

Interlocking pieces for printing tangible Cultural Heritage replicas

G. Alemanno, P. Cignoni, N. Pietroni, F. Ponchio and R. Scopigno

Visual Computing Lab, CNR-ISTI, Pisa, Italy

Abstract

We propose a technique to decompose a 3D digital shape into a set of interlocking pieces that are easy to be manufactured and assembled. The pieces are designed so that they can be represented as a simple height field and, therefore, they can be manufactured by common 3D printers without the usage of supporting material. The removal of the supporting material is often a burdensome task and may eventually damage the surface of the printed object. Our approach makes the final reproduction cheaper, accurate and suitable for the reproduction of tangible cultural heritages. Moreover, since the proposed technique decomposes the artwork in pieces, it also overcomes the working space limits of common printers. The decomposition of the input (high-resolution) triangular mesh is driven by a coarse polygonal base mesh (representing the target subdivision in pieces); the height fields defining each piece are generated by sampling distances along the normal of each face composing the base mesh. A innovative interlocking mechanism allows adjacent pieces to plug each other to compose the final shape. This interlocking mechanism is designed to preserve the height field property of the pieces and to provide a sufficient degree of grip to ensure the assembled structure shape to be compact and stable. We demonstrate the effectiveness of our approach and show its limitations with some practical reproduction examples.

1. Introduction

Fabrication technologies are a common resource for rapid prototyping and industrial production processes. More recently, the appearance of low-cost devices made these technology available to larger communities and also usable for small-scale production.

Most commercial 3D printing system are now based on the *additive* process, i.e. the shape is created by laying down successive layers of material. Each slice is printed on top of the previous to compose the 3D object. Some method uses meltable materials to produce the layers (SLS, FDM), while others lay liquid materials that are hardened with different technologies. 3D printers may produce complex shapes usually within a few hours, depending on the size of the object. The printing techniques are nowadays enough accurate to reproduce copies of tangible cultural heritages. However there are still some limitations that should be overcome to make this technologies suitable to this purpose:

Size Due to a limited workspace, the produced objects are, usually, very small. Usually, the workspace of a common 3D printer is between 20 and 40 cm size cube. This limi-

tion restrict the usage of printing technique to the reproduction of only small objects. Many tangible cultural heritages (e.g. sculptures or low reliefs) have relevant dimensions, printing them in a reduced scale may lead to a very inaccurate copy and/or a wrong perception. Several details may eventually disappear in the printed object if the printing scale becomes too small. A recent approach overcomes this limitation by decomposing the original object into different components and by glueing them together [LBRM12]. The same approach was used (manually) in an early paper on the 3D fabrication of a 1:1 replica of a wall portion containing inscriptions [BCF*04].

Physical constraints. Most complex models may require additional structures to support the object while being produced. This support structure must be removed after the printing. The materials involved in the printing process may be quite fragile before solidification. Hence, most of printing techniques requires additional structures to support the object while being produced. The usage of supporting material may increase the overall cost of the printing procedure and production time. Moreover, the re-

removal of the support structures may damage the small scale details on objects's surface.

1.1. Contribution

The approach proposed in this paper is based on the idea of subdividing the input shape into several pieces. Each adjacent pair of pieces is designed to interlock in an easy manner.

Although the pieces are in 3D, there are no overhangs. Indeed, opposite to [LBRM12], our pieces are defined to be represented as height-fields and therefore can be printed without the usage of supporting material.

We designed a new interlocking mechanism to plug pairs of adjacent pieces. This allows to assemble the complete object without any supporting structure and in a much larger interval of reproduction scales. As opposite to [LBRM12], the interlocking mechanism offers a sufficient grip to avoid the usage of glue. This is created by shaping compatible teeth profiles between adjacent pieces (on the border of the junction, under the constrain of following the overall input shape).

2. State of The Art

Rapid prototyping techniques [DSdB06] have been created to support the design industry. As discussed in Section 1, even if 3D printing techniques improved a lot in the last years, they still present many practical limitations.

To overcome *workspace limits* Linjie Luo et al. [LBRM12] decompose the original shape into different components, which are fabricated separately and glued together to produce the desired shape. More in detail, they designed an optimization framework named *Chopper* that uses planar cuts to decompose object in smaller parts, so that each part fits into the specific printing volume of a 3D printer. The printing requires a 3D printing machine, and the pieces must be glued together. Still on the theme of constraining objects into a limited space in [ZSMS14] the authors present a method for transforming a 3D object into a unfolding box.

In order to improve the *physical resistance* of the printed objects a method to automatically fortify the input shape throughout the use of additional geometric structures is proposed in [SVB*12]. A sophisticated analysis framework to identify fragile points in printed models was introduced in [ZPZ13]. An optimization method to provide the right balancement of a given shape, to make it stand in equilibrium once printed, was presented by [PWLSH13]. These methods may help the modeler to design more robust and sound fabricable objects.

Other papers focus on better reproduction of *appearance properties* in printed objects: [CGPS08] proposed a technique to enhance colors for rapid prototyping; [WPMR09]

and [MAG*09] reported a method for the improved reproducibility of surface reflectance properties by adding micro geometry; [HFM*10] and [DWP*10] proposed a technique to print specific subsurface scattering characteristics.

All the method we cited above have been designed for additive 3D printers. Another class of methods focus on the generation of low- and high-reliefs, and therefore those methods usually target subtractive fabrication techniques. To our knowledge the first paper to introduce a methodology to automatically derive low- and high-relief for a given 3D scene was [CMS97]. For a more detailed overview of relief construction methods in computer graphics we remind to [KWC*12].

Another class of printing methods have been designed to employ cheaper materials and hardware devices; usually these approaches aims to generate a more approximate copy of the given object and often they exploit cheaper 2D materials and more simple printing techniques. Several methods [MS04, STL06, MGE07] reproduce the input model by means of a set of paper strips (or similar materials) which can be folded and glued together to create a 3D representation. Holroyd et al. [HBLM11] described a method to fabricate a three-dimensional shape through a stack of 2D colored slices. The use of planar material was exploited also by Hildebrand et al. [HBA12]; they proposed a method to semi-automatically fabricate objects made up of interlocking planar slices. This method produces a wide range of nice results, however it does not fit well with complex geometries and was extended to generate better approximation in [SP12, SP13, CPMR14].

3. Fabrication Pipeline Overview

The proposed fabrication pipeline takes as input a 3D model and produces as output a decomposition of the model into interlocking pieces. As we stated previously, each piece can be manufactured with a common 3D printer without the usage of supporting material and then manually assembled. This process can be summarized into three stages (see Fig. 1):

Building the Inner Approximation. We build a polyhedral approximation of the input shape. The approximation must lie inside the volume of the input shape and it must be entirely composed by planar faces.

Decomposing the 3D Shape. We exploit the polyhedral approximation to partition the volume of the input shape. The partitioning defines the interlocking mechanism. The shape of the interlocking system is constrained to have each piece of the decomposition defined as a height field.

Mesh Generation. The final stage of the pipeline produces a 3D model for each relief panel.

These pipeline stages will be described in details in the next sections.






Figure 1: Overview of the pipeline stages to convert a 3D model into a set of interlocking pieces.




(a)




(b)


Figure 2: Image (a): model *A* is the original triangular 3D mesh; model *B* is its polyhedral inner approximation; model *C* presents a combined view of the two meshes together. Image (b): the diagram represents an input mesh *S* and the corresponding polyhedral approximation *P*; the planes of *P* define a decomposition of the embedded volume *H*.




(a)




(b)



(c)



(d)



(e)

Figure 3: The coloured cube is a polyhedral approximation of the generic surface rendered in gray (a); image (b) shows the planes associated to each face; in image (c) we highlight the region F_i that belong only to a single face f_i , while image (d) shows the regions E_i that are shared between two faces and image (e) the regions V_i associated to a vertex where multiple regions overlap.

4. Building the inner approximation

The algorithm starts by modeling a very coarse polyhedral approximation of the input triangular 3D mesh. This approximation represents the internal surface of the hollow reproduction. It gives us a structure to guide the decomposition

of the original shape in panels. One example of polyhedral approximation is shown in Figure 2.a.

The polyhedral mesh approximation that we have to build must satisfy these requirements:

1. Each polygon/face has to be planar.
2. It must be completely embedded in the volume of the original shape.

Currently, there are no automatic re-meshing techniques that provide a shape abstraction with the geometric guarantees we need. Searching for such an automatic solution [CSAD04, BLP*13] is out of the scope of this paper. In order to construct a very coarse approximation of the original input shape we used an open source modelling system, Blender [Ble]: we started from the input mesh as an overlay to guide the modeling process.

Once we have a polyhedral P base mesh, we define the printing volume as $H = S - P$ (figure 2.b). That printing volume H is the solid object that will be decomposed into pieces, which should interlock and will be fabricated separately.

5. Decomposing the shape

To create the interlocking system we need to associate to each face f_i of P a portion H_i of H (and therefore of the original surface S). The H_i portion of H is defined by a prism built on the planar offset of the each face f_i . These volumes cover the entire domain H and they overlap in some common regions around edges and vertices of the polyhedral decomposition. As shown by Figure 3, we consider the different intersections of H_i and we use them to partition H into three sub volumes classes: F_i , E_i and V_i that correspond, respectively, to the portions of space belonging to just a single face, to two faces adjacent along an edge and to the space around a vertex shared by three or more faces. Each region E_i , that corresponds to the shared volume between two adjacent faces, will be used to build a comb shaped joinery mechanism: the surfaces of E_i will be distributed among the two pieces corresponding to the two faces incident on e_i (see Fig.4). Conversely, each region V_i will be associated to only one of the faces according to parity of the joinery mechanism and priority of each piece (see Fig. 3).

5.1. Joints

Joints are necessary to properly interconnect the whole structure of panels and to provide the required structural stiffness/robustness. The idea is to provide the interlocking mechanism by splitting the shared space E_i between each pair of adjacent faces of the base mesh. The simple case illustrated in Fig. 4 shows two orthogonal pieces composing part of the rounded cube of Figure 3. The Volume of E_i is distributed over the two faces (see the corresponding green and yellow chunks); in this example the two V_i regions are simply assigned to the yellow face.




Figure 4: The image shows the interlocking mechanism.




Figure 5: Different cases based on dihedral angles: (a) **convex** regions with dihedral angle ≤ 90 ; (b) **convex** regions with dihedral angle > 90 ; (c) **concave** regions with dihedral angle ≤ 90 ; (d) **concave** regions with dihedral angle > 90 ; (e) **flat** regions with dihedral angle = 180.

In general, depending on the dihedral angles of the polyhedral approximation we may encounter different situations. Those cases are shown by Fig. 5. Notice the interlocking system must take into account some constraints, in the case we would ensure that each piece should be defined and reproducible as a height field (lifted along the direction perpendicular to the associated inner plane). Coping with this constraint may be tricky in the general case.

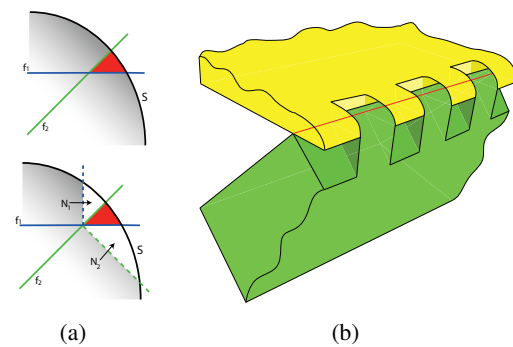


Figure 6: Dihedral angle < 90 degrees on local convex geometry. E_i region (in red) becomes smaller; N_1 and N_2 are holes left by f_1 and f_2 .

Convex regions with angle $\varphi \leq 90$ degrees.

Let's first consider a local convex geometry with an acute dihedral angle (see Fig. 5.a). The surface of E_i (represented

in red in Fig. 6.a) is distributed alternately between the faces incident on e_i and offers the grip to keep the two pieces firmly connected. Notice that some portions of S , indicated as N_1 and N_2 , would be excluded alternatively from F_1 and F_2 . As it can be noticed in Figure 6.b this leaves some unpleasant hole in the final produced model. However, this limitation is only mandatory if we want to keep the structure of each piece to be a 2.5D height field.

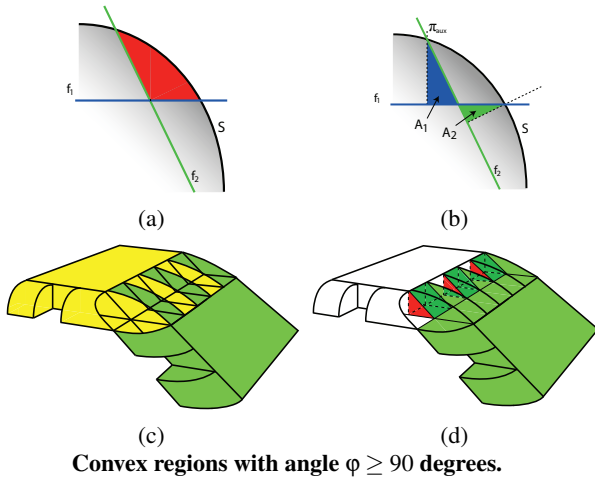


Figure 7: Convex situation when dihedral angle is higher than 90 degrees (a) E_i becomes bigger as the angle increases in degrees (b) The regions A_1 and A_2 need to be modeled to provide support to the interlocking mechanism; (c) and (d) show the resulting pieces with the internal supporting structures (red). No gaps are produced in this case.

Convex regions with angle $\varphi \geq 90$ degrees.

When the polygon is locally convex and the dihedral is greater than 90 degrees, we have the situation illustrated in Figure 7. In this case, we have to construct an additional support structure to accommodate each teeth to its correct position, avoiding the piece to slide to an undesired position. The supporting structure (represented as A_1 and A_2 in the side-view of Figure 7.b) is built exploiting the opposite face. Figure 7 shows the resulting pieces together with the support structure shown in red.

Concave regions with angle $\varphi \leq 90$ degrees.

For concave regions the partitioning becomes more involute because we have to define an additional cutting plane π_{cut} (see Figure 8.a) . This plane cuts dihedral angle for a concave region into two halves. In order to limit the size of the tooth, for each polyhedral face, we rely on an additional plane π_{aux} . This additional plane is perpendicular to the face and it stays as close as possible to the intersection of π_{cut} with original surface. An illustrative

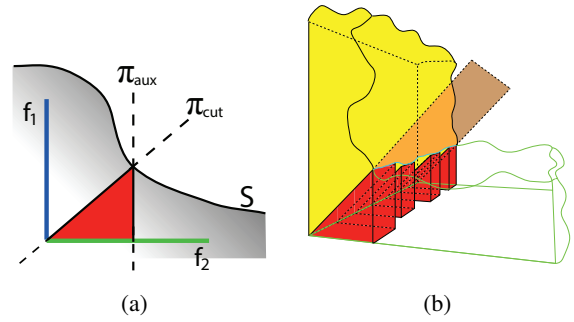


Figure 8: Concave geometry with dihedral angle ≤ 90 degrees: (a) in this case, we have to define a splitting plane π_{cut} and separate the regions E_1 and E_2 ; (b) the resulting teeth (in red) over the two adjoining pieces interlock each other.

representation of the produced pieces is shown in Figure 8.b.




Figure 9: Concave geometry with angle > 90 degrees.

Concave regions with angle $\varphi \geq 90$ degrees.

Similarly to the previous case, we split E_i region using the cutting plane π_{cut} , but the auxiliary plane π_{aux} is perpendicular to the opposite face sharing the edge (Fig. 9.a). This strategy maximizes the grip surface offered by the interlocking mechanism. The teeth may become protruding and visually unpleasant. However, these artifacts are visible only from inside the volume of the object, and therefore they do not produce a visual effect on the final assembled structure.




Figure 10: Local flat geometry.

Almost flat regions with angle ≈ 180 degrees.

When the angle is close to 180 degrees, we rely on two auxiliary planes π_{aux} defined by the side of the prisms to define

the interlocking mechanism. This is a quite simple situation and it is illustrated in figure 10.

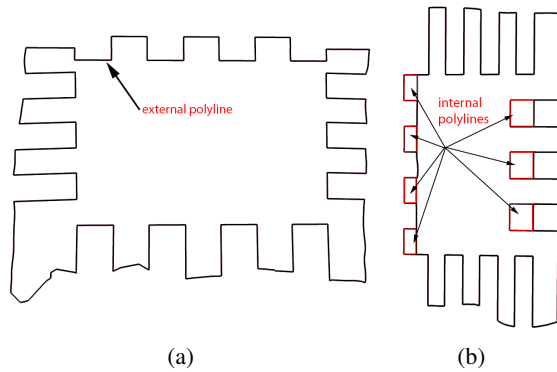


Figure 11: The two drawings show: the external polyline (a); the internal polyline (marked in red in drawing (b)).

6. Mesh generation

Once we defined the subspace sampled by each prism, we have to derive the final tessellation of each piece. A straightforward solution consists in performing a sequence 3D boolean operations: define each piece as the intersection of original mesh with its prism, then subtract the teeth profiles for each piece to derive the interlocking profile. The boolean operations has to be performed precisely, simple relying on distance fields leads to undesirable artifacts on the assembled surface. Unfortunately, computing exact 3D boolean operations may be a tricky task to implement. Conversely, we choose to derive the meshing of each piece through a sequence of three simple steps:

1. Derive the 2D profiles of its base
2. Sample the height field with respect to the original mesh.
3. Triangulate the points to gather the final mesh.

This strategy allows us to perform boolean and tessellation operations in 2D. This choice simplify the mesh generation procedure and makes it more robust.

6.1. Computing 2D profiles

The 2D profiles are computed from the intersection of the planes of P with the original surface S . The 2D profile has to cope with the interlocking mechanism defined in the previous section. We rely on 2D boolean operation using *Clipper*, an open source polygon clipping library [Joh].

The basic procedure to extract the 2D profile is straightforward: given a face of the base mesh P , we starts from the first halfedge pointed by the face and proceeding with the next, following the counter-clockwise order. Each edge e_i is subdivided into a fixed number of parts to generate a set of 2D polygonal masks to remove parts of E_i and V_i , V_j through a sequence of 2D boolean operations.




Figure 12: The sequence of operations to compute the external polyline: (a) the external profile, the face of the base mesh and the direction of the half edge; (b) a set of rectangular masks are alternated along the profile; (c) the final profile is obtained throughout a sequence of boolean operations.




Figure 13: (a) $PCUT$ results from the intersection of π_{cut} with the surface S ; (b) The masks iM are finally used to derive the external polyline; (c) and (d) The region EXT created by π_{aux} is used to define the extension of the polyline.

In the case of concave edges we have to manage the intersection with the additional planes π_{cut} and π_{aux} . This situation is illustrated in figure 13.a and 13.b, the intersection of the π_{cut} with the original surface, and projected on the base face, generate a polyline $PCUT$. Similarly to the convex case, we use this polyline to define a subtractive mask and derive the external polyline.

Additionally, we use the π_{aux} to define the extension of the profile in case the angle is > 90 (see figure 13.c and 13.d).

In the hypothesis that we would be able to print pieces perfectly, no tolerance should be considered while modelling the interlocking pieces. But in this case a high friction between joints may compromise the assembling procedure. On the other hand, a rough printing process may considerably decrease the friction between pieces, making the construc-

tion structurally weak. This issue can be resolved by simply offsetting the external polyline. The sign and magnitude of the offset have to be estimated by considering the tolerance of the printing machine.




Figure 14: The procedure used to generate depth maps.

6.2. Sampling the height field

We use OpenGL and rendering cycles for the generation of the height field. We place a camera in the barycenter of the inner face of each interlocking piece, then we extract the height map through an orthographic projection (see Fig. 14). This simple sampling process cannot be directly applied to local concave geometry. In this case we have to clip the rendering with the plane π_{cut} .

6.3. Triangulating the height field

The first stage of the remeshing process consists of tessellating the base of the piece. This entire process is illustrated in Figure 15. The external polyline is first enriched with additional samples (both in the interior of the polygonal region and on its border), and then we use a 2D constrained Delaunay triangulation library *Triangle* [She]. The triangulated base is cloned and displaced as specified by the height map, then the two profiles are joined by using a triangle strip. An example of final result is shown in Figure 15.d.

As previously stated in Section 5, the situation of a convex geometry with angle > 90 requires an additional support structure to lay down the opposite piece and slide it to its correct position. This implies the use of an additional constraint when tessellating the upper part.

7. Results

We have tested the printed solution on few datasets. We report here the results obtained on three sample digital artworks:

- The first model represents the Ruthwell Cross (figure 16), an Anglo-Saxon tall stone cross (slightly more than 5 meters in height), 8th cent. AD, conserved in the Ruthwell




Figure 15: The tessellation of the base of a piece: (a) we take as input the external profile and the internal sampled points; (b) we enrich the border with additional points to produce a nice tessellation; (c) we tessellate the polygon using the samples and the boundary edges; (d) the resulting tessellated surface.

| Ruthwell Cross | | |
|-----------------|------------------------|----------------|
| # pieces | Average tris per piece | Modelling time |
| 75 | 45.000 | 30 min. |
| Generation Time | Assembly Time | Volume Saving |
| 3 min | 35 min. | 55% |
| Gothic Columna | | |
| # pieces | Average tris per piece | Modelling time |
| 26 | 100.000 | 30 min. |
| Generation Time | Assembly Time | Volume Saving |
| 6 min | - | 40% |
| Urna | | |
| # pieces | Average tris per piece | Modeling time |
| 44 | 120.000 | 20 min. |
| Generation Time | Assembly Time | Volume Saving |
| 8 min | - | 75% |

Table 1: Statistics and running times for computing the subdivision in interlocked pieces of the three replicas shown in the figures.

Church, Dumfriesshire, Scotland. The Ruthwell cross has a quite regular shape, but it is characterised by many carvings covering all its surface and featuring both holy human figures, decorations and verses of an ancient Anglo-Saxon poem written in runes. The reproduction in a small scale will therefore wipe out all the carved details, making a small replica totally improper as a reproduction of the

artwork. The digital model was built with active scanning (laser-based). The model was fabricated using a Fused Deposition Modeling 3D printer. The shape of the pieces guaranteed the absence of support and was printed using the default setup of the printer. The resulting model was quite strong and we succeeded to get the right tolerances for the interlocking mechanism so that it was a bit hard to assemble but quite resistant in the end without the use of any glue.

- The second model represents a portion of a Gothic column (figure 17), part of the Arrigo VII mausoleum, now conserved at the Museum of the Opera Primaziale, Pisa, Italy. It is a good example of a quite complex surface, with a lot of carved details protruding (high-relief) from the twisted basic shape of the column. Also this model was built with active scanning (laser-based).
- The third model represents an Etruscan urn (figure 18), discovered in an underground grave located near Perugia, now conserved at the Archeological Museum of Perugia, Italy. It is an example of an artifact with a quite simple shape, but again characterised by a carved decoration (bas-relief). 3D model acquired with laser scanning.

Table 1 reports some statistic on the entire process. The term *Modelling time* refers to the time required by a non expert user to model the internal base mesh (that is the user-assisted phase); the term *Generation time* is the time required for the automatic generation of the interlocking pieces.

It can be noted that the process is not considerably time consuming, since the time spent for actual printing the replica object is at least one order of magnitude higher than the time needed to produce the fragmented digital representation from the input mesh (the fabrication time is not reported in the table, since it depends on the specific printer used). A further benefit of our method is a considerable saving of material required to print the overall object, due to the hollow space inside the replica (from 40% to 75% in the proposed examples).

8. Conclusions and Future Work

We have presented a new flexible technique for supporting the production of replicas whose overall size could be by far larger than the working space of the available fabrication device. The approach proposed is based on the decomposition of the shape to be printed in a number of interlocking pieces, where each one of these satisfies the working space limit of the 3D printer. Those pieces can be easily manually mounted by means of the interlocking mechanism provided.

While the proposed approach successfully adapts the usage of printing technique to create effective replicas of tangible Cultural Heritage (extending the size flexibility), the process requires some user-assisted intervention. This is mainly related to the initial inner shape approximation, that has been

manually generated using a 3D modeling system for the provided examples. Currently the problem of finding a very coarse polygonal inner approximation of a geometric shape is still challenging and it still needs further investigation to be solved in an automatic and robust manner on complex 3D models. A second point that could be improved is the shape assigned to the joinery mechanism: especially close to sharp features and narrow edges, the shape of the teeth is sometimes visible. A different teeth shape would make this junction much less visually perceivable.



Figure 17: The column from the Mausoleum of Arrigo VII, Pisa, Italy.

Acknowledgments

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement 313193 (EC INFRA "ARIADNE" project).

References

- [BCF*04] BALZANI M., CALLIERI M., FABBRI M., FASANO A., MONTANI C., PINGI P., SANTOPUOLI N., SCOPIGNO R., UCCELLI F., VARONE A.: Digital representation and multimodal presentation of archeological graffiti at Pompei. In *The 2nd EUROGRAPHICS Workshop on Graphics and Cultural Heritage, VAST04: The 5th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage* (2004), Eurographics Association, pp. 93–103. 1



Figure 16: The Ruthwell Cross reproduction.

- [Ble] BLENDER FOUNDATION: Blender - the open source, cross platform suite of tools for 3d creation. <http://http://www.blender.org/>. 4
- [BLP*13] BOMMES D., LÉVY B., PIETRONI N., PUPPO E., SILVA C., ZORIN D.: Quad-mesh generation and processing: A survey. *Computer Graphics Forum* 32, 6 (2013), 51–76. 4
- [CGPS08] CIGNONI P., GOBBETTI E., PINTUS R., SCOPIGNO R.: Color enhancement for rapid prototyping. In *The 9th VAST International Symposium on Virtual Reality, Archaeology and Cultural Heritage* (Aire-la-Ville, Switzerland, 2008), Eurographics Association, pp. 9–16. 2
- [CMS97] CIGNONI P., MONTANI C., SCOPIGNO R.: Computer-assisted generation of bas-and high-reliefs. *J. Graphics, GPU, & Game Tools* 2, 3 (1997), 15–28. 2
- [CPMR14] CIGNONI P., PIETRONI N., MALOMO L., ROBERTO S.: Field aligned mesh joinery. *ACM Transacion on Graphics*. 33, 1 (2014), art.11–1..12. 2
- [CSAD04] COHEN-STEINER D., ALLIEZ P., DESBRUN M.: Variational shape approximation. In *ACM SIGGRAPH 2004 Papers* (New York, NY, USA, 2004), SIGGRAPH '04, ACM, pp. 905–914. 4
- [DSdB06] DIMITROV D., SCHREVE K., DE BEER N.: Advances



Figure 18: The Etruscan urn, Perugia, Italy.

in three dimensional printing – state of the art and future perspectives. *Rapid Prototyping Journal* 12 (2006), 136–147. 2

- [DWP*10] DONG Y., WANG J., PELLACINI F., TONG X., GUO B.: Fabricating spatially-varying subsurface scattering. *ACM Trans. Graph.* 29 (July 2010), 62:1–62:10. 2
- [HBA12] HILDEBRAND K., BICKEL B., ALEXA M.: crdbd: Shape fabrication by sliding planar slices. *Comp. Graph. Forum* 31, 2pt3 (May 2012), 583–592. 2
- [HBLM11] HOLROYD M., BARAN I., LAWRENCE J., MATUSIK W.: Computing and fabricating multilayer models. *ACM Trans. Graph.* 30, 6 (Dec. 2011), 187:1–187:8. 2
- [HFM*10] HAŠAN M., FUCHS M., MATUSIK W., PFISTER H., RUSINKIEWICZ S.: Physical reproduction of materials with specified subsurface scattering. *ACM Trans. Graph.* 29, 4 (July 2010), 61:1–61:10. 2

- [Joh] JOHNSON A.: Clipper - an open source polygon clipping library. <http://www.angusj.com/delphi/clipper.php>. 6
- [KWC*12] KERBER J., WANG M., CHANG J., ZHANG J. J., BELYAEV A., SEIDEL H.-P.: Computer assisted relief generation - a survey. *Computer Graphics Forum* 31, 8 (2012), 2363–2377. 2
- [LBRM12] LUO L., BARAN I., RUSINKIEWICZ S., MATUSIK W.: Chopper: partitioning models into 3d-printable parts. *ACM Trans. Graph.* 31, 6 (Nov. 2012), 129:1–129:9. 1, 2
- [MAG*09] MATUSIK W., AJDIN B., GU J., LAWRENCE J., LENSCH H. P. A., PELLACINI F., RUSINKIEWICZ S.: Printing spatially-varying reflectance. *ACM Trans. Graph.* 28, 5 (Dec. 2009), 128:1–128:9. 2
- [MGE07] MASSARWI F., GOTSMAN C., ELBER G.: Paper-craft models using generalized cylinders. In *Proceedings of the 15th Pacific Conference on Computer Graphics and Applications* (Washington, DC, USA, 2007), PG '07, IEEE Computer Society, pp. 148–157. 2
- [MS04] MITANI J., SUZUKI H.: Making papercraft toys from meshes using strip-based approximate unfolding. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 259–263. 2
- [PWLSH13] PRÉVOST R., WHITING E., LEFEBVRE S., SORKINE-HORNUNG O.: Make It Stand: Balancing shapes for 3D fabrication. *ACM Transactions on Graphics (proceedings of ACM SIGGRAPH)* 32, 4 (2013), 81:1–81:10. 2
- [She] SHEWCHUK J. R.: A two-dimensional quality mesh generator and delaunay triangulator. <http://www.cs.cmu.edu/~quake/triangle.html>. 7
- [SP12] SCHWARTZBURG Y., PAULY M.: Design and optimization of orthogonally intersecting planar surfaces. In *Computational Design Modelling (Proc. of Design Modelling Symp. 2011, Berlin)* (2012), pp. 191–199. 2
- [SP13] SCHWARTZBURG Y., PAULY M.: Fabrication-aware design with intersecting planar pieces. *Computer Graphics Forum* 32 (2013). 2
- [STL06] SHATZ I., TAL A., LEIFMAN G.: Paper craft models from meshes. *Vis. Comput.* 22 (2006), 825–834. 2
- [SVB*12] STAVA O., VANEK J., BENES B., CARR N., MÈCH R.: Stress relief: improving structural strength of 3d printable objects. *ACM Trans. Graph.* 31, 4 (July 2012), 48:1–48:11. 2
- [WPMR09] WEYRICH T., PEERS P., MATUSIK W., RUSINKIEWICZ S.: Fabricating microgeometry for custom surface reflectance. *ACM Trans. Graph.* 28, 3 (July 2009), 32:1–32:6. 2
- [ZPZ13] ZHOU Q., PANETTA J., ZORIN D.: Worst-case structural analysis. *ACM Trans. Graph.* 32, 4 (July 2013), 137:1–137:12. 2
- [ZSMS14] ZHOU Y., SUEDA S., MATUSIK W., SHAMIR A.: Boxelization: Folding 3d objects into boxes. *ACM Trans. Graph.* 33, 4 (July 2014), 71:1–71:8. 2