

Rough Carving of 3D Models with Spatial Augmented Reality

Ammar Hattab
Brown University
Providence, RI
ammar_hattab@brown.edu

Gabriel Taubin
Brown University
Providence, RI
gabriel_taubin@brown.edu




Figure 1: We assist users in roughly cutting a 3D model. (a) Given a target 3D model, we generate an optimized sequence of cutting steps and use a projector to project them to a block of material to guide the user in performing them. (b) Rough carving result. (c) Fine carving result.

ABSTRACT

Carving is a subtractive process where we get the shape by removing materials. While most people can get roughly the right intended shape, it is usually challenging not to over-cut the model. We propose a method that helps an unskilled user to carve a rough physical replica of a 3D model using the minimum number of cuts while only using manual cutting tools. The method starts by analyzing the input 3D model and generates the minimum set of cutting steps that remove most of the material. Then using a projector, we project the instructions sequentially onto a block of material to guide the user in performing them. We use the projector-camera setup to 3D scan the object after cutting and automatically detect the changes to reflect them on the digital model. We demonstrate a complete system to support this operation and show several examples of manually carved 3D models while using the system.

CCS CONCEPTS

- Human-centered computing → Human computer interaction (HCI).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SCF '19, June 16–18, 2019, Pittsburgh, PA, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-6795-0/19/06...\$15.00
<https://doi.org/10.1145/3328939.3328998>

KEYWORDS

Personal Fabrication, Assisted Carving, Rough 3D Carving, Spatial Augmented Reality

ACM Reference Format:

Ammar Hattab and Gabriel Taubin. 2019. Rough Carving of 3D Models with Spatial Augmented Reality. In *SCF '19: Symposium on Computational Fabrication (SCF '19), June 16–18, 2019, Pittsburgh, PA, USA*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3328939.3328998>

1 INTRODUCTION

A large number of people practice the manual carving craft. But for beginners, there are two challenges; the first is to get the right proportions of the carving in 3D, the second is not to over-cut the model. Most people don't have the skills to carve the intended shape from a mental image without over-cutting parts of it. In this paper, we introduce a method to help novice users carve 3D models by using a projector to project guidance colors onto the material.

Our approach is similar to the previous methods like the "Shape Shift" [Skeels and Rehg 2007] or the "Sculpting by Numbers" [Rivers et al. 2012], but in this paper, we focused on the rough carving stage. Carving is usually done in stages. First, the rough stage where the user wants to remove most of the material quickly, without over-cutting the model. Then, the fine carving stage, where the user moves to a different tool to remove smaller and smaller amounts of material, while adding in the fine details of the model. For the rough stage we introduced a novel method to provide the user with cutting instructions to carve a rough physical replica of the model, while for the fine stage we followed the "Sculpting by Numbers" [Rivers et al. 2012] approach to project colors onto the evolving block of



Figure 2: Two stages of carving: rough cutting and fine carving.

material as visual feedback to indicate to the user how far from the intended shape the current shape is, see figure 2.

The rough stage is usually done by using a saw to quickly cut large chunks of the material. This helps the user to reduce the dust that is generated by the milling or sanding operations. For example, the 3d model in the left image of figure 3 takes only 17.5% of the total volume of the shown block, the rest will become dust with the milling operation. While a single cut could remove 40% of the total volume of the block as in the right image of figure 3.

Some 3D models are considered rough in general and need minimal fine carvings. In other cases, for example during prototyping, all the designer wants is the rough shape of the model to feel it in his hands, or to see how it fits in the environment and make adjustments. Fabricating a low-fidelity approximate model has many applications, for this reason a large number of the previous papers focused on that. Automatic fabrication machines like 3D printers or CNC milling machines are too slow for this purpose.

There are two ways people currently perform the rough carving stage. The first method is by marking the model silhouettes (from the front or side view) on the initial block, then cutting these silhouettes using some power saw machine, for example using a band-saw. The second method is to use a normal saw or a chainsaw

to cut large chunks of the material quickly. The problem with the first method is that it only works well for models that align well with the three views (top, front, and side). Also, once the silhouettes are cut on one side, it is harder to draw the silhouettes on the other sides. For more complex models, it is hard for most people to figure out the cutting steps in 3D.

In this paper, we will focus on the second method of using a normal saw and let the user cut a large chunk of the material quickly. These cuts are called Blocking cuts. By examining carving guides created for beginners, we notice that there are two types of blocking cuts; straight cuts that go all the way through the material, and stop cuts, where the cut is performed by two straight cutting steps, see figure 4.

Motivated by these carving guides for beginners, we present a method that can automatically analyze a 3D model and generate a series of optimized cutting steps. At each step, the method computes the cut that removes the largest amount of material. Then using a projector to project these cutting steps as a guide, the user will perform the cut by following the guide showing which parts of the material should be removed until it reaches the target shape.

The system has two main components; a method to generate manual cutting steps, and a projector-camera calibrated setup to project those steps, and to perform 3D scanning.

2 RELATED WORK

2.1 Personal Interactive Carving

The approach of using a projector to help novice users in shaping their 3D models was initiated by "ShapeShift" [Skeels and Rehg 2007], where they used depth guidance to help users in sculpting clay. "Sculpting by Numbers" [Rivers et al. 2012] further developed this technique and used edge guidance in addition to the depth guidance. In [Marner and Thomas 2010], they used this technique to guide users in carving foam and enhanced the setup by adding real-time tracking sensors to both the cutting tool and the material, which allowed them to capture user's carving in real-time rather than doing that through the 3D scanning as in [Rivers et al. 2012; Skeels and Rehg 2007] which might be slow. The main difference between these papers and our work is that they provide general




Figure 3: Left: The bunny 3d model takes only 17.5% of the total volume of the shown block, the rest will become dust. Right: A single cut removes 40% of the total volume of the block of material.




Figure 4: Chainsaw Carving guide for beginners. Two types of cuts shown; straight cuts like cut number 1, and stop cuts like cuts number 2 and 3. Image source: ChainsawCarve.ca.

guidance to the user by projecting the depth or the target object on the material and letting the user figure out how to reach that goal, which might work better for more experienced users. Our proposed method leads the users one step at a time through the process and tells them where to cut. Being more prescriptive than the other approaches, our method is more suitable for novice users with little experience in carving. Since this paper is focused on the rough carving stage, it is considered to complement the previous works which are more suitable for the fine details stage. Once the user gets his carving close enough to the model; then, the other approaches of edge and depth guidance might be more suitable.

Another important difference is that we used 3D augmented reality approach because we need to project 3D instructions, while previous works like [Rivers et al. 2012] for example, only worked in the image domain, where the projected depth and edge guidance are generated as images from the depth map of the scene captured by the 3D scanner. The 3D approach adopted by this paper gives more projection abilities like projecting 3D instructions or rendering the silhouettes from three different views at the same time as in Figure 12, instead of only using the projector direction to show the depth guidance.

Besides using a projector, in the "FreeD" project [Zoran et al. 2013] they proposed a hybrid carving technique, where they developed a hand-held digital milling device that helps the user in carving the 3D model. Our work is more focused on cutting rather

than milling, which we believe is the best approach to removing large chunks of materials quickly without generating a lot of dust from the milling operation.

Another previous work of the authors that used a similar setup but focused on carving CAD models of the CSG models for applications in mechanical design [Hattab and Taubin 2019]. While in this paper we focus general 3d models of the polygonal mesh format.

2.2 Approximate personal fabrication

There are many previous studies that tried to help the user in quickly fabricating a low-fidelity approximate physical replica of a 3d model. These methods tried to avoid the use of 3d digital fabrication machine which might be slow, hard to use or not available to the user. Instead, they relied on the user to perform part of the fabrication like assembling blocks or using a faster fabrication machine like a laser-cutter.

Using a laser-cutter allows a faster fabrication, for example in [Beyer et al. 2015] they proposed a method to decompose the model into planar faces that could be cut by laser and assembled together. In [Song et al. 2016] they proposed a method to decompose the model into parts that can be fabricated using a laser cutter and parts that should be 3d printed to preserve the high details. In [Holroyd et al. 2011] they proposed a method to slice the 3d model, and then using a laser cutter to cut the slices and then gluing the slices together.

Some other papers relied on the user to do some of the fabrication operation, for example in [Mueller et al. 2014] they proposed a method to generate fabrication plans for the user to assemble standard building bricks to fabricate a 3d model. In [Liu et al. 2017] they converted the 3d model to an approximate wire structure that could be fabricated using a wire-bending machine. In [Massarwi et al. 2007; Mitani and Suzuki 2004] they proposed a method to decompose the surface of the model to be flattened and printed on paper, then cut and assembled to get the 3d model.

2.3 Spatially Augmented Reality

In this paper, we used a projector to project cutting instructions onto a block of material. This creates an augmented reality on top of a real physical object. An approach that is called spatially augmented reality. It could be achieved using the 3D rendering of graphics on top of the physical object as in [Raskar et al. 2001] for example. Or using the image domain for example by generating a depth map and a depth guidance as in [Rivers et al. 2012]. In this paper, we followed the 3D approach, where we used the projector intrinsic and extrinsic parameters for the rendering cameras, such that the rendering is performed from the same position and orientation of the projector with respect to the environment. The projector itself is considered as an inverse camera. Spatial augmented reality has many applications. Some papers used this approach for guiding the user to perform certain tasks, such as sculpting [Rivers et al. 2012; Skeels and Rehg 2007], painting [Flagg and Rehg 2006], or placing objects in a certain position. The benefit here is that the user does not need to wear any special devices to see the virtual reality.

The traditional augmented reality approach (head mounted display), was also used to guide users in performing tasks. For example, more recently it was used to guide novice users to perform 3D wire

sculpting [Yue et al. 2017]. And to help the user in interactive fabrication together with a 3D printer [Peng et al. 2018].

2.4 Rough Carving Algorithms

To the best of our knowledge, there are no papers on the rough carving of 3D models using straight cutting tools like the saw. But there are many papers with algorithms for the roughing stage of CNC milling machines, and a few papers with algorithms for the hot-wire CNC cutting machines. Milling is a different operation than cutting, while hot-wire cutting is closer to our work. The difference is that the user can change the direction of the cut in the middle of the cut. While in our work we assume that the user will only follow straight cuts. As an example for hot-wire cutting, an algorithm was developed for roughly cutting models using a hot-wire 8-axis machine [Zhu et al. 2006]. In that paper, the 3D model was simplified first and then the machine was guided to cut each of the simplified model faces, one at a time, starting from the top. There are two advantages of our method over their algorithm; first, our method uses optimization to compute the cuts that remove most of the material first. Secondly, our algorithm could perform optimized concave cuts using stop-cuts, while their method faces a problem with concave regions.

3 METHOD

3.1 System Design

In terms of hardware, the system is composed of a camera and a projector calibrated and fixed with regard to each other. The projector is used to project the cutting instructions on a block of material where it will illuminate the regions that should be removed by the user. The camera is used to capture the 3D shape of the object together with the projector. Similar to the work of [Rivers et al. 2012], we also used a stage that could be rotated accurately by 90 degrees. This allows the user to use four orientations. We mark these orientations by colors on the stage to be referenced later in the carving process. For example, the method might ask the user to rotate the stage to a certain orientation using the marked color. If the user moves the model to one of these orientations, that should be indicated in the software as well. The projector and camera are fixed at a certain angle with the stage to allow projecting the instructions to one of the corners of the block of material covering three of its faces at one time, see figure 5.

In terms of software, the system is composed mainly of three components:

3.1.1 Cutting Steps Generator. Starting from the input 3D model, we need to generate the cutting steps to be projected to guide the user. Section 3.2 is dedicated to this part.

3.1.2 3D Scanning. We need to capture the shape of the current material for two purposes. First, initially we need to find the correct alignment in order to project the cutting instructions on the right position on the physical model. Secondly, to detect physical changes and to reflect them to the digital 3D model. There are many 3D scanning techniques, but since we are using a projector to project the instructions, we can use it for 3D scanning as well. So, we decided to use structured-light gray-code 3D scanning, and we used the method of [Moreno and Taubin 2012] to calibrate the




Figure 5: A projector-camera setup: (a) Block of material. (b) The carving stage that could be rotated by 90 degrees. (c) Cutting tools. (d) Camera. (e) Projector.

projector-camera setup. The result of 3D scanning is a 3D point-cloud sampling of the object.

Notice in Figure 5 that we have three coordinate systems, one for the camera, one for the projector and one for the 3D model. We choose to use the projector coordinate system as the main coordinate system and to convert all points and transformations to this coordinate system. Because the projector is the common thing between the two operations; 3D scanning with the camera, and 3D rendering with the 3D model. The point cloud resulting from scanning is converted to projector coordinate system by a simple transformation using the rotation and translation matrices (These matrices are defined between the camera and the projector and we obtain them during the calibration process):

$$P_p = RP_c + T$$

While the 3D model is transformed to the projector coordinate system by the 3D alignment step. The alignment is initially performed by picking 3 planes from the model block, and the corresponding 3 planes from the 3D scan of the physical block of material. Then we use the method of [Hattab and Taubin 2018] to perform the alignment. We do that for each of the four basic orientations, then the system becomes ready. Where each 3D cutting instruction is transformed to the projector coordinate system and aligned with the physical model as captured by the scanner.

3.1.3 3D Annotations. Now, after alignment of the scanned point cloud and the current instruction model. The model is ready for projecting the next instruction. So, we use WebGL to render the next instruction while using the projector intrinsic parameters for the rendering camera [Simek 2013], and we show the rendering on the projector screen. At the same time, we show a rendering of the whole model with the current instruction highlighted on the computer screen to allow the user to move around it and explore it.




Figure 6: Our method starts by computing the silhouettes of the 3D model from a uniformly distributed set of directions. Then it computes the best cutting steps from the silhouettes.

3.2 Generating Rough Cutting Steps

The input to this step is a 3D model and the dimensions of the block of material, and the goal is to generate a series of optimized cutting steps to roughly cut that model out of the block of material. Each cutting instruction should remove a part of that block. And each cutting instruction should be either a straight planar cut or a stop-cut defined by two planes as explained in the introduction.

The straight planar cuts can only produce convex shapes, while the stop-cuts could remove material from the concave regions of the model. To compute the possible straight planar cuts we could first compute the 3D convex hull of the model and use its faces to define the planar cuts. But to compute the stop cuts is more challenging. The main idea of this method is to find the possible concave regions of the model by looking into it from different directions and use these concave regions to define the stop-cuts.

3.2.1 Method Steps.

- First, scale the 3D model to fit inside the box with the user-specified dimensions.
- Choose a uniformly distributed set of directions to look into the model from them and generate possible cuts. see figure 6.
- For each direction we compute the possible cuts using the following steps (see figure 7):
 - Project the mesh vertices onto a plane aligned with that direction.
 - Compute the silhouette of the model from the projected vertices and edges.
 - Find the convex hull of the silhouette.
 - Offset the silhouette and the convex hull by a specific amount to help the user in avoiding over-cutting the model. see figure 11 for some examples of over-cuts as captured by the 3D scanner. These over-cuts happened when using the system without the offset. We noticed that offsetting the instructions before projecting them helped in reducing them.

- Decompose the silhouette into concave regions, by finding the points where the silhouette meet with the convex hull.
- For each concave region, define a possible stop-cut from that region. First, we find the point on the silhouette of that region with the furthest distance from the convex hull, we call it the maximum point. This point splits silhouette of that region into two parts, for each of the two parts we find another point with the furthest distance from the line defined by the maximum point and one end of the silhouette of the region. These three points define a wedge shape that could be extruded to define the stop-cut.
- Generate the possible planar cuts using the convex hull of the silhouette. Each line of the convex hull defines a plane together with the specified direction. For each of these planar cuts, we compute the volume of the removed piece by that cut. see figure 8.
- From all the planar cuts from all directions, we select the cut that produces the largest removed volume (removes the largest amount of the material). When the removed volume by the best planar cut is smaller than a threshold, we move to the stop-cuts. The user can also manually switch to stop-cuts.
- Generate the possible stop cuts (two-planes cuts) using the concave regions. For each concave region, we first extrude the wedge of the region in the specified direction. Then we intersect the extrusion with the current shape of the material using constructive solid geometry (CSG). The intersection in red represents the material to be removed. see figure 9.
- From all the stop-cuts from all directions, we select the cut that produces the largest removed volume (removes the largest amount of the material). When the removed volume by the best stop-cut is smaller than a threshold, we stop the algorithm.

The system also supports two modes; the normal mode where the user follows suggested instruction orientation, orient the part first then cut. Or, we could let the user select an orientation, then let the system generates all possible cuts that could be done from that orientation.

4 EXPERIMENTS AND DISCUSSION

4.1 Carving Process

Initially, the user mounts the block of material on the stage. Then the user inserts the dimensions of the block into the software. After which, the system 3D scans the block from the four orientations and align the 3D scan with the 3D model. Now the system is ready for the carving.

The system then analyzes the 3D model and generates a list of 3D instructions, for each cutting instruction we need to perform the following steps (see figure 10):

- 4.1.1 *Finding the best orientation.* We want to find the best orientation to show the instruction that gives the best projector visibility, such that we can show the instruction to the user. We simply want to choose from the four initial orientations. We do that by computing the angle between the viewing direction vector of each of the orientations and the normal vector of the cut (or the average




Figure 7: For each chosen direction: (a) Vertices projection. (b) Computed Silhouette. (c) Convex hull of the silhouette. (d) Extreme points that define the stop-cuts. (e) Offset the silhouette and the convex hull. (f) Concave regions (stop-cut). (g) Blue regions show material that can be removed by planar straight cuts. While the red regions show material that can be removed by stop-cuts.



Figure 8: Planar cuts are generated from the convex hull of the silhouette. (a) Silhouette in a specific direction. (b) The red region represents the material to be removed by a planar cut. (c) The model after removing the material.

of the two normal vectors in the case of stop-cuts) And choosing the orientation with the minimum angle. After choosing the best orientation we show it to the user and we ask him or her to orient the part in the specified orientation.

4.1.2 3D Projection. We then use the projector to project the instruction, the user can also choose to highlight the edges of the instruction to make it more clear. The stop-cut instructions are split from the middle and the two parts are shown in red and blue colors, see figure 15.

4.1.3 Marking. As an optional step, before performing the cut, the user could use a pencil or other tool to mark the cut lines. A process that is usually done in general before any manual cutting.

4.1.4 Cutting. The user performs the actual cutting. The projection can be stopped if it is easier for the user. The user can also take off the stage to perform the cut from other orientations, and put it back after the cut.

4.1.5 3D Scanning. We start 3D scanning on the object, to capture the physical cuts that we made. After 3D scanning, the system automatically compares the resulting point cloud of the object with the previous point cloud (before doing the cut) in the same orientation. And finds the difference which represents the cut. To compute the difference, first we build a KD-Tree of the points in the previous point cloud. Then for each point in the current point cloud, we use the KD-Tree to find the nearest point in the previous point cloud and compute the distance between these two points. If the distance is larger than a threshold, that means that this point belongs to the changed region. Now we have the changed points that represent the cut, if the cut is of the planar type, we simply find the best fitting plane to the changed points using least squares, and refine the fitting with few iterations. If the cut is of the stop-cut type, we pick a random point of the changed points and fit a plane to its neighbors (nearest points around it) and we refine it until it settles on one of the two cuts planes. We fit another plane to the remaining points among the changed points. We use these detected




Figure 9: Stop-cuts are generated from the concave regions.
(a) Silhouette in a specific direction, and one of the concave regions shown in red. **(b)** We extrude that concave region in the specified direction. **(c)** We intersect the extrusion with the current model (red color). **(d)** The model after the cut.

planes to perform the actual cut and update the model. Then the user could start the next cut operation. And repeat the previous steps.

4.2 Rough Model Examples

In Figure 14, we show several 3D models and the rough carving result 3D model after some number of cutting steps. Refer to Table 1 for some statistics on these models.

In Figure 15, we show an example of the cutting steps for the bunny 3D model.

4.3 Carving Experiments

In Figure 13, we show some examples of rough carving results. To perform these experiments, we used a low-resolution Optoma projector (960 X 540 pixels) and a Basler camera (1928 X 1208).

4.4 Discussion

During the experiments, we measured the time taken by each of the carving steps and found that on average it takes the user 1.5 minutes to perform the carving step on a block of material with 4-1/2 X 5 X 3" dimensions. Usually, the first few steps takes more time due to larger surface area being cut. Also stop-cuts and in general steps with hard angles takes more time.

From this info, we estimated the time needed to roughly carve several 3D models shown in Figure 14. We also used a 3D printing software to estimate the time it needs to 3D print each of these models. See Table 1 for these statistics. In this table, we also show the number of steps performed, and the percentage of remaining




Figure 10: General process step. For a specific instruction, **(a)** First we find the best orientation, and we ask the user to orient the model to the specified orientation. **(b)** We project the instruction. **(c)** The user marks the places to be cut (optional). **(d)** The user performs the carving using hand tools. **(e)** We 3D scan the object, and compute the physical changes and apply them to the 3D model **(f)** Ready for the next instruction.



Figure 11: Examples of over-cuts as captured by the 3D scanner (left and middle). Offsetting the instructions helped in reducing them (right model).

volume to the total difference volume. This number gives us a hint on how much material is left to reach the target model and how rough is the carving.

The table shows that this rough carving method is on average 10 times faster than 3D printing, although it produces much lower fidelity model. The advantage of carving methods is that the fidelity of the output carving is proportional to the amount of time we spent working on them, while 3D printing time is the same whether the model is rough or one with many fine details.

	Steps	Left Volume	Carving Time	3D printing
1	15	10.25 %	22.5 mins	4 hrs 53 mins
2	22	11.54 %	33 mins	2 hrs 12 mins
3	26	13.97 %	39 mins	7 hrs 27 mins
4	13	13.42 %	19.5 mins	6 hrs 1 min
5	25	14.24 %	37.5 mins	6 hrs 27 mins
6	19	19.54 %	28.5 mins	6 hrs 11 mins
7	22	34.30 %	33 mins	5 hrs 2 mins
8	12	21.68 %	18 mins	4 hrs 20 mins
9	23	9.49 %	34.5 mins	4 hrs 47 mins
10	17	16.97 %	25.5 mins	5 hrs 7 min
11	15	12.54 %	22.5 mins	5 hrs 45 mins

Table 1: Statistics for models in figure 14

Another important factor is the size of the model. The larger the model the more important the rough cutting stage. Since we save more time and produce less dust by removing larger chunks of materials quickly compared to milling or carving operations. The advantage of our projector-camera setup is that it could be easily scaled up to larger objects by placing them at an appropriate distance to cover the large object without any modification to the method or the software.

After a while of trying the system, when the user is trained to follow the guide correctly, 3D scanning becomes less necessary for the different steps. Only projection is enough, which make the system even faster.

In this paper, we used a greedy approach to generate the cutting instructions for the remaining volume, where at each iteration we try to find the largest volume that could be cut by the user, this greedy approach doesn't guarantee that we get the best instructions overall, but only the best instruction at a given time.

For the complexity of the algorithm, we can separate it to the two steps, the first step of the algorithm requires projecting the polygon mesh (projecting the vertices) on different planes directions, the run-time of this step depends on the number of vertices n , the number of edges e , and the number of projection direction d . The run-time of this step is of order $O(nd \log e)$. In the second step of the algorithm we depend on the resulting number of edges r resulting from the first step $O(r^2)$ since for each iteration we need to check the volume resulting from extruding each edge of the edges and intersecting it with the current material and then select the cut/edge that removes the largest volume, and we have to run at most r iterations, usually much lower than that. But since each edge check would require a Boolean intersection operation, which is an expensive operation, the total run-time of the algorithm becomes slow. To make the algorithm runs faster we simplify the input mesh first (reduce the number of vertices and edges using standard simplification algorithm), this makes our algorithm runs in the order of seconds. This simplification gives the same carving output result since the algorithm is expected to generate rough cutting results, so the fine details of the model are not important at this stage. Table 2 shows the average computation time for one carving instruction step for different simplified versions of the same 3d model. Notice that we have do this computation before each caring step.

	Number of vertices	Number of faces	Computation time
1	20505	41006	289.762 seconds
2	6502	13000	36.463 seconds
3	2002	4000	6.970 seconds
4	394	784	1.818 seconds

Table 2: Average computation time for one carving instruction step. All the computations are done for different simplified versions of the same 3d model.

4.5 Limitations

First, there are some limitations in the geometry of the input 3d model, since we only used straight cuts. For example, our method would not support geometries with holes or with concave regions that are not accessible from the outside with straight cut by a saw. The user might reach these regions with a different carving tool. In the future, a possible extension of the method of this paper would generate instructions for more carving tools and not only the saw.

For the proposed algorithm, we used a greedy approach to generate cuts that remove the largest amount of material at each step. In general, greedy approaches might not give the optimal solution for the whole sequence of cuts, for example there might be a different sequence of cuts that removes larger amount of materials in fewer steps than our greedy approach.

Since we used a carving stage that could be rotated by 90 degrees, the user is limited in four positions only. From the carving experiments, we found the user can easily perform most cutting instructions from one of these positions. But, a few cuts were hard to perform, for example, there are cuts with difficult angles, that intersect with the carving base. However, the user can choose to skip these cuts.

In terms of the system design, since we used one projector and camera, the system is limited to one projection direction and the user has to rotate the stage to cut from the other sides. A possible extension of the system is to use a multi-projector approach like [Raskar et al. 2001] for example.

Another limitation to using the projector is the shadows created by the user hands and the carving tools, sometimes the user has to move his hands or the tools to see the projected guidance colors. Using a multi-projector approach can also solve this problem.

At last, the used 3D scanning approach is slow, but could be made faster to almost real-time by adding synchronization to the projector-camera system. Or by using real-time tracking sensors.

4.6 Other Applications

4.6.1 Projecting Silhouettes. The same system could be used for the first method mentioned in the introduction, where we project the silhouettes of the model from the three views; top, side, and front. Or to combine projecting silhouettes and suggested cutting instructions. see figure 12.

4.6.2 General Assistance in Measuring and Marking. The same system could be used to assist the user in generic wood-working, for measurements and marking instead of using rulers and other tools.



Figure 12: The same setup could be used for projecting Silhouettes.

5 CONCLUSION

We have demonstrated a system that helps unskilled users in roughly carving 3D models while avoiding over-cuts. We found that this system is much faster than the 3D printing, especially for large models. We presented a method to generate optimized cutting steps for an input 3D model to get a rough shape, by removing the most amount of material.

The presented system has the potential to be further developed to allow for a seamless synchronization between the digital and the physical worlds, especially at the prototyping stage.

6 ACKNOWLEDGMENTS

The work described herein was partially supported by a Brown Fellowship and by NSF grants IIS-1717355 and IIP-1500249.

REFERENCES

Dustin Beyer, Serafima Gurevich, Stefanie Mueller, Hsiang-Ting Chen, and Patrick Baudisch. 2015. Platener: Low-fidelity fabrication of 3D objects by substituting 3D print with laser-cut plates. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1799–1806.

Matthew Flagg and James M Rehg. 2006. Projector-guided painting. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*. ACM, 235–244.

Ammar Hattab and Gabriel Taubin. 2018. 3D Rigid Registration of Cad Point-Clouds. In *Computing Sciences and Engineering (ICCSE), 2018 International Conference on*. IEEE, 1–6.

Ammar Hattab and Gabriel Taubin. 2019. Interactive Fabrication of CSG Models with Assisted Carving. In *Proceedings of the Thirteenth International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, 677–682.

Michael Holroyd, Ilya Baran, Jason Lawrence, and Wojciech Matusik. 2011. Computing and fabricating multilayer models. In *ACM Transactions on Graphics (TOG)*, Vol. 30. ACM, 187.

Min Liu, Yunbo Zhang, Jing Bai, Yuanzhi Cao, Jeffrey M Alperovich, and Karthik Ramani. 2017. WireFab: Mix-Dimensional Modeling and Fabrication for 3D Mesh Models. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 965–976.

Michael R Marnet and Bruce H Thomas. 2010. Augmented foam sculpting for capturing 3D models. In *3D User Interfaces (3DUI), 2010 IEEE Symposium on*. IEEE, 63–70.

Fady Massarwi, Craig Gotsman, and Gershon Elber. 2007. Papercraft models using generalized cylinders. In *null*. IEEE, 148–157.

Jun Mitani and Hiromasa Suzuki. 2004. Making papercraft toys from meshes using strip-based approximate unfolding. *ACM transactions on graphics (TOG)* 23, 3 (2004), 259–263.

Daniel Moreno and Gabriel Taubin. 2012. Simple, accurate, and robust projector-camera calibration. In *3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT), 2012 Second International Conference on*. IEEE, 464–471.

Stefanie Mueller, Tobias Mohr, Kerstin Guenther, Johannes Frohnhofer, and Patrick Baudisch. 2014. faBrickation: fast 3D printing of functional objects by integrating construction kit building blocks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 3827–3834.

Huaishu Peng, Jimmy Briggs, Cheng-Yao Wang, Kevin Guo, Joseph Kider, Stefanie Mueller, Patrick Baudisch, and François Guimbretière. 2018. RoMA: Interactive Fabrication with Augmented Reality and a Robotic 3D Printer. (2018).

Ramesh Raskar, Greg Welch, Kok-Lim Low, and Deepak Bandyopadhyay. 2001. Shader lamps: Animating real objects with image-based illumination. In *Rendering Techniques 2001*. Springer, 89–102.

Alec Rivers, Andrew Adams, and Frédo Durand. 2012. Sculpting by numbers. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 157.

Kyle Simek. 2013. Calibrated Cameras in OpenGL. http://ksimek.github.io/2013/06/03/calibrated_cameras_in_opengl. Accessed: 2018-09-19.

Christopher Skeels and James M Rehg. 2007. ShapeShift: A projector-guided sculpture system. In *Proc. UIST*, Vol. 7.

Peng Song, Bailin Deng, Ziqi Wang, Zhichao Dong, Wei Li, Chi-Wing Fu, and Ligang Liu. 2016. CoffFab: coarse-to-fine fabrication of large 3D objects. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 45.

Ya-Ting Yue, Xiaolong Zhang, Yongliang Yang, Gang Ren, Yi-King Choi, and Wenping Wang. 2017. Wiredraw: 3d wire sculpturing guided with mixed reality. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 3693–3704.

Jiang Zhu, Tomohisa Tanaka, and Yoshio Saito. 2006. Rough Machining Process and Its Simulation for Robot Integrated Surface Sculpturing System. In *Proceedings of the Fourth IASTED International Conference on Circuits, Signals, and Systems (CSS2006)*, 20–22.

Amit Zoran, Roy Shilkrot, and Joseph Paradiso. 2013. Human-computer interaction for hybrid carving. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*. ACM, 433–440.



Figure 13: Some examples of rough physical carving experiments.



Figure 14: Examples of rough carving 3D models after a number of steps.




Figure 15: Example of the cutting steps for the bunny 3D model. The number at the top shows the percentage of material removed out of all the material that should be removed. The first row shows the model before the cut, the middle row shows the carving instruction, note that the blue-red steps are the stop-cuts. The bottom row shows the model after the cut.