Accuracy of Perceiving Precisely Gazing Virtual Agents

Sebastian Loth
Social Cognitive Systems, CITEC, Bielefeld University
Bielefeld, Germany
sebastian.loth@uni-bielefeld.de

Gernot Horstmann
Neurocognitive Psychology, Faculty of Psychology and Sports Science, Bielefeld University
Bielefeld, Germany
gernot.horstmann@uni-bielefeld.de

Corinnna Osterbrink
Social Cognitive Systems, CITEC, Bielefeld University
Bielefeld, Germany
co.osterbrink@uni-bielefeld.de

Stefan Kopp
Social Cognitive Systems, CITEC, Bielefeld University
Bielefeld, Germany
skopp@techfak.uni-bielefeld.de

ABSTRACT

Eye gaze of is an informative social signal in interactions with other humans and also with virtual agents (VA). But for a successful communication, users have to accurately perceive the VA's point of gaze (POG). In our study, participants sitting opposite to a VA at a table indicated its POG by positioning a token on the table surface. We measured the perceptual accuracy within and between participants as well as the participants' response times and eye movements for five horizontally aligned gaze targets. We demonstrated that perceiving the VA met perceptual benchmarks from human lookers: a) variances within and between participants were only slightly larger, b) the VA's visual angle was linearly overestimated, and c) variances increased with the visual angle. Finally, participants showed large individual differences but were consistent in their own gaze behaviour and response times across trials and gaze targets.

CCS CONCEPTS

• **Human-centered computing → Empirical studies in interaction design**; *User studies*; *Accessibility design and evaluation methods*;

KEYWORDS

perception, perceptual accuracy, perceptual function, eye gaze, virtual agents

ACM Reference Format:

Sebastian Loth, Gernot Horstmann, Corinnna Osterbrink, and Stefan Kopp. 2018. Accuracy of Perceiving Precisely Gazing Virtual Agents. In *IVA '18: International Conference on Intelligent Virtual Agents (IVA '18), November 5–8, 2018, Sydney, NSW, Australia.* ACM, New York, NY, USA, Article 4, 6 pages. https://doi.org/10.1145/3267851.3267852

IVA '18, November 5–8, 2018, Sydney, NSW, Australia

ACM ISBN 978-1-4503-6013-5/18/11…$15.00
https://doi.org/10.1145/3267851.3267852

1 PERCEIVING THE EYE GAZE OF OTHERS

Eye gaze is a strong social cue. The ability to correctly identify the gaze target of another human is fundamental to development, learning and interactions [24, 26]. Eye gaze helps understanding intentions [15] and instructions [4], and can even override explicit verbal references [25]. Thus, researchers aimed at modelling human-like gaze behaviour for virtual agents (VA) [5, 23], e.g., for attention management [21] and reducing error rates [7]. However, observers have to identify the VA's point of gaze (POG) sufficiently precisely in order to interpret it properly and avoid misunderstandings.

Humans can identify the POG of other humans with great accuracy and precision if the *looker* gazes straight at the observer's face [12]. But as the gaze is directed more sidewards, precision and accuracy are reduced [8], i.e., the visual angle is overestimated and the POG misjudged [8, 10, 12, 13]. Estimating the POG depends on the observer's ability to judge the looker's eye and head rotation. Wollaston demonstrated this by presenting an image of a pair of eyes with noses pointing in different directions [14, 27]. Agents presented in 2D who gaze at the camera are subject to the *Mona Lisa*-effect whereby observers perceive Mona Lisa's gaze or Lord Kitchener's finger as tracing them as they move in front of the picture [6, 22]. We avoided the Mona Lisa effect by using virtual reality (VR) and focussed on human estimates of a VA's POG.

We investigated how and to what degree of accuracy and precision human observers perceive a VA's eye gaze in VR. After introducing the methods, we discuss benchmarks of perceiving human lookers, the large individual differences in our results and the implications for designing VAs and their environments.

2 METHODS

Participants were seated at a virtual table opposite of a VA and indicated their perception of the VA's POG by positioning a token on the table providing a measure of perceptual accuracy [6]. In other studies, participants moved an indicator along a fixed axis, e.g., a ruler with an adjustable marker [8]. However, if the perceived gaze vector does not intersect with that object, estimating some point of shortest distance to the perceived gaze vector introduces an additional source of noise. We avoided this by following [6].

Participants. The participants (8f, 2m, $M_{age} = 22[19, 30]$, 3 left & 7 right handed) from the Bielefeld University Psychology student pool received course credits and candy in exchange for their 30 min participation. All procedures were approved by University

Figure 1: The chair for the participants, a table and the VA. Participants indicated their perceived POG with the token.

Bielefeld's Ethics Committee under approval №2017-148. A written informed consent was obtained prior to the experiment.

Apparatus. A virtual room was created in Unity3D [3] for a HTC Vive with Tobii eye tracker integration (1080x1200 px per eye, 90 Hz eye tracking). The VA [1] was seated face-to-face with the participant at a table [11], see Fig. 1. The low-contrast eyes were replaced for improved perceptibility [2] and controlled by a realistic model [16–18]. The VA's Cyclopean eye (for computing the gaze vector) was 66 cm behind the table's centre line and 54 cm above the table surface. The participants were seated on a chair matching the VR. There was no real table in order to avoid collisions.

Materials. Five target points were at about half-distance between VA and participant. They were distributed evenly with distances of -30, -15, 0, 15, and 30 cm from the mid-point.

Procedure. The participant was seated with VR glasses and controllers and the eye tracker was calibrated. The session started with five randomly selected practice trials, followed by 100 experimental trials in random order (20 trials per target).

Each trial started with the token (a small blue cube) in front of the participant. The VA's head and eye movements were covered by a divider (large green rectangle) because they could be informative with respect to the POG [21], but see [12]. The divider moved upwards, the participants grabbed the token with the hair trigger button of the controller and positioned it on the table surface indicating their perceived POG. The position could be re-adjusted. They completed the trial with the top button of the controller. This lowered the divider for the next trial.

3 RESULTS

We excluded 14 trials because the participants placed the token next to the table. We report on the 986 remaining data points.

3.1 Perceived Visual Angles

3.1.1 Estimation. We estimated the central tendency and variance for each participant. *Within* participant variances are due to perceptual uncertainty about the POG and imprecisions in positioning the token. The SD_{within} indicates the required distance between two objects for identifying one of them by gaze, cf. [8]. The $SD_{between}$ estimates the difference in visual angle required for achieving agreement on the gaze target between observers.

The perceived horizontal visual angle was estimated with separate linear models for each participant and gaze target. The model's fitted coefficient is the gradient of the gaze vector and in turn, the visual angle. The intercepts of the models were small and statistically not significant, i.e., no systematic perceptual offsets. Thus, the models were re-computed without intercept.

3.1.2 Handedness and Target Side. Participants typically positioned the token with their preferred hand irrespectively of the gaze target's position. Thus, a left handed participant might save effort by positioning the token closer to the the central line with a right than with a left hand side target. This implies an interaction of target side and handedness. Thus, we conducted a Bayesian ANOVA [20] with the estimated perceived horizontal visual angles as dependent variable, and target location (*central*, *semi lateral*, and *lateral*), target side (*left* and *right*), and handedness (*left* and *right*) as fixed independent variables, and participants as random independent variable. There was no evidence for or against a main effect of target side, $BF_{10} = 1.162$, handedness, $BF_{10} = 1.747$, or their interaction, $BF_{10} = 2.281$. The corresponding analysis of the vertical angle revealed evidence against a main effect of target side, $BF_{10} = 0.256$, no evidence with regard to handedness, $BF_{10} = 0.489$, and evidence against their interaction, $BF_{10} = 0.211$. We concluded that handedness had no effect and that the VA's gaze was perceived symmetrically. Thus, we combined data points of the corresponding left and right gaze targets.

3.1.3 Perceptual Accuracy. The perceived visual angles were computed by averaging across the central tendencies of each participant. The variance between the central tendencies is reported as $SD_{between}$. The SD_{within} was computed by comparing each data point to the central tendency of the respective participant. For comparability to previous findings, an overall *SD* was computed by comparing each data point to the central tendency of the pooled set, see Table 1.

A linear model using the veridical angle as predictor for the averaged horizontal central tendencies of the participants, $R^2 = 0.999, F = 32590, p < .001$, revealed an overestimation by factor 1.18, see Fig. 3. Because the gaze targets were located on a straight line between VA and observer, the vertical angles differed slightly for central and lateral targets. There was a statistically significant effect on the perceived vertical angle, $BF_{10} = 6449$, indicating that the participants were sensitive to these small differences, see Fig. 2.

3.2 Time Course and Observer Gaze Behaviour

The response times (RT) were measured from the onset of the divider's upward movement until the participants confirmed their response, see Table 2. The analysis included the target location (*central*, *semi lateral*, and *lateral*) and target side (*left* and *right*) as independent variables, and participants as random variable. There was strong evidence against a main effect of target location, $BF_{10} = 0.095$, target side, $BF_{10} = 0.107$, and their interaction, $BF_{10} = 0.003$.

The eye tracking failed with two participants and we report data of eight participants. The participants' gaze was semantically annotated as *divider*, *controller*, *table*, *VA* and *other*. Gazing at the token was combined with the controller because it often occluded the token. The gaze data were normalised onto a trial length of

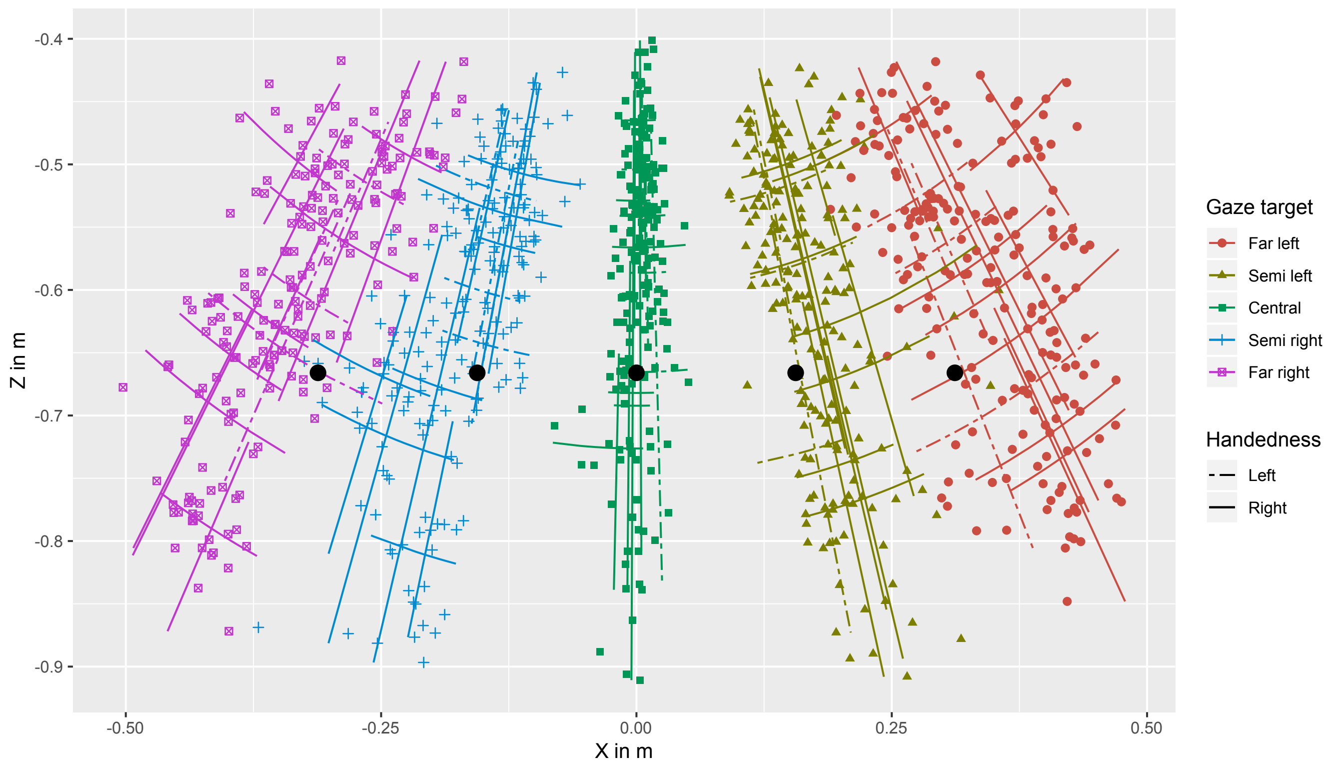

Figure 2: The VA is the origin of the coordinate system. Black dots are the gaze targets. Each point is one measured position of the token. Lines are the fitted perceived gaze angle per participant and their intersection the central tendency for each target.

Table 1: Veridical and perceived visual angles and their variances.

Angle	Target	Veridical	Perceived	Difference	$SD_{Between}$	SD_{Within}	SD_{Pooled}
Horizontal	Central	0.0°	0.3°	0.3°	0.93°	1.17°	1.51°
	Semi lateral	13.2°	15.5°	2.3°	2.05°	2.89°	3.56°
	Lateral	25.1°	29.5°	4.4°	2.86°	3.28°	4.29°
Vertical	Central	39.2°	42.6°	3.4°	4.30°	2.97°	5.24°
	Semi lateral	38.4°	40.3°	1.9°	4.42°	2.82°	5.27°
	Lateral	36.4°	38.4°	2.0°	4.26°	2.25°	4.80°

Table 2: Mean response times and *SD*s for positioning the token as a function of the gaze target in ms.

	Far left	Semi left	Central	Semi right	Far right
Mean	7976	8196	7470	8029	7901
SD	3320	3771	3258	3337	3567

1000 slices providing ratios per time slice. The high *within* participant correlations of these ratios between gaze targets (*central*, *semi lateral*, and *lateral*), see Table 3 showed that each participant had a consistent gaze pattern across targets. Next, data were averaged across gaze targets and a *between* participants correlation was tested, see Table 3. These were much lower than the *within* participants correlations. The strongest *between* correlations were on the divider and the VA. But a) the divider covered the VA's face and was only visible at the start of a trial, and b) all participants mainly gazed at the VA. Thus, these two correlations reflect design aspects whereas a general strategy would affect all task relevant objects. Thus, we concluded that participants were consistent in their gaze behaviour but the patterns differed between participants, see Fig. 4.

4 DISCUSSION

The participants' perceptual accuracy and precision of estimating a VA's POG on a table surface was comparable to experiments with human lookers with regard to the variances, the overestimation of the visual angle and the absence of a systematic offset. This was

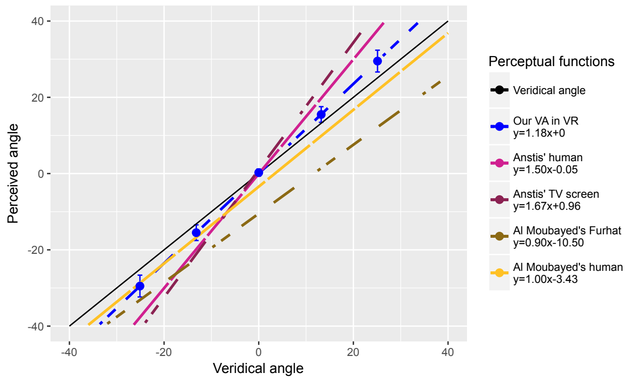

Figure 3: Perceptual functions of our VA and [6, 8] for comparison. Error bars show $SD_{between}$.

Table 3: Mean ratio of gazing at relevant objects and their averaged Pearson rs (using Fisher's z-transformation).

	Mean gaze ratio	*within* Pearson's r	*between* Pearson's r
Virtual Agent	0.47	0.91	0.64
Table	0.27	0.83	0.31
Controller + Token	0.16	0.83	0.07
Divider	0.07	0.96	0.73
Other	0.02	0.06	0.08

despite the additional challenges of VR compared to a real human looker: a) the HTC Vive has a visibly low resolution, b) the shape and rendering of the VA may distort the perceived POG, and c) the precision was subject to the participants' spatial coordination when positioning the token. Thus, while already indicating a very good performance, the variances in Table 1 provide a lower boundary of the perceptual accuracy and precision.

4.1 Accuracy of the Perceived Point of Gaze

In order to unequivocally identify one out of two objects by the VA's eye gaze, the distributions of the perceived eye gaze should not overlap. Otherwise, an observer would identify one and in some proportion of cases the other object as the gaze target, i.e., gaze was perceived ambiguously. The *SD*s of the perceived visual angles provide a metric for a separation between objects that enables sufficiently precise perception [8]. Within an individual, the separation has to be greater than the sum of the two SD_{within}s for the respective gaze targets. However, in order to achieve agreement between two observers, the individual differences have to be accounted for. Thus, the minimum separation increases to the sums of the respective $SD_{between}$s and SD_{within}s. For example, our lateral gaze targets (15 cm and 30 cm to sides and 66 cm in front of the looker) require a minimum separation of 11.06° or 13 cm. The actual separation was 15 cm and thus, the distributions of the perceived POGs in Fig. 2 show very little overlap. In contrast, the accuracy was greater with central targets. The SD_{within} and $SD_{between}$ were about 1° of visual angle equivalent or just 3 cm on the table surface.

4.1.1 Individual and Group Variances. Typically, the *SD* increases numerically with the measured units. We observed such an increase with greater visual angles, see Table 1. We would argue, however, that the variances were due to imprecisions in positioning the token, perceptual uncertainty, and individual differences. These sources of variance were not expected to vary with the visual angle. First, imprecisions in positioning the token contributed to the SD_{within}. Inherent differences in the participants' ability to position the token would have resulted in an interaction of side and handedness. But results were symmetric indicating that position errors added a constant to the SD_{within}. Secondly, the SD_{within} was affected by perceptual uncertainty. As with data from human lookers, it increased if the looker's eyes were rotated to the sides, even if the gaze target remained central (head countered eye rotation), and increased further the more lateral the gaze target was located [10]. Thirdly, systematic individual differences were identified in $SD_{between}$. The linear overestimation of the visual angles resulted in greater offsets between veridical and perceived POGs with larger angles. Thus, a greater numerical effect of individual differences in the degree of overestimation increased the $SD_{between}$ with sidewards POGs.

4.1.2 Overestimation of the Visual Angle. The linear overestimation of the perceived visual angle was attributed to the eyelids creating an aperture or 'peephole' for the observer [8]. If the eyeball rotates half-way to one side (45°), the pupil has traversed almost to the corner of the eyelids. This is more than half of the observable area of the eyeball. Thus, the rotation appears to be greater than it actually is. In contrast, the head rotation was underestimated with humans lookers [8, 10]. Greater convexity of the head's shape increases the perceived rotation and less convexity (image on a flat screen), or even concavity (in some sculptures), reduces it [8].

Combining head and eye rotation in a VA resulted in an overestimation of the visual angle as with a human looker but to a smaller degree, see Fig. 3. The individual overestimation coefficients of the horizontal angle for the semi lateral and far lateral gaze targets were correlated, Spearman's $\rho(8) = 0.758, p = 0.016$, indicating consistency within individuals. This was also the case for vertical angles, $\rho(8) = 1.000, p < 0.001$. However, there was no systematic relation between the individual horizontal and vertical coefficients for the semi lateral, $\rho(8) = 0.115, p = 0.759$, and lateral gaze targets, $\rho(8) = 0.152, p = 0.682$.

In sum, the individual differences in the overestimation coefficients were the main driver for the *between* participants variance. Attempts to correct this by providing initial feedback for a 'calibration' towards the looker showed no effect [12].

4.1.3 Benchmarks of Accuracy. In [13] participants judged whether a human was looking into their face (straight ahead) with an SD_{pooled} of 2.8°. This compares to our central condition with an SD_{pooled} of 1.51° and suggests that our VA was perceived at superhuman levels. However, the participant's face is an area leading to more frequent impressions of being looked at compared to a point and thus, leading to an underreporting of precision in [13].

Participants were asked to identify a gaze target on a transparent disk in [10]. At the central position, the SD_{pooled} was 0.73° (our VA 1.51°) and, at a lateral position of 12°, SD_{pooled} increased to 3.05° (our VA 4.29°). The design using a transparent disk allowed

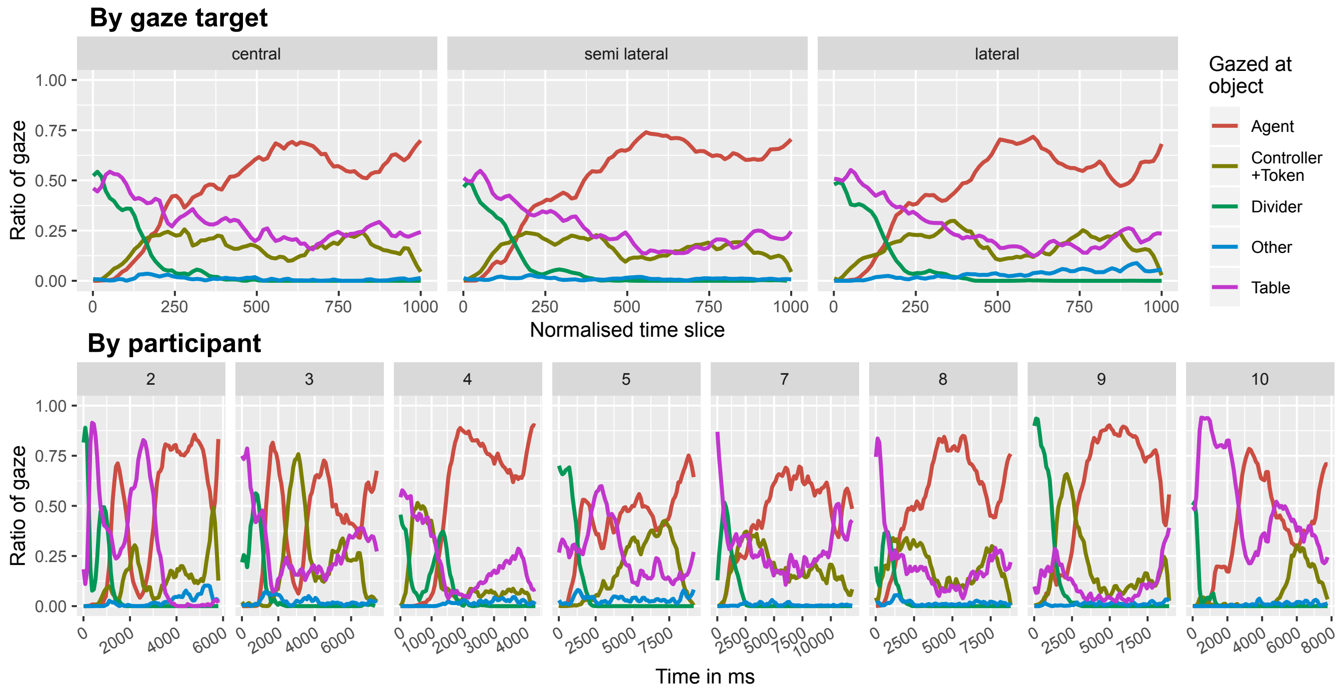

Figure 4: Mean ratio of gazed at items per time slice of standardised (top panel) and mean (bottom panel) trial duration.

a straight gaze. In contrast, targets on the table surface required a downward gaze. Thus, the perceptual precision could be greater in [10] than in our study due to, e.g., partly closed eyelids or a greater distance from a straight gaze line. [10] report only the pooled *SD*. Thus, we cannot distinguish between individual differences, e.g., a participant always used the overestimated and another always the underestimated response, or by imprecisions within participants. The reported variances are smaller than our *pooled* and similar to our *between* variances. In sum, [10] shows a small advantage of the human looker over the VA.

In [12], participants indicated the POG of a human looker with fixed response categories. Some participants were tested with three initial calibration points, SD_{pooled} = 3.93°, and some without, SD_{pooled} = 2.67°. The variances increased if the same looker was presented through a video connection, with calibration SD_{pooled} = 4.25° and without SD_{pooled} = 3.59°. As above, pooled data hinders a precise comparison. But the reported variances are in a similar magnitude.

In [6], participants were asked to put a token at the location of the perceived POG. We believe that [6] reports a pooled mean average difference (MD) of 3.18° on the horizontal axis with a human looker and 3.66° with the Furhat robot. This compares to 0.90°, 2.29° and 2.64° for our *central*, *semi lateral* and *lateral* targets respectively. The impression of our VA's superhuman precision is attributable to the reasons mentioned above. Additionally, the grid on the table might have mislead the participants by presenting potential gaze targets. A straight gaze of Furhat was perceived as looking to the side at an angle of 10.5° corresponding to 12.3 cm in our setting. In contrast, our linear model showed a statistically non-significant intercept of 0.3° that is irrelevant for any application.

In sum, the comparison of the perceptual accuracy and precision of the VA showed that it was similar to a human looker.

4.2 Individual Patterns

The error rates and RTs often increase in conjunction as the task becomes more difficult. Thus, RTs could be slower to lateral than to central gaze targets but they were not. In fact, the participants showed individual differences but their RT patterns were consistent across all conditions. We quantified this by normalising *SD*s of the RT distribution by the respective mean value and averaging them. Doing so by target position (0.44), by participant (0.30) and by position and participant (0.29) revealed a small effect of position and a large effect of participant. This pattern excludes a speed-accuracy trade-off where the participants responded faster with lateral targets than with central targets and thereby lost precision. Thus, the variance in RT was not driven by the location of the gaze target but by individual differences.

The eye tracking data were also analysed with regard to the target position and the participants, see Fig. 4. The participants could have looked more frequently up and down with lateral compared to central gaze targets. However, the eye tracking patterns were similar across targets. In all trials, participants initially looked at either the table surface where the token was located or the divider. As the divider moved up, the probability of looking at the VA for estimating the POG and at the controller for positioning the token rapidly increased. All probabilities were similar across gaze targets and remained comparably stable until the end of the trial. As with the RT, differences in gaze patterns were due to individual differences. However, the participants were consistent and looked at the same object at the same time in each trial with above 80 % probability as illustrated in the lower panel of Fig. 4. For example, participant 9 started with the divider, then collected the token, focussed on the VA almost until the end of the trial and quickly checked the position of the token indicated as table/controller.

In sum, accuracy and precision decreased with more sidewards targets. In contrast, RTs and gaze patterns showed individual differences but great consistency within individuals across all trials. Our results emphasise individual differences in task performance. When modelling behaviours, an individual is preferred over an average across individuals [9, 19]. In fact, we have not observed an average participant. Thus, averaging would create an unobserved chimera.

5 CONCLUSION

The directed eye gaze of a VA was perceived accurately. It meets benchmarks of perceiving the eye gaze of human lookers: a) the perceived visual angles were linearly overestimated b) accuracy and precision decreased with lateral targets, and c) the ability to estimate the VA's gaze is comparable to a human looker.

Eye gaze as a cue supports references to an object, addressing a user in a multi-party dialogue, or providing social signals in turn-taking. In order to ensure that two individuals agree in their judgements about the gaze target, two potential targets should have a minimum separation that can be computed from the *SD* of the perceived visual angles. This threshold is determined by *within* and substantial *between* participants variance and the targets' position relative to the VA and the observer.

There were large individual differences despite the relatively simple task. The participants differed in their response time, how they performed the task, which objects were relevant in time, and their estimates of the visual angle. However, each participant was consistent with their own pattern. Thus, a VA should model the behaviour of one individual rather than an average. The RT and performance patterns showed that a VA does not have to wait longer for a response but rather use the required spatial separation between targets in order to ensure correct identification.

ACKNOWLEDGMENTS

We would like to thank Conrad Alting, Marie Renan Bohle, Annika Buchsteiner, Marie Führing, Jonas Panhuysen, and Clarissa Sävecke, for their contributions in discussions and testing participants. This work was funded as part of the Cluster of Excellence Cognitive Interaction Technology "CITEC" (EXC 277), Bielefeld University.

REFERENCES

[1] 2015. Autodesk Character Generator. (July 2015). https://charactergenerator.autodesk.com/
[2] 2016. Eye Advanced. (Dec. 2016). http://tanukidigital.com/eyeadvanced/
[3] 2017. Unity 3D. (Sept. 2017). https://unity3d.com
[4] Henny Admoni, Anca Dragan, Siddhartha S. Srinivasa, and Brian Scassellati. 2014. Deliberate delays during robot-to-human handovers improve compliance with gaze communication. ACM Press, 49–56. https://doi.org/10.1145/2559636.2559682
[5] Henny Admoni and Brian Scassellati. 2017. Social Eye Gaze in Human-Robot Interaction: A Review. *Journal of Human-Robot Interaction* 6, 1 (March 2017), 25. https://doi.org/10.5898/JHRI.6.1.Admoni
[6] Samer Al Moubayed and Gabriel Skantze. 2012. Perception of gaze direction for situated interaction. In *Gaze-In '12 Proceedings of the 4th Workshop on Eye Gaze in Intelligent Human Machine Interaction*. ACM Press. https://doi.org/10.1145/2401836.2401839
[7] Sean Andrist, Michael Gleicher, and Bilge Mutlu. 2017. Looking Coordinated: Bidirectional Gaze Mechanisms for Collaborative Interaction with Virtual Characters. ACM Press, 2571–2582. https://doi.org/10.1145/3025453.3026033
[8] Stuart M. Anstis, John W. Mayhew, and Tania Morley. 1969. The Perception of Where a Face or Television 'Portrait' Is Looking. *The American Journal of Psychology* 82, 4 (Dec. 1969), 474. https://doi.org/10.2307/1420441
[9] Kirsten Bergmann, Stefan Kopp, and Friederike Eyssel. 2010. Individualized Gesturing Outperforms Average Gesturing – Evaluating Gesture Production in Virtual Humans. In *Proceedings of the 10th International Conference on Intelligent Virtual Agents (LNCS 6356)*, J Allbeck, N Badler, T Bickmore, Catherine Pelachaud, and A Safonova (Eds.). Springer International Publishing, Berlin, Heidelberg, 104–117. https://pub.uni-bielefeld.de/publication/1993187
[10] Marvin G. Cline. 1967. The Perception of Where a Person Is Looking. *The American Journal of Psychology* 80, 1 (March 1967), 41. https://doi.org/10.2307/1420539
[11] Epix. 2014. 3D model of Samsung SUR40. (March 2014). https://3dwarehouse.sketchup.com/model/3532707a0addc62e13680f9c7026bd5b/SAMSUNG-SUR40-Interactive-Surface-Table
[12] Caroline Gale and Andrew F. Monk. 2000. Where am I looking? The accuracy of video-mediated gaze awareness. *Perception & Psychophysics* 62, 3 (April 2000), 586–595. https://doi.org/10.3758/BF03212110
[13] James J. Gibson and Anne D. Pick. 1963. Perception of Another Person's Looking Behavior. *The American Journal of Psychology* 76, 3 (Sept. 1963), 386. https://doi.org/10.2307/1419779
[14] Heiko Hecht, Evgenia Boyarskaya, and Akiyoshi Kitaoka. 2014. The Mona Lisa effect: Testing the limits of perceptual robustness vis-à-vis slanted images. *Psihologija* 47, 3 (2014), 287–301. https://doi.org/10.2298/PSI1403287H
[15] Chien-Ming Huang, Sean Andrist, Allison Sauppé, and Bilge Mutlu. 2015. Using gaze patterns to predict task intent in collaboration. *Frontiers in Psychology* 6 (July 2015). https://doi.org/10.3389/fpsyg.2015.01049
[16] Laurent Itti, Nitin Dhavale, and Frederic Pighin. 2003. Realistic avatar eye and head animation using a neurobiological model of visual attention. In *Proc. SPIE 5200, Applications and Science of Neural Networks, Fuzzy Systems, and Evolutionary Computation VI*, Bruno Bosacchi, David B. Fogel, and James C. Bezdek (Eds.). SPIE Digital Library, San Diego, CA, USA, 64. https://doi.org/10.1117/12.512618
[17] Tore Knabe. 2017. Realistic Eye Movements. (Sept. 2017). http://tore-knabe.com/unity-asset-realistic-eye-movements
[18] Sooha Park Lee, Jeremy B. Badler, and Norman I. Badler. 2002. Eyes alive. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques (ACM Transactions on Graphics (TOG))*, Vol. 21. ACM New York, NY, USA, San Antonio, 637–644. https://doi.org/10.1145/566654.566629
[19] Sebastian Loth. 2017. Beyond Likeability: Investigating Social Interactions with Artificial Agents and Objective Metrics. *Frontiers in Psychology* 8 (Sept. 2017), 1662. https://doi.org/10.3389/fpsyg.2017.01662
[20] Jonathan Love, Ravi Selker, Josine Verhagen, Martin Smira, Anneliese Wild, Maarten Marsman, Quentin Gronau, Richard D Morey, Jeffrey N Rouder, and Eric-Jan Wagenmakers. 2018. JASP. (2018). https://jasp-stats.org/
[21] Tomislav Pejsa, Sean Andrist, Michael Gleicher, and Bilge Mutlu. 2015. Gaze and Attention Management for Embodied Conversational Agents. *ACM Transactions on Interactive Intelligent Systems* 5, 1 (March 2015), 1–34. https://doi.org/10.1145/2724731
[22] Sheena J. Rogers. 1995. Perceiving Pictoral Space. In *Perception of space and motion*, William Epstein and Sheena J. Rogers (Eds.). Academic Press, San Diego, 119–164.
[23] Kerstin Ruhland, Christopher Edward Peters, Sean Andrist, Jeremy B. Badler, Norman I. Badler, Michael Lee Gleicher, Bilge Mutlu, and Rachel McDonnell. 2015. A Review of Eye Gaze in Virtual Agents, Social Robotics and HCI: Behaviour Generation, User Interaction and Perception: A Review of Eye Gaze. *Computer Graphics Forum* 34, 6 (Sept. 2015), 299–326. https://doi.org/10.1111/cgf.12603
[24] Atsushi Senju and Mark H. Johnson. 2009. The eye contact effect: mechanisms and development. *Trends in Cognitive Sciences* 13, 3 (March 2009), 127–134. https://doi.org/10.1016/j.tics.2008.11.009
[25] Maria Staudte and Matthew W. Crocker. 2009. The effect of robot gaze on processing robot utterances. In *Proceedings of the 31st Annual Conference of the Cognitive Science Society*. Amsterdam, Netherlands. http://www.coli.uni-saarland.de/~masta/cogsci09.pdf
[26] Michael Tomasello and Malinda Carpenter. 2007. Shared intentionality. *Developmental Science* 10, 1 (Jan. 2007), 121–125. https://doi.org/10.1111/j.1467-7687.2007.00573.x
[27] William Hyde Wollaston. 1824. On the Apparent Direction of Eyes in a Portrait. *Philosophical Transactions of the Royal Society of London* 114 (Jan. 1824), 247–256. https://doi.org/10.1098/rstl.1824.0016