

CGI INTERVIEWS TECH VFX VR/AR/MR

Paul Debevec: A Name You Absolutely Need to Know in CG, VFX, Animation, and VR

By IAN FAILES | 12/01/2016 12:29 pm |  1

If you've been involved in lighting and rendering in visual effects, or in the creation of cg characters, then you've probably heard of Paul Debevec. But even if you haven't, you're still likely to have seen the influence of his computer graphics research in many different places.

Several films, including *The Matrix*, *Spider-Man 2*, and *Avatar*, along with games and real-time rendered content, have drawn upon Debevec's insights into virtual cinematography, image-based lighting (IBL), and the crafting of photoreal virtual humans. He even helped 3D-scan President Obama with a version of the famed light stage capture system.

Much of this work was undertaken at the University of Southern California Institute for Creative Technologies (USC ICT), where Debevec is now an adjunct research professor. In June this year he began as a senior staff engineer in the GoogleVR Daydream team. It's here that Debevec is working at the intersection of virtual reality and real-time rendering.

PARTNER MESSAGE

Next week, Debevec will be a keynote speaker at **SIGGRAPH Asia 2016** in Macau, a conference that Cartoon Brew is also attending. In this Q&A, we caught up with Debevec for a look back at why his research has been such a major influence in computer graphics, animation, vfx, and vr. He even has a few tips for how to deliver a compelling academic presentation.

Cartoon Brew: Your research into virtual cinematography with your model *Chevette* and then *The Campanile Movie* ended up having a direct connection with Hollywood visual effects. Can you talk about what you thought 'virtual cinematography' could be back at that time?

Paul Debevec: I originally thought of 'virtual cinematography' as creating a dynamic camera shot moving through a realistic scene first by digitizing the real scene through photogrammetry [a process that involves using multiple photographs taken of a single scene to reconstruct a CG model] and then using image-based rendering to animate any camera motion desired. My first idea to do this came while writing stereo reconstruction algorithms in Prof. Ramesh Jain's computer vision class at the University of Michigan, and I finally saw it through in 1991 by making a photogrammetric model of my *Chevette* and animating it to fly across the screen like the DeLorean in *Back to the Future*.






For *The Campanile Movie*, we used the [Façade system I presented at SIGGRAPH 96](#) to create a textured model of the Berkeley campus and also to matchmove our rendered footage to the live-action campus and the cardboard model I was walking around with. When the vfx supervisor for *The Matrix* movies John Gaeta saw our film at the SIGGRAPH 1997 Electronic Theater, he realized this technique could solve how to render matchmoved virtual backgrounds in the 'bullet-time' shots, and hired a few of the students from our Berkeley team to make it happen. It was completely amazing to see the effect woven into such a great visual effects movie.

How did image-based lighting become one of your passions? What were some of the first things you were able to demonstrate with your research in this area?

Paul Debevec: I was amazed how visual effects companies like ILM were able to make digital characters like the dinosaurs in *Jurassic Park* look like they were really there with the actors in the live-action photography, since this is a key element of movie magic. From my computer vision classes it was conceivable how you could match the virtual camera motion to the real cameras, but getting the lighting to match was a mystery. When I asked some industry contacts how it was done, they said sometimes they'd try to write down where the lights were in the scene and replicate them in cg, and then iterate the lighting manually through artistic skill until things looked like they matched. I didn't like this answer, since it couldn't produce the 'correct' answer and the process didn't seem like something that I would be able to make happen in my own animated films.

As I worked on the [SIGGRAPH 97 paper](#) on high dynamic range (HDR) photography, I realized that the pixel values of an HDR image are actually proper measurements of the amount of red, green, and blue coming from every direction in the scene in the photograph.



Light probe image of Grace Cathedral, San Francisco. Image courtesy Paul Debevec.

And from learning about global illumination (GI) techniques for my qualifying exam at UC Berkeley, I understood that light comes from all directions and thus should be photographed panoramically, and I'd been experimenting with 360 photography for a few years at that point. So I captured omnidirectional HDR 'light probe' images in my kitchen at Berkeley, in the breezeway outside Soda Hall, and in the Eucalyptus Grove nearby, and eventually in Grace Cathedral across the Bay and the Galleria della Uffizi in Florence.

I considered writing my own IBL renderer, but Greg Ward showed me how to import the light probe images I was taking into his 'Radiance' GI renderer. I made the first renderings lit by HDR images for my [SIGGRAPH 98 IBL paper](#) using the kitchen, eucalyptus, and grace light probes, often using the nice candlestick model Radiance had from Greg's SIGGRAPH 93 cover image. The scene in the paper with a circle of candlesticks on a pedestal transformed into

the spheres on the stands for the animation **Rendering with Natural Light**, and the Uffizi Gallery light probe would eventually illuminate the middle scene of **SIGGRAPH 99's Fiat Lux**.

When you and your co-researchers produced the first light stage, what were you setting out to do? How did this progress and change with different light stage iterations?

Paul Debevec: The original goal for a light stage was to illuminate a real human face with image-based lighting, making them appear as if they were in some other location, so that you could realistically composite them into a background plate or a virtual set. My first thought was to put the person into a modestly-sized inward-pointing spherical screen and project a spherical image of the environmental lighting onto the screen all around them. But I couldn't find a way to affordably cover a full sphere with imagery and still get enough contrast ratio and brightness for the effect to be worthwhile.




This 1-minute exposure shows how Light Stage 1 moved a small spotlight around a person's head during the capture process. Image courtesy Paul Debevec.

So the Light Stage from our **SIGGRAPH 2000 paper** used an image-based relighting process instead: we built a gantry to rotate a single spotlight all around the subject, recording how they appear lit from every lighting direction, and then multiplied the color channels of each image by the intensity of that color of the light coming from the corresponding direction in the environment, and added all of these scaled image channels together. In effect, that relights the actor's face with the light of any light probe image you can imagine.

Light Stage 2, which was used to create digital actors for *Spider-Man 2* and *Superman Returns*, did the same thing faster with a 3 meter arc of strobe lights. But in 2001, as soon as Color Kinetics iColor MR lights became available, we built Light Stage 3 as a sphere of 156 inward-pointing RGB-controllable light sources, and this finally realized the vision of physically performing image-based lighting on an actor in the stage.



Actor Alfred Molina, who played Doctor Octopus in *Spider-Man 2*, sits in Light Stage 2.





Real and synthetic versions of Alfred Molina by Imageworks. Image courtesy Paul Debevec.

In 2004, when LED display panels became available a few years later, I designed an inward-pointing stage out of them to perform a test for Digital Domain for *The Curious Case of Benjamin Button*, and this technique is essentially what illuminated Sandra Bullock and George Clooney for the movie *Gravity*, where we worked closely with Framestore and Lightstage, LLC for the R&D phase of the project.

By far, the most commonly used light stage process we have developed is called polarized gradient illumination facial scanning, which we developed in 2005 as a way of deriving sub-millimeter 3D facial scans with high-resolution texture and displacement maps. When bright white LEDs became available, we added a second set of lights and new control circuitry to Light Stage 3, calling the new device Light Stage 5 (since we had proposed building large version of Light Stage 3 called Light Stage 4).

For our [SIGGRAPH 2005 paper](#), we programmed the lights to rapidly illuminate the actor from different directions as they were recorded at thousands of frames per second by a high-speed camera, allowing us to change the lighting on them as a post-production process. But I wondered if this programmable lighting system could be able to generate lighting patterns which would help us derive high-resolution 3D geometry of the face. I had two good ideas for this: one was to light the face with a succession of gradient lighting conditions based on the first four spherical harmonics, which allowed us to derive per-pixel surface normal maps of the face from any camera positioned around the sphere.




The second idea was to polarize the entire sphere of incident light in a special pattern that allowed us to tune the specular reflections in and out based on flipping a polarizer in front of the camera. Cross-polarized allowed us to capture a flat-lit, shadow-free, and shine-free texture map for the face. Subtracting this from the parallel-polarized condition allowed us to isolate the specular map for the face, which even more importantly gave us a way to calculate surface normals based on only specular reflections, yielding the sharp, sub-millimeter detail I wanted free from the effects of subsurface scattering. We first applied this technology for the movie *Avatar*, starting in 2006, scanning the principal cast for Weta Digital to help them craft the Na'vi versions of the characters as well as human digi-doubles, and to well over a hundred actors since.

Can you talk about some of the most memorable characters that have come to life with light stage captures, from your point of view? Any surprises in terms of where or how the light stage has been used?

Paul Debevec: Two of my favorite projects are [Digital Emily](#), from 2008, which used high-resolution light stage scans to achieve one of the first photoreal digitally rendered animated faces, and Digital Ira from 2013, which did the same for a face that was rendered interactively in real time. In both cases, the most rewarding part of the work was collaborating with world experts working on

complementary groundbreaking technologies: Image Metrics (now Faceware) provided the facial animation technology for Digital Emily, and Activision provided amazing real-time rendering techniques for Digital Ira.



And, in both cases, we did our absolute best with the scanning technology, but the final renderings came from our collaborators, and at some point the level of realism achieved completely surprised us and changed what we thought was possible. Activision has now purchased its own Light Stage X system from USC ICT and is using it to scan actors for many of its video games.

For movies, it was very exciting to work with Sony Imageworks adapting the Light Stage 2 system into movie production for *Spider-Man 2*, which ended up being a really good movie too, and working hard to bring our high-res scanning techniques into production use for *Avatar* (and we got film credits on both!). *The Curious Case of Benjamin Button* working with Digital Domain used a variety of light stage captures and both image-based lighting and relighting in achieving the first photoreal digital main character in a film, which was amazing to be a part of. [Vfx supervisor] John Dykstra surprised us by asking to use a particularly ambitious version of image-based relighting capture for *Hancock*, having us record Will Smith and Charlize Theron in Light Stage 5 from 100 viewpoints in 156 lighting conditions in an extensive set of expressions with an array of high-speed cameras.




Digital Domain utilized the Light Stage 5 to help build a CG model of an aged Brad Pitt for *The Curious Case of Benjamin Button*. Image courtesy Paul Debevec.

And probably the most emotional project we helped with was *Furious 7*, where we helped Weta Digital create a completely believable digital version of Paul Walker to allow the movie to be completed after his tragic accident. For this, we scanned Caleb and Cody Walker at the highest resolution possible, even recording their subtle facial movements as they read through lines originally said by their brother Paul. Because of the family resemblance, this data helped Weta Digital create Paul's digital double, which appears in about half of his character's scenes, most of which no one even guessed were rendered.

How has the digital human and light stage research moved forward with developments in rendering, and in real-time rendering in gaming and vr?
What do you see digital faces and humans actually being used for in vr?

Paul Debevec: That's a very interesting question. Our first light stage paper came out a year before Stanford's landmark paper on practical sub-surface scattering [a way of representing the way light scatters as it hits skin or a similar object], and one of the advantages it offered was that it could render a realistic human face without having to perform any light transport computation, subsurface or otherwise. Instead, the light transport of light falling on the face from every direction was captured through photography and simulated from the image data, allowing *Spider-Man 2* and *Superman Returns* before it was possible to have human-quality implementations of subsurface scattering algorithms.



Digital Ira was a real-time rendering project and collaboration between USC ICT and Activision. Image courtesy Paul Debevec.

This image-based relighting process from one-light-at-a-time data is still used in visual effects for lighting visualization and shader validation, but now that global illumination light transport can be done so efficiently with Arnold, Octane, V-Ray, and the like, it's the high-resolution light stage scan geometry which gets used for building and rendering the characters. Our Digital Ira collaboration with Activision was an amazing chance to work with some of the world's best real-time shader writers, including the masters screen-space subsurface scattering and environmental lighting techniques. This showed that you can render a full-screen entire realistic head at 60 fps on what is now a 5-year-old gaming laptop. Based on that, I have no doubt that the VR hardware we'll see in 2017 and 2018—even mobile systems—will be capable of rendering photoreal digital human characters.

You've also been researching light fields [light field cameras, for example, rely on both the intensity of light in a scene and also the direction that light rays are traveling in space to make imagery]. Where does that fit in with your digital humans and VR research?

Paul Debevec: I stumbled into light fields back in the Façade photogrammetry work, where I often had shot many images of the front of a building and needed to decide which one to use as the texture to map onto the 3D model. The simple answer was to choose the photo that was shot the most from straight on, but I realized that if the viewer is going to be seeing the building from off to the side a bit, then the best photo to use would be one taken from around the same angle. This way, unmodeled geometric detail like window sills sticking out would still look approximately right, and the different shines and glints the building makes would be reproduced somewhat as well. So the SIGGRAPH 96 Façade paper proposed 'view dependent texture mapping' as a way of blending between different images based on viewing angle, effectively rendering each model polygon as a surface light field.

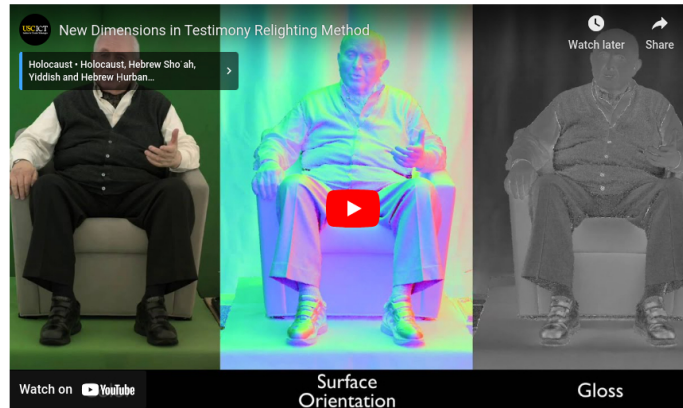


Above: Debevec recently consulted with cloud graphics company OTOY on the capturing of light fields, seen demonstrated in this video.

Our first light stage paper proposed the 'reflectance field' as a representation of

how each incident light field is transformed by a face or object into a radiant light field. Since a light field is four dimensional—a collection of 2D images shot from two dimensions of different angles—a reflectance field is actually eight dimensional, which is a massive amount of data when high-resolution imagery is involved. So in practice, the first light stage recorded 4D slices of the reflectance field from each camera: sequences of 2D images lit from a 2D set of lighting angles. It's been awesome to see other researchers build systems to record and simulate 6D and 8D reflectance fields since then.

We first combined light stage capture with light fields in earnest in 2006 for our first project using our full-body 8-meter Light Stage 6 system. We couldn't afford (or deal with the data from) a hundred high-speed cameras, so we used a vertical array of just three high speed cameras looking into the light stage and recorded the subject on a rotating treadmill as they made 36 walk or run cycles. As a result, we effectively could get 108-view 360-degrees around light fields for these repetitive motions, and by interleaving 33 lighting conditions at 990 frames per second, we could relight the person to any lighting environment as well as rendering them seen from any angle.




In 2012, our lab at USC ICT began working with USC's Shoah Foundation to use the light stage to record testimony of survivors of the World War II Holocaust. For this, we used Light Stage 6 to create flat symmetrical lighting from above and placed over 100 video cameras around the stage to record each survivor's stories from all the way around and some latitude of up and down as well. When shown on the 216 video projectors of our lab's 3D projector array display, the survivor appears life-sized, as a hologram-like, automultiscopic, glasses-free 3D image. By far the most interesting dimension recorded was their responses to over 1,000 questions asked to each survivor, allowing query recognition techniques from the ICT's Natural Language Group to play back appropriate responses to questions asked from anyone wishes to learn more about the survivors and their experiences. The lab is currently working on bringing this experience into virtual reality.

Finally, one of the skills you've mastered over the years is communicating complex concepts. Can you talk about what it takes to make a compelling SIGGRAPH presentation, and to deliver this research in papers and talks in digestible and interesting ways?

Paul Debevec: Thanks so much for that. I think giving a good talk or making a successful video is mostly a matter of storytelling. Everything you say and show needs to be done in consideration of how it will be perceived by the audience. How do you make them realize they care about the problem? Can you make them frustrated that it can't be solved in a straightforward way? Can you give them the vicarious enjoyment of figuring out a new way? Can you put results on the screen that make it clear that the technique really works, without even having to tell them that it really works? Can you mentally model what you're sure they already know, and only build from that? Can you say each sentence in as simple a way as possible, with as simple words as possible? Can you successively set up expectations and pay them off?





Paul Debevec (left) observes President Obama being scanned by a Light Stage at the White House for a Smithsonian 3D-portrait project. Image courtesy Paul Debevec.



The resulting scans. Image courtesy Paul Debevec.

I think the most important part of a talk is how you transition between your slides. I always make sure I know what my next slide will be, either through a written list of slide titles or using ‘presenter view,’ so that I can say a sentence or two to trigger some expectations and then switch to it so that the audience is in a good frame of mind to process all of that new visual information you’re throwing at them. It’s so much better than switching to a slide, giving the audience an initial bout of confusion as to what everything means, and then helping them recover from it. I also always go through my slides the night before a talk, or earlier in the day, to pre-load all the stories I’ll need to tell.

I almost never have equations on my slides, and once I had written an entire SIGGRAPH paper without an equation and ended up scrambling to find a way to include one just so the reviewers wouldn’t get worried. I think the best presentation is pictures with captions and the occasional video clip, and not even a complete sentence written down. But the most important part of speaking is to get a good night’s sleep! I’ve never given a great talk without a fully rested brain.

*Paul Debevec’s SIGGRAPH Asia **keynote** is entitled ‘Achieving Photoreal Virtual Humans In Movies, Games, And Virtual Reality’ and will take place on December 6th at The Venetian Macao.*



Read More:

#Avatar #Digital Domain #Framestore #Furious 7 #Google Daydream #Gravity #Industrial Light & Magic #Paul Debevec #photogrammetry #SIGGRAPH #SIGGRAPH Asia #Sony Pictures Imageworks #Spider-Man 2 #The Curious Case of Benjamin Button #The Matrix #Weta Digital



Ian Failes [More Articles](#)

Ian Failes is a writer covering visual effects and tech for *Cartoon Brew*.

Comments for this thread are now closed

1 Comment

[Login](#)

• Share

[Best](#) [Newest](#) [Oldest](#)



GW

7 years ago

Glad to hear about his work. What do you know of how current modeling methods work? I don’t work in animation but I’d like to stay up to date with the times. OSU has a website but it’s over a decade out of date. I’d like to know about the recent advances in modeling, particle physics, flocking, and so on but I can’t find a resource to explain it all.

0 0 Share

[Subscribe](#) [Privacy](#) [Do Not Sell My Data](#)

DISQUS

MORE IN CGI:

- [Pixar, Adobe, Apple, Autodesk, And Nvidia Ally On OpenUSD](#)
- [Christopher Nolan Forgot To Credit Over 80% Of VFX Crew On ‘Oppenheimer’](#)
- [SIGGRAPH 2023 Offers Feast Of Talks On New Animation, Including ‘Spider-Verse,’ ‘Nimona,’ And ‘Elemental’](#)
- [Epic Games Unveils Metahuman Animator, A Potential Game Changer For Automated Facial Animation And Lip Sync](#)

Sponsored Stories from Cartoon Brew



Tonje Skar Reiersen On The Epic Journey Behind 'Titina'



How To Enhance Your 2D Animation With Vector Art Available In Cartoon Animator 5



Praxinos Is Looking For Beta Testers For Its Unreal Engine Animation Software, Odyssey

Latest News from Cartoon Brew



Chinese Historical Epic 'Chang'An' Is The Fourth-Highest-Grossing Animated Film Of 2023 So Far



Commentary: AI Should Not Be Everything, Everywhere, All At Once




Bruce Lee Will Be Animated In The New Fantasy Series 'House Of Lee,' Coming In 2024



Know Your Indie Filmmaker: Stephen Irwin



Johnny Hardwick, 'King Of The Hill' Writer, Producer, And Dale Gribble Voice Actor, Dies At 64




Dreamworks Story Artist Darren Webb Talks Through His Boards From 'Ruby Gillman, Teenage Kraken'




Animator Spotlight: Irven Spence




The 75th Emmys Have Been Moved To January 15



An Early Look At This Year's Wide Open Animated Feature Oscar Race



Pixar's President Explains Why The Studio's Films Cost So Much - And It's Not For The Reason You Think



Familiar Faces Feature Among Animated Films Selected For TIFF's 2023 Short Cuts Lineup



'The Boy and the Heron Crosses' 5 Billion Yen Milestone In Japan, But Trails Other Miyazaki films



[About Cartoon Brew](#) • [Staff](#) • [Submit a Tip](#) • [Comment Guidelines](#) • [Contact us](#)

[Advertise on Cartoon Brew Today](#)

©2023 Cartoon Brew, LLC. All Rights Reserved. [Privacy Policy](#)

