

Towards Expressive Gaze Manner in Embodied Virtual Agents

Brent Lance

Stacy Marsella

David Koizumi

University of Southern California

Information Sciences Institute

4676 Admiralty Way, Suite 1001

Marina del Rey, CA 90292-6695

(310) 448-9124

brent@isi.edu

Abstract

Empathy can be viewed in largely cognitive terms as the ability to do role taking, specifically to perceive, imagine and take on the psychological point of view of another [Piaget, 1965]. It also can be seen in more affective terms, specifically to react emotionally to another's emotional state. [e.g., Stotland, 1978]. In either case, revealing a person's, or a virtual character's, internal state is an important aspect of inducing an empathic response in another. This paper considers the problem of designing a model of expressive gaze manner for virtual characters. By expressive gaze manner, we mean how what a person is thinking and feeling is conveyed through the physical manner by which they gaze. For example, angry glares, gapes, stares, furtive glances and peeks are different in their physical properties and in what they reveal about the person gazing. A model of expressive gaze needs to describe how to change the physical properties of a character's gaze shift in order to exploit this ability to express, and not just describe when and where the agent should gaze.

Ultimately, the purpose of this model is to provide Embodied Conversational Agents (ECA) with expressive gaze. This paper describes an exploratory study that is a first step in collecting the data required to build a model of expressive gaze. We extracted gaze data from Computer Graphic (CG) animated motion pictures and performed a preliminary analysis of a portion of this data. Animated films are of particular interest here given the animator's skill in creating obviously artificial characters that nevertheless evoke emotion and empathy in the audience.

1. Introduction

The process of establishing an empathic relation between two people is often intimately tied up with their gaze. Empathy requires the ability to identify and understand another person's internal state, including their emotions and their motivations [Eisenberg et al. 2003]. Thus, if we want a virtual character to induce empathy in the user, we must explore appropriate ways for the character to reveal its internal state¹. Conversely, for a virtual character's behavior to suggest empathy for another, the behavior must suggest an attentiveness to the other's internal state and situation. In either case, gaze plays an important role in both a character *inducing* empathy in another and in suggesting a character *has* empathy for another. For gaze both reveals emotion as well as signals the gathering of emotional information and overall social involvement.

This paper addresses the problem of designing a robust general model of expressive gaze manner in virtual characters. By "a model of expressive gaze" we mean how the physical properties of a character's gaze shift change in order to reveal its internal state. Such a model of gaze is unique in the sense that it focuses on how the physical properties of individual gaze shifts should be implemented, in addition to when and where the agent should gaze. In order to develop this model, we need to determine the mapping between the internal and external factors that affect the gaze, and the physical properties of the gaze.

¹ Brown [Brown, 1986, p.190-191] argues that the salience of a person's expressive behavior will focus an observer's interpretation on the person's psychological perspective.

In the discussion that follows, we describe our first attempt at collecting the data required to determine this mapping, as well as a preliminary analysis of this data. We based our work on both the psychology literature, which provides a theoretical framework, and Computer Graphic (CG) animated motion pictures, which provides data on gaze manner. We discuss our approach and also discuss interesting patterns our preliminary analysis has revealed in the data, such as how the animators used neck and back angles to reveal emotional state.

2. Background

Gaze serves a variety of functions [Argyle, 1976; Kendon 1990]. It is used to gather information. It is used to regulate the information gathered (for example to manage cognitive load). It is used to regulate conversations, especially by managing the turn-taking between participants. Gaze also serves to signal interest and express emotional state. In particular, the manner in which any gaze is performed can be revealing. Slowly turning the head to gaze at some event may reflect a lethargic or depressed state while the rapid movement of head and body toward an event may suggest an alert, surprised state. Thus, we see every gaze from the perspective that it's manner is expressive. And, an Embodied Conversational Agent (ECA) needs the ability to perform expressive gazes [Rickel 1999]. This way, the user can obtain important information, such as the emotional state of the agent, or feedback on the user's actions, from the way that it gazes. However, this also means that errors in the gaze model will cause the user to attribute the wrong mental state to the agent.

Since ECA's are designed to interact with users in real-time, any gaze models will also have to function in real-time and in response to the context of the agent's internal state and external situation. While a great deal of success has been achieved in generating specifically crafted nonverbal communicative behavior in pre-rendered non-interactive environments such as CG animated films, this success has not carried over into real-time interaction. The reason for this is that the creation of expressive nonverbal behavior by an animator has traditionally been both time consuming process and tailored to a specific context. This process also requires a great deal of craft, where the animator works until a behavior "looks right." Specifically, the animator can iteratively examine and modify the animation as necessary until it achieves the desired effect. For an animation generated in real time by an ECA, there is no such interpretive oversight or possibility for modification. Although this craft has been expressed at a

high level [Thomas, 1981], it has not been detailed to the level of the physical, dynamic properties of the gaze necessary to automate an ECA's expressive gaze.

Unfortunately, the psychology literature on human interactive gaze also has shortcomings that make it impossible to extract this mapping directly from data in the literature. Much of the relevant information in the psychology literature is about when and where to gaze, and there is not a great deal of information about how specific gaze shifts should be executed. For instance, we were unable to find specific details on the velocities of gaze shifts, and how these velocities change as the internal state of the character does. We were also unable to locate data describing the relationship of the head and back velocities to the velocity of the eye movement. While there are several psychological sources that discuss the speed of gaze shifts, they tend to do so within very highly constrained situations, and no attempt is made to link the speed of the shifts with internal states, such as emotion. This paper describes an initial attempt to collect the necessary data to determine a mapping between the way gazes are executed, and the way that they are interpreted.

After careful consideration of the options, we chose to collect the required gaze data from CG animated motion pictures. CG animation was chosen over videos of humans for several reasons. In animation, all signals are intentionally designed by the animator to convey meaning to the viewer. In fact, one of the main benefits of using animation over humans is that the gaze shifts in animation are explicitly crafted to be understood by the audience. Every gaze shift has intentional communicative value. Additionally, while functioning under the somewhat similar constraints of virtual characters in a virtual environment, the animators develop extremely colorful characters. In fact, animators readily get us to empathize with their celluloid artifacts. Some of these films have won Oscar Awards, due in no small part to the ability of the animators to make the characters in the films expressive and believable. Additionally, these characters showcase a broader range of emotion than can usually be captured from filming discussions between human subjects. Finally, studying how problems with expressive gaze were solved for one domain may reveal information applicable to the other.

Thus, we decided to extract the information we needed from CG animated motion pictures instead of from videos of interacting human beings. In particular, we chose to explore *Toy Story 2* and *Antz* because they were major releases with fairly human-like characters.² These films are from different studios in order to explore different animation styles.

After deciding on the films, we selected a number of scenes from each film of characters interacting. Two scenes were chosen from each film. The scenes were selected to give us a range of character's acting in several different situations. Within each film, one character appeared in both scenes. This was done to determine whether or not the characters' personality would have a greater effect on their gaze shifts than the circumstances they were in. From these scenes, we collected a number of statistics, including head and body angular velocity, dwell time, and eye, head and body angle away from target, for each gaze shift. We also attempted to determine the environmental factors which either cause or affect individual gaze shifts, or the overall sequence of gaze behavior. Finally, we performed statistical analyses on this data.

The rest of this paper provides greater detail about the approach that we used. Section 3 examines related work from both the Computer Science and Psychology communities. Section 4 describes the data collection approach, detailing how we mined the necessary data from the animated motion pictures. Section 5 will discuss the factors, the statistics collected, and interesting results obtained from the data so far. Finally, our future plans for this model will be given in Section 6.

3. Related Work

There have been several attempts to build a robust gaze model for ECA's and other virtual humans. [Chopra-Khullar, 1999] is one example of a model of visual attentive gaze. This system uses an arbitration system that associates a set of motor activity primitives with a set of gaze behavior primitives. Two queues of potential gaze targets are sent to the gaze manager. One of these queues holds objects that require conscious attention, while the other holds immediate attention capture targets. The gaze manager arbitrates between the two queues, and executes a gaze shift using the proper gaze behavior primitive for the type of shift. If both queues are empty, a spontaneous looking behavior is

generated, in which attention is drawn to novel or complex items in the environment.

A similar system to the one described above is detailed in [Gillies, 2002]. This system also uses a gaze manager system that arbitrates between several queues of possible gaze targets. The algorithm arbitrates between two queues, one holding immediate requests, and the other holding monitoring requests. It produces an undirected attention behavior in the event that both queues are empty. The main difference between this algorithm and the previous one is that this one does not appear to use gaze behavior primitives to the same extent that the previous system does. Instead, this system has the gaze manager generate values for properties such as the length of the gaze and the interval between gazes in real-time, allowing for such properties to be controlled to a greater extent by the controller that generates gaze requests.

A different approach can be found in [Peters, 2003]. A four-part bottom-up attention model is implemented in this system. It utilizes a saliency model with an inhibition of return response to filter renderings of the scene from the point of view of an agent within the scene. In addition, this system uses a memory model to keep a record of what and where objects are in the environment. The gaze generation process takes all of this into account in order to generate the final gaze animation.

A gaze model for conversational agents can be found in [Pelachaud, 2003]. This algorithm combines a communicative function model along with a statistical model. The communicative model computes the gaze to convey specific meanings. The statistical model was constructed from gaze data collected from interactions between two individuals, and is used to compensate for missing factors in the communicative function model. In addition, the paper introduces temporal gaze parameters as a method of personalizing gaze shifts to characters.

However, our approach to gaze model generation differs from the above studies in several ways. Most of these models generate task-related attentive behavior, instead of communicative gaze behavior. The model that does generate communicative gaze behavior is similar to our work, but uses a statistical model to compensate for "missing factors" such as social aspects, whereas we are attempting to gather a wide range of contextual factors. Further, the emphasis is more on high level gaze behavior as opposed to the low level properties of gaze manner we are investigating (such as neck angle, back angle, velocity, etc.).

Each of these gaze models, and ours as well, draws heavily from psychological gaze studies, such as [Duncan, 1977; Argyle, 1976; Kendon, 1990]. These

² The use of animated films does raise a concern about non-human bodies and the challenge of generalizing gaze shifts across body types

Head/Neck Variables		
Head Vertical Angle	Head Vertical Velocity	Head Vertical Angle From Target
Head Horizontal Velocity	Head Horizontal Angle	Head Horizontal Angle From Target

Back/Body Variables		
Back Vertical Angle	Back Vertical Velocity	Back Vertical Angle From Target
Back Horizontal Velocity	Back Horizontal Angle	Back Horizontal Angle From Target

Eye Variables	
Eye Vertical Angle	Eye Vertical Angle From Target
Eye Horizontal Angle	Eye Horizontal Angle From Target
Dwell Time	

Figure 1: Eye/Neck/Back Physical Properties

studies capture gaze data from the interaction between a pair of individuals. One such study is described in [Kendon, 1990]. In this reference, a study where data is collected from the videotaped interaction between two college students is described. The paper mostly discusses aggregate data, such as the proportion of time spent gazing at the other individual during various circumstances, like when speaking or listening. The mean lengths of the gazes observed in these circumstances are also given. [Kendon, 1990] also postulates several theories that were useful in deciding what contextual features to examine in our study. One regards the usage of gaze at the beginning and ending of spoken phrases. Another discusses the likelihood of using mutual gaze to keep a steady level of emotional arousal. The information from this resource caused us to examine various speaking and listening roles, as well as emotional factors in our study.

The study done in [Duncan, 1977] contains a great deal of information on many different types of nonverbal communication. There were two sections that were of primary use to us. The first gave figures such as number of gazes, rate of gaze, gaze length, and percent of time spent gazing, both for speaking and listening. The second described correlations found between these variables and between other nonverbal communicative behaviors, giving us several other contextual factors to consider when examining our data.

One of the primary works describing human communicative gaze is [Argyle, 1976]. This is a comprehensive review of human communicative gaze. It covers a variety of topics from pupil response and

blinking, to the amount of gaze and mutual gaze during conversation, to the difference in gaze between normal individuals and individuals with psychiatric disorders. The theories and predictions found in this text provide a foundation for the work described in below.

4. Data Collection Overview

We extracted the data from two CG animated feature films, Toy Story 2 and Antz. We chose two scenes from each film. In Antz we chose a scene (Scene 1) near the beginning where the characters Cutter and Mandible discuss Mandible's vision for the colony. This scene runs from 00:07:00 to 00:08:00 in the movie. In Toy Story 2, the first scene (Scene 2) is about halfway through the film, and shows Woody (the main character) attempting to steal his missing arm from Al's pocket. The second scene from Antz (Scene 3) is closer to the end, where Cutter intimidates some bugs at Insectopia, before talking with another character, Bala, about returning to the hive. The second Toy Story 2 scene (Scene 4) is shortly after the first, and is an argument between the characters Woody, Pete, and Jessie when Woody finds out that they're going to Japan. Each of these four scenes contains very different types of gaze shifts, including dialog, task, and monologue related shifts. In addition, a number of the scenes share characters, allowing us to examine how an individual character's gaze patterns change as the circumstances they are in change.

Gaze Attraction
Gaze Aversion
Emotional
Task Related
Speech Related
Social Status

Figure 2: Subset of Contextual Factors

We collected data that described the velocity, angle, and time of gaze shifts for all of the various body parts in a gaze shift (such as the eyes, neck, and back). We also considered the angles that these joints turned through, and the angle that they faced away from the target at the end of the shift to be important values. Finally, we attempted to collect the contextual factors that affect each shift. These different data points can be seen in Figure 1. This set of physical properties is more detailed than available in the psychology resources. In order to collect all of these measurements, multiple passes were made through the scene.

The first pass collected data on the time it took to do the gaze shift and the angle that the shift swept through for the eyes, neck and back. Angles were described as a vertical and horizontal component as viewed from a local coordinate system inside of the character. All angular measurements were based on estimates derived from visual examination of the character's gaze shifts. This approach has been used in several psychological studies of gaze [Kendon, 1990, Argyle, 1976] (in particular see [Argyle, 1976], pp 35-57, for a discussion of gaze measurement). Once we had both the angle that each body part had turned through, and the time that it took to turn, the various velocities were also calculated.

We then collected a second set of angular information. We referred to these as the angle-from-target, or AFT measurements. These measurements represent the distance in degrees between the current direction of a specific joint, and the final target of the gaze shift. The AFT's for the eyes, head, and body were estimated by examining how the character faced the target of the gaze shift at the end of the shift. These measurements allow us to easily differentiate between a gaze where a character turns completely towards the target of the shift, and one where the shift is a subtle motion of the eyes towards the target.

The final pass for numerical data was used to collect the dwell times of the gaze shifts. The dwell time of the shift was defined as the length of time that the character gazed at a specific target, from when the

character's eyes first encountered the target, to when the character's eyes were no longer on the target. If the character ceased to be on camera before the gaze had ended, it was not given a dwell time measurement. We did not make the assumption that a character would continue to gaze at the target once the character was off-screen. This limited the number of shifts that had a dwell time value associated with them, but we felt that it was important to only record data that could be visibly verified, and did not need to be assumed.

After the numerical data was collected, one last pass through the scene was made. During this pass, each individual gaze shift was associated with a number of contextual features that could have influenced the properties of the gaze shift. Examples include who is currently speaking, what physical tasks are performed simultaneously with the gaze, and the emotional status of the character performing the gaze. Figure 2 provides a list of the highest level categories. The rest of the factors were more specific instances within these categories. The "Gaze Attraction" and "Gaze Aversion" categories refer to whether or not a specific gaze shift was towards a given target, or away from a given target, and are not related to emotional theories of attraction or aversion. Additionally, the "Social Status" category refers to whether or not a character is of a higher or lower social status than any other characters they may be conversing with. The "Emotional" category contained possible emotional states for a character to be in. The "Task Related" and "Speech Related" categories contained various factors such as whether or not a character was physically manipulating an object, or whether the gaze shift occurred at the beginning or ending of a spoken phrase. For a factor to be considered, it had to either occur within +/-0.5 seconds of the gaze shift (based simply on using [Kendon 1990] to provide a rough guideline) or be a persistent situational factor (such as the social status between speaker and addressee).

The list of factors was distilled from two places: first, a preliminary list of potential factors was assembled from various sources in the psychological literature [Kendon, 1990, Argyle & Cook, 1976; Duncan & Fiske,

Factor	Approach/Avoid
Anger	5
Deceit	-3
Disagreement	2
Emphasis	1
Fear	-5
Joy	2
Persuasion	4
Reproach	2
Sorry For	3
Surprise	-3

Figure 3: Latent Variable Mapping

1977], as well as gaze models and emotional models from the computer science literature. Second, this preliminary list was used in conjunction with the film, and any necessary new factors that appeared in the film and were not on the list were added.

5.1. Preliminary Analysis and Results

The data set collected to date contains a number of limitations. Because this was an exploratory investigation into what physical properties would turn out to be effective predictors of what internal states, we collected many physical properties and contextual factors. Indeed, the data set has more variables (236) than cases of gaze shifts (169). The data set is also very sparse. Data collection was complicated by the extremely time consuming and labor intensive approach used. We must emphasize, however, that this was merely an exploratory study, meant to test the utility of the data collection approach and drive further research. In addition, while this exploratory study was insufficient to generate a full model of expressive gaze manner, there were interesting data subsets that could be analyzed. We discuss some of that analysis in this section. The knowledge acquired in this exploratory work should also serve us well in collecting additional data, as discussed in section 6.

One of the issues we were interested in exploring was the relationship between the expression of emotion and gaze. We labeled each gaze shift with an emotional state (if possible). The emotion categories that we used were based on the work of Elliot [1992]. Following this, we used the emotion categories, and a few other applicable categories to generate a latent

variable called Approach/Avoidance. This latent variable represents whether the character gazing desires a closer or more distant association with the target of the gaze. (A latent variable is an unobserved variable, usually used to explore the relationships between observed variables. Here we are specifically using it to arrange emotion categorical variables into an action tendency scale of approach versus avoidance.) The categories used to generate this latent variable can be seen in Figure 3. The variables are combined through simple summation. Using this derived approach/avoidance variable, we ran regressions against the vertical angles that the neck and back joints turn through. The results can be viewed in Figure 4.

Note that back vertical angle and neck vertical angle both significantly correlate with approach-avoidance. However, back vertical angle is negatively correlated while neck vertical angle is positively correlated.

5.2 Discussion

It is somewhat surprising that Back Vertical Angle is negatively correlated with approach-avoidance while Neck Vertical Angle is positively correlated. One might expect that they would both be positively correlated – that approach would, overall, be associated with movement towards the gaze target. However, a little thought suggests what this reverse correlation might achieve. The neck vertical angle is directly revealing approach and avoidance. However, positive back angles will tend to reduce the height and apparent size of the character, as if they are hiding or reducing their exposure. More negative vertical angles on the other

Pearson Correlation Matrix

	<i>Approach/Avoidance</i>	<i>Neck Vertical Angle</i>	<i>Back Vertical Angle</i>
Approach/Avoidance	1.000		
Neck Vertical Angle	0.231	1.000	
Back Vertical Angle	0.017	0.892	1.000

Tests of Correlation Significance

<i>Variable</i>	<i>Coefficient</i>	<i>Std. Error</i>	<i>Std. Coefficient</i>	<i>Tolerance</i>	<i>T</i>	<i>P (Two Tail)</i>
Constant	-2.002	1.144	0.000	0.000	-1.750	0.131
Neck Vertical Angle (NVA)	0.152	0.055	1.080	0.204	2.740	0.034
Back Vertical Angle (BVA)	-0.275	0.103	-1.059	0.202	-2.680	0.037
NVA *	0.014	0.003	0.770	0.978	4.275	0.005
BVA						

Analysis of Variance

<i>Source</i>	<i>Sum of Squares</i>	<i>DF</i>	<i>Mean Square</i>	<i>F-Ratio</i>	<i>P</i>
Regression	88.173	3	29.39	8.508	0.014
Residual	20.727	6	3.454		

Figure 4: Regression Results

hand will tend increase the height of the character and may even tend to create an effect of “looming over” the target, much like animals in combat tend to increase their height and size. Of course, as we collect and analyze more data, an important aspect of model validation will be to corroborate such analyses with animators. Another possible approach would have been to speak with the animators first to discuss their high-level design approach and then ideally obtain the low-level gaze data directly from the animation files. However, when we explored this possibility, we discovered that this resource was not available to us during the course of this research.

5.3 Additional Observations

While performing the data analysis, we noticed a number of interesting observations that resulted from collecting the data from animated features. One of the primary observations that we made is that the camera usually focuses on the speaker during any sort of interaction between characters. Often the listener is not even in frame. When the camera does cut to the listener it is in order to get a reaction to what was just said. This skews our data towards speakers and reactions, and away from maintenance gazes performed by listeners [Argyle, 1976]. The camera movement also interfered with our collection of the dwell time statistic, which describes how long a character gazes at a target. The camera would often cut away from a character while he was gazing, and it was impossible to know if they continued to gaze at that target.

Another interesting observation is the slight discrepancies we noticed between data we observed in the films, and data recorded from actual humans in the psychology literature. For example, [Kendon, 1990] describes the humans he observed often made a gaze aversion before they began an utterance. We did not find this to be the case. In fact, when making a gaze shift at the beginning of an utterance, it is a gaze attraction 81% of the time. One possible explanation is that, since the functional use of gaze for regulating interaction, such as turn-taking, is not strictly required, the animator is ensuring that the audience clearly understands to whom the speaker and listener roles are assigned. Of course, at this point we have insufficient data to create a full model. In the conclusion, we discuss how we plan to go beyond this limited data set.

6. Future Work and Conclusion

The craft of animators and actors is a rich source of data for the design of ECAs that can evoke our emotion and empathy. Although there are many high level, rough guidelines of this craft, it has not been explicitly detailed down to the level that ECA designers need to animate their characters, no doubt in large measure because the artist's perception and interpretation of the artifact she is creating is an integral aspect of the iterative creative process. Nevertheless, to exploit this resource, we need to go beyond the high level guidelines. In this paper, we presented an approach that uses animated films to extract a model of expressive gaze manner and discussed preliminary results of an exploratory data collection and analysis. Although the analysis is very preliminary, we are proceeding to implement the partial approach/avoidance model discussed here using our virtual humans [Rickel, 2002] and then will evaluate its behavior. Results from the analysis done so far is also helping us focus our future data collection efforts in a more efficient manner.

Moving forward, we are exploring less labor-intensive, more automated, approaches to collect data. As part of that effort, we plan on extending our data to include standard films. The goal of future data collection and analysis will be to collect sufficient data to create a full computational model of expressive gaze manner. We will then proceed to validate it in two ways. We will validate the model's characteristics with animators and actors. Second, we will use the model to control virtual human characters and have subjects evaluate the resulting behavior.

Acknowledgements: This work was started by the first author and the late Dr. Jeff Rickel. The authors are deeply indebted to Dr. Rickel for his ideas and guidance in this research. Special thanks to Dasarathi Sampath for his help.

References:

Argyle, M.; M. Cook. *Gaze and Mutual Gaze*. Cambridge University Press, New York, NY, 1976.

Bentler, P. Multivariate Analysis with Latent Variables: Causal Modeling. *Annual Review of Psychology* 31:419-456, 1980.

Brown, R. *Social Psychology: The Second Edition*. The Free Press, New York.

Chopra-Khullar, S.; N.I. Badler, Where to Look? Automating Attending Behaviors of Virtual Human Characters, Proceedings of the Third Annual Conference on Autonomous Agents, 16-23, Seattle, Washington., 1999.

Duncan, S. Jr.; D.W. Fiske, *Face to Face Interaction: Research, Methods, and Theory*. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1977.

Eisenberg, N., Losoya, S. & Spinrad, T. Affect and Prosocial Responding. In Davidson, Scherer and Goldsmith, *Handbook of Affective Sciences*. Oxford University Press, 2003, pp. 787-803.

Geiselman, R.; J. Woodward, and J. Beatty, Individual Differences in Verbal Memory Performance: A Test of Alternative Information-Processing Models. *Journal of Experimental Psychology: General*, Vol. 111, No. 1, 109-134, 1982.

Elliott, C. "The Affective Reasoner: A Process Model of Emotions in a Multi-agent System". Ph.D. dissertation, Northwestern University. The Institute for the Learning Sciences Technical Report No. 32, 1992

Gillies, M.F.P.; N.A. Dodgson, Eye Movements and Attention for Behavioral Animation. *The Journal of Visualization and Computer Animation*; 13: 287-300, 2002.

Kendon, A. *Conducting interaction: patterns of behavior in focused encounters*. Cambridge University Press, 1990.

Pelachaud, C.; M. Bilvi, Modelling Gaze Behavior for Conversational Agents. Proceedings of the Intelligent Virtual Agents 4th International Workshop, Kloster Irsee, Germany, 93-100, 2003.

Peters, C.; C. O'Sullivan, Bottom-Up Visual Attention for Virtual Human Animation, Proceedings of Computer Animation and Social Agents (CASA 2003), Rutgers University. 2003.

Piaget, J. *The moral judgment of the child*. Glencoe, IL: Free Press, 1965.

Stotland, E., Mathews, K., Sherman, S., Hansson, R., & Richardson, B. *Empathy, fantasy, and helping*. Beverly Hills: Sage Publications, 1978.

Rickel, J.; W.L. Johnson. Extending Virtual Humans to Support Team Training in Virtual Reality. In G. Lakemeyer and B. Nebel (Eds.), *Exploring Artificial Intelligence in the New Millennium*. San Francisco: Morgan Kaufmann Publishers, 2002, pp. 217-238. (This book collects extended versions of award-winning papers.)

Rickel, J.; W.L. Johnson. Animated Agents for Procedural Training in Virtual Reality: Perception, Cognition, and Motor Control. *Applied Artificial Intelligence* 13:343-382, 1999. (Special issue on Animated Interface Agents.)

Rogosca, D. R. Causal Models in Longitudinal Research: Rationale, Formulation, and Interpretation. In *Longitudinal Methodology in the Study of Behavior and Development*, J. R. Nesselroade and P. B. Baltes, Eds. New York: Academic Press, 263-302, 1979.

Thomas, F.; O. Johnston. *Disney animation: the Illusion of Life*. Abbeville Press, 1981.