

Perception of Gaze Direction for Situated Interaction

Samer Al Moubayed

Department of Speech Music and Hearing
KTH, Stockholm, Sweden
sameram@kth.se

Gabriel Skantze

Department of Speech Music and Hearing
KTH, Stockholm, Sweden
gabriel@speech.kth.se

ABSTRACT

Accurate human perception of robots' gaze direction is crucial for the design of a natural and fluent situated multimodal face-to-face interaction between humans and machines. In this paper, we present an experiment targeted at quantifying the effects of different gaze cues synthesized using the Furhat back-projected robot head, on the accuracy of perceived spatial direction of gaze by humans using 18 test subjects. The study first quantifies the accuracy of the perceived gaze direction in a human-human setup, and compares that to the use of synthesized gaze movements in different conditions: viewing the robot eyes frontal or at a 45 degrees angle side view. We also study the effect of 3D gaze by controlling both eyes to indicate the depth of the focal point (vergence), the use of gaze or head pose, and the use of static or dynamic eyelids. The findings of the study are highly relevant to the design and control of robots and animated agents in situated face-to-face interaction.

Categories and Subject Descriptors

H.1.2 [Models and Principles]: User/Machines Systems;
I.3.6 [Computer Graphics]: Methodology and Techniques

Keywords

Situated Interaction, Gaze Perception, Head Pose, Eyelids, Talking Head, ECA, Furhat, Robot Head

1. INTRODUCTION

Gaze is a very important component in situated interaction between humans. Gaze direction and dynamics have been found to serve several different functions, including turn-taking control, deictic reference, and attitudes [1]. In a multi-party or situated dialogue, gaze may be an essential means to regulate dialogue flow, address a person in a crowd, or to point to a specific object out of many [2].

Compared to other primates, the color of the sclera ("the white of the eye") in human beings is clearly distinct from the iris and the skin around the eyes, due to the lack of pigment. Relative to body size, human eyes are also larger. This contrast in human eyes is useful for determining the direction of other people's gaze and thereby their focus of attention. Experiments have shown that apes tend to pay more attention to head pose, while infants pay more attention to eye gaze when following the attention of an adult human [3]. It has been hypothesized that this unique feature of humans' eyes has evolved due to the fact that humans to a larger extent are dependent on "reading" other people's minds and follow

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Gaze-In'12, October 26, 2012, Santa Monica, California, USA.
Copyright 2012 ACM 978-1-4503-1516-6 ...\$15.00.

their gaze while communicating or solving joint tasks [3].

Recently, a lot of research has been done on the role of gaze in human-robot interaction [4] and embodied conversational agents [5], in order to exploit the functions of gaze in multimodal human-machine interaction models in different interaction scenarios.

Several studies have explored situated human-robot interaction, where the interlocutors sit around a table with objects that can be referred to, thus constituting a shared space of attention [6,7,8]. These studies provide significant and important findings on the advantageous role of gaze on interaction. However, since they focus on more high level effects of gaze on the overall interaction, they do not provide enough detail and resolution to study and quantify the precision and properties of the perception by humans of the gaze direction generated by robots.



Figure 1. A snapshot of the setup with Furhat looking at the grid and a close-up on Furhat's face details.

In this paper, we explore how accurately humans can perceive the target point in space of the gaze of a robot head in a setting that is typical of situated interaction, and how the rotation of the eyes can be calibrated to better conform to the characteristics of this perception.

We have investigated several factors that may influence the perceived gaze target:

1. The use of only head pose vs. only gaze when shifting between targets.
2. The relative position of the perceiver: frontal view, or from a 45 degrees angle side view.
3. The use of static or dynamic eyelids (i.e., whether the eye lids follow the vertical movement of the eyes).
4. Convergence of the eyes (horizontally) to convey depth.
5. The perception of human gaze in the same setting, for comparison.

2. METHODOLOGY

2.1 Furhat: a Back-projected Robot Head

The use of facial animation for conversational agents, displayed on a flat screen, has been investigated over many years. However, when it comes to situated, multi-party interaction, the use of a flat

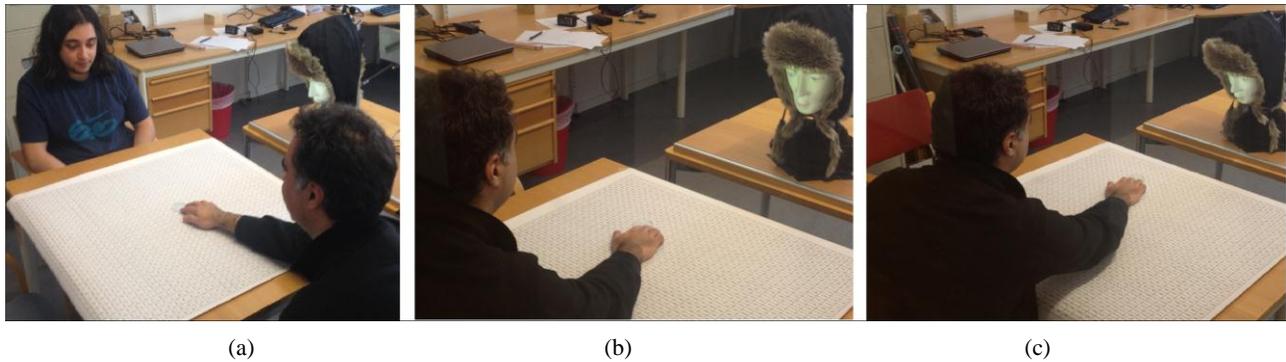


Figure 2. Snapshots of (a): the human condition. (b) Frontal viewing of Furhat moving only the eyes. (c) Frontal viewing of Furhat moving the head instead of the eyes.

screen to communicate with an animated agent suffers from what is known as the Mona Lisa effect [9]: Since the agent is not spatially co-present with the user, it is impossible to establish exclusive mutual gaze with one of the observers – either all observers will perceive the agent as looking at them, or no one will.

While mechanical robot heads are indeed spatially co-present with the user, they are expensive to build, inflexible and potentially noisy. Furhat, shown in Figure 1, can be regarded as a middle-ground between a mechanical robot head and animated agents [10]. Using a micro projector, the facial animation is projected on a three-dimensional mask that is a three dimensional printout of the same head model used in the animation software. The head is then mounted on a neck (a pan-tilt unit), which allows the use of both head pose and gaze to direct attention. It has previously been shown in an experimental setting that such a 3D projection increases the system’s ability to regulate turn-taking in multi-party dialog, as compared to the same animated agent shown in a 2D screen [11].

We use Furhat in this experiment to generate gaze stimuli in a situated setup. Furhat has a three dimensional model of the both eyes with a pupil, iris and eyelids that can be controlled separately.

2.2 Situated Interaction Setup

We employ in this experiment a configuration that to some degree simulates a situated human-robot task solving setup. In the setup, the human and Furhat sit around a squared table that is assumed to be the center of attention of Furhat. In theory, the robot is hypothesized to be looking at areas of interest on the table, such as physical objects. A relatively similar setup, to test the accuracy of gaze direction between different shaped-displays, using a virtual grid, on the vertical dimension is deployed in [12].

In order to be able to make high resolution quantifications of the perception of gaze targets, the table was covered by a squared grid (partly shown in Figure 1 and 2). The grid was divided into 35 rows and 35 columns, resulting in 1225 equally sized squares, each with a side of 2.2 cm.

The experiment is designed so that Furhat and the human will sit at two different sides of the table, and the rotation of the head or the eyes of Furhat is controlled to look at different parts of the grid. The subjects are instructed to place a small *tagging object* (a glass disk, shown in Figure 1) on the square where they best believe the robot is gazing at. The object is used as a guide to help the subject determine where the eyes are looking, substituting the need for physical objects on the table. In order to cover the differ-

ent areas of the grid, without generating a very large number of gaze stimuli, the grid was divided into 9 equally sized virtual regions (split into 3 rows and 3 columns), and Furhat would choose a random square from each of the regions to look at. Thus, for each experimental condition, and for each subject, Furhat generated 9 gaze stimuli.

In order to automate the experiment, and to increase the level of interaction, Furhat prompted each answer by moving the gaze to a new target and by saying something like: “where am I looking now?”. The subject would then place the tagging object on the perceived target square and verbally dictate the number of the square to the experiment instructor who in turn entered the number into the system, which then triggered the next stimuli.

Although intuitive, it is important to note that in such a setup, it is expected that it is more difficult to precisely detect the target square on the grid the further it is from the eyes of the robot, because the same rotation in degrees of the eyes would cover a larger area if the robot is looking towards the further edge of the grid compared to looking towards the closer edge.

2.3 Gaze Generation

A geometric computer model of the setup was implemented, taking into account the design of the head, the facial parts, the mask, the neck, and the physical dimensions and placements of the table and the grid. The model was used to estimate the required horizontal and vertical rotations of the eyes and/or the neck of Furhat, to gaze at any given square on the grid, using a simple vector based geometrical calculation.

Although this model is estimated as accurately as possible, and has taken into account as many details of the setup as possible, it should be noted that this is only an estimate of the exact angle between the eyes of Furhat and the focal point. However, since this model is used for all the experimental conditions, it is not the exact location of the answers of the subjects that is under interest, but rather the differences in the perceived direction of gaze across the conditions.

3. EXPERIMENTAL CONDITIONS

We here describe the different conditions that were investigated. Figure 3 shows snapshots of some of the conditions taken from the same animated agent used in Furhat.

3.1 Human vs. Furhat

In this condition, we explore the gaze perception accuracy when observing a human. This condition is important as a baseline of the perceptual granularity of gaze direction in a human-human

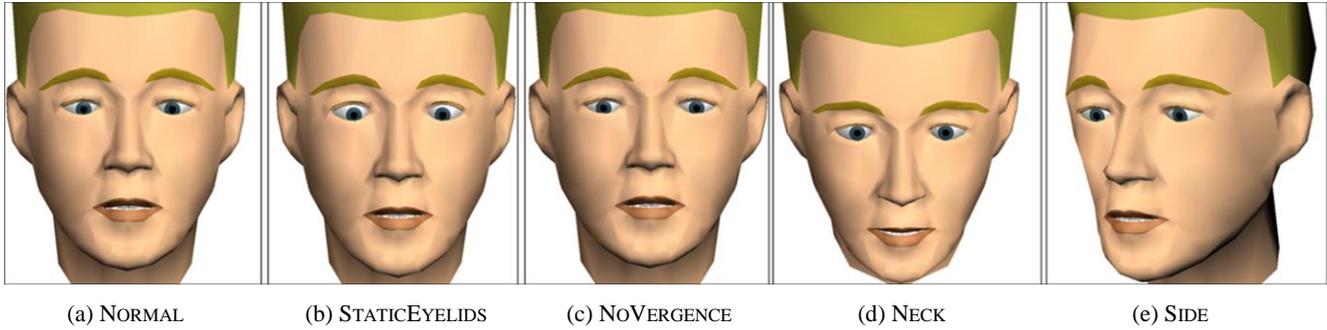


Figure 3. Example photos of the different conditions. The snapshots were taken from the same animated agent used in Furhat but were taken using a flat screen, for better visibility.

setting along the x and y axes. A human was seated opposite to the subject, at a similar height and distance from the table as Furhat (see Figure 2). Just as with Furhat, subjects were instructed to place the tagging object on the square where they perceived the human was looking. During this condition, the human agent always sustained a frontal head pose, while looking at the different target points on the grid with the eyes only. The condition included 9 different target points on the grid, corresponding to the 9 different regions.

3.2 Front vs. Side Viewing

In this condition, we explore the difference between perceiving gaze from the front of Furhat vs. looking at Furhat from the side. For the side condition, the subjects were seated at a 45 degrees angle from the head, hence having a side view of the eyes. As mentioned in 2.1, due to the Mona Lisa effect, showing the eyes of the agent on a 2D surface would mean that the eyes would always point in the same *relative* direction to the observer. Hence, if subjects move around the display, the direction of the eyes will move accordingly. This condition is aimed at measuring whether the eyes of Furhat will indeed keep being perceived gazing at the same absolute spatial point on the grid no matter whether the subject is sitting to the front of, or to the side of, Furhat.

In addition to that, the perceived direction of gaze might also depend on whether the eyes of Furhat are viewed frontal or from the side. For example, the viewing angle might affect the visibility of the eyes and their details and hence perhaps alter the perceived target point of gaze.

3.3 3D Gaze (Vergence) vs. Parallel Eyes

When looking at a face gazing frontal on a 2D screen, the eyes of the face will be looking at the observer no matter where the observer is standing, *and* no matter how far the observer is standing away from the screen. In a physically situated face-to-face setup, this need not be the case. The eyes of a human will be looking in parallel lines (both equally rotated) only when gazing at infinity. Humans focus on spatial points in space by rotating the eyes differently, so that both eyes are gazing the focal target, and hence, depending on the distance of the focus target, the eyes will rotate with two different angles. However, to our knowledge, it is no yet known whether humans can and do take advantage of this cue in order to more precisely locate the gaze target other humans are gazing at. It is not trivial to investigate that using human subjects, but the current setup, using an animated face, allows us to do that due to the easy control of the animated models.

Thus, the distance (depth) of the 9 different gaze targets on the grid is calculated using the geometrical model of the setup, then the direction of the gaze of the robot is tested when the depth information is taken into account (using different angles for each eye) and when it was not (using the same angle for each eye). This subtle difference can be seen between (a) and (c) in Figure 3.

3.4 Gaze vs. Head Pose (Neck)

Humans may look at objects using both head pose and eye gaze. Thus, they may fixate at points either by moving only their eyes while keeping their head static, by moving their head while keeping their eyes directed forward, or by a combination of both. However, it is has not been quantified whether one of them provides a more precise perception of the target focus point than the other. This condition compares the use of only gaze (the head always tilted 15 degrees down, facing straight ahead), vs. the use of only head pose (the gaze always looking straight ahead, relative to the head, while the head is facing the gaze target point).

3.5 Static vs. Dynamic Eyelids

When the human eyes move vertically, the eyelids automatically follow the eyes in the same direction, with the upper eyelid being localized exactly on top of the iris. This property of the eyes, although very easy to employ, it has mainly been ignored in the design and synthesis of animated agents and robots [13]. Although this movement in humans is automatic, and hence hard to control for in experimental setups, it is easy and straightforward to test it in analysis by synthesis setups using animated or robotic eyes.

In this condition, we therefore compare static eyelids (calibrated for the eyes looking frontal) vs. dynamic eyelids following the vertical movement of the eyes (the difference between (a) and (b) in Figure 3).

Table 1. The combination of conditions used in the experiment

Conditions	Coding
Furhat \times Gaze \times Vergence \times Dynamic \times Front	NORMAL
Furhat \times Gaze \times Vergence \times Static \times Front	STATIC EYELIDS
Furhat \times Gaze \times No Vergence \times Dynamic \times Front	NO VERGENCE
Furhat \times Gaze \times Vergence \times Dynamic \times Side	SIDE
Furhat \times Neck \times Vergence \times Dynamic \times Front	NECK
Human \times Gaze \times Vergence \times Dynamic \times Front	HUMAN

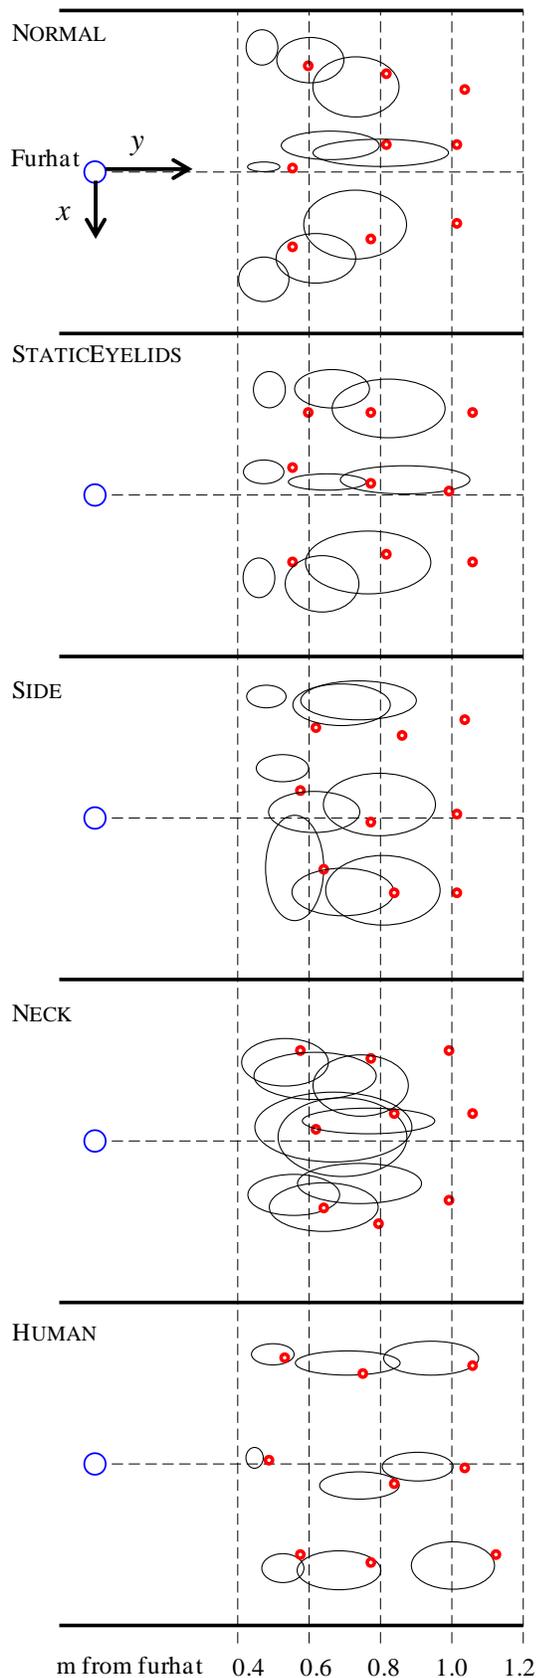


Figure 4. Distributions (st.dev.) of the answers (ellipses) and original targets (dots), plotted per condition and region.

4. EXPERIMENT

In order to investigate the different conditions described above, without having to explore all possible combinations of them, we first defined a “normal” condition and then defined five combinations of conditions. Each condition differed with regard to one specific aspect from the normal condition. This allows us to compare each of these combinations with the normal condition to investigate each factor separately. The different combinations are shown in Table 1.

The 6 conditions were presented in random order to 18 subjects. All subjects had normal, or corrected to normal vision. All subjects were volunteer students or employees, and carried out the experiment for free. Each condition consisted of 9 stimuli gaze points, each corresponding to one of the nine regions on the grid (divided by three rows and three columns). Thus, each subject answered to 54 stimuli points (9 stimuli * 6 conditions), and every condition received 162 answers (9 stimuli * 18 subjects). The experiment for each subject was done in one session, with no breaks, and with an average length of 20 minutes.

5. ANALYSIS AND RESULTS

The answers to all the gaze stimuli were transformed from the numbers of the squares on the grid, into their physical rotation angles from the eyes of Furhat, and into centimeters distance from Furhat (on the table plane), for both the horizontal and vertical axes (x will correspond to the horizontal dimension of the grid from the point of view of Furhat, and y will correspond to the vertical dimension of the grid from the point of view of Furhat).

Figure 4 shows a bird’s-eye view of the results for all conditions except NOVERGENCE, since it did not deviate significantly from the NORMAL condition (as we will come back to). The dots in the plot represent the original gaze targets, while the ellipses represent the estimate distributions of the answers of the subjects for each of the 9 regions. The horizontal radius corresponds to one standard deviation along the y axis, and the vertical radius corresponds to one standard deviation on the x axis. The distance on either axes between the center of the ellipses and the dots equals the shift between the average location of the perceived gaze and the original gaze target.

Table 2 shows the averages of the answers for all subjects and regions for both axes, calculated in degrees. The left columns show the distance between the average answers and the original targets (positive values meaning exaggerated answers), and thus correspond to the distance between the center of the ellipses and the dots in Figure 4. The right columns show the average distance between each answer and the mean of the distribution, thus corresponding to the size of the ellipses in Figure 4. All values were compared against the NORMAL condition and tested for significant difference (as indicated by the shades of the cells), using two tailed t-tests ($p < 0.05$; $dF = 161$). It is important to note that the left columns only show the shift in the location between the perceived gaze and the original gaze target – they do not indicate the precision of the gaze perception. However, they do indicate how the conditions affect the location of the perceived direction of gaze, and thus the need for calibration. On the other hand, the right columns provide a measure of the agreement between subjects. Thus, lower values in these columns indicate a higher precision. We will here discuss the results of each differentiating feature in the conditions, as described in section 3.

Table 2. Averages of the answers (in degrees) for all subjects and regions, for both axes. Grey cells indicate significant differences from the NORMAL condition.

	Diff from original target		Abs. diff from the perceived mean	
	x	y	x	y
NORMAL	6,07	-8,71	3,66	4,44
SIDE	5,48	-7,61	4,79	5,10
STATIC EYELIDS	5,82	-7,24	3,48	4,27
NO VERGENCE	6,40	-8,36	4,06	4,67
HUMAN	2,79	-3,54	3,18	3,13
NECK	-1,08	-6,44	3,24	5,86

5.1 Human vs. Furhat

The results indicate that the perception of the human gaze has a higher precision (at least on the x axis), as compared to Furhat (regardless of condition). The mean deviation of about 3 degrees corresponds surprisingly well with the findings on the perception of human gaze reported in [14]. It is a bit surprising to see that the target was in general perceived as being about 3.5 degrees below the actual target and 2.8 degrees to the side of the target. Since humans should have a lot of experience in perceiving each other's gaze, there is no obvious explanation for this.

These findings are important for the design of similar setups in human-human interaction settings, and provide estimates on how well objects can be identified using gaze depending on their relative placements in space. But they also provide a golden standard for the design of robot heads. One should bear in mind, however, that the human condition only used one human to generate gaze targets. The accuracy and precision of perceiving the direction of human eyes might be different across different humans, possibly depending on the shape and color of the eye; a question that mandates further research to investigate.

5.2 Static vs. Dynamic Eyelids

Using dynamic eyelids (i.e. lowering the eyelids in relation to the vertical position of the eyes) does not seem to affect the precision of the gaze perception. However, it does affect the perception of the target location, lowering it down, in average 1.47 degrees (difference between the NORMAL condition and the STATIC EYELIDS condition over the y axis). This is important to take into account when calibrating the psychometric gaze function, as we will come back to in 5.6.

5.3 Front vs. Side Viewing

As can be seen in Figure 4, looking at Furhat's gaze from a 45 degrees angle skews the perceived gaze location in an asymmetrical way. This affects the perceived target's location, as well as the precision, which becomes lower on the x-axis (4.79 degrees vs. 3.66 degrees). One possible explanation for this is that the visibility of both eyes is worse when looking from the side, especially when the head looks at the farther side of the table from where the subject is sitting (as seen in Figure 4). This means that, in the case of Furhat, the gaze function needs to be calibrated differently depending on where the observer is sitting.

5.4 3D Gaze (Vergence) vs. Parallel Eyes

No significant differences were found between generating gaze that accounts for the distance of the gaze target from the eyes, and between generating equally rotated eyes, neither in displacement nor in precision. There are several possible explanations for this.

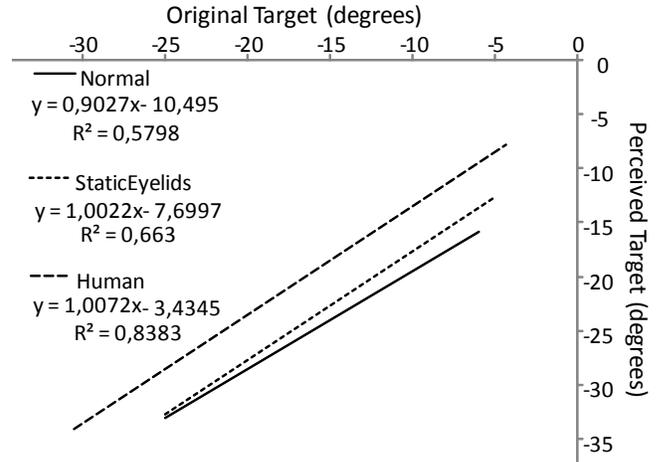


Figure 5. Psychometric functions for three of the conditions on the y-axis.

The spatial dimensions of the grid area could be too small to affect large perceivable differences between the relative rotations of both eyes. Another possibility is that the design of Furhat's eyes is not accurate enough to account for these subtle differences. A third possibility is that humans do not utilize depth information in eye rotations when determining gaze direction.

5.5 Gaze vs. Head Pose (Neck)

The main difference between the NECK and the NORMAL condition is that the precision on the y-axis in the perception of where Furhat is looking when moving the neck is significantly worse than when moving the eyes, which is strikingly obvious when looking at Figure 4. However, the precision on the x-axis is not affected at all (it is actually the condition that is closest to the HUMAN condition). Also, the displacement on the x-axis is the least of all conditions. Thus, head orientation seems to be a very accurate indication of where a person is looking on the x-axis, but very inaccurate on the y-axis. We do not have a good explanation of this, but it might have to do with the fact that the head is symmetrical on the x-axis but not on the y-axis.

5.6 The Psychometric function of gaze

The measure of precision presented earlier is informative on which factors in gaze perception result in better agreement between subjects, and hence the robot eyes should employ the one that provides the most precise one, that is the one that provides the smallest ellipses in Figure 4. This is however not the case for the location of the ellipses. The deviance of the location of the perceived gaze in relation to the original gaze target can be compensated for by recalibrating the robot eyes using data collected in the way we have done here.

By comparing the perceived gaze location and the original gaze, it is possible to estimate a linear fit to the data that results in a linear function that would provide the appropriate rotation of the eyes depending on required the location of the gaze target.

Looking at Table 2 and Figure 4, it is apparent that most conditions need some kind of calibration. However, different functions might be needed for different conditions. An example of such psychometric functions is presented in Figure 5. The figure shows the linear fits for the vertical gaze rotation (y-axis) for the NORMAL and the STATIC EYELIDS conditions, together with the estimated functions and how well they fit. We also added the HUMAN condition, which again illustrates the displacement of about 3.5 degrees on the y-axis.

6. CONCLUSIONS AND DISCUSSION

While the accuracy of the HUMAN condition is better than the other conditions (at least for the *y*-axis), the difference is not very big, even for the SIDE condition. This shows that there is virtually no Mona Lisa effect when observing Furhat – a finding that is in line with previous studies we have done on the use of gaze for multiparty dialogue [11]. The experimental setting used here could be useful for evaluating the gaze perception accuracy of different eye designs in robots.

The findings presented here have several implications for the design of robots that are supposed to take part in situated interaction with humans. When discriminating between objects on the *x*-axis, it might be useful to use head orientation, while discrimination on the *y*-axis might best be done using gaze. Another important implication is that it seems like different psychometric calibration functions should be used depending on the design of the robot, for example depending on whether the eyelids are dynamic and move with vertical gaze shifts or not. The possibility to do this is of course dependent on whether the eyes of the robot are actually used for robot vision (which is not the case with Furhat). Another complicating factor is that the psychometric function seems to be different depending on the viewing angle. Should the system take this into account and compensate for this? This question becomes even more challenging if we take into account that even human gaze would ideally need some calibration to convey the intended target. Thus, the designer has to decide whether the robot should mimic human behavior as much as possible, or maximize the correspondence between intended and perceived target.

Of course, the study also gives rise to several new questions that call for further investigation. There are many combinations of conditions that we haven't tested here. For example, does the perception of human head pose follow the same patterns as we saw here with Furhat (accurate on the *x*-axis but inaccurate on the *y*-axis)? Does side-viewing of human gaze show the same asymmetrical patterns that we saw with Furhat? There are also other individual factors that we haven't explored here. For example, it has been shown in studies of perception of human gaze that the distance of the perceiver may affect the perceived target [15]. Interestingly, they also found that the eyebrows play a role, since eyebrows correlate positively with the lowering of the eyeballs – it seems that humans may use the eyebrows to expand or restrict the availability of their gaze direction to others, which in turn affects the accuracy of detection the gaze direction [15]. The experimental setting presented here could be used to test if similar patterns can be replicated in human-robot interaction.

7. ACKNOWLEDGMENTS

This work is partly supported by the European Commission project IURO (Interactive Urban Robot), grant agreement no. 248314, as well as the SAVIR project (Situating Audio-Visual Interaction with Robots) funded by the Swedish Government (strategic research areas). We would like to thank all the volunteer subjects for taking part in the study.

8. REFERENCES

- [1] Argyle, M., Ingham, R., Alkema, F., & McCallin, M. 1973. The different functions of gaze. *Semiotica*, 7(1), 19-32.
- [2] Mirenda, P. L., Donnellan, A. M., & Yoder, D. E. 1983. Gaze behavior: A new look at an old problem. *Journal of Autism and Developmental Disorders*, 13(4), 397-409.
- [3] Tomasello, M., Hare, B., Lehmann, H., & Call, J. 2007. Reliance on head versus eyes in the gaze following of great apes and human infants: the cooperative eye hypothesis. *Journal of Human Evolution*, 52(3), 314-320.
- [4] Mutlu, B., Forlizzi, J., & Hodgins, J. 2006. A storytelling robot: Modeling and evaluation of human-like gaze behavior. In *Proceedings of 6th IEEE-RAS International Conference on Humanoid Robots* (pp. 518-523).
- [5] Torres, O., Cassell, J., & Prevost, S. 1997. Modeling gaze behavior as a function of discourse structure. *Proc. of the First International Workshop on Human-Computer Conversation*.
- [6] Boucher, J. D., Pattacini, U., Lelong, A., Bailly, G., Elisei, F., Fagel, S., Dominey, P. F., & Ventre-Dominey, J. 2012. I reach faster when I see you look: gaze effects in human-human and human-robot face-to-face cooperation. *Frontiers in neurobotics*, 6.
- [7] Yoshikawa, Y., Shinozawa, K., Ishiguro, H., Hagita, N., & Miyamoto, T. 2006. Responsive robot gaze to interaction partner. In *Proceedings of robotics: Science and systems*.
- [8] Johnson-Roberson, M., Bohg, J., Skantze, G., Gustafson, J., Carlson, R., Rasolzadeh, B., & Kragic, D. 2011. Enhanced Visual Scene Understanding through Human-Robot Dialog. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- [9] Al Moubayed, S., Edlund, J., & Beskow, J. 2012. Taming Mona Lisa: communicating gaze faithfully in 2D and 3D facial projections. *ACM Transactions on Interactive Intelligent Systems*, 1(2), 25.
- [10] Al Moubayed, S., Beskow, J., Skantze, G., & Granström, B. 2012. Furhat: A Back-projected Human-like Robot Head for Multiparty Human-Machine Interaction. To be published in *Cognitive Behavioural Systems. Lecture Notes in Computer Science*. Springer.
- [11] Al Moubayed, S., & Skantze, G. 2011. Turn-taking Control Using Gaze in Multiparty Human-Computer Dialogue: Effects of 2D and 3D Displays. In *Proceedings of AVSP*. Florence, Italy.
- [12] Delaunay, F., de Greeff, J., and Belpaeme, T. 2010. A study of a retro-projected robotic face and its effectiveness for gaze reading by humans. In *Proceedings of the 5th ACM/IEEE international conference on Human-robot interaction (HRI '10)*. IEEE Press, Piscataway, NJ, USA, 39-44.
- [13] Bailly, G., Raidt, S., & Elisei, F. 2010. Gaze, conversational agents and face-to-face communication. *Speech Communication*, 52(6), 598-612.
- [14] Gibson, J. J., & Pick, A. D. 1963. Perception of another person's looking behavior. *The American journal of psychology*, 76(3), 386-394.
- [15] Watt, R., Craven, B., & Quinn, S. 2007. A role for eyebrows in regulating the visibility of eye gaze direction. *The Quarterly Journal of Experimental Psychology*, 60(9), 1169-1177.