A Head-Eye Coordination Model for Animating Gaze Shifts of Virtual Characters

Sean Andrist, Tomislav Pejsa, Bilge Mutlu, Michael Gleicher

Department of Computer Sciences, University of Wisconsin–Madison
1210 West Dayton Street, Madison, WI 53706, USA
{sandrist, tpejsa, bilge, gleicher}@cs.wisc.edu

ABSTRACT

We present a parametric, computational model of head-eye coordination that can be used in the animation of directed gaze shifts for virtual characters. The model is based on research in human neurophysiology. It incorporates control parameters that allow for adapting gaze shifts to the characteristics of the environment, the gaze targets, and the idiosyncratic behavioral attributes of the virtual character. A user study confirms that the model communicates gaze targets as effectively as real humans do, while being preferred subjectively to state-of-the-art models.

Categories and Subject Descriptors

H.1.2 [**Models and Principles**]: User/Machine Systems—*Human factors*; H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Evaluation/methodology, User-centered design*

General Terms

Algorithms, Design, Experimentation, Human Factors

Keywords

Gaze shifts, head-eye coordination, virtual characters, control models

1. INTRODUCTION

Directed gaze shifts—the intentional redirection of gaze toward a particular piece of information in the context of interaction—are a fundamental building block of human gaze behavior. Through subtle variation in timing and movement of the head and eyes in shifting the gaze, individuals construct a range of complex communicative behaviors. When animating a virtual agent, control mechanisms must synthesize the wider range of such movements so that the agent displays natural communicative behaviors, yet provide sufficient control over the subtleties of the movements to allow for individual variation and expressions. Creating fine-granulated control mechanisms for gaze that achieve a combination of communicative effectiveness, naturalness, and parametric control remains an open challenge.

In this paper, we present a parametric control model that can be used to animate gaze shifts of virtual characters. Our model builds on findings from neurophysiology and procedurally specifies combined head and eye movements to create humanlike gaze shifts. The neurophysiological basis of our model helps achieve effective communication and subjective naturalness, while the procedural implementation allows for parametric control over the gaze shifts generated by the model. An empirical study confirms that the model meets our goals of creating gaze shifts that effectively communicate gaze direction and appear natural and realistic. The primary advantage of our model over previous ones is that it provides equivalent or better measured performance while providing additional control. For example, in a follow-up study which includes a less comprehensive description of the model and empirical study presented here, we show that manipulations in the model achieve high-level outcomes such as increased affiliation and learning [1].

2. BACKGROUND

This section reviews existing models for gaze synthesis from research in computer graphics and human-computer interaction (HCI), as well as the neurophysiological research that informs our gaze model.

2.1 Models for Gaze Synthesis

Numerous gaze models have been proposed in the literature, each with different methods, goals, and contributions. For example, data-driven models [3, 11] are a common and powerful approach. However, it is often difficult to manipulate parameters or incorporate known constraints in these models without providing new hard-to-find or hard-to-create examples. The "expressive gaze model" [16] takes a hybrid data-diven/procedural approach, focusing on the communication of emotion. Because this approach handles the head and eyes separately, it does not capture the complexities of eye-head synchronization during gaze shifts as our model does. Other models consider *where* the agent should be looking [12, 13], determine *why* an agent might be looking toward particular targets [6, 17, 23], or generate *idle* gaze behavior when the agent is not actively gazing [2, 18].

A current state-of-the-art model takes a similar approach to ours, procedurally generating gaze shifts while taking idiosyncratic head propensity into account [22]. Head propensity is

Gaze-in'12, October 26, 2012 Santa Monica, California, USA

defined in this model as the extent to which the head is employed in generating a gaze shift. While the head propensity variable can achieve some of the effects of our parameters, the existing model does not involve head latency and has not been shown to accommodate the wider range of gaze behaviors. Our evaluation confirms that our model generates more natural and realistic gaze shifts than the head propensity model does. Our model further achieves high-level outcomes such as improved affiliation with the agent and improved learning through principled manipulations of the subtle low-level variables provided by the model [1].

2.2 Neurophysiological Models of Gaze

Research in neurophysiology has studied how humans carry out gaze shifts by coordinating head and eye movements in a tightly connected dynamic process. In most existing models of directed gaze shifts, kinematics of saccadic movements, such as duration and peak velocity, are simplified, producing movements that depend only on the amplitude and direction of directed eye movements towards the target. In reality, concurrent head movements affect eye movements significantly [5]. This causal relationship holds in both directions; for example, head movement amplitude decreases as the eyes movements start at increasingly contralateral positions (i.e., oriented away from the target in relation to head direction). Shifts that start at such positions require the eyes to increase their contribution to the shift [19].

Head movement latency in relation to the onset of the eye movement can vary from person to person and from task to task. Factors such as target amplitude, the predictability of the target, target saliency, vigilance of the subject, and whether the gaze shift is forced or natural affect this latency [21, 24]. The modality of the gaze target makes a difference as well; studies that compare auditory and visual targets show that eyes tend to lead the head most often when people orient to visual targets, whereas the head tends to lead the eyes most often when people orient to auditory targets [8, 9].

Humans are mechanically limited in their ability to rotate their eyes, a limitation referred to as the oculomotor range (OMR). The human OMR has been estimated to be between 45° and 55°. However, encoding these OMR values as static parameters into virtual humans is not sufficient, as the effective OMR may fluctuate during the course of a single gaze shift. This fluctuation is a product of a neural (as opposed to mechanical) limitation imposed on eye motion [10].

The degree to which individuals use their heads in performing a gaze shift is highly idiosyncratic. The neurophysiological research literature describe some people as "head-movers," i.e., individuals who move their head fully to align with the gaze target every time, and some as "non-head-movers" [7]. From a bio-mechanical standpoint, humans should universally be "non-head-movers", as fully moving the head—which is almost a hundred times heavier than the eyes—is not an economical solution [14]. This neurally based idiosyncratic characteristic of head and eye movements is not captured by naïve inverse kinematics solutions for animating gaze shifts.

3. MODEL OF HEAD-EYE COORDINATION

In the model we present here, the parameters of the specific gaze shift (e.g., the target direction) and parameters of the character (e.g., maximum head velocity) are used to compute a number of internal timing parameters at the onset of the gaze movement. Table 1 lists all parameters included in our model. Once the internal parameters are computed, the gaze shift begins, and eyes and the head are rotated directly towards the target at dynamically-changing angular velocities. The velocities are recomputed in each frame of animation based on the progress of the gaze shift and the current rotations of the eyes and the head, allowing the model to react to perturbations of the head position or target during motion. The model calculates the rotations for the eyes independently in order to achieve convergence. A visual representation of this model is provided in Figure 1.

3.1 Internal Parameter Computation

The first phase in generating gaze shifts is to determine the latency of the onset of head movement, hl, in relation to the onset of the eye movement (Figure 1b) . Whether an individual follows a head-first or an eyes-first approach is determined by factors such as the vigilance of the agent (AG_{vig}), the target salience (GT_{sal}), the eccentricity of the target ($GAmp$), the predictability of the target location (GT_{pred}), the modality of the target—visual or auditory—(GT_{mod}), and the intent of the agent—forced or natural shift—(AG_{int}) (Figure 1a). Each of these factors is associated with a different likelihood ratio of leading to a head-first versus eyes-first gaze shift. A ratio value r is defined as $\frac{P_h}{1-P_h}$, where P_h is the probability of a head-first gaze shift. These ratios are summarized in Table 2. For example, gaze shifts with large amplitudes (greater than 30°) are 3.05 times more likely to involve a head-first approach. When considered in isolation from other parameters, auditory targets always produce head-first gaze shifts, hence the ratio of ∞ [24].

Rewriting for P_h in the ratio definition above, we get

$$P_h = \frac{r}{r+1}.$$

We can compute the final probability of a head-first gaze shift by linearly interpolating between the probability of a head-first gaze shift (P_h) and an eyes-first gaze shift ($1 - P_h$). Probabilities for each factor are sampled to determine whether each suggests a head-first or an eyes-first movement. These various "votes" on the type of movement are combined to determine the head latency. If f is the number of factors that vote to determine a head-first gaze shift and n is the total number of parameters, then the ratio s can be computed as $\frac{f}{n}$. We can use s to compute the head latency (hl) for the gaze shift by linearly interpolating between the head latency of a purely head-first gaze shift, hl_h, and the head latency of a purely eyes-first gaze shift, hl_e. We chose to use -100 ms for hl_h and 100 ms for hl_e in the implementation of our model based on the range of values proposed in the neurophysiology literature [24].

3.2 Generating Gaze Motion

Once the model determines the head latency, it initiates the movements of the eyes and/or the head. Each eye and the head move towards the target, following velocity profiles that resemble standard ease-in and ease-out functions (Figure 1c). The maximum velocity of the eyes and head, EV_{max} and HV_{max}, are computed based on positive linear relationships with the amplitude of the intended gaze shift [10]. We derived a piecewise polynomial function to approximate the full velocity profile for both the head and the eyes determined from the literature [18, 14]. This polynomial function can be expressed as follows, where g is the proportion of the gaze

Table 1: All parameters of the model

Parameter	Type	Symbol	Description
Gaze Target Position	Input	GT_{pos}	The destination of the gaze shift.
Gaze Amplitude	Internal	$GAmp$	The rotational difference between the current and the final gaze alignment. Determined from the gaze target.
Target Predictability	Input	GT_{pred}	A value between 0 and 1, where 0 and 1 correspond to unpredictable and predictable targets, respectively. Could be manually manipulated or driven by the agent's cognitive system.
Target Saliency	Input	GT_{sal}	A value between 0 and 1, where 0 and 1 correspond to low and high saliency, respectively. Could be manually manipulated or driven by the agent's vision system.
Target Modality	Input	GT_{mod}	A binary value that indicates whether the target is visual or auditory.
Agent Vigilance	Input	AG_{vig}	A value between 0 and 1, where 0 and 1 correspond to low and high vigilance, respectively. Could be manually manipulated or driven by the agent's cognitive system.
Agent Intent	Input	AG_{int}	A value between 0 and 1, where 0 and 1 correspond to natural and forced gaze shifts, respectively. Could be manually manipulated or driven by the agent's cognitive system.
Head Latency	Internal	hl	A value that varies between -100 ms and 100 ms. Stochastically determined from the above inputs.
Oculomotor Range	Input	OMR	The maximum pitch and yaw of the agent's eyes in degrees.
Effective OMR	Internal	OMR_{eff}	Neurally determined limits on eye motion. Updated throughout gaze shift.
Head Alignment	Input	h_{align}	A value between 0% and 100%, where 0% corresponds to minimum head alignment and 100% corresponds to full head alignment.
Maximum Eye Velocity	Internal	EV_{max}	Computed based on a positive linear relationship with the amplitude of the intended gaze shift, which saturates at about 500°/sec for gaze shifts at or beyond the OMR of the character [10].
Eye Velocity	Internal	EV	Follows the velocity profile described in the text.
Maximum Head Velocity	Internal	HV_{max}	Computed based on a positive linear relationship with the amplitude of the intended gaze shift [10]. For example, a gaze shift of 24° will result in a maximum head velocity of 50°/sec in our implementation.
Head Velocity	Internal	HV	Follows the velocity profile described in the text.
Initial Eye Position	Input	IEP	Initial gaze configuration of the eyes and head. Presented as a rotational offset of the current eye rotation from a central (in-head) orientation. Employed only when the rotational offset is contralateral (on the opposite side of center) to the target.

shift completed, V_{max} is the maximum velocity, and V is the current calculated velocity.

$$V = \begin{cases} 2V_{max} \cdot g & g \in [0, 0.5) \\ 4V_{max} \cdot g^2 - 8V_{max} \cdot g + 4V_{max} & g \in [0.5, 1] \end{cases}$$

Human eye rotations have limitations defined by the oculomotor range (OMR) (Figure 1d). A virtual character's baseline OMR can be geometrically determined based on the size of the eye cavities of the character model. At the onset of a gaze shift, OMR_{eff}—a *neurally* determined limit on eye motion—is computed based on the initial eye position (IEP) and the OMR. IEP is measured in degrees as a rotational offset of the current eye orientation from a central (in-head) orientation. This value is only non-zero when the rotational offset is contralateral (on the opposite side of center) to the target. When the eyes begin the gaze shift at these angles, the OMR_{eff} has a value close to the original baseline OMR. When the eyes begin the gaze shift closer to a central orientation in the head, the OMR_{eff} diminishes [4]. We approximated this relationship in our implementation with the following function:

$$OMR_{eff} = OMR \cdot (\frac{1}{360} IEP + 0.75).$$

OMR_{eff} is also updated throughout the gaze shift at every time step according to the concurrent head velocity, HV. As the head moves faster, the OMR_{eff} diminishes [10]. This relationship was approximated in our implementation with the following function, where OMR_{IEP} is the value for OMR_{eff} that was computed in the previous equation at the onset of the gaze shift.

$$OMR_{eff} = OMR_{IEP} \cdot (\frac{-1}{600} HV + 1)$$

Head alignment, h_{align}, is a user-defined parameter which

Table 2: Ratio of *likeliness of head-first shift / likeliness of eyes-first shift*

Amplitude $\frac{large}{small}$	Intent $\frac{forced}{natural}$	Predictability $\frac{high}{low}$	Vigilance $\frac{high}{low}$	Target Saliency $\frac{high}{low}$	Target Modality $\frac{auditory}{visual}$
3.05	2.20	1.6	1.33	0.53	∞

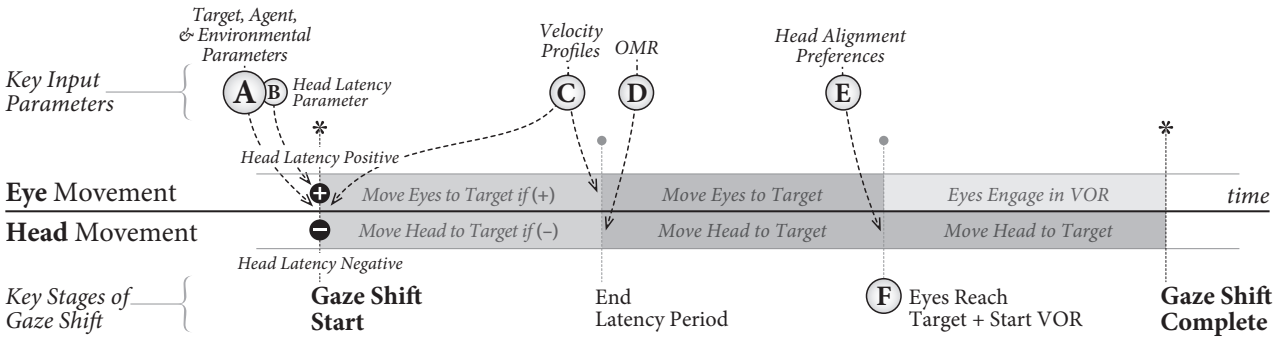

Figure 1: A visual representation of our model of gaze shifts mapped onto time. Key input variables and processes include (A) target, agent, and environmental parameters, (B) head latency, (C) velocity profiles for head and eye motion, (D) oculomotor range (OMR) specifications, (E) head alignment preferences, and (F) the vestibulo-ocular reflex (VOR).

specifies a significant amount of idiosyncratic variability in the behaviors of the character, namely, whether the character is a "head-mover" or a "non-head-mover" (Figure 1e). A parameter value of 0% for head alignment indicates that once the eyes have reached the gaze target, the head stops moving. On the other hand, at a 100% parameter value for head alignment, the head continues rotating until it is completely aligned with the target. Head alignment values between these two extremes can be computed using a linear interpolation between the two corresponding rotational values.

The last phase of the model involves the vestibulo-ocular reflex (VOR) (Figure 1f). When the eyes reach the gaze target, they remain locked to the target for the remainder of the gaze shift, rotating in the opposite direction of head motion until the head completes its portion of the gaze shift.

Ancillary Components of Our Model – Our model also includes a blink controller that serves two key functions: generating gaze-evoked blinking as described by Peters [22] and idle blink behavior. In addition, when the agent is not actively engaging in a gaze shift following our model, the eyes are controlled by an implementation of the model presented by Lee et al. [18]. This implementation creates random saccadic eye movements in a principled way and dramatically increases the realism of the character.

4. EVALUATION

Development of our model was followed by an empirical evaluation of the communicative accuracy and perceived naturalness of gaze shifts that our model generated. We conducted a study to compare gaze shifts generated by our model against those generated by a state-of-the-art model [22] as well as against gaze shifts displayed by human confederates. In addition we explored the effect that participant gender and the gender of the virtual character had on our measures, as gender is known to strongly shape gaze perception in artificial agents [20].

Experimental Setup and Task – Participants observed a series of videos in which either an animated virtual character or a human confederate shifted gaze toward one of sixteen objects arranged on a desk. This simplified scenario allowed us to focus our evaluation on the effectiveness and naturalness of gaze shifts, while minimizing contextual and interactional factors and facilitating the matching of animated and real-world conditions. Participants observed the agent from the perspective of a collaborator seated across from the agent or human confederate. The objects on the desk were separated into four groups, and were distinguished by color and shape. Still images from the videos are shown in Figure 2.

Study Design – We conducted a 2x2x8 factorial experiment with split-plot design. Our factors included participant gender (two levels varying between participants), gender of the virtual agent (two levels varying between participants) and model type (eight levels varying within participants). The model type independent variable included comparison conditions for gaze shifts generated by the head propensity model [22] and those produced by our model. For each model, we manipulated the model parameters to create *three* distinct comparison conditions with different head alignment/propensity levels, 0%, 50%, or 100%, with the goal of determining how changes in the model parameters affected communicative accuracy and perceived naturalness.

The model type independent variable also included two control conditions. In the first control condition, a male or female human confederate presented gaze shifts toward the object on a desk in front of him/her. The second control condition involved a virtual agent maintaining gaze toward the participant without producing any gaze shifts.

Study Procedure – Each participant was shown 32 videos of a virtual character or human. In the videos, the agents or the confederates gazed toward the participant, announced that they are about to look toward an object with a specific color on the table, shifted their gaze toward the object, and moved their gaze back toward the participant. Following each video, the participants filled out a questionnaire for subjective evaluation. Participants observed each gaze model generating gaze shifts toward all object types, colors, and positions. We randomized the order in which the participants observed the videos to minimize transfer effects. Each video was 10 seconds long, with the overall study lasting approximately 20 minutes.

Participants – Ninety-six participants (46 females and 50 males) took part in the study. The participants were recruited through Amazon.com's Mechanical Turk online marketplace, following crowd-sourcing best practices to achieve a wide range of demographic representation and mitigate issues from differences in participants' viewing parameters, such as screen resolution [15]. Participants received $2.50.

Measurement – The study used two dependent variables:

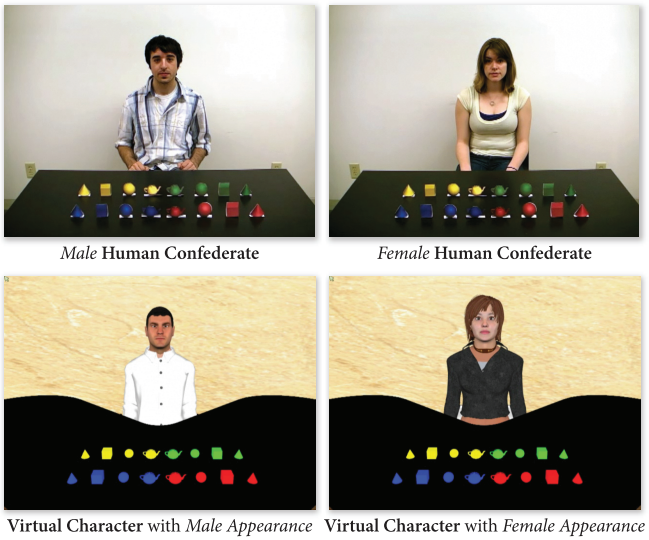

Figure 2: Still images from the videos presented to the participants.

communicative accuracy and *perceived naturalness*. Communicative accuracy was measured by capturing whether participants correctly identified the object toward which the gaze shift of the human confederates or the virtual characters directed. To measure perceived naturalness, we constructed a scale using four items that measured *naturalness*, *human-likeness*, *lifelikeness*, and *realism*. These items were highly correlated (Cronbach's $\alpha = .94$). Participants rated each video for each item using a seven-point rating scale.

4.1 Results

We conducted a mixed-model analysis of variance (ANOVA) to determine the effect that different gaze models had on how accurately participants identified the object that the agents or the human confederates looked toward and the perceived naturalness of the gaze shifts. Overall, model type had a significant effect on communicative accuracy, $F(7, 1958) = 16.77, p < .001$, and perceived naturalness, $F(7, 1958) = 151.03, p < .001$. Detailed comparisons across conditions for each factor are described in the next paragraphs. We used Tukey-Kramer HSD to control the experiment- wise error rate in all post-hoc tests.

Communicative Accuracy – Pairwise comparisons across conditions found no significant differences in the communicative accuracy of the that gaze shifts produced by our model, aggregated across all levels of head alignment, and those produced by human confederates, $F(1, 1958) = 0.03, p = ns$. Similarly, no differences in accuracy were found between our model and the head propensity model, aggregated across all levels of head propensity, $F(1, 1958) = 0.17, p = ns$. The results suggest that the gaze shifts generated by our model are as accurate as those performed by human confederates and those generated by the state-of-the-art model.

Perceived Naturalness – Comparisons across conditions showed that participants rated gaze shifts generated by our model, aggregated across all levels of head alignment, as marginally more natural than those generated by the head propensity model, aggregated across all levels of head propensity, $F(1, 1958) = 3.34, p = .068$. Comparisons over the realism scale (one of the items included in the perceived naturalness scale) found that gaze shifts produced by our model were rated as significantly more realistic than those generated by the head propensity model, $F(1, 1958) = 5.75, p = .017$. Finally, pairwise comparisons across the two models with corresponding head alignment/propensity values showed that, at 100% alignment/propensity, participants rated gaze shifts produced by our model to be significantly more natural than those generated by the state-of-the-art model, $F(1, 1958) = 9.40, p = .002$. Results on the communicative accuracy and perceived naturalness measures are illustrated in Figure 3.

Gender – Participants rated gaze shifts performed by the agent with female features as significantly more natural than those performed by the agent with male features, $F(1, 1958) = 17.17, p < .001$. On the other hand, communicative accuracy of the gaze shifts performed by the agent with male features was significantly higher than that of the shifts performed by the agent with female features, $F(1, 1958) = 4.85, p = .028$. Finally, the analysis found a marginal interaction between participant gender and the gender features of the virtual character, $F(1, 1958) = 3.20, p = .074$. Figure 4 illustrates these results.

5. DISCUSSION

The results from the study show that gaze shifts generated by our model communicate gaze direction as accurately as do human gaze shifts and those generated by the state-of-the-art model. The small gaze targets in our task make it challenging, with performance between 60-70%. However, this is considerably better than the 25% performance expected by chance, and exhibited in the agent control condition.

The results also show that gaze shifts generated by our model are perceived as marginally more natural and significantly more realistic than those generated by the state-of-the-art model. A key advantage of our model is that it provides parametric controls that may be used to achieve communicative outcomes. In a follow-up study we show how the use of the alignment parameter can be used in an educational scenario to provide a tradeoff in the high-level outcomes of affiliation and learning [1]. Finally, while our results suggest a gender effect from the virtual character, the measured difference may be due to any one of a number of differences in the characters' designs.

Our model is focused on the important but specific case of

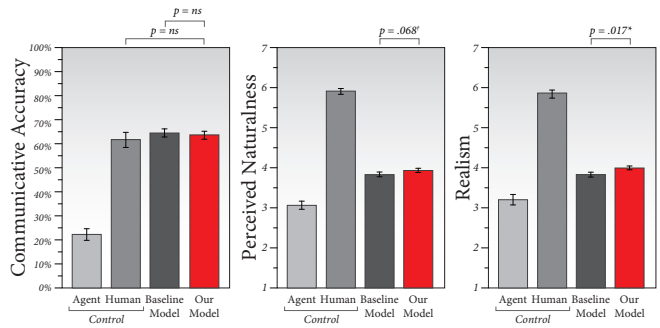

Figure 3: Results from the communicative accuracy, perceived naturalness, and realism measures. The baseline model refers to the head propensity model we used for comparison.

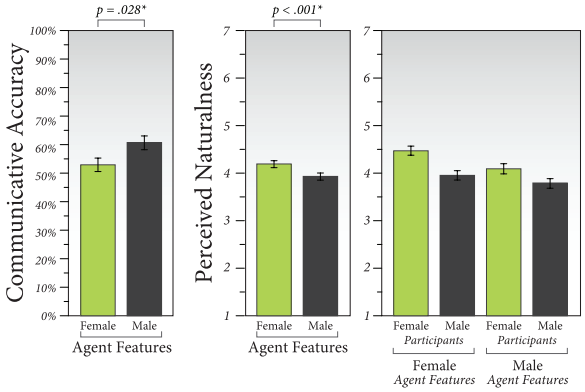

Figure 4: Communicative accuracy and perceived naturalness across agents with female and male features.

head-eye coordination in directed gaze shifts. Future extensions of this model should consider the movement of other body parts, such as employing the neck in performing head movements, as well as other low-level gaze mechanisms such as the opto-kinetic reflex (the saccadic pursuit of gaze targets in motion). Future work should also focus on how to dynamically target different communicative goals, such as communicative accuracy or realism, based on context. While the work presented here explored our gaze shift model on a limited range of characters, we intend to explore its application to different character designs (e.g., stylized cartoons vs. realistic 3D models) and embodiments (e.g., robots vs. on-screen characters).

6. CONCLUSION

Gaze is a complex and powerful nonverbal cue. Subtle changes in how people direct their gaze can help them achieve a wide range of social and communicative goals. In this paper, we have considered a specific and important aspect of gaze, producing gaze shifts. By creating a mechanism for synthesizing gaze shifts in a natural, yet parameterized fashion, we have provided a building block for creating effective social and communicative behaviors.

7. ACKNOWLEDGMENTS

This research was supported by National Science Foundation award 1017952. We would like to thank Allison Terrell, Danielle Albers, and Erin Spannan for their help in creating the videos used in the experiments described in this paper.

8. REFERENCES

[1] S. Andrist, T. Pejsa, B. Mutlu, and M. Gleicher. Designing effective gaze mechanisms for virtual agents. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '12, pages 705–714. ACM, 2012.

[2] A. Cafaro, R. Gaito, and H. Vilhjálmsson. Animating idle gaze in public places. In *Intelligent Virtual Agents*, pages 250–256. Springer, 2009.

[3] Z. Deng, J. Lewis, and U. Neumann. Automated eye motion using texture synthesis. *IEEE Computer Graphics and Applications*, pages 24–30, 2005.

[4] E. Freedman and D. Sparks. Activity of cells in the deeper layers of the superior colliculus of the rhesus monkey: evidence for a gaze displacement command. *Journal of neurophysiology*, 78(3):1669–90, 1997.

[5] E. Freedman and D. Sparks. Coordination of the eyes and head: movement kinematics. *Experimental Brain Research*, 131(1):22–32, 2000.

[6] A. Fukayama, T. Ohno, N. Mukawa, M. Sawaki, and N. Hagita. Messages embedded in gaze of interface agents—impression management with agent's gaze. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 41–48. ACM, 2002.

[7] J. Fuller. Head movement propensity. *Experimental Brain Research*, 92(1):152–164, 1992.

[8] J. Goldring, M. Dorris, B. Corneil, P. Ballantyne, and D. Munoz. Combined eye-head gaze shifts to visual and auditory targets in humans. *Experimental brain research*, 111(1):68–78, 1996.

[9] H. Goossens and A. Opstal. Human eye-head coordination in two dimensions under different sensorimotor conditions. *Experimental Brain Research*, 114(3):542–560, 1997.

[10] D. Guitton and M. Volle. Gaze control in humans: eye-head coordination during orienting movements to targets within and beyond the oculomotor range. *Journal of neurophysiology*, 58(3):427–59, 1987.

[11] R. Heck. *Automated authoring of quality human motion for interactive environments*. PhD thesis, University of Wisconsin–Madison, 2007.

[12] L. Itti, N. Dhavale, and F. Pighin. Photorealistic attention-based gaze animation. In *2006 IEEE International Conference on Multimedia and Expo*, pages 521–524. IEEE, 2006.

[13] S. Khullar and N. Badler. Where to look? automating attending behaviors of virtual human characters. *Autonomous Agents and Multi-Agent Systems*, 4(1):9–23, 2001.

[14] K. Kim, R. Brent Gillespie, and B. Martin. Head movement control in visually guided tasks: Postural goal and optimality. *Computers in Biology and Medicine*, 37(7):1009–1019, 2007.

[15] A. Kittur, E. Chi, and B. Suh. Crowdsourcing user studies with Mechanical Turk. In *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 453–456. ACM, 2008.

[16] B. Lance and S. Marsella. The expressive gaze model: Using gaze to express emotion. *Computer Graphics and Applications, IEEE*, 30(4):62–73, 2010.

[17] J. Lee, S. Marsella, D. Traum, J. Gratch, and B. Lance. The rickel gaze model: A window on the mind of a virtual human. In *Intelligent Virtual Agents*, pages 296–303. Springer, 2007.

[18] S. Lee, J. Badler, and N. Badler. Eyes alive. In *ACM Transactions on Graphics (TOG)*, volume 21, pages 637–644. ACM, 2002.

[19] M. McCluskey and K. Cullen. Eye, head, and body coordination during large gaze shifts in rhesus monkeys: movement kinematics and the influence of posture. *Journal of neurophysiology*, 97(4):2976–91, 2007.

[20] B. Mutlu, J. Forlizzi, and J. Hodgins. A storytelling robot: Modeling and evaluation of human-like gaze behavior. In *Humanoid Robots, 2006 6th IEEE-RAS International Conference on*, pages 518–523. IEEE, 2006.

[21] J. Pelz, M. Hayhoe, and R. Loeber. The coordination of eye, head, and hand movements in a natural task. *Experimental Brain Research*, 139(3):266–277, 2001.

[22] C. Peters and A. Qureshi. A head movement propensity model for animating gaze shifts and blinks of virtual characters. *Comput. Graph.*, 34(6):677–687, 2010.

[23] I. Poggi, C. Pelachaud, and F. De Rosis. Eye communication in a conversational 3d synthetic agent. *AI communications*, 13(3):169–181, 2000.

[24] W. Zangemeister and L. Stark. Types of gaze movement: variable interactions of eye and head movements. *Experimental Neurology*, 77(3):563–577, 1982.