GazeDrone: Mobile Eye-Based Interaction in Public Space Without Augmenting the User

Mohamed Khamis[1], Anna Kienle[1], Florian Alt[1,2], Andreas Bulling[3]
[1] LMU Munich, Germany
[2] Munich University of Applied Sciences, Germany
[3] Max Planck Institute for Informatics, Saarland Informatics Campus, Germany
mohamed.khamis@ifi.lmu.de, florian.alt@ifi.lmu.de, bulling@mpi-inf.mpg.de

ABSTRACT

Gaze interaction holds a lot of promise for seamless human-computer interaction. At the same time, current wearable mobile eye trackers require user augmentation that negatively impacts natural user behavior while remote trackers require users to position themselves within a confined tracking range. We present GazeDrone, the first system that combines a camera-equipped aerial drone with a computational method to detect sidelong glances for spontaneous (calibration-free) gaze-based interaction with surrounding pervasive systems (e.g., public displays). GazeDrone does not require augmenting each user with on-body sensors and allows interaction from arbitrary positions, even while moving. We demonstrate that drone-supported gaze interaction is feasible and accurate for certain movement types. It is well-perceived by users, in particular while interacting from a fixed position as well as while moving orthogonally or diagonally to a display. We present design implications and discuss opportunities and challenges for drone-supported gaze interaction in public.

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords

Gaze Interaction; Drones; Active Eye Tracking; UAV

INTRODUCTION

Being a fast and natural modality, gaze holds a lot of potential for seamless human-computer interaction.

To date, mobile and remote eye tracking are the predominant technologies to enable such interactions [18]. Mobile eye trackers rely on head-mounted cameras to track users' absolute point of gaze and movements of the eyes. In contrast, remote eye trackers use cameras placed in the environment, e.g., attached to a display. While mobile trackers allow for

Figure 1. GazeDrone is a novel system for gaze interaction in public space. We use the drone's camera to allow users to interact via gaze from random positions and orientations relative to the system and even while moving.

free movement and continuous tracking, they currently require heavy user augmentation, which makes users not behave naturally in public [23]. Remote trackers do not require augmentation, but their tracking range is limited to about 60 to 90 cm in front of the tracker [13]. Gaze estimation accuracy then degrades as users move away from the tracker [12].

At the same time, Unmanned Aerial Vehicles (UAV), also known as drones or quadcopters, have entered the mainstream consumer market. Drones have become increasingly equipped with a multitude of sensors, such as GPS, accelerometers, and recently also high-resolution cameras. While drones are associated with privacy concerns [28], which we discuss later in this paper, they also present opportunities for seamless pervasive interactions. Previous work explored interacting with drones via explicit input such as mid-air gestures [7], and via implicit input such as body motion [22] or facial recognition [6]. While these works focused on interaction *with* the drone, in this work we focus on gaze interaction *through* the drone. Drones can be regarded as portable interactive platforms that can sense the user's input and channel it to surrounding pervasive systems, such as public displays.

Figure 2. The drone continuously streams the video feed to a NodeJS server. After detecting the face landmarks, we measure the distance between the inner eye corner and the pupil for each eye (D_{left} and D_{right}). The ratio determines if the user performed a gaze gesture. The values were decided based on a series pilot studies with 11 participants.

To address the limitations of mobile and remote eye tracking, we present GazeDrone, the first system that combines a camera-equipped aerial drone with a computational method to detect sidelong glances for spontaneous (calibration-free) gaze-based interaction. GazeDrone inspires fresh thinking about a whole new range of gaze-enabled applications and use cases. For example, rather than requiring users to move into the tracking range of a remote tracker and position themselves properly, the drone could instead approach the user and conveniently track their eyes in their current location. This would enable hands-free gaze interaction with physically unreachable systems such as mid-air displays [24] or large distant displays. GazeDrone advances the state of the art in eye tracking by allowing gaze interaction (1) without augmenting the user, (2) from arbitrary positions, distances, and orientations relative to the interactive system, and (3) without restricting movements (i.e., users could interact via gaze while on the move).

The contributions of this work are threefold. First, we introduce the concept and implementation of GazeDrone, a novel system that enables pervasive gaze-based interaction in public spaces through an aerial drone. Second, we report on an evaluation of GazeDrone to investigate its performance and user acceptance. Third, we present four design implications that inform future drone-supported gaze interactions.

RELATED WORK

Researchers have investigated ways for enabling gaze-based interactions beyond the desktop, specifically also displays in public. For example, previous works explored gaze gestures [9] and smooth pursuit eye movements [27] for public display interaction. Similar to our work, Zhang et al. introduced SideWays, a gaze-based system that responds to the user's gaze gestures to the left and to the right [29]. For all of these techniques, enabling interactions from different positions relative to the display remains one of the most important and under-investigated challenges [13, 14].

One approach to address this is to actively guide users into the tracker's range. Zhang et al. investigated ways for guiding users to position themselves in front of the center of a public display [30]. It took their users 4.8 seconds to align the face correctly based on an overlaid outline. In GravitySpot, visual cues implicitly guide users to a public display's sweet spot (e.g., eye tracker's range) [2]. Another approach is to rely on mobile eye tracking to continuously track the relative position

of the display in the tracker's scene camera. For example, GazeProjector utilizes feature matching to detect surrounding displays and map gaze points onto them; the gaze points are transferred through a local WiFi network, to which the displays and the eye tracker are connected [15]. A third approach is active eye tracking, which refers to systems that use, for example, pan-and-tilt cameras to adapt to the user's head position [20, 25]. While all of these approaches allow for more freedom in user positioning, they either require user augmentation or their range is still limited and prohibitive for interactions from far away from the eye tracker. The only exception is EyeScout [14], where an eye tracker was mounted on a rail system to allow the tracker to follow the user along the display. However while this approach significantly increases the lateral range of eye tracking, it is confined by the eye tracker's range, which is typically 60-90 cm [13].

In contrast, GazeDrone does not require user augmentation and users are also not required to walk into the eye tracker's range. While gaze has been used to remotely operate drones from a desktop computer [11], GazeDrone is first to leverage aerial drones for active eye tracking and thereby enable interactions with nearby interactive systems. Figure 1 illustrates a sample use case, in which users can interact from an arbitrary position relative to a public display.

GAZEDRONE

GazeDrone consists of three main components: a server, a client, and an aerial drone. Previous solutions from industry and research have already demonstrated the feasibility of tracking and following users through drones [6, 21, 22]. In this work we focus exclusively on gaze interaction through the drone's camera. As illustrated in Figure 2, we use a Parrot AR.Drone 2.0[1] to continuously transfer the video stream to the server via WiFi. The video stream is then processed on a NodeJS server. The server runs the tracking algorithm and estimates the user's gaze. The gaze data is then pushed to the client. The client can be programmed as desired, depending on the use case. For example, it could run Android Things for IoT applications, or a web browser for web-based interfaces.

We detect gaze gestures (left and right) in real time to evaluate GazeDrone's suitability for interactive applications and while users are moving. Using the front camera of the AR.Drone,

---

[1]PARROT AR.DRONE 2.0 **https://www.parrot.com/global/drones/parrot-ardrone-20-elite-edition**

we stream the video feed (640 × 360 px at 7–9 fps) to the server. Facial landmarks are detected using the Conditional Local Neural Fields (CLNF) model [4], which is extended using multiple training datasets [5, 10, 16]. These extensions were integrated in the OpenFace framework [3]. The detected facial landmarks are the inner eye corners and the pupils of each eye. We measure the distance between the pupil's center and the inner eye corner for the left and right eyes (D_{left} and D_{right} separately). The ratio of D_{left} to D_{right} is calculated to determine whether the user is looking to the left or to the right (see Figure 2). For example, if the user looks to the left, D_{left} increases while D_{right} decreases, which results in a higher D_{left} to D_{right} ratio. A series of pilot studies with 11 participants revealed that thresholds of 1.05 and 0.9 are appropriate in our setup for detecting left and right gaze gestures. We use a window of 5 frames for gesture detection. For example, we conclude that the user is looking to the left if we receive 5 frames in which $\frac{D_{left}}{D_{right}} \geq 1.05$.

Commercial eye trackers often employ IR-sensors to exploit infrared-induced corneal reflections for improved tracking quality. However, these trackers typically have a range of 60–90 cm; using them for GazeDrone would require users to stand too close to the drone. Hence we opted for video-based eye tracking through the drone's front camera. We expect that the range of IR-based trackers will cover a wider area in the near future. At that point, they can be integrated into GazeDrone, increasing the range of detected eye movements.

USER STUDY

We evaluated GazeDrone for stationary and moving scenarios on an 86" projected display in our lab.

Design

Inspired by prior work [29], we defined three basic gaze manipulation tasks: selection, scrolling, and sliding. In the selection task, participants had to perform a gaze gesture to the left or to the right in response to an arrow shown on the display (Figure 3A). In the scrolling task, participants had to scroll through a set of figures via discrete gaze gestures until the target figure (shown at the top) is at the center of the display for 2 seconds (Figure 3B). In the sliding task, participants had to move a slider towards the center; the task was completed after the slider had stayed for 2 seconds at the center of the display (Figure 3C). In the latter two tasks, participants always started at a state where the target was two steps away from the starting position. In half of the cases the participant had to perform two steps to the right, and in the other half two steps to the left were needed. Previous work reported that users found it challenging to use their peripheral vision to judge if the target was reached [29]. Hence, audio feedback was provided at the recognition of input.

To cover cases where users are moving, we experimented with four user movements: (1) *Stationary*: as a baseline condition, participants stood 3.5 meters in front of the display's center, i.e., position 5 in Figure 4. (2) *Orthogonal*: walking towards the display, i.e., 8 → 2 in Figure 4. (3) *Parallel*: walking parallel to the display, i.e., 4 ↔ 6 in Figure 4. (4) *Diagonal*:


Figure 4. We experimented with different user movement conditions: (a) stationary - position 5, (b) orthogonal movement - 8 → 2, (c) parallel movement - 4 ↔ 6, and (d) diagonal movement - 7 → 3 and 9 → 1.

walking diagonally towards one corner of the display, i.e., 7 → 3 and 9 → 1 in Figure 4.

The study was designed as a repeated measures experiment. Participants performed 4 blocks (stationary, orthogonal, parallel, diagonal), each block covered the three tasks. Each participant performed 4 runs per condition, resulting in a total of 48 trials per participant (4 user movements × 3 tasks × 4 runs). Participants always started with the selection task, since it is the most basic interaction. Scrolling and sliding were performed second and third at an alternating order across participants. For parallel and diagonal movements, participants moved from left to right (4 → 6 in parallel, and 7 → 3 in diagonal) in two of the four runs, while the other two runs were from right two left (6 → 4 in parallel, and 9 → 1 in diagonal). The order of the movement conditions and the starting position were counter balanced using a Latin-square.

Participants and Procedure

We recruited 17 participants (6 females) with ages ranging from 23 to 35 years ($M = 26.29$, $SD = 3.24$). All participants had normal or corrected-to-normal vision. The experimenters started by introducing the study and asking participants to fill in a consent form. According to the Latin square arrangement, participants were told the expected movement, task, and starting position. To exclude possible effects of unstable hovering, the drone was held and moved by an experimenter. We concluded with a questionnaire and a semi-structured interview to collect qualitative feedback.

Limitations

While GazeDrone is capable of tracking the user's gaze while hovering independently, the drone's stability, its speed and its distance to the user influence the perception of GazeDrone. Users are concerned about their safety if a flying drone is not far enough from their face or is not perfectly stable [8]. Hence, in our evaluation of GazeDrone, an experimenter manually carried the drone to overcome the influences of the state-of-the-art technology limitations on user perceptions. Nevertheless, progress in research and industry promises solutions through advancements in camera resolutions, proximity sensors, and processing power of on-drone chips. We expect that in the

Figure 3. Participants performed 3 tasks: (A) Selection: performing a gaze gesture towards the direction shown on the screen. (B) Scrolling: scrolling through the objects until the target (shown on top in blue) is at the center of the screen for 2 seconds. (C) Sliding: moving the slider to the center and keeping it there for two seconds.

near future, a field deployment of GazeDrone will be feasible without safety concerns.

Quantitative Results

We measured the *correct input rate*, which we define as the number of times the system recognized a user's gaze towards the expected direction. An incorrect input could be a result of the system mistakenly detecting a gaze gesture towards the wrong direction (incorrect system detection), or a result of the user mistakenly gazing towards the wrong direction (incorrect user input). We analyzed the data using a repeated measures ANOVA. This was followed by post-hoc Bonferroni-corrected pairwise comparisons.

We found a significant main effect of the user movement type ($F_{3,45} = 5.551$, $p < 0.01$) on correct input rate. Significant differences in correct input rates ($p < 0.01$) were found between stationary ($M = 70\%$, $SD = 33\%$) and parallel movement ($M = 52.2\%$, $SD = 40.3\%$). This means that input was significantly more accurate when stationary compared to when moving parallel to the display. No significant differences were found between any other pair, which means that we could not find any evidence of differences in performance among the other movement conditions. Figure 5 shows that the highest accuracy is achieved when users are stationary. The figure also suggests that accuracy is almost as high when moving orthogonally towards the display, and drops slightly when moving diagonally towards the display. However, a sharp drop is noticed when moving parallel to it.

We attribute the lower accuracy in the moving conditions to motion blur. The low accuracy of the parallel movement condition can be explained by the participants' feedback. Participants reported that interacting while moving towards the display (orthogonally or diagonally) is more natural, compared to when moving parallel to the display; some reported being often confused when they had to move parallel to the display in a direction, while performing gaze gestures to the other directions. This suggests that there are more "incorrect user inputs" in the parallel movement condition.

We also found a significant main effect of the task type ($F_{2,30} = 4.662$, $p < 0.01$) on correct input rate. Pairwise comparisons ($\alpha = 0.05 / 3$ comparisons $= 0.0167$) indicated a significant difference between the selecting task ($M = 69\%$, $SD = 37.5\%$) and the sliding task ($M = 54\%$, $SD = 37\%$). This means that performing selection tasks is easier compared to sliding tasks, which is in-line with previous work [29].

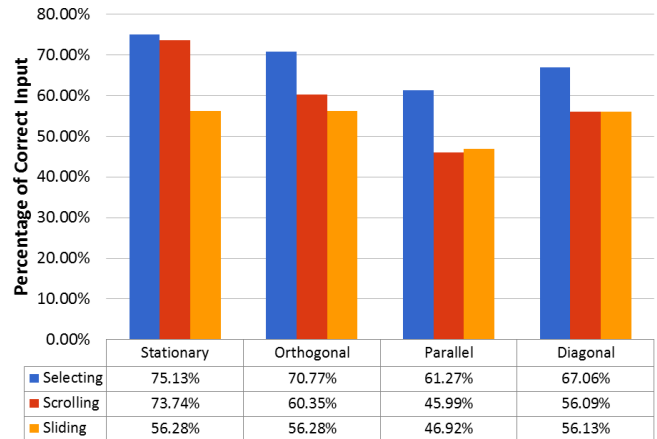|	Stationary	Orthogonal	Parallel	Diagonal
Selecting	75.13%	70.77%	61.27%	67.06%
Scrolling	73.74%	60.35%	45.99%	56.09%
Sliding	56.28%	56.28%	46.92%	56.13%

Figure 5. Performance is highest when selecting while stationary. Performance is high in orthogonal and diagonal movements, but significantly lower in parallel movements. This is attributed to: 1) reduced quality due to motion blur, and 2) incorrect input by participants when moving parallel to the display.

Subjective Feedback

When asked how often the system recognized their gaze gestures accurately on a 5-point scale (5=always correct;1=never correct), feedback from the participants matched the quantitative results. They found interaction very accurate when stationary ($Mdn = 4$, $SD = 1.12$) and when moving orthogonally towards the display ($Mdn = 4$, $SD = 1.05$). While accuracy is moderate when moving diagonally ($Mdn = 3$, $SD = 1.03$), participants perceived accuracy to be lower when moving parallel to the display ($Mdn = 2$, $SD = 0.8$).

In the interviews, 11 participants reported that they would use GazeDrone if deployed in public, while six were concerned about the privacy implications of being tracked by a drone in public. All participants mentioned that they like the flexibility and hands-free interaction enabled by GazeDrone. Four particularly highlighted that they found the system innovative and interesting. On the flip side, four reported feeling uncomfortable when drones are close to their face. Two of which stated they would only use it if it was too far from them and if the drone was small. One participant complained about the noise caused by the propellers of the drone. A participant mentioned she would rather control *when* to provide gaze data by, for example, launching a smartphone app.

In addition to using GazeDrone for pervasive interactive systems, participants reported other use cases in which they would imagine GazeDrone being used. For example, a participant suggested employing GazeDrone near touristic attractions as an audio guide for the objects users are looking at. A partici-

pant proposed utilizing GazeDrone in large warehouses with high-bay areas; GazeDrone can detect which product a worker is trying to reach, another bigger drone could then bring it to the worker. Participants suggested collecting gaze data at situated advertisements for market research, assisting the disabled, and hands-free interaction with distant and mid-air displays in sports, e.g., biking, jogging, or skiing.

DISCUSSION

The results indicate that GazeDrone is well perceived by users. Performance is highest when the user is stationary, and almost as good when the user is moving orthogonally and diagonally towards the display. Performance however drops sharply when moving parallel to the display. These findings are supported by qualitative feedback from participants and quantitative results.

Free vs Enforced Ergonomic Movement

Participants reported that interacting while moving parallel to the display is unnatural and demanding. A possible reason is that users are not used to looking to their sides for a long time while walking. This suggests that some walking patterns are not well-suited for gaze interaction while on the move. While making eye tracking more robust against motion blur is an important direction for future work, systems should also support and guide users to interact in an ergonomic way.

Previous work investigated guiding users to position themselves in a target location, the "sweet spot", from which interaction is optimal [2]. This was done by using visual cues that, for example, gradually brighten the content on the display as the user approaches the sweet spot. Similar approaches can be employed to influence the user's movement towards the interactive system. GazeDrone can be used to support interaction while on the move, but when higher accuracy is required by the application, GazeDrone would then gradually guide users to become stationary or move in a particular pattern that optimizes performance (e.g., towards the display rather than parallel to it). Previous work has shown that behavior of robots can influence the user's proximity [19]. Future work could investigate if the drone's behavior can similarly make the user slow down or move in a certain pattern.

Size of the Drone

Some participants reported that the bigger the drone the less likely they are comfortable interacting with GazeDrone. This suggests that drone-supported gaze interaction should utilize smaller drones. Although not reported by our participants, a further disadvantage of big drones is that they might block the user's view of the display.

Hence, we recommend hovering the drone with an adjustable camera angle at an altitude below the user's height, and to use small drones, for example, the Aerius quad-copter (3 cm × 3 cm × 2 cm [1]).

Privacy Implications

Feedback from six participants is in-line with previous work that showed that users are not comfortable with drones storing data about them [28]. Although we can technically store the gaze data recorded by GazeDrone for offline processing and,

hence, higher accuracy, we opted for real-time processing in our implementation. This means that we do not store data, but rather process them on the fly.

Even when processing in real time, it is still recommended that users are warned before drones collect data about them [26]. Hence, a field deployment would require mechanisms to inform the user that GazeDrone is tracking their eye movements, and allow them to opt out when desired. For example, GazeDrone can use auditory or visual announcements (e.g., LED lights [17]) to communicate that it is in eye tracking mode. Previous work proposed using gestures to signal a "stop" command to drones; this feature can be utilized by users to indicate that they do not wish to be tracked [7]. Similarly, and one participant suggested, the drone could enable gaze interaction on demand only after the user's request.

CONCLUSION

In this work we proposed a novel approach for gaze-based interaction in public pervasive settings. GazeDrone employs drone-supported gaze-based interaction, hence our approach does not require augmenting the user, and does not restrict their movements. We described the implementation of GazeDrone, and reported on a lab study. The results show that GazeDrone is well perceived and can indeed track users' eyes while moving despite motion blur. Performance is highest when the user is stationary. Gaze interaction while moving orthogonally or diagonally towards the display yields high performance, but performance drops when moving parallel to the display. We concluded with four design implications to guide further research in drone-supported eye tracking.

REFERENCES

1. AERIUS. Aerius. Webpage, 2017. Retrieved March 21st, 2017.

2. Alt, F., Bulling, A., Gravanis, G., and Buschek, D. Gravityspot: Guiding users in front of public displays using on-screen visual cues. In *Proc. UIST'15* (2015).

3. Amos, B., Ludwiczuk, B., and Satyanarayanan, M. Openface: A general-purpose face recognition library with mobile applications. Tech. rep., CMU-CS-16-118, CMU School of Computer Science, 2016.

4. Baltrusaitis, T., Robinson, P., and Morency, L.-P. Constrained local neural fields for robust facial landmark detection in the wild. In *Proc. ICCV'13 Workshops* (June 2013).

5. Belhumeur, P. N., Jacobs, D. W., Kriegman, D. J., and Kumar, N. Localizing parts of faces using a consensus of exemplars. *IEEE Trans. Pattern Anal. Mach. Intell 35*, 12 (Dec 2013), 2930–2940.

6. Camera, H. Hover camera. Webpage, 2017. Retrieved March 21st, 2017.

7. Cauchard, J. R., E, J. L., Zhai, K. Y., and Landay, J. A. Drone & me: An exploration into natural human-drone interaction. In *Proc. UbiComp'15* (2015), 361–365.

8. Chang, V., Chundury, P., and Chetty, M. "spiders in the sky": User perceptions of drones, privacy, and security. CHI '17 (2017).

9. Drewes, H., and Schmidt, A. *Interacting with the Computer Using Gaze Gestures*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007, 475–488.

10. Gross, R., Matthews, I., Cohn, J., Kanade, T., and Baker, S. Multi-pie. *Image Vision Comput. 28*, 5 (May 2010), 807–813.

11. Hansen, J. P., Alapetite, A., MacKenzie, I. S., and Møllenbach, E. The use of gaze to control drones. In *Proc. ETRA'14* (2014), 27–34.

12. Hennessey, C., and Fiset, J. Long range eye tracking: Bringing eye tracking into the living room. In *Proc. ETRA'12* (2012), 249–252.

13. Khamis, M., Alt, F., and Bulling, A. Challenges and design space of gaze-enabled public displays. In *Adj. Proc. UbiComp'16* (2016).

14. Khamis, M., Hoesl, A., Klimczak, A., Reiss, M., Alt, F., and Bulling, A. Eyescout: Active eye tracking for position and movement independent gaze interaction with large public displays. In *Proc. UIST'17* (2017), 155–166.

15. Lander, C., Gehring, S., Krüger, A., Boring, S., and Bulling, A. Gazeprojector: Accurate gaze estimation and seamless gaze interaction across multiple displays. In *Proc. UIST'15* (2015), 395–404.

16. Le, V., Brandt, J., Lin, Z., Bourdev, L., and Huang, T. S. *Interactive Facial Feature Localization*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, 679–692.

17. LightCense. Lightcense. Webpage, 2017. Retrieved March 21st, 2017.

18. Majaranta, P., and Bulling, A. *Eye Tracking and Eye-Based Human-Computer Interaction*. HumanComputer Interaction Series. Springer London, 2014, 39–65.

19. Mumm, J., and Mutlu, B. Human-robot proxemics: Physical and psychological distancing in human-robot interaction. In *Proceedings of the 6th International Conference on Human-robot Interaction*, HRI '11 (2011), 331–338.

20. Ohno, T., and Mukawa, N. A free-head, simple calibration, gaze tracking system that enables gaze-based interaction. In *Proc. ETRA'04* (2004), 115–122.

21. Pestana, J., Sanchez-Lopez, J. L., Campoy, P., and Saripalli, S. Vision based gps-denied object tracking and following for unmanned aerial vehicles. In *2013 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)* (Oct 2013), 1–6.

22. Plus, H. Hexo plus. Webpage, 2017. Retrieved March 21st, 2017.

23. Risko, E. F., and Kingstone, A. Eyes wide shut: implied social presence, eye tracking and attention. *Attention, Perception, & Psychophysics 73*, 2 (2011), 291–296.

24. Schneegass, S., Alt, F., Scheible, J., Schmidt, A., and Su, H. Midair displays: Exploring the concept of free-floating public displays. In *Proc. CHI EA '14* (2014).

25. Sugioka, A., Ebisawa, Y., and Ohtani, M. Noncontact video-based eye-gaze detection method allowing large head displacements. In *Engineering in Medicine and Biology Society, 1996. Bridging Disciplines for Biomedicine. Proceedings of the 18th Annual International Conference of the IEEE*, vol. 2 (Oct 1996), 526–528 vol.2.

26. Telecommunications, N., and Administration., I. Voluntary best practices for uas privacy, transparency, and accountability. Tech. rep., National Telecommunications and Information Administration, 2016.

27. Vidal, M., Bulling, A., and Gellersen, H. Pursuits: Spontaneous interaction with displays based on smooth pursuit eye movement and moving targets. In *Proc. UbiComp '13* (2013), 439–448.

28. Yao, Y., Xia, H., Huang, Y., and Wang, Y. Privacy mechanisms for drones: Perceptions of drone controllers and bystanders. In *Proceedings of the 35th Annual ACM Conference on Human Factors in Computing Systems*, CHI '17 (2017).

29. Zhang, Y., Bulling, A., and Gellersen, H. Sideways: A gaze interface for spontaneous interaction with situated displays. In *Proc. CHI'13* (2013), 851–860.

30. Zhang, Y., Chong, M., Müller, J., Bulling, A., and Gellersen, H. Eye tracking for public displays in the wild. *Personal and Ubiquitous Computing 19*, 5-6 (2015), 967–981.