from *SIAM News*, Volume 36, Number 3, April 2003

Step by Step, Dancer/Scientist Choreographs A Unique Career

By Erica Klarreich

Kathy Pullen's saga has all the elements of an inspirational Hollywood tearjerker. The talented young dancer who dreams of making it big in New York City. The knee injury that forces her to give up dance, for good, she thinks. The gradual recovery and reentry into the world of dance.

And, oh yes, a PhD in physics from Stanford University along the way.

At first glance, Pullen's two loves—dance and physics—seem totally disparate. But when the knee injury blocked Pullen from dancing, she ended up forging a remarkable synthesis of the two fields. In her dissertation, completed in July, Pullen offers a solution to one of computer animation's thorniest problems: how to endow the mechanical motion of a computer-animated human with the nuance and subtlety of real human motion.

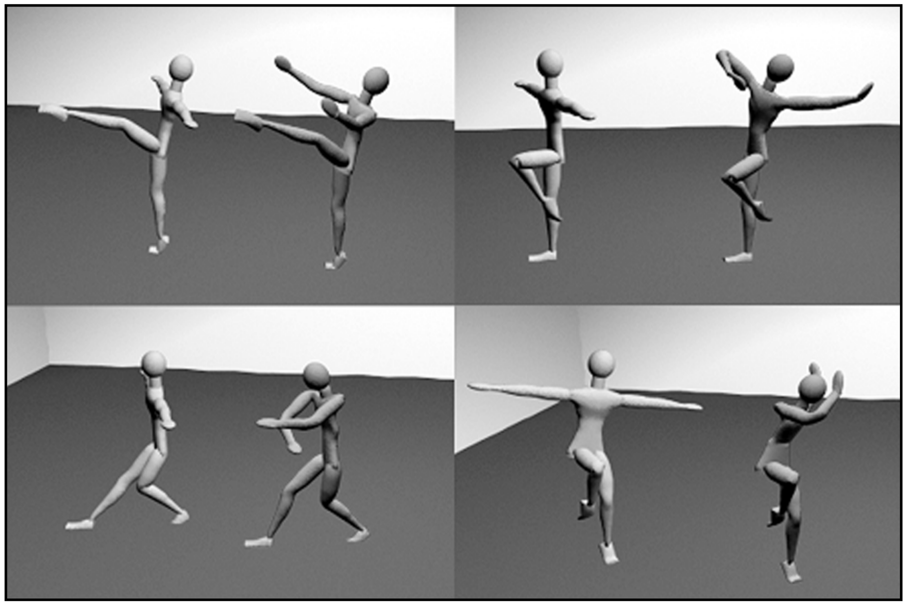

Kathy Pullen's algorithm uses motion-capture data to add nuance to an animation. In each of these frames, the figure on the left is a computer-generated animation in which only the movement of the lower body has been choreographed. For the figure on the right, Pullen used motion-capture data to generate an appropriate motion for the upper body, and to layer "texture" onto both the upper and the lower body.

For her dissertation research, Pullen choreographed and performed a modern dance, using motion-tracking equipment to capture her exact movements. She also choreographed—but didn't perform—a new dance. She then used traditional computer techniques to make a simple, no-frills animation of the new dance. Finally, she designed an algorithm that would allow her to add motion-capture data from the performance, and thus her own personal style, to the animation. The result is a graceful dance whose feel is uncannily like that of the dance she actually performed.

"When I watched the animation, I thought, 'Wow, that looks like me dancing, but it's not something I did,' " Pullen says. "It was freaky."

Two Camps of Animators, Two Very Different Mindsets

Human motion is fiendishly difficult to animate. Capturing a particular individual's style is especially challenging: People are so attuned to bodily cues that they will notice instantly if a movement is even slightly off. "When you see a friend walking down the street, you can recognize who it is instantly from the way he walks, even if you can't see his face, because different people move in such different ways," says Christoph Bregler of New York University, who was Pullen's dissertation adviser. Bregler discussed Pullen's research as part of the I.E.Block Community Lecture he gave in July at the SIAM 50th Anniversary Meeting.

The animation community's approach to human motion is polarized around two extremes. Some companies, such as Disney and Pixar, use artists to design animations frame by frame. Others, such as Industrial Light & Magic and Electronic Arts, take motion-capture data, obtained from tracking an actual human movement, and incorporate it directly into an animation. While both approaches work reasonably well, each has significant drawbacks.

Animations created by artists can be full of nuance and individuality, but the process is labor-intensive and agonizingly slow. In traditional animation, the artist first maps out the motion to be created, then sketches the "key frames"—the defining moments of the motion. In an animation of a human walking, for instance, key frames might include one in which one leg passes the other, another in which a foot hits the ground. The animator next fills out the motion by drawing the frames in between the key frames.

Nowadays, animators frequently save some labor by drawing fewer frames and using a computer to interpolate between them. But computer assistance can go only so far. Computer interpolations are smooth curves; human motion is irregular and noisy. If the hand-drawn frames are too far apart in time, the interpolated motion will seem robotic. To create realistic human motion, animators must still design by hand many frames in between the key frames. Consequently, even with the aid of computers, it can take a team of animators days to create a short phrase of motion for a single character.

Instead of artists' sketches, the other approach to animation relies on motion-capture data, which is usually collected either by reflecting light beams off markers on an actor's body or by having the actor move through a magnetic field with sensors attached to his limbs. Motion-capture data instantly produces lifelike motion, and easily conveys the style of the specific person whose motion has been tracked. In exchange, however, the animator gives up some control over the animation. Once the motion has been captured, it is difficult to alter; if an animator changes her mind about what she wants, another motion must be captured from scratch. And because animators generally start from key frames, they often find it hard to incorporate a captured motion whose details they haven't been able to control at every step. "There are many professional animators who would never touch motion-capture data because they want to keep their control," Pullen says.

The two camps have very different mindsets, Bregler says. "One of the agendas of our group is to blur the boundaries, to take the best qualities of each of the opposite extremes and see what we can do with that."

Pullen's project epitomizes this approach. Her algorithm starts with the animator sketching out the key frames, which gives the animator flexibility and control over the motion. The program then uses motion-capture data to layer what Pullen calls "texture" onto the animation, quickly achieving the subtleties of style that normally take such intensive work on the part of animators.

An Edge Recaptured

Pullen never expected to work on computer animation. A physics and dance major in college, she intended to move to New York after graduation and immerse herself in the world of dance. But her knee problems forced her to reassess that plan. "I'd had four knee surgeries, and my body was just falling apart," she says.

Dance was an all-or-nothing affair for Pullen. "If I couldn't do it eight hours a day, I didn't want to do it at all," she says. Because she "loved the interaction between light and matter," she decided to enroll in a graduate program in optics at Stanford. But after a couple of years, her adviser, Mark Kasevich, accepted a tenure-track offer from Yale University, and Pullen didn't want to follow him there. At the same time, something else happened: Her knee started to recover.

Pullen joined a biophysics group that was using lasers to track single molecules in cells, and enrolled in a ballet class. It was now three years since she had set foot in a dance studio. "I thought maybe now I could do it just for fun, that I had lost that edge, that need to do it constantly," she says. But dance quickly reasserted its claim. "I had thought it was impossible, but before long I was back to the level where I'd been when I quit dancing," she says. Dancing six hours a day, she soon began to neglect the biophysics experiments. "My heart wasn't in it," she says.

Pullen never intended the algorithm to be fully automatic. "I always wanted it to be something the animators explore back and forth with, that helps them with their creative process."

Pullen was torn. She thought about taking a leave of absence and dancing full time, but hated the idea of turning her back on so many years of scientific study. Then some friends suggested that computer animation might be a way to combine her interests in physics and movement. Pullen started reading papers about animation on her own. "I hadn't even known what motion capture was," she says. "I knew nothing about computer science." By the end of a summer of reading, Pullen had a good idea of what she wanted to do. Bregler was about to join the computer science faculty at Stanford, and Pullen decided to become a member of his fledgling movement group.

"Before Chris was even there, I had moved into an office in the computer science department," she says.

An Individual's Style Layered onto an Animation

In music, when an instrument plays or a vocalist sings, each note consists of a base pitch, which follows the line of the melody, and overtones, the higher frequencies that give the instrument or voice its characteristic sound. In a similar way, human motion can be decomposed into frequency bands that represent different aspects of the movement; as in music, the higher frequencies are responsible for an individual's style.

A motion can be described by graphs representing changes over time in the angles of the elbows, knees, and so forth. Pullen's algorithm uses a Laplacian pyramid to decompose each graph of a joint angle; the result is a hierarchy of graphs, each corresponding to a different frequency band. Pullen's algorithm layers an individual dancer's style onto an animation by adding the high frequencies from motion-capture data to the animation's joint angles.

The process starts with an animator's key frames. A computer interpolates between the frames to produce a motion along the lines the animator had envisioned. As a smooth computer interpolation, however, the result lacks the subtlety of human motion.

In theory, any bit of motion-capture data could then be used to texture any bit of the animation. But in practice, the resulting animation is most pleasing to the eye and most like real human motion if each bit of the animation is textured by a bit of motion-capture data that corresponds to a similar movement. For instance, to texture a motion in which a person is stomping his foot, it is best to use motion-capture data of a stomp, if such data is available. Thus, one of the main tasks for Pullen's algorithm is to decide how to pair fragments of the animation with fragments of motion-capture data.

Pullen's algorithm chops the animation and the motion-capture data into tiny snippets by making a break every time the motion of one of the joints changes direction. All the snippets are rescaled to take the same amount of time. One of the lower frequencies is selected for matching purposes, and the snippets of animation are compared with the snippets of motion-capture data at that frequency. For each snippet of the animation, Pullen's program finds the K snippets of motion-capture data that match it most closely (that is, the ones that minimize the sum of the squared differences between the snippet of animation and the snippet of motion capture). Different values of K, a parameter that can be chosen by the animator, will produce different animations.

The idea, ultimately, is to choose a single snippet of motion-capture data to match each snippet of the animation. Pullen's algorithm finds not one but K matches, to give the algorithm the flexibility needed to carry out the following step. Pullen found that the results of the texturing were most pleasing when, as often as possible, consecutive snippets of the animation were textured by consecutive snippets of motion-capture data. So the next step is to look through all the possible ways of matching each snippet of animation to one of its K best matches, to find the matchings that connect as many consecutive snippets of motion-capture data as possible.

The process results in one, or perhaps several, possible pathways through the database of motion-capture snippets. Any of these paths can be used to texture the animation, and it is up to the animator to decide which one produces the most pleasing result.

For the actual texturing, the algorithm simply replaces the high frequencies of the animation with the high frequencies from the matching snippets of motion-capture data. Finally, because the animation snippets may no longer fit together smoothly, the algorithm smooths each transition from one snippet to the next by blending the graph with its best quadratic fit in a small neighborhood of the transition.

The animator has many choices to make as the algorithm progresses—which low frequency to use for matching, which joints to texture, and how many matches to keep for each snippet of the animation. By changing such parameters, the animator can control completely the degree to which the texturing process alters the motion he sketched, and can choose the textured animation that suits his purpose best. Pullen never intended the algorithm to be fully automatic. “I always wanted it to be something the animators explore back and forth with, that helps them with their creative process,” she says.

Pullen tested her algorithm not only on a modern dance, but also on two other animations. For the first, a cartoon of an otter, she used what she calls a “funky” walk to give the otter a distinctive bob of the head and twist of the spine as it walked. In the second animation, Pullen textured each of two walkers from a different set of motion-capture data. The resulting walks are clearly distinguishable: One walker ambles along in a fairly ordinary way, while the other moves with the mincing step of a drag queen.

Pullen’s algorithm can be used not just to add nuance to existing motion, but also to synthesize new motion. In one application of the method, Pullen choreographed the leg motions of a dance but didn’t design any arm movements to go with them. As usual, she used the high frequencies from motion-capture data to add texture to the joint movements of the lower body. But for the upper body, she used not just the high but also the low frequencies of the motion-capture data. The result was an animation in which the figure danced with her whole body, often making arm gestures Pullen hadn’t anticipated. “It was things I would have done, but put together differently,” she says. “I thought, ‘I can get choreography ideas from this.’ ”

In July, Pullen described her research in a talk at the SIGGRAPH meeting—the main professional meeting for computer graphics experts. Afterward, SONY and Industrial Light & Magic discussed bringing her in as a consultant. But Pullen’s goals for the immediate future are not very computer-oriented. A little consulting work might be fun, she concedes. For now, though, her main plan is to join a circus troupe.

“What I really want to do is dance,” she says.

Erica Klarreich is a freelance writer based in Berkeley, California.