Privacy Considerations for a Pervasive Eye Tracking World

Daniel J. Liebling
Microsoft Research
One Microsoft Way
Redmond, WA 98052
USA
danl@microsoft.com

Sören Preibusch
Microsoft Research
21 Station Road
Cambridge CB1 2FB
United Kingdom
spr@microsoft.com

UbiComp '14, September 13 - 17 2014, Seattle, WA, USA
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3047-3/14/09…$15.00.
http://dx.doi.org/10.1145/2638728.2641688

Abstract

Multiple vendors now provide relatively inexpensive desktop eye and gaze tracking devices. With miniaturization and decreasing manufacturing costs, gaze trackers will follow the path of webcams, becoming ubiquitous and inviting many of the same privacy concerns. However, whereas the privacy loss from webcams may be obvious to the user, gaze tracking is more opaque and deserves special attention. In this paper, we review current research in gaze tracking and pupillometry and argue that gaze data should be protected by both policy and good data hygiene.

Author Keywords

Eye tracking; gaze tracking; biometrics; privacy

ACM Classification Keywords

K.4.1 Public Policy Issues: Privacy

Introduction

Pervasive eye gaze tracking provides not only new interaction techniques on a variety of devices, but also the unprecedented ability to understand human attention at scale. While much of the literature focuses on psychometrics or human computer interaction, little if no attention has been paid to the privacy implications of collecting vast amounts of this new kind of personal information.

Other technologies have reached this inflection point before; the location community saw GPS exponentially grow from a military application to near ubiquity in only a few decades. Practitioners in that field created important work examining the implications of pervasive location tracking from personal, institutional, and societal approaches. Genomic privacy is another example. The falling costs of sequencing lead to voluntary, commercially built databases of genomes that allow identifying relatives of patients and their health status. Comparatively, modern gaze trackers have existed for a few decades and prices have dropped exponentially in the early 2010s.

We believe the gaze community also needs to have these conversations. In our own interest, we should clarify acceptable procedures for collecting and using gaze data at scale. The promise that lies in ubiquitous eye tracking will only materialize when consumers comfortably adopt gaze as a new mode of human-computer interaction. Technology assessment is needed.

This paper is a first step towards a privacy impact assessment of eye tracking. We argue that gaze is unique because it reveals the subconscious in ways that are difficult to control. We review current research on personal information that can be disclosed by analyzing eye tracker data. We frame the discussion on gaze tracking privacy in terms of *what* data is collected, *who* is involved and the *scale* of the system. Finally, we draw analogies to other biometrics in suggesting possible avenues for addressing privacy concerns.

Gaze Data is Unique

Gaze data differs from other signals of human activity precisely because gaze and associated data like blinks and pupillometry are not fully under volitional control. We can disguise our voices to fool speech recognizers; alter our appearances with clothing and makeup, and change our keystrokes to defeat keyloggers; however, we have only partial control of our gaze. Although we can avert our eyes and direct our vision, the subtleties of eye movement are not under conscious control. It is hard to prevent a fleeting glance; pupil dilation is even more difficult, if impossible to control. It is cognitively and physically tiring to maintain top down control over one's gaze. Further, many of the sensitive attributes derivable from gaze data are not borne from *what* we look at, but *how* we look, which is harder to control.

Gaze Reveals Sensitive Attributes

At its core, eye movement data is usually coded as (*x*, *y*, time) tuples. Depending on the gaze tracker, pupil dilation may also be reported. This time series data is noisy due to biological noise (ocular microtremor (OMT), nystagmus, etc.) as well as ambient illumination and sensor uncertainty. In many applications, eye movements are condensed into fixations that approximate the focus of attention. Ordered in time, the sequence of fixations sequence comprises a scanpath. Such data may be coupled with details about the underlying stimuli (e.g., areas of interest displayed on screen), creating a richer notion of both *what* was attended to and *how* attention varied. Even without knowledge of the stimuli, some scanpaths are highly stereotyped and recognizable. Knowing how and certainly at what people gaze provides a wealth of understanding. Yet gaze data intended for a single purpose such as evaluating a new user interface can unwittingly reveal sensitive attributes of participants when it is analyzed more deeply.

Physical attributes

As humans age, their scanpaths change. At birth, genetically programmed scanpaths guide the way infants look at the world. Later in life, we assess scenes in ways that are influenced by prior knowledge, preference, and task [1][3]. Biological aging also affects how saccades occur during certain tasks [15]. Thus, one may infer approximate age of the subject given appropriate data. Besides age, symptoms of various neurological and behavioral disorders can present as eye movement abnormalities [4][17].

Interests and social attributes

Faces are extremely salient parts of a scene. Eyes scan faces in highly stereotyped ways [1]. Given a single face, we can guess whether the viewer is familiar with that person. Given a set of faces, we are drawn to each face in an order that is influenced by race, gender, sexual preference, and socialization. Starting in infancy, Own-Race Face Bias affects how individuals view faces of their own race versus other races [1]. Similarly, given a set of faces judged to be attractive, one tends to attend more to faces from the gender to which one is attracted. Finally, face processing is mediated by cognitive processes that are modulated by brain pathology. For example, autistic individuals' facial scanpaths differ from controls [5].

Measurement of pupillary response (pupillometry) can also provide information about the participant. In earlier research, Hess et al. found that pupil size is related to degree of interest in a scene [9]. Recent research shows that females' pupil diameter changes in response to viewing images of their partner or an attractive actor is modulated by their hormonal cycles [13]. Body mass index (BMI) may even be estimated by presenting a set of images of foods of varying nutritional content [8].

Attribute	Sources
Age	Scanpath, microtremor
Gender	Scanpath
Race	Scanpath
Affect	Scanpath, pupil dilation
Sexual preference	Scanpath, pupil dilation
Body mass index	Pupil dilation
Hormonal cycle	Pupil dilation
Health	All
Task focus	Scanpath, pupil dilation

Table 1. Summary of some sensitive attributes discernable from gaze data. Sources are not exhaustive.

Pupil diameter need not change only in response to visual stimuli; for example, the response can change when subjects encounter emotionally charged sounds [16]. Pupil response also appears to be modulated by subconscious processing; it is thought to provide a signal of cognitive load. Change in pupil diameter has been shown to be evidence of off-task mind wandering [22] as well as "offline" processing of information during non-task time [19].

Activity and expertise

From gaze data, systems can also infer skill-level differences on certain tasks [14]. Even without knowing the stimulus, some tasks are easy to identify. Reading activity is quite clear; the eyes make most progress in a single direction with a small amount of regressions, concluding with return sweeps at the end of a line. However, the behavior is culturally sensitive; bilingual readers have different saccade and fixation patterns while reading their dominant language [17].

Gaze Uniquely Identifies Individuals

Since our attention is modulated by both bottom-up visual features (e.g., saliency maps) and top-down volitional control, influenced by culture and life experience, it should come as little surprise that gaze patterns can uniquely identify individuals. Bednarik et al. obtained 60% accuracy using just pupil diameter measurements over one second periods while viewing a still object [2]. Following the trend of machine learning contests, the first Eye Movement Verification and Identification Competition [10] took place in 2012. The four datasets included 250-1000 Hz recordings from two eye trackers, as participants followed a jumping dot. The best models achieved accuracy from 58% to almost 98% using models of movement speed and direction.

Even if the individual's gaze data has not previously been recorded and associated with his or her identity, one could use attributes derived from the gaze patterns to approximate that identity. For example, gender identification effectively halves the global search space. By carefully selecting combinations of attributes, the identity entropy can be greatly reduced. This is particularly applicable to research studies where small participant groups, typically in close geographic proximity, are measured. Research groups releasing "anonymized" gaze tracking output may inadvertently reveal the (at least plausible) identities of their subjects.

Resulting loss of privacy

The data that is collected through gaze tracking results in privacy losses of two kinds: first, the *identity* of an individual, as the user's unique gaze pattern allows fingerprinting. Part of the identity inference are also several bio-indicators, including health status, both momentary and longer term, age, and fatigue. Second, the inference of *interests*, which need to be understood broadly. When annotated by the semantics of the content displayed, the measurement of interest in items displayed on screen reveals political, sexual, cultural or other lifestyle preferences.

The leakage of information about identity and interests violates the privacy principle of informational self-determination. There is a twofold loss in users' ability to "determine for themselves when, how, and to what extent information about them is communicated to others" [24]. In the United States, actionable privacy protection is achieved through *notice* and *choice*. Ubiquitous gaze tracking puts both principles at risk. First, users are unable to voluntarily control their gaze; they are thus disempowered to make choices to withhold their data. Second, there is no effective mechanism to communicate to users what information their gaze is leaking. Optical cameras such as webcams can project back the image they are recording. In video telephony, such as Skype, a small playback of themselves allows users to control their "image." Projecting the gaze on screen would be highly distracting at best;

projecting the inferred interests and biological attributes seems unlikely.

Who is tracking the eye

Behind the lens

The proliferation of eye tracking will mean that only some of the devices able to record users' gaze will be under their control. Imagine an eye tracker built into a smartphone. When used as an input device, the user may have an expectation of privacy to the extent that the data is only recorded to navigate the user interface. However, this expectation may be ill-conceived and the gaze data may leave the device and be repurposed.

A car equipped with gaze tracking could use the data to warn the driver when fatigue is setting in, but also pass on these observations to a pay-as-you-drive car insurer. Recorded fixations could be sent to the car manufacturer that uses the telemetry data to improve the layout of the cockpit. In a connected world of owned, co-owned and third party devices, users will often be unable to determine the recipients of their gaze data. At the same time, dual use of the recordings (as input mode and observer) will make it impractical to shut down or occlude the tracker.

Scale

The scale of tracking data is a key consideration; if a participant performs one laboratory experiment, then that data is unlikely to be released and there is probably little impact to the subject. But when eye trackers are pervasive and sharing their data with a central infrastructure, the opportunity to track individuals across time and place now becomes real. As an analogy, consider moving from a single building under video surveillance to city-wide CCTV coverage. Privacy concerns heighten when real-time tracking of individuals becomes possible.

Collections of gaze recordings already exist and could potentially be joined. Vrzakova and Bednarik [23] introduced the EyeCloud concept, calling for a corpus of eye movement. Such corpora are beginning to come together; since 2012, The Eye Movement Verification and Identification Competition has provided a set of gaze data across subjects with the goal of facilitating a competition to most accurately identify individuals based on their gaze patterns. Larger corpora gleaned from a variety of tasks will no doubt be released in the future. Already, several public eye tracking installations exist; the Eye-Follower display at Cité des Sciences et de l'Industrie in Paris began in 1986; a three-month exhibit at the National Gallery in London captured data from almost 10,000 individuals [25]. The privacy concern is not so much with the collection of data, but with the potential for identification and sharing in ways that go beyond the users' expectations.

Even if data stays on personal devices, there is still some disclosure risk. Depending on the device, gaze data could be discovered as evidence during court proceedings. In the United States, unrestricted seizure of digital content on mobile phones without warrant is prohibited due to a recent Supreme Court case [19]. However, data collected at public kiosks or access points, if retained indefinitely, could inadvertently become public.

Privacy and Potential Remedies

The gaze community should begin to consider privacy affordances in eye tracking systems as they become more pervasive. It is unreasonable to expect a user to understand the mapping of raw data to sensitive attributes. Instead, it is up to the developers to inform the user in a comprehensible way about the data being collected and its potential implications, and let the user limit the data in sensible ways. We now discuss some ways in which gleaning of sensitive attributes from unintended disclosure of eye tracking data can be mitigated.

Affordances for self-introspection
Although viewing the gaze tracker output in real-time is problematic (due to the feedback problem, or not wishing to disrupt the experience), letting users observe the aggregate data collected by the browser can be informative. Initiatives like "affective mirrors" provide replay functionality so that individuals can observe themselves as others see them, augmented with predictions made based on sensor observations. Desktop eye trackers or operating systems that use them could provide similar interfaces, increasing awareness and grounding informed consent. Users could "keep an eye" on their exposure through a dashboard that summarizes recorded data and inferences from it, akin to bandwidth quota or performance monitors.

Levels of abstraction
Although storing the raw tracker output at each timestamp is tempting, in most cases, fixations suffice to know the focus of attention. Abstracting higher than fixations, regions of interest provide a coarser notion of attention. The latest Tobii EyeX API takes this approach; developers specify regions of interest and are notified when the gaze enters. In this fashion, the consuming application is unaware of the underlying fine eye movements. The underlying principle of hit/no hit checks has been applied in other privacy-preserving architectures, including the Prüm Convention [4] that foresees the sharing of biometric databases with fingerprints and DNA samples amongst European states.

Fuzzing
The location community is well aware of the privacy implications of sharing location data [12] since it reveals not only our home and workplaces but also sensitive personal and family attributes, expressed through a sequence of spatial locations. Individuals are rightfully uncomfortable about sharing visits to sensitive places like doctors' offices or political offices.

One potential remedy to extracting sensitive features is to add noise to the original data before passing it down the application chain. Given that eye tracking has many downstream applications from reading detection to word identification to control activation, it seems that the noise might need to be generated with the application in mind. For reading detection, additional gaze or fixation points could be included in the stream, generated by models parameterized from saliency maps or cross-population empirical data. Ideally the noise is not separable from the true points; they should be robust to state estimation techniques. Additional research to determine the impact of fuzzing is necessary.

Physical barriers
Shielding is a simple yet effective way to avoid eavesdropping. Since most eye trackers use infrared illumination, IR-filtering eyeglass lenses could be used to block tracker functionality. Although this empowers the user to opt out of tracking, if gaze interaction techniques become ubiquitous, then there is a large

potential downside. Advances in tracking hardware and data processing algorithms may also make it possible to detect the gaze with visible light only. Users of filtering eyeglasses can protect themselves from open and covert trackers, but whoever wears them could be subject to social stigma (“tin foil hat”).

Policy and regulation
Policy remedies can take many forms. At the most basic level, documents like privacy statements detail the extent of data collection and use. Companies in the United States and European Union are often required by regulation or law to disclose these in online services through acts like the Data Protection Directive (EU). They embody the principle that users should know when data is collected about them and what will happen with the data. Applied to video surveillance, this same principle requires notices for places under CCTV. Some countries require standardized pictograms [6]. Covert video surveillance is an exception. Similarly, notices should be displayed in areas where eye tracking is deployed and covert eye tracking should be avoided.

Lawmakers should acknowledge gaze data as a form of biometric data that warrants particular protection. Based on the possible inferences such as race and sexual preference, gaze data should be included in the list of data for which the law foresees special safeguards [6]. As with other privacy-invasive technologies, vulnerable user groups need special attention.

Tracking indicators and controls
Modern laptops have status LEDs to show that the microphone or the webcam are recording. Admittedly, such status indicators are missing on mobile devices, although they display an icon onscreen when location tracking is active. It is common for malware to enable sensors without letting users know.

Devices, preferably in the physical hardware, should provide feedback to users to let them know their gaze is being tracked. If possible, the depth of tracking should be communicated, for example, whether the eye tracker is turned on, if multiple persons are within range, and if the gaze is actively being tracked.

System- and application-level privacy controls allow users to control what and to whom information is disclosed. The information may not be explicit; Facebook recently allowed users to control implicit profile attributes inferred from explicit user activity. Kinect fitness games on the Xbox One video game console sense heart rate through video and users can delete this information by visiting a Web site [17]. Similarly, eye tracking hardware and software should offer controls to turn off tracking or disable certain features.

Gaze can enhance privacy

Although gaze tracking data is a rich data source with sensitive privacy issues, gaze trackers may enable new privacy protection affordances for individuals. For example, gaze tracking combined with displays that provide multiple views (through shutters or by exploiting viewing angle [11]) could safeguard private information across multiple users of the same display.

Eye tracking with personal identification could also be used to subtly change content based on the viewer such that the displayed message is altered during saccades (when the visual system is effectively “paused”), thereby imperceptibly delivering private messages to the intended user. As with other non-invasive biometrics, eye tracking can also improve security through continuous user authentication.

Conclusion

With decreasing cost of gaze trackers, pervasive eye tracking is likely to become reality. Although many input modalities have side channels that can identify or classify the user, gaze tracking is unique because it reveals personal attributes that are difficult to disguise. With the coming explosion of gaze tracking data, researchers and practitioners must be conscious of inadvertently exposing their experimental subjects and application users to unintentional privacy leaks. By taking a minimal approach to only process, store, analyze and share only the data necessary to accomplish a task, privacy risks can be moderated. Finally, informing the user and crafting policies that inform and grant control help restore the user's agency. We believe the benefits of pervasive eye tracking are vast, but as with most technology, a measure of caution upfront will benefit the public in years to come.

References

[1] Bar-Haim, Y., Ziv, T., Lamy, D., and Hodes, R.M. Nature and Nurture in Own-Race Face Processing. *Psychological Science 17 (2)*. 2006.

[2] Bednarik R., Kinnunen T., Mihaila A., and Fränti P. *Proc. 14th Scandanavian Conference on Image Analysis*. LNCS 3540, pp. 780-789. 2005.

[3] Borji, A. and Itti, L. Defending Yarbus: Eye movements reveal observers' task. *Journal of Vision, 14 (3)*. 2014.

[4] Convention between the Kingdom of Belgium, the Federal Republic of Germany, the Kingdom of Spain, The French Republic, the Grand Duchy of Luxembourg, the Kingdom of the Netherlands and the Republic of Austria on the stepping up of cross-border cooperation, particularly in combating terrorism, cross-border crime and illegal migration. UNTS 2617. 2005.

[5] Dalton, K.M., Nacewicz, B.M., Johnstone, T., Schaefer, H.S., Gernsbacher, M.A., Goldsmith, H.H., Alexander, A.L., and Davidson, R. J. Gaze fixation and the neural circuitry of face processing in autism. *Nature Neuroscience, 8*. pp. 519-526, 2005.

[6] DIN Deutsches Institut für Normung e. V. DIN 33450. *Graphisches Symbol zum Hinweis auf Beobachtung mit optisch-elektronischen Einrichtungen (Video-Infozeichen)* [Graphical symbol for information about surveillance with optical-electronic devices (video-info signs)], 2014.

[7] Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. 24 October 1995.

[8] Graham, R., Hoover, A., Ceballos, N.A., and Komogortsev, O.V. Body mass index moderates gaze orienting biases and pupil diameter to high and low calorie food images. *Appetite, 56 (3)*. 2011.

[9] Hess, E.H. and Polt, J.M. Pupil size as related to interest value of visual stimuli. *Science, 132*, pp. 349-350. 1960.

[10] Kasprowski, P., Komogortsev, O.V., and Karpov, A. *Biometrics: Theory, Applications and Systems (BTAS) 2012 IEEE 5th Int'l. Conf. on Biometrics Compendium.* 2012.

[11] Kim, S., Cao, X., Zhang, H., and Tan, D.S. Enabling Concurrent Dual Views on Common LCD Screens. *Proc. ACM Conference on Human Factors in Computing Systems (CHI) 2012*.

[12] Krumm, J. A survey of computational location privacy. *Personal and Ubiquitous Computing, 13 (6)*. 2009.

[13] Laeng, B. and Falkenberg, L. Women's pupillary responses to sexually significant others during the hormonal cycle. *Hormones and Behavior*, *52* (4). 2007.

[14] Liu, Y., Hsueh, P.-Y., Lai, J., Sangin, M., Nüssli, M.-A., and Dillenbourg, P. Who is the expert? Analyzing gaze data to predict expertise level in collaborative applications. *Proc. IEEE Int'l. Conf. on Multimedia and Expo 2009*.

[15] Munoz, D.P., Broughton, J.R., Goldring, J.E., and Armstrong, I.T. Age-related performance of human subjects on saccadic eye movement tasks. *Experimental Brain Research 121 (4).* 1998.

[16] Partala, T., Jokiniemi M., and Surakka V. Pupillary responses to emotionally provocative stimuli. *Proc. ETRA 2000.*

[17] Privacy in Xbox One and Kinect. http://www.microsoft.com/security/online-privacy/xbox.aspx. Retrieved 1 July 2014.

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

[19] Riley v. California, 573 U.S. ___ (2014).

[20] Smallwood J., Brown, K.S., Tipper, C., Giesbrecht, B., Franklin M.S., et al. Pupillometric Evidence for the Decoupling of Attention from Perceptual Input during Offline Thought. *PLoS ONE, 6 (3)*. 2011.

[21] Umphress, D., and Williams, G. *Int'l. J. Man-Machine Studies, 23 (3)*. 1985.

[22] Uzzaman, S., & Joordens, S. The eyes know what you are thinking: eye movements as an objective measure of mind wandering. *Consciousness and cognition*, *20* (4). 2011.

[23] Vrzakova, H., and Bednarik, R. EyeCloud: Cloud Computing for Pervasive Eye-Tracking. *Proc. 3rd Int'l. Workshop on Pervasive Eye Tracking and Mobile Eye-Based Interaction (PETMEI) 2013.*

[24] Westin, A. *Privacy and Freedom*. New York: Atheneum, 1967.

[25] Wooding, D.S., Mugglestone, M.D., Purdy, K.J., and Gale, A.G. Eye movements of large populations: Implementation and performance of an autonomous public eye tracker. *Behavioral Research Methods 34 (4)*. 2002.

[26] Yarbus, A.L. *Eye Movements and Vision*, New York, USA: Plenum. 1967.