

UC Berkeley

UC Berkeley Previously Published Works

Title

Do geometric models affect judgements of human motion?

Permalink

<https://escholarship.org/uc/item/8zj4x1tj>

Authors

Hodgins, JK

OBrien, JF

Tumblin, J

Publication Date

1997

Peer reviewed

Do Geometric Models Affect Judgments of Human Motion?

Jessica K. Hodgins, James F. O'Brien, and Jack Tumblin

College of Computing and Graphics, Visualization, and Usability Center
Georgia Institute of Technology
Atlanta, GA 30332-0280
[jkh|obrienj|ccsupjt]@cc.gatech.edu

Abstract

Human figures have been animated using a wide variety of geometric models including stick figures, polygonal models, and NURBS-based models with muscles, flexible skin, or clothing. This paper reports on experiments designed to ascertain whether a viewer's perception of motion characteristics is affected by the geometric model used for rendering. Subjects were shown a series of paired motion sequences and asked if the two motions in each pair were "the same" or "different." The two motion sequences in each pair used the same geometric model. For each trial, the pairs of motion sequences were grouped into two sets where one set was rendered with a stick figure model and the other set was rendered with a polygonal model. Sensitivity measures for each trial indicate that for these sequences subjects were better able to discriminate motion variations with the polygonal model than with the stick figure model.

Keywords: motion perception, motion sensitivity, perceptual study, computer animation, geometric model.

1 Introduction

Few movements are as familiar and recognizable as human walking and running. Almost any collection of dots, lines, or shapes attached to an unseen walking figure is quickly identified and understood as human. Studies in human perception have displayed walking motion using only dots of light located at the joints and have found test subjects quite adept at assessing the nature of the underlying motion[13].

From the proceedings of *Graphics Interface '97*, pages 17–25. Sponsored by the Canadian Human–Computer Communications Society. Held May 21–23, 1997, in Kelowna, B.C., Canada.

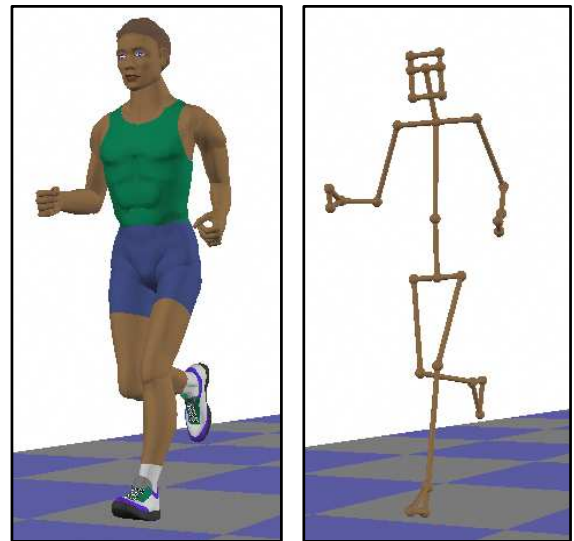


Figure 1: Image of an animated runner rendered with a polygonal model (left) and with a stick figure model (right). These images are typical of those used in this study.

In particular, subjects could identify the gender of a walker and recognize specific individuals even when no other cues were available[6, 16, 17].

In part because people are skilled at detecting subtleties in human motion, the animation of human figures has long been regarded as an important but difficult problem in computer animation. Recent publications have presented a variety of techniques for creating animations of human motion. Promising approaches include techniques for manipulating keyframed or motion capture data[29, 28, 4, 25], control systems for dynamic simulations[11, 18], and other procedural or hybrid approaches[1, 3, 15, 19, 5, 21]. Each method has its own strengths and weaknesses, making the visual comparison of results essential, especially for the evaluation of such subjective qualities as "natu-

ralness” and “realism.”

Our ability to make judgments about human motion from displays as rudimentary as dot patterns raises an important question: does the geometric model affect the viewer’s judgment of the motion or can the viewer make accurate judgments independent of the model used for rendering? There are three plausible but contradictory answers to this question.

Possibility 1. Simple representations allow finer distinctions when judging human motion. Simpler models may be easier to comprehend than more complex ones, allowing the viewer’s attention to be more completely applied to the details of the movement rather than the details of the model. A stick figure is an obvious abstraction and rendering flaws may be easily ignored. When more detailed models are used, subtle flaws in rendering, body shape, posture, or expression may draw attention away from the movements themselves.

Possibility 2. Complex, accurate representations allow finer distinctions. People have far more experience judging the position and movement of actual human shapes than they do judging more abstract representations such as stick figures. A viewer, therefore, may be able to make finer distinctions when assessing the motion of more humanlike representations. Furthermore, complex representations provide more features to identify and track. Each body segment in a polygonal human model has a distinctive, familiar shape, thereby making it easier to gauge fine variations in both position and rotation.

Possibility 3. Both simple and complex representations allow equally fine distinctions. The human visual system may use a displayed image only to maintain the positions of a three-dimensional mental representation. Judgments about the motion may be made from this mental representation rather than directly from the viewed image. Displayed images must of course supply enough cues to keep the mental representation accurate, but additional detail and accuracy may be irrelevant. Just as joint positions shown by light dots are sufficient to control the mental representation, connecting the dots with a stick figure might not improve the viewer’s perception. Similarly, encasing a stick figure within a detailed human body shape might likewise prove unnecessary.




Figure 2: The dot pattern on the left shows the joint locations of a human runner at a single point in time. On the right, these joint locations are shown over the course of one step in the running cycle. Although it is difficult to determine the nature of these patterns from a still image, studies show that most people are able to recognize the motion, and even make fine judgments when shown moving sequences of similar images.

Objective evidence is needed to determine which of these three possibilities is correct. We argue that definitive experiments to select between possibilities 1 and 2 are impractical. The question of which style of geometric model is more useful for judging motion is likely to be context dependent and endlessly complex, affected by all of the variables of both the motion and the rendering. If possibility 3 were correct, and model style were largely irrelevant, then we would be able to perform critical comparisons of motion synthesis techniques using substantially different geometric models. This paper provides experimental evidence to disprove possibility 3 by showing that viewer sensitivities to motion variation are significantly different for the stick figure model and the polygonal model shown in figure 1. In particular, for the types of motion variation we tested, the viewers were more sensitive to motion changes when displayed through the polygonal model than they were with the stick figure model. We discuss the implications of this result in section 5.

2 Background

Several researchers have used light-dot displays to study perception of human movements and to investigate the possibility of dynamic mental models[10]. The light-dot displays showed only dots or patches of light that moved with the main joints of walking figures (figure 2), but even these minimal cues were sufficient for viewers to make detailed assessments of both the motion and the nature of the fig-

ure. In related work, light dots placed on the fingertips and hands were sufficient for skilled readers of American Sign Language(ASL) to read some signed messages[22].

The ability to perceive human gaits from light-dot displays has been widely reported to be acute and robust. Early experiments by Johansson[13] reported that 10-12 light dots “evoke a compelling impression of human walking, running, dancing, etc.” Because such displays provide motion cues independent of form or outline, other investigations have used them to study human motion perception. Work by Cutting and Kozlowski[6] showed that viewers easily recognized friends by their walking gaits on light-dot displays. They also reported that the gender of unfamiliar walkers was readily identifiable, even after the number of lights had been reduced to just two located on the ankles[16]. In a published note, they later explained that the two light-dot decisions were probably attributable to stride length[17]. Continuing this work, Barclay, Cutting, and Kozlowski[2] showed that gender recognition based on walking gaits required between 1.6 and 2.7 seconds of display, or about two step cycles. Our experiments used pairs of 4 second stimuli displaying about six step cycles, but we noticed that test subjects often marked their answer sheets near the midpoint of the second stimuli which is consistent with Barclay’s results.

Motion is apparently essential for identifying human figures on light-dot displays. The Cutting studies reported that while still light-dot displays were not recognized as human, moving light-dot displays of a walking figure were recognized immediately. Poizner and his colleagues[22] also noted that movement is required for accurately reading reading ASL gestures.

This capacity to recognize moving figures was shown to be robust in the presence of masking by additional light points. In a modified experiment, subjects were shown light-dot displays of walkers facing either left or right, and asked to determine walking direction. Only very complex masks of extraneous light dots moving in patterns that were similar to those of the walking figure were able to disrupt the viewer’s judgment[7].

Although human walking movements are accurately sensed from light-dot displays, synthetic human movements are easily accepted as human. Using measurements from light-dot displays, Cutting and colleagues[8] found that apparent torso structure and rotation were strongly correlated with judgments of walker gender. Cutting then constructed a

simple mathematical model of light-dot motion for human walkers and computed displays of synthetic walkers. Viewers easily identified the synthetic displays as human walkers and accurately determined the intended gender of the walkers. These experiments clearly showed that variations in torso rotation are important for gender judgments. Accordingly, we chose to measure viewer sensitivity to torso rotations in our experiments.

Proffitt and his colleagues[24] found that occlusion of light dots by clothing or human body segments plays an important role in gait judgment and may also provide information about body outlines. Synthetic displays without occlusion yielded poorer subject performance. These experimental observations suggest that extremely simple models of human figures, such as thin stick figures may present similar difficulties for the viewer.

Surprisingly, the perception of rigid body segments between moving light dots at joints does not generalize to movements of isolated pairs of light dots. Ishiguchi[12] showed test subjects one fixed light dot and a second one that moved on an arc of ± 15 degrees as if it were on the end of a pendulum with the first light dot as the pivot joint. Viewers perceived the dots as attached to a flexible bar held fixed at the first light dot rather than as a rigid bar moving as a pendulum. Thus the perception of rigid body segments in the largely pendulum-like movements of human walking is exceptional; perhaps the ensemble of dot movements is important, or perhaps the movements are so intimately familiar that the perception of a more basic flexible bar is overridden.

3 Experimental Method

While it is impossible to exhaustively test all of the variables that may affect the perceived motion, we can form a preliminary assessment of whether the geometric model affects a viewer’s motion perception with A/B comparison tests. We evaluated two different variables in separate experiments described below: torso rotation and additive noise. For both these experiments, the modifications to the motion were controlled by a normalized parameter t that varied between $t = 0$ and $t = 1$. Samples of the animated motion are shown in figure 3 and motion traces in figure 4. In each case, subjects viewed pairs of animated sequences rendered using the same geometric model and were asked whether the motions in the two sequences were the same or different. We computed a sensitivity measure for each type of geometric model. The difference between the sensitivity values is a measure of whether a particular subject

was better able to discriminate between the motions when they were rendered with a polygonal model or with a stick figure model.

3.1 Experiment One: Torso Rotation

This experiment measured whether a subject’s ability to differentiate between larger and smaller yaw rotations of a runner’s torso was affected by the geometric model used for rendering. The motion sequences were generated by making kinematic modifications to data obtained from a physically based dynamic simulation of a human runner[11]. The vertical, or yaw, rotation of the torso at the waist was exaggerated. The neck was counter rotated to compensate for the torso rotation so that the facing direction of the head remained unchanged.

The magnitude of the exaggeration in torso rotation was controlled by a normalized parameter, t . A value of t equal to zero gave a magnification factor of $1\times$, so that the modified motion was identical to the original data. Larger values of t correspond linearly to higher magnification factors, with $t = 1$ yielding a $10\times$ magnification of the torso rotation. The motion of body segments below the waist was left unchanged.

The test consisted of a series of 40 pairs of motion sequences divided into two sets of 20 pairs each. One set was rendered with the stick figure model and the other with the polygonal model (figure 1). Except for the geometric models, all parameters used to render the animations were identical for the two sets. Within each set, half of the pairs were randomly selected to show two different motion sequences (different t values). Of those that were different, the pairs with the largest disparity in t were placed toward the beginning of each set so that the questions became progressively more difficult. To minimize bias due to fatigue or learning effects, we varied the order in which the two sets were presented.

Twenty-six student volunteers served as subjects. All had normal or corrected-to-normal vision. Subjects were tested in groups of two or three in a quiet room. The test stimulus was presented on a 20-inch monitor approximately three feet from the subjects. All animations were pre-rendered and shown at 30 frames per second in NTSC resolution.

Subjects were told that they would be shown a series of 4-second computer-generated animations of a human runner and that the animations would be grouped in A/B pairs with five seconds of delay between the presentation of each pair. Subjects were asked to view each pair and then indicate on a response sheet whether the two motions were the same

or different. They were also informed that the variations would be confined to the motion of the runner’s upper body and that the questions would become progressively more difficult. A monetary reward was offered as an incentive to the subject who had the most correct responses. Subjects were not told the purpose of the experiment.

3.2 Experiment Two: Additive Noise

The format of the second experiment was identical to that of the first, except for the manner in which the running motion was modified. For this experiment, time-varying noise was added to the joint angles for the waist, shoulders, and neck. The noise was generated using a sinusoidal wave generator[26] with frequency varying randomly about that of the runner’s gait (approximately 3 Hz). The amplitude of the additive noise was controlled by a normalized parameter t , as in the torso rotation test. A value of $t = 0$ resulted in motion data that was identical to the original data (zero noise amplitude). The maximum noise amplitude used, given by $t = 1$, produced a variation of ± 0.15 radians about the original joint angles.

Twenty-six additional volunteers who had not participated in first experiment were subjects for this second experiment. Testing procedures were identical to those used in the first experiment.

4 Results

To analyze the data from the two experiments, the responses were used to compute the Choice Theory sensitivity measure for each subject on each test set[20]. The sensitivity measure, $\log(\alpha)$, is defined as

$$\log(\alpha) = \frac{\log(H/(1-H)) - \log(F/(1-F))}{2}, \quad (1)$$

where H is the fraction of pairs in a set that were *different* and which the subject labeled correctly, and F is the fraction of pairs in a section that were *the same* and which the subject labeled incorrectly[20]. This measure is zero when the subject’s responses are uncorrelated with the correct responses (i.e. 50% correct) and increases as response correlation improves (figure 5). Additionally, the measure is symmetric, naturally invariant with respect to response bias, and suitable for use as a distance metric[20].

In Section 1, we proposed three possible answers to the question of whether the geometric model used for rendering affects a viewer’s judgment of motion. The third possible answer implied that subjects would achieve similar sensitivity measures when asked identical questions about the motion of



Figure 3: Sample image sequences from the animations used in the experiments. **First Row:** Original running motion rendered with the polygonal model. **Second Row:** Torso rotation with $t = 1.0$. **Third Row:** Additive noise with $t = 1.0$. **Fourth Row:** Original running motion rendered with the stick figure model. **Fifth Row:** Torso rotation with $t = 1.0$. **Sixth Row:** Additive noise with $t = 1.0$. Images are sampled at 0.067 second intervals.




Figure 4: The top graph shows a trace of the motion of the left and right shoulders and the head for one stride of the original running data. The camera is positioned directly above the runner. The middle graph shows the same traces after the data were altered for the torso rotation test ($t = 1$). The bottom graph shows the traces for the additive noise test ($t = 1$).

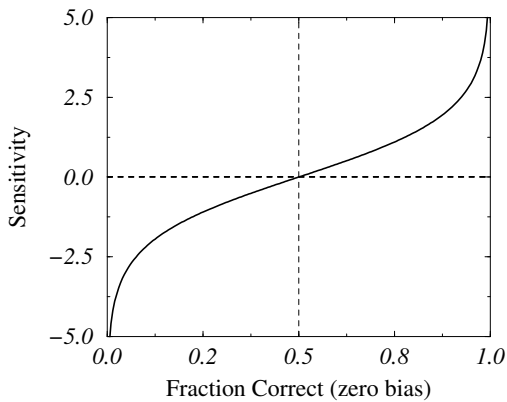


Figure 5: Plot of sensitivity scores, $\log(\alpha)$, versus fraction correct. Shown at zero bias.

stick figure models or polygonal models. To test this hypothesis, we computed the difference in sensitivity for each subject:

$$\Delta \log(\alpha) = \log(\alpha_{poly}) - \log(\alpha_{stick}). \quad (2)$$

For the torso rotation test, the mean of the difference in sensitivities across all subjects was 0.427 with a standard deviation of 0.77. Student's t test for paired samples [23] shows this difference to be significant, $p \leq 0.012$. When we reduced the subject pool by removing those who had been unable to judge differences accurately in either set ($\log(\alpha_{poly}) < 1.0$ and $\log(\alpha_{stick}) < 1.0$, or equivalently less than 73% correct for both the stick figure model and the polygonal model), the mean rose to 0.73 while the variance fell to 0.68. Again, the t test for paired samples shows this difference to be significant, $p < 0.001$.

For the additive noise test, the mean of the difference in sensitivities across all subjects was 0.74 with




Figure 6: A graphical representation of the difference in sensitivity for the torso rotation test. The left graph shows the sensitivity values for all subjects; the right graph shows values for only those subjects who performed well on the torso rotation test with a sensitivity of $\log(\alpha) \geq 1.0$ (approximately 73% correct or better) on at least one of the two sets. If a subject has a positive slope for the line connecting the two sensitivities, then that subject was more sensitive to motion changes when polygonal models were used.

a standard deviation of 0.69, a difference significant at $p < 0.001$. When we reduced the subject pool by removing those who had been unable to judge differences accurately in either set, the mean was 0.72 and the variance was 0.73, a difference significant at $p < 0.001$.

The results of the torso rotation test are displayed graphically in figure 6. Sensitivity values for the set rendered with the stick figure model are plotted on the left vertical axis, while values for the the set with the polygonal model are on the right. The values for each subject are connected with a straight line, the slope of which indicates the difference in sensitivities measured between the two geometric models. Figure 7 shows a histogram of the sensitivity differences for the torso rotation test. As with the slopes in figure 6, positive values correspond to higher sensitivity for the set rendered with the polygonal model. Figures 8 and 9 show similar plots for the data from the additive noise test.

These results indicate that, for the two types of motion variation tested, subjects were better able to discriminate motion variations using the polygonal model than they were with the stick figure model.

5 Discussion

Though the differences in sensitivity measures show that our subjects were more sensitive to motion




Figure 7: Histogram of sensitivity differences for the torso rotation test. The upper graph shows the occurrence frequency for sensitivity differences ($\log(\alpha_{poly}) - \log(\alpha_{stick})$) across all subjects. The bottom graph shows the data for subjects who achieved a sensitivity score of $\log(\alpha) \geq 1.0$ (approximately 73% correct or better) on either the set using the polygonal model or the set using the stick figure. Positive values of the sensitivity difference indicate a higher sensitivity to changes in the motion with the polygonal model. (Bucket size = 0.5.)

changes when a polygonal model was used for rendering, our results should not be generalized to mean that polygonal models are always better than stick figure models for judging motions. Rather the two types of geometric models are distinctly different and in the cases we tested polygonal models allowed better discrimination. We measured sensitivity for only two types of variation in running motions. There may well be variations for which the difference in sensitivity has the opposite sign, implying that stick figures might be a better type of model for making fine discriminations about that particular type of motion variation.

Our results, however, do show that stick figures and polygonal models are not equivalent for tasks that require making fine discriminations about motion. Any useful comparison of motion sequences requires that the same models and rendering methods be used for each. Comparing motions of a stick figure model to those of a more complex model may be meaningless because viewer sensitivities can differ substantially. Furthermore, presenting the results of animation techniques using stick figures, or other simple models, should be avoided because it is likely that viewers would have different sensitivities to the

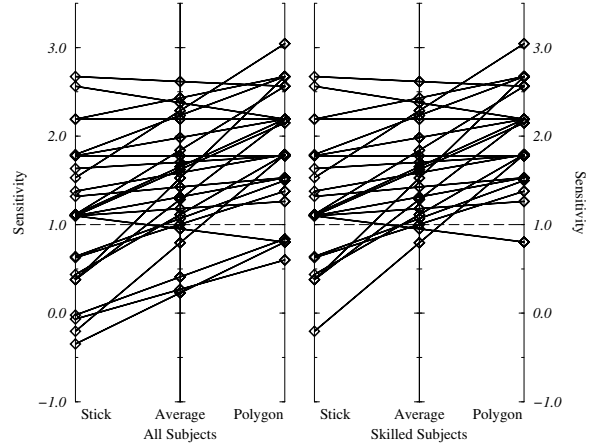


Figure 8: A graphical representation of the difference in sensitivity for the additive noise test. (See figure 6.)




Figure 9: Histogram of sensitivity differences for the additive noise test. (See figure 7.)

more realistic models used in the final rendering.

Considerable familiarity with the motion appears to make differences in the geometric models less significant. For example, when the authors of this paper took the tests, they answered nearly all questions correctly. (Of course, the authors were not included among the subjects whose data is reported above.) If a larger subject pool showed that subjects who were very familiar with particular animated motions showed equal sensitivity to the two models, then we would have evidence that using stick figures for preliminary *pencil tests* of motion sequences will provide good information about the motion. The subject, in this case the animator, is very familiar with the motion and may be able to make subtle observations independent of the geometric models used for ren-

dering.

A potential problem with the experimental design used in this study is that the test must be of an appropriate difficulty. If the test is too difficult then all of the subject's responses will be guesses regardless of which model is presented. Conversely, if the test is too easy then all of the subject's responses will be correct. In either case the data gathered will not be useful. We can increase or decrease the difficulty of a test by changing the spacing of the t values for the trials or the amount of information given to the subjects about the alterations to the motion. Unfortunately, it can be difficult to devise a test sequence of appropriate difficulty. This dilemma could be overcome by using tests that adaptively adjust difficulty level by selecting subsequent questions based on past responses. Alternatively, selection criteria can be used to cull subjects whose responses are not significantly correlated with the test stimuli.

The additive noise test was designed to be easier than the torso rotation test, and the plots in figures 6 and 8 show that in general the subjects' average scores were indeed higher for the additive noise test. While our assessment that the polygonal models allow greater sensitivity holds irrespective of culling, it is interesting to note how culling based on performance criteria does affect the data. For the more difficult torso rotation test, culling notably alters both the mean of $\Delta \log(\alpha)$ as well as the shape of the histogram in figure 7. For the easier additive noise test, culling has essentially no effect. Moreover, the effect of culling on the torso rotation data appears to make it more closely resemble the data from the additive noise test, supporting the notion that lowering the difficulty of the test and culling based on performance criteria are approximately equivalent.

Although we did not formally measure the subjects' perceptions of how well they did on the test, it appeared that their perceptions did not always match their performance. Several subjects were certain that they had scored higher on the section with the stick figure model when in fact they had a higher sensitivity to motion changes with the polygonal model.

To create the animation sequences for these tests, we altered only the motion and the geometric models used; all other aspects of the rendering were held constant. It would be interesting to explore whether, and how, other aspects of the rendering affect the perception of motion. For example, we have informally observed that the motion of the simulated run-

ner appears more natural when the tracking camera has a constant velocity rather one that matches the periodic accelerations of the runner's center of mass. If the camera motion matches the motion of the center of mass exactly, then the running motion appears jerky. More sophisticated models that incorporate clothing and skin may help to smooth out rapid accelerations of the limbs and make the motion appear more natural. Motion blur probably plays a similar role. Textured ground planes and shadows help to determine motion of the feet with respect to the ground and may provide important clues about the details of the motion.

If we had enough psychophysical results to build a model of how people perceive motion, we could optimize the rendering of animated sequences by including only those factors that would make the greatest differences in how a viewer perceives the sequence. This approach of using results from the psychophysical literature to refine rendering techniques has already been used successfully for still images[14, 27, 9].

Acknowledgments

The authors would like to thank Jacquelyn Gray, John Pani, and Neff Walker for their valuable comments. This project was supported in part by NSF NYI Grant No. IRI-9457621, Mitsubishi Electric Research Laboratory, and a Packard Fellowship.

References

- [1] N. I. Badler, C. B. Phillips, and B. L. Webber. *Simulating Humans: Computer Graphics Animation and Control*. Oxford University Press, New York, 1993.
- [2] C. D. Barclay, J. E. Cutting, and L. T. Kozlowski. Temporal and spatial factors in gait perception that influence gender recognition. *Perception & Psychophysics*, 23(2):145–152, 1978.
- [3] A. Bruderlin and T. W. Calvert. Goal-directed, dynamic animation of human walking. In *SIGGRAPH '89 Conference Proceedings*, volume 23 of *Annual Conference Series*, pages 233–242. ACM SIGGRAPH, July 1989. Held in Boston, Massachusetts, 31 July – 4 August 1992.
- [4] A. Bruderlin and L. Williams. Motion signal processing. In *SIGGRAPH '95 Conference Proceedings*, Annual Conference Series, pages 97–104. ACM SIGGRAPH, August 1995. Held in Los Angeles, California, 6–11 August 1995.

- [5] M. F. Cohen. Interactive spacetime control for animation. In *SIGGRAPH '92 Conference Proceedings*, volume 26 of *Annual Conference Series*, pages 293–302. ACM SIGGRAPH, July 1992. Held in Chicago, Illinois, 226–31 July 1992.
- [6] J. E. Cutting and L. T. Kozlowski. Recognizing friends by their walk: Gait perception without familiarity cues. *Bulletin of the Psychonomic Society*, 9(5):353–356, 1977.
- [7] J. E. Cutting, C. Moore, and R. Morrison. Masking the motions of human gait. *Perception & Psychophysics*, 44(4):339–347, 1988.
- [8] J. E. Cutting, D. R. Proffitt, and L. T. Kozlowski. A biomechanical invariant for gait perception. *Journal of Experimental Psychology: Human Perception and Performance*, 4(3):357–372, 1978.
- [9] J. A. Ferwerda, S. N. Pattanaik, P. Shirley, and D. P. Greenberg. A model of visual adaptation for realistic image synthesis. In *SIGGRAPH '96 Conference Proceedings*, Annual Conference Series, pages 249–258. ACM SIGGRAPH, August 1996. Held in New Orleans, Louisiana, 4–9 August 1996.
- [10] J. Freyd. Dynamic mental representations. *Psychological Review*, 94(4):427–438, 1987.
- [11] J. K. Hodgins, W. L. Wooten, D. C. Brogan, and J. F. O'Brien. Animating human athletics. In *SIGGRAPH '95 Conference Proceedings*, Annual Conference Series, pages 71–78. ACM SIGGRAPH, August 1995. Held in Los Angeles, California, 6–11 August 1995.
- [12] A. Ishiguchi. The effect of orientation on interpolated elastic structure from dot motion: Its occurrence and persistence. *Perception & Psychophysics*, 44(4), 1988.
- [13] G. Johansson. Visual perception of biological motion and a model for its analysis. *Perception & Psychophysics*, 14(2):201–211, 1973.
- [14] J. K. Kawai, J. S. Painter, and M. F. Cohen. Radioptimization - goal based rendering. In *SIGGRAPH '93 Conference Proceedings*, Annual Conference Series, pages 147–154. ACM SIGGRAPH, August 1993. Held in Anaheim, California, 1–6 August 1993.
- [15] H. Ko and N. I. Badler. Straight line walking animation based on kinematic generalization that preserves the original characteristics. In *Proceedings of Graphics Interface '93*, pages 9–16, Toronto, Ontario, Canada, May 1993. Canadian Information Processing Society.
- [16] L. T. Kozlowski and J. E. Cutting. Recognizing the sex of a walker from a dynamic point-light display. *Perception & Psychophysics*, 21(6):575–580, 1977.
- [17] L. T. Kozlowski and J. E. Cutting. Recognizing the gender of walkers from point-lights mounted on ankles: Some second thoughts. *Perception & Psychophysics*, 23(5):459, 1978.
- [18] J. Laszlo, M. van de Panne, and E. Fiume. Limit cycle control and its application to the animation of balancing and walking. In *SIGGRAPH '96 Conference Proceedings*, Annual Conference Series, pages 155–162. ACM SIGGRAPH, August 1996. Held in New Orleans, Louisiana, 4–9 August 1996.
- [19] B. Laurent, B. Ronan, and N. Magnenat-Thalmann. An interactive tool for the design of human free-walking trajectories. In *Proceedings of Computer Animation '92*, pages 87–104, 1992.
- [20] N. A. Macmillan and C. D. Creelman. *Detection Theory: A User's Guide*. Cambridge University Press, New York, 1991.
- [21] K. Perlin. Real time responsive animation with personality. *IEEE Transactions on Visualization and Computer Graphics*, 1(1):5–15, March 1995.
- [22] H. Poizner, U. Bellugi, and V. Lutes-Driscoll. Perception of american sign language in dynamic point-light displays. *Journal of Experimental Psychology: Human Perception and Performance*, 7:430–440, 1981.
- [23] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C*. Cambridge University Press, New York, 1992.
- [24] D. R. Proffitt, B. I. Bertenthal, and R. J. Roberts Jr. The role of occlusion in reducing multistability in moving point-light displays. *Perception & Psychophysics*, 36(4):315–323, 1984.

- [25] C. F. Rose, B. Guenter, B. Bodenheimer, and M. F. Cohen. Efficient generation of motion transitions using spacetime constraints. In *SIGGRAPH '96 Conference Proceedings*, Annual Conference Series, pages 155–162. ACM SIGGRAPH, August 1996. Held in New Orleans, Louisiana, 4–9 August 1996.
- [26] C. Schlick. Wave generators for computer graphics. In Alan Paeth, editor, *Graphics Gems 5*, pages 367–374. Academic Press, 1995.
- [27] P. Teo and D. Heeger. Perceptual image distortion. In *First IEEE International Conference on Image Processing*, pages 982–986, November 1994. Held in Los Alamitos, California, 11–16 November 1994.
- [28] M. Unuma, K. Anjyo, and R. Takeuchi. Fourier principles for emotion-based human figure animation. In *SIGGRAPH '95 Conference Proceedings*, Annual Conference Series, pages 91–96. ACM SIGGRAPH, August 1995. Held in Los Angeles, California, 6–11 August 1995.
- [29] A. Witkin and Z. Popović. Motion warping. In *SIGGRAPH '95 Conference Proceedings*, Annual Conference Series, pages 105–108. ACM SIGGRAPH, August 1995. Held in Los Angeles, California, 6–11 August 1995.