THE JOURNAL OF VISUALIZATION AND COMPUTER ANIMATION
J. Visual. Comput. Animat. 2000; **11**: 105–112

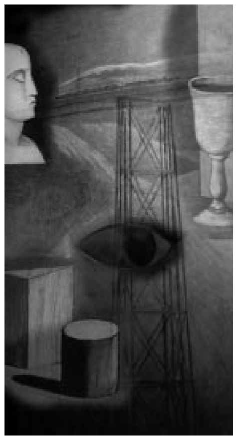

Psychological Evidence for Unconscious Processing of Detail in Real-time Animation of Multiple Characters

By Markus Oesker, Heiko Hecht and Bernhard Jung*

Detailed animation of 3D articulated body models is in principle desirable but is also a highly resource-intensive task. Resource limitations are particularly critical in 3D visualizations of multiple characters in real-time game sequences. We investigated to what extent observers perceptually process the level of detail in naturalistic character animations. Only if such processing occurs would it be justified to spend valuable resources on richness of detail. An experiment was designed to test the effectiveness of 3D body animation. Observers had to judge the level of overall skill exhibited by four simulated soccer teams. The simulations were based on recorded RoboCup simulation league games. Thus objective skill levels were known from the teams' placement in the tournament. The animations' level of detail was varied in four increasing steps of modelling complexity. Results showed that observers failed to notice the differences in detail. Nonetheless, clear effects of character animation on perceived skill were found. We conclude that character animation co-determines perceptual judgements even when observers are completely unaware of these manipulations. Copyright © 2000 John Wiley & Sons, Ltd.

Received: 12 August 1999; Revised: 17 December 1999

KEY WORDS: interactive visualization; human character animation; perceptual processing

Introduction

The Role of Visual Animation

Many findings of perceptual psychology confirm the remarkable ability of the human visual system to extract complex 3D information from very diminished visual displays. A few moving point-lights that correspond to the major body joints, for instance, suffice to create a vivid impression of a person dancing or carrying a heavy object.[1] We can see such things as mass, friction, intention, etc. although the stimulus basis of these impressions is two-dimensional and rather diminished. With sufficient experience, human observers are able to reconstruct the proper dynamic world of 3D objects solely based on the 2D retinal image.[2]

In the above-mentioned cases, animation is often vital for successful perception. Animation also bears witness of and activates knowledge that the visual system has internalized about the world but that is not accessible by conscious reflection. For example, a large proportion of the adult population has misconceptions about the laws of classical mechanics, but the misconceptions are not always reflected in their perceptual judgements. Up to 50% of university students that were tested predict that an object dropped from a moving carrier falls straight down, but they admit that it looks rather strange when presented with an animated version of their mistaken predictions.[3] Here implicit visual knowledge is correct while explicit knowledge is mistaken. And it is only through visualization that the implicit knowledge becomes useful.

However, graphics animation is not always so clearly advantageous. There are examples where visual and explicit cognitive 'knowledge' are similarly mistaken. Both can deviate in astonishing ways from the laws of classical mechanics. For example, projectiles such as a cannonball or a baseball appear most natural and their animation looks most realistic when the projectile continues to accelerate after it has left the

*Correspondence to: M. Oesker, Fakultat Fur Biologie, Universität Bielefeld, UniversitatsstraBe 25, D-33615 Bielefeld, Germany.

cannon or the thrower's hand.[4] In other cases visualization seems altogether gratuitous. It matters little whether we drop a planar disc or a detailed rendition of a beer keg from the moving carrier. Thus it is an empirical question under what circumstances complex visualization is beneficial and under what circumstances it might be detrimental.

Our study was devised to gain some insight into the challenge of determining the amount of detail that is best visualized in the context of real-time 3D character animation. In particular, we sought to assess from the point of view of perceptual psychology how the visual system processes aspects of 3D character animation. We also investigated whether the degree to which observers notice a given visualization technique is related to its effectivity. To do so, we selected RoboCup soccer simulations, a dynamic event that allows a large range of variation in the style of visualization.

The Simulation

To freely vary the complexity of character animation, we chose a game sequence of individual agents that provided a maximal range of sophistication. A soccer game can be visualized strategically by mere dots representing the players, or the players can be animated in indefinite degrees of detail (see Figure 1). The artificial agents designed for simulated soccer games provided the backdrop for our experiment. We built upon the Robot World Cup Initiative (RoboCup), which has proposed robotic soccer as a new standard problem to foster a wide range of artificial intelligence (AI) and intelligent robotics research. Although the ultimate goal of RoboCup is to build soccer teams consisting of real robots, it offers a software platform for the development and interactive competitions of distributed multi-agent soccer simulations.[5] RoboCup soccer games are distributed simulations consisting of up to 22 players, which are realized as independent control processes, and a central server that maintains the current simulation state. The simulation environment defined by the RoboCup server is two-dimensional; that is, players and ball are represented as circles. Every 100 ms the RoboCup server generates a snapshot of the game state describing the current positions of players and ball as well as some additional information about players' kicking actions and scored goals. These snapshots constitute the input for 2D or 3D visualization systems. Until recently, this platform had only visualized the soccer games by representing individual players as circular patches on a uniform background corresponding to a soccer field viewed from above (2D RoboCup soccermonitor). Onto this platform we have built Virtual RoboCup, an animation system that provides real-time 3D visualizations of RoboCup simulation league soccer games with task-level character animation.[6,7] Virtual RoboCup allows us to view the soccer field from arbitrary angles. More importantly, it creates character animations for the players on the fly.

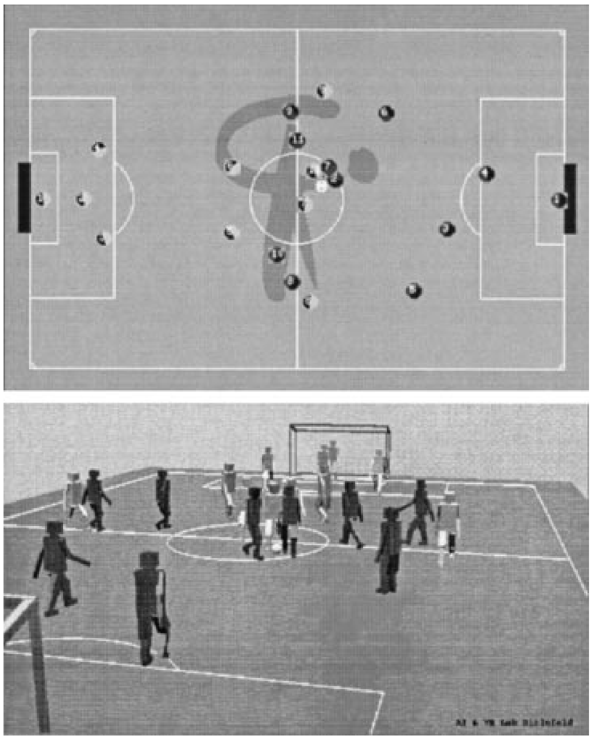

Figure 1. RoboCup provides a 2D environment for simulated soccer games (top). Given 2D input from the soccerserver, Virtual RoboCup generates real-time 3D visualizations where players are animated as anthropomorphic figures (bottom).

Each articulated body model of Virtual RoboCup players consists of a hierarchically ordered set of 15 box-shaped elements representing the head, neck, torso, upper and lower arms, thighs and calves, hands and feet (see Figure 2). The segments can rotate around their joints with three degrees of freedom each. However, to reduce computational complexity, only a fraction of all possible rotations were modelled for the purpose of animation.

Running is effected by leg rotation in the hip joint and by flexing the knee joints. The shoulders mirror leg angles. Fast processing is realized through interpola-

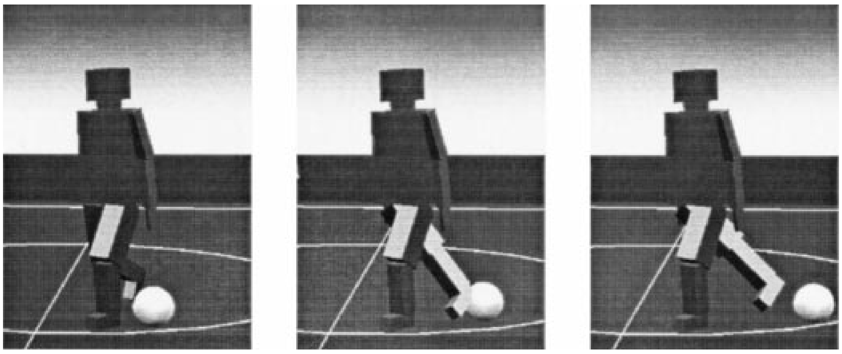

Figure 2. An animated character shoots the ball: preparation–contact–follow-through.

tion between predefined body poses that consist of a cyclic sequence of 30 parametric keyframes (see Figure 3 for examples). The whole sequence shows a player moving two steps forward, one step per leg. Because simulated players move with changing speed, body postures are computed with respect to the player's rate of displacement. A 'posture index' in the keyframe table is changed each animation cycle by adding a player's displacement to its position during the last cycle. A new body pose is thus generated by weighted interpolation between the two keyframes that are closest to the posture index. The functional description running devised by Hodgins[8] was taken into account in the design of the keyframes. As a result, the dynamics of animated running appears fairly natural.

Kicking is effected by co-ordinated action in the hip, the knee and the foot. Inverse kinematics techniques are used to ensure that players first recoil their foot, then extend it forward to the desired location, and then establish contact with the ball. In a final follow-through phase the active leg continues to swing in the direction of the ball.

To generate a smooth and detailed animation of running and kicking actions, the original 2D RoboCup simulator output is enhanced with intermediate states. Multiprocessing techniques are used to prevent disadvantageous linking of the rendering task with reception and evaluation of incoming simulation data.

Issues of Visualization

In recent years the progress in speed of computational hardware has led to a considerable increase in graphics performance, and 3D animations have become feasible on ordinary workstations. In spite of this development, it is still very hard to reach reasonable frame rates in graphics applications with real-time requirements. It is therefore important to avoid wasting valuable resources on modelling detail that has no perceptual utility.[9]

Visualization can function as an important and intuitive source of information if it succeeds in conveying information about the soccer players' actions that go beyond pure positional values. These include anticipation of movement, assessment of strength, intention, etc. The development of effective animations hinges on visualizing those features that carry critical information while omitting non-critical features. Their criticality can only be determined empirically. It is for instance known that observers are particularly sensitive to the modelling of human features and detail.[10]

No matter what arguments for or against visualization can be made, it is an empirical question whether or not human observers do in fact incorporate visualization detail into their perceptual judgements. For a meaningful empirical test the experimental variations in the degree of visualization must not interfere with the judgements of the observers. We

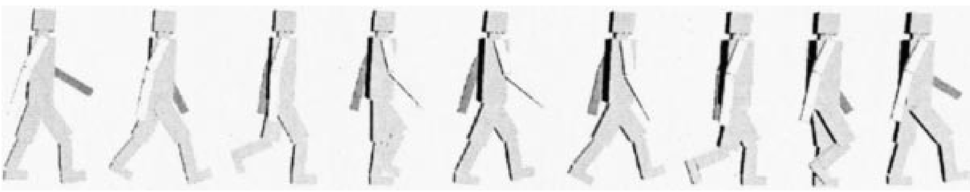

Figure 3. Some of the 30 predefined locomotion keyframes.

thus needed a task that could be performed identically regardless of the visualization used. At the same time the range of animated detail had to cover a representative range.

The Experiment

Design

To test the influence of articulated body animations on observers' judgements of the abilities of simulated soccer teams, the following experiment was designed. First, a factor of objective skill level was created. Four teams were selected to span a large range of accomplishment. The teams' playing skills ranged from a RoboCup tournament winner in 1998 to a team that was a few years back in evolution. All possible matches between the four teams were used in the experiment, also considering what team was on the left and the right side of the soccer field. Short sequences were isolated from each of the resulting 12 recordings of simulation soccer matches. In all cases the players and the ball were visible such that a meaningful judgement of skill could be made. The 12 sequences were fully crossed with four different levels of animation detail, resulting in a total of 48 sequences. For each sequence, observers were asked to specify the level of skill of both participating teams separately on a linear scale ranging from 0 to 12. We were thus able to contrast the level of playing skill with the degree of character animation.

Stimuli Composition of the Sequences

Four teams with known and heterogeneous abilities were used. Each team consisted of five identical players controlled by the same algorithm but starting at different positions on the field. We required as a main selection criterion for selection of teams that they be objectively discriminable by means of scored goals. Table 1 gives an insight into the selected teams' performance as exhibited during a competition. Teams Sopra1, Sopra3 and Sopra4 were judged to be reasonably intelligent by the experimenters while being clearly discernible in their level of skill. The team composed of Krislet players not only scored no goals at all but also behaved strangely. Krislet players tend to move very straight towards the ball without paying attention to their team-mates' actions. Often they meet in one location, disturbing each other's movements.

Team	Goals
Sopra1	34 : 00
Sopra3	14 : 10
Sopra4	03 : 10
Krislet	00 : 31

Table 1. Four soccer teams rank-ordered according to the number of goals scored during a competition

Three teams had participated in a competition that took place in 1998 at the University of Bielefeld, Germany and had reached the first, third and fourth place. The simple client named Krislet designed by Kryzsztof Langner was taken from the Internet.[11] All four clients were implemented in Java and ran smoothly using Soccerserver version 3.28, which was used for the experiment. Three competitions were recorded in which each of the four teams played against each other. From these recordings we cut sequences of 20 s duration corresponding to 400 animation cycles each. Every sequence showed a promising attack, which was defined as driving the ball in the direction of the opposing team's goal or at least the attempt thereof. Whenever possible, we chose sequences that contained the scoring of a goal. The selected set of sequences contained scenes of all 12 possible combinations of teams. Each team staged one attack against each of the other three teams.

Animation conditions were as follows (see Figure 4).

1. Running and shooting actions both animated.
2. Animation of kicking only. When shooting the ball, players recoiled their leg, bending the knee and foot, and then continued to extend the leg after contact with the ball. However, there was no running action. Players moved rigidly across the field.
3. Animation of running only. During locomotion, players moved their legs with action in the hip and in the knee. Players did not perform a special shooting action. Consequently, the direction and speed of the ball changed although neither legs nor feet touched it.
4. No animation of running or kicking. The players moved and turned completely rigid on the field. Their limbs did not change orientation with respect to one another.

Figure 4. Animation of player movement. Each row shows the same temporal sequence while the animation style is varied. From top to bottom, full animation (top row), kicking only (second row), running only (third row) and no animation (bottom row) are depicted.

Apparatus and Procedure

The sequences were presented using a slightly modified version of Virtual RoboCup allowing choice between the different levels of animation detail. Each of the above game sequences was presented four times using the different levels of animation detail. The resulting 48 sequences were presented in random order to each observer. The defending team was always named A, the other one B. Players of team A were coloured yellow, the ones of team B red. The observer's point of view corresponded to a position near the corner to the right of the defender's goal (see Figure 5). The direction of gaze was directed at the ball. This presentation ensured that observers could not identify teams except by the clients' actions.

The experiment was implemented as a Perl script that played sequences using the logplayer of Soccerserver version 3.28 and Virtual RoboCup's 3D soccermonitor. It was carried out as a full screen application on an SGI Indigo 2 XZ machine with a 21'' monitor. The sequences were generated using the source code for the original teams. For the experimental session the digitally recorded sequences were presented at a frame rate of 7 Hz. The refresh rate of the monitor was 72 Hz.

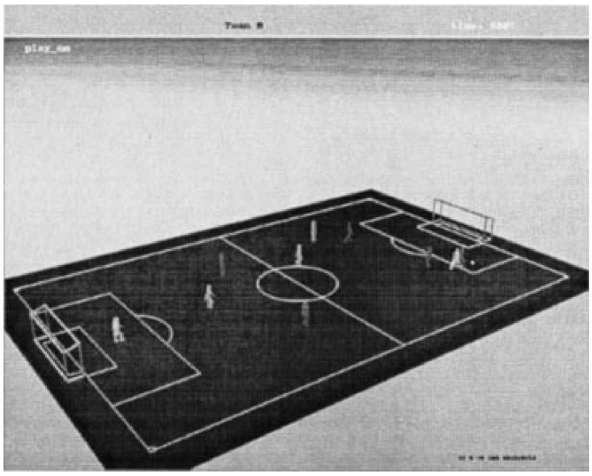

Figure 5. Observers' view of the simulated soccer field.

Eight student observers (four men, four women) were paid for their participation. After viewing the 20 s sequence, each observer had the opportunity to review the entire sequence if so desired. Then she was asked to first decide which of the two teams was more apt and skilful in its overall play. Once this decision had been made, she had to assign a grade to each team. A grade of 0 corresponded to pitifully poor skill, a grade of 12 to exceedingly adept. To familiarize observers with the task, about 10 practice trials were randomly selected from the pool of trials and displayed.

Results and Discussion

After all data had been collected, participants were asked what in their opinion distinguished the teams. They were also asked whether they had noticed any changes in the animation of the bodies between trials. Amazingly, none of the observers reported changes in the animation style. Even when directly asked whether in some trials players had moved or shot differently, observers failed to report differences in animation. Instead, they typically noticed that in some teams players would be rather dumb and all move straight towards the ball without any further strategy.

Observers fairly reliably recognized the objective skill of the teams. In 75·1% of all cases they correctly assigned the higher skill grade to the higher-ranking team (see Table 1). The skill grades assigned to the different teams were correlated positively with their objective skill ($r = 0{\cdot}51$, $p < 0{\cdot}0001$). To determine the Spearman rank correlation, the sum of squared ordinal distances between rankings for each objective–subjective pair of assessments was normalized.[12] A value of 1 would indicate perfect, a value of 0 absence of correlation. The p value indicates the likelihood of the correlation to be produced by chance. Thus, as expected, the main strategic differences produced by the clients were reflected in the judgements.

For the purposes of analysing the effects of the unnoticed changes in character animation, the grade scores were entered into a repeated measures analysis of variance (ANOVA) with four levels of animation as independent factor. For the dependent variable the average judged grade for both teams of a given sequence was computed. Thus, the effect of perceived skill as a function of animation detail was isolated from the quality of the game. Note that the animation manipulation was always the same for both teams. For a perfect observer there should obviously be no differences for games that are identical except for the animation of all players. Thus the ANOVA only reflects the influence of character animation. A significant main effect was found for this factor ($F(3,21) = 3{\cdot}28$, $p = 0{\cdot}041$). The F ratio is the ratio of treatment variance to variance due to chance. The likelihood that the results are due to chance is 4·1% as indicated by the p value. The effect of animation on skill judgements is depicted in Figure 6. Individual contrasts reveal that the difference between the full animation and the absence of all character animation was most pronounced ($F(1,7) = 7{\cdot}81$, $p = 0{\cdot}0268$). Fully animated characters lead to the highest ratings of playing skill for both teams. They are also marginally better than characters with shooting animation only ($F(1,7) = 4{\cdot}22$, $p = 0{\cdot}079$) but do not differ significantly from the run-only condition. The difference between no animation and shooting action only was not significant.

The fact that adding shooting detail to the running animation had very little effect on the results indicates that the limit of meaningful level of detail may have been reached. It is unlikely that additional hand or head movements would add much to the type of skill assessment in the present type of experiment. The poor effect of kick animations may not originate in the limits of observers' perceptual capacities but may instead be caused by features of the animation itself. Soccer players are running nearly continuously for the time of a match while shooting actions are performed only at times and by single players. Moreover, our animation technique may be suboptimal for visualization of short and accentuated events such as kicking actions. As mentioned above, Virtual RoboCup uses multi-processing to ensure a continuous flow of animation.

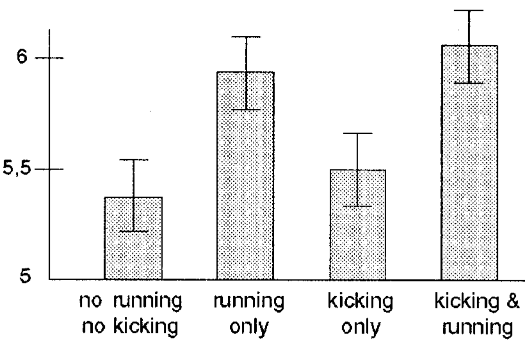

Figure 6. Skill ratings averaged over all performance variations, plotted by degree of character animation. The ratings reflect perceived skill regardless of how correctly observers identified the teams. Error bars indicate standard errors of the mean.

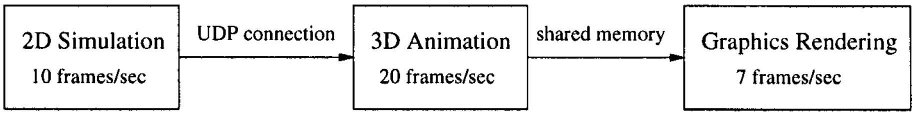

Figure 7. The animation process generates 3D scenes at a fixed rate of 20 frames per second. Rendering speed depends on the underlying graphics hardware.

The animated characters' changes in location and body pose were synchronized with incoming simulation data but were not synchronized with the rendering process. Since the animation process continuously updates shared memory, some data may have been lost owing to the slow rendering process. This was necessary for the 3D animation to keep up with the 2D soccer simulation. It also had the advantage of producing very smooth animation. However, some simulation steps might not have been visualized when rendering was slow (see Figure 7). This loss may have negatively affected the perception of kicking actions, which lasted only between four and six simulation cycles. Contact with the ball was as short as a single time step.

Conclusion

We were concerned with the question of where the visualization of human character details is advantageous and where it might be nonsensical. Four levels of character animation were added to a given display of two teams playing soccer. Rather astonishingly, human observers failed to consciously notice these manipulations in animation style. This lack of awareness, however, did not prevent the animation from influencing observers' judged skill levels of the teams. Everything else controlled for, a team whose characters were animated in their running and shooting actions were judged to be more skilful. The running action tended to be most important in this context. These findings reveal the important influence of the level of animation detail on perceptual variables. They also reveal that explicit judgements, such as obtained by questionnaires or by mere inspection of the displays, are insufficient to assess the importance of level of detail in character animation. Detail is processed unconsciously.

The limitations of the present study are twofold. First, our within-subjects design is unnatural because of the repeated observations that were demanded from the observers. Since observers are usually only confronted with one particular type of character animation, a between-subjects design might have been more informative. Unfortunately, a number of observers large enough to avoid repeated viewings was beyond our means. Also, there is no reason to believe that the sequential presentation of game scenarios did anything else but accentuate effects that would have been harder to detect in a between-subjects design. Second, our results may only speak to non-interactive visualization. Our observers clearly were in a passive mode and might have been more attentive if they had been able to manipulate the game. Further study is needed to assess the generalization of unconscious processing of detail to interactive situations. Work with a simulated interactive squash game indicates that this is not only the case, but that extraneous variables such as intended actions and the success of the action enter the picture.[13]

In sum, we have found that human observers are quite sensitive to detail in 3D character animation even when unaware of its manipulation. Consequently, we should consider adding other features to character animations for simulation league soccer games. For example, player numbers could be added on the characters' backs, or body postures and degrees of acuteness of shoulder and elbow angles could be used to reflect the players' stamina. To ensure the effectiveness of such additions, they should of course be tested empirically.

In our visualization experiment the 2D soccer game simulation was strictly independent of the 3D character animation. Because of this constraint, limits exist with respect to the degree of naturalness that can be reached in the visualization as compared to what might be possible with an integrated 3D simulation and visualization. Nonetheless, we have demonstrated that variations in animation style of the add-on visualization do influence observer judgements of player skill. As a general rule, psychological studies are necessary to determine what level of detail is functional and where better to omit detail that is costly but goes unnoticed and unprocessed by the visual system.

References

1. Johansson G. Visual perception of biological motion and a model for its analysis. *Perception and Psychophysics* 1973; **14**: 201–211.
2. Runeson S, Frykholm G. Kinematic specification of dynamics as an informational basis for person-and-action perception: expectation, gender recognition, and deceptive intention. *Journal of Experimental Psychology: General* 1983; **112**: 585–615.
3. Kaiser MK, Proffitt DR, Whelan S, Hecht H. Influence of animation on dynamical judgements. *Journal of Experimental Psychology: Human Perception and Performance* 1992; **18**: 669–690.
4. Hecht H, Bertamini M. Understanding projectile acceleration. *Journal of Experimental Psychology: Human Perception and Performance* 2000; **26** (in press).
5. Kitano H, Asada M, Kuniyoshi Y, Noda I, Osawa E. RoboCup: the Robot World Cup Initiative. Proceedings of the First International Conference on Autonomous Agent (Agent-97), 1997.
6. Oesker M. Simulation und Animation anthropomorpher Figuren in dynamischen virtuellen Umgebungen. *Master's Thesis*, Universität Bielefeld, 1998.
7. Jung B, Oesker M, Hecht H. Virtual RoboCup: real-time 3D visualization of 2D soccer games. *Proceedings of RoboCup99, Stockholm*, 1999; 121–126.
8. Hodgins JK.Three-dimensional human running. *Proceedings of the IEEE Conference on Robotics and Automation*, 1996.
9. Kaiser MK, Montegut MJ. Of red planets and indigo computers: Mars database visualization as an example of platform downsizing. *Behavior Research Methods, Instruments and Computers* 1997; **29**: 48–53.
10. Magnenat-Thalmann N. Visualizing humans by computer. *Educational Media International* 1992; **29**: 213–220.
11. Langner K. Krislet, a sample client for RoboCup simulation league. 1997. Available: http://ci.etl.go.jp/~noda/soccer/client/index.html [23 March 1999].
12. Siegel S. *Nonparametric Statistics for the Behavioral Sciences*. McGraw-Hill: New York, 1956.
13. Hecht H. Subjective realism in a simulated squash game. In *Advances in Multimedia and Simulation. Human–Machine Interface Implications*, Holzhausen K-P (ed.). Fachhochschule Bochum: Bochum, 1997; 365–370.

Authors' biographies:

Bernhard Jung is an Assistant Professor in Artificial Intelligence and a member of SFB 360 ''Situated Artificial Communicators'' at the University of Bielefeld. He received his Ph.D. in Artificial Intelligence from the University of Bielefeld in 1996. His current research interests include knowledge representation, virtual reality, intelligent multimedia, multi-agent systems, and cognitive robotics.

Markus Oesker received a Diplom (M.Sc.) in Computer Science from the University of Bielefeld, Germany in 1998. As Ph.D. student he is a member of a project on computer-based training for information processing in biological neuronal networks (Rubin) at the University of Bielefeld. Research interests include visual simulation and perception, intelligent multimedia, and adaptive systems.

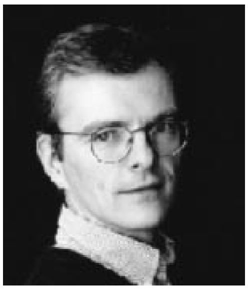

Heiko Hecht is currently a visiting scholar at the Man-Vehicle Lab, Massachusetts Institute of Technology. He received his Ph.D. in Cognitive Psychology in 1992 from the University of Virginia. His research interests include criteria for reality in visual displays, pattern goodness, apparent motion, structure from motion, ecological perception and action, and intuitive physics.