Scientists in Pixar's New Looks Project work out a long-standing CG problem

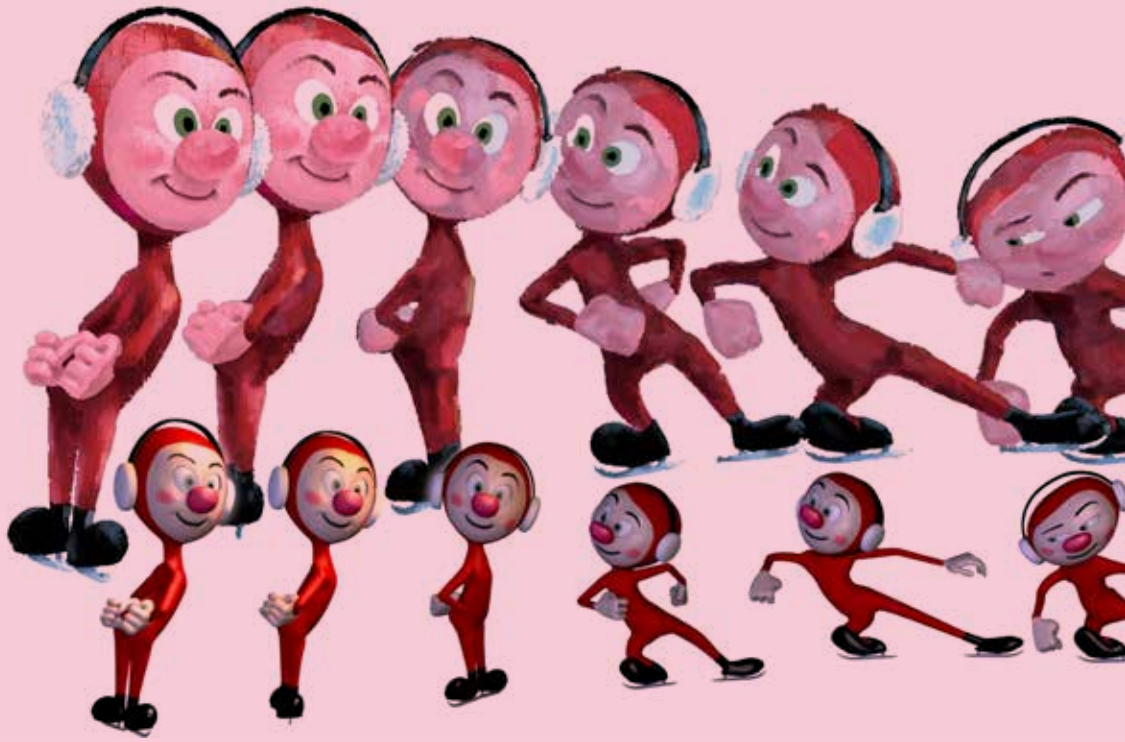

SMOOTH M

By Barbara Robertson

On July 24, at SIGGRAPH 2013, a team of CG scientists led by Pixar's Michael Kass will introduce technology that could change the look of animated feature films.

It has been 18 years since Pixar's *Toy Story*, the first 3D animated feature, made computer graphics history. In the years since, CG films have grown more visually complex with greater numbers of characters, locations, and simulations, and with more realistic lighting. But, no matter how photorealistic 3D animated features have become, even the most recent films aren't that far from *Toy Story* in style.

Most CG characters in animated features have a brightly colored cartoon look or, more rarely, a softer style that imitates storybook characters. Exceptions such as Pixar's *Wall-e*, Animal Logic's *Owls of Ga'hoole*, and ILM's *Rango* showed us characters with a grittier look that, in the case of Animal Logic and ILM, drew on those visual effects studios' experiences in fitting CG characters into live-action films. Other 3D features have teetered on the edge of the uncanny valley with performances captured from actors applied to CG characters designed to look more or less like humans.

But, none of the characters in animated features, whether beautiful, fanciful, cute, gritty, or gross, are painterly. Nor are the stars of hand-drawn animated features, where simply painted characters perform against lush backgrounds. None look like oil paintings or watercolors.

When it comes to portraying CG film stars with a non-photorealistic or illustrative style, all animated features have the same limitation.

It's a technical problem.

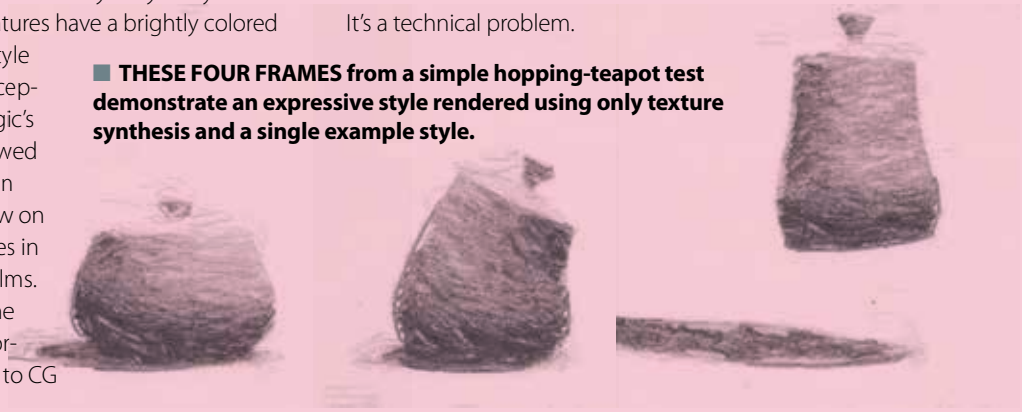

THESE FOUR FRAMES from a simple hopping-teapot test demonstrate an expressive style rendered using only texture synthesis and a single example style.

THE TWO ARTIST-PAINTED keyframes in the top row, at far left and right, provided the style for the automatically generated in-betweens. The bottom row shows the original CG-rendered sequence that guided the interpolation.

Video: Go to "Extras" in the July/August 2013 issue box

"The paper describes two important contributions," Kass says. "One is the temporal coherence, the fluidness of images from frame to frame, and the other is the artistic control that you get by painting keyframes."

Pixar Senior Scientist Kurt Fleischer, who served as a technical director for the project, describes its evolution. "Michael [Kass] was working on image filtering when some people in production became interested in the artistic images. That was the genesis. We experimented with different kinds of techniques that looked nice and continuous over time, but we didn't have artist-specific control. We only cracked that in the last couple of years."

A SIGGRAPH 2001 paper titled "Image Analogies" by Aaron Hertzmann, Charles Jacobs, Nuria Oliver, Brian Curless, and David Salesin provided one starting point. In that paper, the authors introduced the idea of using two images to define a style transformation.

"Their work was with still images, and if you did that transformation on a series of images, you'd get temporal incoherence," Kass says. "So, one of the first things we did was to try to make that approach work in animation, to make it work over time."

Making that leap to animation was difficult.

"Michael [Kass] shepherded the work over several years with graduate students and postdoctoral students," Fleischer says. "Toward the end, it was Forrester [Cole] and Pierre [Bénard]. There was a nice moment when Pierre and Forrester were here at the same time. I felt like we had critical mass."

Artistic Control

About two years ago, the technology had moved far enough toward solving the problem of temporal coherence that the team could consider the second goal: building in artistic control.

"The project was at an advanced stage when I started working on it two years ago, but the art direction aspect hadn't happened yet," says Cole, then a postdoc student who joined Pixar in July 2011. "We wanted to move it from a research prototype."

To provide the art direction essential to move the technology from the lab into production, the team decided to introduce keyframes.

"We use select keyframes from a traditional 3D animation that are lit and textured in a simple way," Kass says. "The artists can paint over them in any way and the technique comes up with a series of in-betweens that maintain the style while achieving temporal coherence."

In practice, an artist might select keyframes from a simply rendered animated sequence, load them into Adobe's Photoshop, and paint over them to create an output layer. "The simple renderings and the painted images form style transformation pairs," Cole explains. "There can be an arbitrary number. You might have 10 or 20

OVES

Jitter Begone

"One of the big challenges with any non-photorealistic rendering (NPR) technique is creating temporal coherence," says Michael Kass, senior scientist at Pixar. "If you take any style and process the frames individually, you get jitter, shake, boils, and pops. You need coherence between frames. That's why hand-drawn animation with texture has excess movement. Each frame is slightly different. You might want a little difference for artistic purposes, but with traditional techniques, you can't avoid a huge amount. For a film the length of a feature, it's difficult to watch. You want something easier on the eyes."

Now, Kass and a team of scientists in Pixar's New Looks Project group have replaced that problem with an opportunity. An opportunity for artists and animators to create CG features in entirely new, expressive styles.

At SIGGRAPH, the team will show an animation clip in which a 3D character that looks as if an artist had painted it with watercolors, smoothly skates through an atmospheric background. There are no jitters, shakes, boils, or pops.

The video accompanies a landmark technical paper titled "Stylizing Animation By Example" by Pierre Bénard, Forrester Cole, Michael Kass, Igor Mordatch, James Hegarty, Martin Sebastian Senn, Kurt Fleischer, Davide Pesare, and Katherine Breeden.

AT LEFT, A FRAME of Charles Muntz from Disney/Pixar's *Up* with simple CG shading. At right, the same frame rendered using a Rembrandt etching as the example style.

different keyframe pairs in a sequence of 400 or 500 frames. All the frames not painted will be transformed by the algorithm."

This means that artists needn't draw on every frame to produce an illustrative style; they simply paint a few keyframes. The technique takes care of the in-betweens. And, the temporal coherence.

Under the Hood

"One of the underlying existing technologies we use is texture synthesis," Kass explains. "The idea is that here's a texture, make more of it. More specifically, we use guided texture synthesis: Make more of it and let me tell you where it should be and roughly how it should look. That's one of the foundation technologies. We use it to maintain the style of the in-betweens."

"Then, since we have an underlying 3D animation," Kass continues, "we render a velocity field for every frame that tells us where every pixel came from and where it will go. We use that information to modify traditional texture synthesis algorithms. The textures we synthesize have the smooth temporal textures we need."

There's more.

"We use the painted keyframes as constraints on the texture synthesis," Kass says. "We have certain pixels on some images the artist painted that we lock down. The texture synthesis can't change those. Essentially what the technique does is try zillions of ways to chop up the images and rearrange them to come as close as possible toward meeting goals."

There are a series of goals. One goal is to make points on different frames corresponding to the same point on an object look as much like each other as possible. Similarly, another goal involves making regions of final frames look like parts of the corresponding examples from keyframes. The technique combines the series of such goals as these into a final score with an aim toward achieving the highest score possible.

"The underlying 3D animation carries it along," Fleischer says. "It might start with a character that doesn't have much texture or lighting, and from that derive position and motion. The velocity vectors determine how to move the pixels from frame to frame."

The solution isn't calculated one frame at a time, however. The technique tries to solve the whole problem – all the frames in the sequence.

"The technique makes a series of sweeps over the entire animation sequence, going from beginning to end, back and forth, until it gets a solution for the whole sequence," Kass says.

In Practice

With the current algorithms, a 30-second sequence might require two hours to calculate and render, although that time would depend in part on how the artists organize the project. "The artists can break the sequence into layers and process the layers independently," Forrester says. "The key reason for that is because we want to give them freedom."

For example, using layers, an artist might add keyframes only for the eyes or create outlines that extend beyond the edges of the underlying 3D animated sequence to create a looser style.

"To do this well, you need to run rendering tests to get an idea of what the style looks like," Kass says. "We started with the animation and then painted the keyframes because if you change the animation, you need to change the painting."

The idea for the skater in the animated clip created to demonstrate the technique came from Director Teddy Newton, who had combined 3D and hand-drawn animation techniques for his Oscar-nominated short film "Day & Night." In the clip, the lighting on the skater changes as he glides through a spotlight.

"In the underlying CG animation, the skater gets brighter and yellower as he moves through the light," Fleischer says. "When we do the synthesis, that underlying rendering affects the transitions, so he gets brighter."

With this technique, artists can paint the illumination on keyframes, choose to have underlying lighting contribute to the final style, or use some combination of both.

"You might want to have different styles for different elements or for part of a scene," Kass says. "You can have one style transformer, or rather than one abstract style transformer, you can nail styles to particular elements. The system can work with a single keyframe, or you can add more and more keyframes."

Now that the technique is ready for its close-up, the team hopes that someone at Pixar will adopt a painterly style for their next short film. That's the next test.

"That's what we've been hoping for from the beginning," Fleischer says.

And then, perhaps, we will see CG features in many new artistic styles, at last. ■ CGW

Barbara Robertson *is an award-winning writer and a contributing editor for* CGW. *She can be reached at BarbaraRR@comcast.net.*