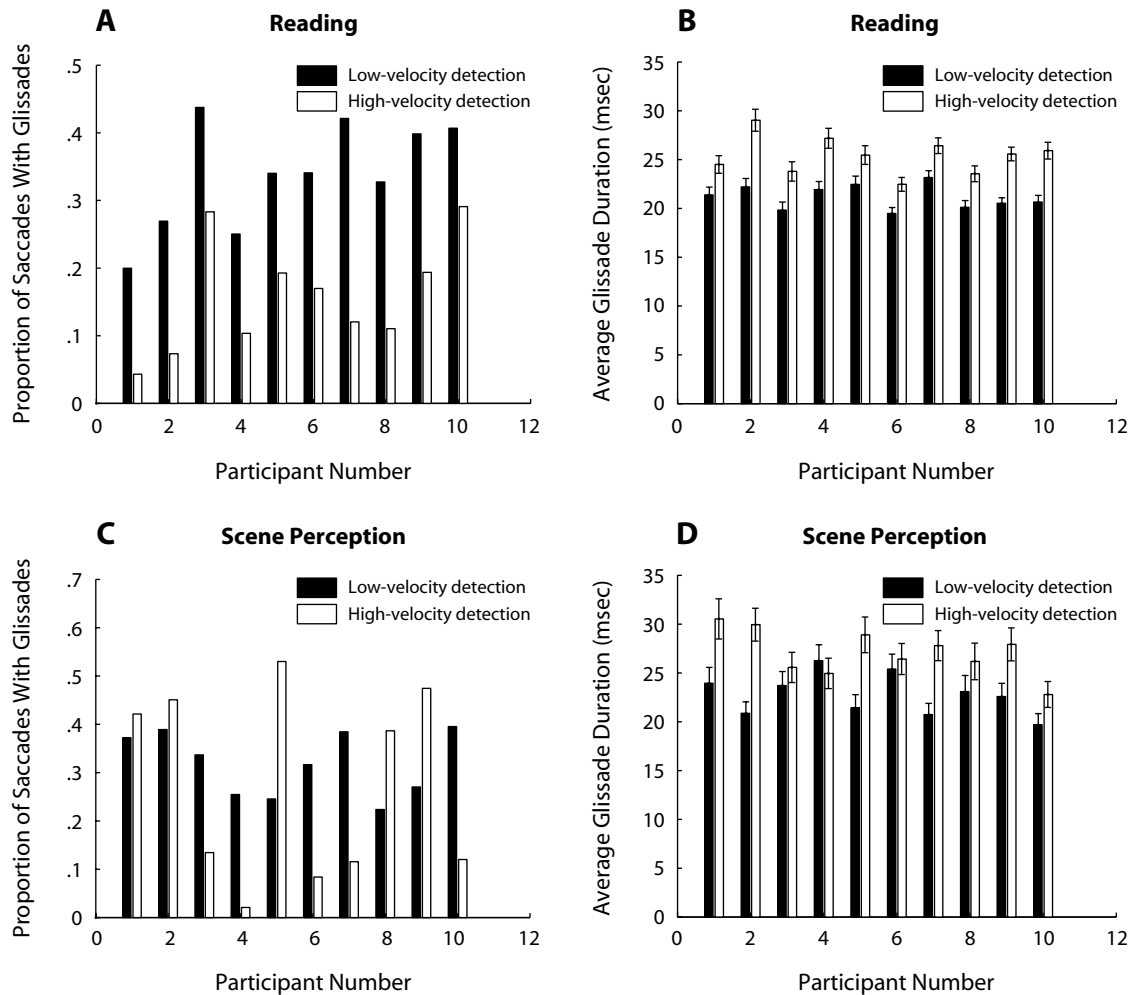
Figure 10. Proportions and durations of glissades during reading and scene perception (as detected by the proposed algorithm). (A) Reading. (B) Reading. (C) Scene perception. (D) Scene perception.

For each recording, the experimenter rated the quality of the collected data on a scale of 1–5. Among the data that got the highest quality rating (5), data from 10 participants were chosen at random.

Scene Perception

Eye movement data collected during scene perception originate from the study by Nyström and Holmqvist (2008), where details about the experiment can be found. In short, gaze positions were recorded at 1250 Hz with the SMI HiSpeed eyetracker while 10 subjects watched digitized photographs of "real-world" images on a computer screen. The task was either to "inspect the images carefully" or to identify the semantically most informative part of the image. The setup was otherwise very similar to that of the reading experiment.

RESULTS AND DISCUSSION

Algorithmic Comparisons

Table 3 summarizes detection results for the three tested algorithms: The I-DT algorithm (as defined in Salvucci

& Goldberg, 2000), the algorithm by Smeets and Hooge (2003; i.e., the velocity algorithm), and the algorithm proposed in this article. For the first two algorithms, we report data calculated from the commercial implementations found in the BeGaze 2.2 software developed by SMI. To allow for a fair comparison against the proposed algorithm, I-DT thresholds were set to 0.5° for dispersion and 40 msec as minimum fixation duration, whereas the velocity algorithm used a peak velocity threshold of 35°/sec and a minimum saccade duration of 10 msec—that is, as closely matched as possible to the thresholds used in the proposed algorithm. Since values of 1° for dispersion (optimal threshold according to Blignaut, 2009) and 100 msec for minimum duration (default setting in ASL Eyenal) seem to be a common setting in many studies, results for this configuration of the I-DT are also included.

There are some striking differences in Table 3. First, fixation duration, saccade duration, and peak saccade velocity differ significantly across different algorithms, although they were tested on exactly the same data. Second, peak saccade acceleration values differ by a factor of 10 between the commercial implementations and the
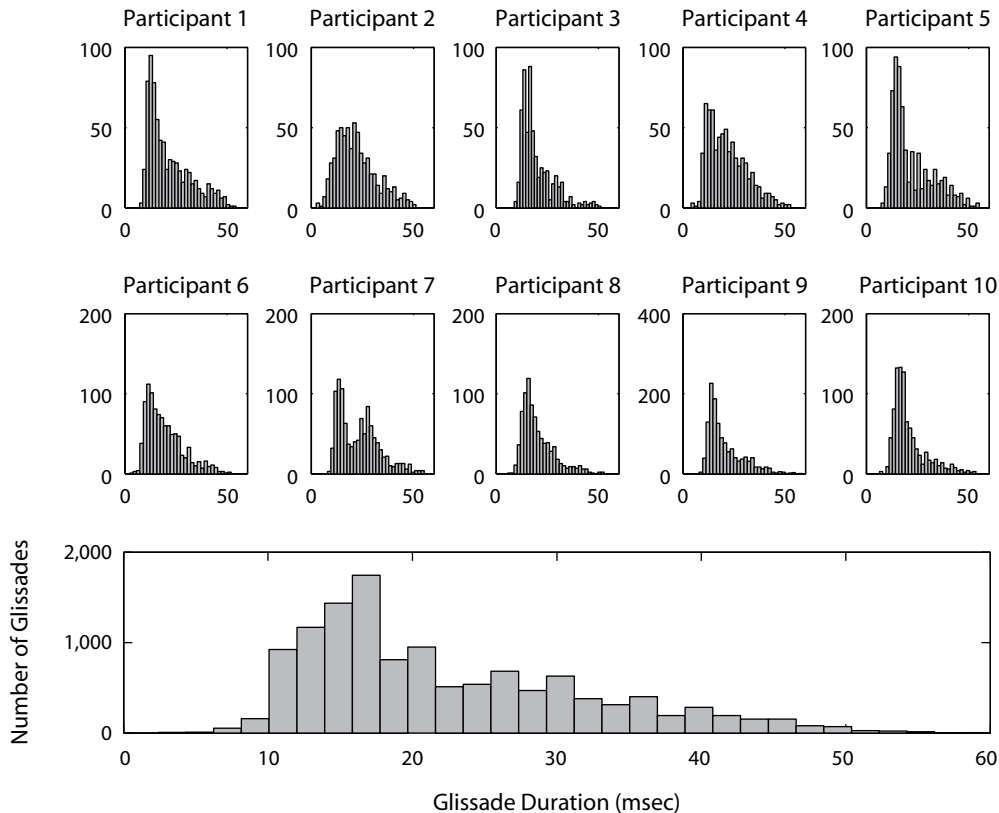
Figure 11. Distribution of glissades with small amplitudes (low-velocity detection criteria) collected from 10 people during reading. Bin size is 2 msec.

proposed one. An acceleration above 100,000°/sec^2 is unlikely to originate from real eye movements, according to Bahill et al. (1981). It is interesting to note, however, that saccade peak acceleration has previously been shown to vary significantly across studies (Duchowski, 2003, p. 126), possibly due to variation in the recording apparatus and filters used when acquiring the eye movements and calculating the acceleration profiles (see also Inchingolo & Spanio, 1985).

Clearly, such large differences are troublesome, since the choice of algorithm may verify or falsify your hypothesis. Moreover, comparisons of these important dependent variables across studies, labs, and algorithms seem futile. Looking just at Table 3, it may be difficult to understand why the results are so different. The histograms of the variables offer further insights. Figures 7 and 8 depict histograms for a number of eye movement parameters collected during both reading and scene perception.

Beginning with Figure 7, we can see that something is deeply wrong with the histograms produced both by the velocity algorithm and, in particular, by the I-DT algorithm. Maybe the most disturbing feature is that there are dominant peaks for duration values close to zero (bins are 4 msec for duration data), which means that a large portion of the fixations and saccades lie on the 0–10 msec interval. The same trend is discernible also in scene perception data (Figure 8).

So, where do the I-DT and velocity algorithms fail? As was explained in the Current Algorithms and Their Shortcomings section, the I-DT algorithm initiates a window of a length equal to the minimum fixation duration (40 msec) over the gaze data. Then the window is expanded until the intrafixational dispersion exceeds a predefined threshold (e.g., 1°). Due to noise inherent in the eyetracker, in addition to intrafixational eye movements, the dispersion is typically reached very soon after the onset and, therefore, terminates and outputs a fixation near the minimum allowed duration. This way, a long fixation of the eye may be divided into many shorter fixations in the output of the algorithm. This undesired behavior of the I-DT algorithm is even more prominent when data are analyzed at high sampling rates (>200 Hz) (especially with small saccades), rather than at lower rates, since gaze points are compared against the dispersion threshold more frequently (and thus have more chances to exceed it). Some versions of the I-DT have introduced noise resilience by requiring that more than one sample needs to exceed the threshold before the fixation is terminated (Applied Science Laboratories, 2001).

The velocity algorithm produces somewhat more reliable histograms, but also there are portions of very short fixations (1–10 msec) identified by the algorithm. Many of these fixations are detected, since the algorithm is not endowed with a systematic method to handle glissades.

A   **Participant 3 (Reading)**

B   **Participant 5 (Reading)**

C   **Participant 3 (Scene Perception)**
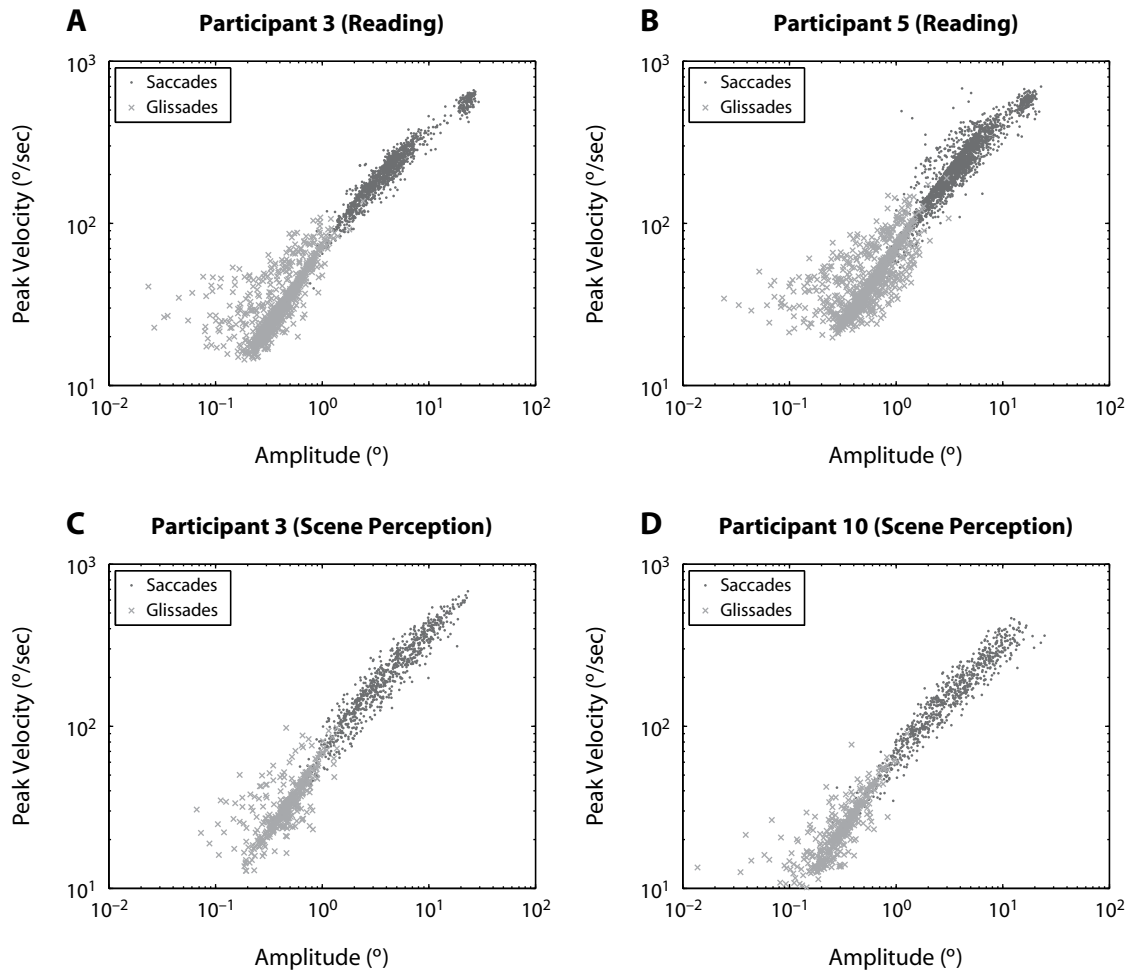
D   **Participant 10 (Scene Perception)**

Figure 12. Main sequence for saccades and glissades (data from 4 people during reading or scene perception). Amplitude is calculated as the Euclidean distance between the start and end locations of a saccade. Dots and crosses represent saccades and glissades, respectively. (A) Participant 3 (reading). (B) Participant 5 (reading). (C) Participant 3 (scene perception). (D) Participant 10 (scene perception).

Instead, the algorithm terminates the saccade before the glissade, identifies the glissade as another, short saccade, and inserts an absurdly short fixation in between. Moreover, unproportionally many short saccades are detected by the velocity algorithm, most likely due to recording imperfections. This can happen, for example, when the algorithm processing the eye video finds a false corneal reflex that is confused with the real one and then rapidly switches between the two.

The proposed implementation overcomes all of the problems above, as can be seen from the histograms.

In all fairness, many researchers choose to exclude a very short fixation after it has been run through the event detector (such as the velocity algorithm). Some manufacturers even recommend their users to do so: "Postprocessing or data cleanup may be needed to prepare data during analysis. For example, short fixations may need to be discarded or merged with adjacent fixations, or artifacts around blinks may have to be eliminated" (SR Research, 2007, p. 85). Obviously, this would make the fixation du-

ration histogram generated from the velocity algorithm more similar to that of the proposed algorithm.

Glissade Detection

Figure 9 shows the results of event detection on eyetracking data from 1 person making five saccades while reading. The first saccade originates from a large movement when the reader changes to a new line, whereas the other four saccades reflect forward reading on the new line. Note how two out of the five saccades (1 and 2) end with glissades.

In our evaluation data, the average duration of a glissade was about 24 msec and occurred in over half of all the registered saccades. There were, however, large individual differences, as Figure 10 illustrates. Some participants had relatively few glissades (see, e.g., Participant 1, who had about 25% glissadic saccades during reading), whereas other participants made a glissade for the majority of the saccades. Glissade duration also differed across participants, but to a smaller degree, and seemed to be fairly stable around the overall average value.

A   Small-Amplitude Glissades (Reading)

Low-Velocity Glissade Detection

B   Large-Amplitude Glissades (Reading)

High-Velocity Glissade Detection

**C   Small-Amplitude Glissades
(Scene Perception)**

Low-Velocity Glissade Detection

**D   Large-Amplitude Glissades
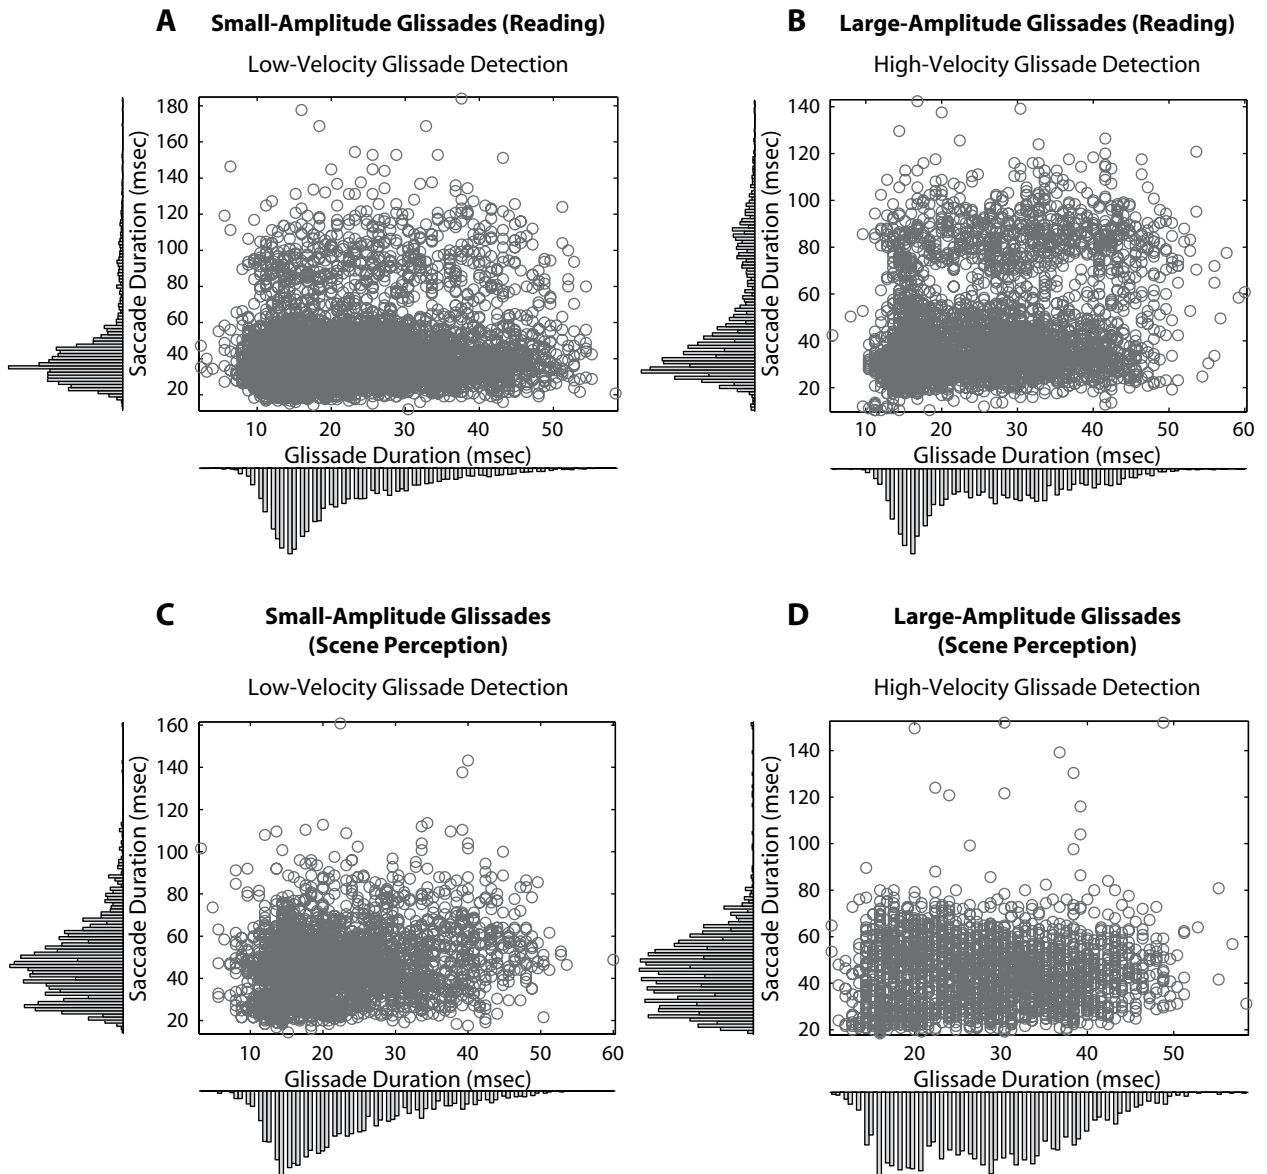(Scene Perception)**

High-Velocity Glissade Detection

Figure 13. Glissade versus saccade durations. Large glissades are defined as having a peak velocity higher than the saccade detection threshold. No correlation was found between saccade and glissade durations. (A) Small-amplitude glissades (reading). (B) Large-amplitude glissades (reading). (C) Small-amplitude glissades (scene perception). (D) Large-amplitude glissades (scene perception).

Figure 11 shows histograms of glissade durations. Note how the left-skewed distribution of glissades resembles that of saccades (and fixations). Interestingly, Kapoula et al. (1986) argued that glissades have all the features of saccades; they follow the *main sequence* (a systematic relationship between peak velocity and amplitude) and are triggered by bursts in neural activity in the same regions as saccades. The data analyzed in this article indicate that glissades lay on the main sequence; Figure 12 shows how glissades extend the peak-velocity–amplitude relationship of saccades.

If, in addition, the visual intake was closed, as it is under saccades, glissades could be argued to be a variety of saccades. Deubel and Bridgeman (1995) showed

brief flashes during glissades and could show that they were perceived by participants. Flashes were intentionally strong enough, however, to produce afterimages on the retina, so possibly the reported perception came from the afterimage, rather than from the flash itself. Using the dual Purkinje eyetracker, McConkie and Loschky (2002) suggested that perceptual onset begins 6 msec after the velocity peak of a glissade, since this is when participants start detecting a display change (switching from a low- to a high-resolution version of an image). The changes were not robustly detected until 32 msec after the peak, however, suggesting that a complete visual intake is not reached until then. Moreover, glissades typically do not occur in both eyes simultaneously (according to Kapoula
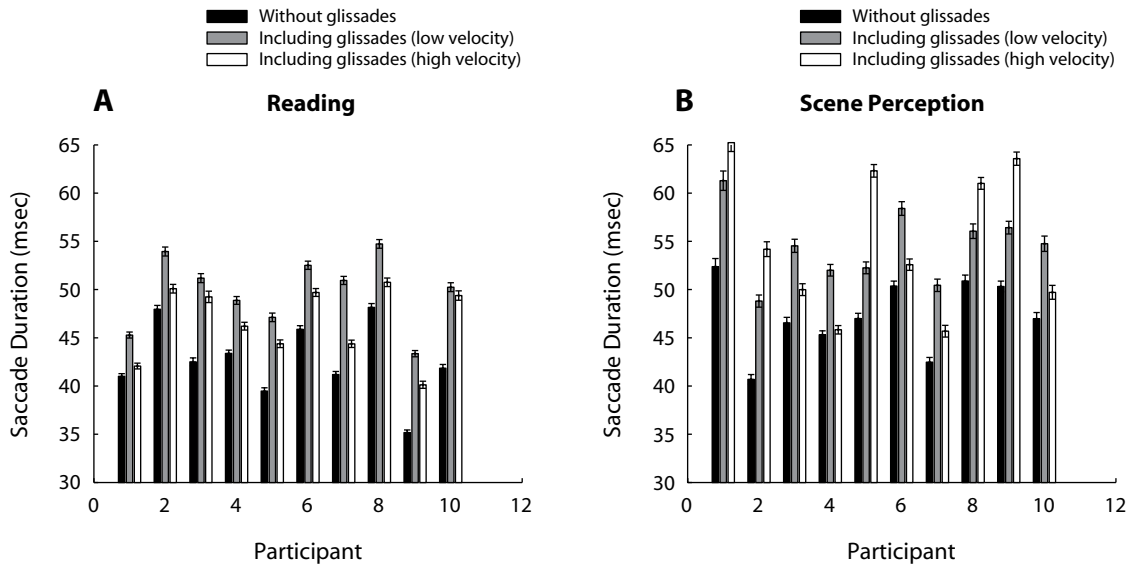
Figure 14. Influence of saccade duration depending on whether the glissades are included in the saccade or not. (A) Reading. (B) Scene perception.

et al., 1986), but one eye is still and could therefore take in visual information while the other eye is glissading.

Kapoula et al. (1986) reported that glissade amplitude increases with saccade size (but only moderately), using data where participants made saccades toward jumping dot-like targets. Figure 13 shows the same relationship for our reading and scene perception data. Figures 13A and 13C illustrate this relationship for glissades with small amplitude (as defined in Figure 9), whereas Figures 13B and 13D consider only glissades with a peak velocity larger than the saccade detection threshold. There was no significant correlation between glissade and saccade durations in either of the cases.

Although it is still unclear whether glissades are a variety of saccades or should be part of the following fixations, glissades occupy a significant portion of time in between saccades and fixations. From the viewpoint of the event detection algorithms and their output, it matters whether we assign glissades to the fixation or to the saccade. Figures 14A and 14B show how saccade duration is affected when glissades are added to the saccade and when they are not. With glissades, saccade duration increases by 5–15 msec. On average, glissades during reading and scene perception were found to add 24.5% and 31.3%, respectively, to the saccade duration and, if instead added to fixations, 5.4% and 5.6%, respectively, to the fixation duration.

Kapoula et al. (1986) observed, by testing 3 participants with jumping dot stimuli, that glissades are more common for small saccades than for large saccades. In our larger data from reading and scene perception (see Figures 15A and 15B), we found no support for a simi-
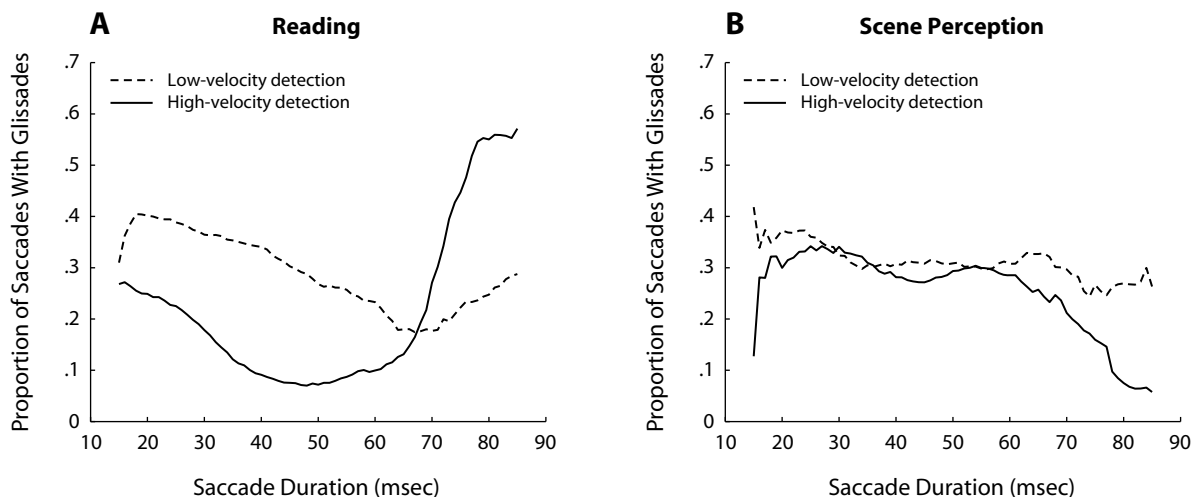

Figure 15. Proportions of glissades for different saccade durations. (A) Reading. (B) Scene perception.

lar trend in reading and scene perception. Interestingly, however, in the reading data, the long (return sweep) saccades have a much higher percentage of glissades than do equally long saccades during scene perception. Could this be evidence that the oculomotor system adapts the saccadic retardation to the velocities of the short saccades of a line and then fails to increase retardation to the much higher velocity of the return sweep? In scene perception, where saccadic velocities are not bimodally distributed and are much more random, the effect is rather the reverse.

CONCLUSIONS

We have developed a new event detection algorithm to overcome some limitations of current velocity-based algorithms. It has two major components. First, the new algorithm is endowed with a noise-dependent velocity threshold that is adaptive with respect to different viewers, different recording situations, and/or varying levels of noise during a whole experiment or a single trial. The automatic adaption of velocity threshold makes our algorithm settings-free for users. In most current algorithms, the variable level of noise is disregarded and not automatically adjusted for. Instead, the end user must make the important decision of how to deal with noise via the velocity threshold, using his/her subjective judgment or just trusting the default settings of the analysis software.

Second, the new algorithm was designed with the explicit goal of detecting glissades. Detecting glissades is appropriate because they exist and have a statistical distribution, like many other eye movement events. If we use algorithms that are not explicit about glissades, their durations will randomly spill over to the fixation or the saccade or will trick the algorithm into outputting absurdly short fixations, making results less reliable and less comparable. Glissades were found to occur frequently (in over half of all the saccades) and had an average duration of about 24 msec. This would add more than 25% to the average saccade duration and about 5% to the average fixation duration. Our algorithm solves all these problems and gives the user control over glissades in the data analysis. In addition, acknowledging and studying glissades as a separate class of eye movements has the potential of opening up new insights on oculomotor control and perception, and data on glissades may be linked to fatigue and neurological disorders (see, e.g., Cuiffreda & Tannen, 1995, p. 39).

One limitation of this article is that the proposed algorithm has not been tested against the SR EyeLink detection algorithm, which—according to the algorithm description in their manual and to Stampe (1993)—may be less inconsequential to glissades than are most algorithms, with the intent to remove glissades by filtering and assign their durations to the fixation. Another weakness is that the algorithm cannot deal with glissade-like movement preceding a saccade (this situation is, however, rarely found in the data). Finally, the algorithm does not detect smooth pursuit as a separate class of eye movements.

REFERENCES

Applied Sciences Laboratory (2001). *Eyenal manual: Version 1.0* [Computer software manual]. Bedford, MA: Author.

Bahill, A. T., Brockenbrough, A., & Troost, B. T. (1981). Variability and development of a normative data base for saccadic eye movements. *Investigative Ophthalmology & Visual Science*, **21**, 116-125.

Blignaut, P. (2009). Fixation identification: The optimum threshold for a dispersion algorithm. *Attention, Perception, & Psychophysics*, **71**, 881-895.

Cuiffreda, K. J., & Tannen, B. (1995). *Eye movements for the clinician*. St. Louis: Mosby.

Deubel, H., & Bridgeman, B. (1995). Perceptual consequences of ocular lens overshoot during saccadic eye movements. *Vision Research*, **35**, 2897-2902.

Duchowski, A. T. (2003). *Eye tracking methodology: Theory and practice*. New York: Springer.

Engbert, R., & Kliegl, R. (2003). Microsaccades uncover the orientation of covert attention. *Vision Research*, **43**, 1035-1045.

Frens, M. A., & Van der Geest, J. N. (2002). Scleral search coils influence saccade dynamics. *Journal of Neurophysiology*, **88**, 692-698.

Inchingolo, P., & Spanio, M. (1985). On the identification and analysis of saccadic eye movements—A quantitative study of the processing procedures. *IEEE Transactions on Biomedical Engineering*, **32**, 683-695.

Irwin, D. E. (2004). Fixation location and fixation duration as indices of cognitive processing. In J. M. Henderson & F. Ferreira (Eds.), *The interface of language, vision, and action: Eye movements and the visual world* (pp. 105-133). New York: Psychology Press.

Kapoula, Z. A., Robinson, D. A., & Hain, T. C. (1986). Motion of the eye immediately after a saccade. *Experimental Brain Research*, **61**, 386-394.

Karsh, R., & Breitenbach, F. W. (1983). Looking at looking: The amorphous fixation measure. In R. Groner, C. Menz, D. F. Fisher, & R. Monty (Eds.), *Eye movements and psychological functions: International views* (pp. 53-64). Hillsdale, NJ: Erlbaum.

Matin, E. (1974). Saccadic suppression: A review and an analysis. *Psychological Bulletin*, **81**, 899-917.

McConkie, G. W., & Loschky, L. C. (2002). Perception onset time during fixations in free viewing. *Behavior Research Methods, Instruments, & Computers*, **34**, 481-490.

Nyström, M., & Holmqvist, K. (2008). Semantic override of low-level features in image viewing—both initially and overall. *Journal of Eye Movement Research*, **2**, 2:1-2:11.

Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, **124**, 372-422.

Rottach, K. G., Das, V. E., Wohlgemuth, W., Zivotofsky, A. Z., & Leigh, R. J. (1998). Properties of horizontal saccades accompanied by blinks. *Journal of Neurophysiology*, **79**, 2895-2902.

Rötting, M. (2001). *Parametersystematik der augen- und blickbewegungen für arbeitswissenschaftliche untersuchungen*. Aachen, Germany: Shaker.

Salvucci, D. D., & Goldberg, J. H. (2000). Identifying fixations and saccades in eye-tracking protocols. In *Proceedings of the 2000 Symposium on Eye Tracking Research and Applications* (pp. 71-78). New York: ACM Press.

Savitzky, A., & Golay, M. J. E. (1964). Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, **36**, 1627-1639.

Shic, F., Chawarska, K., & Scassellati, B. (2008). The amorphous fixation measure revisited: With applications to autism. In B.C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual*

Meeting of the Cognitive Science Society (p. 1964). Austin, TX: Cognitive Science Society.

SHIC, F., SCASSELLATI, B., & CHAWARSKA, K. (2008). The incomplete fixation measure. In *Proceedings of the 2008 Symposium on Eye Tracking Research and Applications* (pp. 111-114). New York: ACM Press.

SMEETS, J. B. J., & HOOGE, I. T. C. (2003). Nature of variability in saccades. *Journal of Neurophysiology*, **90**, 12-20.

SR RESEARCH (2007). *EyeLink user manual. Version 1.3.0* [Computer software manual]. Ottawa, ON: Author.

STAMPE, D. M. (1993). Heuristic filtering and reliable calibration methods for video-based pupil-tracking systems. *Behavior Research Methods, Instruments, & Computers*, **25**, 137-142.

WEBER, R. B., & DAROFF, R. B. (1972). Corrective movements following refixation saccades: Type and control system analysis. *Vision Research*, **12**, 467-475.

NOTES

1. The latter will henceforth be referred to as the *velocity algorithm*.
2. To cope with saccades with dual velocity peaks, the search is not initiated from the central saccade sample.