DEPARTMENT: VISUAL COMPUTING: ORIGINS

On Raising a Virtual Human

Norman I. Badler, *Cesium GS, Philadelphia, PA, 19106, USA*

I have been fortunate to have experienced and contributed to the software development of virtual humans. Whether they are Avatars, Conversational Agents, game NPCs, or Virtual Influencers, they all have origins in computer graphics (CG). Indeed, the "DNA" of these human models encompasses a large set of CG component technologies including 3-D modeling, animation techniques, and interactive systems. Finding (or purporting to know) "the" path from the origins to the present would require writing a CG and art history book about representing the human form. Here I will try to document portions of an evolutionary path—and my personal journey—to elucidate the context and processes that led up to the explosion of virtual human software "DNA" by the turn of the millennium.

Visual Computing: Origins stories are often fascinating tales of connections, discoveries, and effort. In computer graphics, we celebrate the ideas and people who brought real and imaginary worlds into the realm of synthetic visuals. Tracing origins of new ideas by infinite regression through the citation fields of technical papers is a classic exercise for graduate students embarking on a Ph.D. Most often, however, the process by which big ideas develop and flourish are hidden in the personal experiences of the creators along the evolutionary chain. This generative process is fraught with the same processes and consequences contributed by natural biological evolution: an unexpected mating of ideas, a meaningful mutation of a core concept, a withering of less capable offspring, and even the butterfly effect of random environmental influences. As we look backwards, we can better recognize this evolutionary process to seek or at least understand the triggering conditions. While in the midst of this evolution, we may not clearly see those "aha" moments that will shape the future; they may be discovered in retrospect to trace the crucial influences along the path.

0272-1716 © 2023 IEEE
Digital Object Identifier 10.1109/MCG.2023.3320319
Date of current version 7 November 2023.

MY PH.D. THESIS

The fundamental premise of my Computer Science Ph.D. thesis from the University of Toronto, in 1974, is simple to state. People exist in a dynamically changing 3-D world and effortlessly transform visual observations into textual descriptions. How can I design a computer process that achieves the same goal? In humans, this capability is linked with language acquisition, notions of object permanence, and visual processing. Describing what we see seems a completely natural and well-established skill long before we undertake more structured pedagogy.

I hypothesized that this descriptive facility might be computable. I wanted to jumpstart the process by working from synthetic visual data rather than camera images. To describe moving objects, I needed a reliable source of temporally connected image frames: a movie. Therefore, my first task was to implement a 3-D object animation system. I programmed a system that fed 3-D object movement parameters into a simulation engine to time-slice the movements and transform everything on a per-frame basis. Objects consisted of solid polyhedral and 3-D line models. Since the idealized point camera was a first-class 3-D object, I needed to spherically project object points about the camera onto the image plane. That meant straight lines mapped to arcs on the sphere. As objects could have solid surfaces and everything existed in 3-D, I needed to write a visible line algorithm to properly render the scenes directly onto a 16-mm microfilm scanner (youtu.be/UZ12RaI1Opo). In

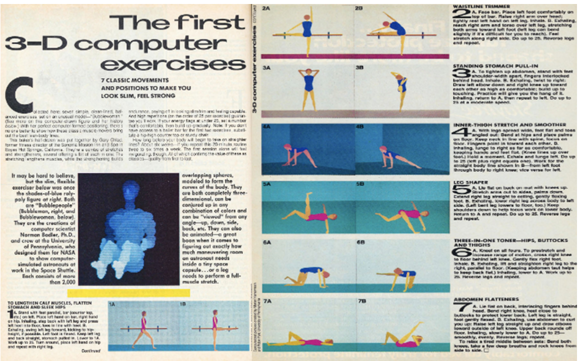

The first 3-D computer exercises

7 CLASSIC MOVEMENTS AND POSITIONS TO MAKE YOU LOOK SLIM, FEEL STRONG

FIGURE 1. The *SELF* magazine exercise article with bubblewoman figures. (Reprinted with permission).

1972, I did not think that this animation system was itself worthy of publication, and it ended up being barely mentioned in my thesis as an appendix.

COLOR GRAPHICS PEOPLE

I joined the University of Pennsylvania ("Penn") as an Assistant Professor in 1974. Knowing how costly raster graphics rendering of 3-D polygons could be using the early (1976) color frame buffers, I sought a different way to exploit the raster format to obtain images that had more human-like shape than the stick figures in my thesis animations. This gave birth to our first synthetic humans: Bubbleman and Bubblewoman. They were constructed with overlapping sets of spheres. The spheres allowed more natural solidity and body shape. This idea derived from discretizing the 3-D medial axis transformation of a surface. The armatures on which the spheres hung were moved by 3-D kinematic joint rotations. The sphere rendering was a reasonably efficient back-to-front sort of projected spheres into shaded disks with some depth attenuation. By 1979, we had published papers on Bubblepeople with Ph.D. student Joseph O'Rourke and Masters student Hasida Toltzis.[1]

With Masters student Mary Ann Morris, we described a joint rotation transformation scaled along a body segment to proportionally transform sphere centers and portray realistic shape twisting.[2] Bubblewoman peaked in utility when she was featured in an exercise article in the September 1984 issue of *SELF* magazine, reproduced in Figure 1. Masters students Marion Hamermesh and Jon Korein used Bubblewoman to replace photos of an actual model's poses. We used the 3-D computer graphics model to re-color the model and "re-shoot" improved camera angles. To my knowledge, this is the first instance of a computer graphics human model starring in a major consumer-level magazine article.

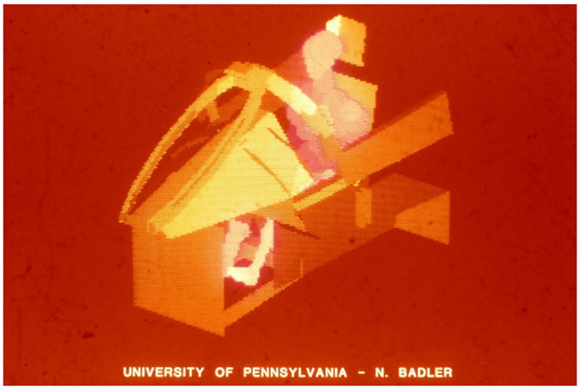

FIGURE 2. Bubbleman in a polygonal cockpit model. SIGGRAPH '80 program credits image and SIGGRAPH slide set 13 (4) #72.

These shadowy but solid-looking human figures had already caught the attention of human factors and ergonomics engineers. By 1980, we received a small U.S. Navy grant to embed Bubbleman into a 3-D polyhedral workspace environment. For human factors engineers, this was the first time they could use color raster graphics, such as shown in Figure 2, to draw solid renderings of how someone would fit, for example, in a cockpit.

NASA AND TEMPUS

With our late 1970s work on human modeling, I was ready to focus on human model improvements. One of my Ph.D. students, Steve Platt, constructed a polyhedral facial model and used physics-based springs to animate the skin mesh. We were starting to see the computational and display advantages of polygons over spheres. O'Rourke, Platt, and Morris summarized our work on human movement understanding for the First Annual National Conference on Artificial Intelligence.[3]

That paper proved pivotal in my research career. After its session, a gentleman named Jim Lewis came up to me and introduced himself. He was from NASA Johnson Space Center (JSC) and said his Crew Station Design group was interested in human model visualization for the Space Shuttle. I soon visited Houston JSC and formulated a plan for a software tool to aid their human factors studies.

In 1981, we received our first grant from NASA JSC. We were to provide an *interactive* software system to evaluate human factors questions such as fit, reachability, and clearance. Among the features we proposed to build were polygonal human models, anthropometric scalability, inverse kinematics limb positioning, and real-time interaction in a 3-D polygonal environment. The resulting system, called TEMPUS, used a color raster graphics device for interaction.

To construct the polygon models, we needed a 3-D modeling system that would support joints and object hierarchies. There was no off-the-shelf modeling system in the lab in 1980, so we had to build our own. Ph.D. student James Korein and Masters student Graham Walters jointly designed their own representation called "Peabody." Objects were first drawn on paper and 3-D coordinates hand-entered into text files. Masters student Jane Rovins constructed our first polygonal human model suitable for the raster display. Built in Peabody and composed of triangles, the androgynous figure used the minimum shapes necessary to show solid forms in the torso, head, and limbs. It had articulated joints centered on pyramidal apexes. Other better-shaped models soon followed.

Within a year of starting work with NASA, we had delivered mostly working software. We delivered on all our promises over the following year. TEMPUS was an interactive tool engineers could now use to check fit, reach, and clearances in the space shuttle cockpit and payload bays with both "shirt-sleeved" and EVA-suited articulated human figures. I now started to realize the research advantages of working with people who had to solve real problems. The Crew Station Design Section at NASA JSC were fantastic partners. Key staff included Barbara Woolford, Geri Brown, Jim Maida, Linda Orr, and Abhilash Pandya. My research team had learned how to create an anthropometrically scalable human model in conjunction with practicing NASA engineers.[4]

The TEMPUS project, our papers, and our successes began to attract further attention in the human factors arena. In November 1982, I assembled a set of papers on human modeling for a Special Issue of *IEEE Computer Graphics and Applications*. Among these contributions was a paper by William Fetter of Boeing, who had popularized the term "computer graphics." Fetter created a line representation and animations of a movable human for human factors analyses.[5] The history Fetter documents in that paper provides an essential early strand of virtual human DNA.

JACK

With ample research funds, I began to attract more graduate students. One particular individual stood out, Ph.D. student Cary Phillips. Phillips arrived at Penn just when we had purchased one of the new Silicon Graphics 1400 Workstations. I had already met Silicon Graphics founder Professor James Clark and looked forward to working with the earliest available system. This was a new beast for us. Because SGI had built-in rendering capabilities, we could concentrate on modeling and interaction in 3-D. I essentially dedicated this workstation for Phillips' use and let him explore it on his own. Soon I realized that my programming days were over; Phillips was just so efficient that there was no need for me to try programming anymore.

Phillips was especially interested in how one might interact in 3-D using only the three-button mouse and keyboard. The corresponding 3-D cursor location was portrayed with a small "*Jack*"-shaped icon ✳ of three short mutually perpendicular line segments that looked like a piece from a kids' game of jacks; hence the genesis of the *Jack* software system. The mouse buttons could be used to select dimensions. Together with keyboard keys, the mouse and buttons initiated global or local translations or rotations within a joint hierarchy. When applied to an articulated human model, interactive posing was easy. Ph.D. student Jianmin Zhao's inverse kinematics code provided an interactive reach capability as a hand that, for example, chased a 3-D input cursor position. Soon the name *Jack* became synonymous with the human figure itself and provided convenient and memorable branding.

Phillips finished his Ph.D. in 1991. During one of my more productive academic sabbaticals, I took his text as the backbone for my book with Phillips and Professor Bonnie Webber as co-authors.[6] This volume was the first computer graphics book to describe human body models in detail. Other Ph.D students contributed work in Natural Language.

Jack software development continued in earnest between 1988 and 1996. One of my opening lines in any talk I gave was, "At the Penn Computer Graphics Research Lab, we divide the human body into graduate students." Since human models had to work in tight and awkward spaces, Masters student Gary Monheit developed an accurate articulated spine model.[7] We had developed the premier interactive human factors modeling and analysis software system of its time. Its fame was spreading among the ergonomics community through talks, papers, publicity, and word-of-mouth. In 1990, we changed the name of the Computer Graphics Research Lab to the "Center for Human Modeling and Simulation" to better reflect our focus and value.

FIGURE 3. Three different size *Jack* models in an automobile cockpit. The range of size variability from 5th to 95th percentile males is evident. The rightmost image is scaled down slightly to fit the image frame.

With fame came jealousy. While some members of the Human Factors community welcomed new interactive computer graphics tools such as *Jack*, not everyone was happy with us "upstarts." I remember meeting a senior human factors researcher at a conference, and he groused about me being a "dilettante." I think he missed the point of *Jack*: we did not *do* human factors experiments, we just made tools to enable analysts to do their jobs. Figure 3 shows three different anthropometrically scaled *Jack* bodies in a vehicle cockpit.

This point came home in another context when we were meeting with General Motors. They were interested in evaluating factory workcell and vehicle cockpit comfort. Therefore, we added some features to provide suitable vision and comfort measures. But they made it clear that *Jack* was not to determine *whether* something was comfortable or not. We were to return numerical data to allow the analyst (or the management) to assess whether the action was indeed "comfortable" to their own corporate standard.

Unlike some other proprietary human models, *Jack* was not bound into any particular CAD system. Geometry converters let *Jack* and his female counterpart, *Jill*, inhabit most any 3-D imported mesh model. Being agnostic to computer environment (as long as one had a Silicon Graphics workstation) was a huge advantage because we could license *Jack* to nearly anyone. As *Jack*'s capabilities grew, we also obtained funding from multiple organizations within the U.S. government and industries. An unexpected obstacle appeared: we were asked why government branch X should also support work that government branch Y was funding. Our response was simply that they get what they fund *plus* all the features that everyone else funds. That nailed it. Everyone received the same unified version.

With so many disparate customers and applications, one of my Ph.D. students, Welton "Tripp" Becket, had a brilliant idea. He wrote a LISP-API for *Jack*, so that all its functionality could be accessed and extended through external software. *Jack* was now not only the premier interactive virtual human system, but it was also embeddable into other applications. The LISP interface was eventually rewritten in C++ by Dr. Paul Diefenbach.

By 1994, I felt as if I was running a small business within the University. At the peak of activity, I had 24 Ph.D. students, four full-time staff programmers, two international licensing agents, multiple licensed sites, ongoing software maintenance fees, and a mailing list in the hundreds. I became "salesman *Jack*" and raked up considerable travel to license the system to new users.

Jack became a generic brand. I had various people tell me they saw *Jack* being used when in fact I knew that they were talking about someone else's human model. This success was both great and terrible. The time was ripe to move *Jack* out of the University and into its own company.

TRANSOM TECHNOLOGIES

The spin-off process began with Penn's Center for Technology Transfer. They had been involved around 1990 and had offered a license to all the *Jack* intellectual property to a robotics company for about $60K. I was relieved when I found out that the company rejected the offer. I felt that *Jack* was worth far more.

By 1994, a new effort began to create a spin-off company rather than sell it outright. Working through a local venture capital firm, we realized that the window to make a deal was starting to close. We were not secret about our algorithms and gave out free trial licenses, so it would not take that long for a powerful corporate interest to reproduce a *Jack*-like system on their own. We already knew of at least one company that was actively engaged in making a *Jack* clone because we kept running into them and their demos at the same trade shows and meetings.

The process of generating a spin-off company in the software space had never been attempted at Penn. Every step of the process was novel; every aspect of the eventual licensing agreement had to be generated from scratch. The venture firm folks and Penn staff knew that the spin-off required finding an able CEO to run it. I quickly realized that being a CEO would be different than being a University faculty member, so I did not want to be a candidate. In fact, I really wanted to disengage completely. I was longing to get back to being a researcher and teacher and retire from the salesman role.

The CEO search started, and eventually they identified an MBA from the University of Michigan, Jim Price. We arranged a *Jack* user meeting at Penn and invited all our sponsors and licensees. It was not the first time we had done such a user meeting. However, this one's real purpose was to give Price a firsthand opportunity to hear from and converse with our customers and see their applications. By the end of the day, Price was convinced that *Jack* had "legs." Jim agreed to become the CEO of a new company he would form called "Transom Technologies."

Price wanted to keep the Transom operation in his hometown of Ann Arbor, MI, USA, but all the knowledgeable *Jack* programmers were at Penn. He opened a satellite office in Philadelphia, PA, USA, and hired several of my student-based staff to maintain continuity. The choice to keep the main office in Ann Arbor was ultimately a smart one because there was a strong local ergonomics program at the University of Michigan. Price hired Dr. Ulrich Raschke to join Transom. *Jack* finally had an "authentic" human factors researcher on staff. *Jack* had now grown up and moved to Transom in 1996. *Jack* now had a corporate home in both Ann Arbor and Philadelphia.

In 1998, Transom was sold to Engineering Animation (EAI). EAI was bought by Unigraphics, and then Unigraphics was acquired by UGS, a General Motors subsidiary. Not long after, Siemens bought UGS and integrated *Jack* into its product lifecycle management tools. *Jack* is used extensively today for large-scale factory and workplace design and evaluation.

Jack was certainly not the first virtual human model. Fetter's line models date from 1964. Solid polygon hands and faces were developed in the 1970s by Ed Catmull and Frederic Parke at the University of Utah. Nadia Magnenat-Thalmann, Daniel Thalmann, and colleagues were creating animated characters, including a virtual Marilyn Monroe. Animation studios explored 3-D characters, including Jeff Kleiser and Diana Walczak's 1988 "Nestor Sextone for President." What made *Jack* unique in the early 1980s was its grounding in the scientific communities of anthropometry and human factors, its color graphics and on-screen interactivity, and a growing user community.

FROM HUMAN MODELS TO VIRTUAL BEINGS

It would be disingenuous to claim that the *Jack* human models were the precursor to the realistic virtual beings of today. As one strand of the DNA, however, we did contribute on a number of interactive, structural, and algorithmic fronts. By the 1980s, many other researchers, corporations, and artists were making equally important advances and contributions to the virtual human DNA pool. Motion sensing devices came on the market that allowed real-time human motion capture. Artists vastly improved surface, face, hair, and clothing models. A wide range of artificial intelligence and cognitive science researchers were trying to add human sensing, thinking, reasoning, planning, gestures, and emotions to graphical body models. New conferences sprung up, such as autonomous agents and multiagent systems (AAAS), intelligent virtual agents (IVA), and symposium on computer animation (SCA). The human factors community began the digital human modeling conference (DHM). Artists and graphics practitioners, as well as researchers, excitedly reported their progress at the ACM SIGGRAPH Conference. The creation of modern digital humans was an interdisciplinary, integrative, and evolutionary process. The first time these interconnections were seriously recognized was in an annual series of three meetings in 1996–1998 simply called, "Virtual Humans."

FIGURE 4. George (left) and Gilbert (right) have an animated discussion about cashing a check at the bank.

One DNA thread can be traced to a seminal paper led by Justine Cassell and presented at SIGGRAPH 1994.[8] This paper and video, two frames of which are shown in Figure 4, demonstrated two *Jack* human models, George and Gilbert, conversing, gesturing, and speaking to each other about cashing a check at a bank. Catherine Pelachaud created the animated faces. Professor Mark Steedman designed the dialog planner. Ph.D. students Scott Prevost and Matthew Stone worked on the speech synthesis and language components. Ph.D. student Brett Achorn built the gesture engine. Ph.D. students Tripp Becket and Brett Douville managed the graphics. The LISP-API played a key linking role. We had demonstrated that many crucial language, planning, and graphics components could be integrated to create a digital human.

JACK JOINS THE MILITARY

Although we had conceived of *Jack* as an ergonomics and human factors model, its real-time interactivity opened doors to the world of simulations. Military simulations for vehicles and aircraft had been in existence for a while, but they lacked realistically modeled people. As compute and display capabilities improved, scaling up content and increasing simulation modeling granularity (down to the individual level) became desirable. A special government office was set up to coordinate efforts across the various branches: The Defense Modeling and Simulation Office (DMSO). DMSO produced specifications and standards to ensure scalability, interoperability, and extensibility as simulations incorporated new weapon systems and additional people.

Key to distributed interactive simulation (DIS) were protocol data units (PDUs). DIS allowed multiple computers to run independent but coordinated aspects of a large-scale simulation, thus obviating the need to do everything on one massive machine. The PDUs communicated state to entities across the network. For example, the PDUs for a soldier might be walk, run, kneel, drop to prone, crawl, rise to stand, aim gun, fire, run, and be dead. Parameters specified speed and direction. This was a nice control system: PDUs were issued to virtual humans who executed and animated the desired state changes and actions.

Jack became the ideal 3-D virtual soldier. We changed his clothing to Army fatigues. We already had a walking locomotion model.[9] With *Jack*'s real-time inverse kinematics, he could hold a gun with two arms and aim it any desired direction. The only animations we had to hand-build and add were a crawling motion and the posture changes from standing to kneeling, standing to prone, and vice versa. The PDUs were essentially mapped to the *Jack* API. We could transition from any state PDU to any other by animating arcs in a "posture graph."[10] For example, to transition from a walk to a crawl, *Jack* would stop, go prone, then crawl. The posture graph ensured that all transitions made sense and pushed the performance burden onto *Jack* rather than the PDU-issuing agent.

The culmination of our military simulation effort was a demonstration for the 1994 Interservice/Industry Training, Simulation and Education Conference (I/ITSEC). Professor Mike Zyda and colleagues at the Naval PostGraduate School in Monterey, CA, USA, and Sarcos Robotics in Salt Lake City, UT, USA, coordinated with us.[11] The demonstration setup is shown in Figure 5. Sarcos built a virtual reality platform called iPort so the user could virtually navigate the 3-D environment and see it through a VR display. The field of view was controlled so that the forward view was always centered even if the iPort was turned. A servo constantly and gently rotated the seat back to the neutral, forward position. Of course, the point of the demo was to find and track down the PDU-controlled *Jack* virtual humans who appeared, ran, hid, and shot at the operator.

FIGURE 5. The iPort soldier interface. Three virtual *Jack* models are visible, each dressed in camouflage. On the left is the user's avatar, showing his position and pose in the interface. This avatar was movable anywhere in the scene by user swiveling and pedaling in the iPort. In the upper right, two *Jack* "combatants" are visible targets for the military simulation. Image courtesy of Mike Zyda.

The Naval Air Warfare Center Training Systems Division (NAWCTSD) began embedding *Jack* in its systems. To remedy *Jack*'s lack of certain realistic motions, NAWCTSD contracted with a private company, Boston Dynamics, to custom-build and PDU animate an alternate human model. That model was called DI-Guy. DI-Guy soon took over all the *Jack* military simulation implementations by providing a superior range of clothed models and better movements. The DI-Guy motions were actually motion captured from soldier subjects. I found out some years later from Marc Raibert, founder of Boston Dynamics, that the original DI-Guy motion capture sessions were run by one of my former undergraduate students. He had learned to use our Ascension motion capture system at Penn at a time when few such systems were available in university labs. More ironic than tragic, I eventually appreciated that Boston Dynamics did us a great favor by doing the necessary virtual soldier technology transfer. Doing so kept us out of a costly maintenance and support mission.

I only saw *Jack* in an actual DIS setting once more. I was in a conference room watching a small 3-D battle. A *Jack* soldier was standing a short distance away from a burning tank. I was talking to my host and not paying much attention to the screen since he was no longer giving control commands. Suddenly the *Jack* soldier took several steps away from the tank. "Oh," my host said, "he started to feel the heat from the fire and moved away." I knew *Jack* was a computer graphic model, I knew the fire was an animation, and I knew *Jack* did not have any reasoning ability. But at that moment all that I knew to be true was vanquished by the sheer naturalness and human-like response *Jack* showed to an external stimulus. The era of building spontaneous human behaviors into simulated beings had begun.

LESSONS LEARNED

Over my 50-year career in academe, I managed to learn a number of skills that were not taught in school. Besides being fortunate to get started early in the evolution of computer graphics, I gained an understanding that effective research required management, human resources, and marketing skills. I learned more from my students than they did from me. Especially at a university such as Penn, where Ph.D. students are primarily supported by research funds, being able to obtain and maintain funding is a critical requirement for any multiyear, multicollaborator enterprise.

One of the most helpful principles for me was to have at least one very long-term goal. The goal had to be easily articulated and ripe with potentially interesting problems and approaches (e.g., "how to make a computer 'understand' human movement"). This was no all-or-nothing "moonshot." It defined a goal and not an absolute requirement. Along the way of trying to achieve that goal, my students and I learned a lot about ancillary but crucial parallel fields—sometimes outside computer science—such as anthropometry, movement notations, dance theory, cognitive science, human physiology, and human factors engineering. We never received funding to "just" do computer graphics. Support was always to further multiple developments toward building interactive virtual people, mostly for engineering use and evaluation. We learned to listen to what experts and users wanted to do but could not or found difficult to do, such as building physical cockpit mock-ups. Publishing our work in venues other than just the top conferences or journals in computer science was necessary to reach specialized communities to whom our work would matter.

As an educator as well as a researcher, I learned to motivate students (both undergraduate and graduate) by giving them interesting and challenging problems rather than just telling them what to do. Usually, better solutions and approaches were their own doing. I empowered them and helped them confront novel situations later in their own careers. The one quote I kept on my desk for my entire Penn career was from General George Patton: "Never tell people how to do things. Tell them what to do, and they will surprise you with their ingenuity."

Finally, I realized—several times—that the best long-term strategy was to let go of our work. Launching *Jack* as a start-up was just one obvious case. There were various projects where we ceded incipient ideas and territory to others for their exploitation, such as the DIS *Jack* that motivated DI-Guy. Ph.D. students graduated and took their systems with them. We avoided the Pygmalion effect and did not fall in love with our own creations. New ideas had room to flourish.

REFERENCES

1. N. Badler, J. O'Rourke, and H. Toltzis, "A spherical representation of a human body for visualizing movement," *IEEE Proc.*, vol. 67, no. 10, pp. 1397–1403, Oct. 1979.
2. N. Badler and M. Morris, "Modelling flexible articulated objects," in *Proc. Comput. Graphics, Online Conf.*, 1982, pp. 305–314.
3. N. Badler, J. O'Rourke, S. Platt, and M. Morris, "Human movement understanding: A variety of perspectives," in *Proc. 1st Annu. Nat. Conf. Artif. Intell.*, 1980, pp. 53–55.
4. M. Grosso, R. Quach, and N. Badler, "Anthropometry for computer animated human figures," in *State-of-the Art in Computer Animation*, N. Magnenat-Thalmann and D. Thalmann, Eds. Berlin, Germany: Springer, 1989, pp. 83–96.
5. W. Fetter, "A progression of human figures simulated by computer graphics," *IEEE Comput. Graphics Appl.*, vol. 2, no. 9, pp. 9–13, Nov. 1982.
6. N. I. Badler, C. B. Phillips, and B. L. Webber, *Simulating Humans: Computer Graphics, Animation, and Control.* London, U.K.: Oxford Univ. Press, 1993.
7. G. Monheit and N. Badler, "A kinematic model of the human spine and torso," *IEEE Comput. Graphics Appl.*, vol. 11, no. 2, pp. 29–38, Mar. 1991.

8. J. Cassell et al., “Animated conversation: Rule-based generation of facial expression, gesture and spoken intonation for multiple conversational agents,” in *Proc. 21st Annu. Conf. Comput. Graphics Interactive Techn.*, 1994, pp. 413–420.
9. H. Ko and N. Badler, “Animating human locomotion in real-time using inverse dynamics, balance and comfort control,” *IEEE Comput. Graphics Appl.*, vol. 16, no. 2, pp. 50–59, Mar. 1996.
10. J. P. Granieri, J. Crabtree, and N. Badler, “Production and playback of human figure motion for visual simulation,” *ACM Trans. Model. Comput. Simul.*, vol. 5, no. 3, pp. 222–241, Jul. 1995.
11. D. Pratt et al., “Insertion of an articulated human into a networked virtual environment,” in *Proc. Conf. AI, Simul. Plan. High Autonomy Syst.*, 1994, pp. 84–90.

NORMAN I. BADLER is the head of research with Cesium GS in Philadelphia, PA, 19106, USA. His interests span computer animation, human modeling, cultural heritage, and interactive systems. Badler received his Ph.D. degree from the University of Toronto on “Temporal Scene Analysis”. He is a member of IEEE Computer Society and ACM SIGGRAPH. Contact him at norm@cesium.com or badler@upenn.edu.

Contact department editor Dave Kasik at dave.kasik@gmail.com, department editor Chris Johnson at crj@sci.utah.edu, or department editor Mary C. Whitton at mcwhitton@gmail.com.