MAJOR QUALIFYING PROJECT

Submitted to the faculty of
WORCESTER POLYTECHNIC INSTITUTE
in partial fulfillment of the requirements for the
degree of Bachelor of Science

Digital sculpture restoration

Submitted by:

Nathan S. Kaplan

Submitted to:

Professor Cosme Furlong

NEST – NanoEngineering, Science, and Technology
CHSLT- Center for Holographic Studies and Laser micro-mechaTronics
Mechanical Engineering Department
School of Engineering
Worcester Polytechnic Institute
Worcester, MA 01609-2280

May 6th, 2021

Objectives

This project sought to help the Worcester Art Museum restore Edward Augustus Brackett's 1851 sculpture, *Shipwrecked Mother and Child*, through digital restoration techniques. To achieve this goal, we had to recreate five missing pieces which had broken away from the statue in the 170 years since its inception. These parts not only needed to be faithful to Brackett's original vision and aesthetic but also be durable enough to last the next three decades.

Abstract

The process of sculpture restoration, the repairing of environmental damage in a work of art, has long been a careful and meticulous art. Capturing the aesthetics and techniques used by an artist is challenging in and of itself, but translating them to fill the voids left by missing pieces is a skill possessed only by a handful of individuals. To aid conservators in the restoration process and lower the skill floor required to carry out restoration work, we collaborated with the Worcester Art Museum to develop a modern set of techniques for sculpture restoration. Our case study for the project was *Shipwrecked Mother and Child,* an 1851 sculpture by New England sculptor Edward Augustus Brackett. After 80 years in the WAM's storage, the sculpture was missing five pieces, which we set out to recreate with our workflow. Our process begins by using 3D scanners to digitize the work of art. The scans are then imported into a 3D modeling program to manually recreate the missing portions of the statue based on the scan data. Finally, these parts are 3D printed and fitted onto the statue to replace what was once lost. Due to time constraints caused by the COVID-19 pandemic, we could not create a finished version of each of the five components. However, each component saw at least three printed iterations that integrated nearly seamlessly with the statue, proving that our workflow could recreate missing components in an artist's style and produce components that cleanly attach to the break sites on the statue. Throughout the project, we experimented with more sophisticated printing technologies such as PolyJet and ceramic SLA and used augmented reality for rapid iteration without the need to 3D print at every stage. While the project did not see a definitive conclusion, it has provided a solid foundation for future work into digital sculpture restoration. This future work could involve more sophisticated printing methods, more advanced augmented reality, and artificial intelligence for the automated generation of parts.

Acknowledgments

This project team was composed of a single member: Nathan Kaplan. First, he would like to thank Professor Cosme Furlong-Vazquez, for his guidance, patience, and wisdom throughout the project. Alongside the professor, he would like to thank the Center for Holographic Studies and Laser micro-mechaTronics (CHSLT) at WPI for acquiring equipment and software used by project, as well as providing the use of its facilities. The acquisitions made by the Professor went on to support all of CHSLT. Next, he would like to thank the Worcester Art Museum (WAM) and its objects conservator, Paula Artal-Isbrand, for providing him with a unique project and opportunity which allowed him to contribute to a vital restoration project; while, at the same time, also allowing him to experiment with bleeding-edge modern technologies. Furthermore, he would like to thank her for providing some of the images seen throughout the report. Finally, he would like to thank Colin Hiscox, whose foundational work on the project and continued support throughout allowed the scope and depth of the work to reach new levels.

Table of Contents

List of Figures

1. Introduction

Completed in 1851 by Edward Augustus Brackett (1818-1908), *Shipwrecked Mother and Child* was the self-taught sculptor's magnum opus [1]. The statue started life in 1848 as a series of clay and marble mock-ups, but to fulfill his vision, Brackett spent his life savings on a block of Vermont marble from which he sculpted the anatomically correct life-size figures. Following its completion, Brackett displayed the sculpture in New York before moving it to Boston in 1854, where it was on loan to the city's Atheneum until the early 1900s before returning to Brackett. In 1904, he gave the statue to its current home, the Worcester Art Museum (WAM), where it was displayed for almost four decades before moving into permanent storage. In 2019, after nearly 80 years in the museum's archives, the statue was moved to the Jeppson Idea Lab, where it is currently undergoing conservation and restoration before settling into a permanent home on the museum floor. In August of 2020, when work began on this project, the statue was still undergoing cleaning due to delays caused by the COVID-19 pandemic. This step entails cleaning years of pollution, dirt, and other grime off the marble surface to get it ready for restoration work. WPI began working in parallel on restoration efforts to combat these delays, which involve repairing any damage taken on by the statue. With *Shipwrecked Mother and Child,* we focused on filling in five primary areas of loss: the woman's right big toe, left pinky toe, right index finger, the tip of her left pinky finger, and a missing chunk of the pedestal. Although more damage was present across the sculpture, the WAM identified these five areas as they represented the most prominent missing pieces of the statue. The following chapters will begin by detailing the background information on the history of sculpture restoration and previous digital restoration efforts. The background will also discuss the technical information needed for understanding the technologies used throughout the project, such as 3D scanning, 3D modeling, and 3D printing. From there, the paper will thoroughly detail the methodology we used to recreate the lost areas of the statue in such a fashion that someone with no knowledge of the subject matter could reproduce the work done. Next, the paper will show the project's results and discuss what could be improved and further developed in future work on *Shipwrecked Mother and Child*. Finally, the paper will conclude with a reflection on the project and thoughts on areas of future work in digital restoration as a whole.

2. Background

This chapter introduces the background information necessary to understand why the Worcester Art Museum was interested in new restoration methods and the technology behind said methods. The chapter explores a brief and non-comprehensive history of sculpture renovation, from classical approaches to the pioneering digital restoration innovations developed to restore Tullio Lombardo's *Adam* in the mid-2000s. The second Subsection will detail the technical background behind the primary technologies used in this project, 3D scanning, 3D modeling, augmented reality, and 3D printing.

2.1. Sculpture restoration

This Subsection provides a brief and incomplete history of sculpture restoration to help the reader understand why this type of research would benefit the art conservation world. It summarizes the timelines into three eras: Classical (pre 21st century), Modern (the standard techniques used today), and Digital (bleeding edge restoration techniques).

2.1.1. Classical (pre 21st century) restoration methods

In the classic approach to stone sculpture restoration, widely used before the turn of the 21st century, a conservator aimed to erase or disguise as much damage as possible through various aggressive, non-reversible means [2]. Abrasives, acids, and chisels all attacked the statue's surface, intending to remove any foreign deposits or stains. Unfortunately, this would come at the cost of damaging the sculpture's original finish and detail. Some classic conservators showed such little care for a piece's history that, in some cases, they recut the work to 'improve' the aesthetic of the piece. With such a mindset to conservation, the statues became the product of the restorer rather than that of the original artist. With that, these pieces lost their history and the intention of their original sculptor. This approach to conservation began to change at the end of the 20th and beginning of the 21st centuries when conservators started respecting the art for what it was.

2.1.2. Modern restoration methods

Unlike their classical counterparts, today's conservators focus on preserving the original artist's vision through delicate cleaning of the surface and potential repair of significant damage [2]. Museums begin by researching the provenance and history of a piece to understand its material composition and what the surface grime contains (Artal-Isbrand, Personal Communications, October-December 2020). Based on this research, a conservator will formulate a set of solvents and begin cleaning small pieces of the statue to find the solvent which can remove grime without damaging the statue's surface. Then, the conservator will begin the meticulous process of cleaning the entire work. For *Shipwrecked Mother and Child,* the cleaning process is taking Artal-Isbrand over a year to complete. With older works of art, a conservator might even choose to keep surface deposits on the statue, either to enhance its age or perhaps because removing the sediment deposit would damage the original work [2]. Finally, in modern conservation, reversibility is critical. Often, conservators will not repair significant damage if the fix involves invasive treatments (such as drilling into the stone or cutting the stone away). When repairing damage, the conservator does so to ease potential future restoration work (*ibid*. Artal-Isbrand, 2020). With *Shipwrecked Mother and Child*, Artal-Isbrand intends to use reversible conservationist adhesives to glue the restored pieces back onto the statue. These adhesives will safely separate from the statue without leaving a mark, allowing pieces to be replaced when necessary.

Unfortunately, these careful methods do not solve all the issues that come with damage repair and loss restoration. For example, on *Shipwrecked Mother and Child*, half of the figure's right index finger is missing (from the tip of the finger to the first knuckle). To restore the finger, Artal-Isbrand began by making a mold off the left index finger. Because only the backside of this finger exists, she had to sculpt the remaining half while being true to Brackett's aesthetic. With a completed left index finger, she began work on mirroring it into the right index finger which would replace the missing appendage. While this method certainly works and has been used many times for previous restoration efforts, the reliance on hand sculpting is prone to human error. As such, museums are working to pioneer new conservation techniques to recapture what might have been.

2.1.3. Digital restoration & The conservation of Tullio Lombardo's *Adam*

On the evening of October 6th, 2002, Tullio Lombardo's *Adam,* an essential work from the Italian renaissance, fell to the floor of the Metropolitan Museum of Art in New York City and shattered into pieces. Over the next 12 years, the statue would undergo a complete restoration to return to its original appearance [3]. *Adam* broke into 28 large pieces and a few hundred smaller fragments, with the arms, legs, and tree trunk suffering the brunt of the impact and, therefore, much damage. Because of the nature of the break, many of the pieces had thin and fragile edges; thus, the conservators needed to develop a way to work around them and reassemble *Adam* without further damaging the sculpture. In early 2003, the Met began using 3D scanners to digitize the major fragments so that the statue could be reassembled entirely in software. The scans allowed conservators to use Finite Element Analysis (FEA) techniques to determine how and where to pin the figure and what kinds of adhesives could rebind *Adam*. Finally, the software model allowed the team to build a sort of 'exoskeleton' that precisely held the statue together during the reassembly process. This digital conservation laid much of the groundwork for this project, as we used many of the same techniques to approach restoring *Shipwrecked Mother and Child.*

2.2. Technical background

This Subsection provides the necessary technical background to understand the concepts discussed in the paper. The bulk of the work done by the project was surrounding 3D scanning, 3D modeling, and 3D printing and their applications for art restoration. Nevertheless, the project also explored how the emerging field of augmented reality could be applied to conservation.

2.2.1. 3D scanning

3D scanning is the process of digitizing real-world objects through automated non-contact methods. While many different scanning processes exist, the one we used in this project is a technology known as structured-light 3D scanning, which works by projecting QR code-like patterns known as fringes and using cameras at known offsets to the object to measure the pattern's deformation [4]. This deformation is digitized as a point cloud, with each point representing a point on the surface of the scanned object. The scanner used during this project, a Creaform Academia 50, typically generates between two and four of these point clouds per second [5]. However, each point cloud is wholly independent, with no common origin or defined references. As such, to create a coherent 3D model, the point clouds need to get transformed into a single unified coordinate system through a process known as registration, which looks for common elements between point clouds to stitch them together [6]. Registration is a continual process that works as the scanner scans; this way, if the scanner "knows" what it last scanned, it only needs to compare the newly generated point cloud to a small subset of the overall data, improving scan efficiency. We used 3D scanning for this project as it allowed us to create a precise 3D model of the statue necessary for recreating the missing components.

2.2.2. 3D modeling

3D modeling is the process of creating three-dimensional objects in specialized software. Often known as CAD, or Computer-Aided Design, in the Mechanical Engineering world, we chose to use the 3D modeling terminology throughout this project as it fits more in line with the artistic nature of the work. Compared to CAD, 3D modeling is a less rigid process where shapes are defined by wireframes, collections of edges and vertices that define the surface of an object, rather than being defined by a rigid system of dimensions. All 3D modeling in this project was done with a free and open-source piece of software known as Blender, which (among many features) allows users to manipulate wireframes and organically sculpt shapes. These features were crucial to the project as they presented a low-skill floor for the rapid iteration of the generated pieces of the sculpture.

2.2.3. Augmented reality

Augmented reality is an emerging technology that works to enhance or augment a user's vision with holograms. On the following page are two images of the Microsoft Hololens; the first, Figure 2.1, is simply an image of the device worn on the head with a glass screen going in front of the users' eyes. While the HoloLens does not produce 3D images, the content displayed to the wearer appears in the world as if it were there: imagine virtual paintings hung on real walls. The second image, Figure 2.2, is an example of what the wearer of the HoloLens might see. The motorcycle in that picture does not exist in real life; instead, it is being projected in front of the designer's eyes so that she may see what the bike would look like in person. The HoloLens knows where to project objects through four cameras mounted on the headband, which track high-contrast points around the room to determine where it might be in reference to everything else [7]. From there, the HoloLens sizes and moves the projected objects to appear as if they were in a specific part of the room. Compared to simply viewing the parts in modeling software before printing, augmented reality gives users a tangible idea of how large parts are in real life and how they may look when integrated with the statue.

Figure 2.1: Product render of the Microsoft HoloLens 2

Figure 2.2: Promotional material for the Microsoft HoloLens 2

2.2.4. 3D printing

3D printing, or additive manufacturing, creates objects through the computer-controlled combination of material layer by layer [8]. This Subsection will explain how the two printing technologies used during this project, Fused Deposition Modeling (FDM) and PolyJet, work. Much of the printing in this project was done on the Zortax M200, an FDM printer, representing the most common 3D printing processes as of 2020 [9]. The technology works by feeding a continuous thermoplastic filament, typically PLA or ABS, into a moving heated printer head. There, the thermoplastic is heated and extruded out of the printer head with a gear. Next, two stepper motors will move the printer head on a plane to deposit a horizontal slice of an object, known as a layer. When the FDM printer finishes creating the layer, the head will move up to print the next one, and the process will repeat until the object has finished printing. Compared to newer printing technologies, FDM does not offer the same levels of quality nor speed, but it makes up for this in its significantly reduced cost and ease of use. Since this project was based out of the CHSLT lab at WPI, which had the Zortrax M200 FDM printer on hand, many initial prototypes were printed on it for convenience. Furthermore, all the fast high-quality printers at WPI had a long print queue, which slowed iteration compared to using the Zortrax.

This project also used a newer printing technology known as PolyJet, which works on the same fundamental principles as a standard 2D inkjet printer [10]. Instead of depositing ink on paper, a PolyJet printer works by depositing a UV curable photopolymer one layer at a time onto the object being printed. Once each layer finishes being printed, it is cured with a flash of UV light. In a significant improvement over FDM printers, high-end PolyJet printers can mix different resins to create unique colors, allowing for full-color 3D printing. Since the scanners used during the project captured color data, this would theoretically allow us to print parts for the statue that perfectly match the finish of the marble from which it is carved.

3. Approach

Before statue restoration could begin, it was essential to identify the most important missing parts on *Shipwrecked Mother and Child* and create restoration plans for each. While the statue had many minor nicks, five areas of damage were identified by the Worcester Art Museum to be the focus of the restoration; these are: the right big toe, the left pinky toe, the left pinky finger, the right index finger, and a chunk of the statue's pedestal.

At the beginning of the project, we determined two factors that would constitute a completely restored part—first, the recreated appendage needed to seamlessly fit into the rest of the statue. Not only would the gap between part and statue need to be invisible, but everything we created needed to match Brackett's aesthetic and sculpting style. Secondly, the recreated parts must be made out of a material with similar visual properties to the original marble. Not only did this entail matching color and reflectivity, but since 3D printed resins often yellow with age and *Shipwrecked* needs on the museum floor for the next 30 years without further restoration work, we needed to find materials that did not age.

Over the following Subsections, the paper will detail the five areas and independently describe how each should be filled in to meet the standards of the Worcester Art Museum. Then in the methodology, the paper will describe how we went about achieving both of these goals for every piece.

3.1. The right big toe

Figure 3.1: The broken surface of the right big toe.

As shown in Figure 3.1, the right big toe is wholly missing from the statue. Therefore, an entirely new toe needed to be created to restore this section, then glued over the break site to seamlessly complete the foot. Fortunately for the project, the left big toe is entirely intact, as visible in Figure 3.2. However, solely using scans of the existing left big toe would be insufficient as its interior geometry was occluded by the other toes, making scanning impossible. Luckily, Worcester Art Museum's Objects Conservator, Paula Artal-Isbrand, made a plaster cast of the left toe, visible in Figure 3.3, for use in previous digital restoration efforts by Colin Hiscox when he was using Fringe Projection scanning to digitize the sculpture before the start of this project. Therefore, the plan was to combine as much scan data as possible from the original left toe with the scan of the plaster version to create a single piece that could fit onto the missing portion of the statue.

One unique consideration when working with the big toe was the remains of a metal pin on the break site. At the turn of the 20th century, when the toe likely separated from the statue, “pinning” was a common practice for reattaching lost appendage to statues. It is unknown when this original restoration attempt took place or who carried it out, but at some point in the last century, the pin failed, and the toe went missing. As will be discussed in the methodology, the pin helped us constrain the requirements of the scanner, as it needed to register the millimeter thick feature accurately.

Figure 3.2: The intact left toe

Figure 3.3: A plaster cast of the left toe made by Artal-Isbrand for use in the project.

3.2. The left pinky toe

Figure 3.4: The missing left pinky toe.

As shown in Figure 3.4, the woman's left pinky toe is missing from the statue. Much like with the right big toe, the objective here was to recreate the missing portion of the statue. Our first thought was to scan the right foot and use its little toe as a reference for the recreation. Unfortunately, this idea was quickly extinguished. Due to the right foot's position, Brackett never sculpted a right pinky toe, so there is no available reference on the woman from which to recreate the loss. Luckily, the child's left foot is visible, providing us with an anatomical reference of a toe that we thought we could scale up to recreate the missing appendage. An image of the child's left foot is visible in Figure 3.5, while an image of a plaster cast of his left pinky toe is in Figure 3.6. Again, much like the right big toe, the plan for the left pinky toe was to combine all the scan data we could get from the statue with that from the plaster cast to recreate the missing portion. If this proved insufficient, the neighboring toes would provide references for small details such as the nail. Once more, these tiny features helped establish the minimum resolution of our scanner.

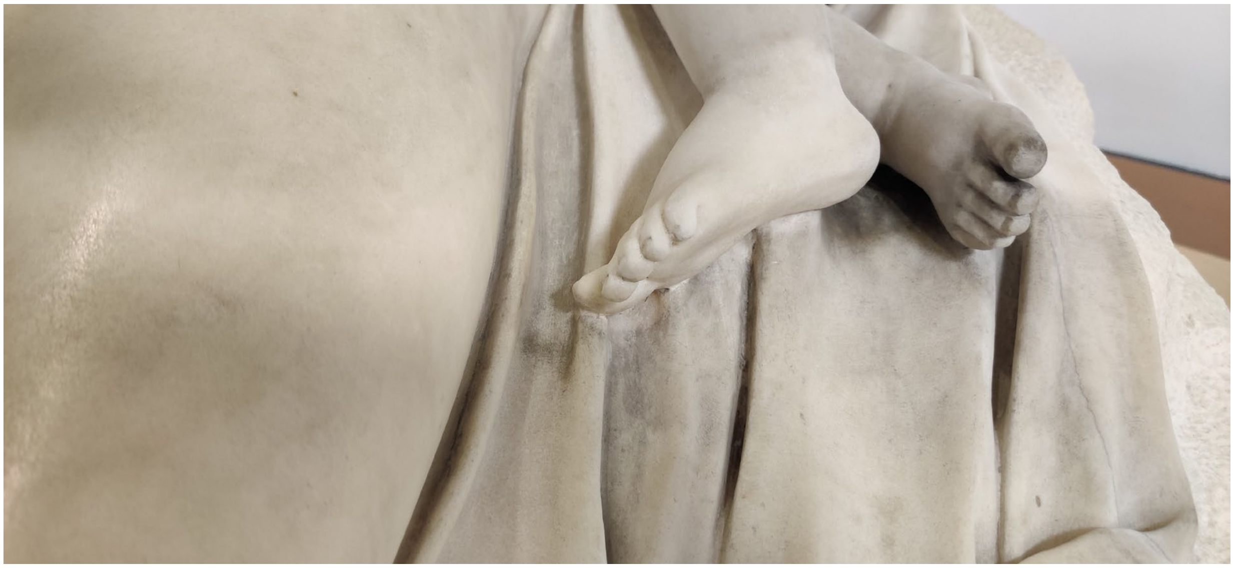

Figure 3.5: The child’s left foot.

Figure 3.6: A plaster cast of the child’s left pinky toe, with a ruler for scale.

3.3. The tip of the left pinky finger

Figure 3.7: The woman's left hand, with the loss on the tip of the left pinky finger visible.

As shown in Figure 3.7, the tip of the woman's left pinky finger is missing from the statue. To restore the finger, it would need to be "filled-in" and appear no different to its neighbors. Because the other fingers on her left hand were the benchmark, we chose to use the scan data of the neighboring fingers to restore the area of loss.

Due to the simplicity of restoring the fingertip compared to the toes or the right index finger (no need to combine scan data from plaster casts with that from the statue, heavy constraints which make its positioning more objective than subjective, and an abundance of reference data), we identified the pinky finger as an "easy to fix" areas of loss that should be targeted first when developing the restoration workflow. The idea being issues with the process would be much easier to fix when parts are quicker to generate and easy to fit and finish.

3.4. The right index finger

Figure 3.8: The woman's right hand, with the loss between the tip and knuckle of the right index finger visible.

As shown in Figure 3.8, half of the woman's right index finger is missing. To restore the finger, we would need to create a piece that not only fit perfectly onto the break site but looked correct in the context of the rest of the hand. Not only did this mean the index finger needed to be the correct length (shorter than the middle finger, longer than the ring), but it needed to curl naturally into a relaxed position.

As with the toes, Artal-Isbrand made a plaster cast of the finger before this project began, which we would combine with scan data from the statue to recreate the piece. Again, we knew that occlusion would be a constant consideration when scanning and that we would need methods to create what could not be seen by the scanner. So for the index finger, the plan was to use the scan of that plaster cast visible in Figure 3.9, the scan of the neighboring middle finger, AND the scan of the backside of the left index finger as seen in Figure 3.10 to create a single, composite right index finger.

Figure 3.9: The “Finger on a stick,” a plaster cast of the left index finger whose underside was sculpted by hand

Figure 3.10: The woman’s left hand. The top half of her left index finger was used to restore the missing right index finger.

3.5. The pedestal

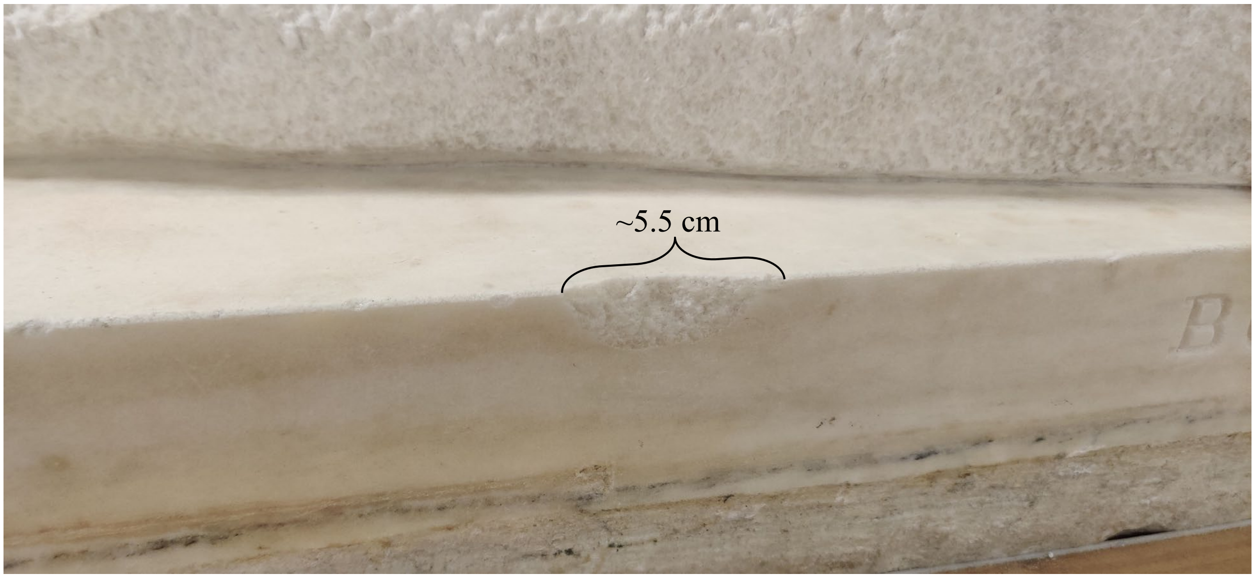

Figure 3.11: The missing piece from the base of the statue, located on the backside below the woman's right foot.

As shown in Figure 3.11, there is a small chunk missing out of the statue's pedestal. Compared to human appendages, the restored pedestal piece was by far the most straightforward to define. Once restored, the hole should be entirely invisible, with both the top and side planes smoothly continuing. This definition also meant a very trivial restoration plan, which did not require reference geometry from the statue. Since we can simplify the pedestal into a rectangle with a slight fillet (a rounded edge), the plan to restore this piece of the statue was to create the analogous rectangular prism in 3D, then "subtract" away the existing material to be left with a small part which could fill the damage. Much like with the tip of the pinky finger, we decided that due to the simplicity of restoring the pedestal compared to other portions of the statue, we would focus on it first during the restoration process. Furthermore, due to this piece's objective nature, we decided to use it as a control when testing new printing technologies and determining whether or not they looked good on *Shipwrecked*.

4. Methodology

In this project, we aimed to generate and 3D print five small missing components of *Shipwrecked Mother and Child* by Edward Augusts Brackett to aid in the statue's restoration. In doing so, we developed a workflow that involved digitizing the statue through 3D scanning, part generation in Blender (a 3D modeling program), and part realization through 3D printing. The following Sections detail and explain the process taken to create these pieces. In order, they outline:

1. How the statue was scanned
2. How the scans were post-processed and cleaned of noise
3. How the missing pieces were generated from the scan data using Blender
4. How the project used augmented reality for rapid iteration
5. How the parts were 3D printed and fitted on the statue

The remainder of this subjection will provide a summarized version of the following chapter, highlighting the essential ideas and numbers from each of the five Subsections. The specific details surrounding workflows, processes, data types, and installation processes will be covered in those Sections, so we urge the reader to dive into them if more detail is required.

Before the digitization of *Shipwrecked Mother and Child* could begin, we needed to acquire a 3D scanner capable of accurately detecting the most delicate features of the break sites. These small valleys and ridges provided the constraint for the minimum measurement resolution of the scanner. Measuring both the statue and scans previously made by Colin Hiscox, we determined that the scanner must have a resolution of at least 500-micron (.5mm) to detect the most minor features present on the break surfaces. Below in Figure 4.1 is a cast Artal-Isbrand made of the right big toe break surface, with a ruler providing scale. Note that this cast represents the negative of the surface. As can be seen, the small divet in the lower left is approximately one millimeter wide, but to scan it, we need something with a half mm resolution. Otherwise, we would see two peaks without a valley between, making a flat plane. Of course, the higher the resolution, the better and the more precise our scans would be, but this value provided a hard minimum constraint.

Figure 4.1: A plaster cast of break site from the left big toe

The second constraint in selecting the scanner was cost. The project had a budget of $10k, which severely constrained eligible scanners. After a significant amount of research and some demos, we narrowed the choice down to two scanners: the Creaform Academia 10 and the Einscan Pro 2x Plus. While the later scanner had the higher measurement resolution (200-micron vs. 500-micron) and a lower cost ($6,899 vs. $7,990), we opted to go for the Academia 10 scanner due to its more powerful scanning software and WPI's established relationship with

Creaform. However, after discussing these findings with the project advisor, Professor Furlong-Vazquez, he pushed to purchase the Creaform Academia 50 at twice the cost of 10. Compared to its counterpart, the 50 had a 250-micron measurement with a resolution and featured the ability to scan texture (color data). After discussions with both WPI's mechanical engineering department and the Worcester Art Museum, we ended up buying the Academia 50 as it would provide the best tools for the project and serve as a great educational tool in future classrooms. Note that the 250-micron resolution of the Academia 50 was the "lowest" resolution point of the entire workflow. Blender's resolution is "infinite," with users being able to add vertices up to the limit of their computer's processing power, while the worst quality 3D printer used, the Zortrax M200, had a resolution of 90 microns. As such, we can say that the produced parts were accurate within 250 microns of the original statue. Please reference Appendix B.1 for the full technical specifications of the Creaform Academia 50.

With the scanner purchased and having spent a week practicing on various objects in the CHSLT lab, we could begin the process of scanning the statue, which is detailed in Subsection 4.1. The process, which took approximately 15 hours over five days, involved scanning the statue twice. The first scan was at a low mesh resolution of 2mm, allowing the entire statue to be digitized for future promotional use (this data was never used for part generation but was a fantastic tool for creating stunning renders of *Shipwrecked Mother and Child*). The second set of scans was at a high mesh resolution of 0.5mm with a measurement resolution of 0.25mm (the highest quality setting available on the Academia 50). With these scans, we only digitized the specific areas of interest for the restoration, as there was no need to have an ultra-high-resolution version of the face when the toe was what was broken. It is also from these scans that part generation would take place.

With scanning complete, it came time to recreate the areas of loss in 3d modeling software. Relatively simply, we had to take this step because we needed to recreate what was once lost. The scan data only provides information on what is currently on the statue: nothing of value can be printed based on what exists; instead, we need what does not exist. As detailed in Subsection 4.3.1., we decided on Blender, an open-source 3D modeling program, as the program of choice for this step in the workflow. The concept behind recreating a lost appendage is simple. Taking the example of the right big toe: the statue had an intact left toe, which we could digitally copy, mirror, then place onto the right toe. From there, we used boolean subtraction tools to remove the existing portion of the statue from the newly created left toe, leaving us with a part that could be 3D printed and fit perfectly onto the statue.

Finally came part realization through 3D printing. Unfortunately, due to significant time constraints at the end of the project, we were forced to do most of our printing on FDM printers instead of the preferred high-resolution technologies like PolyJet or SLA. Despite this, we could get lovely prints out of an FDM printer, as detailed in Section 4.5. Again, since these parts were printed at 90 microns compared to the 250-micron resolution of the scanner, they did not end up negatively affecting the resolution of the produced parts.

It is important to note that this workflow is circular; while we completed scanning before starting on modeling, poor fit issues on all iterations of big toe meant we needed to go back and rescan the statue. Likewise, we constantly went between Blender and the printer: modeling a piece, fitting it onto the statue, then making necessary adjustments per Artal-Isbrand's recommendations. Each iteration shed light on how we could improve the parts and the digital restoration workflow itself.

4.1. Scanning the Statue

To scan *Shipwrecked Mother and* Child, we used a Creaform Academia 50 3D scanner alongside VXElements, a proprietary software suite which converts the scan data into a 3D model. We used a Dell XPS 15 9560 with 64 gb of RAM as the scanning computer during the process. To find a PC suitable for scanning, please reference Appendix A.1 to find the minimum recommended specs for VXElements. Both the Dell XPS 15 and the Creaform Academia 50 are visible below in Figure 4.2, which shows Nathan Kaplan in the midst of scanning *Shipwrecked Mother and Child*.

Figure 4.2: Nathan Kaplan in the process of scanning Shipwrecked Mother and Child. Image taken by Paula Artal-Isbrand [11]

Before scanning can begin, we must first configure and calibrate the Academia 50, the process for which is detailed in the steps below:

1. Plug the scanner in. The Academia 50 has two plugs, one USB plug which goes into any USB 3.0 port on the scanning computer and a power port which needs to plug into a wall socket.
2. Open VXElements, Creaform's proprietary scanning software; this project used version 8.1.1 of the software suite, though further updates have been released since. If necessary, go through the configuration steps by registering the scanner's serial number with the software.
3. Within VXElements, open VXScan, the scanning portion of the program. If the scanner is connected and powered on, a green bar in the top right of the window should say 'connected.' If it says 'disconnected,' use the help guide to troubleshoot the issue. This project discovered an error with VXElements that prevented an AMD machine from detecting a plugged-in scanner; make sure to use Intel devices when using the Academia 50.
4. Under the positioning parameters, drop-down in the top right set the positioning method to "Targets / Geometry / Texture." Next, check both the "automatic shutter" and "capture texture" boxes in the scanner parameters drop-down. The last parameter to configure is the resolution under the scan parameters drop-down to the left of the screen; set it to 2mm for low-resolution scanning or 0.5mm for high-resolution scanning. Note: in the 3D scanning world, "texture" refers to the color information of the scanned object.

5. With VXScan configured, remove the calibration plate from the scanner's carrying case and place it on a flat surface near the object we want to scan. Next, press control+shift+D to begin the calibration process and follow the instructions displayed on the screen to calibrate the Academia 50.

With configuration and calibration complete, scanning with the Academia 50 is a straightforward process. Begin by pointing the scanner at the target object, press the trigger, then slowly move the scanner over the piece's surface to scan it. As we do so, the scanner will project a QR code on the statue's surface, representing what it is currently scanning. On top of the device are two status lights. If the top light is illuminated in red, the scanner is too close to the object and needs to be pulled back. If the bottom light is illuminated, the scanner is too far from the object and needs to be brought closer. If both lights are illuminated, the scanner has lost tracking. When the scanner loses tracking, immediately stop moving the scanner and retrace its motion until both status lights turn off. If this process does not work, refer to VXScan to see where the scanner lost tracking, represented by a blue square that shows what was last seen, and resume from there.

Due to the statue's size, we could not scan the entire statue at a high (0.5mm mesh) resolution without overloading VXScan. Thus, we first began by scanning the entire statue at a lower mesh resolution of 2mm, a process that took approximately 10 hours. The completed scan is visible in Figure 4.3, Figure 4.4, Figure 4.5. While the low-resolution model was not used to generate any parts, it was invaluable to cross-reference when looking at the smaller, higher detail scans. Furthermore, it was used to create high-quality renderings, such as the one in Figure 4.6.

Once the low-resolution scan of the statue was complete, we made four additional high-resolution scans of the areas of interest for the restoration: the left hand, the break on the pedestal, the right hand, and the feet, visible in Figure 4.7, Figure 4.8, Figure 4.9, and Figure 4.10 respectively. These high-resolution scans were then exported into Blender to create the missing components of the statue. Compared to the low-resolution scans, they contained the necessary detail about the break sites of the missing parts and the reference geometry we planned on creating the missing components out of.

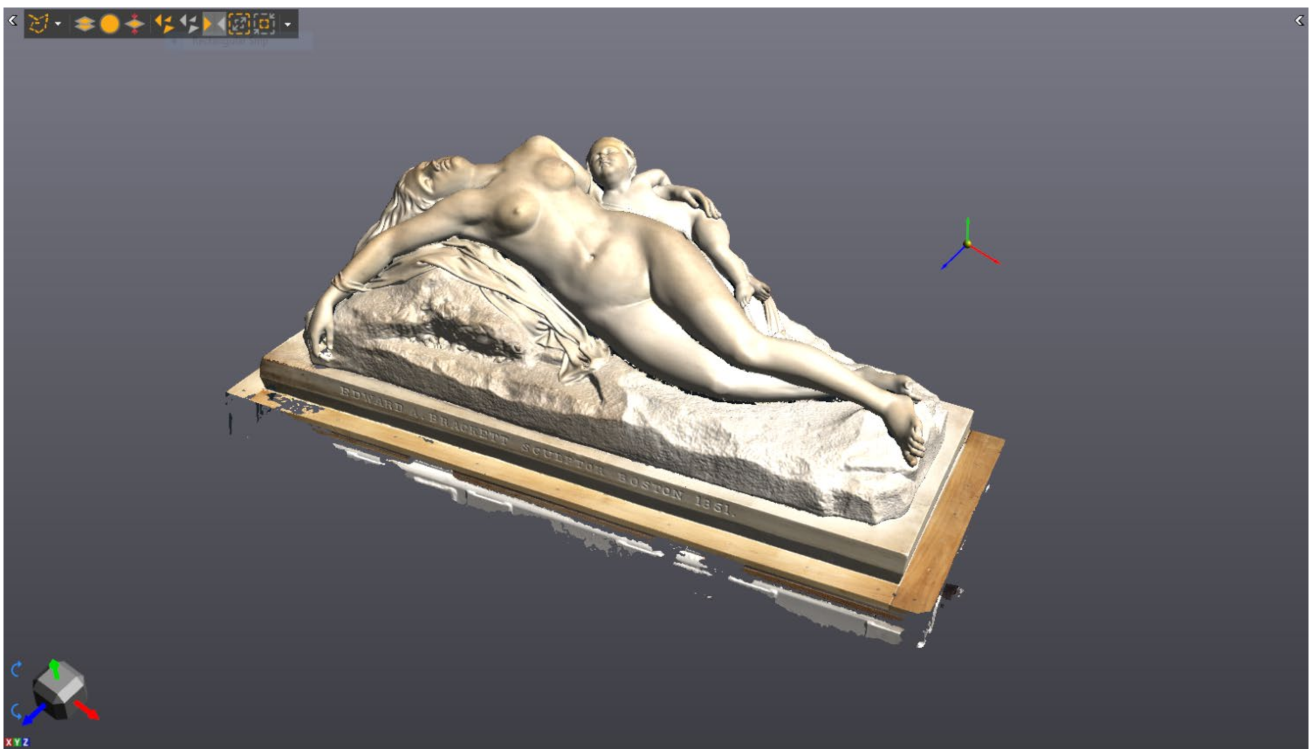

Figure 4.3: The statue in VXElements after scanning. Front right perspective

Figure 4.4: The statue in VXElements after scanning. Back Right Perspective

Figure 4.5: The statue in VXElements after scanning. Close-up of the woman's torso and the child

Figure 4.6: A rendering of Shipwrecked Mother and Child rendered in Keyshot, a professional rendering environment.

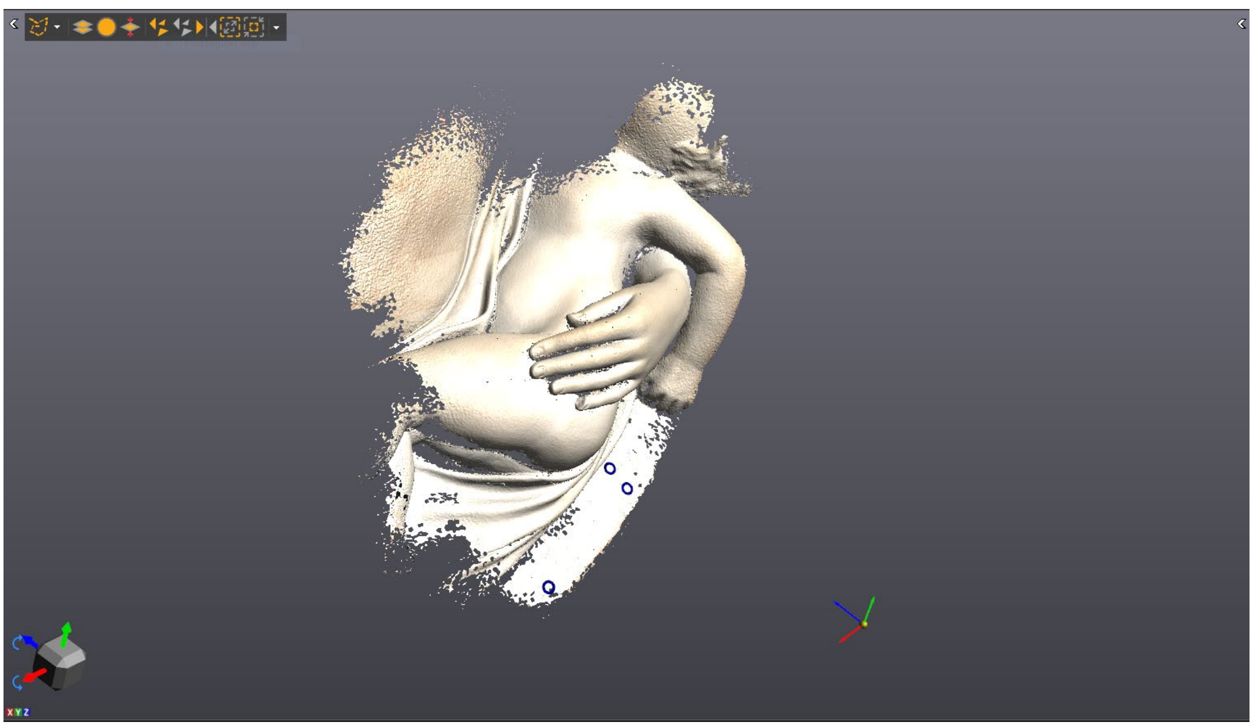

Figure 4.7: The high-resolution scan of the woman's left hand with the missing pinky visible

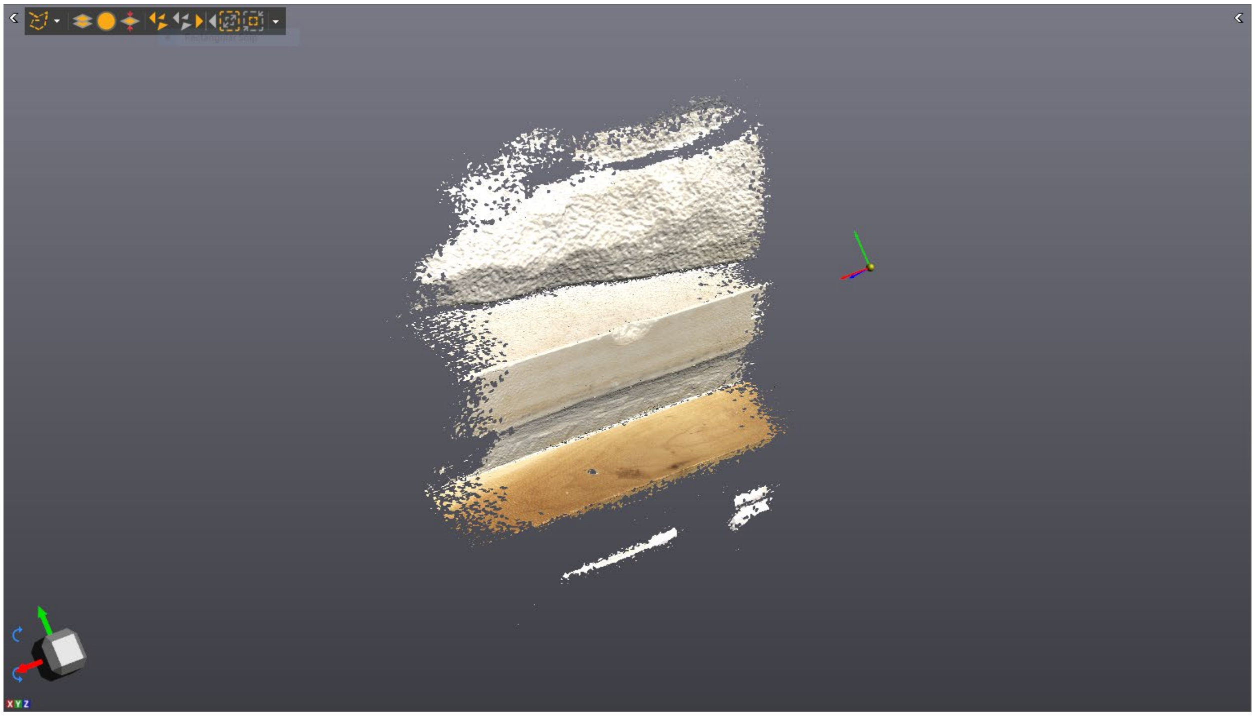

Figure 4.8: The high-resolution scan of the pedestal, with the missing chunk visible

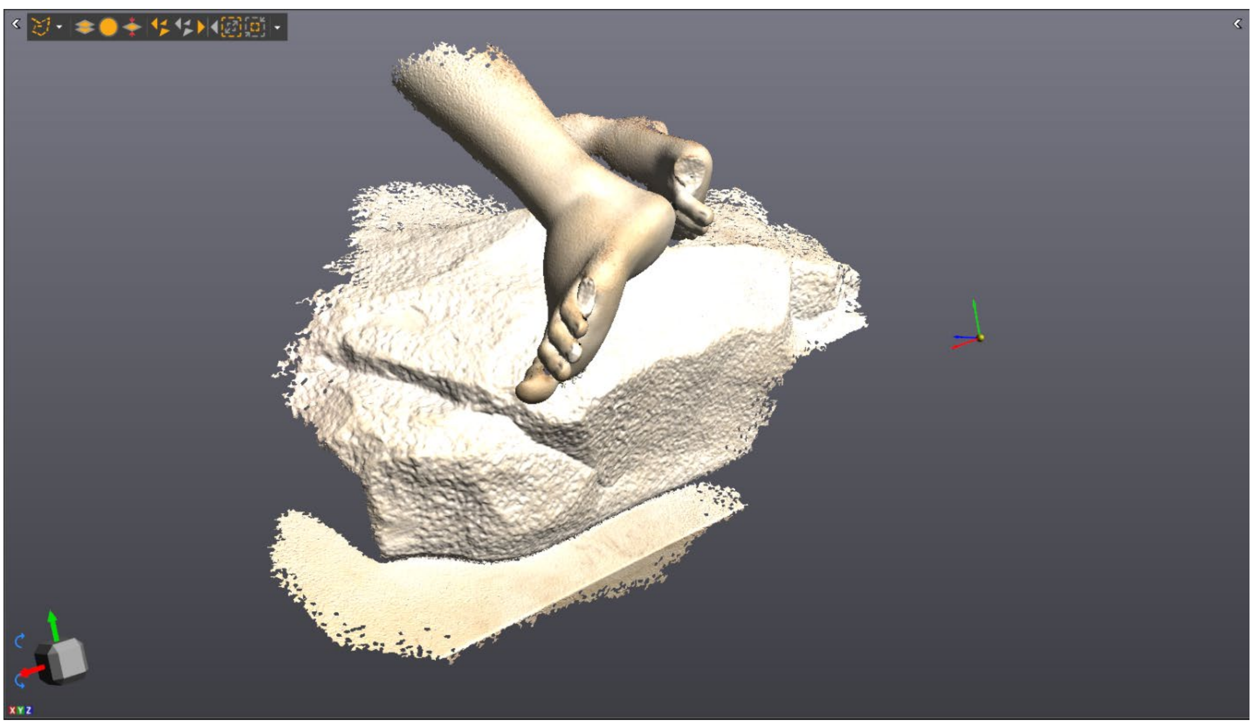

Figure 4.9: The high-resolution scan of the feet, with both missing toes visible

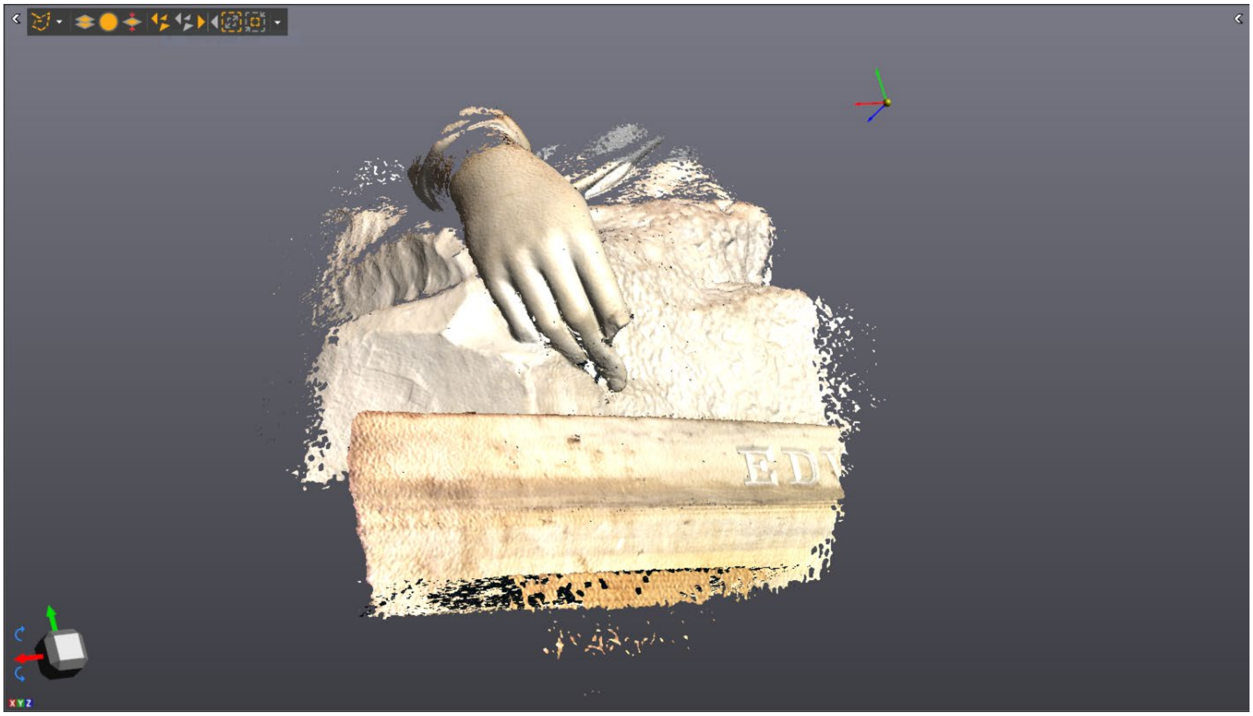

Figure 4.10: The high-resolution scan of the right hand, with the missing portion of the index finger visible

4.2. Post-processing the scans

Once scanning was completed, all scans needed a quick post-processing step to remove noisy data (refer to the steps below for how this was accomplished). To illustrate the reason for post-processing, refer to Figure 4.9 and Figure 4.10 that show two different objects: the former scan has been post-processed, while the latter scan has been. The unprocessed scan shows many floating ‘dots’ around the area of interest. They, among other artifacts, detract from the quality of the scan. To remove them from the model, we use VX Scan’s built-in toolchain using the following steps:

1. Press control+alt+c to put VXScan into the ‘connect’ selection mode, which will select an object based on what is connected to it.
2. Control+click on the area of interest: this will select the portion of the model we want to keep. Suppose the model contains multiple distinct sections that we would like to keep; control-click on them in turn to select them all.
3. Invert the selection using control+i; this will select all the floating portions of the model we want to remove
4. Finally, hit the delete key, which will delete all the noise from the model.

After the model is cleaned of noise, it is imported into VXModel, a modeling program inside Creaform’s VXElements suite of programs. In VXModel, we non-destructively patch any small holes in the scan using the ‘auto-fill holes’ menu and setting tolerance to a small number (<5 mm). These holes are formed when the scanner misses a particular detail, likely due to occlusion by other elements. This patching process uses the curvature of surrounding elements to assume the curvature of the missing pieces. It is crucial to make sure this step does not affect the elements being restored. If it does, go back and rescan the areas of interest to minimize gaps. Do note that while this same hole patching step can be done in VXScan’s scan parameters menu, VXModel has a more robust visual approach. Once this process is complete, the scan will begin to look more like the one in Figure 4.9, free of noise and messy data. Finally, use the export menu (File -> export –> Mesh or Ctrl + Shift + F) to export the scans as .OBJs for use in Blender, where the missing part generation will take place.

4.3. Generating the areas of loss in Blender

With the model scanned, we need to generate the parts that will fill in the areas of loss on the statue. To do this, we will use Blender, a free and open-source 3D creation suite, allowing us to model the missing components based on the scan data. The Blender models will then be exported as .STLs for 3D printing.

4.3.1. Why Blender?

After experimentation with various programs, we chose the latest version of Blender (2.91 at the time of writing, please reference Appendix A.2 for software requirments) as the ideal 3D modeling software for the digital restoration workflow for three primary reasons:

1. Blender's ability to work with soft body meshes and other wireframes allows for sculpting and tweaking the generated parts without going to yet another software piece. Minimizing the number of programs in the workflow was critical to keep the iteration time between prototypes as low as possible. If, for example, the missing right index finger needed to curve a little more or a little less, it would be straightforward to change in Blender's mesh editing suite.

2. Blender’s powerful Boolean tool, which allows meshes to be subtracted or joined together, providing an invaluable tool for creating negatives of the break surfaces, as will be detailed later in this Subsection.
3. Blender’s Python integration allows users to write automated scripts for placing and manipulating meshes; while the project never used this functionality, it was thought to be an essential feature to have early on.

4.3.2. Generating an area of loss: the pedestal

We first began by attempting to fill in the missing chunk of the pedestal. It proved to be the perfect piece to test and develop the workflow as there was no ambiguity about what the filled-in chunk should look like. The remainder of this Section will be written as a step-by-step guide to reproducing the mesh created for this project. Refer to this example before following the instructions in “Generating an Area of Loss: The Right Big Toe.”

1. Begin by importing the .obj we exported from VXElements into Blender. Do this with file->import->wavefront. In the new window that just opened, navigate the file explorer to find the exported pedestal mesh. Select it, then hit ‘import OBJ’ to import the model.
2. Blender will import our model in meters when the export scale was in millimeters. To fix this, use the ‘transform’ options under the ‘object properties’ menu (an orange square icon located in a vertical menu on the right side of the window) to scale the model down by a factor of a thousand. Before leaving the object properties menu, set the object’s location to be 0, 0, 0 as it will likely import a few hundred meters from

the Blender origin. Once imported and scaled, we should see something similar to Figure 4.11: the pedestal mesh with the missing chunk in Blender.

3. To begin filling it in, place a cube into the scene using the add->mesh menu in Blender's object mode.
4. With the cube added to the scene, use the move tool (mapped by default to the G key) to move the cube so that its center roughly lines up with the gap's center. Then use the rotate tool (mapped by default to the R key) to rotate the cube to line it up with the two flat faces of the pedestal. If necessary, continue moving and rotating the cube until it fits nicely. Finally, use the scale tool (mapped by default to the S key) to scale the cube until it fills the entire gap in the pedestal.
5. With the cube correctly positioned, it should fill the void in the statue; however, we must still match the hard edge of the cube to the rounded edge of the statue's pedestal. To do this, go into the modifier properties of the cube, a blue wrench icon located in a vertical menu to the right of the Blender window. Once there, hit the 'add modifier' drop-down, and click on the bevel option located under the generate menu. Set the number of segments to 12 (the more segments, the rounder the edge), and drag the 'amount' slider until the cube's corner is the same radius as that on the statue. Finally, hit control+A on Windows or command+A on Mac to apply the transformation. The generated cube is visible in Figure 4.12, independent of the pedestal, and then in Figure 4.13 as part of the pedestal.
6. Once positioned correctly, go back into the modifier properties menu on the cube and select the Boolean modifier from the generate menu. Ensure the "difference" button is selected in the Boolean properties, then click the eyedropper and select the pedestal.

Again apply the modifier to the model. The boolean difference will subtract the part of the cube intersecting with the rest of the mesh, resulting in the negative of the break surface being imprinted onto the piece. The results of this process are visible in Figure 4.14 and Figure 4.15, with the negative of the broken surface imprinted on the rear of the piece. At this point, the piece that will fill in the hole on the pedestal is now ready for printing and fitting onto the statue.

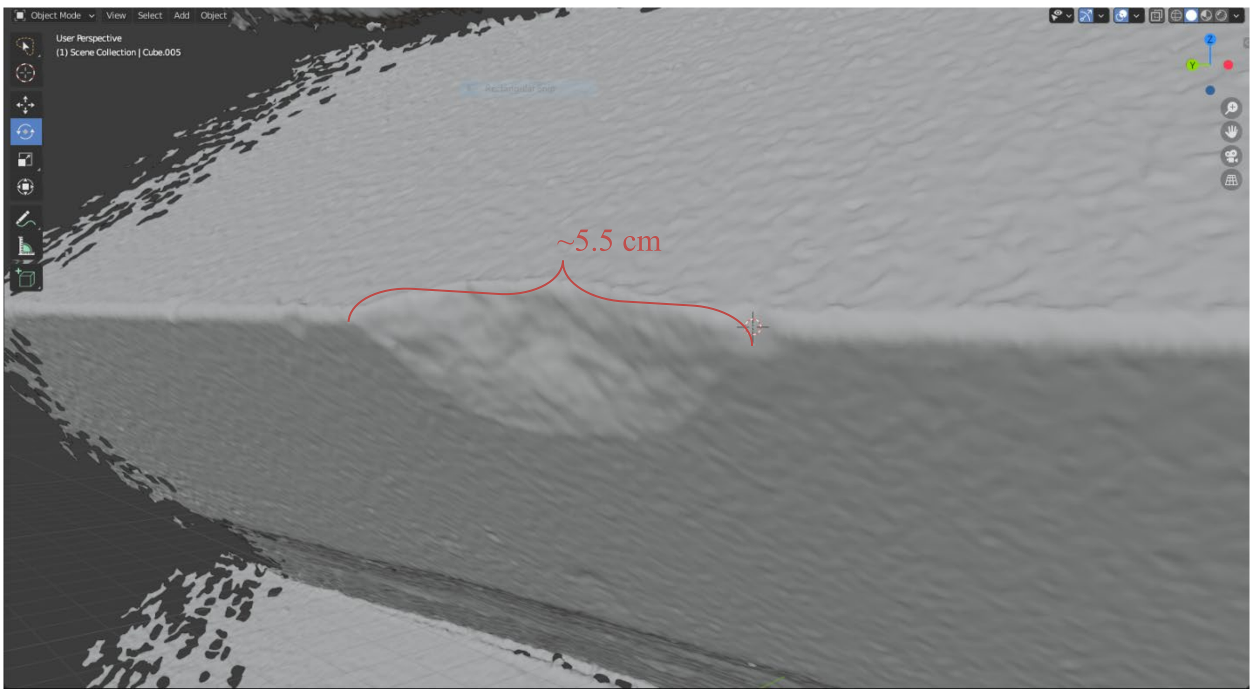

Figure 4.11: The missing chunk of the pedestal as seen in Blender.

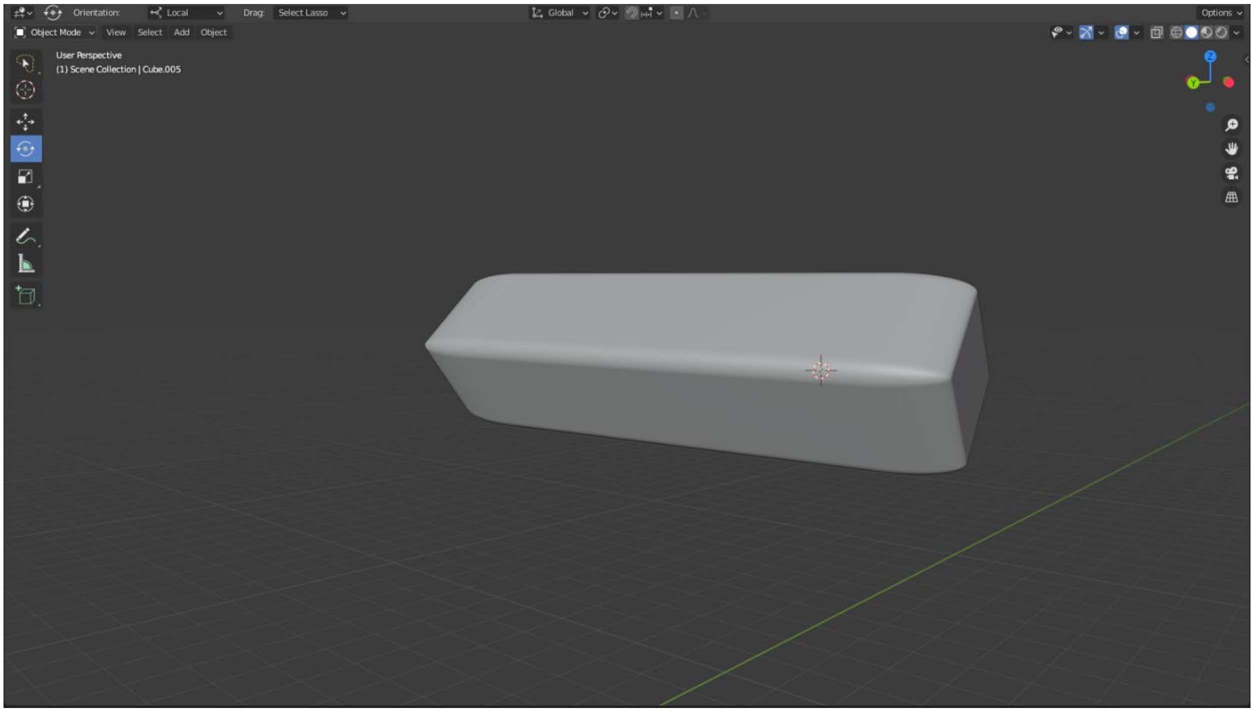

Figure 4.12: The rectangular prism sitting alone.

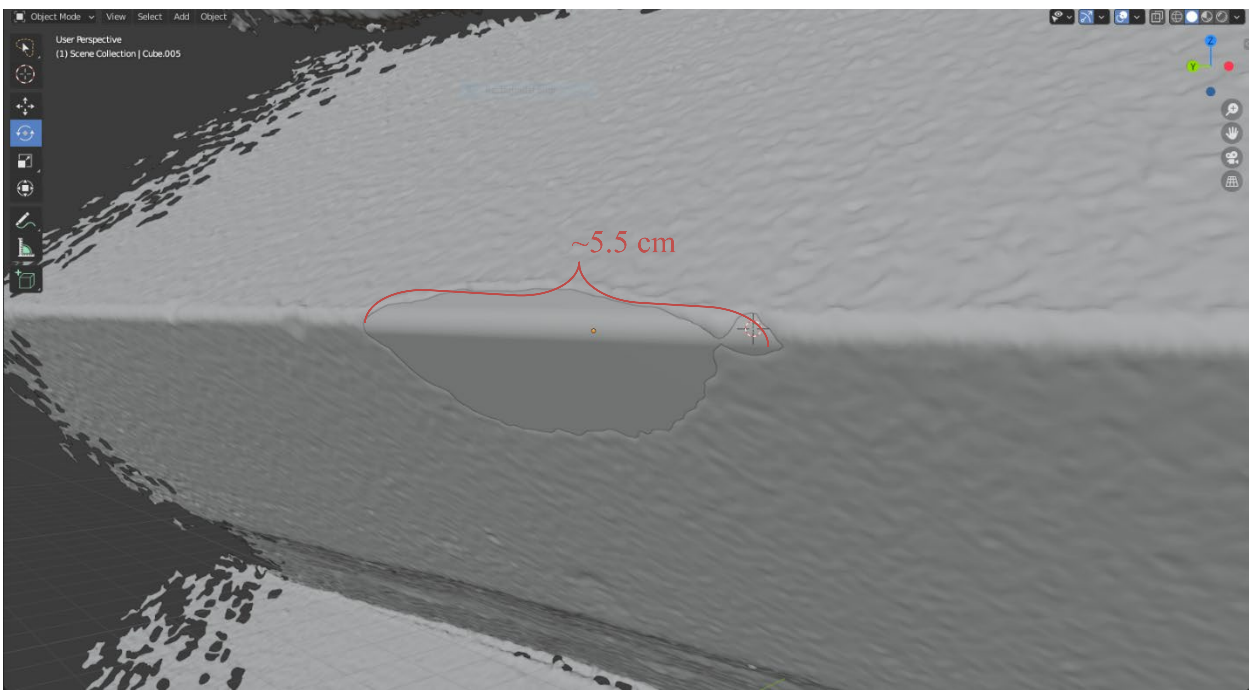

Figure 4.13: The missing chunk of the pedestal is filled in with an analogous rectangular prism.

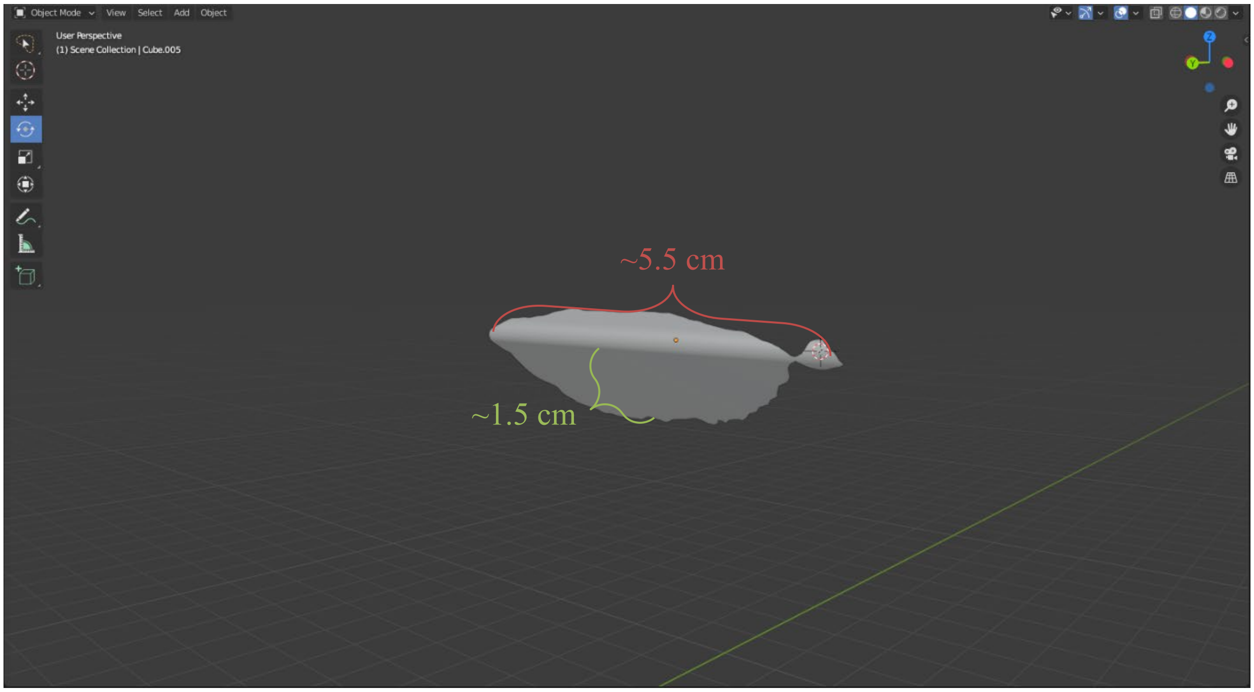

Figure 4.14: The front side of the generated pedestal piece. Width and height are marked to provide the reader with a sense of scale. It measures 5.5cm wide x 1.5cm tall x 0.5cm deep

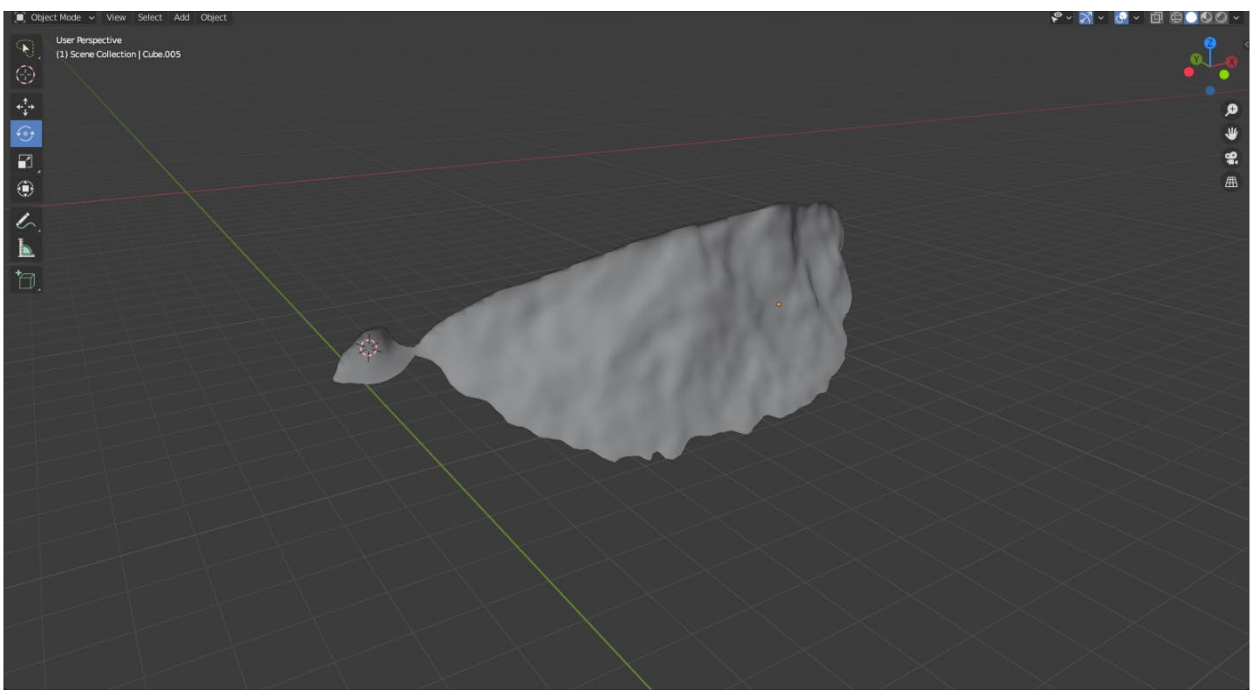

Figure 4.15: The backside of the generated pedestal piece, showing the negative of the break surface.

4.3.3. Generating an area of loss: the right big toe

Generating a big toe is significantly more complex than generating the pedestal piece. It involves the need for hand sculpting, combining multiple scans into one model, and an eye for how the toe should look on the woman's foot. For this project, we were able to get Paula Artal-Isbrand to guide us through iterations of the toe, where one printed prototype would inform how the next should be modified. Before continuing, it is important to realize that with Blender, there exist many ways of doing any given task. If the reader is an experienced user who knows how to use the program, feel free to generate a right toe using any method, the following instructions are designed for users with little to no Blender experience. With that in mind, the remainder of this Section will be written as a step-by-step guide to creating a left toe. Refer to the "Generating an area of loss: the pedestal" before following these instructions. Appendix C provides a high-level overview of the described process.

0. Unlike the missing piece from the pedestal, the toe needs reference geometry from elsewhere on the statue; in this example, we will be using the opposite toe as our starting point. Before starting, duplicate the imported scan. To do this, click on the model in scene collection, the 'tree' of items in the top right, and hit control+c then control+v to duplicate the object. Once done, hide the duplicated copy using the eye icon. This step provides a safe reference to go back to in case any mistakes are made.
1. Begin by cutting out the right toe from the high-resolution scan of the feet.
   a. With the visible object selected, use the drop-down in the top left of the window to go from 'object mode' to 'edit mode;' this will overlay the object's wireframe

onto the scan of the feet. Each dot on the wireframe represents a vertex, while each line represents an edge connecting vertices.

   b. Go into Lasso selection mode, either by hitting W on the keyboard until a squiggly circle with an arrow appears in the top left corner of the window or by pressing shift+spacebar, then hitting the L key. Also, hit alt+z to put the mesh into 'x-ray' mode, which will allow you to select vertices that are occluded by the faces of the mesh.

   c. Left-click and sweep a rough circle over the left toe. It does not matter if neighboring toes are selected as long as the entirety of the left toe is highlighted. With the toe entirely selected, hit P, then S to separate the selected vertices into a separate object.

2. With the toe cut out from the rest of the mesh, go back into object mode with the top-left drop-down. Double click on the newly separated toe, then go back into edit mode; this allows us to edit the new toe mesh instead of the foot mesh. At this point, the 'artistry' begins, so be careful and make regular duplicates of the active piece so that backups exist.

3. With the Lasso select tool enabled, begin selecting and deleting every part of the cutout that is not the left toe. Again, select by pushing and holding left click, dragging around an area, and delete by pressing the 'delete' key, then the V key. We need to use a little bit of judgment to determine what constitutes 'toe' and what does not. If in doubt, a toe should look like a cylinder with a dome on it.

4. Once only the toe is left, go into object mode and add a mirror modifier to the toe in the modifier properties menu. Apply it, go back into edit mode, and delete the

original copy of the toe (select the original copy with the lasso tool, then hit 'delete' and V). Once done, tab back to object mode.

5. Use the move and rotate tools to place the mirrored toe over the break surface of the right toe; a little artistry is required to get orientation, sizing, and placement correct.
6. Next, we must switch into edit mode to fill in the large holes on the toe model and make it a closed surface; otherwise, we cannot use the boolean subtract feature to imprint the break surface onto our new toe.
   a. Begin by deleting all vertices that protrude any more than a little bit into the foot; we need a mirrored toe that slightly interests the existing foot. Again, we need to use our best judgment for what 'slightly' constitutes, but the intersection must exist.
   b. NOTE: while Blender's built-in hole patching tools are good, the hole on the toe will likely be too big for Blender to fill it in correctly. If a more robust method than this Section is about to describe is required, the Blender community has many free and paid extensions that can help patch the hole.
   c. When patching the toe, there was one large hole on the model, split into two distinct planes: the back of the toe (where the mirror of the break surface goes) and the inside of the toe. To get Blender to patch a large hole such as the one described, we need to section off these two portions from each other by creating a new face between them. To do this, find four vertices that, when connected, will make a face that separates the two sections. Click on each vertex while holding the control key, then hit the F key to create a face; once done, we should be left with two holes, each on a distinct plane.

d. Now that the hole has been separated into distinct planes, we can solidify the model by patching the smaller respective holes. Begin by clicking on a vertex on the edge of a hole, hold alt, then double click on a neighboring vertex also on the edge of the same hole, which will cause Blender to select the loop of vertices that run around the edge. Finally, hit “alt+F” to close the hole. Do note, using alt+F can be a little finicky and create floating vertices; make sure to go around the edge of the loop and delete any vertexes cut off by the patch.
e. Repeat the same process for the second hole and any other hole on the toe.
f. With that done, now would be an excellent time to duplicate what we have to save our work.

7. With the hole patched, we need to remesh our object to make it sculptable. Skipping this step will make touching up the model an arduous and painful process.
   a. Begin by going back into object mode, then go to the object data properties in the right vertical menu (the icon is three green dots in an upside-down triangle). There click on the ‘remesh’ drop-down, select the ‘Quad’ option, and hit the ‘QuadriFlow Remesh’ button, opening a new window.
   b. Within this window, leave the default checkboxes alone (Note: tick on ‘Preserve Mesh Boundary’ if detail is lost during remeshing), set the face count to something larger than 20,000 (the higher, the better, but the more work for the computer), and hit the ‘Ok’ button.
   c. Note: If a yellow error pops up saying something like “the object needs to be a manifold,” double-check that you have no floating vertices or self-intersections on the model; revisit the hole patching step if necessary.

8. If the remesh worked, then only one step remains: sculpting the toe to match the foot. Again, this was done in collaboration with Paula Artal-Isbrand from the Worcester Art Museum, who watched over the process and gave notes about where material could be added and removed.
    a. To begin, swap from object mode into sculpt mode using the same drop-down in the top left of the Blender window,
    b. In sculpt mode, we will use the various tools to push and pull the mesh to match the flow of the toe. We recommend playing around with the different tools to find what works best for you. When making the pieces for WAM, we used the 'draw,' 'inflate,' and 'smooth' brushes the most. A brush's size and strength can be controlled with the sliders at the top of the window.
    c. The majority of the sculpting work will happen around the base, shaping the toe to match nicely with the break site on the mother's foot. Be careful only to make minor modifications around the base of the toe where it meets the foot; we are not trying to enforce a vision of how the toe should look, rather ensuring a smooth integration between the parts.
    d. With the sculpting work complete, it is recommended to pass over the entirety of the toe with a low-strength smooth brush to get rid of any major imperfections (careful not to destroy the nail when doing this).
    e. Finally, swap back into object mode and apply a boolean difference to the toe, and imprint the break surface on it, just as was done with the pedestal piece.

It should be pointed out that there are a few differences with the process used in the project compared to the ‘simple’ one described above. Firstly, as shown in Figure 4.17, the inside of the right toe was a noisy mess of data, so it had to be combined with the scan of the toe on the stick to create a smooth and continuous toe. This step was done using Blender’s additive Booleans, first taking an average between the mirrored right toe and the toe on the stick, then stitching them together. This process created a rather “ragged” surface which was then smoothed out using the smooth brush. The result of this combined then smoothed toe is visible in Figure 4.18. However, this process is not strictly necessary for filling in occluded areas. Sculpting the missing features is a fantastic solution as long as care and attention are placed on maintaining the artist’s style.

We used the same techniques described in this Subsection to generate the left pinky toe, the right index finger, and the fingertip. By cutting and manipulating existing portions of the mesh, Blender provides all the necessary tools and more to digitally recreate lost portions of the statue. Skilled artists with a firm grasp of the software might even be able to recreate a missing appendage without using scanned reference data, allowing for more artistic interpretations of the loss that provide a sort of “naturalness” that simple copy-pasting can not.

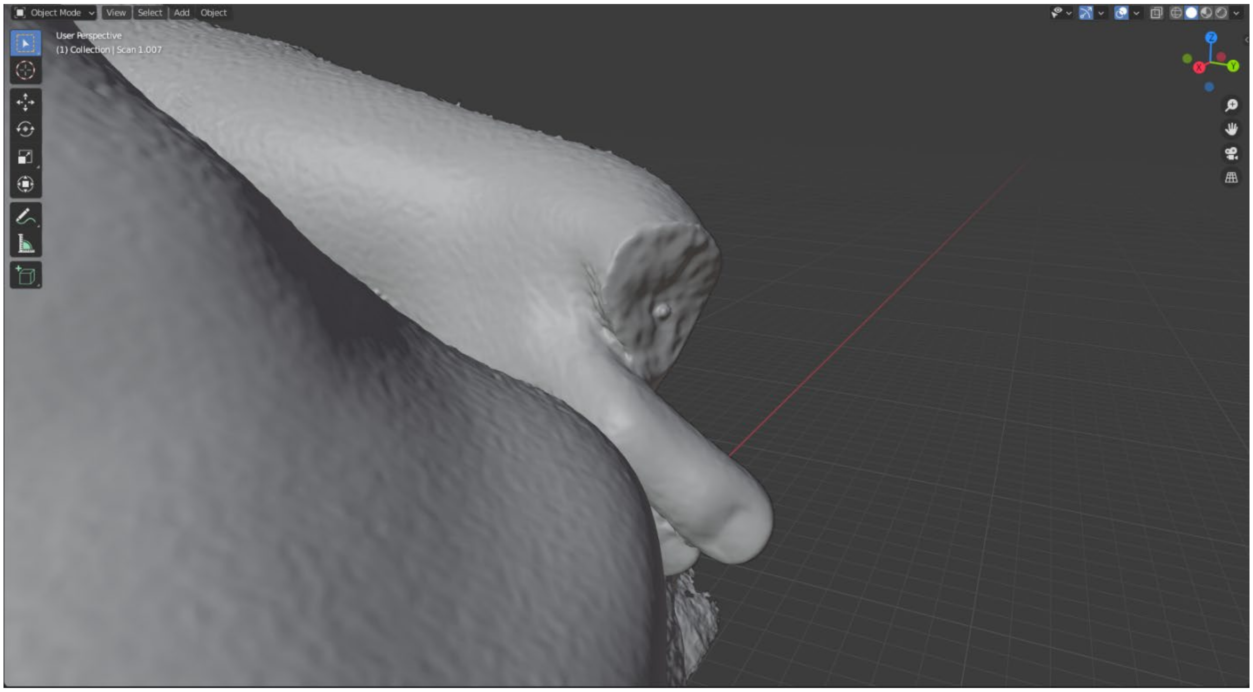

Figure 4.16: The missing right big toe as seen in Blender.

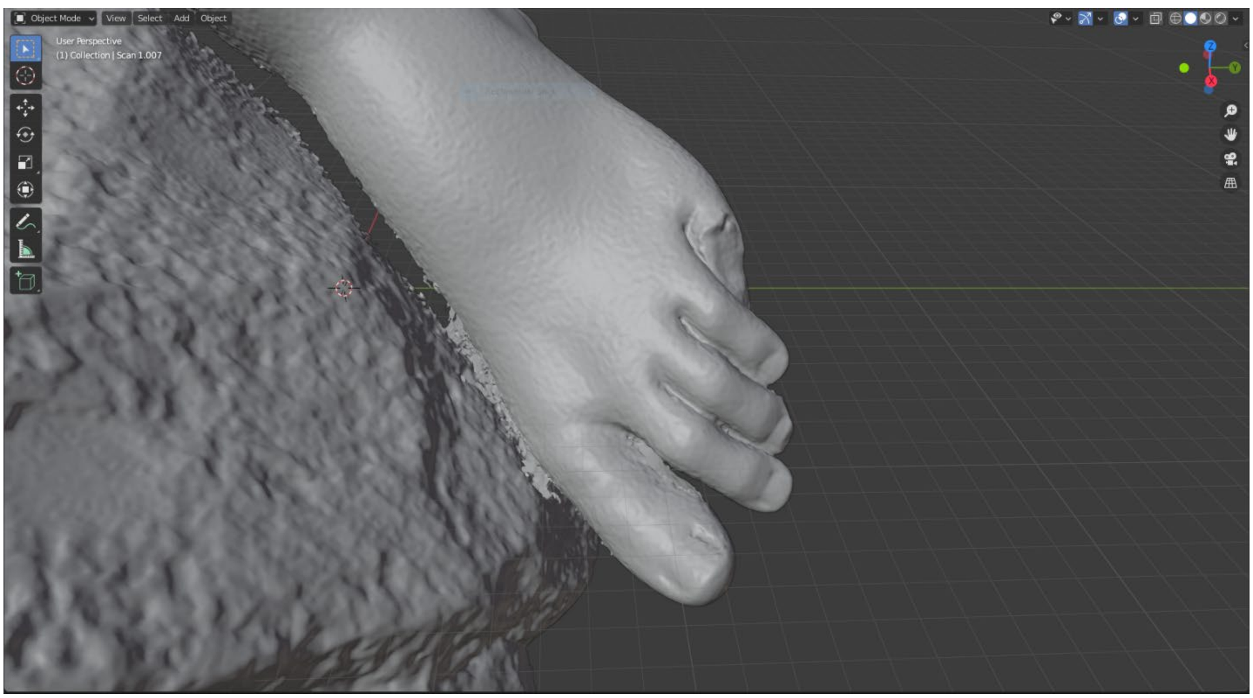

Figure 4.17: The existing left big toe, which was used to recreate the right toe.

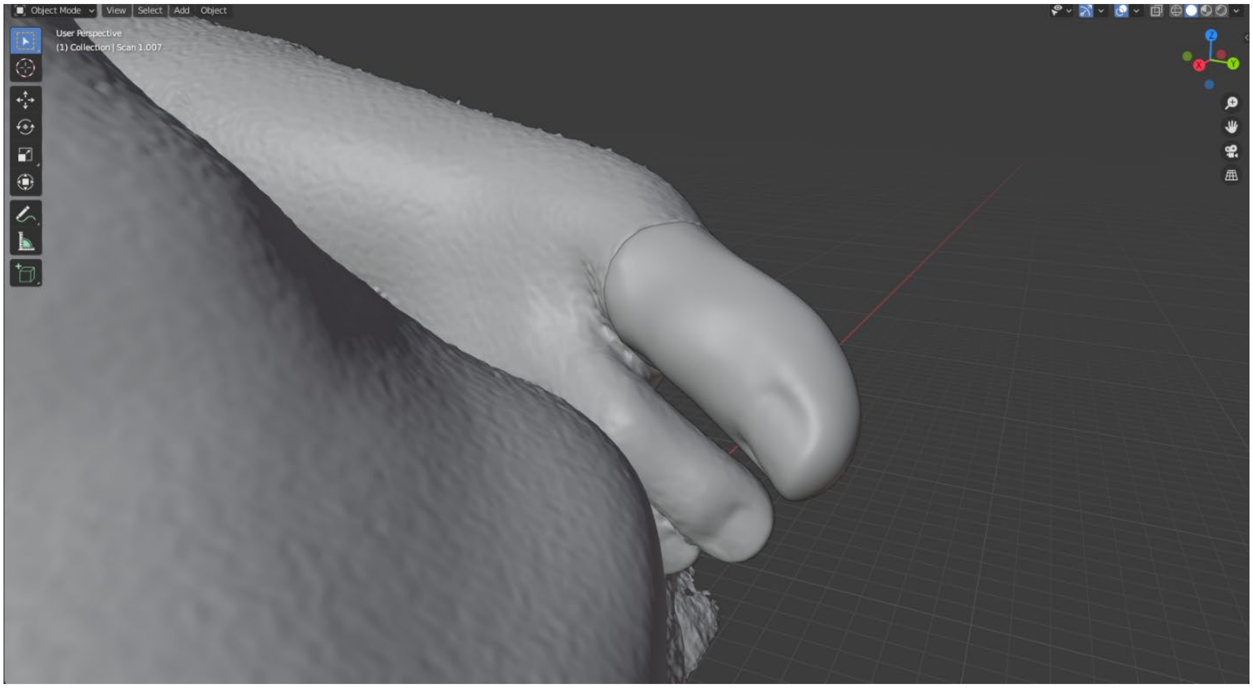

Figure 4.18: The recreated left toe, based on a mirror of the right toe, stitched to the left foot.

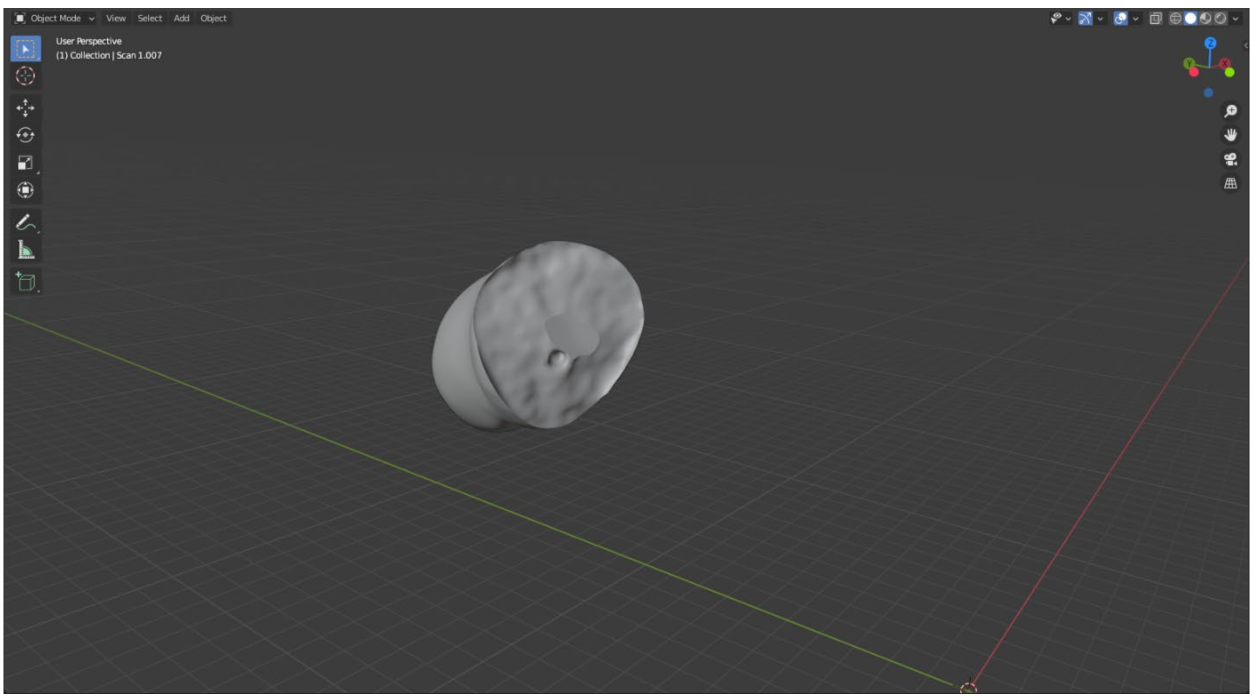

Figure 4.19: The backside of the recreate toe, showing the negative of the break surface.

4.4. Augmented reality: an experiment in rapid iteration

Per recommendation from the project's advisor, Professor Cosme Furlong-Vazquez, we attempted to use a Microsoft HoloLens 2, a pair of augmented reality goggles, to see the generated pieces on the statue before printing them out. See Appendix B.2. for the full technical specifications of the HoloLens 2. The idea behind the recommendation was that iteration cycles could go much faster when print time is removed from the equation. Before loading the models on the HoloLens 2, export them from Blender using the following steps:

1. Export the models as glTFs (in the file->export menu); this will open a new window to set export properties.
2. Check the 'limit export to selected objects' box (located under the 'include' drop-down) and select the object we would like to export from the scene collection in the top right of the main Blender window.
3. Finally, in the top right of the export window, a drop-down menu swaps between various glTF export types; exporting as .glTF is preferable as the file size is much smaller and will not cause the HoloLens to lag. However, .glTFs will not always load onto the HoloLens, so we must swap to the alternate .glB file format in those cases.

We chose to use the glTF 2.0 file format for two reasons. First, the more common wireframe formats (.obj, .stl) are limited to 10,000 vertices on the HoloLens; since the models used in the project have over 20,000 vertices and remeshing an in-progress part before each export is an unnecessarily slow step, it made sense to find an alternate solution. The second reason was that the HoloLens' built-in model viewer could not scale the models to their actual size.

Please note that before loading the models onto the HoloLens, it is necessary to complete the first-time setup and connect a Microsoft account to the device. Once the device is configured, use the following steps to view the 3D models on the HoloLens:

1. Turn the HoloLens on and plug it into a computer, then open the file explorer. It will appear as a new drive, much like a USB stick would. Next, click on the HoloLens folder and move the exported files into the '3D Models' folder.
2. Next, put on the device and open the Microsoft Store by performing the start gesture (holding your wrist in front of your face, then tapping it), then selecting the icon that looks like a shopping bag.
3. Using the search bar, find the 'glTF Viewer' application from Mike Taulty and install it using the 'get' button. This app allows the viewing and manipulating of glTF files on the HoloLens; without it, we would be limited by the HoloLens' 10k vertices constraint. With the glTF Viewer app installed, perform the start gesture once to open it.
4. To open the models, say the word "open" out loud; this will open the HoloLenses file browser into the '3D models' folder.
5. From there, select the model we want to view, and it will automatically be placed into the world. To move the model, grab onto it with a hand and move it around, as in Figure 4.20.

Figure 4.20: Interacting with generated chunk in the HoloLens.

Figure 4.21: The user’s hands as seen through the HoloLens.

Figure 4.22: An outside perspective of Nathan Kaplan using the HoloLens. Image taken by Paula Artal-Isbrand [11]

4.4.1. Shortcomings of using augmented reality for sculpture restoration

After a two-hour session where we tried iterating both a piece of the pedestal and the little toe, both Mr. Kaplan and Ms. Artal-Isbrand found the HoloLens frustrating to use for such a use case. An example of a common problem is visible in Figure 4.23 and Figure 4.24: pieces that appear aligned from one angle are not aligned from a different viewing angle. Furthermore, precisely placing items in space is frustrating and challenging, even if the alignment issue did not exist. The Hololens seems to lack the necessary sensitivity/ resolution to place objects in space accurately. In other words, trying to line up virtual objects with physical ones precisely is an arduous task. Because of these issues, there was no way to confirm whether the break surfaces of the pedestal correctly lined up with each other or whether the pinky toe fit appropriately into the rest of the foot. While there is definite potential for such technology as applied to art conservation, time constraints prevented us from exploring the idea further. We concluded that augmented reality is still too much in its infancy to be applicable for sculpture restoration.

Figure 4.23: The generated chunk of the pedestal, in red, as seen through the Hololens. At this viewing angle, it appears to line up with the rest of the statue.

Figure 4.24: The generated chunk of the pedestal, in red, as seen through the Hololens. At this second viewing angle, it is clearly out of line with the rest of the statue.

4.5. Printing out of PLA and physical iteration

Once all parts had been generated in software, we began printing prototypes out of PLA on FDM printers. These parts were printed on the Zortrax M200, located in the CHSLT lab at WPI, and sliced in Zortrax's proprietary slicer, Z-Suite. See Appendix A.3 for Z-Suite's software requirements and Appendix B.3 for the M200's full technical specifications. The files were printed with the highest setting available on the printer: a 0.09mm layer thickness and a 'high' print quality. Also, break surfaces were printed as vertically as possible, which would prevent layer height restrictions from over-constraining their shape. With careful slicing and print settings, we could get rather stunning quality out of the printer. Figure 4.25 and Figure 4.26 both show the printed pedestal piece filling the gap on the statue. This part marked the first successful restoration of a missing piece of the sculpture.

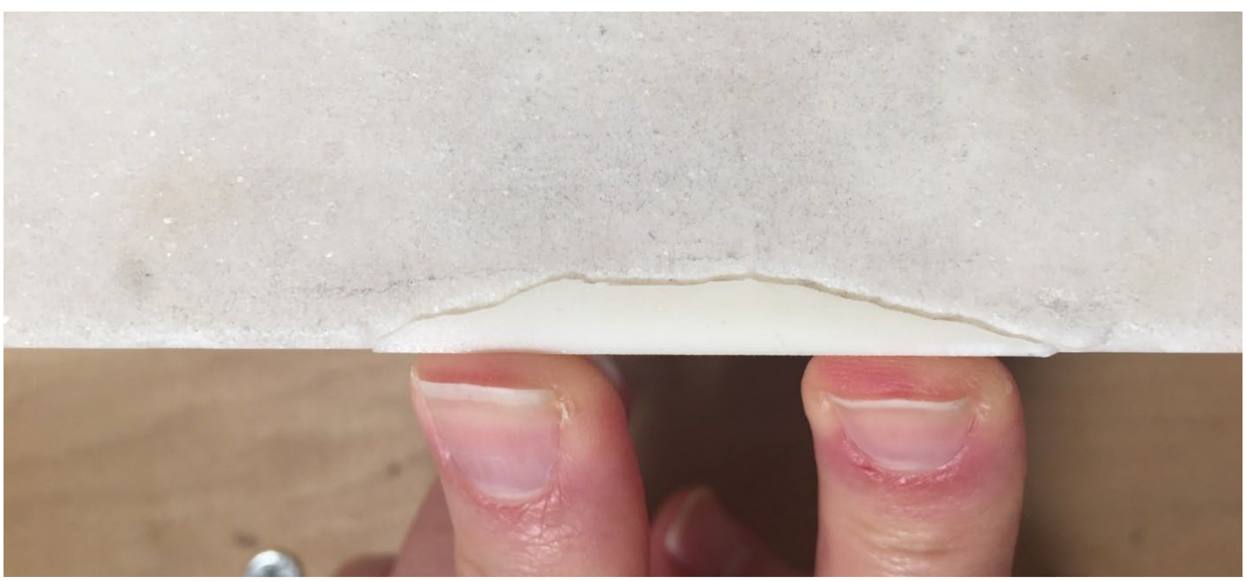

Figure 4.25: Printed pedestal piece fitting into the missing chunk of the statue. Top View.

Figure 4.26: Printed pedestal piece fitting into the missing chunk of the statue. Front View.

Not all parts were perfect from the get-go, and some required several iterations before they even began to look right. Figure 4.27 shows one such example with the right index finger. There, the finger is undersized in comparison to the rest of the hand, with a clear gap visible on the left side of the seam in the left image and a missing area above the seam on the right image. This feedback was taken back to Blender through the developed iteration workflow, and the finger was remade. Figure 4.28 and Figure 4.29 show off the next iteration.

Figure 4.27: Poor fit on an iteration of the right index finger

Figure 4.28: Improved Finger compared to the previous iteration. Back view.

Figure 4.29: Improved Finger compared to the previous iteration. Side view

5. Results and discussion

Due to restrictions caused by the COVID-19 pandemic and a late start to the project because of difficulties acquiring the Academia 50 scanner, the project was cut off before completion. This chapter will present each of the five generated pieces and discuss whether they were finished, where we succeeded, and where we needed improvement. Afterward, there will be a brief discussion on the 3D printing research that did not make it into the project due to the aforementioned time constraints. Note: All final software iterations can be obtained through Professor Cosme Furlong-Vazquez at the CHSLT lab at WPI.

5.1. The right big toe

We begin with the right big toe, deemed one of the three ‘complicated’ pieces due to its size, complexity, and reliance on plaster casts to be generated. The big toe was the nearest to completion of the three complex pieces, fitting the broken surface well. However, it needed one or two more iterations before being signed off. As marked by the red circles in Figure 5.1, the final iteration of the toe was too thin and did not correctly fill the break site on the mother’s foot. These images show that on both sides of the toe, a bit of break surface is visible. The error likely occurred due to over-aggressive smoothing in Blender, which ended up flattening the surface and causing the gap. To address the fit issues, we did another iteration in Blender (shown in Figure 4.18 and Figure 5.2), henceforth referred to as the “final software iteration,” as it was never printed out. Furthermore, the final software iteration also addressed a unique challenge presented by the metal pin in the middle of the big toe break site that was poorly dealt with in prior toe versions. In early iterations of the toe, like the one visible in Figure 5.3, the pin caused trouble due to low-quality scan data, visible in Figure 5.4, which entirely missed the pin and other major features of the break site geometry, leading to the poor fit issues. To fix this, we rescanned the statue's feet; however, whether it was the geometry or the material of the pin, some aspects prevented it from being adequately scanned. This data inaccuracy resulted in a 'hook' shape, visible in Figure 5.5, which had to be manually sculpted based on reference images into the more accurate shape visible in Figure 5.6.

Despite being the largest of the generated pieces, the big toe was a successful example of how our developed digital restoration workflow could be applied to an actual statue restoration. Even though the final printed iteration had fit issues, it flowed well with the overall geometry of the foot. Moreover, as described in the "generating the right big toe" Section of the methodology, it was easy to generate based on the left toe's reference geometry despite the lengthy process. The changes between the final printed and final software iterations also proved how Blender's sculpting suite of tools would allow for quick refinement of the model without needing to go back to square one.

Figure 5.1: The final printed iteration of the toe, with circles marking the fit issues around the broken surface.

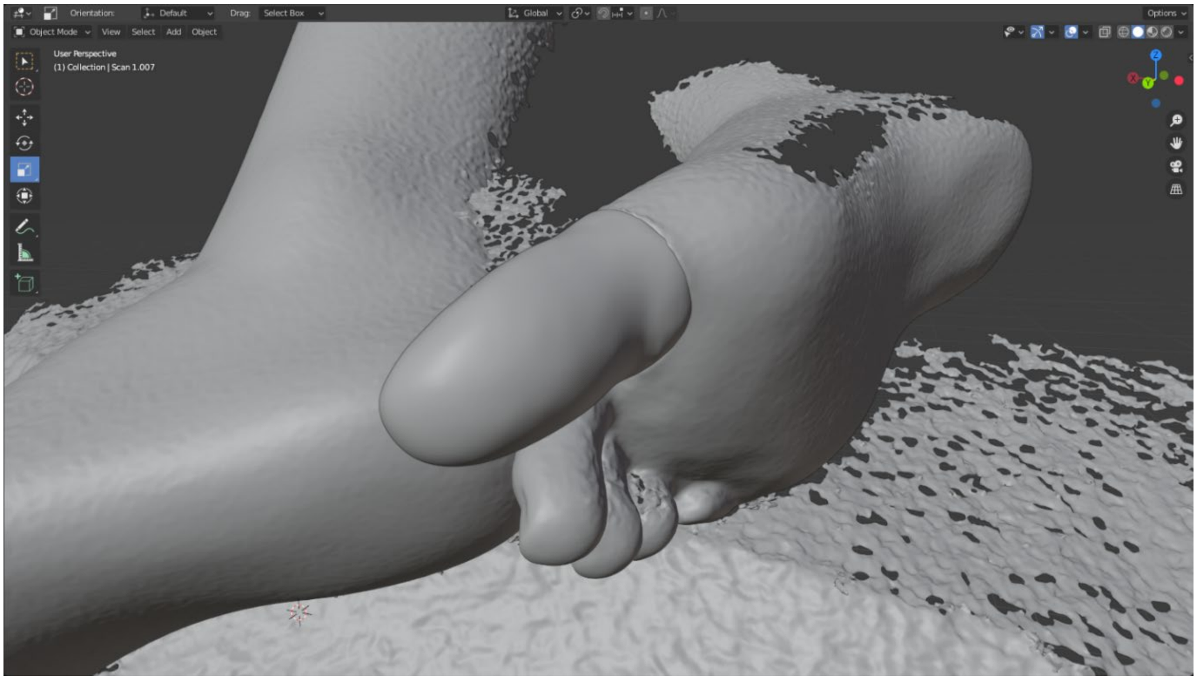

Figure 5.2: Final software iteration of the right big toe, with no fit issues around the break surface

Figure 5.3: An early iteration of the right big toe, with fit issues around the break surface marked in a red circle.

Figure 5.4: Low-quality scan data used to generate early iterations of the big toe.

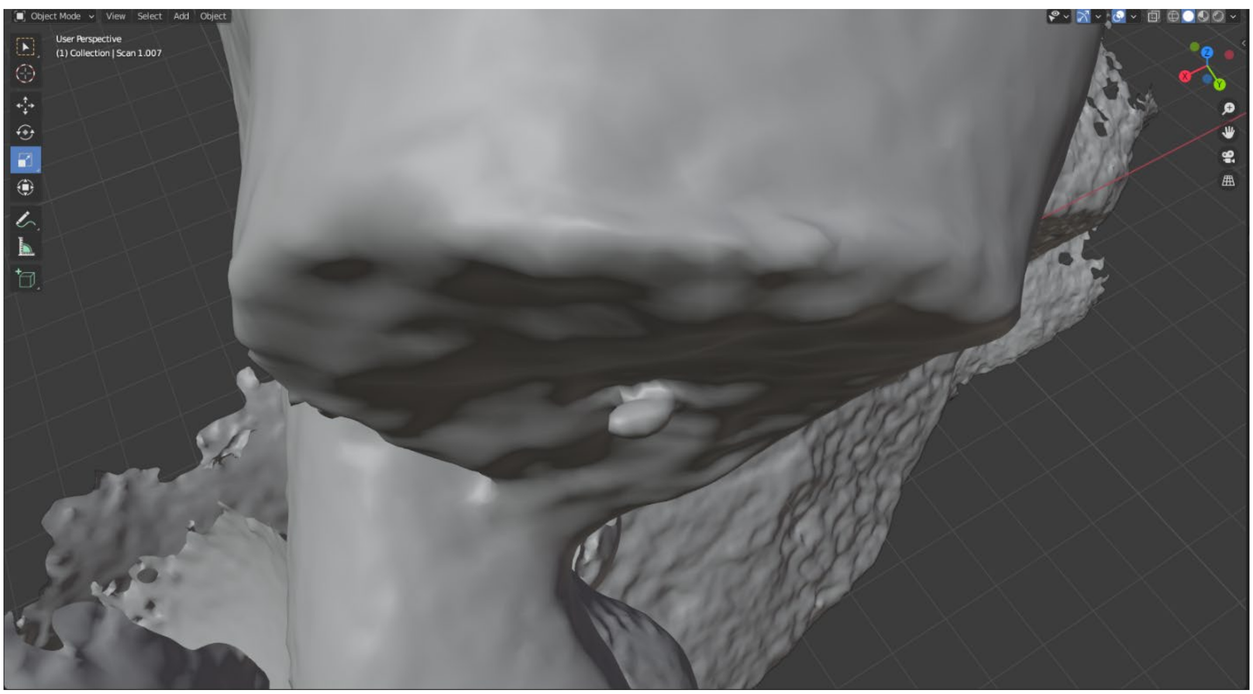

Figure 5.5: The ‘hook’ of the metal pin caused by poor quality scan data.

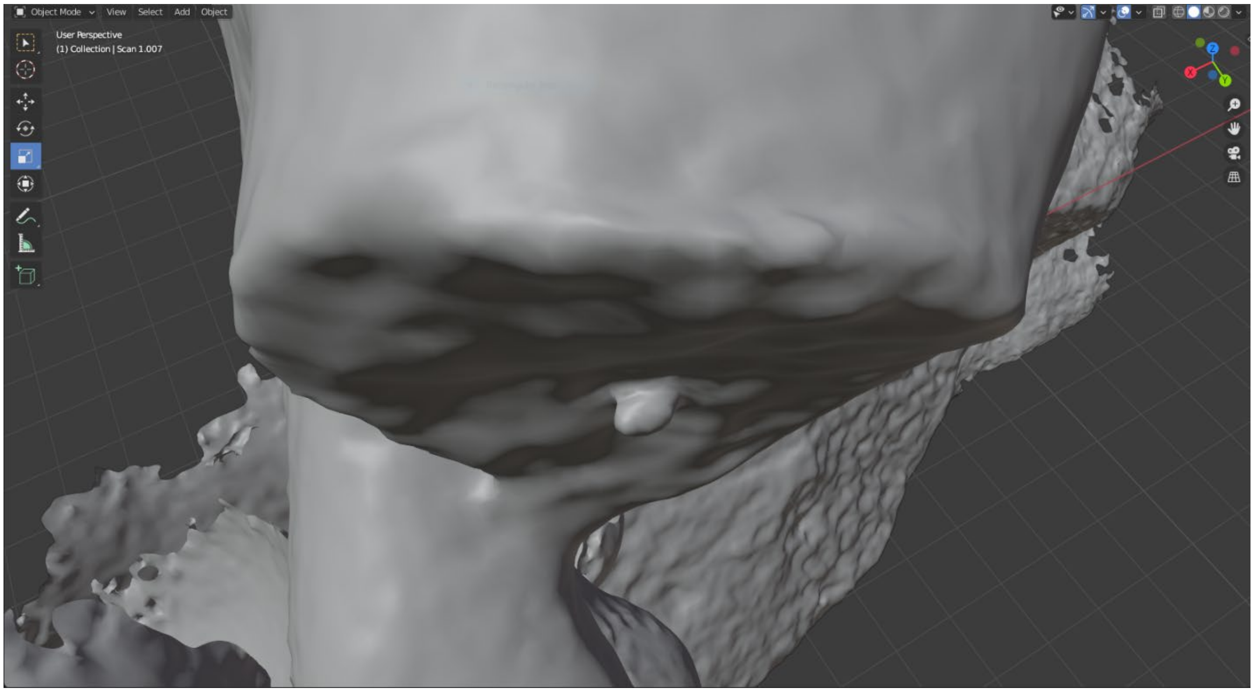

Figure 5.6: The hand-sculpted version of the pin based on reference images of the toe.

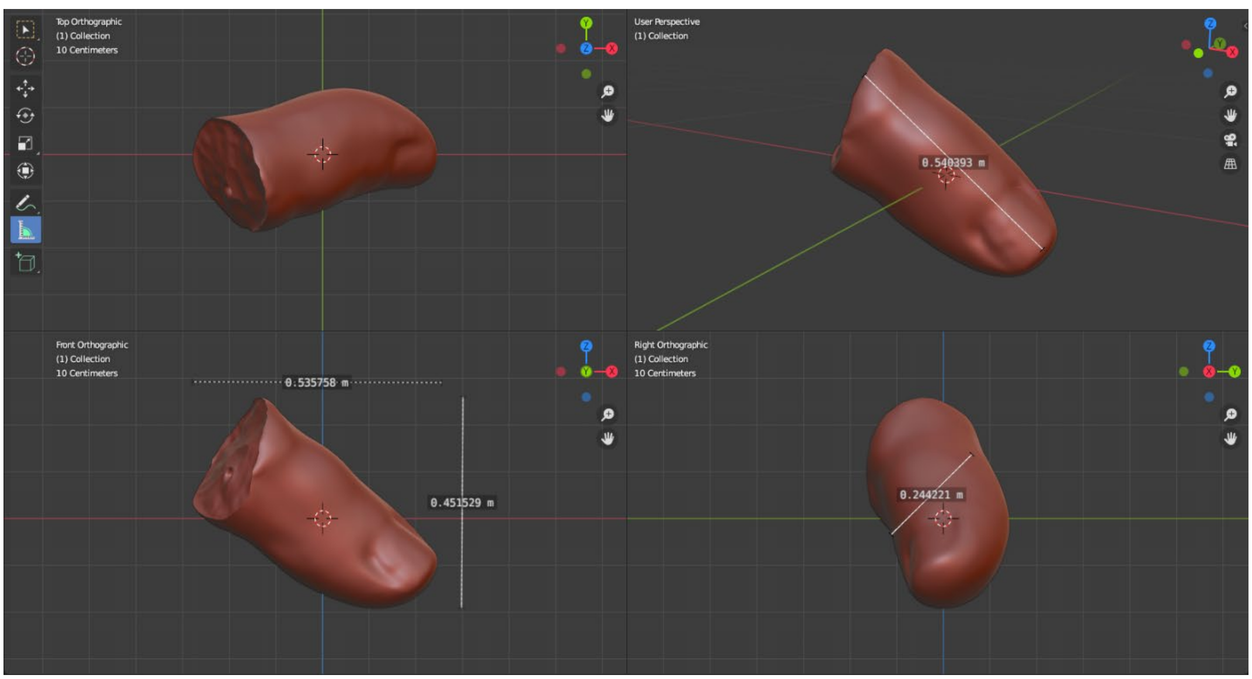

Figure 5.7: Orthographic quad view of the big toe. Critical dimensions marked.

5.2. The left pinky toe

The least complete of any of the five components, the left pinky toe only saw two printed iterations before the project's conclusion. Of the three complicated pieces, the pinky toe proved to be the most challenging to generate due to a lack of suitable reference geometry. However, per a recommendation from Artal-Isbrand, we took the scan data from the child's foot and used its little toe as a reference for the woman's. Doing this created a unique issue where the scale of generated object differed from that of the reference, adding another variable (on top of position and rotation) to control during the generation process. Despite this, the last printed iteration of the pinky toe, visible in Figure 5.8, did not look out of place on the statue. It curled in much the same way as the other three small toes and even maintained the bump after the nail distinct of Brackett's style. Due to concerns surrounding the time constraints near the end of the project, many early iterations of parts were sliced with large layer heights and low precision settings to expedite printing. While these rough parts did a decent job of capturing the overall shape of the generated part and allowed quick decisions about their sizing, placement, and rotation, they did not fit correctly onto the statue, as visible in Figure 5.9 and Figure 5.10, requiring them to be reprinted using new settings. The left pinky toe was the only generated component that did not see a high-quality printed iteration, nor did it see a final software iteration. As such, what is visible in Figure 5.11 represents where the toe was left off: fitting nicely onto the statue after some manual cleaning of the inside but not fitting snuggly enough to hold itself without added pressure from a human.

In the face of these issues, the left pinky toe proved how our digital restoration workflow could recreate a piece of the statue without ideal reference geometry. In the future, this foundational work could open avenues to recreating fingers from toes or arms from legs while remaining true to the original artistic vision.

Figure 5.8: Second printed iteration of the right pinky toe. Despite the low print quality, the pinky toe looks aesthetically correct with the rest of the foot.

Figure 5.9: Poor fit around the break site caused by the low-quality print settings. Top View.

Figure 5.10: Poor fit around the break site caused by low-quality print settings. Bottom view.

Figure 5.11: The toe fitting better after cleaning the inside.

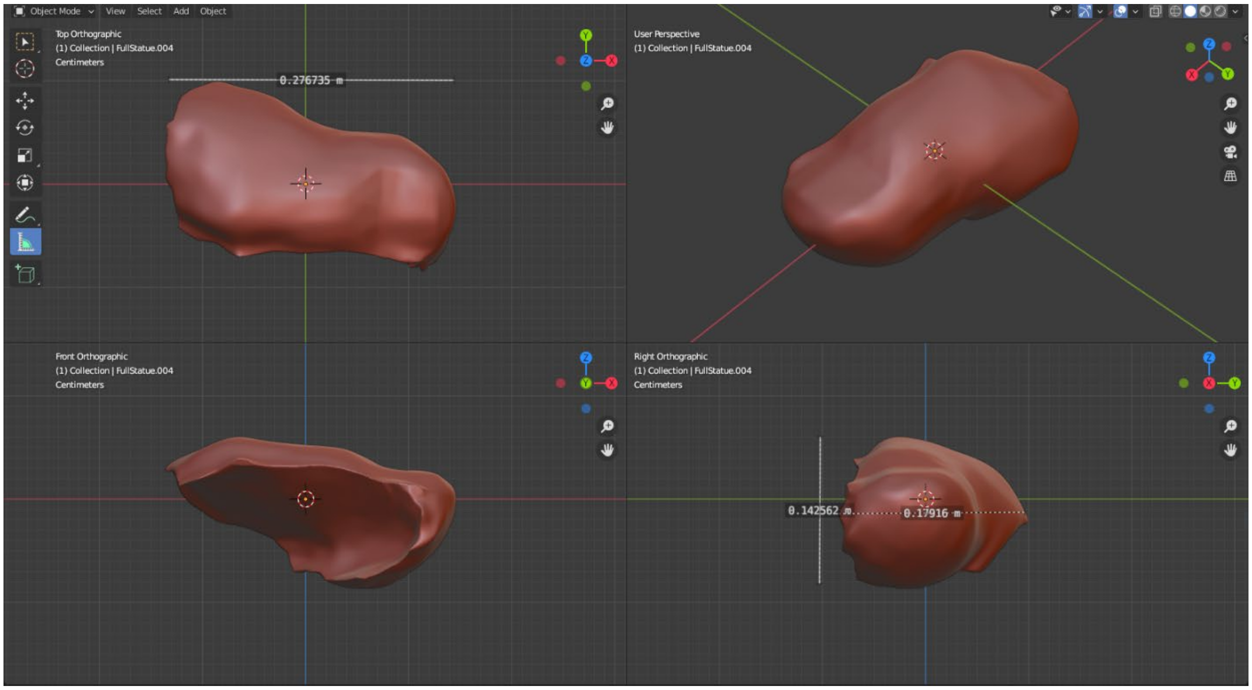

Figure 5.12: Orthographic quad view of the little toe. Critical dimensions marked.

5.3. The tip of the left pinky finger

Following the initial experiments with the missing chunk of the pedestal, the tip of the pinky finger was the first generated piece to use reference geometry from elsewhere on the statue. To do this, we sliced off the tip of the left ring finger, following the same process used to slice the left toe in the methodology Section. Since the ring finger is slightly larger than the pinky finger, the tip had to be scaled as well as rotated and transformed into position. When scaling, we used both the fingernail and the existing finger to size the generated piece. Finally, we used Blender's sculpting tools to touch up the nail so that it followed a single continuous curve; as with the rest of the touchups, it was crucial to be light-handed to not enforce our vision of what the sculpture should look like. Figure 5.13 shows the finger after sculpting, with the curve of the fingernail nicely meeting at the seam between the generated part and the statue. Figure 5.14 shows this realized in high-quality PLA, with the tip of the pinky snuggly fitting onto the statue. However, the snug fit was not achieved on the first try due to an excess of material underneath the finger, preventing it from sitting flush with the break site (Figure 5.15). Nevertheless, through the developed Blender workflow, lowering the amount of material underneath the finger was an effortless process, allowing it to fit flush against the child's thigh and fit perfectly onto the mother's finger.

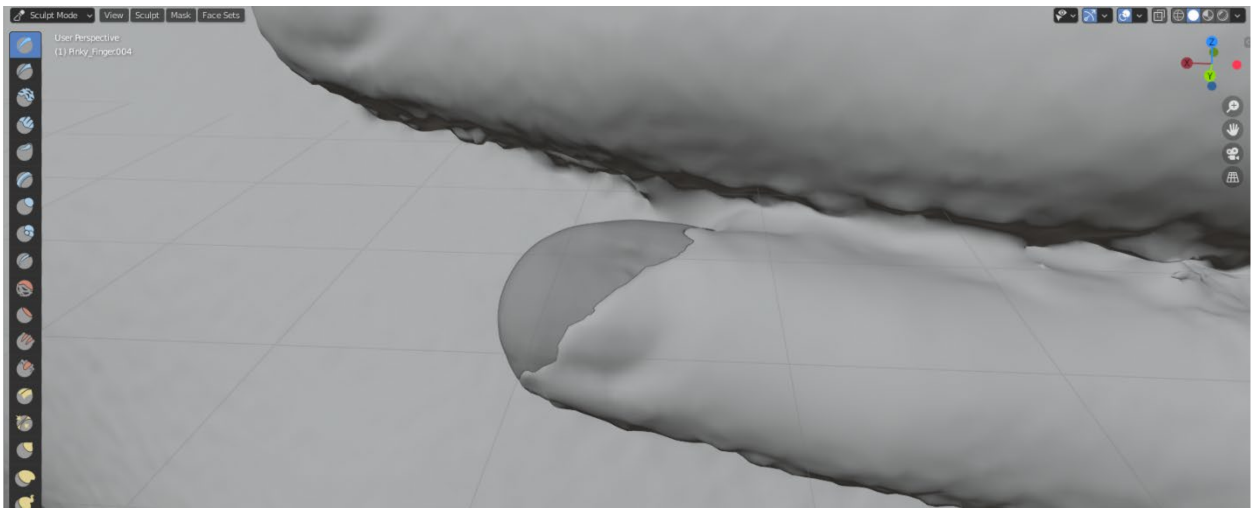

Figure 5.13: The fingertip after sculpting touchups in Blender.

Figure 5.14: The last printed iteration of the fingertip nicely integrating with the rest of the hand, a small seam is visible.

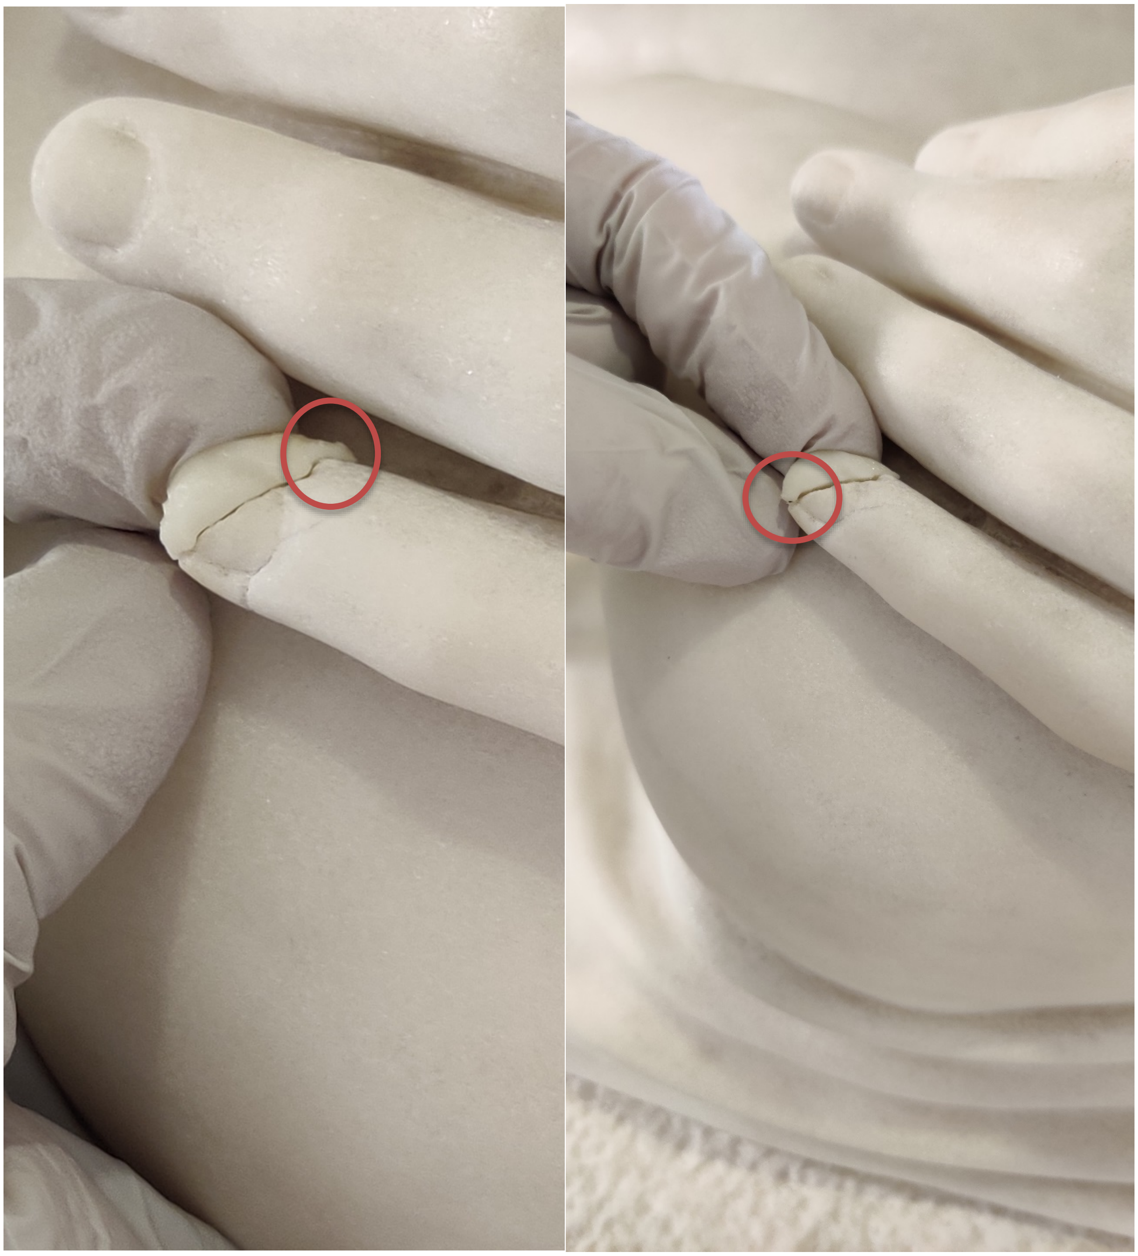

Figure 5.15: Poor fit issues with the first iteration of the fingertip.

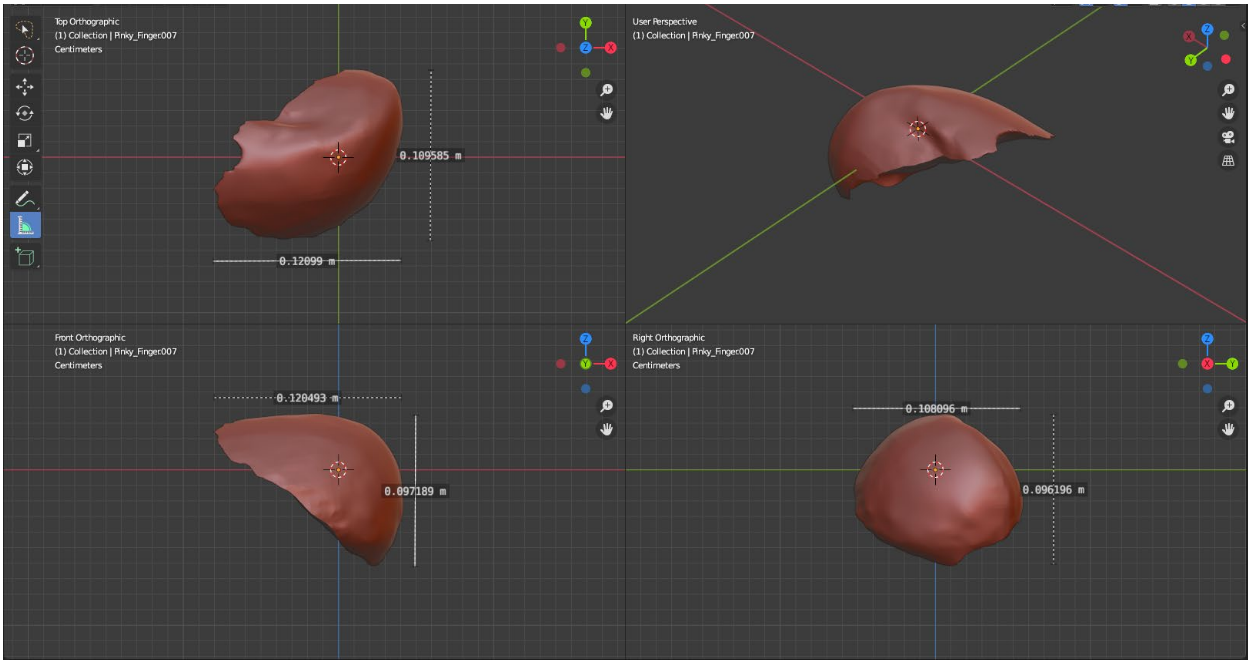

Figure 5.16: Orthographic quad view of the pinky finger. Critical dimensions marked.

5.4. The right index finger

The last of the complex pieces, the right index finger, was the most challenging piece to generate because of its complicated break surface and overall geometry. Unlike the left pinky toe, many references existed for the finger design: from the neighboring fingers on the right hand to the left index finger and the plaster cast made by Artal-Isbrand. Instead, the difficulty with the right index finger came from needing to integrate it smoothly with the existing knuckle, requiring a significant amount of iteration and thought into how Brackett would have sculpted the finger. One hint we had was the cascading curling pattern of the other fingers: starting at the pinky and traveling inward, each finger was straighter than the last. Furthermore, the left hand provided insights into how much shorter the index finger should be than the middle finger. As previously discussed in the 'printing' Section of the methodology, our initial attempt at generating the finger, visible below in Figure 5.17, missed the mark. Not only was the generated finger thinner than the knuckle, but it was also missing a section that needed to cover the break site on the knuckle. Two iterations later, we arrived at the one visible in Figure 4.28, Figure 4.29, Figure 5.18, Figure 5.19, and Figure 5.20. While a considerable improvement over the starting point, the final printed iteration of the finger still had three issues, each of which was resolved in the final software iteration, visible in Figure 5.21 and Figure 5.22. First, despite adding material to fill in the knuckle, there was still a gap between the edge of the finger and the edge of the break site, which is best shown in Figure 5.19. Unfortunately, knowing how much material to add to cover a break surface is more of an art than a science. While it looks like the break site at the knuckle is completely covered and no longer exposed to air in the final software iteration, there is no way to confirm this without doing another printed iteration. Second, the finger was too wide compared to the mother's other fingers: visible in Figure 5.18, the inside of the finger seems 'inflated' compared to how it should be. A red dotted line marks what should have been the edge of the finger in Figure 5.20. The final issue present in the printed iteration is the seam just above the fingernail, caused by an oversight during an iteration of the index finger. While much work remained to be done, the finger proved how the developed workflow could successfully generate complicated geometries.

Figure 5.17: First iteration of the right index finger. Undersized and missing a portion that should have covered the break site on the knuckle.

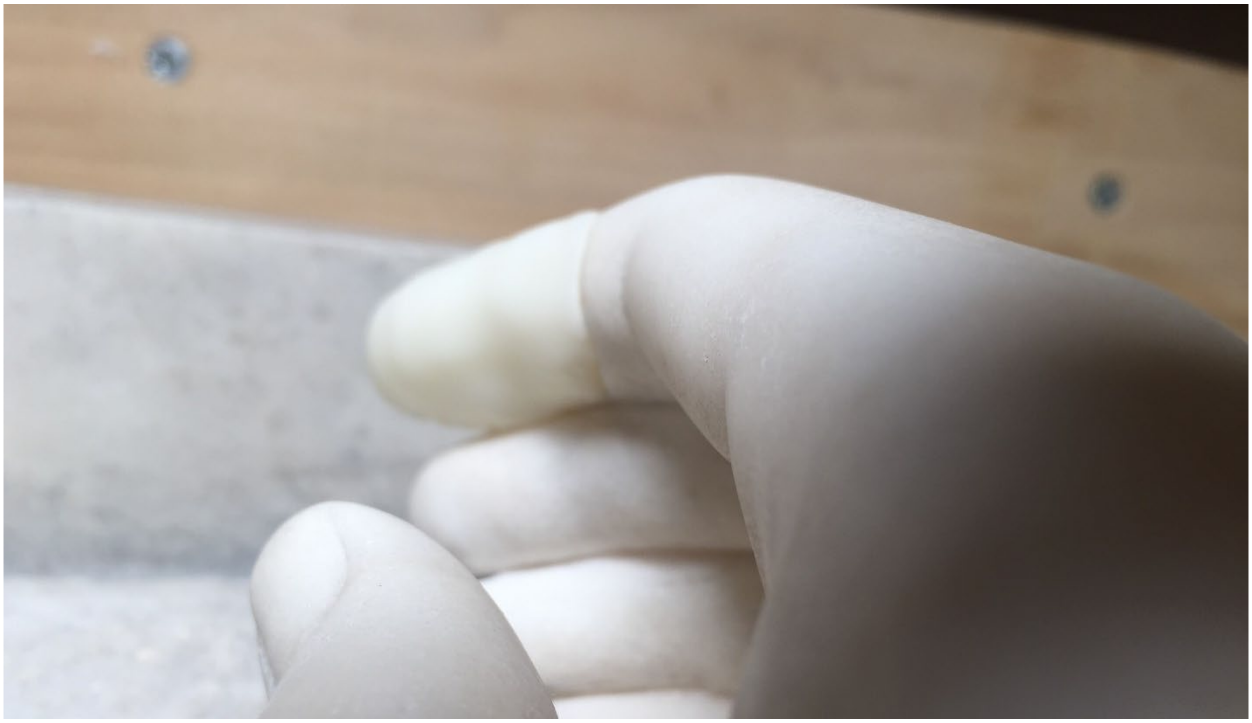

Figure 5.18: Final printed iteration of the right index finger. Inside Top View.

Figure 5.19: Final printed iteration of the finger. Front View.

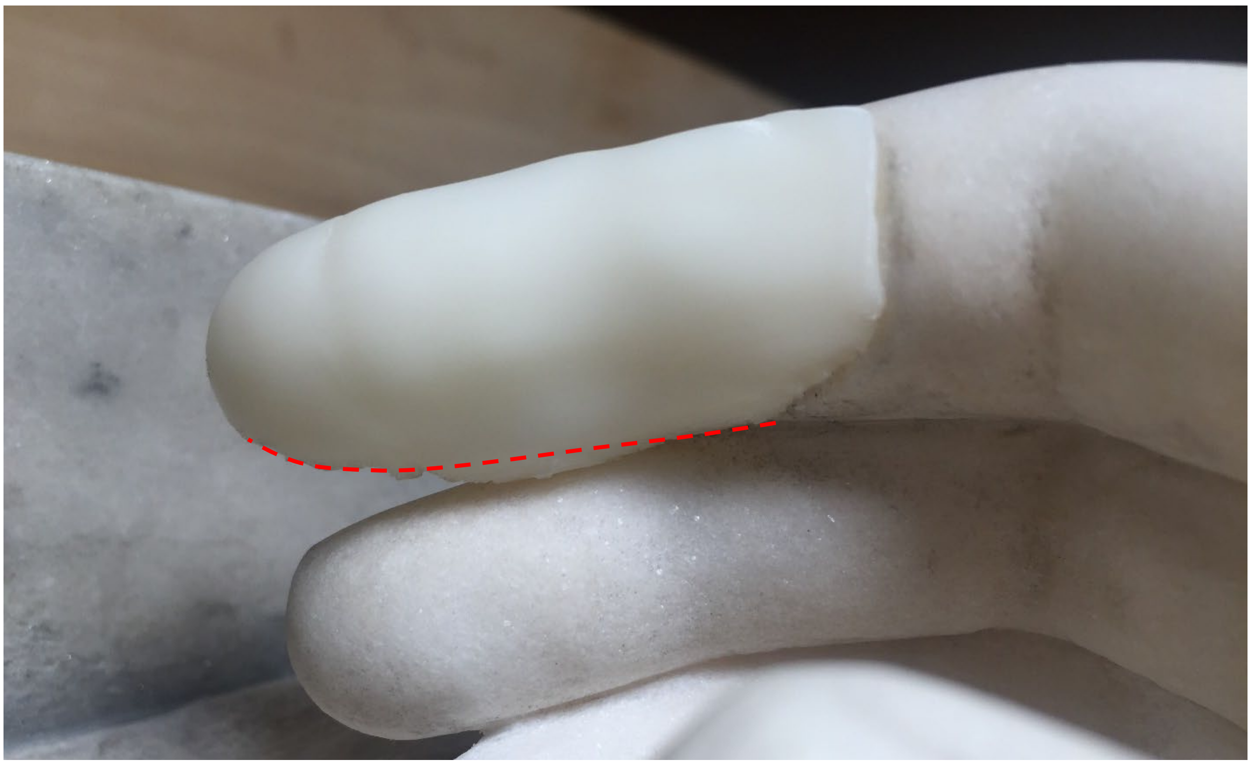

Figure 5.20: Final printed iteration of the finger. The red dotted line marks what should have been the edge of the finger.

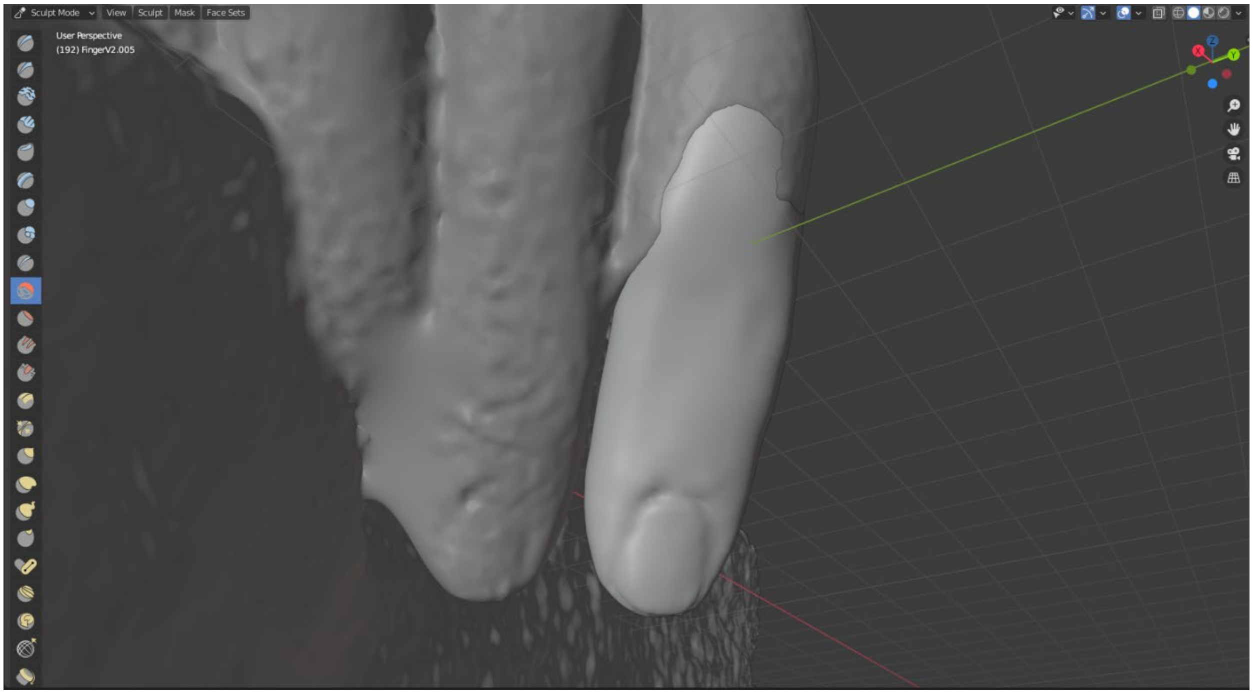

Figure 5.21: Final software iteration of the finger. Front View.

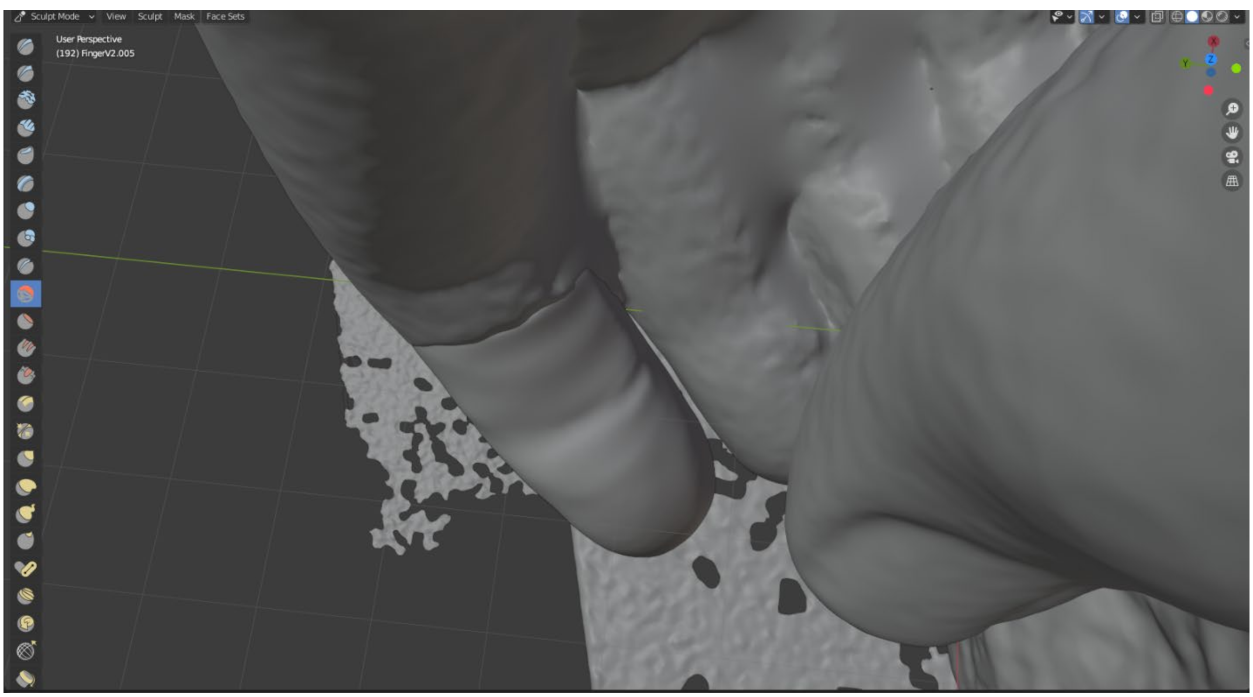

Figure 5.22: Final software iteration of the finger. Back View.

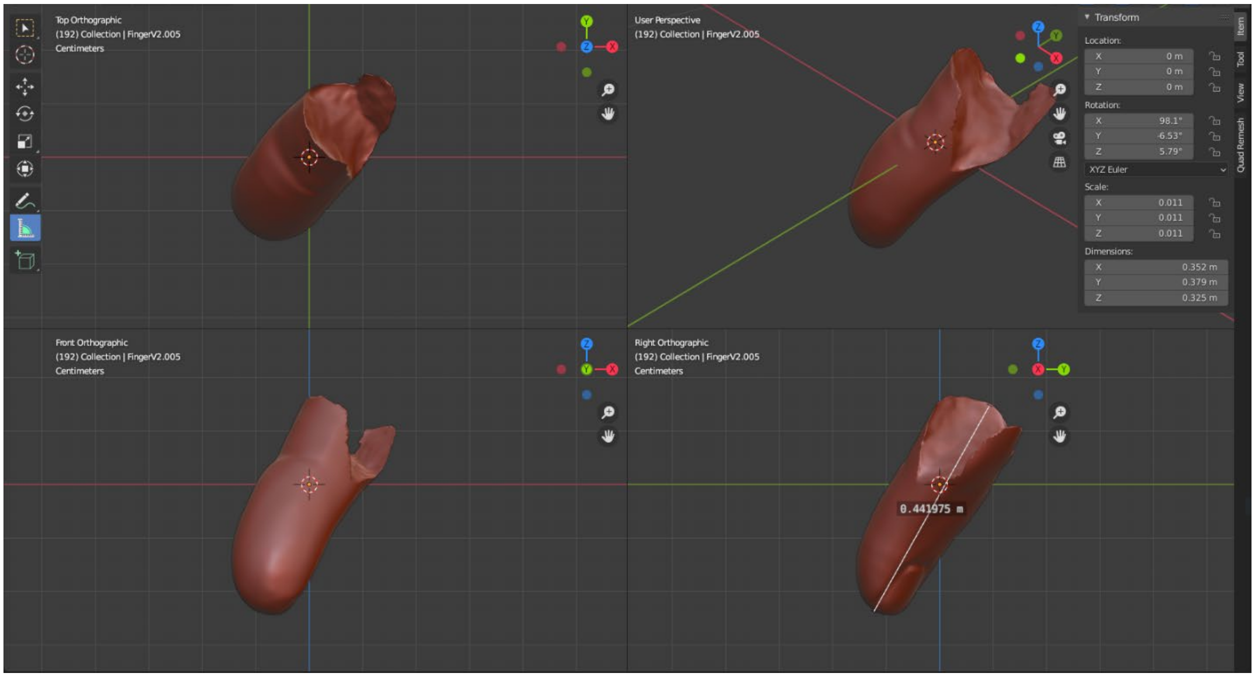

Figure 5.23: Orthographic quad view of the finger. Critical dimensions marked.

5.5. The pedestal

Creating the missing portion of the pedestal was a straightforward process as it did not require any reference geometry from elsewhere on the statue. As detailed in the methodology, the printed piece was based on a rectangular prism with rounded edges. Because of this simplicity, the design integrated well with the statue from the first iteration. Once the model had been confirmed, the piece was used to test the updated slice setting on the Zortrax printer and later the new printing technologies, which will be discussed in the next Section. Figure 5.24 shows how the final printed version of the pedestal piece fits almost seamlessly into the statue.

Figure 5.24: A printed piece cleanly filling in the missing portion of the pedestal

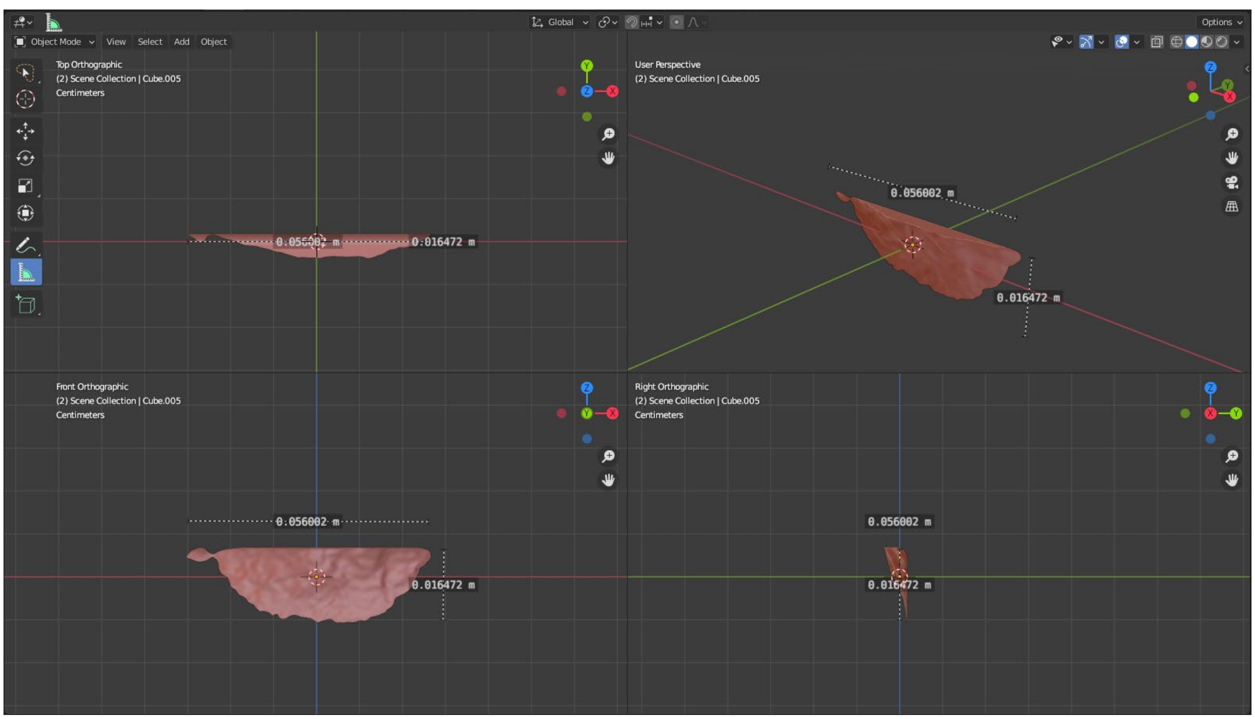

Figure 5.25: Orthographic quad view of pedestal. Critical dimensions marked.

5.6. Discussion on novel 3D printing methods

Near the end of term, work began with high-resolution versions of the parts on the Stratasys Objet PolyJet printer located on WPI's campus and multi-color PolyJet printers at AET labs, a local company from which we acquired our scanner. These PolyJet printers were meant to create the final stand-ins before confirming and creating the final ceramic parts. A unique 3D printing technology, PolyJet printers work like 2D inkjet printers, except instead of dropping ink, they drop a photopolymer that is cured by UV light. Compared to a similar technology like Stereolithography Apparatus (SLA) printing, which cures the surface of a pool of resin one layer at a time, PolyJet allows for multi-color printing through resin mixing. At the start of the project, the plan was to make negatives off of these printed parts, then cast ceramic versions of each piece. However, as it yellows over time and *Shipwrecked Mother and Child* needs to remain on display with little to no maintenance, the finished pieces could not be resin. Therefore, we began exploring using ceramic printing to generate the final parts for the sculpture. This method works by adding a ceramic powder into the resin pool of a typical SLA printer, then printing the part as usual. The finished print is then fired in a kiln to remove the resin and harden the ceramic into a finished part. With a similar material composition and finish, these 3D-printed ceramic parts would blend seamlessly with the rest of the statue and not age over time. Unfortunately, various issues caused the ordered ceramic samples never to arrive, and the project ended before the technology could be further explored.

6. Conclusion and future work

This project developed a digital restoration workflow for recreating lost portions of sculptures using scan data, modeling software, and 3D printers. The case study was *Ship Wrecked Mother and Child* by Edward Augustus Brackett, which lost five critical pieces since being put into the Worcester Art Museum's archive nearly eighty years ago. We began by 3D scanning the statue using a Creaform Academia 50, an off-the-shelf scanner. From there, the data was imported into Blender, a free and open-source modeling program, where we manually recreated the missing portions of the statue through the transformation of existing portions of the statue. Finally, these recreated pieces were brought into existence through 3D printing, which, when attached to the sculpture, allowed us to see the statue as it may have been a century ago. The work done in this project proves that digital restoration workflows can aid conservators in restoring art to the original artist's vision.

Future work should focus on three primary areas: 1) new printing technologies, 2) augmented reality, and 3) automated modeling through sophisticated algorithms or artificial intelligence. For us, ceramic 3D printing was seen as the ideal solution for recreating the missing pieces, but conceivably, sophisticated five-axis machining techniques can be used to recreate an artist's touch out of marble. If cheaper solutions are required, a material investigation into different resins and how they age should be conducted to find which ones can remain on a museum floor for decades without discoloration.

The second area of focus should be augmented reality. We stopped experimenting with the technology after a frustrating experience attempting to view the generated models on the statue through a Microsoft HoloLens. Perhaps future generations of the technology with more sophisticated spatial tracking and more intuitive user controls might be better tools for rapid iteration than 3D printing.

Finally, we believe that either an algorithmic or AI approach to part generation should be explored to remove the reliance on humans who might subconsciously shape the statue how they see fit. In the past, art conservators enforced their vision of what art should look like on the pieces they worked to preserve. Today, art conservators study and analyze pieces to channel and restore what the original artists created. In the future, digital restoration will bring the original artists' visions back to life so that they may shape the statue as they envisioned.

References

[1] Sheehan, N., 2019, “After 80 Years, Sculpture Emerges from Worcester Art Museum Basement”, Worcester Telegram [Online]. Available: https://www.telegram.com/news/20191109/after-80-years-sculpture-emerges-from-worcester-art-museum-basement.

[2] Rosenthal, A., Insall, D., Podany, J., Larson, J., Bromelle, N., and Zuccari, F., 2017, "Art conservation and restoration", Britannica [Online]. Available: https://www.britannica.com/art/art-conservation-and-restoration.

[3] 2014, “TULLIO LOMBARDO’S ADAM A Masterpiece Restored”, Metmuseum.org [Online]. Available: https://www.metmuseum.org/press/exhibitions/2014/tullio-lombardos-adam.

[4] Emmrich, J., 2021, “What is Structured Light Scanning?”, Medit Blog [Online]. Available: http://blog.medit.com/medit/what-is-structured-light-scanning.

[5] “Educational Software Packages for Teachers: Creaform ACADEMIA”, Creaform [Online]. Available: http://www.creaform3d.com/en/education/teaching.

[6] Arrighi, P., 2020, “3D Scanning Technologies - What Is 3D Scanning and How Does It Work?”, Aniwaa [Online]. Available: https://www.aniwaa.com/guide/3d-scanners/3d-scanning-technologies-and-the-3d-scanning-process/.

[7] Cdedmonds, and DCtheGeek, 2019, “Locatable Camera - Mixed Reality”, Microsoft Docs [Online]. Available: http://docs.microsoft.com/en-us/windows/mixed-reality/develop/platform-capabilities-and-apis/locatable-camera.

[8] Cummins, K., 2019, “The Rise of Additive Manufacturing”, The Engineer [Online]. Available: http://www.theengineer.co.uk/the-rise-of-additive-manufacturing/.

[9] Alsop, T., 2020, "Top 3D printing materials 2020", Statista [Online]. Available: https://www.statista.com/statistics/560323/worldwide-survey-3d-printing-top-technologies/.

[10] Bild, S., 2018, “SLA vs. PolyJet [Technology Review]”, TriMech Blog [Online]. Available: https://blog.trimech.com/sla-vs-polyjet-technology-review.

[11] Artal-Isbrand, P., 2020, “Nathan Kaplan scanning *Shipwrecked Mother and Child*”

[12] 2021, "HoloLens 2 — Overview, Features, and Specs | Microsoft HoloLens", Microsoft.com [Online]. Available: https://www.microsoft.com/en-us/hololens/hardware.

[13] 2021, "Specification of Zortrax M Series 3D Printers | Zortrax Support Center", Zortrax [Online]. Available: https://support.zortrax.com/m-series-specification/

[14] Hiscox, C., 2020, “Fringe Reconstruction”, Worcester Polytechnic Institute, unpublished report.

APPENDIX A. Software Information including hardware requirements

APPENDIX A.1. VXElements

Version Number: 8.1.1

Software included with the scanner.

Recommended Hardware Requirements:

Processor	Intel Core i7 (6+ cores) – 2.6 GHz or more
Operating System	Windows 10 (64 bits)
Graphics Card	Nvidia Quadro T1000, OpenGL 4.5
Memory	32 GB
Hard Drive	SSD with 200 GB of free space
Display	1920 x 1080 (Full HD) display
Other	1x USB 3.0 port, Excel 2016

APPENDIX A.2. Blender

Version Number: 2.91

Download Link: https://www.blender.org/download/

Recommended Hardware Requirements:

Processor	64-bit quad-core CPU
Operating System	Windows 10, macOS 10.13 - 11.0, Linux
Graphics Card	Graphics card with 4 GB of RAM
Memory	16 GB
Hard Drive	Not Specified
Display	1920 x 1080 (Full HD) display
Other	Three button mouse or pen + tablet

APPENDIX A.3. Z-SUITE

Version Number: 2.16.1.0

Download Link: https://support.zortrax.com/downloads/software/

Minimum Hardware Requirements:

Processor	Intel i3 or equivalent AMD (3.0 GHz +)
Operating System	Windows 7 (64-bit) / macOS up to 10.14
Graphics Card	Nvidia GT 730 / AMD R7
Memory	8 GB
Hard Drive	Not Specified
Display	Not Specified
Other	Not Specified

APPENDIX B. Hardware Information

APPENDIX B.1. Creaform Academia 50

Figure B.0.1: Creaform Academia 50 3D scanner. Promotional material. [5]

Technical Specifications as published by Creaform [5]:

Part Size (Range)	0.3 – 3 m
Accuracy	Up to 0.1 mm
Volumetric Accuracy	0.300 mm/m
Measurement Resolution	0.250 mm
Mesh Resolution	0.500 mm
Scanning Area	380 x 380 mm
Stand-off Distance	400 mm
Depth of Field	250 mm
Light Source	White light (LED)
Laser Class	24 bits
Texture Resolution	50 to 150 DPI
Positioning Methods	Geometry and/or targets and/or texture
Measurement Rate	550,000 measurements/s
Weight	0.95 kg
Dimensions (LxWxH)	150 x 171 x 251 mm
Operating Temperature Range	5 – 40°C
Operating humidity Range	10-90%
Certifications	EC Compliance, compatible with rechargeable batteries, IP50, WEEE

APPENDIX B.2. Microsoft Hololens 2

Technical Specifications as published by Microsoft [12]:

Display: Optics	See-through holographic lenses (waveguides)
Display: Resolution	2k 3:2 light engines
Display: Holographic Densify	>2.5k radiants
Display: Eye-based rendering	Display optimization for 3D eye position
Sensors: Head Tracking	4 visible light cameras
Sensors: Eye Tracking	2 IR cameras
Sensors: Depth	1-MP time-of-flight depth sensor
Sensors: IMU	Accelerometer, gyroscope, magnetometer
Sensors: Camera	8-MP stills, 1080p30 video
Audio and Speech: Microphone Array	5 channels
Audio and Speech: Speakers	Built-in spatial sound
Human Understanding: Hand Tracking	Two-handed fully articulated model, direct manipulation
Human Understanding: Eye Tracking	Real-time tracking
Human Understanding: Voice	Command and control on-device; natural language with internet connectivity
Human Understanding: Windows Hello	Enterprise-grade security with iris recognition
Environment Understanding: 6DoF Tracking	World-scale positional tracking
Environment Understanding: Spatial Mapping	Real-time environment mesh
Environment Understanding: Mixed Reality Capture	Mixed hologram and physical environment photos and videos
Compute and Connectivity: SoC	Qualcomm Snapdragon 850 Compute Platform
Compute and Connectivity: HPU	Second-generation custom-built holographic processing unit
Compute and Connectivity: Memory	4-GB LPDDR4x system DRAM
Compute and Connectivity: Storage	64-GB UFS 2.1
Compute and Connectivity: Wi-Fi	Wi-Fi 5 (802.11ac 2x2)
Compute and Connectivity: Bluetooth	Bluetooth 5
Compute and Connectivity: USB	USB Type-C
Fit: Sizing	Single size, fits over glasses
Fit: Weight	566g
Software: Operating System	Windows Holographic Operating System
Software: Included Software	Microsoft Edge, Dynamics 365 Remote Assist, Dynamics 365 Guides, 3D Viewer
Power: Battery Life	2-3 hours of active use
Power: Charging	USB-PD for fast charging
Power: Cooling	Passive (no fans)
Power: Batteries	Contains lithium batteries (use caution)

APPENDIX B.3. Zortrax M200

Technical Specifications as published by Zortrax [13]:

Physical Dimensions w/ Spool	350 x 440 x 505 mm
Technology	LPD (Layer Plastic Deposition) / FDM (Fused Deposition Modeling)
Layer Resolution	90 microns
Minimal Wall Thickness	450 microns
Dimensional Accuracy	+/- 0.2%
Angle Accuracy	+/- 0.2%
Platform Levelling	Automatic measurement of platform points' height
Build Volume	200 x 200 x 180 mm
Material Container	Spool
Material Diameter	1.75 mm
Nozzle Diameter	0.4 mm
Support	Mechanically removed – printed with the same material as the model
Hotend	Single
Connectivity	SD card
Available Materials	Full offer is available at: https://zortrax.com/materials/zortrax-m-series/
External Materials	Applicable
Maximum Printing Temperature	290°C
Build Platform	Heated
Maximum Platform Temperature	105°C
Ambient Operating Temperature	20 - 30°C
Storage Temperature	0 - 35°C
AC Input	110V ~ 4A 50/60Hz 240V ~ 1.7A 50/60Hz
Maximum Power Consumption	200 W
Software Bundle	Z-SUITE
Supported File Types	.stl, .obj, .dxf. 3mf
Supported Operating Systems	Mac OS up to Mojave / Windows 7 or newer

APPENDIX C. Flowchart of part generation process