

CONTACT US

July 04, 2017

Share

An interview with Wenzel Jakob

Foundry Trends recently caught up with Wenzel Jakob, creator of the Mitsuba renderer, for a deep dive into the powerful auto-retopology Instant Meshes algorithm he co-developed.

An assistant professor leading the Realistic Graphics Lab at EPFL in Lausanne, Switzerland, Wenzel's research revolves around rendering, appearance modeling, and geometry processing. Here's what he had to say.

You're perhaps best known for writing the Mitsuba renderer. What led you to jump into the field of auto-retopology?

Mitsuba was initially conceived as a playground for my own research on rendering when I was a graduate student.


During my years at Cornell, I noticed that my research on improving the rendering algorithms in Mitsuba drew me closer and closer to geometry – it turns out that there are many interesting connections between these fields.

After graduating, I saw an opportunity to enter a completely different field to work on pure geometry problems at ETH Zürich. It was during this time that the Instant Meshes algorithm was created in collaboration with Daniele Panozzo, Marco Tarini, and Olga Sorkine-Hornung.

Your Instant Meshes algorithm is much faster than other auto-retopology solutions, with very good quality. What are the trade-offs it makes?





As you mention, many other solutions exist—remeshing is classic problem in the area geometry processing, which has received considerable scrutiny over the years. The current state-of-the-art methods all rely on directional symmetry fields—the idea here is that to quadrangulate the mesh, one first needs to know how each quad should be oriented. These orientations can be visualized as little crosses that cover the surface (with flow lines shown on top):



To create a high-quality mesh, we'll want these crosses to be as similar as possible. The next step involves actually placing polygons, which becomes much easier once the orientations are already known.

The insight of Instant Meshes is that both of these problems can be expressed as a special kind of smoothing operation that allows for a very efficient implementation. We subdivide the mesh into smaller and smaller chunks and repeatedly blur data on these simplified meshes. The neat part about this approach is that all of this runs in parallel and scales to meshes with hundreds of millions of elements. Instant Meshes also has a unique way of dealing with curvature -- even without specifying any constraints, IM can automatically and robustly align to corners and other sharp features.



Unlike many researchers in the field, you not only released your source code, but even released a fully functional auto-retopology application via GitHub. Why did you choose to go this route?

Daniele, Marco, Olga, and I are firm believers in reproducible research, which means providing access to the full source code of a research project so that others can not only develop extensions but also use the implementation to reproduce all results and statistics shown in a paper.

In the case of Instant Meshes, we also felt that we hit a sweet spot in terms of performance and usability where our software could be immediately useful to practitioners outside of academia, for instance in 3D scanning, retopo, or CAD workflows.

The response has been phenomenal! Instant meshes has been downloaded over 60.000 times, and our YouTube video also has nearly that many views -- these are simply incredible numbers to us.

Instant Meshes seems like an ideal choice for auto-retopology on large meshes acquired using 3D scans or photogrammetry. What are your thoughts on this type of application?

As part of our work, we often need to scan objects using laser or structured light scans. While Instant Meshes is definitely an important part of the toolbox, there are still a many pain points with the general scanning pipeline that we've found frustrating over the years. First of all, it involves many separate steps: acquiring a set of scans, aligning them, performing mesh fusion/reconstruction, UV parameterization, remeshing, and so on. That's not even all of them, and often a different tool is used for each step. The worst part is that most steps must be repeated if it later turns out that another scan is needed to add further detail from another viewpoint.



Motivated by these frustrations, we've been working on an new real-time WYSIWYG pipeline for 3D scanning named **Field-Aligned Online Surface Reconstruction** that was just accepted to SIGGRAPH 2017. Its key principle is that the user should always be able to see the final-quality remeshed object complete with texture and displacement maps. The system indicates where further scans are needed and interactively updates the view as they are added. All pipeline stages are fused together, which is interactive thanks to a new streaming version of the Instant Meshes algorithm.

Tell us about your new research group - what are you working on now?

In addition to the work on mesh generation for surfaces and volumetric meshes, the Realistic Graphics Lab focuses on appearance modeling and light transport algorithms. This involves studying the optics of natural and man-made materials that occur in our world and developing efficient mathematical models that describe them. The goal of our work is to considerably accelerate the underlying algorithms, and to improve the accuracy and fidelity of the resulting renderings. We use Modo to model scenes that are then rendered using our research rendering platform.

Subscribe to our newsletter

First Name *

Email Address *

Submit



The Foundry Visionmongers Limited is registered in England and Wales.

