Leveraging the Asymmetric Sensitivity of Eye Contact for Videoconferencing

Milton Chen
Computer Graphics Laboratory and Interactivity Laboratory
Stanford University
miltchen@graphics.stanford.edu

Abstract

Eye contact is a natural and often essential element in the language of visual communication. Unfortunately, perceiving eye contact is difficult in most videoconferencing systems and hence limits their effectiveness. We conducted experiments to determine how accurately people perceive eye contact. We discovered that the sensitivity to eye contact is asymmetric, in that we are an order of magnitude less sensitive to eye contact when people look below our eyes than when they look to the left, right, or above our eyes. Additional experiments support a theory that people are prone to perceive eye contact, that is, we will think that someone is making eye contact with us unless we are certain that the person is not looking into our eyes. These experimental results suggest parameters for the design of videoconferencing systems. As a demonstration, we were able to construct from commodity components a simple dyadic videoconferencing prototype that supports eye contact.

Keywords

Eye contact, gaze perception, videoconferencing

INTRODUCTION

We use our eyes to sense the world and to express ourselves. When two people look into each other's eyes, they experience eye contact. Eye contact is a natural experience of face-to-face communication [2].

A major criticism of video-mediated communication is that most videoconferencing systems do not allow eye contact. The camera is typically mounted above the display; thus, attempts to engage in eye contact are typically perceived as looking down rather than into the remote observer's eyes.

Eye contact can be supported using one of three approaches: (1) warping the video so that it appears to be captured from the remote observer's eyes, (2) merging the camera and display optical path, or (3) mounting the camera close to the display so that they appear to share the same optical path. Computer vision has been used for video warping [8] but it can produce unnatural looking eyes. The camera and the display optical path can be merged either by placing the camera behind a semi-transparent display [10][14] or by placing the camera behind a small hole on a front-projected screen. A disadvantage of the second approach is that commodity displays, such as the ubiquitous desktop monitor, cannot be used. The third approach has been used successfully by [3][15] for a 12-inch diagonal display and by [4] for a 76-inch display; however, it is unknown if this approach can be applied to all displays.

Eye contact may also be supported metaphorically. The GAZE Groupware System allows users to express gaze direction as image orientation in a 3D virtual environment [19]. This approach allows gaze awareness in videoconferencing with many participants; however, since image rotation may not alter the perceived gaze direction of the person in the image, the gaze direction derived from the person in the image may conflict with the gaze direction expressed by the image's orientation.

In the hope of improving the perception of eye contact in videoconferencing, we conducted experiments to determine how accurately people perceive eye contact. Our goal is to provide parameters for the design of videoconferencing systems; specifically, regarding the precision requirements to support eye contact in videoconferencing.

We begin by summarizing the classic findings. Next, we describe our experimental procedure and present our discovery that the sensitivity to eye contact is asymmetric. We are an order of magnitude less sensitive to whether there is eye contact when people look below our eyes than when they look to the left, right, or above our eyes. We conjecture that this asymmetry is due to the anatomical properties of our eyes: it is harder to tell whether people are looking at you when they are looking down. After presenting our results, we propose the theory that we are prone to perceive eye contact: we will think that there is eye contact unless we are certain that the person is not looking into our eyes. Lastly, we suggest design

CHI 2002, April 20-25, 2002, Minneapolis, Minnesota, USA.

parameters for videoconferencing systems, and as a demonstration, we describe a simple dyadic videoconferencing prototype constructed from commodity components. This prototype allows eye contact for the majority of our subjects.

Throughout this paper, we will use the terms adopted by the early gaze researchers: a "looker" is defined as the person sending out the gaze and an "observer" is defined as the person judging the gaze.

PREVIOUS WORK

A common belief is that we can precisely judge the direction of another person's gaze. The exact precision was measured by psychologists who wanted to understand visual communication and by those designers of videoconferencing systems who wanted to support eye contact.

Perceiving Eye Contact

Gibson and Pick performed the first study on the perception of gaze direction [9]. They instructed a looker to assume a passive facial expression and to fixate on seven points on a horizontal line while facing an observer at a distance of 2m. The gaze targets were 10 cm apart, the middle target being the bridge of the observer's nose. For each fixation, the observer judged whether the looker was looking directly at him or not. There were six observers, and they perceived 84 percent of fixations at the bridge of the nose as the looker looking directly at them. More importantly, the standard deviation of the responses over the seven targets corresponded to an angular deviation of 2.8°, and Gibson and Pick defined this standard deviation as the just noticeable deviation of the looker's gaze from the bridge of the observer's nose. A 2.8° rotation of the eyeballs roughly corresponds to 1 mm of linear displacement of the looker's iris. From 2 m, 1 mm corresponds to 1 minute of arc. Since human Snellen visual acuity is typically said to be 1 minute of arc, Gibson and Pick concluded that the acuity of perceiving eye contact is as good as the Snellen visual acuity.

In contrast to Gibson and Pick, who examined the perception of a looker who looked to the left and right of the observer, Cline used a half-silvered mirror to allow his looker to fixate on targets to the left, right, upward, and downward of the bridge of an observer's nose [5]. The gaze targets were 2°, 8°, and 12° in each direction. The looker assumed a passive facial expression and sat 122 cm from the observer. Both the looker's and the observer's heads were held in place with headrests. For each fixation, the observer marked the looker's gaze direction on a transparent response board. There were five observers and the fixations at the bridge of the observer's nose had a standard deviation of 0.75° horizontally and 1.25° vertically; from this, Cline reaffirmed Gibson and Pick's conclusion that the acuity of eye contact is as good as the Snellen visual acuity. When the looker looked below the observer's eyes by 8° and 12°, the perceived directions were on average 1.6° and 3.7° below the gaze targets, respectively.

Gibson and Pick's as well as Cline's conclusion that a gaze directed at the bridge of an observer's nose can be perceived with an acuity matching the Snellen visual acuity was further affirmed by Jaspars et al. [11]. They found that observers could discriminate between gaze shifts of 0.6°.

The studies described so far all used gaze targets separated by large visual angles. Their claim that a gaze deviation of roughly one degree is accurately detected can be tested directly if the gaze targets are more closely spaced. Kruger and Huckstedt performed one such experiment [13]. Their looker fixated on seven points around the observer's eyes: forehead, bridge of the nose, tip of the nose, right and left eye, and right and left face edge. The observers were able to correctly identify the location of the feature points 35 and 10 percent of the time from a distance of 80 and 200 cm, respectively. Ellgring repeated the Kruger and Huckstedt experiment with a homogeneous group of schoolgirls and obtained a higher percentage of correct judgments [6]. However, even the most accurate judgments, fixations at the eyes, were still short of 50 percent accuracy. From 80 cm, the gaze targets were about 1.7° apart. If the acuity of gaze perception matched the Snellen visual acuity, we would expect a higher percentage of correct responses. Perhaps the observers were able to precisely see the iris positions but were unable to precisely judge the gaze direction from the iris positions.

Researchers also found two systematic errors in the perception of gaze. First, if the looker's head is rotated away from the observer, the observer tends to underestimate the angle of this rotation [1][5][9]. For example, if the looker aims his head toward the observer's left, more eye contact is perceived when the looker looks toward the observer's left than when he actually looks between the observer's eyes. Second, at greater distance, observers tend to overestimate eye contact [12][17]. Lastly, Ellgring and Cranach showed that the accuracy of gaze perception for gaze targets around the face could be improved with practice; however, the perception of gaze aimed at the bridge of the observer's nose did not improve [7].

Perceiving Eye Contact in a Videoconference

Bell Laboratories performed the first study on perceiving eye contact in a videoconference during the design of the Mod II PicturePhone [18]. They found that the threshold of losing eye contact is 4.5º for looks to the side of the camera and 5.5º for looks above or below the camera. Unfortunately only the results of their study are known; it is unclear whether the decrease in the sensitivity of perceiving eye contact from around 1º as found by [5][9][11] in the face-to-face condition to around 5º in their video condition is caused by the video medium. The actual visual angle between the PicturePhone camera and the expected eyes on the display is 5.8°. The PicturePhone team also found that people like to view the other party's eyes 40% down from the top edge of the display; thus, the camera should be placed above the display.

The claim that the camera should be placed above the display was challenged by Stapley [16]. Stapley mounted a line of miniature light bulbs on a camera at a spacing of 2.5 cm. The looker, while 1 m from the camera, was instructed to look into the camera or the lighted bulb. An observer judged eye contact while viewing the looker on a monitor from a distance of 1 m. Stapley found that the camera should be placed 1.4° to the right and 1.4° below the display. However, he reassigned the looker to be the observer in each experiment, which meant that the observer knew the expected percentage of eye contact. White has shown that eye contact judgments can be shaped by the experimenter's bias [20].

In contrast to the PicturePhone team and Stapley, who asked the observers to judge whether they felt eye contact, Anstis et al. asked their observers to judge where the looker was looking [1]. They found little difference between the face-to-face medium and the video medium. In both, the observers' eye contact sensitivity was high. They also found no significant asymmetry in acuity regarding the different gaze directions.

The common belief that we can precisely judge the direction of another person's gaze is generally confirmed by the findings of classic gaze experiments; however, the exact precision can be further refined: one degree [1][5][9][11][16] vs. a few degrees [6][13][18]. The classic findings also suggest that the sensitivity to eye contact is roughly symmetric in that there is no one direction that is significantly less sensitive than the other directions [1][5][6][11][13][16][18].

METHODOLOGY

The classic gaze experiments were conducted before video recordings were practical. Each research team used a different looker and since the influence of the looker's eye appearance was unknown, comparing results obtained by different researchers was difficult. To create a controlled dataset for studies in gaze perception, we built a recording studio. The studio consists of a gaze recording room and a gaze measuring room.

Gaze Recording and Measuring Studio

Figure 1 shows a picture of the gaze recording room. The room has a 2.4 m by 1.8 m front-projected display driven by a high-end computer. A 5 cm by 5 cm hole is cut in the middle of the display and a professional-quality video camera looks through this hole from behind the display.

The looker sits 2.4 m from the display with the seat adjusted so that the line from her eyes to the camera is perpendicular to the plane of the display. From this distance, a 10 cm forward or backward movement of the looker's head will cause a shift in visual angle of 0.04° for gaze targets next to the camera and 0.6° for gaze targets 15° away from the camera. The large size of the display allows us to maintain a high precision for gaze recording without using a headrest. The gaze measuring room consists of a 1.5 m by 1.1 m front-projected display driven by a high-end computer. The observer sits 2.4 m from the display.

We conducted four experiments in our studio. The first experiment examined the observer's directional sensitivity to eye contact. The second experiment used more lookers to examine the effect of eye appearance. The third experiment examined the systematic error between perceiving gaze in a recorded video and in an actual videoconference. The last experiment examined the effect of compression artifacts and camera resolution on eye contact.

Experiment 1: Sensitivity to Gaze Direction

In this experiment, we recorded a male looker with medium-sized dark brown eyes and wearing contact lenses in the gaze recording room. A gaze target was shown on the display and the looker was instructed to examine the gaze target. When the looker pressed a key, the computer began to record a head-and-shoulder video of him. The recording stopped after 3 seconds and the gaze target was shown at a new location to begin a new recording cycle. A studio-quality videoconferencing light illuminated the looker. The videos were recorded at 640x480 pixels per frame, 15 frames per second, and compressed using MPEG-4. The videos were of sufficient quality for the observers to see the eyes of the looker clearly.

The gaze target was a 5 cm by 5 cm image of a person's face that was chosen to minimize gaze fatigue. We used a small cross as the gaze target at first; however, lookers had difficulty keeping their gaze focused for the duration of the recording and the looker's eyes sometimes diverged during the forced fixation.

The dots on the display in Figure 1 indicate the gaze target locations. The gaze targets radiate in eight directions from the camera at an incremental step of 1° of visual angle. The downward direction covers a range of 15°, and the other seven directions cover a range of 5°. Fifty-one videos were recorded for this looker.

After the recording, we showed the videos to an observer in the gaze measuring room. For each video, the observer was asked if the looker in the video was looking directly into the observer's eyes. Each video was looped until the observer responded. The videos were shown in random order until each video had been shown three times. Sixteen observers with 20/20 vision after correction participated in this experiment. The observers were chosen from current students and recent graduates of Stanford University.

Experiment 2: Sensitivity to Eye Appearance

In this experiment, we recorded a male looker with light blue eyes and wearing contact lenses, a female looker with

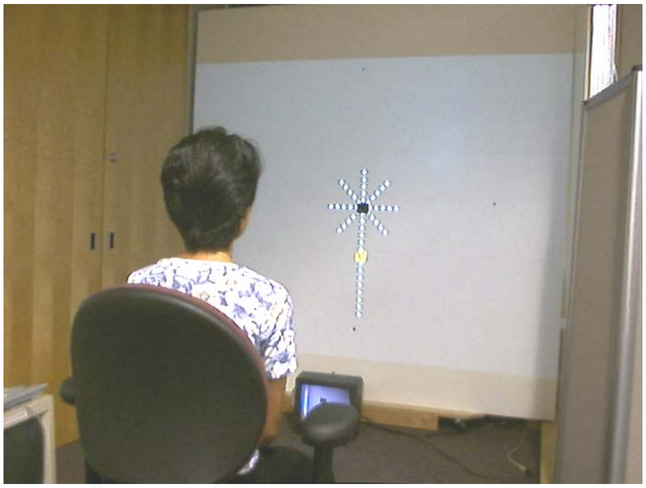

Figure 1. The gaze recording room. The 2.4 by 1.8 m front-projected display has a small hole in the middle that allows a camera to be placed behind the display. The small picture on the display is the gaze target. Radiating in eight directions at an incremental step of 1º of visual angle, the dots on the display indicate the locations at which the gaze target can appear.

dark brown eyes and wearing contact lenses, and a male looker with dark brown eyes and wearing glasses under the same condition as in Experiment 1. All lookers have medium-sized eyes.

We showed the videos of these three lookers to the sixteen observers in the first experiment. The experimental procedure was the same as in the first experiment except only videos where the lookers were looking below the camera were shown.

Experiment 3: Error Due to Recording

In this experiment, we linked the gaze recording and measuring room with live audio and video. The gaze target was replaced by a live video of the observer, and the observer saw a live video of the looker. The transmitted videos were of the same quality as those in the first two experiments.

During the experiment, the looker and the observer engaged in casual conversation. At random times, the looker would ask the observer if she thought that the looker was looking into her eyes. As in the second experiment, only the gaze target at the camera and the fifteen targets below the camera were used. The targets were displayed in random order. The looker from the first experiment and the sixteen observers from the previous two experiments participated in this study.

Experiment 4: Influence of Video Quality

In this last experiment, we repeated the third experiment with uncompressed video and face-to-face conditions. For the uncompressed video case, everything was identical to the third experiment except that analog uncompressed video instead of MPEG-4 compressed video was used in linking the rooms.

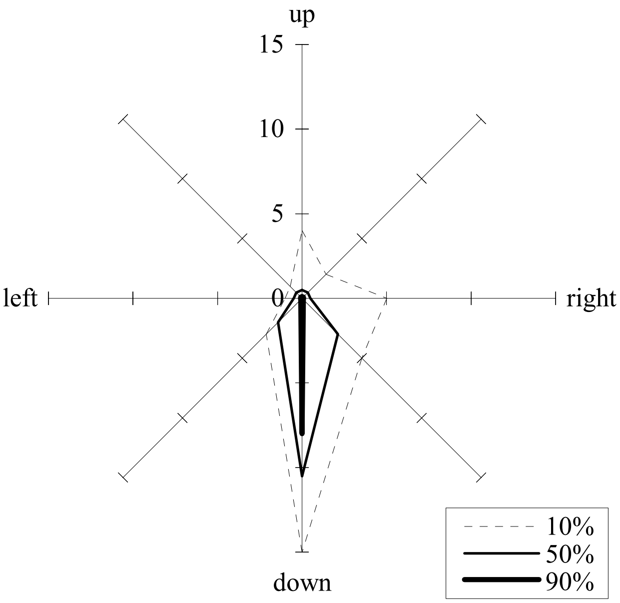

Figure 2. Sensitivity to gaze direction, experiment 1. The contour curves mark how far away in degrees of visual angle the looker could look above, below, to the left, and to the right of the camera without losing eye contact. The three curves indicate where eye contact was maintained more than 10%, 50%, and 90% of the time. The percentiles are the average of sixteen observers. The camera is at the graph origin.

For the face-to-face case, the looker sat 1 m from the observer. While looking at the observer's eyes, nose, mouth, chin, neck, or chest, the looker engaged in casual conversation with the observer. At random times, the looker would ask the observer about eye contact. At the end of this experiment, the distances between the feature points on the observer and the observer's eyes were measured. The looker and the observers from the first experiment participated in this experiment. The order in which the observers participated in each of the four experiments was randomized.

RESULTS

Figure 2 shows the result of the first experiment, the sensitivity of eye contact with respect to the direction in which gaze deviates from the camera. Notice that the observers were very sensitive when the looker looks up, left, or to the right, but less sensitive when the looker looks below the camera. For the up, left and right cases, the looker can look at most 1º away from the camera before perception of eye contact is lost. However, for the down case, observers were much less sensitive to eye contact.

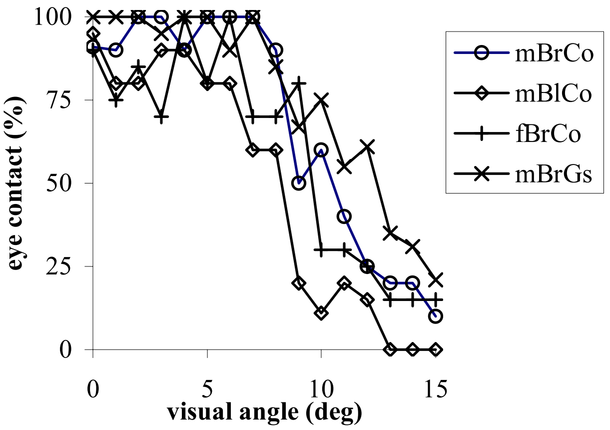

Figure 3. Sensitivity to eye appearance, experiment 2. The curves show the percentage of times that eye contact was perceived for four lookers looking in the down direction in Figure 2. The horizontal axis marks the visual angle in degrees that the looker looked below the camera. The four lookers were a male with dark brown eyes wearing contact lenses (mBrCo), a male with light blue eyes wearing contact lenses (mBlCo), a female with dark brown eyes wearing contact lenses (fBrCo), and a male with dark brown eyes wearing glasses (mBrGs). Each curve is the average of sixteen observers. The average standard deviation is roughly 30% for each looker.

Figure 3 shows the results of the second experiment, the sensitivity to the appearance of the eyes. Notice that for all lookers, the knees of the curves are roughly around 10º. One explanation for the lack of any significant difference between the blue-eyed looker, where the pupil is distinct from the iris, and the brown-eyed lookers, where the pupil is not clearly delineated from the iris, is that the pupil is always centered within the iris, thus our observers did not need to see the pupil to judge gaze direction.

Figure 4 shows the results of the third experiment, the difference between perceiving eye contact in a recorded video and in an actual videoconference. Notice that when the looker was seen in videoconferences, the observers were more likely to perceive eye contact. This effect is especially pronounced around the critical angle where eye contact is lost. One explanation for this phenomenon is that when the observers are not sure whether the looker is looking at them, they will believe there is eye contact if they are engaged in a conversation with the looker since people typically look into each other's eyes during face-to-face conversation.

Figure 5 shows the results of the fourth experiment, the influence of video quality. Notice that high quality compression seems to achieve roughly the same results as uncompressed video; however, the observers were more sensitive in the face-to-face medium. The difference

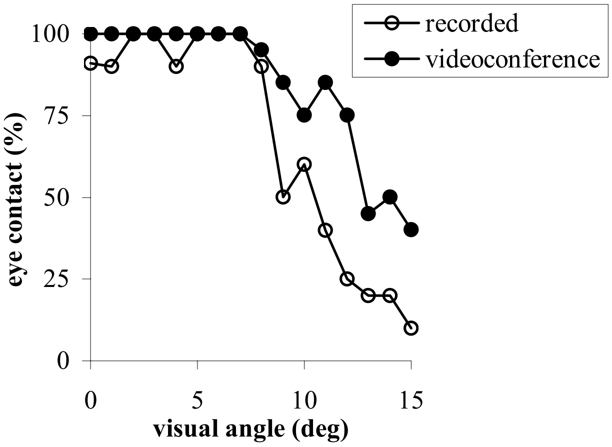

Figure 4. Error due to recording, experiment 3. The curves show the percentage of times that eye contact was perceived when the looker was recorded in advance or was live through videoconferencing. The horizontal axis marks the visual angle in degrees that the looker looked below the camera. Each curve is the average of sixteen observers. The average standard deviations are 31% for recorded and 17% for videoconferencing.

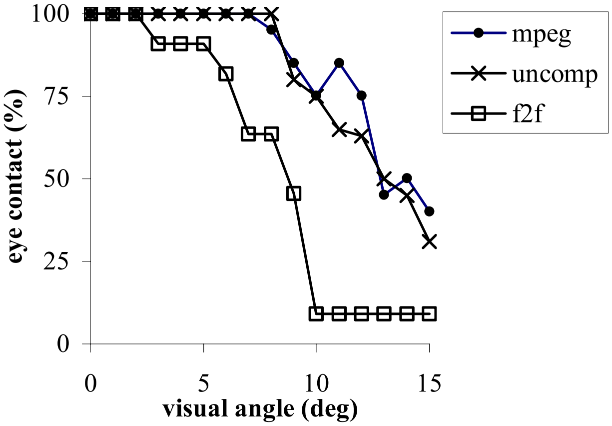

Figure 5. Influence of video quality, experiment 4. The curves show the percentage of times that eye contact was perceived when the looker and observer were in visual conference. The three conditions are videoconferencing with MPEG-4 compressed video, videoconference with uncompressed analog video, and face-to-face conference. The horizontal axis marks the visual angle in degrees that the looker looked below the camera or the eyes of the observer. Each curve is the average of sixteen observers. The average standard deviation is 22% for the face-to-face conference.

between the face-to-face and videoconference conditions could be due to viewing distance, 1 m for face-to-face and 2.4 m for videoconference. When the observer is far from the looker, the observer tends to overestimate eye contact [12][17]; however, we have scaled the video size to match the viewing distance. Another possible explanation is that the limited resolution of the camera and the video capture board limited the observer's sensitivity to eye contact.

THE NATURE OF EYE CONTACT

The claim that our sensitivity in perceiving eye contact is lower when a looker's eyes are looking downward than in other directions may be explained by the characteristics of our anatomy. When a looker looks to the left or right of the camera, his eyeballs rotate within the eye socket, which causes a noticeable change in the position of the iris within the sclera, the whites of the eyes. When the looker looks above the camera, the rotations of his eyes again causes a noticeable change in the position of iris within the sclera: his upper eyelids track the iris position while his lower eyelids remain stationary. When the looker looks below the camera, both his upper and lower eyelids track the iris position, thus there is not a very noticeable change in the position of the iris with respect to the sclera. We have observed this characteristic of anatomy in our lookers. This characteristic was also noticed by [16].

An intriguing observation in both the Gibson and Pick experiment and in our experiment 2 is that even when the looker looks directly into the observer's eyes or the camera, the observers do not perceive eye contact 100 percent of the time. The reported eye contact is 84 percent in Gibson and Pick and roughly 90 percent in our study. We viewed the videos of our lookers frame-by-frame and found that for a significant number of frames, the looker's eyes do not appear to be optically balanced, that is, the eyes do not appear to focus at the same point in space. This is rather surprising since all of our lookers appeared to have optically balanced eyes during an initial face-to-face eye inspection.

We examined the pictures of people in popular magazines who at first glance were looking at us. Much to our surprise, a number of them did not appear to have optically balanced eyes upon close inspection. A good way to see this is to cover up one of the eyes in the picture to judge where it is pointing, repeat the procedure for the other eye, and finally judge both eyes together. After a minute or so of repeating this cycle, we sometimes perceive that the eyes do not converge.

It is possible that our lookers' eyes are optically balanced; however, when they are forced into the unnatural task of staring at a fixed target, their eyes become diverged. We hypothesize that the resulting slight optical imbalance of a looker's eyes is the reason why looking into the camera does not always result in eye contact. We further hypothesize that judging gaze direction is a time-consuming effort and the to-be-judged eyes tend to be constantly in motion. Thus, we typically are unable to see the optical imbalance.

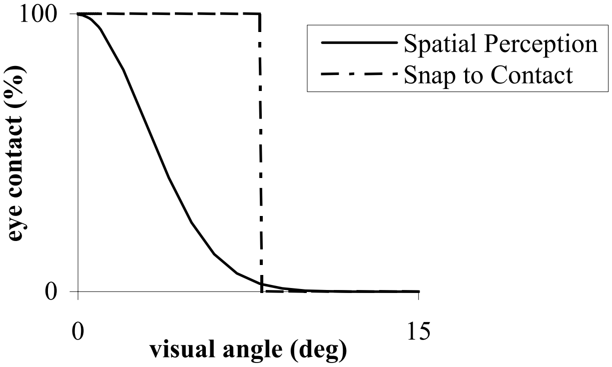

Figure 6. The Snap to Contact theory. The Spatial Perception curve illustrates the classic model for eye contact: the percentage of perceived eye contact can be approximated by a normal distribution. The Snap to Contact curve illustrates the theory that observers are prone to perceive eye contact. The critical angle at which eye contact is lost is influenced by the observer's expectations and viewing conditions.

The Snap to Contact Theory

The perception of eye contact and the more general task of judging gaze direction have often been framed as a spatial perception task. The spatial perception model, as described in the influential work of Gibson and Pick [9], states that an observer estimates a looker's head orientation and eye position within the face; together, this allows the determination of an absolute gaze direction. This model implies that the percentage of perceived eye contact can be approximated by a normal distribution, where the standard deviation can be used to indicate the just noticeable deviation of the looker's gaze from the observer's eyes. Gibson and Pick's data do support the spatial perception model: their data roughly followed a bell-shaped curve; however, their study only measured a looker who looked to the left and right of an observer's eyes.

To explain our findings when the looker looks below the camera, when the looker engages in conversation with the observer, and when the viewing condition is changed from videoconference to face-to-face, we extend the spatial perception model to account for the observer's expectation. Figure 6 illustrates this idea and we call it the Snap to Contact theory. The theory assumes that people cannot always judge gaze direction accurately and they will bias their perception toward contact unless they are certain that the looker is not looking at them. If the looker looks below the camera, the resulting change in appearance is less pronounced than if the looker were looking in other directions, thus more eye contact will be perceived in the

Figure 7. Desktop videoconference prototype. We constructed a mechanical camera holder so that an inexpensive desktop conferencing system can support eye contact.

down direction, as shown in Figure 2. If the observer is conversing with the looker, the looker is expected to engage in eye contact, thus more eye contact will be perceived, as shown in Figure 4. If the viewing condition changes from face-to-face to videoconference, the limited resolution of the conference system will make judging gaze direction more difficult, thus more eye contact will be perceived, as shown in Figure 5.

THE REQUIREMENT FOR EYE CONTACT

Our experimental results suggest the precision requirements for camera positioning in a videoconferencing system and simple improvements. Because our sensitivity in the downward direction is lower than in other directions, the camera should be placed above the display to support eye contact. Figure 3 suggests that a conservative solution is to make the visual angle between the camera and the eyes rendered on the display less than 5º.

For a hand-held device such as a PDA or cell phone, assuming a 1 foot viewing distance, 5º translates to a maximum distance of 1 inch between the camera and the rendered eyes. For a desktop monitor-based conferencing system, and assuming a 3 foot viewing distance, 5º translates to a maximum distance of 3 inches between the camera and the rendered eyes. For an 8-foot wall size display, and assuming an 8 foot viewing distance, 5º translates to a maximum distance of 8 inches between the camera and the rendered eyes. These suggested design parameters appear to be achievable using commodity parts. As a demonstration, we pieced together commodity components to meet the just described design parameters in a desktop conferencing system.

Figure 7 shows our prototype. A 640 by 480 video window is centered along the top edge of a 20” monitor. The size of the window is 10 by 7 inches. The camera, a LogiTech Pro 3000, is mounted so that the centerline of the lens is 1” from the top edge of the monitor screen. Assuming the viewer’s head is 36” from the display, the vertical error in gaze is less than 5º, below the experimentally verified threshold.

We asked eight subjects to converse with the looker from experiment 1 using our prototype. These subjects were different from the sixteen observers in our main experiments. All eight of the subjects perceived eye contact; however, the looker did not perceive eye contact in the case of one of the subjects: he appeared to be looking down. This subject’s upper right eyelid droops a little, making him appear to look slightly downward even face to face. While our prototype seems to demonstrate that only a simple modification of current conferencing systems would allow most people to perceive eye contact, many more lookers should be tested to validate this claim.

The experiments in this paper were designed to minimize subject fatigue; consequently, the number of sample points is limited. We plan to measure a significantly larger number of subjects in order to generalize and expose the shortcomings of our results.

There are many questions that still need to be answered. For example, can eye contact perception be improved with practice? Unlike in other areas of nonverbal communication, there is no database of gaze videos for the research community. We have collected over 1000 high-quality videos of over a dozen lookers and we are releasing this database in the hope of accelerating progress in gaze research [21].

ACKNOWLEDGMENTS

This project is sponsored by the Stanford Interactive Workspace Project, the Stanford Immersive Television Project, the Stanford Learning Lab, and a Department of Defense Graduate Fellowship. I wish to thank Dan Nelson and Brian Luehrs for helping me to construct the gaze recording studio, Jingli Wang for designing and machining the camera holder, and James Davis, Brad Johanson, Yung-Hsiang Lu, Albert Huntington, Linda Wang, Danny Yang, Cindy Chang, Lisa Zhong, Lock Kwan, Terry Winograd, Pat Hanrahan, and the CHI reviewers for critiquing the draft.

REFERENCES

[1] S. Anstis, J. Mayhew, and T. Morley. The Perception of where a Face or Television ‘Portrait’ is Looking. *American Journal of Psychology*, pages 474-489, 1969.

[2] M. Argyle and M. Cook.
Gaze and Mutual Gaze.
Cambridge University Press, 1976.

[3] W. Buxton, A. Sellen, and M. Sheasby.
Interfaces for Multiparty Videoconferences.
Video-Mediated Communication (edited by K. Finn, A. Sellen, and S. Wilbur), Lawrence Erlbaum Associates, pages 385-400, 1997.

[4] M. Chen.
The Design of a Virtual Auditorium.
ACM Multimedia, 2001.

[5] M. Cline.
The Perception of Where a Person is Looking.
American Journal of Psychology, pages 41-50, 1967.

[6] J. Ellgring.
Die Beurteilung des Blickes auf Punkte innerhalb des Gesichtes. *Zeitschrift fur experimentelle und angewandte pshchologie*, pages 600-607, 1970.

[7] J. Ellgring and M. von Cranach.
Processes of learning in the recognition of eye-signals.
European Journal of Social Psychology, pages 33-43, 1972.

[8] J. Gemmell, C. Zitnick, T. Kang, K. Toyama, and S. Seitz. Gaze-awareness for Videoconferencing: A Software Approach.
IEEE Multimedia, Vol. 7, No. 4, pages 26-35, 2000.

[9] J. Gibson and A. Pick.
Perception of Another Person's Looking Behavior.
American Journal of Psychology, pages 386-394, 1963.

[10] H. Ishii and M. Kobayashi.
ClearBoard: a Seamless Medium for Shared Drawing and Conversation with Eye Contact.
Proceedings of CHI, pages 525-532, 1992.

[11] J. Jaspers, et al.
Het observen van ogencontact.
Nederlands Tijdschrift voor de Psychologie, 28, pages 67-81.

[12] D. Knight, D. Langmeter, and D. Landgren.
Eye-contact, Distance, and Affliliation: the Role of Observer Bias. *Sociometry*, pages 390-401, 1973.

[13] K. Kruger and B. Huckstedt.
Die Beurteilung von Blickrichtungen.
Zeitschrift fur experimentelle und angewandte psychologie, pages 452-472, 1969.

[14] K. Okada, F. Maeda, Y. Ichikawaa, and Y. Matsushita.
Multiparty Videoconferencing at Virtual Social Distance: MAJIC Design.
Proceedings of CSCW, pages 385-393, 1994.

[15] A. Sellen.
Remote Conversations: The Effects of Mediating Talk with Technology.
Human-Computer Interaction, pages 401-444, 1995.

[16] B. Stapley.
Visual Enhancement of Telephone Conversations.
Ph.D. Thesis, University of London, 1972.

[17] G. Stephenson and D. Rutter.
Eye-contact, Distance and Affiliation: a Re-evaluation.
British Journal of Psychology, pages 385-393, 1970.

[18] R. Stokes.
Human Factors and Appearance Design Considerations of the Mod II PicturePhone Station Set.
IEEE Transactions on Communication Technology, pages 318-323, 1969.

[19] R. Vertegaal.
The GAZE Groupware System: Mediating Joint Attention in Multiparty Communication and Collaboration.
Proceedings of CHI, pages 294-301, 1999.

[20] H. White, J. Hegarty, and N. Beasley.
Eye Contact and Observer Bias: a Research Note.
British Journal of Psychology, pages 271-273, 1970.

[21] http://graphics.stanford.edu/~miltchen/gaze/