The Official Journal of the International Society on Infant Studies *Infancy*

Infancy, **17**(2), 141–158, 2012

ISSN: 1525-0008 print / 1532-7078 online
DOI: 10.1111/j.1532-7078.2011.00106.x

WILEY-BLACKWELL

Looking Forward: Comment on Morgante, Zolfaghari, and Johnson

Sarah C. Creel
Department of Cognitive Science
University of California

Morgante et al. (in press) find inconsistencies in the time reporting of a Tobii T60XL eye tracker. Their study raises important questions about the use of the Tobii T-series in particular, and various software and hardware in general, in different infant eye tracking paradigms. It leaves open the question of the source of the inconsistencies. Here, observations from a Tobii eye tracker are presented to elucidate possible sources of timing inconsistencies, including those found by Morgante et al. The ramifications of the reported timing inconsistencies are related to various infant paradigms. The focus is on the level of concern a researcher should have if any eye tracker displays these timing characteristics, and what corrective measures may be taken. While posing no problems for some paradigms, timing inconsistencies are potentially problematic (but correctable) when assessing event-related looking behavior. Observed timing contraindicates use in fast gaze-contingent displays ($<$ 100 ms). General suggestions are made regarding timing in eye-tracked data collection.

There are many potential benefits to using an eye tracker in infant studies (reviewed by Aslin, in press), and recent technological advances and miniaturization have resulted in the first generation of eye trackers suited for use with infants. The Tobii family of eye trackers in particular has been used extensively in infant research in the last decade. It is recommended by its ease of calibration and setup, and its design that allows a greater range of

---

Correspondence should be sent to Sarah C. Creel, Department of Cognitive Science, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0515. E-mail: creel@cogsci.ucsd.edu

infant motion than other eye trackers which use close-up viewing of the eye. Morgante, Zolfaghari, and Johnson (in press; henceforth MZJ), however, report findings that this eye tracker is inconsistent in its reporting of event timing. Specifically, discrepancies were found between the timing of eye movements in the text output file from Tobii Studio software vs. a human coding of the movie file of stimulus presentations with gaze position superimposed (Evaluation 1); and between the Tobii Studio movie file and the output file from E-Prime (Psychology Software Tools Inc., Pittsburgh, PA), an external experimental-presentation program (Evaluation 2).

MZJ's findings represent an extremely important demonstration for anyone using, or interested in using, an eye tracker with infant populations. Two questions that may be present in many readers' minds are as follows: How much does this affect the paradigms I use? And, if there is a problem, is there a way to fix it? This report is intended to speak to these questions. I briefly describe some additional, illuminating data from a similar Tobii eye tracker, based on observations of a Tobii T120 eye tracker made in my laboratory in early 2008. I then outline potential corrective measures. The reader should be aware that some of the corrective measures require technical expertise. While MZJ discuss both temporal and spatial inconsistencies, this report only deals with temporal factors.

TEMPORAL RESOLUTION AND EYE TRACKING

The most desirable temporal resolution for any instrument depends on the particular construct the researcher wants to measure. For instance, millisecond resolution would not be sufficient for measuring particles that exist for femtoseconds, but would be overkill for measuring human life spans. This is true of different eye tracking paradigms as well. While some paradigms require high temporal resolution and accuracy, this is unnecessary for other paradigms. First, I discuss various recent options for infant eye tracking, including their timing characteristics. Next, I describe timing information collected on a Tobii T120 in my laboratory. Finally, I outline cases where timing accuracy in eye tracking holds varying degrees of importance—experiments with a weak temporal component, vs. experiments for which timing is more critical—and discuss potential fixes and their feasibility.

VARIOUS EYE TRACKERS AND THEIR PROPERTIES

What are the existing options for infant eye tracking? Older generations of eye trackers designed for use with adults or primates featured either chin/

head rests or heavy head-mounted equipment. These features eliminated the problem of head movement, which adds a degree of complexity in determining gaze direction. However, such features were infeasible for infants and small children. Fortunately, there are now several more appealing options (Table 1). First are *remote* eye trackers like the Tobii, which require little-to-nothing to be placed on the participant. These are now produced by several companies and feature either a camera that is separate from the display monitor or is integrated into the base of the monitor. The major hurdle for these systems is compensating for changes in head position. Some systems, like the Tobii, compensate by using facial recognition software to calculate head angle and distance. The Eyelink 1000 uses a different approach, requiring one piece of "headgear:" a small (approximately 2 cm), high-contrast black-and-white sticker placed on the participant's forehead. This sticker allows the software to calculate head distance and head position. Remote systems have time resolution of 50–500 Hz. Second, several companies now offer eye tracking *glasses*: clear, lightweight safety-goggle-style glasses fitted with small adjustable eye and scene cameras. Some of these have been used with infants (Corbetta, Guan, & Williams, in press; Franchak, Kretch, Soska, & Adolph, 2011). Glasses, like other head-mounted systems, automatically compensate for head movement by moving with the wearer's head. They have the advantage of portability and real-world use, but have slower sampling rates, and data must be hand-coded because the display is not fixed. However, the hand-coding is performed from an eye track superimposed on the scene camera output, which is more straightforward than coding from frame-by-frame images of the eye itself.

Many companies provide software for data collection and visualization, and some interface with third-party software such as E-Prime. Understandably, many companies do not care to disclose their prize-winning software recipes for successful eye tracking or for visually appealing data analysis. However, this is often at cross-purposes with researchers, who—also understandably—need and want to know exactly how the system is working. More "home-brewed" software packages tend to have less appealing graphical user interfaces (although PsyScope and SMART-T [Shukla, Wen, White, & Aslin, 2011] do have graphical interfaces). However, most are more transparent in their operations (and free). Some examples include the PsychToolBox3 for Matlab (Brainard, 1997; Pelli, 1997), which has an embedded Eyelink Toolbox (Cornelissen, Peters, & Palmer, 2002) for use with SR eye trackers; and the Tobii software development kit (SDK), which is now freely available to Tobii users.

One of the more crucial operations that software should perform is keeping accurate time during an experiment. That is, it should accurately match up a running list of gaze locations with event presentations. Without this,

TABLE 1
Properties of some potentially infant-appropriate eye trackers

System	*Hardware*	*Subject wears*	*Sampling frequency*	*Software*	*Software pros*	*Software cons*
ASLD6 High speed remote	Separate camera	Nothing	120–240 Hz	ASL software	Graphical interlace	Unknown
ASL, SMI, Tobii Glasses[a]	Wearable camera	Lightweight glasses and camera	30 Hz	Dependent on brand	Graphical interlace	Must hand-score data
Eyelink 1000 Remote	Separate camera Arm mount available	Sticker	250–500 Hz	Experiment Builder	Graphical interface Infant calibration routines	Unknown
				Matlab	Greater flexibility Timing data collected on eye tracker computer Excellent timing resolution	Programming experience needed
ISCAN ETL-500	Wearable camera	Camera + headgear[b]	60 Hz[c]	ISCAN	Unknown	Unknown
Positive Science	Wearable camera	Lightweight cap, cameras, wireless transmitter	30 Hz	Yarbus software	Graphical interface Already used with infants	Unknown
SMI RED, RED500	Integrated w/monitor	Nothing	RED: 60–120 Hz RED500: 500 Hz	SMI software	Graphical interlace	Unknown

TABLE 1
Continued

System	*Hardware*	*Subject wears*	*Sampling frequency*	*Software*	*Software pros*	*Software cons*
Tobii	Camera separate or integrated w/monitor Arm mount available	Nothing	60–120 Hz (T and X series)	Studio	Easy setup	Slow data output Black-box
				E-Prime	Graphical interface Seems to correct clock drift	Black-box
			300 Hz (TX300 series)	Talk2Tobii	Greater flexibility	No eye tracker time stamp
				PsyScope	Graphical interface	No eye tracker time stamp
				Matlab	Greater flexibility	Must program in Tobii API
Video camera	Humans	Nothing	30 Hz (or camera frame rate)	Humans	Components easily replaced	Must hand-score from eye image Must be trained

Note. For Tobii and Eyelink, data are provided based on the author's knowledge; for other systems, data are provided based on information on manufacturers' web sites. All information should be verified by the reader.

[a] These are three different products.

[b] See Corbetta et al. (in press) for data loss issues with this headgear.

[c] Other ISCAN systems have higher sampling frequencies.

the researcher cannot verify when looks and experimental events occur relative to each other—although one can make rough guesses, which may be sufficient for some purposes. For instance, the Talk2Tobii Matlab system (available at F. Deligianni's web site, http://www.cbcd.bbk.ac.uk/people/affiliated/fani/talk2tobii) that underlies the SMART-T system (Shukla et al., 2011) cannot directly access the eye tracker clock, but Shukla et al. checked the system's delay using Aslin's (in press) testing procedure and found it sufficient for the purposes of their paradigm. Considerations not noted in Table 1 include technical support and cost. Because both vary widely, readers are encouraged to consult colleagues regarding their experiences with technical support and to consult companies to obtain price quotes before making a decision.

OBSERVATIONS ON TIMING ACCURACY

In this section, I discuss potential sources of timing inaccuracy and outline how these relate to my observations on timing in the Tobii T120. One known source of temporal inaccuracy is *clock drift*: any two computer clocks will run at slightly different speeds and will tend to fall out of sync with each other. One solution is to use a single clock to time stamp eye gaze information *and* experimental events. This is how some systems, such as the Eyelink 1000, deal with gaze information: events that the experimenter specifies, such as the onset of a spoken word or appearance of a visual image, are sent from the experimental computer to the computer keeping a running list of gaze data. This of course assumes rapid transmission from the experimental computer to the eye tracker computer—if messages are late, then looks will seem to be *early* relative to stimulus presentation.

A second solution is to have both clocks keep independent time records and reconcile the times afterward. When using Tobii Studio, two clocks are involved: the clock on the eye tracker itself and the clock on the PC running Tobii Studio software. Clock drift is typically reconciled after-the-fact, by measuring the times of each clock at the start point and the end point of an experiment. If the elapsed time interval mismatches, indicating drift, the time on clock is stretched or compressed to fit into the time span of the other clock. Clock drift is one plausible mechanism for MZJ's observed discrepancy between the Tobii Studio output file and the E-Prime output file: E-Prime is correcting for clock drift, while Tobii Studio is not. It is worth noting that clock drift is not usually as large as that seen in MZJ's Experiment 2 (http://www.codinghorror.com/blog/2007/01/keeping-time-on-the-pc.html), although even greater drift has been observed (http://www.

spectron.us/NewWebRoot/ProductivityTools/PcClockDrift/PcClockDriftTheoryMain.php).

A second source of temporal inaccuracy is *communication speed* between the experimental computer and the eye tracking computer. This can occur when using the first solution to clock drift, where external software sends experimental-event information *to* Studio (or other software keeping track of gaze data), which incorporates those events into its data file. Again, this solution requires rapid transmission between not only the computers themselves, but also the software running on those computers. This is where I first noticed timing discrepancies: events sent from PsyScope X (Build 51; Cohen, MacWhinney, Flatt, & Provost, 1993; http://www.psy.cns.sissa.it/) to a PC computer using Tobii's TriggerData interface showed up at unevenly spaced time intervals in the gaze data file. Messages sent at 1-ms intervals appeared in the data file at separations of roughly 4–5 ms (time *x*, time *x* + 4, time *x* + 8, time *x* + 12), with some at sporadically greater delays.

Of course, this could be a result of the experimental-presentation software (PsyScope X), and not the Tobii software. To determine this, the

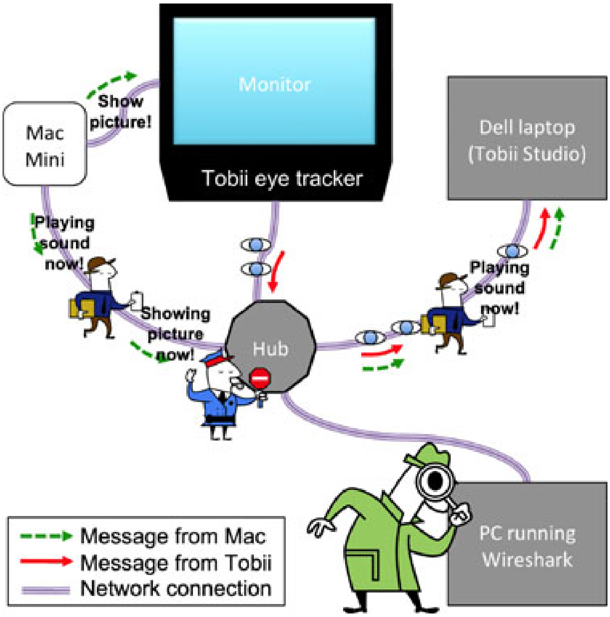

Figure 1 Setup of timing experiment. This is similar to a data collection setup, but the PC running Wireshark would not be present.

communications between experimental, Tobii Studio, and the Tobii T120 eye tracker computer (in 60 Hz mode) were observed over a local area network. Figure 1 depicts the experimental setup. A Mac Mini running PsyScope X sent information about event onsets to a Dell PC laptop running Windows XP Professional 2002 and Tobii Studio, using Tobii's TriggerData interface with the Nagle algorithm (which can delay networked communication) turned off (L. Bonatti, personal communication, 2/15/2008). The eye tracker computer was physically embedded in the experimental-presentation monitor and that monitor was connected to the Mac. The computers were connected through a network hub. The Mac communicated experimental events to the PC, and the eye tracker server communicated gaze information. Another PC laptop connected to the hub ran the program Wireshark (http://www.wireshark.org), which recorded all messages passing through the hub. No participant was being eye tracked (although a track was obtained briefly anyway; see below)—this was simply a test of the timing of the software. The Wireshark computer used its own clock to time-stamp each message's passage through the network.

How rapidly are messages leaving the experimental computer, and how fast are they being inserted into the compiled data file—where are they getting slowed down? Time stamps on messages from the Mac (Figure 2, blue line) appeared on the network at 1-ms intervals (*M* = .999 ms, *SD* = .027 ms, mode = .999 ms, range: 0.898–1.083 ms), as specified in PsyScope X. Thus, the output of TriggerData messages from PsyScope

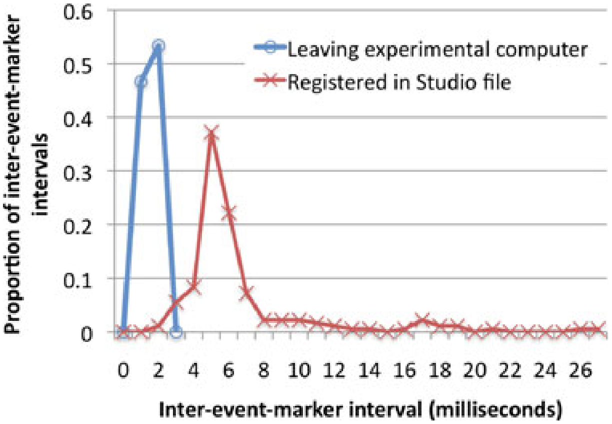

Figure 2 Delays between consecutive event markers, timed in PsyScope X to occur at 1-ms separations. Leaving PsyScope computer (thick line with circles) and appearing in Studio data output file (thin line with X's). Figure does not display full range of IOI's for Studio file. Bin size is 1 ms.

occurred at 1-ms intervals with a small amount of jitter, introducing little temporal variability. However, events were recorded in the PC output file as having occurred anywhere from 1 to 74 ms apart (Figure 2, red line; the graph does not extend this far but the longest delay observed was 74 ms), with a mean of 6.24 ms (*SD* = 7.04; mode = 4). Each successive time stamp in a run of 10 time stamps "pushed" the others forward in time. These delays corresponded closely with the appearance of acknowledgment messages from the PC to the Mac. This suggests that the PC software (Tobii Studio) is imposing the delay. It is not clear why this was happening, particularly since other processor-heavy interfaces (bluetooth, wireless internet) were turned off, but one explanation is that Tobii Studio was occupied with processing other information. Interaction with Tobii technical support (Feb., 26, 2008, personal communication) revealed that the TriggerData interface was "designed for simpler tasks, like sending information on when an image is shown or a movie is started, that do not require millisecond precision."

A constant delay of 6 ms is, on its own, of very little concern for most timing paradigms. What is of greater concern is that the delays are cumulative: this suggests that *other* events occurring temporally close to a message may delay its entry into the Studio data file. Also of concern are the occasional longer delays, which do not just contribute noise; they also create a bias—eyes appear to react *sooner* than they actually do. It should be noted that it is impossible to tell from this analysis whether there is an additional constant delay on the event times, because there is no verifiable clock comparison between the experimental clock and the other two clocks.

Another question is how rapidly gaze information is available. According to the Tobii T60/T120 manual, the maximum delay from the eye tracker acquiring gaze information to making that gaze information available is 33 ms. Is this accurate? I analyzed the same network messages described in the preceding section to find the absolute delay on the availability of gaze data samples. There were two types of messages of interest from the Tobii eye tracker server to the PC: User Datagram Protocol (UDP) messages, which contained Tobii clock time information and are assumed here to be accurate; and Transmission Control Protocol (TCP) messages, which contained both gaze data information and Tobii clock times. By comparing the time stamps in the gaze data TCP messages to the hub clock time (when they passed through the network), I calculated the absolute delay of the gaze data samples. That is, the delay was the current time minus the eye tracker clock time. If this number is positive, there is a delay.

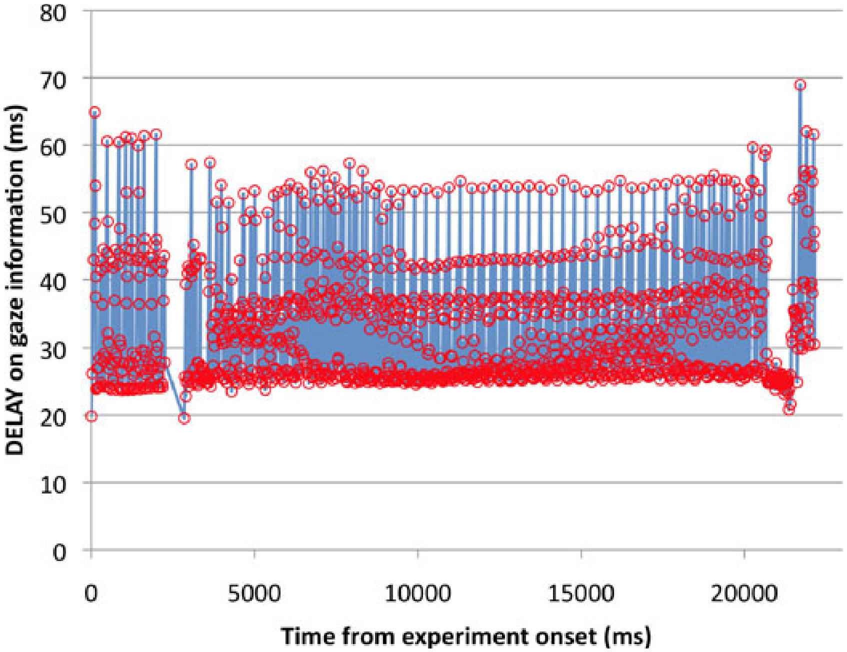

Figure 3 Delays in arrival of each gaze data sample by time during data collection (each circle = one sample). A y-value of zero would indicate no delay.

Delays across the course of the experiment (Figure 3) had a mean of 33 ms and $SD = 8.9$ ms, with 19.7% of samples showing a delay greater than 40 ms (range: 19.5–68.9 ms). Interestingly, the one region where the eye tracker detected gaze (where the blue line becomes flatter briefly around 22,000 ms) shows a smoothing out of these delays ($M = 25.07$ ms, $SD = .90$ ms). This suggests that when gaze position is unknown, delays and delay variability in time stamp arrival increase. In short, analysis suggests that the 33-ms "maximum delay" figure in the Tobii manual is a fairly accurate *mean* value, but not an accurate maximum.

Note that Tobii Studio has been updated since these data were collected. New features include dynamically defined areas of interest, support for the glasses instrument, a new user interface, and a new remote viewer. The SDK is now free, and there is an online application market (http://appmarket.tobii.com/wiki/index.php/Application_Market_for_Tobii_Eye_Trackers). It is not clear whether these fixes address the problems raised by MZJ and me. A more informative document is a recent white paper released by Tobii (http://www.tobii.com/Global/Analysis/Training/White-

Papers/Tobii_Eye_Tracking_Timing_whitepaper.pdf) As Aslin (in press) notes, it is advisable for investigators to test such improvements for themselves.

DISCUSSION

Implications for paradigms with a weak temporal component

Some situations have a high tolerance for temporal noise. If one measures total looking time, temporal noise on the order of that observed in MZJ, or by this author, is unlikely to be a problem. For instance, Turati, Valenza, Leo, and Simion (2005) and Frank, Vul, and Johnson (2009) examined infants' fixations to certain regions (faces or face parts) over others. For Turati et al., differences in total fixation durations were on the order of 500 ms or more; with such a large effect, noise in measured fixation lengths on the order of tens of milliseconds is a relatively low noise level. Others have used eye trackers for *slow gaze-contingent* designs (e.g., Deligianni, Senju, Gergely, & Csibra, 2011; Shukla et al., 2010): presentation of a stimulus at a point where the child is fixating, contingent on the child looking to that location. These are "slow" because fixations typically last 200–300 ms, depending on the task (e.g., Rayner, 1998), giving the computer time to present the stimulus before the eye is able to move again. This relatively lengthy duration allows a good bit of flexibility in eye tracker delay to register an eye movement. Note the distinction between *slow gaze-contingent* changes (in the range of 100–200 ms) and *fast gaze-contingent* changes (< 100 ms) that are designed to occur during a saccadic eye movement. In short, eye trackers with high temporal "noise" (tens of milliseconds) are usable for tasks that measure effects on the order of hundreds of milliseconds. This includes automatic coding of total looking time and slow gaze-contingent display changes. For Tobii eye trackers, use of external software rather than Tobii Studio is warranted, given MZJ's findings of substantial temporal drift in Tobii Studio's recordings relative to E-Prime (see next section).

Implications for paradigms with a strong temporal component

Other situations have less tolerance for temporal noise. For many paradigms, the researcher may want to know the timing of eye movement activity relative to particular stimulus events. Accurate timing of looks relative to stimuli is critical in studies of spoken language processing using the visual-world paradigm (Cooper, 1974; Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995) and, as MZJ note, in studies of visual anticipation. This is because one runs the risk of drawing the wrong conclusion about the cognitive capacities underlying the eye movements if inaccurate timing is present. This concern for temporal resolution also applies more generally to any case where one wants to synchronize time-sensitive information from one source (e.g., a recorded electroencephalogram [EEG]) with another source.

Paradigms that are concerned with timing of eye movements relative to events, as well as those that need almost instantaneous gaze information, are both affected by the observed timing issues.

Tobii Studio, in addition to difficulties reported by MZJ, appears not to keep track of external-software events at the millisecond level, which introduces bias into event-related eye gaze measurements. The Tobii eye tracker *server* shows delays of 20–70 ms in the availability of gaze data samples, possibly due to the algorithm it uses to correct for head position (this is of course speculative as the algorithm is proprietary).

The best solution to relative *and* absolute timing inconsistency is to use the same clock to collect all information and to make sure that this clock is registering events on time. For the Tobii systems tested by MZJ and me, the next-best option must be employed: the experimental clock and the eye tracker clock must be coregistered at both early and late time points during the experiment, by external software. These might include E-Prime or Matlab. E-Prime presumably corrects for clock drift “under the hood.” Matlab (and the freely downloadable Matlab PsychToolBox 3) has the advantage of greater transparency and would require use of the Tobii SDK, which is currently only supported for PCs. This is preferable to using the TriggerData interface (both part of the SDK and incorporated into PsyScope) to send information to the Tobii eye tracker, as TriggerData messages are incorporated by Studio with variable delays and were not designed for millisecond precision. MZJ’s findings (Experiment 2) suggest that using the Tobii Studio movie output (.avi) files, while appealing for conference presentations, is not recommended for analysis because of the introduction of sizable drift of unverified origin.

In principle, clock drift is simple to correct. Clock times should be obtained from the experimental computer and the server at the same time, at both the start and end of the experiment. Further, timing of all experimental events needs to be recorded (or calculated) in terms of the time *from the recorded starting time stamp*, not just from the start of a particular trial. One can do this in various different types of software. For instance, Matlab’s PsychToolBox3 has a GetSecs() function that retrieves the current computer clock time. To get the time in a simple Matlab script, simply type currenttime = GetSecs(). Now the variable currenttime contains the computer clock time. (Note that it is advisable to “prime” this function by calling it once first, so that it is already loaded in memory before using it for consequential time measurements.) Just after running GetSecs(), you would want to get the current clock time from the eye tracker clock. For the Tobii, this is a bit tricky. There is a function in the Tobii SDK to retrieve the eye tracker clock time. However, to my knowledge, this is not implemented in PsyScope X or in Talk2Tobii/SMART-T, which used a non-Tobii-sanctioned Mac version of the SDK. One imagines that this could be implement easily in a new Tobii Matlab toolbox (http://www.tobii.com/eye-tracking-research/global/products/software/tobii-toolbox-for-matlab/features/), which works on PCs, although this toolbox has not been tested.

Once these times are obtained, clock drift is corrected as follows (Figure 4). First, raw clock times may be in seemingly arbitrary units. If so, you will need to subtract out each clock's start time from all following times (Figure 4, top), so that both clocks start at zero. Make sure you know whether each clock is in microseconds or milliseconds. Second, check for drift: if there is drift, then the end time for Clock Y will be different from the end time for Clock X (Figure 4, middle). This can be corrected by applying a linear transformation to adjust one set of times to the units of the other. For instance, if one clock says

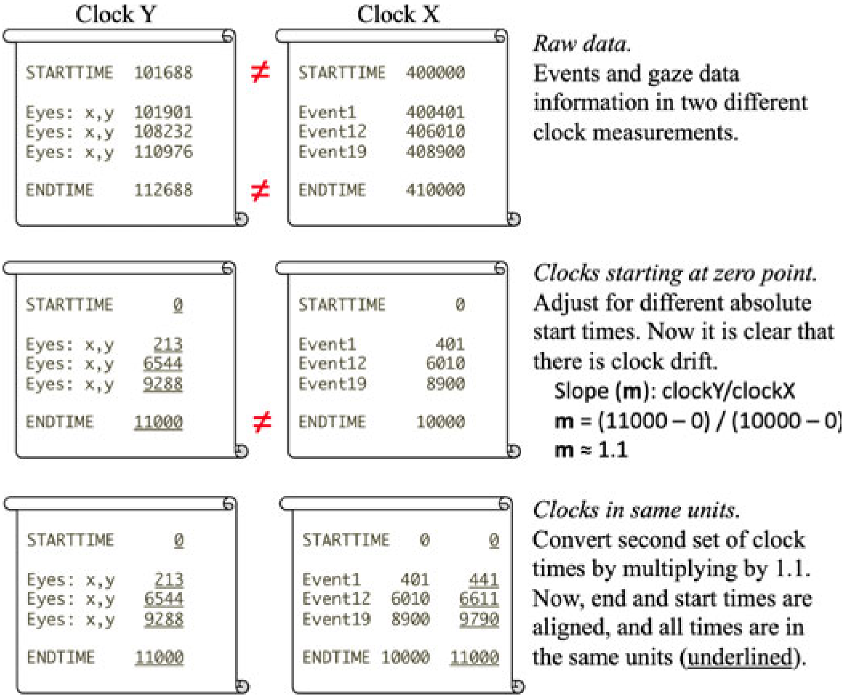

Figure 4 Schematic of how to correct for clock drift. Drift is exaggerated relative to what is typically observed to make effects clear.

that 11,000 ms have elapsed, and the other says that 10,000 ms have elapsed, then multiply the time stamps from the 10,000 ms clock by 11,000 ms/10,000 ms (=1.1). Now the two sets of times are in the same "clock units" (Figure 4, bottom) and can be interdigitated correctly.

Especially great caution is in order when using an eye tracker with event-related measures (EEG/ event related potential [ERP] and magnetoencephalography [MEG]). While the above corrective measures should work with these paradigms as well, they are highly time sensitive: one of the basic assumptions of signal averaging across trials, as is commonly performed in ERP and MEG paradigms, is that signals occur at the same point in time relative to a stimulus. This means that even small amounts of temporal jitter can partially or completely obscure some event-related signals, particularly early ones with brief duration.

Off-line correction for drift and shift should suffice for most purposes. What is not correctable is the speed at which gaze data become available. This would be a concern in fast gaze-contingent presentation, where the computer presenting the experiment needs almost instantaneous access to the current state of visual activity to change a visual display *during a saccade*. Display changes are made during a saccadic eye movement because visual input is temporarily cut off during this time, allowing the change to take place without attracting attention (McConkie & Rayner, 1975; Rayner, 1975). A saccade for children aged 8–18 months lasts for 30–100 ms (Garbutt, Harwood, & Harris, 2006), which is comparable to adults (Rayner, 1998, 2009), and is dependent on the distance the eye must travel—longer distances yield longer saccade durations. Of course, young infants' saccades are hypometric—they saccade in the right direction, but often undershoot (Aslin & Salapatek, 1975; Garbutt et al., 2006), meaning that longer-duration saccades are less likely to occur. To my knowledge, fast gaze-contingent display changes have not yet been used with infants, but in my view it is important to allow for new experimental possibilities in the future. For instance, researchers who study short-term memory and feature binding in infants (e.g., Oakes, Messenger, Ross-Sheehy, & Luck, 2009) might find this a fruitful paradigm.

To change a display during a saccadic "blackout," the experimental program needs not just the most recent gaze information, but also time to plan the screen change (say, 10 ms; although this is merely an estimate) and to execute it. Time also depends on how often the visual display refreshes (a typical rate for modern displays is 60 Hz, or about 16.7 ms, although there are faster displays), when during the eye tracker's sampling period a saccade begins, and the delay on arrival of gaze data information. Backing up from here, the program needs up to 17 ms (refresh time) + up to 17 ms (for a 60-Hz eye tracking sampling frequency) + estimated 10 ms

TABLE 2
Shortest saccade durations for which various combinations of sampling rate and gaze information latency can guarantee execution of a gaze-contingent display change. The sixth column is a total of the second through fifth columns

Sampling frequency (Hz)	*Inter-sample intervals (ms)*	*Maximum gaze info latency (ms)*	*Maximum screen refresh time (ms)*	*Time to execute screen change*	*Shortest saccade tolerated (guaranteed)*
60	16.7	0	16.7	10.0	43.4
60	16.7	25	16.7	10.0	68.4
60	16.7	45	16.7	10.0	88.4
120	8.3	0	16.7	10.0	35
120	8.3	25	16.7	10.0	60
120	8.3	45	16.7	10.0	80
300	3.3	0	16.7	10.0	30
300	3.3	25	16.7	10.0	55
300	3.3	45	16.7	10.0	75

Note. Gaze data latency values were chosen to provide an illustrative range, to show how increasingly later availability of these data increases the saccade length needed to execute display changes.

(execution time) = 44 ms or so, to *guarantee* that screen changes occur before the saccade ends. This is true when gaze data samples are made available as soon as measurements are made. If gaze data are not immediately available, as is the case for the delay in availability adds on to this time. Table 2 presents the fastest saccades that can be *reliably* responded to—that is, all saccades of that and longer durations can be handled—with various combinations of sampling rates and gaze information delay.

Aslin (in press) describes how to use a high-speed camera to simultaneously film both event presentations and eye movements made by an adult. This recording can then be hand-coded and matched up with the corresponding data output from the Tobii or any other eye tracker. The reader is encouraged to consider using this method to pinpoint the degree of noise or bias in their eye tracking systems—in some cases, the level of inaccuracy may not be a problem, and in cases where it does, it is better to know than not to know. It would also be interesting to simulate infant-like movement parameters in this paradigm, to assess how much participant motion affects temporal accuracy of gaze and event data. Of course, unlike other timing issues, speed of data availability is not directly assessed by Aslin's procedure. This is because the software may interdigitate gaze data and event data after the fact based on *recorded* time, not *arrival* time, hiding a substantial delay in real-time data availability. Of course, one could easily adapt Aslin's procedure to test *how rapidly* an eye tracker changed a stimulus display in response to an eye movement.

A final note is that *technically sophisticated* does not necessarily mean *more accurate*. For instance, Aslin's (in press) method of assessing timing is simple but elegant and cost-effective. Inversely, Tobii has developed a Matlab toolkit that would require programming, which on the surface sounds rather complex. However, according to Tobii's web site, this Matlab toolkit is slow to process eye movement data, responding within 45–80 ms after receiving that data. This means the experimental-presentation computer would be somewhat slow in executing fast gaze-contingent display changes based on eye movements (although this may be acceptable for slow gaze-contingent paradigms).

Finally, it should be noted that there are applications for which the Tobii is well suited. These include work with populations who cannot tolerate the presence of head position monitors (glasses and camera on the head) or stickers (Eyelink Remote), such as non-human primates. The Tobii is also well suited to augmenting communication in communication-impaired adults.

CONCLUSIONS

This report briefly presents timing tests with a Tobii T120 eye tracker similar to the T60XL tested by Morgante et al. (in press). I discuss the implications of Morgante et al.'s and my observations for infant experiments using this device to collect data. For paradigms with a weak temporal component, the observed problems were of limited concern. For situations requiring accurate timing relative to presented events, using Tobii Studio to collect timing information on external events is not recommended, because time stamps are not incorporated accurately (this report), and it may not correct for clock drift (MZJ, Experiment 2). A potential fix is described. For situations requiring immediate access to gaze data information, such as fast gaze-contingent display changes, the Tobii eye tracker as tested by this author is not recommended. Note that the newest Tobii eye tracker (the T300) is designed to have a shorter sampling period and gaze information latency, reducing some of the concerns mentioned here. Of course, caution suggests testing of any eye tracker-software combination, rather than taking its reported timing properties at face value.

I hope that this discussion of eye tracking options and timing concerns has been illuminating. I close by noting that there is no single ''best'' eye tracker. The decision of what eye tracker to purchase ultimately rests on the range of tasks the experimenter wishes to use it for, the skill level of researchers in the laboratory, and the cost-effectiveness of one solution over another.

ACKNOWLEDGMENTS

Thanks to Mark Wallen for invaluable technical assistance, to Keith Rayner for comments on an earlier draft, and to Marta Kutas for encouragement.

REFERENCES

Aslin, R. N. (in press). Infant eyes: A window on cognitive development. *Infancy*, doi: 10.1111/j.1532-7078.2011.00097.x

Aslin, R. N., & Salapatek, P. (1975). Saccadic localization of visual targets by the very young human infant. *Attention, Perception & Psychophysics*, *17*(3), 293–302.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.

Cohen, J. D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavior Research Methods, Instruments & Computers*, *25*, 257–271.

Cooper, R. M. (1974). The control of eye fixation by the meaning of spoken language: A new methodology for the real-time investigation of speech perception, memory, and language processing. *Cognitive Psychology*, *6*, 84–107.

Corbetta, D., Guan, Y., & Williams, J. L. (in press). Infant eye-tracking in the context of goal-directed actions. *Infancy*, doi: 10.1111/j.1532-7078.2011.00093.x

Cornelissen, F. W., Peters, E. M., & Palmer, J. (2002). The eyelink toolbox: Eye tracking with MATLAB and the psychophysics toolbox. *Behavior Research Methods, Instruments, & Computers*, *34*(4), 613–617.

Deligianni, F., Senju, A., Gergely, G., & Csibra, G. (2011). Automated gaze-contingent objects elicit orientation following in 8-months-old infants. *Developmental Psychology*, *47*(6), 1499–1503.

Franchak, J. M., Kretch, K. S., Soska, K. C., & Adolph, K. E. (2011). Head-mounted eye-tracking: A new method to describe the visual ecology of infants. *Child Development*, *82*(6), 1738–1750.

Frank, M. C., Vul, E., & Johnson, S. P. (2009). Development of infants' attention to faces during the first year. *Cognition*, *110*(2), 160–170.

Garbutt, S., Harwood, M. R., & Harris, C. M. (2006). Infant saccades are not slow. *Developmental Medicine and Child Neurology*, *48*(8), 662–667.

McConkie, G. W., & Rayner, K. (1975). The span of the effective stimulus during a fixation in reading. *Attention, Perception & Psychophysics*, *17*(6), 578–586.

Morgante, J. D., Zolfaghari, R., & Johnson, S. P. (in press). A critical test of temporal and spatial accuracy of the Tobii T60XL eye tracker. *Infancy*, doi:10.1111/j.1532-7078.2011.00089.x

Oakes, L. M., Messenger, I. M., Ross-Sheehy, S., & Luck, S. J. (2009). New evidence for rapid development of color-location binding in infants' visual short-term memory. *Visual Cognition*, *17* (1–2), 67–82.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442.

Rayner, K. (1975). The perceptual span and peripheral cues in reading. *Cognitive Psychology*, *7*, 65–81.

Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, *124*(3), 372–422.

Rayner, K. (2009). Eye movements and attention in reading, scene perception, and visual search. *Quarterly Journal of Experimental Psychology*, *62*(8), 1457–1506.

Shukla, M., Wen, J., White, K. S., & Aslin, R. N. (2011). SMART-T: A system for novel fully automated anticipatory eye-tracking paradigms. *Behavior Research Methods*, *43*(2), 384–398.

Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, *268*(5217), 1632–1634.

Tobii Technology AB. (2010). *Timing guide for Tobii eye trackers and eye tracking software*. [White paper]. Retrieved from http://www.tobii.com/Global/Analysis/Training/WhitePapers/Tobii_Eye_Tracking_Timing_whitepaper.pdf

Turati, C., Valenza, E., Leo, I., & Simion, F. (2005). Three-month-olds' visual preference for faces and its underlying visual processing mechanisms. *Journal of Experimental Child Psychology*, *90*(3), 255–273.