


# Sculpting by Numbers

Alec Rivers  
MIT CSAIL

Andrew Adams  
MIT CSAIL

Frédo Durand  
MIT CSAIL



**Figure 1:** We assist users in creating physical objects that match digital 3D models. Given a target 3D model (a), we project different forms of guidance onto a work in progress (b) that indicate how it must be deformed to match the target model. As the user follows this guidance, the physical object’s shape approaches that of the target (c). With our system, unskilled users are able to produce accurate physical replicas of complex 3D models. Here, we recreate the Stanford bunny model (courtesy of the Stanford Computer Graphics Laboratory) out of polymer clay.

## Abstract

We propose a method that allows an unskilled user to create an accurate physical replica of a digital 3D model. We use a projector/camera pair to scan a work in progress, and project multiple forms of guidance onto the object itself that indicate which areas need more material, which need less, and where any ridges, valleys or depth discontinuities are. The user adjusts the model using the guidance and iterates, making the shape of the physical object approach that of the target 3D model over time. We show how this approach can be used to create a duplicate of an existing object, by scanning the object and using that scan as the target shape. The user is free to make the reproduction at a different scale and out of different materials: we turn a toy car into cake. We extend the technique to support replicating a sequence of models to create stop-motion video. We demonstrate an end-to-end system in which real-world performance capture data is retargeted to claymation. Our approach allows users to easily and accurately create complex shapes, and naturally supports a large range of materials and model sizes.

**Keywords:** personal digital fabrication, spatial augmented reality

## 1 Introduction

Most people find it challenging to sculpt, carve or manually form a precise shape. We argue that this is usually not because they lack manual dexterity – the average person is able to perform very precise manipulations – but rather because they lack precise 3D information, and cannot figure out what needs to be done to modify a work in progress in order to reach a goal shape. An analogy can be made to the task of reproducing a 2D painting: when given outlines that need only be filled in, as in a child’s coloring book or a paint-by-numbers kit, even an unskilled user can accurately reproduce a complex painting; the challenge lies not in placing paint on the canvas but in knowing where to place it. Motivated by this observation, we present Sculpting by Numbers, a method to provide analogous guidance for the creation of 3D objects, which assists a user in making an object that precisely matches the shape of a target 3D model.

We employ a spatially-augmented reality approach (see e.g. Raskar et al. [1998] or Bimber and Raskar [2005] for an overview of spatially-augmented reality), in which visual feedback illustrates the discrepancy between a work in progress and a target 3D shape. This approach was first proposed by Skeels and Rehg [2007]. In this approach, a projector-camera pair is used to scan the object being created using structured light. The scanned shape is compared with the target 3D model, and the projector then annotates the object with colors that indicate how the object ought to be changed to match the target. The user follows this guidance to adjust the object and rescans when necessary, bringing the object closer to the target shape over time.

Our proposed method provides guidance that illustrates depth disparities between the current work and the target, similar to the approach of Skeels and Rehg [2007], as well as an additional form of guidance, which we call edge guidance, which aids in reproducing high-frequency surface details. We demonstrate an application to reproducing an existing physical object at a different scale or using different materials. We further show how a scene can be re-

peatedly deformed to match a sequence of target 3D models for the purpose of stop-motion animation. Starting from a depth video recorded with a Kinect, we produce claymation with physically correct dynamics. Finally, we present results of user testing with novice sculptors.

## 2 Related work

The advent of digital fabrication technologies, such as 3D printers, laser cutters, and CNC machines, has made it possible to automatically produce 2D and 3D objects. The main tools used in 3D digital fabrication are 3D printers and 5-axis CNC machines. Devices of both kinds are typically expensive, though recent efforts in the DIY community have made low-priced entry-level tools available [Hokanson and Reilly ; Kelly ; MakerBot Industries ; Drumm 2011; Sells et al. 2009]. 3D printing and CNC methods are highly accurate, but suffer from fabrication artifacts, and are limited in the size and materials of the object they can produce. Multi-material fabrication with these devices is particularly challenging.

Recent work in computer graphics has therefore focused on software techniques that enable alternative methods for fabrication, and our technique falls into this category. Two recent works propose methods to generate models of sliding planar slices that approximate target 3D models [Hildebrand et al. 2012; McCrae et al. 2011]. Lau et al. [2011] propose a method to convert 3D furniture models into fabricatable planar parts. Others incorporate fabrication considerations directly into the modeling process, so that fabrication is straightforward by design: Kilian et al. [2008] propose a method for designing shapes that can be produced by folding a single sheet of material, while Mori and Igarashi [2007] propose an interface for designing stuffed animals.

Our approach aims to address the limitations of both manual crafting and current digital fabrication technologies by combining elements of each in a hybrid approach. Hybrid human-computer interfaces have been used before, especially in medical applications, for example to combine a human surgeon’s ability to plan motions and react to changing conditions with a robotic arm’s precision [Kragic et al. 2005; Mako Surgical ]. Interestingly, in these approaches, it is typically the human that does the planning and the machine that does the execution, while in our approach it is the reverse. Rivers et al. [2012] propose a hybrid approach to digital fabrication of 2D shapes in which a user coarsely positions a tool that can then automatically adjust its own position to keep it exactly on a digital plan; our approach shares the philosophical goal of enabling digitally-guided manual fabrication, but applies to 3D shapes, and works by just guiding a user without automatic error correction.

Hybrid manual-automatic methods have also been used for 2D design. Painters have used photographs or a camera obscura as references [Coke 1964; Pennsylvania Academy of the Fine Arts et al. 1994]. Flagg and Reh [Flagg and Reh 2006] presented Projector-Guided Painting, a system for digitally guiding an artist engaged in reproducing a 2D painting. They support the artist in building a painting out of layers, including showing how to orient the brush at certain stages, and explicitly help with the task of mixing paint to achieve a desired color. Dixon et al. [2010] proposed a hybrid method for the task of sketching a human face from a reference photograph in which the user is guided by automatic feedback. The guidance included high-level sketching principles for each part of the face as well as specific feedback on how to adjust the freehand drawings to more closely match the shape of the reference image. ShadowDraw [Lee et al. 2011] proposed a more open-ended system in which a user is allowed to draw freely while a “shadow image” of suggestive contours is updated and displayed behind the sketch to illustrate possible sketch completions.

While we are interested in the challenge of creating physical sculptures from a virtual model, the complementary challenge of creating virtual models using physical objects as an input method has also been explored. Sheng et al. [2006] presented a system in which a user whose fingers were being tracked could manipulate deformable physical props to perform sculpting operations on a virtual 3D model, with the goal of providing a more natural, hands-on modeling experience.

Our approach to guided sculpting involves scanning the work in progress and projecting guidance that annotates the physical object. 3D scanning has been an active area of research for many years, with a variety of optical methods being proposed; see Chen et al. [2000] for a survey. One of the oldest and most robust methods involves projecting a sequence of Gray codes and imaging the result to retrieve depth information [Inokuchi et al. 1984; Caspi et al. 1998; Rusinkiewicz et al. 2002]. This method is simple to implement, robust, and able to capture a full depth field quickly.

## 3 System design

The core operation in guided sculpting is creating a physical object that matches a target 3D model. This is done by projecting different forms of guidance onto the object that illustrate how it must be changed to match a target shape. All adjustments to the model are done manually, which naturally supports a wide range of materials and shapes. The guidance makes it possible to achieve complex and precise shapes.


This approach requires two main technical components: a method for scanning a 3D shape at high spatial and depth resolution, and a projector registered to the same coordinate system to project guidance onto the material. We achieve both goals with the same hardware by using a projector-camera pair to perform structured light scanning with Gray codes (e.g. [Caspi et al. 1998]). While faster scanning methods exist, e.g. the Primesense/Microsoft Kinect, they typically have much lower spatial resolution, which is more of an issue to us than refresh rate. We require the user to explicitly request each new scan, which has two advantages: the guidance does not change unexpectedly, and the user’s hands are not scanned and confused with the target object.

### 3.1 Scanning

The input to the scanning process is a target 3D model, which may be a mesh or a point cloud. It is virtually positioned within the working volume, and a block of raw material (e.g. clay) is placed at roughly the same real position.

Whenever the user initiates a scan, we project a sequence of Gray codes onto the scene and photograph the scene illuminated by each code. These images give us a correspondence from camera coordinates to projector coordinates. We filter out invalid or implausible correspondences, and then invert this map to compute a correspondence from projector coordinates to camera coordinates. From this we triangulate to compute a depth map from the projector’s perspective. We also compute a depth map of the target model rendered from the same perspective. Finally, we compare the depth maps to compute the guidance that we project onto the object. We describe each of these operations in detail below.

**Obtaining pixel correspondences** Gray-code-based structured light scanning operates by projecting a sequence of black and white horizontal and vertical bars onto a scene and recording the resulting images. By examining the illumination of a single camera pixel across all projected patterns, one can determine the X and Y coordinate of the projector pixel that was imaged by that camera pixel.



**Figure 2: Sculpting guidance:** A user is given feedback during sculpting by two forms of guidance projected onto the model, which they can toggle between. Guidance is computed by comparing the target depth map (a) to the scanned depth map of the current shape (b). Depth guidance (c) illustrates the difference between the two maps and is helpful for matching the position of the surface in absolute 3D coordinates. Edge guidance (d) illustrates the second derivative of the target model (i.e. ridges and valleys), and is useful for recreating fine details once the model is close to the target shape. A smoothly-varying range of colors is used: in depth mode green indicates where the surface should be moved closer to the scanner and red farther, while in edge mode green indicates a positive second derivative (a ridge) and red negative (a valley). In either mode, blue indicates background, so any part of the material colored blue should be removed entirely.

While the basic approach is straightforward, a number of factors can create incorrect correspondences. First, some camera pixels may image parts of the scene that are not illuminated by the projector at all. These can be detected and excluded by ignoring pixels that do not vary substantially between images in which the projector displays entirely white and entirely black.

More difficult to detect are cases in which a camera pixel records specular highlights or reflections of light cast by the projector. These camera pixels will vary appropriately between white and black projector images, and a corresponding projector pixel can be computed for them; however, this invalid correspondence will corrupt the depth map at that projector pixel. We are able to reject the majority of these correspondences with a consistency check. Given the homographies for the camera and projector, we can determine the world-space ray that corresponds to each camera and projector pixel. These rays ought to intersect within the working volume in the case of a true correspondence. If the rays do not approximately intersect, as may be the case if the camera is seeing a reflection, we reject the correspondence. Put another way, the scan in X, the scan in Y, and the epipolar constraint combine to overconstrain the pixel correspondence, which lets us reject invalid pixels. This constraint also helps eliminate spurious correspondences detected due to noise.

A final source of inaccuracy can arise when narrow alternating black and white columns in a fine-scale Gray code pattern are blurred into an average gray due to the combination of camera and projector defocus. In our case, the working volume is often close to the camera and projector, so their depths of field are relatively shallow. Loss of the finest Gray code patterns results in losing the least-significant bits of the projector-camera correspondence.

When coarsely shaping an object this loss of resolution does not matter, but towards the end of the sculpting process it becomes important. We therefore allow the user to perform a slightly more time-consuming “precise” scan inspired by laser plane sweep scanning (e.g. [Curlless and Levoy 1995]). In a simple implementation, we could project a single white row or column that sweeps across the entire projector image, and for every camera pixel simply record for which code the illumination was highest. This method is robust to defocus, but slow. We instead project the Gray code patterns as usual to compute a correspondence that may be incorrect in the least significant bits. We then sweep one-pixel-wide vertical or horizontal lines that recur every 8 rows or columns across all 8 positions, and record the position of maximum illumination for every camera pixel to correct the three least-significant bits of the corre-

spondence. This precise scan takes about 7 seconds, and gives us a depth resolution of under 1mm at the center of the working volume.

**Computing depth maps** As we will be displaying guidance using the projector, we need to compute the depth map from the projector’s point of view rather than the camera’s. To achieve this, we must resample the correspondence into the projector’s point of view.

We do so by splatting: each valid camera pixel contributes to a  $7 \times 7$  kernel with exponential falloff about the corresponding projector pixel. The contributions are summed in homogeneous coordinates and then normalized to produce a camera coordinate for some subset of projector pixels. The kernel size chosen is enough to interpolate across small holes in the depth map, which arise from projector pixels not seen by any camera pixel.

Given these resampled correspondences and the homographies for both camera and projector, we can triangulate to compute depth values for each projector pixel. Meanwhile, the depth map for the target model is computed by rendering it from the projector’s point of view using the projector’s calibrated homography as the OpenGL projection matrix. These depth maps are then compared to compute the guidance that we project onto the material.

**Hardware setup** We use a 3M MP160 projector running at  $800 \times 600$  at 60Hz and a Nokia N9 camera-phone using the FCam API [Adams et al. 2010] generating raw data at  $1632 \times 1232$  at 13Hz (75ms exposure time). The FCam API makes it easier for us to synchronize the camera with the projector, and the raw stream prevents us from losing resolution to demosaicing. Gray codes do not rely on color or spatial coherence, so we can treat each pixel independently and ignore the Bayer mosaic. The separation between camera and projector is 5”, and they placed roughly 2’ from the center of the work area. We compute world-space homographies for the camera and projector by scanning a known planar calibration target placed at two known depths in the work volume.

### 3.2 Guidance

We provide two forms of guidance, which the user may toggle between depending on the context, and which are projected directly onto the object being created. Each form is useful for recreating a different aspect of the model’s appearance. The user can also view a rendering of the target 3D model on the screen of the computer running the system.

**Depth guidance** is used to achieve the correct surface position in absolute 3D space. In this mode, projector pixels are colored based on the difference between the scanned depth at that pixel and the target depth (Figure 2 (c)). Green indicates areas where the user must bring the surface closer to the projector and red indicates areas where the user must move the surface farther away. Pixels that do not intersect the target model are colored blue. The range of the scale is user-tunable. The user can also probe the exact depth error in millimeters at a particular point by mousing over that pixel on the laptop controlling the projector.

**Edge guidance** depends only on the target 3D model, and is used to help the user match the high-frequency surface details of the target 3D model. In this mode, projector pixels are colored according to the second derivative of the target depth map (i.e. ridges and valleys, as shown in Figure 2 (d)). This color follows the same scale as the depth feedback mode, and its range can be similarly tuned.

**Rotation.** The method described so far assists a user in achieving the correct shape from a single point of view. However, a complete model must be correct from all sides. The challenge is to rotate the target model and the partial sculpture by the same amount about the same axis of rotation so the user can continue without invalidating work already done. While automatic methods (e.g. ICP [Zhang 1994]) could be employed to deal with this, we handle this more simply by sculpting atop a Lego stage, which can be accurately rotated by 90 degrees around a known axis.

## 4 Applications

Using our system in practice involves an interplay of computer guidance and human intuition, in which both the precision of the computer and the intuition of the artist are required to create an object with the correct proportions and aesthetic qualities.


In this process, the computer’s ability to accurately detect absolute 3D shape in world coordinates means that the guidance can be relied upon to achieve the correct relative positions and sizes of features of the model, which can otherwise be extremely challenging for a novice sculptor. However, the process also relies on the user’s understanding of the shape. One simple reason is that the guidance must be interpreted and internalized, as it will both be blocked by the user’s hands and tools, and is slowly invalidated as the user works on the model between scans. As users become familiar with the system, they quickly become able to perceive how a shape must be deformed and judge when to initiate a rescan.

A less obvious but important way in which we benefit from having an artist (even if a novice one) “in the loop” is that the user can make unprompted changes to the model to achieve a desired aesthetic appearance. Thus, a user may interpret the projected guidance liberally, and perform freehand adjustments to recreate salient details. This can both speed the creation of major features and enable the reproduction of fine-detail surface features that are too small to be individually highlighted and reproduced using the system.

We describe how this process plays out for the basic case of making a shape that matches a target 3D model, and go on to explore additional applications.

### 4.1 Creating an object that matches a 3D model

We found that use of our system fell into a fairly consistent pattern. A typical modeling session will begin in depth guidance mode. Using this the user will establish the coarse shape and silhouette of the model. As the model approaches the target shape, the user will make successively smaller adjustments to the object, shortening the red-to-green guidance range as necessary to highlight the necessary



**Figure 3:** The head of Michelangelo’s David, in the original marble (photograph by Steve Hanna) (a) and in a 6” tall version made in modeling clay with our system (b) using data from the digital Michelangelo project [Levoy et al. 2000]. The clay model was created by novice sculptors.

adjustments. Eventually, the user will have a shape that matches the target 3D model in absolute coordinates to the limit of the user’s ability to model and the system’s ability to scan.

Even after the model has been made to match the target 3D shape as closely as possible in depth guidance mode, the surface may still lack high-frequency surface details, such as wrinkles or the rough surface of hair. These details are often essential to the aesthetic appearance of a model, but are difficult to replicate with simple depth feedback: they may be finer than the scan can resolve, or may be hidden by low-frequency depth discrepancies that, although larger, can be ignored. Therefore, users will typically finish a model by switching to edge guidance mode. In this mode, they can see the locations of ridges, valleys, and other high-frequency details projected onto the surface, and use those as a starting point for an artistic pass on the surface.

The user will usually want to shape an object from all sides. Rather than forming the model perfectly from one side and then rotating, it is generally easier to form the material to match the target model roughly from all sides first, and then make successively higher-detail passes from each direction. This minimizes the problem of alterations to one side affecting already-completed regions of another.

Results of novice sculptors using our system can be seen in Figures 1 and 3. We found that our system enabled unskilled sculptors to achieve far better results than they would otherwise be capable of, while giving them a sense of ownership of the resulting creations that would be absent with an automatically-produced object.

One unexpected possibility allowed by our approach is collaborative sculpting. Because the computer guidance provides a common goal to work towards, it is possible for multiple users to work on a sculpture at the same time. This is both time-saving and highly enjoyable.

### 4.2 Sculpting from a physical example

Our approach can also be used with a simple extension to allow the user to create an object not just from a digital 3D mesh, but also from an already-existing physical object. To do this, the user places an object in front of the projector-camera pair and scans it using the method above, rotating it and scanning it from all four directions. The recovered point clouds are saved. The model is then removed and the stored point clouds are used as the target shape for modeling as above.

Because we use the point cloud recovered by the scanning algo-





**Figure 4: Sculpting from a physical example:** *With a simple extension, we can use a scan of an existing physical object as a target for sculpting. For example, without any prior cake-decorating experience, we remade this toy car as a cake. We are free to make a replica at a different scale and out of different materials: the car has been scaled-up and made out of madeira cake, frosting, marzipan, cookies, and candy.*

rithm directly, we are able to sidestep the difficulties and edge cases of attempting to reconstruct a full 3D mesh from a scan. In addition, we allow the user to change the scale or the materials of the reproduction relative to the original.

To illustrate the flexibility of this process, we scanned a toy car and used our system to recreate it at a scaling factor of  $2.5\times$  as a cake consisting of a mixture of edible materials (Figure 4). A timelapse of the process can be seen in the video that accompanies this paper. The completed cake was eaten and enjoyed.

### 4.3 Stop-motion animation

Our approach also adapts well to being used as a guide for stop-motion animation. By loading a series of targets, instead of just one, and allowing the user to flip between these targets, it becomes possible to generate a series of snapshots of a physical object that correspond to a desired 3D animation.

While a skilled sculptor can already achieve what our system allows in its basic form, this application provides a capability that is otherwise extremely difficult even for expert artists. This is the ability to deform a physical object into a sequence of shapes such that the animation formed by concatenating images of those shapes appears physically correct. Put another way, while a skilled artist may be able to make a plausible reproduction of a single pose of a human in the act of running, even that skilled artist may have trouble determining how much to adjust the model to make it correspond to the



**Figure 5: Stop-motion animation:** *Our approach can be used to help create stop-motion animation. We demonstrate an end-to-end system, in which a 2.5D video (left and center columns) of a real person walking was captured with a Kinect and used to create a claymation sequence (right column).*


pose of that runner exactly  $1/30$ th of a second later.

To test this, we built an end-to-end system in which a performance is captured by a Kinect and used to make a claymation sequence (Figure 5). As with reconstructing an existing physical object, we have the advantage that the point cloud derived from the Kinect can be used directly for modeling, without needing to generate an explicit 3D mesh for each frame. Modeling proceeds as normal, with the addition of the ability to toggle between frames and to take a snapshot of the scene to save as a frame of the animation. We show thumbnails from the resulting animation in Figure 5, and the full animation can be viewed in our accompanying video.

### 4.4 User testing and limitations

We felt that our approach made it possible for users with no prior experience sculpting to create aesthetically-pleasing 3D sculptures. To test this, we asked six first-time users to reproduce a 3D model of the Stanford bunny out of modeling clay. Each user was asked to make one attempt freehand, referring to a rendering (Figure 1(a)), and one attempt using our system, after a 10 minute training period. Users were given 30 minutes for each attempt. To save time, we requested that only the front of the bunny need look correct. Half of the users started with the freehand attempt, and half with our method. The results are shown in Figure 6.

Users commented that the system greatly improved their ability to achieve the correct proportions. The dynamic of the sculpting process was also changed. When sculpting freehand, users typically felt finished after 30 minutes; they acknowledged that the propor-



**Figure 6: User testing results:** After 10 minutes of training and practice, users were asked to sculpt the Stanford bunny out of polymer clay both freehand (top row), and also using our system (bottom row). 30 minutes were allotted for each task. The odd columns first used the system and then sculpted freehand, while the even columns sculpted freehand and then used our system. For all users but the artistically-inclined column 3, the system greatly improved their ability to create a correctly-proportioned bunny.

tions weren't quite right, but also that they weren't sure what to do to fix them. Using our method was slower, but always provided a clear direction in which to proceed; after 30 minutes with our system, users felt that they could have continued to productively improve the model.

Sculptures made with our method were generally better proportioned than the freehand sculptures. Only one user did not benefit from the system (shown in Figure 6 column 3). This user's freehand model was very good. While able to make continual progress using our method, this user was not able to achieve the same quality in the same amount of time. When a sculptor has enough skill to innately know how to adjust a model, automatic guidance just slows them down.

## 5 Conclusion

We have demonstrated a new approach to making a physical object that matches a target 3D shape in which an unskilled user is assisted by guidance projected onto the object as it is being made. We have found that our system greatly increases what novice sculptors are capable of (e.g. Figures 1, 3, 4, and 5). In addition to producing physical replicas of a 3D model, our approach can be used to enable reproduction of existing physical objects, potentially at a different scale and out of different materials. We further demonstrate how our approach can be used to retarget real-world performance capture data to generate stop-motion animation with physically correct dynamics.

We believe that the combination of human actuation and situated feedback from computer-accurate measurements has great potential. In future work, we particularly wish to explore applications to the task of assembling rigid parts to accurately match a target configuration, which may be useful in industrial contexts.

## 6 Acknowledgments

The authors wish to thank Abe Davis, Elena Adams, and the user study participants. The work was partially funded by NSF grant 0964004, Foxconn, and a gift from Cognex.

## References

ADAMS, A., HOROWITZ, M., PARK, S. H., GELFAND, N., BAEK, J., MATUSIK, W., LEVOY, M., JACOBS, D. E., DOLSON, J., TICO, M., PULLI, K., TALVALA, E.-V., AJDIN, B.,

VAQUERO, D., AND LENSCH, H. P. A. 2010. The Franken-camera. *ACM Trans. Graph.* 29, 4 (July), 1.

BIMBER, O., AND RASKAR, R. 2005. *Spatial augmented reality: merging real and virtual worlds*. A. K. Peters, Ltd., Natick, MA, USA.

CASPI, D., KIRYATI, N., AND SHAMIR, J. 1998. Range imaging with adaptive color structured light. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20, 470 – 480.

CHEN, F., BROWN, G. M., AND SONG, M. 2000. Overview of three-dimensional shape measurement using optical methods. *Optical Engineering* 39, 10.

COKE, V. D. 1964. *The painter and the photograph: from Delacroix to Warhol*. University of New Mexico Press.

CURLESS, B., AND LEVOY, M. 1995. Better optical triangulation through spacetime analysis. *Proceedings of the Fifth International Conference on Computer Vision* (June), 987–994.

DIXON, D., PRASAD, M., AND HAMMOND, T. 2010. iCanDraw: using sketch recognition and corrective feedback to assist a user in drawing human faces. In *Proceedings of the 28th international conference on Human factors in computing systems*, ACM, New York, NY, USA, CHI '10, 897–906.

DRUMM, B., 2011. Printrbot. <http://www.printrbot.com/>.

FLAGG, M., AND REHG, J. M. 2006. Projector-guided painting. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*, ACM, New York, NY, USA, UIST '06, 235–244.

HILDEBRAND, K., BICKEL, B., AND ALEXA, M. 2012. crdbd: Shape Fabrication by Sliding Planar Slices. In *Computer Graphics Forum (Eurographics 2012)*, vol. 31.

HOKANSON, T., AND REILLY, C. DIYLILCNC. <http://diylilcnc.org/>.

INOKUCHI, S., SATO, K., AND MATSUDA, F. 1984. Range-imaging for 3D object recognition. *ICPR*, 806–808.

KELLY, S. Bluemax CNC. <http://www.bluemaxcnc.com/Gantry-Router.html>.

- KILIAN, M., FLÖRY, S., CHEN, Z., MITRA, N. J., SHEFFER, A., AND POTTMANN, H. 2008. Curved folding. *ACM Trans. Graph.* 27, 3 (Aug.), 75:1–75:9.
- KRAGIC, D., KRAIG, D., MARAYONG, P., LI, M., OKAMURA, A. M., AND HAGER, G. D. 2005. Human-machine collaborative systems for microsurgical applications. *International Journal of Robotics Research* 24, 731–741.
- LAU, M., OHGAWARA, A., MITANI, J., AND IGARASHI, T. 2011. Converting 3D furniture models to fabricatable parts and connectors. *ACM Trans. Graph.* 30, 4 (Aug.), 85:1–85:6.
- LEE, Y. J., ZITNICK, C. L., AND COHEN, M. F. 2011. Shadow-Draw. *ACM Trans. Graph.* 30, 4 (July), 27:1–27:10.
- LEVOY, M., GINSBERG, J., SHADE, J., FULK, D., PULLI, K., CURLESS, B., RUSINKIEWICZ, S., KOLLER, D., PEREIRA, L., GINTON, M., ANDERSON, S., AND DAVIS, J. 2000. The digital Michelangelo project. In *Proceedings of the 27th annual conference on computer graphics and interactive techniques*, ACM Press, New York, New York, USA, 131–144.
- MAKERBOT INDUSTRIES. MakerBot. <http://www.makerbot.com/>.
- MAKO SURGICAL. RIO Robotic Arm Interactive System.
- MCCRAE, J., SINGH, K., AND MITRA, N. J. 2011. Slices: a shape-proxy based on planar sections. *ACM Transactions on Graphics (TOG)* 30, 6 (Dec.), 168.
- MORI, Y., AND IGARASHI, T. 2007. Plushie: an interactive design system for plush toys. *ACM Trans. Graph.* 26, 3 (Aug.), 45:1–45:7.
- PENNSYLVANIA ACADEMY OF THE FINE ARTS, DANLY, S., AND LEIBOLD, C. 1994. *Eakins and the photograph: works by Thomas Eakins and his circle in the collection of the Pennsylvania Academy of the Fine Arts*. Published for the Pennsylvania Academy of the Fine Arts by the Smithsonian Institution Press.
- RASKAR, R., WELCH, G., AND FUCHS, H. 1998. Spatially augmented reality. In *In First IEEE Workshop on Augmented Reality (IWAR98)*, 11–20.
- RIVERS, A., MOYER, I. E., AND DURAND, F. 2012. Position-Correcting Tools for 2D Digital Fabrication. *ACM Trans. Graph.* 31, 4.
- RUSINKIEWICZ, S., HALL-HOLT, O., AND LEVOY, M. 2002. Real-time 3D model acquisition. *ACM Trans. Graph.* 21, 3 (July), 438–446.
- SELLS, E., SMITH, Z., BAILARD, S., BOWYER, A., AND OLIVER, V. 2009. RepRap: The Replicating Rapid Prototyper-maximizing customizability by breeding the means of production. *Handbook of Research in Mass Customization and Personalization 1*, 568–580.
- SHENG, J., BALAKRISHNAN, R., AND SINGH, K. 2006. An interface for virtual 3D sculpting via physical proxy. In *Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia*, ACM, New York, NY, USA, GRAPHITE '06, 213–220.
- SKEELS, C., AND REHG, J. 2007. ShapeShift: a projector-guided sculpture system (Poster). *UIST 2007 Poster Session*.
- ZHANG, Z. 1994. Iterative point matching for registration of free-form curves and surfaces. *International Journal of Computer Vision* 13, 2 (Oct.), 119–152.