2021

Photogrammetry and Texture Baking: Creating Lightweight 3D Objects for Any VR Headset. In E. Kroski (Ed.) 32 Virtual, Augmented, and Mixed Reality Programs for Libraries

Sarah E. Huber

Andrew Sumner

CHAPTER # 11
PHOTOGRAMMETRY AND TEXTURE BAKING
Creating Lightweight 3D Objects for Any VR Headset

Sarah Huber
Engineering Technology Information Specialist & Assistant Professor
Andrew Sumner
Visual Solutions Engineer

Purdue University (Indiana)

Patrons will learn how to make 3D objects that are viewable across a range of consumer to high-end VR headsets through the processes of photogrammetry and texture baking. Photogrammetry is an affordable and relatively simple process that recreates objects in 3D with great detail. The process includes taking photos of an object with a Smartphone and then bringing them into open-source software that aligns them into what is called a point cloud. A point cloud is the first processing step in creating a 3D object. Point clouds can be complex and data intensive, so many consumer VR headsets do not have the capacity to render the objects for optimal viewing. Texture baking makes 3D objects 'lightweight' in geometry for rendering, but the details are maintained using images for material properties. Lightweight 3D objects can be viewed across a range of VR headsets with stunning results. This activity includes three sections: (1) Photo Capture, 2) Photogrammetry Image Stitching, and (3) Texture Baking. Patrons will create 3D objects they can use for their video games, websites, artist portfolios and more!
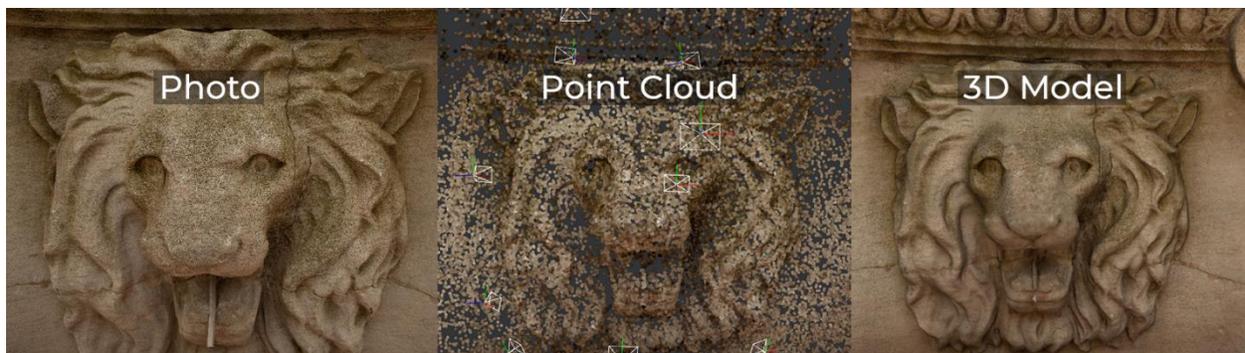
Figure #1: Photo capture, point cloud created from photogrammetry and texture baked 3D model

OVERVIEW

This guide walks patrons through the process of creating 3D objects from photos to be shown in VR/AR/Web. The instructions can be followed by individual patrons, but it is recommended patrons work in small groups one time through. There are enough details to the process that patrons collaborating through each section will be helpful. We recommend two to three separate workshops, 2-3 hours each, to cover the topics: (1) Photo Capture, (2) Photogrammetry Image Stitching, and (3) Texture Baking. Due to the level of detail, patrons would benefit from an instructor working with each group.

Any digital camera that takes quality photos can be used for the photo taking portion. For this activity, however, we focus on the use of Smartphones because of affordability and access. Photogrammetry software 'reconstructs' 3-dimensional models by comparing features between photos to align them into a 3D point cloud. Because the entire workflow relies on the quality of the photos taken, below are general rules to optimize image capture.

For processing images in the photogrammetry and texture baking software, included are directions with troubleshooting suggestions. We encourage patrons to play with the troubleshooting. Due to the state of the software, one set of prescribed directions across different scenarios limits successful 3D object creation. For this activity, we took photos of a lion sculpture on our university campus. **Tutorial images referenced below can be found in a Dropbox folder at https://bit.ly/3gejFsy**. If patrons want to use photos already taken, the sculpture photos can be found in the same Dropbox folder.

Necessary Equipment and Materials
- Smartphone with the ability to take photos
- Desktop computer for processing (Nvidia graphics card is required and is commonly found)
- USB cable
- SD card reader
- Free, open-source software: Meshroom for photogrammetry and xNormal for texture baking

Recommended but Optional Materials
- If patrons do not have access to computers with dedicated graphics cards for Meshroom, MicMac is a less powerful, alternative software
- Sketchfab for storing and viewing the published 3D object

STEP-BY-STEP INSTRUCTIONS
Photo Capture
- It is important for every image to have the same camera settings. For Android, select 'Pro' mode (manual mode) to be able to set consistent ISO and shutter speeds between photos. For iPhone, consider purchasing the app ProCamera for the same purposes.
- Take approximately 75 photos total. It is possible to make models from small amounts of photos but the quality typically is low. For this project, try to choose an object you can walk around.
- Walk around the object 180 degrees, taking one photo at a time, trying to cover as much of the object in photos as you can reach. For example, take photos straight-on with the object and then try to take photos from above and below to cover the areas you cannot see straight-on. Do not worry about covering areas that are similar to other photos you have taken – it is most important to get complete coverage around the point of interest.
- Coverage and consistency matter most in photogrammetry. Overlap between each photo should be (30-50%).
- Variables to consider:
  - Lighting, reflections from shiny surfaces, change in camera settings will all cause inconsistencies in the data and then the software will struggle to process it as intended.
  - Objects with surfaces that are *not* reflective and well-lit will be the easiest to scan.
- Once photos are taken, transfer them to a computer. This is easiest with a USB cable for phones or an SD card reader for cameras.

Photogrammetry Image Stitching in Meshroom
This section includes the most detail and has room for troubleshooting, so it is important to give patrons time and help where needed.
- Install Meshroom
- Orient yourself to the interface (See Image 1: Default Meshroom):
  - Images tab on the left
  - Image viewer in the middle-left
  - 3D Viewer in the middle-right
  - Graph Editor in the bottom-left
  - Node in the bottom-right
- Drag your photos into the Images column on the left. (See Image 2: Added Images).
- Save the project using File > Save As.
- Press the green **Start** at the top-center of the screen. Processing the images can take some time depending on the number of photos and size. Expect longer processing times (>10

minutes) for image counts greater than 50. In the Graph Editor at the bottom of the interface, are the Nodes. When a Node is complete, it turns green.

- Once the **StructureFromMotion** Node turns green, press **Stop** at the top of the screen. Then, you can click on that Node to see the result of your images being turned into a point cloud. (See Image 3: StructureFromMotion Node.)
- Click on the point cloud in the 3D Viewer. Rotate and pan around the object to assess the quality. If the point cloud looks good, move on to the process of **Decimation**. If the point cloud is a poor representation of your object, move to **Troubleshooting**.
- **Troubleshooting**: Stop the process (button at the top next to Start). Here are some Setting options to troubleshoot:
  - On Feature Extraction, you can change Describer Preset from Normal to High. This adds more detail to the process looking for features in images to use for matching.
  - On FeatureMatching, enable Guided Matching. This adds a second stage to matching with new constraints based off the first stage for more accuracy.
  - On StructureFromMotion, increase the minInputTrackLength to 3 or 4. This will only use the most accurate matches instead of attempting to include outliers (with some speed-up).
  - On DepthMap, reducing the Nb Neighbour Cameras (both SGM and Refine) can speed up processing by comparing less cameras for coverage of the same area. This one is heavily dependent on your coverage -- more spread out images can benefit from this.
  - On Texturing, reduce the Texture Size to 4096 and disable Use Udim. This will use less overall resolution for the color information and take much less time to process.
- When the results are complete and the patron is satisfied with the object results, move on to Decimation. Remind patrons to save the project after processing.
- **Decimation**: Once the point cloud looks good, the polygon count will be very high. For example, with the lion sculpture, the polygon count is over 3 million. To make a lightweight 3D object, we need to do the following steps to make the polygon count smaller.
  - Patrons will use the MeshDecimate node; this will 'decimate' or reduce the geometry down to a specified amount. Patrons should Decimate *before* Texturing in order to apply the texture to this new lightweight mesh.
  - Click and drag the Texturing node further to the right in the Graph Editor and click and drag the MeshFiltering node down. This will give the patron room in the Graph Editor to easily see the node connections and provide space to add a node between MeshFIltering and Texturing.
  - Disconnect the MeshFiltering output wire connected to Texturing by right-clicking on the wire and selecting Remove.

- o   Right-click in any open space in the Graph Editor to search for the MeshDecimate node. Select to MeshDecimate to add it as a Node.
- o   Connect the MeshFiltering to MeshDecimate with the output going to inputMesh on Texturing. Click and drag from the circles next to the output and input to connect nodes together. The result will look like a white line has connected them.
- o   On MeshDecimate, set the Max Vertices to 50,000. This limits the total amount of vertices (and then faces). (See Image 4: MeshDecimate Node.)

- **Texturing**: A final step in reducing the polygon count within Meshroom is to Texturize.
  - o   Change the Texture Side to 2048. This lowers the output resolution of the texture.
  - o   Change the Unwrap Method to LSCM and ensure the settings match the recommendation in section 2. LSCM will force the output texture to be a single image instead of splitting up between many.
  - o   Deselect the Use UDIM check mark. UDIMs are an advanced workflow for using multiple texture maps for a single mesh which is not ideal for Virtual Reality. (See Image 5: Texturizing.
- Once these nodes are set up, click the green **Start** again to process the mesh and decimation. Meshroom will compute any nodes that are not green to finish the process.

Texture Baking for Lightweight 3D Object in xNormal
Texture Baking: the process of 'baking' information from the higher resolution version of a mesh to a lower polygon version. Think of it as a high-to-low process where the details from the high quality are what will be saved into images to 'fake' that information on top of a lower resolution mesh.
- Download xNormal
- Patrons can use xNormal to bake a 'Normal Map' which will fake depth in a texture to any rendering engine (Sketchfab, Unity, Unreal, etc).
- For this process, patrons will need the high-quality mesh and the lower-quality mesh with texture. Right-click the MeshFiltering node and select Open Folder to find the high-quality mesh. Then right click the Texturing node to find the lower resolution mesh and the image texture.
- In xNormal, select the High Definition Meshes tab on the top-right.
- Drag and drop the high-definition model from before here (the Mesh Filtering result folder).
- Select Low Definition Meshes and drag the lower polygon mesh from the Texturing output folder from Meshroom here.
- Select Baking Options and change the following settings (see Image 6: xNormalOptions):
  - o   Set the output file to a specific folder (click the 3-dot icon next to the text bar at the top), name the image, and change the format to jpg.
  - o   Change the Size to 4096 by 4096 (lower if high resolution details aren't required).

- o Check that Normal Map is selected.
- o Press Generate Maps in the bottom-right. The process will 'render' the information to a blueish image.

Results

With this process, patrons have captured a physical object, computed a 3D reconstruction from it, and baked information from a high detail into a lightweight model capable of being rendered in any VR/AR/Web application. The quickest way for patrons to use this model with the textures created for visualization would be through Sketchfab, an online repository for browsing and uploading 3D models. Patrons should upload the mesh file and the 2 textures created -- the color map and the normal map. With this platform, most VR headsets can view the models either through a web browser or Sketchfab's mobile application.

- To export the finished creation, find the 3D model made in Meshroom by right clicking the Texturing node and choosing 'Open Folder'. Copy the .obj and the texture image to a location where you want to store your 3D model. Then, copy the resulting texture from the xNormal bake to the same location. These are now your finished model with a color texture and a normal texture.

Requirement
- NVIDIA CUDA-enabled GPU (compute capability >= 2.0)
  Dedicated graphics card (Nvidia), these are very common in desktop PCs and gaming-oriented laptops

Preparation
- Install Meshroom and xNormal

RECOMMENDED NEXT PROJECTS
- Turntable photogrammetry: Object are put on a lazy susan to take photos, rotating the object instead of taking photos walking around the object.
- Material scanning with photogrammetry: baking a surface scan to a material you could use on other digital objects.