

Kintsugi 3D: An Empirically-Based Photogrammetry Production Pipeline

Lou Brown and Charles Walbridge; Minneapolis Institute of Art, Minneapolis, MN USA
Michael Tetzlaff; University of Wisconsin – Stout; Menomonie, WI USA

Abstract

Advancements in accurate digitization of 3D objects through photogrammetry are ongoing in the cultural heritage space, for the purposes of digital archival and worldwide access. This paper outlines and documents several user-driven enhancements to the photogrammetry pipeline to improve the fidelity of digitizations. In particular, we introduce a new platform for capturing empirically-based specularities of 3D models called Kintsugi 3D, and visually compare traditional photogrammetry results with this new technique. Kintsugi 3D is a free and open-source package that features, among other things, the ability to generate a set of textures for a 3D model, including normal and specular maps, based empirically on ground-truth observations from a flash-on-camera image set. It is hoped that the ongoing development of Kintsugi 3D will improve public access for institutions with an interest in sharing high-fidelity photogrammetry.

Motivation

Cultural heritage imaging efforts are turning to photogrammetry of 3D objects in order to acquire enhanced visual documentation for research, conservation and public access. However, many currently available methods are lacking in their ability to reproduce realistic color, texture detail, or material accuracy. For instance, although *Agisoft Metashape*¹ can utilize image masks to reduce error or bounce-lighting from typical uniformly illuminated image-sets, it will often produce visual errors such as blurriness, level of detail instabilities, or texture map crowding as a result of its texture mosaic technique.

We present a software-based solution, featuring a new free and open source platform called *Kintsugi 3D*². Developed by a team at the University of Wisconsin – Stout with support from the Minneapolis Institute of Art (Mia) and Cultural Heritage Imaging (CHI) under a grant from the National Endowment for the Humanities (NEH), *Kintsugi 3D* is capable of synthesizing empirically-based roughness, specularity, normal, and diffuse textures (Figure 1) from an image set containing photographs of an object captured with flash-on-camera illumination. *Kintsugi 3D* is an evolution of its predecessor, *IBRelight* [1], with a user experience redesigned to specifically target the application of building textures and materials for use in lightweight 3D viewer applications. What sets *Kintsugi 3D* apart from other workflows that produce specular maps is that all the textures produced are



Figure 1. A cropped visual juxtaposition of several generated texture maps of "Saint Lucy," by Benedetto di Bindo Zoppo, from the Minneapolis Institute of Art (68.89). These textures were created in Kintsugi 3D Builder from 375 photographs.

empirically-based: they are derived directly from photographic data using classical optimization methods, and the reconstruction error from this optimization process can be recorded as metadata for the object as documentation of the fidelity of the digitized form.

Kintsugi 3D also features its own *Viewer* application for public access to finished digitizations in the highest possible quality, using a custom shader designed specifically for materials derived from photographs. The goal of this viewer is to support the rest of the *Kintsugi 3D* platform with a lightweight app for public access to this robust reproduction quality, while striving for feature parity with comparable viewers such as Sketchfab or Smithsonian

¹ <https://www.agisoft.com>

² <https://github.com/michaelt919/Kintsugi3DBuilder>
<https://github.com/UWStout/Kintsugi3DViewer>

Voyager. *Kintsugi 3D* also supports exporting in standard texture formats to support existing efforts using established viewers like *Sketchfab* or *Voyager*. The simplicity and open access of the *Kintsugi 3D* platform makes this available even to institutions without the infrastructure or support to otherwise develop such hands-on physical or digital interactive experiences.

Kintsugi 3D still relies on *Agisoft Metashape* (or potentially other photogrammetry alternatives in the future) for camera alignment and 3D reconstruction; it merely replaces the final stage of texture generation. As such, it is an extension, not a replacement, for established photogrammetry solutions. However, *Kintsugi 3D* does change the photogrammetry pipeline in certain significant ways. Professional photographers at many institutions currently capture image sets that utilize white backgrounds with uniform lighting on the object in each image set. While this makes it possible to easily mask images from the contrast between the background and object, it also makes achieving the necessary uniform illumination for accurate textures more challenging, and empirically deriving specularities from such images is not possible. There is also a risk of color issues in the textures due to bounce lighting or interreflections that, among other things, dull out colors, reducing the texture fidelity.

In contrast, *Kintsugi 3D*, like its predecessor *IBRelight*, uses a photographic technique that leverages a flash mounted on the camera [1]. There are two primary modes of capture: against a black background in a studio environment, or in-gallery. These two options reduce unintentional bounce light and offer accessibility and flexibility in terms of how the photos are taken, while providing essential reflectivity data for *Kintsugi 3D* to reconstruct specular maps, which most other photogrammetry workflows cannot replicate. Since *Kintsugi 3D* is a new tool being released in 2024, the primary goal of this paper is to provide preliminary guidance on a pipeline for effectively using *Kintsugi 3D*, and compare the results in different viewing applications. We also discuss additional model refinement steps related to texture mapping that can enhance photogrammetry results in general.

Guidelines for flash-on-camera photography

Kintsugi 3D expects that the object is photographed under a very specific kind of illumination – an on-axis / camera-mounted light source. A common example of this is the built-in flash that is available on many digital SLR cameras. Alternatively, at Mia, a Picolite³ was attached to the camera to achieve this configuration. It is best for this light source to be as small as possible so that it is very close to the underlying assumption, built into *Kintsugi 3D*, that the light source is infinitely small. It is also important that the light provided by the flash is significantly brighter than any other ambient illumination in the photography environment – ideally, there would be basically no ambient light (a pitch-dark room when the flash is off). Good photographs for *Kintsugi 3D* will exhibit strong highlights, which is important for *Kintsugi 3D* to be able to capture specularities.

It is also important to understand that on-axis flash photography for *Kintsugi 3D* requires a lower exposure level than is typical for traditional photography. In other circumstances, a photographer would properly expose the non-specular parts of an object, with some “clipping” of the specular highlights being expected. The goal, however, when taking photographs for *Kintsugi 3D*, is to accurately measure the specular highlights, which means that any clipping of

the specular highlights must be minimal. While the photographs may appear “underexposed” to the eye of the experienced photographer, this underexposure is necessary for *Kintsugi 3D* to capture the required specular information. On the other hand, the underexposed parts of the images will usually correspond to diffuse reflectance that can still be determined with an acceptable amount of precision by combining all of the available photographs.

Finally, it is generally considered best practice to not use automatic camera settings such as aperture, shutter speed, and ISO when capturing image sets for photogrammetry. Avoiding such automatic camera settings (including flash brightness in addition to the settings listed above) is even more important when using *Kintsugi 3D*. These settings all affect the effective exposure of the photographs and, when set to “automatic mode,” the exposure may change between photographs. *Kintsugi 3D* makes an implicit assumption that all images have the same exposure, and if this is not the case, there may be unexpected variation or miscalibration of albedo intensity in the textures produced by *Kintsugi 3D*.

Kintsugi 3D Pipeline

Once photographs with flash-on-camera illumination have been taken, *Kintsugi 3D*'s pipeline begins with some off-the-shelf photogrammetry tool, with the primary requirement being that it must generate and export camera calibrations (typically an XML file) after the photographs are successfully aligned and the 3D model generated. While *Agisoft Metashape* is currently the only option supported by *Kintsugi 3D*, there are plans for further development to support alternative photogrammetry tools.

The basic pipeline is as follows:

On-Flash Photography → *Metashape* → *Kintsugi 3D Builder*

A more complex version of this pipeline includes UV-remapping, and remeshing:

On-Flash Photography → *Metashape* →
Instant Meshes and/or Choice of 3D Editing Tool →
Decimate Model → *UV Map Model* → *Metashape* →
Kintsugi 3D Builder

We follow a *Metashape* processing workflow recommended by Cultural Heritage Imaging (CHI), similar to the one described by Over et al. [2] in order to best utilize the latest version of the software.

It is advantageous to generate an ambient occlusion (AO) map within *Metashape* (or the user's photogrammetry / 3D modeling software of choice), as *Kintsugi 3D* cannot currently generate its own. This AO texture stores a surface's precalculated exposure to ambient light, allowing for more accurate ambient lighting in 3D viewers. It is also helpful to generate a diffuse map in *Metashape* to check for obvious misalignments in camera locations or otherwise unusable images within an imageset – however, this diffuse map will not be used by *Kintsugi 3D*.

Remeshing and UV Mapping

Remeshing and UV mapping are optional steps of the workflow, which are not directly connected to *Kintsugi 3D*, but also enhance the texture generation step of the photogrammetry workflow. A UV map is the “flattened” representation of a 3D object, generally cut up and arranged as a 2D image. UV maps are intrinsic to almost all 3D file formats (FBX, OBJ, glTF, PLY, etc.). Remeshing is the

³ <https://broncolor.swiss/products/picolite>

process of regenerating the 3D polygons of a model as cleaner mesh for UV mapping. This step can also reduce detail as a side effect, so it is best to start with a high-poly model. The remeshing software of choice we employ is *Instant Meshes*, a free and open source remeshing tool that is accessible to beginners. It should be noted that other heritage-oriented tools, such as *Smithsonian Cook*⁴, also use the same software for remeshing. For our models, we used incremental target face counts of 250,000, 500,000, and 1,000,000 polygons for remeshing. The mesh can also be decimated further using *Metashape* or another 3D modeling program.

After remeshing, the UV map can either be created by hand or generated in *Metashape*. If using *Metashape*, the UV map will be automatically generated when building the texture for the object. However, the automatically generated UV maps from *Metashape* will often be messy and difficult to read. To improve this, the UV map can also be edited by hand in programs such as Maya, 3dsMax or Blender. (Note that when using external 3D modeling tools, it is essential to never move, rotate, or scale the object at all when moving between *Metashape* and other programs, otherwise *Metashape* and *Kintsugi 3D* cannot continue to use the same camera locations / photo alignments for texture baking.)

The advantage of letting *Metashape* generate the UV maps is that it prioritizes minimizing visual distortion from the original photographs. However, it will almost always create jagged UV seams. On the other hand, hand-made UV maps may contain unfinished seams, holes, and other errors due to manual editing, but tend to be much easier to visualize and edit later in programs such as *Adobe Photoshop*, and have cleaner seams with little to no jagged edges.

It is common practice to generate lower-polygon versions of models (“decimation”) that can be more easily loaded on the web and on mobile devices. It is important to note that UV mapping must be reperformed after decimation, since the process destroys a model’s UV map, in order to accommodate the alterations to the 3D polygons. For the same reason, if using remeshing, decimation, and/or UV mapping in an external program, the ambient occlusion map generation (as described earlier) must be performed *after* these mesh-altering steps.

Kintsugi 3D Builder

With the photogrammetry processing stage completed, the next step is to create a project in *Kintsugi 3D Builder*. The following materials are required to create a project:

- A set of tone-mapped and color-corrected images taken using flash-on-camera, previously aligned and processed using photogrammetry software.
- A file containing camera calibrations from the photogrammetry software’s photo alignment step. (Currently, only *Agisoft Metashape* is supported, but we hope to eventually support exports from alternative tools.)
- A 3D model of the object, aligned with the camera calibrations).

Next, *Kintsugi 3D* will preprocess and load the image set. After loading, the object should appear in the middle of the screen. If it is rotated unusually, the object’s orientation can be adjusted as needed.

The first step after a model is loaded is to perform color calibration using a photograph containing an X-Rite ColorChecker

Classic⁵ with six squares in the black to white gradient, as described in prior work [1]. This photo can be from outside of the primary image set for the project, as long as the photo is from the same photography session: i.e., the same lighting conditions, exposure settings, and on-flash camera setup, and located at approximately the same location relative to the camera as the subject being scanned. (To account for light attenuation, *Kintsugi 3D* assumes that the ColorChecker is the same distance from the camera as the nearest point on the object.) The user will need to manually identify the six grayscale squares using an “eyedropper” tool integrated into *Kintsugi 3D*. Once this has been done, *Kintsugi 3D* is able to use the photograph of the Color Checker chart to convert the pixel values in all the photographs (which are subject to exposure and tonemapping) into absolute reflectance measurements.

The next step is to perform lighting calibration. In this stage, the user interacts with the model to tell *Kintsugi 3D* where the light source is physically located relative to the camera lens axis. To accomplish this, the user leverages shadows in the source photographs, cast by the object upon itself from the mounted flash, as indicators of the light’s position. As the user manually adjusts the light position, *Kintsugi 3D* will show them a projection of the photograph onto the 3D model, using a camera placed at the same position as the currently calibrated light position. When all the shadows are occluded by the object itself from this perspective, the virtual camera is at the same location as the flash. This calibration is performed using a single source image at a time, but may be confirmed using additional images in the image set.

The final step is to configure and run a specular fit to generate the specular model and texture maps, using algorithms outlined in prior work [3, 4]. The user may specify the texture size, with texture resolutions as high as at least 8192 x 8192 supported. For objects



Figure 2. Saint Lucy, by Benedetto di Bindo Zoppo, from the Minneapolis Institute of Art (68.89), processed from 375 photographs. Top left: rendering in Sketchfab using textures from Kintsugi 3D. Top right: rendering in Sketchfab using texture from Metashape. Bottom left: source photograph taken with flash-on-camera. Bottom right: catalog photograph for reference.

⁴ <https://smithsonian.github.io/dpo-cook/>

⁵ <http://xritephoto.com/colorchecker-classic>

made out of materials like jade or which are otherwise translucent, a translucency option is available, which introduces an additional non-Lambertian constant term in the fit as an approximation of translucency.

Results

The results of this workflow exhibit several key advantages over textures generated by a conventional photogrammetry process. The most obvious is *Kintsugi 3D*'s ability to generate empirically-based specular textures, both a specular color map and a roughness map. Figure 2 shows an example of a model processed by *Kintsugi 3D*, compared with an image of the same model rendered with *Metashape*'s diffuse texture and no specularity, and two ground-truth images: one flash-on-camera photo from the input image set, and a reference image of the object from Mia's catalog. Although these textures can be created by hand or by digitally manipulating a color texture map (which we understand to be the primary method used to add specularity to most heritage objects currently on *Sketchfab*), the difference here is that *Kintsugi 3D*'s textures are derived from the original photographs using classical and well-understood optimization methods. The residual error of fitting the photographs to these specular textures can even be archived with the object as metadata, along with a description of the method used to fit the textures and a shader program for correctly rendering the object from the textures.

An additional advantage we have found is that *Kintsugi 3D* can often produce even diffuse textures with better color accuracy than *Sketchfab*. Conventional photogrammetry produces textures in a manner that is agnostic to diffuse and specular reflection and tends to blend some specularity into the diffuse map, resulting in lighter and more desaturated color. In contrast, *Kintsugi 3D* explicitly models the diffuse and specular parts of the object's reflectance, allowing it to reconstruct not only the specularity, but also the diffuse texture more accurately. Figure 3 is a good illustration of this, where the *Metashape* texture is significantly




Figure 3. "The Age of Bronze," by Auguste Rodin, from the Minneapolis Institute of Art (54.1), processed from 399 photographs. Top left: rendering in Sketchfab using textures from *Kintsugi 3D*. Top right: rendering in Sketchfab using texture from *Metashape*. Bottom left: source photograph taken with flash-on-camera. Bottom right: catalog photograph for reference.



Figure 4. "Enshrined Buddha" from the Minneapolis Institute of Art (89.55), processed from 222 photographs. Top left: rendering in the Sketchfab web viewer using specular textures generated from *Kintsugi 3D Builder*. Top right: rendering in *Kintsugi 3D Viewer*. Bottom left: rendering in the Sketchfab 3D web viewer using a color texture generated by Agisoft *Metashape*. Bottom right: a catalog photo of the object for reference.

brighter and more desaturated than the object's true diffuse albedo, while the textures produced by *Kintsugi 3D* do not have this issue. In addition, *Kintsugi 3D*'s use of the *ColorChecker* chart ensures that the color values in the texture map are normalized to match the object's real-world reflectance. In this sense, the *ColorChecker* serves in a role for calibrating material appearance similar to what the use of scale bars in a photogrammetry workflow provides for calibrating scale.

Integration with 3D Viewers

The results of the *Kintsugi 3D* workflow can be rendered in several viewing applications. The textures and materials generated by *Kintsugi 3D Builder* will look best within *Kintsugi 3D Viewer*, a new viewer implementation in the open-source Godot game engine based on earlier prototypes [5]. However, in contrast with prior work [1, 5], *Kintsugi 3D Builder* is also able to support other existing 3D viewers, such as *Sketchfab*⁶ and *Voyager Story*⁷, although these viewers may require extra conversion steps, and are not able to utilize all the specularity textures generated by *Kintsugi 3D Builder*. Figure 4 compares *Kintsugi 3D Viewer* with *Sketchfab* (still using textures produced by *Kintsugi 3D Builder*), along with a *Sketchfab* render using the *Metashape* texture, and a ground truth

⁶ <https://sketchfab.com>

⁷ <https://3d.si.edu/voyager-story-standalone>

image from the museum catalog. The differences between *Kintsugi 3D Viewer* and *Sketchfab* (when using textures from Kintsugi 3D Builder) are subtle. Beyond any concerns of the quality of the rendering, commercial platforms like Sketchfab also have the disadvantage of being subject to licensing terms, in contrast with the free and open source license for the *Kintsugi 3D* platform. In this sense, *Kintsugi 3D Viewer* is preferable, allowing an institution to retain control over its own depository of objects and future-proof against potential changes to the terms and conditions on commercial platforms.

Kintsugi 3D is still in active development and does lack some rendering features offered by other viewing tools, such as post-processing. Future plans for *Kintsugi 3D Viewer* include support for some of these effects and other rendering options.

Conclusion

Kintsugi 3D is a novel platform that complements existing photogrammetry tools and 3D viewers to offer better support for empirically-based textures and materials. It offers benefits for cultural heritage institutions wanting to generate models with improved color accuracy, and advanced material representation such as empirically derived specular maps. At the same time, its accessibility reduces the entry barrier for smaller institutions by producing high quality results with the simpler flash-on-camera photography technique. While technical results have been demonstrated in a research context for some time [1, 3, 4, 5], we hope that the public and open-source release of *Kintsugi 3D*, with a user experience streamlined from earlier efforts to focus on the use case of sharing models online, makes this empirically-based, flash-on-camera workflow for capturing material appearance available and accessible to a broader community of heritage professionals.

References

- [1] Michael Tetzlaff, Gary Meyer, and Alex Kautz, "IBRelight: An Image-Based 3D Renderer for Cultural Heritage," *IS&T Archiving Conference*, 2018, pp. 93-98.
- [2] Jin-Si R. Over, Andrew C. Ritchie, Christine J. Kranenburg, Jenna A. Brown, Daniel Buscombe, Tom Noble, Christopher R. Sherwood, Jonathan A. Warrick, and Phillippe A. Wernette, *Processing coastal imagery with Agisoft Metashape Professional Edition, version 1.6—Structure from motion workflow documentation*, U.S. Geological Survey Open-File Report 2021–1039, 2021, 46 p., <https://doi.org/10.3133/ofr20211039>
- [3] Giljoo Nam, Joo Ho Lee, Diego Gutierrez, and Min H. Kim, "Practical SVBRDF Acquisition of 3D Objects with Unstructured Flash Photography," *ACM Transactions on Graphics*, vol. 37, no. 6, 2018, pp. 267:1-267:12.
- [4] Michael Tetzlaff, "High-Fidelity Specular SVBRDF Acquisition from Flash Photographs," *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, vol. 30, no. 4, 2024, pp. 1885-1896.
- [5] Tyler Garcia, Zhangchi Lyu, and Michael Tetzlaff, "An Online Model Viewer for Cultural Heritage in Unity 3D," *IS&T Archiving Conference*, 2022, pp. 50-55.

Acknowledgements

This work was supported by a Research and Development Grant from the National Endowment for the Humanities [PR-290101-23].

The authors wish to recognize the UW-Stout development team for *Kintsugi 3D* – student developers Isabel Smith, Tyler Betanski, Jacob Buelow, Luke Denney, and Victor Mondragon, and UX consultant Darcy Hannen – as well as Dan Dennehy and Pierre Ware from Mia, and Carla Schroer and Mark Mudge from CHI for their participation in this project.

Author Biographies

Lou Brown is a BFA graduate from the University of Wisconsin–Stout. Lou's role was sponsored through a National Endowment of the Humanities (NEH) grant. Lou learned photogrammetry pipelines, conducted quality comparisons, wrote pipeline documentation, and liaised between the Minneapolis Institute of Art team and Michael Tetzlaff's team at UW-Stout.

Michael Tetzlaff is an Assistant Professor of Computer Science at the University of Wisconsin – Stout. He received his B.S. from Bethel University in St Paul, and his M.S. and Ph.D. from the University of Minnesota. His research interests include image-based computer graphics techniques and especially their applications in cultural heritage.

Charles Walbridge is Lead Collections Photographer at Minneapolis Institute of Art (Mia). With more than 20 years experience documenting varied collections of fine art objects under exacting standards, Charles has been at the leading edge of digital capture, color management and studio workflow processes for cultural heritage. Since 2013, he has played a key role in developing 3D modeling techniques for Mia's Visual Resources studio, forming cross-discipline collaborations with developers, animators, and maker spaces.