From 3D models to 3D prints: an overview of the processing pipeline

Marco Livesu[1]   Stefano Ellero[2]   Jonàs Martínez[3]   Sylvain Lefebvre[3]   Marco Attene[1]

[1]IMATI-CNR, Genoa, Italy   [2]STAM SRL, Genoa, Italy   [3]INRIA, France

Abstract
Due to the wide diffusion of 3D printing technologies, geometric algorithms for Additive Manufacturing are being invented at an impressive speed. Each single step along the processing pipeline that prepares the 3D model for fabrication can now count on dozens of methods, that analyse and optimize geometry and machine instructions for various objectives. This report provides a classification of this huge state of the art, and elicits the relation between each single algorithm and a list of desirable objectives during model preparation – a process globally refereed to as Process Planning. The objectives themselves are listed and discussed, along with possible needs for tradeoffs. Additive Manufacturing technologies are broadly categorized to explicitly relate classes of devices and supported features. Finally, this report offers an analysis of the state of the art while discussing open and challenging problems from both an academic and an industrial perspective.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modelling—Geometric algorithms, languages, and systems J.6 [Computer-Aided Engineering]: Computer-aided manufacturing (CAM)—

1. Introduction

Digital 3D models have a central role in modern product development where Additive Manufacturing (AM) is taking up in industrial practice. An industrial product has a typical lifecycle constituted of the four successive phases of *conception*, *design*, *realization*, and *service*: once a product *concept* is designed by a CAD expert, the resulting 3D model must be realized through proper fabrication tools, and then put into service for actual use. The sequence of operations required to move from *design* to *realization* is known as *Process Planning* (PP). For AM technologies, a typical PP is depicted in Fig. 1 and includes a tessellation step, the calculation of a building direction, the calculation of parallel slices, and the conversion of the slices into commands for the printer. In some cases additional operations can be necessary, such as the calculation of support structures or the partitioning of large objects that do not fit the printing chamber.

1.1. Survey contents and objectives

The objective of this report is to give the reader a comprehensive overview of the Process Planning pipeline in the context of Additive Manufacturing. We provide a deep picture of the algorithms, data-structures and shape representations that are involved in turning an input virtual model into a set of instructions that the AM machine can understand. We cover key publications from the infancy of AM to the most recent advances. We categorize algorithms based on the PP problem they solve, and discuss their impact on various desirable properties of the final parts.

Process Planning only. Recent surveys cover AM in a much broader sense, from a user's perspective. In particular, the reader can refer to the excellent survey by Gao and colleagues that proposes an overview of the status and challenges of AM in engineering [GZR*15]. It covers fundamental principles and available technologies, design and modelling for AM, and its impact on the industry and society at large. The survey by Medeiros and colleagues [MeSEPC16] provides a clear picture of recent contributions in *functional* fabrication that is, parts that embed a functionality (e.g. articulated or deformable). We do not cover these topics, or only mention them from the perspective of process planning. We also do not cover topics regarding the design of the part. We recognize that a knowledge of the eventual manufacturing device can drive the design itself (e.g. [Ros07, BDDH11]). However, while we mention some methods that tightly interact with the process planning, a complete overview of the methods in the "grey area" between design and process planning would lead us too far from the scope of this survey. The reader can refer to [HPE16, TMV*16, SAZ15] and references therein for discussions on new trends in shape design for AM. Furthermore, we do not provide in depth technical reviews of the current 3D printing technologies. In Section 1.3 we briefly go through the most important AM paradigms available. Our goal, however, is not to provide a comprehensive list, but rather to discuss the main properties – and limitations – that impact the process planning pipeline.

Motivation and timeliness. Researchers operating in the area of process planning have been quite active in recent years, producing a

```
G90 G21
G00 X-37.939 Y8.793 Z16.000
M03 F200
G1 X-37.939 Y8.793 Z16.000
G1 X-37.939 Y8.793 Z6.000
G1 X-35.617 Y8.793 Z6.000
G1 X-34.780 Y8.748 Z6.000
G1 X-34.060 Y8.622 Z6.000
G1 X-33.448 Y8.379 Z6.000
G1 X-32.890 Y7.992 Z6.000
```
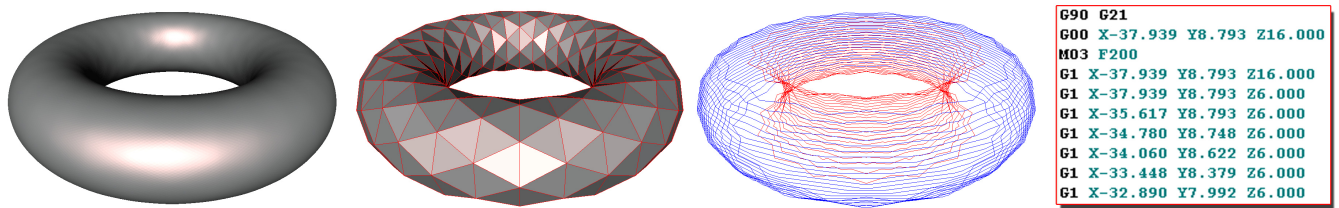
Figure 1: Typical process planning phases in additive manufacturing: a design model (left) is tessellated to enter the Process Planning phase (centre-left). Such a tessellation is sliced (centre-right), and each slice is converted to a sequence of machine instructions (right). Other operations are typically required depending on the specific technology at hand.

tremendous amount of publications. The last (and only) survey focusing on the process planning for AM was published sixteen years ago [KMD00] and, although the process planning pipeline has not changed much from an ideal point of view, many advances in algorithms have been proposed since. The same goes for surveys that focus on a more specific subject, such as [DM94] and [PVRD03], that review slicing procedures for AM and were published in 1994 and 2003, respectively. One exception is the recently released survey on the role of build orientation in layered manufacturing [TJ13] which, however, reviews contributions coming from the mechanical engineering community, and does not cover recent literature coming from other fields (e.g. computer graphics).

Target reader. Indeed, additive manufacturing poses problems that cross several fields. Among classical players, such as designers, mechanical engineers and material scientists, new communities have become interested in it. For instance researchers in computer graphics, with their expertise at creating and manipulating digital 3D objects; but also applied mathematicians, with their expertise in fields such as topology optimization which is becoming increasingly important in the generation of support structures and infills. This report aims to be an entry door for those who want to contribute to AM, avoiding to re-invent the wheel, and promoting more interdisciplinary approaches. Both industry people and academics are targeted. From an industry perspective, this survey wishes to be a support to answer questions such as: I have a design, which algorithms are suitable to process it? What geometric processes would I need to implement? From an academic perspective, scholars can use this report to validate the originality of their own ideas in the area, to look for existing solutions to similar problems, and to know which open challenges deserve their attention.

1.2. General introduction to process planning

Process Planning (PP) is the set of operations performed after designing a model and before its actual manufacturing. PP techniques have been mostly developed in traditional subtractive manufacturing (SM) such as machining, and it is worth having a picture of these methods before entering the realm of modern Additive Manufacturing. Indeed, if in AM most of the steps can be performed algorithmically, in SM the process can be so complex that the experience of a skilled manufacturer is unavoidable. For example, an appropriate machining must be identified for each model (i.e. turning, milling, broaching, drilling/boring, etc.) depending on its geometry: cylindrical components require turning at least for the

production of the overall body, flat or irregular surfaces must be milled, features such as slots and holes require specific operations such as broaching, drilling, tapping, boring, etc. Depending on its complexity, a single model is often machined in several steps, occurring in different machines (e.g. turning, then milling, drilling and finally grinding). The next step of the process is the definition of the stock, that is the piece of raw material to be machined. The choice is based on the material of the object, on its overall dimensions and on the characteristics of the machine (e.g. only cylindrical parts can be loaded on lathes). For CNC machines, the defined stock, together with the characteristics of the machine (e.g. cutting speed, forwarding speed, maximum tool displacements, etc.) are input to a CAM software tool that generate the instructions for the CNC machine.

Innovative additive manufacturing processes have revolutionised the whole pipeline of manufacturing; process planning is the step which has benefited the most from this innovation. In fact, except for part finishing, the whole manufacture occurs in one machine. The user does not have to define any stock and also the toolpath computation is considerably simpler.

The shape of the object is not fabricated as a whole: it is manufactured as a union of layers (or slices), and layers are built one by one, one on top of the other, to form the final geometry. This involves a reduction of the PP problem dimension, because 2D toolpaths are generated (within each slice) instead of complex 3D paths. However, the result is now orientation dependent: gravity plays an important role in the manufacture of each layer. In addition, its influence must be kept into account in the production of overhangs (which may require external supports), as stresses might be induced to the layer, potentially bringing to the failure of the production. Finally, the slice-by-slice approach makes it easier to fabricate multi-color and multi-material prototypes: indeed, since a slice is 2D, all its points are always accessible by the printing tool, independently of the overall 3D object complexity.

1.3. 3D printing technologies

In this Section we briefly describe the main additive manufacturing technologies. Our goal is not to provide an extensive list, but rather to discuss the main properties – and limitations – that impact the process planning pipeline.

All the technologies we consider build an object layer after layer. They mainly differ, however, by whether they actually *locally deposit* material or whether they *solidify material* within an otherwise

AM	Additive Manufacturing		LDI	Layered Depth Images
AMF	Additive Manufacturing file Format		LDNI	Layered Depth Normal Images
CAD	Computer Aided Design		PP	Process Planning
CAM	Computer Aided Manufacturing		PSW	Process SoftWare
CLI	Common Layer Interface		RP	Rapid Prototyping
CNC	Computer Numerical Control		SM	Subtractive Manufacturing
CSG	Constructive Solid Geometry		SLA	Stereolithography
DLP	Direct Light Processing		SLM	Selective Laser Melting
FDM	Fused Deposition Modelling		SLS	Selective Laser Sintering
FFF	Fused filament fabrication		STL	STereoLithography file format

Table 1: *Acronyms reference table.*

non-solid substance. This has important implications on the process planning and this distinction therefore is the basis of our two main categories, *material deposition* and *layer solidification*.

A second fundamental distinction between these technologies is whether they deposit/solidify material along continuous paths (vector) or whether they rely on a discrete device (raster). This directly drives whether the output of the processing pipeline is a set of continuous paths (vector) or a set of images (raster).

1.3.1. Material deposition

Material deposition refers to methods that create the next layer by *locally* depositing material on a previously printed layer. This encompasses techniques such as material extrusion (e.g. fused filament deposition [Cru89]), material jetting (e.g. UV sensitive resin droplets [Pol98, SARW*15]), directed energy deposition (e.g. laser cladding [Vil99]).

Vector or raster. Filament fused deposition is a vector approach as it deposits continuously along paths. The motion of the extruder is achieved through either a three axis gantry, or a delta robot configuration. Processes relying on resin droplets are usually discrete, similarly to inkjet printers. The print head is often attached to a two axis gantry, with the build plate moving up or down along the layering direction.

Properties. The key advantages of material deposition are the ability to combine multiple materials, a printing time that mostly depends on the part volume, and the ability to fully enclose voids. A major inconvenient however is the strong requirement for support structures, since material can only be deposited on top of an already existing layer. The process planning therefore has to automatically generate disposable support structures, which we discuss in Section 3.3.1.

Finally, some of these technologies are able to print *out-of-plane*, for instance generating a continuous spiralling path from bottom to top (e.g. *spiralize* feature of the *Cura* open source slicer) or even wire-frame structures [MIG*14]. We discuss this in more details in Section 3.4.2

1.3.2. Layer solidification

Layer solidification refers to all the processes that build the object by solidifying the top (or bottom) surface of a non-solid material (powder, liquid), typically within a tank. This starts by lowering the tank, adding a full layer of non-solid material, and then using a process that solidifies the material in specific places. This encompasses

technologies such as vat photo-polymerization (e.g. stereolithography or SLA), powder bed fusion (e.g. selective laser sintering or SLS), binder jetting (e.g. plaster powder binding [SHCW94]), and sheet lamination (e.g. paper layering–cutting [Mco05]).

Vector or raster. SLA processes have both variants, relying either on a laser beam (vector) or on a projected image using a DLP projector (raster). SLS processes are typically driving a laser beam through continuous motions, following contour paths. Both for SLS and SLA, beam motions are obtained using mirrors and galvanometer mechanisms – forming a so-called laser scanner – providing fast and precise movements.

Properties. A major advantage of layer solidification is the reduced need for support structures on complex geometries, enabling a much wider range of parts to print without any support. Note that supports may still be necessary to stabilize the part (see Section 3.3.1) as it may be able to move within the non-solidified material (in particular with liquid resins, but also in powder depending on part weight). Another need for support arise from heat dissipation issues, in particular with metal powder melting.

A drawback of within-layer solidification is that it is more challenging to mix different materials. This is achieved on some technologies by locally depositing additives, such as pigmented inks on powder binding 3D printers such as [SHCW94] or the HP Jet Fusion 3200, or by masking techniques using different resin tanks on SLA printers [ZCYK11]. Another drawback is the necessity for non-solidified materials to exit cavities: this prevents the formation of fully closed empty voids within the part.

Layer solidification presents an interesting tradeoff regarding printing times. Each layer starts by a full-tank layer filling (powder) or resin sweep (SLA) which usually takes the same, constant time. This implies that printing time is much more impacted by the height (number of layers) of the object, than by the solidified volume. As a consequence, printing a single small object is generally time consuming, while printing objects in batches can lead to significantly reduced print time per-parts, as the constant per-layer time is amortized. We discuss printing in batches in Section 3.6.

1.4. Pipeline

1.4.1. Overview

The traditional process planning pipeline (Figure 2) starts from design specifications: these define not only the geometry of the part to be printed, but also additional requirements such as dimensional tolerances, need for surface finishing, materials to be used, and many other characteristics of the part. Industrial practice in this phase is still dominated by 2D technical drawings, even if software tools (e.g. modellers from Autodesk and Dassault) and standard format specifications (e.g. STEP ISO 10303) exist to produce and represent accurate designs directly in 3D: in this latter case the so-produced models are known as CAD (Computer-Aided Design) models. Either 2D technical drawings or 3D CAD models must normally be cast to a corresponding CAM (Computer-Aided Manufacturing) representation to undergo the fabrication process, though some integrated CAD/CAM systems are emerging to make this transition as transparent as possible.

	Vector/ Raster	Multiple Materials	Support Structures	Cavities	Build Time
Material deposition (e.g. fused filament, resin droplets)	Both	Supported	Overhangs	Supported*	Depends on part volume
Material solidification (e.g. liquid resin, various powders)	Both	Supported (technology dependent)	Stability / Heat dissipation	Not supported	Depends on number of layers

*: assuming the boundary of the internal cavities can be printed without support structures.

Table 2: Relation between broad AM technological solutions and supported features.

Thus, in a standard product development pipeline, moving from the CAD to the CAM world represents the switch from the design to the process planning phase. Based on the CAM model, the final goal of Process Planning is to determine the machining instructions that the fabrication tool must execute to build the part. Due to the layer-by-layer nature of Additive Manufacturing, it is important to select an appropriate building direction and to slice the model accordingly. Each of the slices must then be converted to a proper *toolpath*, that is, to a sequence of movements that the building tool must follow to fabricate the slice: these movements track the outer boundary of the slice, but also its inner parts and possible support structures. Figure 1 summarizes these main steps.

It is worth mentioning that process planning is typically an iterative procedure, and it might happen that the design specifications are not compatible with the fabrication technology (see Section 3.1): in this case the control must go back to the design phase for the necessary updates. Fortunately, Additive Manufacturing technologies pose much fewer constraints on the design geometry (wrt to, e.g., CNC milling with limited degrees of freedom), but still there are cases where the aforementioned re-design cannot be avoided.

1.4.2. Current practice

Modern 3D printing companies and services accept designs in different representations, each leading to a different complexity for the conversion to an effective CAM model. In mechanical engineering, the vast majority of design models come as a collection of NURBS patches with possible trimming curves: such a *nominal geometry* is normally tessellated to form a CAM model and start the process planning phase. When the shape is simple enough, constructive solid geometry (CSG) tools are normally preferred because they guarantee that the resulting model is solid. In this latter case, most CAD softwares provide a tessellation module, though a new trend of tools is emerging to completely avoid the tessellation and perform the whole process planning on the native representation, including the slicing phase. In other areas (e.g. cultural heritage), the input model is usually produced by a 3D digitization campaign which often leads to a triangular mesh: such a representation is already tessellated, though many operations might be necessary to make it actually enclose a printable solid (see Section 3.1.2).

Whatever the format of the design geometry, the dominating format for the tessellated models used in CAM is the STL. STL represents an unstructured collection of triangles but is mostly used to represent structured meshes: thus, the coordinates of any vertex are encoded once for each of the triangles incident at that vertex, which leads to a highly redundant representation. Furthermore, STL files describe only the surface geometry without any representation of colour, texture or other common CAD model attributes. For these reasons, modern standardization efforts (e.g. the AMF format within ISO/ASTM 52915:2013) use indexed representations to avoid redundancy and allow encoding many useful attribute information such as colours, materials and textures.

Nowadays, however, a typical AM process planning pipeline includes the following steps:

- *Check and possible adaptation of the input geometry to fabrication requirements*. The (possibly tessellated) geometry must enclose a solid that the target printing technology can actually fabricate (Section 3);
- *Building direction*. The model must be correctly oriented to fit the printing chamber, minimize surface roughness and printing time, reduce the need for support structures, ... (Section 3.2);
- *Creation of support structures*. Depending on both the fabrication tool and the shape, additional geometry may be necessary to support overhanging parts and to keep the part from moving during printing (Section 3.3);
- *Slicing*. The model must be converted to a set of planar slices whose distance might be either constant or adaptive (Section 3.4);
- *Machine instructions*. Each slice must be converted to either a sequence of movements of the fabrication tool (vector-based) or a grid of pixels that define the solid part of the slice (raster-based) (Section 3.5).

2. Metrics / Desiderata

Setting up the process planning to fabricate an object is a matter of finding a good tradeoff between different objectives, often depending on the applicative scenario. The PP pipeline can be tuned to strive to optimize for one, or a combination of them. Different criteria have been proposed in the literature. We recap here the most relevant and widely used.
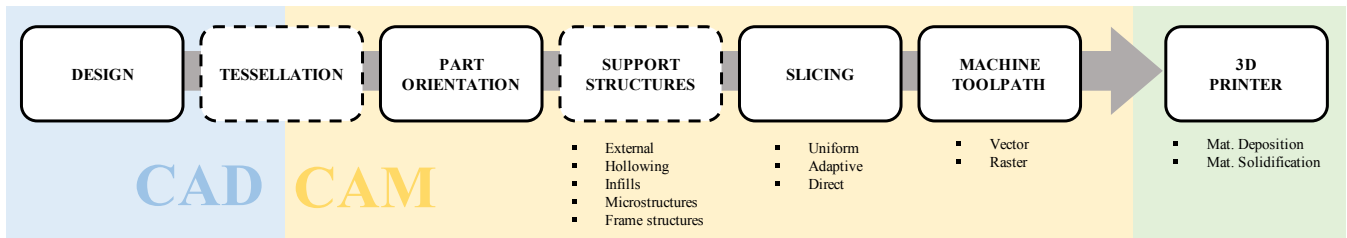
Figure 2: The product developement pipeline for Additive Manufacturing discussed in this paper at a glance. Boxes with dashed boundaries are to be considered optional.

2.1. Cost

In industrial environments it is quite important to keep the production cost as low as possible. The Generic Cost Model [AAD98] puts together all the variables that control the production cost for a single object, and is designed to be general enough to embrace any layered manufacturing process. Similar cost models have been proposed in [BL06b,BL06a,TPR04,PDG99,XLW99]. The typical cost model is defined as the sum of three major components: pre-build, build, and post-processing.

Pre-build cost measures the cost necessary to turn a design into a set of machine instructions to send to the printer. It also accounts for the labour cost (e.g. load the powder into the machine, control or supervise the process planning software) and the time necessary to setup the printer (e.g., cleaning, testing, warming). Methods that aim to minimize the pre-build cost strive for the efficiency of the process, which can be achieved either by using computationally more efficient algorithms or by reducing the user interaction, favouring automatic methods.

Build cost comprises the material cost (for both the part and the supports) and the cost of using the machine for the time necessary to complete the job. The build cost can be reduced in two ways: acting on the printing time, or acting on the material waste. Printing time can be reduced orienting the shape so as to minimize its height (Section 3.2), reducing the number of slices (Section 3.4), or using efficient machine toolpaths (Section 3.5). Material can be reduced by minimizing the volume of support structures, either with a proper choice of the build direction (Section 3.2) or by inserting cavities in the interior of the shape [SDW*16,LSZ*14,WWY*13]. Recent works aim to reduce material waste and achieve a better surface finish by splitting the part into components that can be printed without supporting structures at all (Section 3.7.3). Notice that these methods often trade minimal build cost for the structural strength of the part and, therefore, are not always suitable for industrial production processes.

Post-processing Cost measures the labour, material and time cost necessary to polish the part. This includes: detaching the support structures from the object at the end of the print (Section 3.3.1), and applying some surface finish technique, either manually or through chemical and machine driven processes. These components heavily depend on the process plan. Particularly relevant is the choice of the build direction (Section 3.2) which, in turn, determines the amount and positioning of support structures (Section 3.3.1), and

the extent to which the staircase effect introduced by the layered manufacturing process will affect the part quality.

Notice that in industrial design a shape may undergo several iterations before it is ready for production. To this end, the cost of each print must be multiplied for the time of iterations necessary to finalize the design. Indeed, time is a predominant factor in the estimation of the production cost. To reduce the global cost and enable a faster design optimization loop, [MIG*14] introduced a method to 3D print low-quality wireframe replicas of a shape, reducing the time to complete one iteration by a factor of ten.

2.2. Fidelity

Fidelity is the degree of exactness with which the part has been reproduced starting from its design. Indeed, layered manufacturing is hardly capable of producing a perfect replica of a given design. Parts of the shape that do not align with the building direction expose a typical *staircase* effect (Figure 3a). We distinguish between *form* and *texture*, where the former refers to the overall shape of the prototype, which, to a certain extent, can be quantitatively estimated before printing, and the latter to more local variations of the surface (or high frequencies), which can be approximately estimated only on the printed object since they depend on factors like the printer resolution and the material used. We first discuss the most widely used metrics to evaluate the form approximation error introduced by the staircase effect. Being dependent only on process plan parameters, such as the layer thickness and the shape orientation, these metrics are at the core of the algorithms that strive to optimize the PP pipeline (see e.g. Section 3.2). Then, we introduce metrics to evaluate surface texture, and also discuss some of the practices used in literature to alleviate surface artifacts. Finally, we briefly discuss design compliancy, a criterion that is of fundamental importance in industrial applications, where the printed object is asked to stay within the tolerances set by the designer.

Form. In additive manufacturing objects are created by approximating a given design with a set of layers stacked one on top of the other along the build direction. The goal of form fidelity is to detect the difference in shape between these two entities. We present here the two most widely adopted metrics in literature to evaluate the error introduced by layered manufacturing: the cusp height error and the volumetric difference.

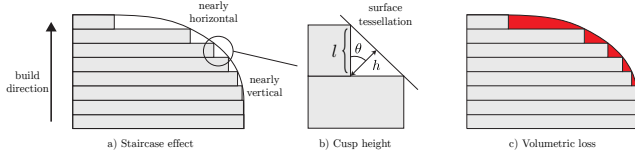- The *cusp-height error* concept was introduced in the context of

Figure 3: *Approximating a curved surface with a stack of layers piled along the building direction introduces the typical staircase effect and reduces fidelity (a). As can be noticed nearly horizontal surfaces introduce more error than nearly vertical ones. Cusp height (b) and volumetric difference (c) are among the most widely used proxies to estimate fidelity.*

adaptive slicing for layered manufacturing [DM94]. It is defined as the maximum distance between the manufactured part's surface and the design surface [AAD98]. It depends on the layer thickness l and the angle θ between the local surface orientation and the build direction (Figure 3b), namely

$$h = \begin{cases} l|\cos\theta| & \textbf{for} \quad |\cos\theta| \neq 1 \\ 0 & \textbf{for} \quad |\cos\theta| = 1 \end{cases} \quad (1)$$

Note that $|\cos\theta|$ is very small for nearly orthogonal angles and grows up to 1 when θ is close to 0 degrees. This well encodes the big approximation error difference between nearly vertical and nearly horizontal surfaces, as shown in Figure 3a. The integral of the cusp height on the whole surface is a good estimator of the fidelity of the printed model. To this end, for triangulated surfaces the cusp height is computed separately on each facet as the dot product between the build direction b and the triangle normal n, $|\cos\theta| = |b \cdot n|$. It is then scaled by the triangle area and normalized by the total mesh area so as to accommodate uneven tessellations and make it scale independent [WZK16];

- The *volumetric difference* is the difference between the volume enclosed by the design and the volume of the printed object (red area in Figure 3c). In [MRI00] Masood and colleagues explain how to evaluate the volumetric error for simple geometries such as cylinders, cubes, pyramids and spheres. Such work was then extended to complex geometries in [MRI03].

Note that, although they are both used as proxies to minimize the staircase effect, volumetric difference and cusp height are not equivalent. Taufik and colleagues [TJ14] express a preference for volumetric error. They observe that, if the shape to be printed contains steep slopes, to a little variation in cusp height may correspond a large variation in volumetric difference, thus making the latter more accurate than the former.

Texture. Also known as *surface roughness*, or *surface finish*, texture aims at measuring tiny local variations on the surface, which determine visual (e.g. the way the object reflects the light), haptic (e.g. the porosity of the surface) and mechanical (e.g. friction) properties of the shape. Unlike form, texture cannot be estimated prior printing, as it mostly depends on parameters like the printer resolution and the printing material. Usually it is measured directly on the printed object, by using sampling techniques [TSB*16].

In industrial environments the design specifies both the desired surface roughness and the metric that should be used to estimate it. The arithmetic mean surface roughness (Ra) is by far the most widely adopted metric, as stated in [TSB*16]. Delfs and colleagues [DTS16] observe that the surface roughness depth (Rz) is a better proxy to measure surface finish, as it is well representative of how a human eye assesses surface quality. Both Ra and Rz are roughness metrics defined in the ISO 4287 standard [ISO97]. They are computed on profile curves obtained by cutting the surface with an orthogonal plane. We point the reader to [Mit09] for a nice explanation of how these metrics are defined and can be estimated in practice; much other information can be found online.

Different strategies have been proposed in the literature to achieve the best surface quality possible. In [RC97] Reeves and Cobb evaluate the *meniscus smoothing* to alleviate the impact of the staircase effect in stereolithography and produce smoother surfaces. It consists of an edited build cycle in which each layer, after solidification, is lift above the upper surface of the resin tank to stretch a meniscus of liquid between each polymerized layer. The resin meniscus is then solidified by using scan data from the previous layer, producing a smoother transition between adjacent layers. Other authors have recently observed that additional artifacts that affect surface finish may be introduced while detaching support structures from the object. In fact, tiny features may be too weak and break during this process, leaving residual support material attached to the surface. Zhang and colleagues [ZLP*15] propose a perceptual method that optimizes for the location of the touching points between surface and supports. Their system tries to *hide* support removal artifacts by placing them at the least salient parts of the shape, as far as possible from its perceptually relevant features. In general, the extent to which support removal may affect surface quality depends on the printer and the material used. In metal printing supports removal is extremely challenging due to the properties of the material involved. Some recent methods to alleviate this problem are discussed in Section 3.3.1.

Design compliancy. In industrial AM environments objects are requested to be compliant with the original design, meaning that both the form and the texture must stay within precise error bounds set by the designer. To this end, a number of form and orientation tolerances are often used in industrial design. Typical form requirements regard the straightness, planarity, circularity or cylindricity of the components. Regarding orientation, typical requirements are parallelism, orthogonality or angularity. Metrology is a vast field and we do not discuss here details regarding how this quantities can be estimated. We point the reader to [TSB*16,HWH15,JQH15,HNX*14] for further details.

2.3. Functionality

When the object to be printed is the result of a shape optimization process, it is typically asked to meet some prescribed functionality requirements. We consider four broad categories: requirements on the robustness of the shape, such as resiliency with respect to previously known or unknown external forces; requirements on the mass distribution, for example to achieve static or dynamic equilibrium; requirements on thermal and mechanical properties, for example
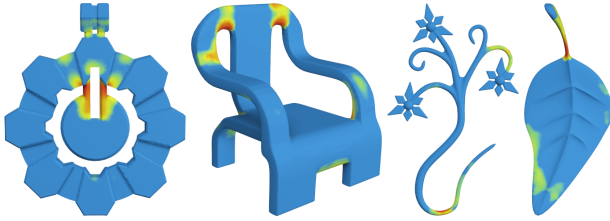
Figure 4: Structural analysis to detect weak object parts. Image courtesy of [ZPZ13].

regarding the heat dissipation or the stiffness of industrial components and requirements in terms of light and waves propagation.

Structural soundness / Robustness. With 3D printing consumers can directly produce their own objects, but not always the digital shapes they started with were meant to be fabricated, and can therefore reveal to be excessively fragile and easily break under cleaning, transportation or handling. Zhou and colleagues [ZPZ13] introduced the worst case structural analysis to detect the most fragile parts of an object (see Figure 4). In [LSD*16] a stochastic finite element method to compute failure probabilities is presented. A number of methods [XLL16, LDJC15, SVB*12] combine a lightweight structural analysis with automatic systems that enforce weak features through operations such as hollowing, thickening, strut insertion and inner structures.

Mass distribution. The distribution of material and cavities inside a 3D printed model has been the subject of recent research. In [PWLSH13] a method to optimize the balance of a shape to make it stand in a given pose is proposed. The work was further extended in [BWBSH14], where a novel optimization of the mass distribution to make an object spinnable around a given axis was proposed. The optimization of the buoyant equilibrium is the subject of [WW16], where a method to create floating objects in a prescribed pose is presented. In [PBJSH16] the authors consider the standing, suspension, and immersion balancing problems for 3D printed objects containing embedded movable masses. [WKW16, MAB*15] propose frameworks for optimizing the interior design and mass distribution of 3D printed objects to achieve static, rotational and buoyant equilibrium.

Thermal / Mechanical properties. Additive manufacturing enables the fabrication of shapes that would be impossible to produce with classical subtractive techniques. Shapes that in the past were interesting only from a pure theoretical standpoint can now be printed and their functionality exploited. To this end, additive manufacturing has fostered a lot of research in fields like topology optimization [DJZ15, BAH11], where the goal is to generate shapes which optimize performances in terms of some physical requirement (e.g. weight, heat dissipation, stiffness). This is important in fields like aerospace industry, where the size and weight of components is a crucial factor, or for the generation of injection moulds, where optimized cooling systems may increase both the productivity and the overall quality of industrial products. Alongside, researchers have been investigating the correlation between the structure of a shape and its physical properties. Recent works

aim to optimize the local structure of a printed object to produce controlled deformations or to meet precise rigidity requirements. We review this body of literature in Section 3.3.2.

Light / Sound. A number of methods use multi-material 3D printing to fabricate curved displays with embedded optical fibers, optimizing fibers design and exploiting internal reflection to guide light inside an object [PRM14, BPH13, WBHP12]. Other methods focus on sound. In [MHR*16] the shape design is posed as non-linear problem that aims to optimize the natural frequencies of a shape, for example to make it sound in a controlled manner, as a musical instrument. Both [UPSW16] and [LLMZ16] propose interactive systems for the design of 3D printed wind musical instruments.

3. Process Planning steps

3.1. Meeting fabrication requirements

Herewith we distinguish between requirements of the shape and requirements of its representation. Shape requirements are printer-specific, and define rules for the compatibility of the geometry with the printing hardware. Representation requirements guarantee that the (tessellated) geometry to be printed actually encloses a solid without ambiguity.

3.1.1. Shape requirements

Checks When the input model comes in raster form (e.g. a voxelization), [TJ11] provides the means to analyze the shape and identify problematic regions whose size drops below the printing resolution (e.g. thin walls or other tiny features). In a slightly more general setting, [RNDA13] describes an approach to estimate the thickness of triangulated models.

Automatically fixing the input While the previous works are meant to *detect* possible incompatibilities, in [WC13] an algorithm is proposed to actually thicken sheet-like structures so as to make them printable. In [SVB*12] a similar approach is presented with a focus on the structural characteristics of the printed prototype: note that this method is not meant to make the model printable, but it shares several aspects with the previous one. If the model cannot fit into the printing chamber due to its size, [LBRM12] proposes an approach to split the model into parts that can be printed separately and reassembled after printing. Apparently no publication deals with the automatic placement of drainage channels for models with internal cavities to be printed with powder bed technologies.

3.1.2. Input representation requirements

Geometry repairing has received increased attention in recent years, not only for 3D printing, but in general for all the scenarios where a "well-behaving" mesh is required (e.g. Finite Element Analysis, advanced shape editing, quad-based remeshing, ...). Some repairing methods transform the input into an intermediate volumetric representation and construct a new mesh out of it [Ju04] [BPK05] [CW13]. In a new trend of methods specifically tailored for 3D printing, a 3D mesh is converted into an implicit representation, and all the subsequent operations (including the slicing) are performed on this representation [HWC13] [ice]. These methods are

very robust but necessarily introduce a distortion. Robustness and precision are indeed major issues in this area, in particular when self-intersections must be removed [Att14]. In this case some approaches rely on exact arithmetics [HKM07], while some others can losslessly convert the input into a finite precision plane-based representation, and then reconstruct a provably good fixed mesh out of it [CK10a] [WM13a]. When used for 3D printing applications, however, the aforementioned exact approaches are useful only if the input actually encloses a solid, while they are not really suitable to fix meshes with visible open boundaries [Att10]. For a more comprehensive overview of mesh repairing methods, we point the reader to [ACK13] and [Ju09].

Since 3D printing can only produce solid objects, a repairing algorithm must ensure that the resulting mesh actually encloses a solid 5. If the input mesh has open boundaries, a typical solution is to fill the holes in advance and then rely on some of the previously mentioned repairing methods. This approach is employed by one of the most popular web-based mesh fixing services [MN13], but it makes sense only if the boundaries are actually delimiting surface holes: in this case, recent techniques [JKSH13] can properly fill even complex holes with non-simply connected boundaries. Unfortunately, in some cases the designer uses zero-thickness surfaces to represent sheet-like features (e.g. a flag), and for models of this kind a hole-filling approach would produce rather coarse results. Another widely used software that performs mesh repairing for 3D printing is Autodesk's Meshmixer [Aut11], where the input STL can be successfully fixed even if it has open boundaries but at the cost of an overall approximation due to the global remeshing approach employed.
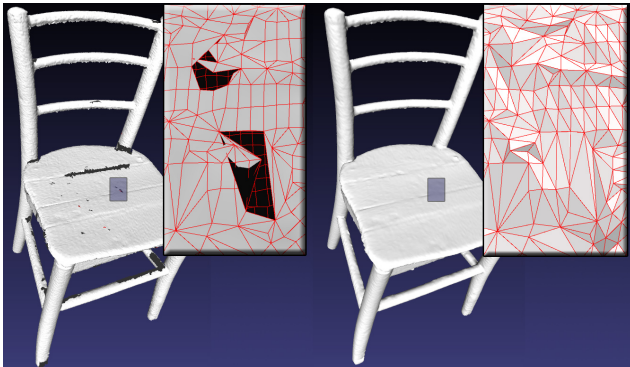

Figure 5: A raw digitized mesh may not enclose a solid due to various defects (left). Mesh repairing algorithms perform little modifications that turn such a raw model to an actual polyhedron that bounds a solid without ambiguity (right). Image courtesy of [Att10].

3.2. Orientation

The choice of the building direction is crucial in layered manufacturing as it directly influences the time necessary to print the object, the amount of support structures necessary to sustain the part during the print, the surface quality and its functionality. The first algorithms to select a proper part orientation date back to the

mid 1990's [AD94, RC95, FF95, CFN*95, LCCG97, HL98]. Several other methods have been proposed ever since, each one striving to optimize either for one, or a combination of the criteria discussed in Section 2.

Directly optimizing for the build direction in the space of all possible orientations is often too complex due to the non smooth nature of the metrics involved [EME15]. Former approaches used to consider a small number of candidate orientations, either predefined or computed on a shape proxy (e.g. the convex hull). Many recent methods (e.g. [WZK16, MCJ*16, EME15]) share a similar heuristic: they start with a regular sampling of the possible orientations; shortlist the orientations that perform best according to some quality metric, and, starting from each of them, navigate the space of solutions looking for the closest local minimum. The building orientation is eventually defined as the one corresponding to the lowest of the local minima explored by such heuristic.

Table 3 summarizes the properties of the orientation optimization techniques we discuss next.

3.2.1. Optimize for Cost

Recent works attempt to optimize for the orientation of the part using the volume of the support structures as only metric. Ezair et al. [EME15] showed that the resulting function is continuous but nonsmooth with respect to the orientation angles. In [EME15, KM06] two GPU-based volume estimation of supports structures are proposed. Morgan and colleagues [MCJ*16] define the support volume as the sum of the volumes of the prisms generated by extruding the down-facing triangles up to the building plate.

3.2.2. Optimize for Fidelity

Delfs and colleagues [DTS16] propose an orientation system that optimizes for surface roughness, using the mean roughness depth as a proxy to optimize for surface finish. One of the key features of the proposed approach is the ability to prescribe local accuracy requirements, attaching a target surface finish to each triangle in the tessellation. Masood and colleagues [MRI03, MRI00] proposed two systems that aim to find the building orientation that minimizes the volumetric error (Section 2.2).

In [HBA13] the authors address the problem of finding the proper build orientation for objects to be printed at low resolution via laser cut (cardboard or plywood). Their contribution is the definition of an optimal orthonormal frame that is suited for the decomposition into smaller parts, each of which to be sliced along one of the three directions with small volume loss error.

Other works focus on the artifacts that may be generated when detaching support structures. In [AKL07] the authors investigate how to orient the shape so as to minimize post processing (i.e. supports removal and surface finish). In [ZLP*15] Zhang and colleagues introduced a perceptual model to find preferable building directions in 3D printing, so as to place support structures in the least salient parts of the object.

3.2.3. Optimize for Functionality

Due to the layered nature of the process, the build orientation can significantly affect the performance of the resulting objects, intro-

ducing structural anisotropy. Ulu and colleagues [UKY*15] propose a FEM-based building orientation optimization that maximizes the minimum factor of safety (FS) under prescribed loading and boundary conditions. Umetani and colleagues [US13] propose a cross-sectional structural analysis based on bending momentum equilibrium. This is used in particular for orientation optimization. Their method avoids computationally expensive FEM simulations, and can be plugged into interactive modeling tools to allow users consider structural robustness during incremental trial-and-error design.

3.2.4. Optimize for mixed factors

In [PP12] and [TPR04] two genetic algorithms to find the optimal part orientation are presented. They both formulate the subject of their optimization as a weighted sum of multiple quality criteria, regarding both cost and fidelity. The weights can be finely tuned by the user to set the importance of each criterion. Similar approaches have been presented in [BL06a, BL06b], however these methods do not scale well with complex mechanical shapes (and non-mechanical shapes) as they consider a very restricted set of candidate orientations computed on the convex hull of the part.

Decision support systems to aid RP users choose the best building direction according to their needs are presented in [PDG99] and [HL98]. Both systems consider multiple criteria that can be prioritized according to the user needs, such as: overhang area, supports volume, build time and cost. On the negative side, these systems are specialized for CAD shapes and do not scale well on free-form shapes.

3.3. Support Structures

Support structures are a key component of process planning. They are used to compensate for some limitations of the manufacturing processes, in particular maximum overhang angles beyond which deposited material falls, and the large increase in time due to printing inner volumes of a part.

In the following we categorize supports into two main categories: disposable *external* supports that assist the fabrication process and are removed afterwards (Section 3.3.1), and *internal* supports that modify the inside of the object to achieve a trade-off between material cost, print time and physical properties (Section 3.3.2).

3.3.1. External

External support structures are sacrificial structures that are fabricated alongside the object. After fabrication completes they are chemically or mechanically removed. This usually involves human intervention and therefore is a time consuming, expensive step. An example of a complex part printed with and without support is shown in Figure 6.

Part orientation is a major factor in external support requirements, and therefore is often optimized to reduce the need for supports (see Section 3.2). Researchers have also proposed methods to slightly deform the design so has to reduce support [HJW15] and even to design models that are guaranteed to print without any
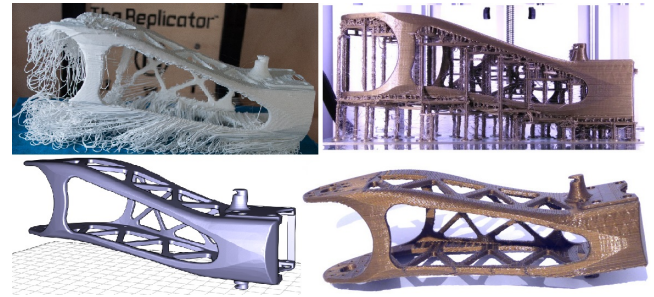

Figure 6: **Top left**: A robot upper leg (3D model below) printed without support. Filament falls due to excessive overhangs. **Top right**: A sacrificial external support structure, and the cleanup model (below). (3D model from the Poppy project https://www.poppy-project.org/en/, image from [DHL14]).

support [RL16]. However in most cases the process planner has to comply with the input model and some amount of supports remains necessary.

Different types of external supports serve different purposes:

- Local deposition technologies can only deposit material on existing surfaces below. Thus, surfaces appearing mid-air and surfaces at an excessive overhang require support just below them.
- Shapes may move or deform during the fabrication process. This typically happens when fabricated objects are imbalanced and when the raw material (powder, resin) cannot sustain the weight of the print. Another source of distortion are stresses from thermal gradients. To reduce these issues supports acting as fixtures are necessary.
- Some processes can generate a large quantity of heat, in particular metal printing. This excessive heat accumulation results in shape distortions and residual stresses. In such a case, additional supports may act as heat diffusers.

We next review approaches from the literature and discuss their use for each type of support. First, however, let us take a closer look at when and why supports are required.

Islands, overhangs, and self-supporting surfaces. A first situation that requires support is illustrated in Figure 7, left. After dividing the shape into layers, one of the slices contains an island – a solid region that appears while not being supported from below. During additive fabrication from bottom to top, the material forming the island will not attach to already solidified material. For technologies using material deposition, the material forming the island will fall. As a consequence, material deposited onto the next layer above will also fall, and this will cascade into a catastrophic failure. For technologies using layer solidification, non solidified material below (e.g. powder) will usually be able to support the island – however weight might accumulate over several layers and a heavy disconnected component may start to sink down. For technologies using resins (SLA), the island is problematic: it will typically end up floating in the viscous resin, cascading into a complete print failure.

Overhangs also can produce problematic cases with material de-

METHOD	OPTIMIZES FOR					
	Cost		**Fidelity**			**Functionality**
	Slice number	Supports volume	Cusp height	Volume loss	Surface finish	Stress resiliency
[DTS16]	○	○	○	○	●	○
[MCJ*16]	○	●	○	○	○	○
[WZK16]	○	●	●	○	○	○
[EME15]	○	●	○	○	○	○
[UKY*15]	○	○	○	○	○	●
[ZLP*15]	○	○	○	○	●	○
[US13]	○	○	○	○	○	●
[HBA13]	○	○	○	●	○	○
[PP12]	●	●	●	○	●	○
[AKL07]	○	○	●	○	●	○
[BL06a]	●	●	●	○	○	○
[BL06b]	●	●	●	○	○	○
[TPR04]	●	●	●	○	○	○
[MRI03]	○	○	○	●	○	○
[MRI00]	○	○	○	●	○	○
[PDG99]	●	●	●	○	○	○
[HL98]	●	●	●	○	○	○

Table 3: Techniques for orientation optimization and their properties, see Section 3.2. Legend: ●: yes; ○: no

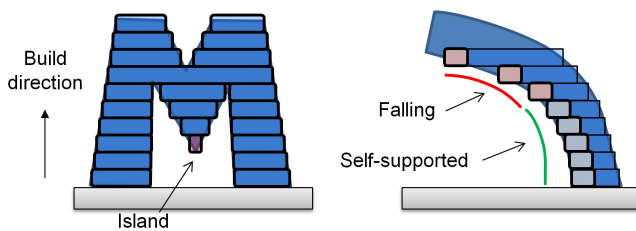Figure 7: **Left**: The M letter is decomposed into layers. The downward facing tip becomes an island with respect to the build direction. **Right**: Each layer is solidified from the outside towards the inside. The hatched rectangles are the cross sections of material deposition paths. After a critical angle, the deposited material no longer bonds to the layer below and falls.

position. The material is typically added progressively along deposition paths. The right illustration of Figure 7 shows a cross section of the deposited paths for an overhang region. Due to the layering, at some excess overhangs the deposited material will simply fall (red rectangles in the Figure). It is interesting to note, however, that until some threshold angle the deposited material will have a sufficient bonding surface with the layer below. This self-supporting property stems from the bonding between successive layers and allows for overhangs to exist up to some maximum angle without requiring support. Overhangs are less of an issue for technologies employing layer solidification – even though excessive overhangs may distort due to change in material properties during solidification, or to auxiliary motion during preparation of the next layers.

Detecting surfaces requiring support. Generating supports for areas in overhang requires two main steps: the detection of the surfaces in need for support, and the generation of the support structure itself. For detecting surfaces in overhang a first family of approaches consider the down-facing facets of the input mesh having an angle too steep to print correctly (e.g. [KJAB*91, AD95]). A second family of approaches consists in performing a boolean difference between two successive slices (e.g. [ACC*88, CJR95]). The width of the difference determines which regions are self-supporting [CJR95]. This can also be detected conveniently and efficiently with two-dimensional morphological operations in image space [HYML09, CLQ13, HWC14]. As discussed in [HWC14] care must be taken however with some specific configurations, where some protruding regions might be mis-classified as supported (see Figure 1.8 in this publication). [DHL14] performs the detection directly at the toolpath level, verifying whether each deposited segment is supported by at least half its width from below.

This first analysis generally leads to a compact set of points to be supported. Several approaches then select a subset by downsampling [ER07, CLQ13, DHL14, HWC14].

Generating a support structure. Once the surfaces requiring support are determined, the support geometry is computed. The main trade-off is between print time, material use, and reliability.

For instance, for filament printers the traditional approach consists in extruding the mesh facets requiring support downward, thus defining a large *support volume*. The support volume is usually printed with a weak infill pattern (see also Section 3.3.2). This still uses a significant amount of material and time, but it is very reliable: the support typically has a large area of contact with both the part and the print bed, ensuring the print stability in most cases. In this context several approaches modify the support volume to reduce its size. Huang et al. [HYW*09] use sloped walls instead of straight walls for the sides, shrinking the support volumes in their middle sections. Heide [Hei10] reduce the support volume by de-

creasing its size and complexity as the distance below the supported model increases.

Other 3D printing technologies can print complex and thin structures more reliably. In the context of SLA, Eggers and Renap [ER07] form a support structure by starting from a regular rhombus mesh filling the print bed. The 3D model is subtracted from the initial structure, removing intersected mesh edges. Points requiring support are attached to the mesh by downward angled beams. Huang et al. [HWC14] produce a support made of a sparse set of vertical pillars connected by angled beams for structural strength. The position of the pillars is optimized and their pairwise connections follows a minimal spanning tree (as seen from above) to keep the support structure small. *MeshMixer* [SU14] builds a thin structure supporting the part in a sparse, limited number of points, generating a support structure resembling a tree. This approach has been used successfully on both FDM and SLA machines. Vanek et al. [VGB14a] propose an algorithm to optimize for similar tree support structures. Wang et al. [WWY*13] optimize truss structures for the primary purpose of strengthening 3D printed objects, and extend their approach for support generation. Support beams are added by tracing rays downwards. Dumas et al. [DHL14] generate bridge scaffoldings that rely only on vertical and horizontal bars, improving reliability and part stability while generating small support structures. Remarkably, the horizontal bars can be printed efficiently on FDM printers, but are also well suited for SLA. The static shape balance at all stages of the process is taken into account, enlarging the support structure whenever necessary. Calignano et al. [Cal14] discuss the design of support structures for SLM (selective laser melting) where supports act both as fixtures and heat sinks.

Support removal. External support structures are meant to be detached from the part after the print is completed. Depending on the materials involved (e.g. printing on metal) this operation can be extremely challenging. Furthermore, residual material can remain attached to the surface, badly affecting the surface finish. In polymer AM processes, to alleviate the supports detaching problem, thermoplastic materials that dissolve in alkaline baths are used [PJB04]. Hildreth et al. [HNCS16] have recently shown that similar approaches are possible even for stainless steel printing, were the differences in the electrochemical stability between different metals can be exploited to dissolve carbon steel supports. Jhabvala et al. [JBAG12] exploited the pulsed laser radiation to print metal supports that are both faster to print and easier to remove; the system supports only SLM printers. A valid alternative consists in trying to orient the part in such a way that the supports necessary to sustain it during the print will stick only to the least salient portions of the shape, as proposed in [ZLP*15] (Section 3.2.2).

3.3.2. Internal supports

The interior of an object is a key factor regarding the material use, print time and mechanical properties of the final result. The impact on material use and print time is easily explained by the fact that the inner volume grows to the cube of the scaling factor (i.e. doubling the size of an object multiplies its volume by eight). Therefore, most of the time and material is spent on the inside of the object. Carving the inside can lead to large savings. However, care must

be taken to ensure that the object can still be fabricated (e.g. avoiding the introductions of overhangs or islands, not forming pockets) and that the end result will remain rigid enough. In addition, modifying the structure of the inside of an object gives the opportunity to change its global mechanical behaviour, for instance making it flexible or rigid in different places, or changing its balance.

We organize this section as follows. We first consider techniques that focus on creating large empty pockets within objects. Next, we discuss approaches that focus on how to infill the object interior, both with dense and sparse patterns. These techniques have an emphasis on material/time savings and exploit specificities of the processes. Then, we focus on frame structures, which are typically beam or cellular structures optimized to create a strong structure within the object interior. We discuss techniques that focus on changing the mechanical behaviour of the object by filling its inside with micro-structures. Since this is a very large topic, we keep the focus on techniques that tightly integrate with the process planning pipeline.

Hollowing. Most of the techniques we discuss here treat the general geometric problem of computing an inner cavity at a fixed distance from the object surface. The surface of the inner cavity is called the *offset surface* of the model [Far85, RR86]. These approaches can be used to compute inner cavities that are either left empty or filled with the some infilling pattern, as described in the following paragraphs.

Early approaches obtain a superset of geometric primitives of the offset surface that are trimmed and filtered to form the final offset boundary [For95]. Qu and Stucker [QS03] presented a vertex convolution method for STL files without explicit treatment of self-intersections. Campen and Kobbelt [CK10b] introduced an exact convolution approach. Hollowing can also be performed by computing the *distance field* of the model, and extracting the offset surface from it. Frisken et al. [FPRJ00] presented and adaptively sampled distance field. Varadhan and Manocha [VM06] approximate the offset surface with a distance field isosurface extraction, guaranteeing a Hausdorff distance bound on the approximation. Pavić and Kobbelt [PK08] traverse an octree distance field and split each cell which is potentially intersected by the offset surface. Liu and Wang [LW11] extract the isosurface of a narrow band distance field.

Another class of hollowing methods consider ray representations of solids such as the *dexel structure* [Hoo86] or the layered depth normal images (LDNI) [CW08]. For a single direction and a uniform grid of rays parallel to that direction, a ray-representation stores the intervals of the rays lying inside the solid. Hui [Hui94]
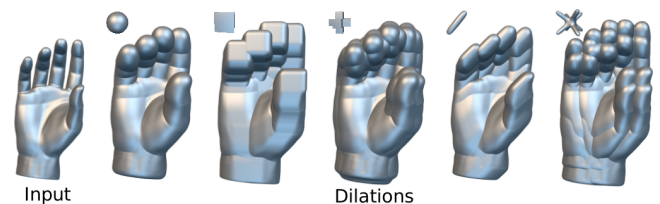

Figure 8: Morphological dilation of an input point cloud, considering different structuring elements. Image courtesy of [CB14].

computes the sweeping of a solid along a trajectory by considering the union of a finite set of ray representations of the solid. Chiu and Tan [CT98] computes the morphological erosion of each dexel, taking into account its neighbourhood dexels. Hartquist et al. [HMS*99] compute the union of spheres over the boundary of the input model. Wang and Manocha [WM13b] place spheres on the samples of a LDNI structure, and compute their union efficiently on the GPU. Chen and Wang [CW11] generate a convolution geometric primitives of the offset surface, constructs their LDNI, and filters the points of the superset that belong to the offset surface. Martínez et al. [MHCL15] consider the dilation of a dexel structure along different successive ray directions, and exploit the winding number of specially constructed meshes. Other methods generate a voxelization of the offset surface. Li and Mc-Mains [LM14] present a GPU approach to compute the Minkowski sum of a polyhedra by computing pairwise Minkowski sums, and obtain a voxelization of their union. Hollowing can also be done by considering the offset of surface points. Lien [Lie08] computes the Minkowski sum between two surfaces sampled by points, and distinguishes the interior and boundary points. Calderon and Boubekeur [CB14] introduced a set of morphological operations for point clouds (see Figure 8).

A few methods hollow the model during slicing, that is at slice level. McMains et al. [MSWS00] consider regularized boolean operations of each slice contour, in order to approximate the offset surface. In order to achieve uniform hollowing thickness, Park [Par05] considers the erosion of a circle swept over the slice contours.

We focused here on hollowing at the process planning stage. As mentioned in Section 2.3 hollowing can also be used at design time to change the mass distribution of an object and optimize for various properties such as balance.

Dense infills. Most technologies such as SLS, SLA, binder deposition and fused filament deposition support dense infilling. On raster devices, it suffices to produce an image of the filled layer contour. On vector devices, densely filling the object requires to follow a space filling curve when solidifying the material. The curve represents the path followed by the deposition device, while it deposits (or solidifies) a wide and thick track of material. The thickness matches the layer height, while the width depends on the technology (typical values are $20\mu m$ for focused lasers, $400\mu m$ for plastic extrusion). Thus, the spacing between two neighbouring paths has to match the deposition width. A smaller spacing produces excess deposition/solidification (*overflow*) while a larger spacing leaves gaps (*underflow*).

The shape of the space filling curve can have an impact on both the print time and the final object strength. Some possible patterns are illustrated in Figure 9.

For most technologies – including fused filament deposition – a popular pattern is the so called *direction parallel* (or *zig-zag*) toolpath [MSWS00], which consists in filling the slice area with a set of equally spaced segments parallel to one another and linked at one of their extremities. The spacing between the hatches is adjustable to control the final density, and the direction of the hatches changes every two slices to avoid strong mechanical biases. To gain

speed or save material, the parameters used during infilling may be different, e.g. reducing binder flow, or changing the laser focus to solidify larger tracks.

Contour parallel infill [YLFW02, HLA94, KP98] is a common alternative to direction parallel infill. The contour parallel pattern consists in a set of concentric closed curves that emanate from the outer boundary of the slice and propagate inwards (see Figure 9). These are in fact the offset contours from the layer boundary. This is an appealing pattern since it closely follows the outer contour of the slices. Unfortunately it also tends to leave gaps within the slices. Indeed, the offset curves from the outer contour meet around the medial axis, and it is unlikely that the remaining space matches the deposition width. These gaps are difficult to handle. For this reason most slicers rely on a hybrid approach that combines direction parallel and contour parallel [MSWS00]. A few contour parallel curves are produced near the boundary of the slice, and the remaining inner polygon is filled the interior with a zig-zag pattern. The most critical part in the implementation of such hybrid approaches is the handling of the meeting point between the contour parallel and the direction parallel toolpaths, in which detaching, under and overfilling may occur [JHF13].

Infill patterns can lead to two issues: First, they may result in many stops and restarts of material deposition, which can lead to several problems (e.g. many small gaps, material feeder failure). Second, sharp turns may slow down the motion or introduce vibrations. For instance, direction parallel toolpaths can lead to many small segments if the slice contains thin walls orthogonal to the deposition direction – inducing a large number of turns. The speed of different dense infill strategies have been compared, including variants that smooth out the sharp turns between hatches [KC02, EMET06]. To obtain a more continuous pattern, Ding et al. [DPCL14] decompose the slice in different regions that use different directions of parallel infill. The patterns from one region to the next are smoothly connected. Zhao et al. [ZCG*16] advocate for the use of spirals. A key advantage is to reduce the number of sharp turns thereby enabling faster motions, while achieving a fill pattern that resembles the contour parallel infill.

Other alternatives to these dominant patterns are Hilbert spiral [Gri94], Moore's spirals [CTF*94], Peano spirals [MA07]. A common drawback for both spirals and contour parallel patterns is the lack of *direction bias*, which prevents the generation of the cross weaved layouts used by direction parallel approaches.

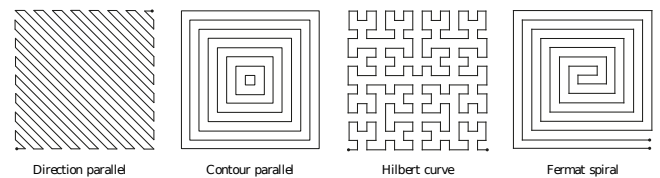| Direction parallel | Contour parallel | Hilbert curve | Fermat spiral |

Figure 9: Four different filling patterns used in additive manufacturing. From left to right: direction parallel, contour parallel, the Hilbert space filling curve and the recently proposed Fermat spiral.

Sparse infills. Since the cost of additive manufacturing is essentially driven by time and material use, sparse infilling is an important feature. The geometry of a sparse infill is often determined

by the target process. In particular, powder and resin systems (e.g. SLS/SLA) cannot create closed voids: non solidified material is trapped and cannot exit, the sparsity would be lost. However, they can print geometries that are significantly more complex than, e.g. fused filament fabrication. This has led researchers to propose internal frame structures as well as micro–structures, that we both describe in the following paragraphs.

Fused filament fabrication (FFF) and similar technologies (e.g. contour crafting) require dedicated infill patterns, due to the strong overhang constraints. In addition, the thin slanted beams that are used with other processes are slow and less reliable to print. Most slicing software for FFF supports sparse infills. These infills are usually 2D zig-zag hatching patterns that are vertically extruded [DGLC15], again possibly changing orientation every two slices to avoid strong mechanical biases. [KPA09] explores hierarchical versions of space filling curves to grade the density of the infill pattern.

These patterns are efficient and simple to compute during slicing, and offer a good support for the roofs of the inner cavity. However, due to the vertical extrusion they are not mechanically strong if pressured on the sides. To obtain stronger patterns, Steuben et al. [SIM16] define the infill as the iso-contours of a scalar field, or the principal directions of a vector field within each slice. For instance, the infill paths can follow the principal directions of stress from a finite element simulation, resulting in stronger patterns for a given load scenario.

There has also been several attempts to move beyond simple vertical extrusions. 3D printing enthusiasts have experimented with interesting 3D infill patterns [Nic11]. The software *Slic3r* proposes an infill pattern that produces a 3D honeycomb pattern. This recently led to a new type of pattern called *rhombic infill*. These are formed by the intersection of at least three sets of parallel planes in space. These infill have a number of interesting properties. First, they can be efficiently generated during the slicing process, and by carefully choosing the angles of the inner planes they can be printed as fast as vertical extrusions [Lef15]. Second, when printed with uniform density they are very strong thanks to the inner 3D cell structures. Third, they can be subdivided to locally increase the infill density, as shown in Figure 10, right. Wu et al. [WWZW16] propose a criterion based on the overall rigidity and balance of the object to perform this subdivision. Lee and Lee [LL16] subdivide closer to the cavity roofs in order to create large empty cavities. They further reduce the size of the structure by removing faces that are not required to support a structure above.

In an attempt to make infill as sparse as possible, Hornus et al. [HLDC16] propose a method that creates maximal inner carvings while ensuring that they remain fabricable with filament deposition (the cavities are self-supporting). This is achieved through morphological operations on the slices, growing a self-supporting inner cavity from the object tops.

Internal frame structures. As mentioned earlier, frame structures are especially well suited for SLA/SLS technologies, even though they have also been successfully demonstrated on fused filament fabrication. Wang et al. [WWY*13] propose to fill an object interior with a sparse truss structure. The structure is optimized to
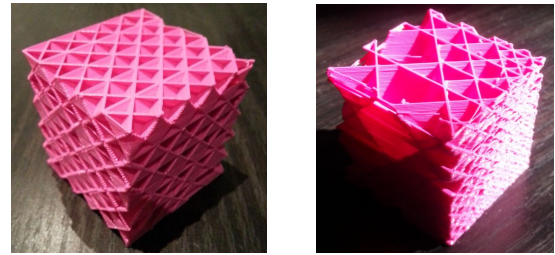

Figure 10: Rhombic infill and its hierarchical version [Lef15].

reduce the number of beams while preserving rigidity. Zhang et al. [ZXW*15] similarly produce an inner structure made of beams but instead exploit the medial axis of the object, using it as a backbone structure. Medeiros et al. [MeSMEC15] generate an adaptive tessellation of the interior, and offset the edges of either the primal or the dual to produce an inner beam structure. Thanks to the adaptive tessellation, the structure is denser along the shape boundary than on the inside. Lu et al. [LSZ*14] optimize for a Voronoi diagram inside the print, which faces form an infill pattern.

A difficulty in designing 3D infill patterns is that their complexity might lead to increased print time, for the same density (this is less true on systems such as SLA/DLP where the entire layer is exposed at once). Yaman et al. [YBSH16] consider how to print efficiently the faces of a Voronoi diagram, following an Euler cycle to solidify in sequence the segments forming the faces in each slice.

Microstructures. Microstructures are internal infill patterns that seek to change the macroscopic physical behavior of the final object. For instance, even when printing with a single material certain microstructures modify the elastic behavior of the object, making it more or less flexible. It is often possible to grade and control the change across the final object.

The design of microstructures with tailored properties was introduced in the 1990's [Sig95, CZCL13]. A large range of techniques deal with optimizing functional microstructures, such as functionally graded materials for CAD applications [JLP*99, KT07, Oxm11] and porous scaffold design for bioengineering [Hol05], among many other applications. A complete review of the field falls out of the scope of this document. Instead, we focus here only on the techniques working in conjunction with the process planning – that is methods within the scope of design for additive manufacturing [Ros07]. In particular, as the size of microstructures becomes smaller, approaches explicitly storing the microstructure geometry become computationally infeasible.

Chen [Che07] defines micro-structures as periodic tiles that are then efficiently mapped into the volume, similarly to volume texture mapping. By deforming the mapping, the infill locally adapts to a density field. Pasko et al. [PFV*11] considered procedural definitions of periodic microstructures. The parameters of the microstructures can vary spatially to produce graded materials [FVP13], for instance to reinforce an object following a cross-sectional stress analysis [LDJC15]. OpenFAB [VWRKM13] provides a specialized language to describe procedural microstructures. The geometric details are efficiently evaluated at slicing time,
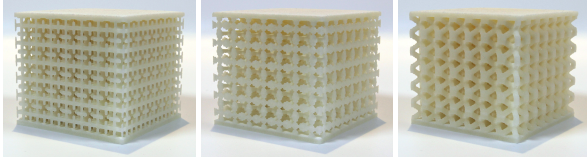
Figure 11: Printed microstructures with a precomputed elastic behavior. Image courtesy of [SBR*15]
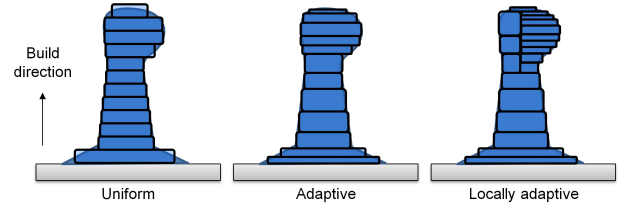

Figure 12: A same object sliced with different approaches. **Left:** uniform slicing (12 slices), **middle:** adaptive slicing (12 slices, volume error is reduced), **right:** locally adaptive slicing, where the top region is split into two sub-regions sliced independently.

streaming voxels to the printer. The advantages of procedural representations for the process planning have been identified early by Park and colleagues in the context of multi-material fabrication [PCB01] (see also Section 3.4.2).

Different works seek to produce microstructures that can be fabricated [ZL08, ALS14] and produce a prescribed elasticity. [SBR*15] and [PZM*15] consider periodic tilings of precomputed microstructures that cover a large spectrum of elastic behaviours (see Figure 11). [MDL16] considers procedural Voronoi-based microstructures, that can be fabricated with SLA/SLS.

Micro-structures are a very promising field of research, with many potential applications [RL15]. We envision that efficient techniques for designing and fabricating micro-structures will tightly integrate with process planning.

3.4. Slicing

Slicing is central to the process planning pipeline, as it is the step where the 3D geometry is divided into a set of planar contours. These contours will be later manufactured by material deposition. In the following we assume that the build direction is along the z axis, aligned with the height of the object. Each slice is a plane intersecting the shape at a given height. Assuming the shape is a solid, then its intersection with a slice plane is a closed 2D contour.

There are two important questions to solve when considering slicing: how to determine the set of slices to use and their vertical position, and how to efficiently compute the contour within each slice given the input shape and the set of slices. The two following sections discuss each of these steps.

3.4.1. Uniform and adaptive slicing

The simplest approach to divide the object into slices is to subdivide it uniformly, as illustrated in Figure 12, left. Given a manufacturing layer thickness τ and an object height H, the object is divided into $N = \lceil \frac{H}{\tau} \rceil$ slices. Each slice i is then located at height $z_i = \frac{i+0.5}{N}$, which is the position of the plane that will be intersected with the object. This approach is widely adopted and most softwares offer it as a standard approach.

However, many objects have a shape that varies greatly along its height. Therefore, in some regions uniform slicing might use too many slices, while it does not properly capture the shape in others. In particular, surfaces that are slanted with respect to the build direction produce a staircase defect through additive manufacturing. These areas require to use very thin slices (small τ). Uniform slicing forces the same small value of τ to be used throughout the

entire part, thus producing a large number of slices and increasing manufacturing time.

Most technologies are able to change the layer height during manufacturing. *Adaptive slicing* approaches exploit this property, by adapting the thickness of each slice to the shape geometry, as illustrated in Figure 12, middle. Given a model for the geometric error (see Section 2.2), these approaches refine or coarsen the slices to meet a quality constraint while reducing print time. This can be achieved by locally determining the slice thickness from the error [DM94, SW*94], by subdividing the slices from the coarsest uniform set of slices [SHB96, KD96, HRJ97], by merging slices starting from the thinnest uniform slices [HA13], or by formulating a global optimization problem [WCT*15, SBK15].

While adaptive slicing is able to adapt to shape changes along the build direction, it still cannot adapt to a change in part complexity *within* the layer. Consider an object with a vertical wall on the left, and a slanted surface on the right. The vertical wall could be printed with thick slices, however the slope on the right imposes the use of thin slices to limit the staircase effect. To reduce this issue, *locally adaptive slicing* has been proposed [TB98]. The key idea is to first subdivide the object into different regions, each region being sliced independently, as illustrated in Figure 12, right. Depending on the target technology, the different regions can be built together by locally changing the layer height [TB98, WCT*15], or they can be printed independently and later assembled [HBA13, WZK16]. The main issue with this approach is that seams appear along the surface where different layer thicknesses meet. For techniques where parts are printed separately a manual assembly step is required.

Interestingly a similar approach was used on the object interior by [SHB97]: Since the inside is never visible, it can be sliced using a larger thickness than the exterior. This idea can be combined with the aforementioned techniques, for instance performing local adaptive slicing only on the exterior shell while the interior uses the maximal thickness [MKD99].

3.4.2. Slice contouring

Once the set of slices is determined, each slice plane has to be intersected with the input geometry. This operation strongly depends on the representation of the input. We first discuss slicing of triangle meshes and ray-representations (ray-reps). Both are well studied and successful approaches, which are often described as *indirect* approaches : the input CAD model has to be converted into a tri-

angle mesh (tessellated) or a ray-rep (ray-tracing or rasterization). Both conversions require the user to set a precision parameter and may loose information. We discuss these issues in more details below. Therefore, a number of *direct* slicing techniques have been proposed, that avoid any re-sampling of the initial CAD model. We discuss these contouring techniques last.

Contouring triangle meshes. A general scheme for contouring triangle meshes consists in first extracting all intersection segments between the slice plane and the triangles, and then forming loops [KJA92]. If the input correctly defines a non self-intersecting solid, the loops will be closed and non-intersecting. Otherwise, a mesh repair step is required (see Section 3.1.2). Alternatively the slicer may attempt to close holes between nearby segments and resolve intersections.

Implementations mainly differ by how segments and loops are formed. Kirschman and Jara-Almonte [KJA92] propose a parallel implementation that intersects each slice plane with all triangles. McMains and Séquin [MS99] propose an efficient algorithm that exploits the mesh connectivity to sweep a slicing plane through the triangles. The observation is that the topology of the 2D contours remains the same between vertices of the input mesh. Therefore, the contours can be very efficiently produced for all slicing planes in between two vertices. The update to be performed at each mesh vertex is often limited and fast. The open source software *CuraEngine* traverses triangles first, and each is intersected by the slicing planes it covers (see `Slicer::Slicer` and `project2D` in `slicer.cpp/.h`). The segments are identified by the faces to which they belong, and contours are formed by looping over segments following mesh connectivity (see `makeBasicPolygonLoop` in `slicer.cpp`). Zhang and Joshi [ZJ15] argue that mesh connectivity might be expensive to obtain and propose to incrementally construct the contours while triangles are traversed. Linked lists of segments are augmented by adding the new segment to the head or tail (or starting a new list). The segments are not identified directly by intersection points, but instead by the triangle edge to which they belong – making the approach robust to numerical errors in intersection computations.

Contouring of ray-representations. Ray-reps are techniques were the geometry is captured by solid/empty intervals along a set of rays. This technique was pioneered by Hook [Hoo86] who proposed the *dexel-buffer*. This data-structure is built by intersecting axis-aligned rays along one direction with the geometry. Given a closed geometry, the number of intersections is even and each interval can be classified as inside (solid) or outside (empty). The dexel-buffer is closely related to the A-buffer [Car84]. There are therefore interesting similarities between ray-reps and A-buffer techniques for order independent transparency [MCTB11]. For instance, dexel-buffers can be constructed by rasterization, recording all fragments drawn in every pixel [LHL14]. Other efficient approaches are based on Layered depth images (LDI) [SGHS98], which capture the geometry as a set of depth images each storing a surface sheet in space (a *depth peel*). This data structure has been extended to solid modelling for additive manufacturing by Chen et al. [CW08, CW08], adding normal information along with depth information (layer depth *normal* images, or LDNI). Fast
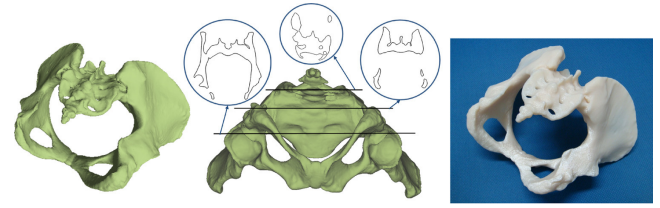

Figure 13: Slicing and contouring of a ray representation LDNI while preserving its topology. Image courtesy of [HWC13].

construction methods have been proposed, e.g. using single pass voxelization [LW13] and compaction for increase memory efficiency [ZWCJ11]. Most ray-reps used multiple directions to better reproduce geometries [BM97].

Once obtained, ray-reps may be directly rendered [Hoo86, WLC10, Lef13] or converted into meshes through efficient procedures [ZL09, WLC10, Wan11]. However, they may also be directly contoured to extract slices. Different strategies have been used: single set of rays from the object side [ZY01, ZLQ*11], rays from two sides [QZY13], and rays from the object bottom [Lef13]. Contouring is then performed by marching along the rays, forming polygonal loops [ZPLZ07, YZRL08]. A key issue when using ray-reps for contouring is to decide upon the resolution required to properly capture the geometry and its topology [HWC13] (see Figure 13). The contours extracted from ray-reps have typically many small segments – each slice is an image and contours are extracted as the outline of the solid pixels. Huang et al. [HWC13] describe a topology preserving contour simplification, and a full image based pipeline for additive manufacturing [HWC14]. Another difficulty is the large memory requirements, which is roughly proportional to the surface area. To avoid saturating memory, tiling schemes may be used [CW13].

Ray-reps have other significant advantages for additive manufacturing, for instance to perform CSG between complex geometries [BM97, WLC10, Lef13], for regulating solids having self-intersections [CW13] or for computing offset surface for hollowing parts (detailed in Section 3.3.2).

Direct slicing. To avoid having to re-sample the CAD model into a triangle mesh or ray-rep, several techniques have been proposed to extract contours directly from the initial geometry. These approaches directly output slice data to the printers. The slice file format is often vendor-dependent, but some independent formats such as CLI (Common Layer Interface) can be used. Open source printers often accept slice data (e.g. G-Code for filament printers or images for DLP printers such as the B9Creator or Autodesk Ember). Jamieson and Hacker [JH95] provide an in-depth discussion of the pros and cons of direct slicing of CAD models for different input types.

Non-uniform B-spline surface (NURBS) are a common surface representation in CAD software. Therefore, approaches have been proposed for the direct slicing of NURBS models in order to avoid a global tessellation of the geometry [VKF*92]. Later approaches consider specialized adaptive slicing [MBH04] and orientation procedures [SLS*05].
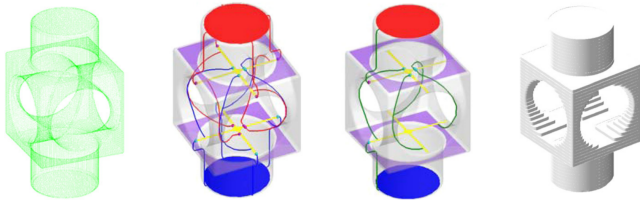
Figure 14: Direct slicing (right) of a point cloud (left) via topological analysis (middle). Image courtesy of [YLQ10].

Figure 15: Printed 3D wireframes, using a 5DOF printer. Image courtesy of [WPGM16].

Techniques have also been proposed for point clouds, which are often obtained from 3D scanners or vision algorithm. They are challenging to print since the connectivity and topology of the surface is unknown. Early methods project points around the 2D slicing plane and reconstruct a contour [LWZL03, WWLZ04, SPP04]. Yang and Qiang [YQ08] propose to rely on the moving least square method to implicitly define the surface from the point cloud. Qiu et al. [QZQ11] and Yang et al. [YLQ10] refine this approach by considering the global topology of the shape to detect and capture extremal points in an adaptive slicing strategy (see Figure 14). Chen et al. [CLQ13] propose a complete system for scanning an object and fabricating it in a different location, based on point clouds. This includes a novel support generation algorithm.

Finally, Rosen [Ros07] proposed to apply direct slicing on microstructure lattices, thereby bypassing the mesh generation step. The lattice beams are directly interesected by the slice plane, and thus only the lattice graph is stored in memory. This makes slicing much more efficient and accurate on such objects.

Non-planar approaches. There has been a number of interesting attempts at moving beyond standard layers for additive fabrication. These attempts usually focus on a specific technology since they exploit properties such as the ability to perform Z-motion during deposition (fused filament fabrication) or the ability to partially cure material (SLA/DLP).

In the context of stereo-lithography (SLA) Pan et al. [PZZC12, PC15] exploit the formation of a meniscus when an object moves out of the resin tank. The meniscus is cured to fill the creases between two layers. Repeating this process produces smooth, accurate surfaces. Park et al. [PHL11] show how dithering can cure resin partially and produce slanted surfaces along a layer.

In the context of filament deposition, Chakraborty et al. [CRC08] proposed curved layer deposition with the objective of strengthening shell-like parts by aligning the toolpaths with the surface. The mechanical properties of the parts are discussed in [SRDH12]. An interesting question is then to combine flat and curved layers, as discussed in [HS12, AT15]. A key challenge of curved layers deposition is that the curvature of the paths is large as they flow along the surface several millimetres up and down. This has been demonstrated for specific parts, however in a general setting this degree of freedom is challenging to exploit. First because a novel type of slicers have to be developed, and second because the current design of deposition nozzle complicates the task: collisions between already printed paths and currently deposited paths become possible [CRC08, AT15].
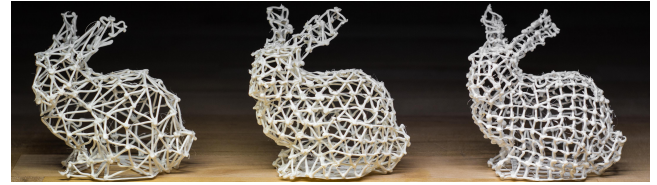
Beam and truss structures are not very well suited for layer by layer fabrication, as the beams are sliced into many small cross-sections. Mueller et al. [MIG*14] propose a dedicated approach to print wire-mesh structures, exploiting the fact that extruded filament hardens quickly to print truss-like structures in mid-air. This raises many challenges regarding path planning of the extrusion device. Wu et al. [WPGM16] address some of these challenges using a 5DOF printer (see Figure 15): two additional degrees of freedom allow the platform to rotate. The path planning problem is formulated as an ordering of the edges of a graph that captures the spatial constraints of extruder motions. Huang et al. [HZH*16] address a similar problem, using a 6DOF robotic arm with a customized extrusion head. In addition to collision avoidance their algorithm considers the stability of the printed part during the whole fabrication process, ensuring that each layer is in a stable equilibrium over its preceding layers.

With the objective of printing hair, fibers, and bristles Laput [LCH15] introduced a technique that exploits the stringing phenomena of filament deposition technologies.

Multiple materials. Some technologies afford for multi-material fabrication. Weiss et al. [WMP*97] describe a process for fabricating multi-material objects. Kumar et al. [KD98, KKD98] describe a modeling representation and adaptive slicing algorithm for multi-material objects. Zhu and Yu [ZY01] propose a dexel slicer for multi-material objects. The solid ray intervals are used as solidification paths. Park et al. [PCB01] describe a system for modeling and fabrication of multi-material objects inspired by procedural volume texturing approaches in computer graphics, noting the advantages for the process planning in terms of memory compactness and resolution independence. Shin et al. [SNDM03] focus on process planning for fabricating objects with continuously varying material properties on a direct metal deposition system.

Machines based on filament extrusion can mount multiple extrusion heads, in which case motion planning during deposition becomes more complex. Choi and Cheung [CC05] input a triangle mesh per material (packed in a single colored STL file) that are sliced independently. The contours in each layer are grouped by inclusion order, so as to reduce redundant motions during deposition. This is later extended to motion planning for multiple, independent deposition devices. The challenge is to coordinate the movements of the nozzle depositing different materials, so as to avoid collisions or even deadlocks [CC06, CZ10]. In the context of computer graphics, Hergel and Lefebvre [HL14] also consider the case of slicing and toolpath planning for multiple materials. Instead of inputing multiple meshes, materials are selected by a *slice shader* which is

executed on every point of each layer. Contours are extracted using an image space approach (see Section 3.4.2). The motion planner is optimized for visual quality, hiding potential defects in less visible regions of the part. Reiner et al. [RCM*14] achieve a visual grading of colors on multi-filament printers, interleaving the deposited filaments in sine wave patterns along the surface (see Figure 16).

Some technologies apply colors on entire layers (inkjet on powder [SHCW94] and inkjet on laminated paper [Mco05]). At slicing time, a raster RGB color image is applied to the layer contour. Since the colors are often specified from the surface (e.g. vertex colors or 2D texture map), they are propagated inside from the contour within a thin shell. However, the technology would be able to color the entire volume, even though material opacity limits the potential.

Multi-jet technologies – the deposit droplets of different resins – have a wider variety of materials and colors. Open-FAB [VWRKM13] allows to models complex multi-material geometries with voxels, and streams slices to a high-resolution multi-jet printer. This forms the basis of a complete modeling system for multi-material modeling [VKWM16]. Wu and colleagues [WSP*00, CSP*01] discuss process planing to convert continuously varying material information into a limited set of base materials by a half-toning technique. Brunton et al. [BAU15] also rely on half-toning in the context of color reproduction – to the point that print-out of scanned objects can be confused with their model. Earlier work also studied how to control optical properties of the final print by combining multiple base materials [HFM*10, DWP*10] – even though these approaches are not integrated within the process planning pipeline.

3.5. Machine instructions

When talking about machine instructions it is necessary to distinguish between machines that operate on each slice like a plotter, that is, connecting pairs of points with straight lines, and machines that operate on each slice like an inkjet printer, that is, interpreting the whole slice as a discrete 2D image. The former require the slices to be defined in vector format, whereas the latter require the slices to be defined in raster format (see also Section 1.3).

3.5.1. Vector case

In the vector case the machine instructions amount to a piecewise linear toolpath (in some cases arcs) along which the printer must deposit, melt or sinter the printing material. The machine toolpath

Figure 16: Producing continuous tone imagery by color mixing on multi-filament printers. Image courtesy of [RCM*14].

must be prepared for each slice, both for its outer contour and the interior. If the part comes in the form of a boundary representation (e.g. a STL file), particular attention must be paid to distinguish between the interior of and the exterior of the shape [VFLS13]. The generation of a machine toolpath for AM has a clear analogy with CNC pocket milling, where the material inside an arbitrarily closed boundary on a flat surface is removed to a fixed path [MA07, Hel91]. In the general case a machine toolpath intended for milling is, however, not suitable for additive manufacturing. In FDM the deposition of material along the path poses some additional challenges. The path has to be designed in such a way that the deposition of material is as regular as possible, thus avoiding under and over deposition. For the same reason the overlap between adjacent paths needs a much finer control, in order to cover the slice area with a uniform layer of deposited material [HJS03] (Figure 17). In powder bed technologies such as laser sintering and melting the heat control is fundamental to guarantee quality results. Many printers generate multiple melting pools at a time, in order to better distribute heat and avoid huge thermal gradients, that would generate rough surfaces and warping [DPCL14]. We discuss here the requirements for a quality machine toolpath and the major differences between available algorithms and technologies.
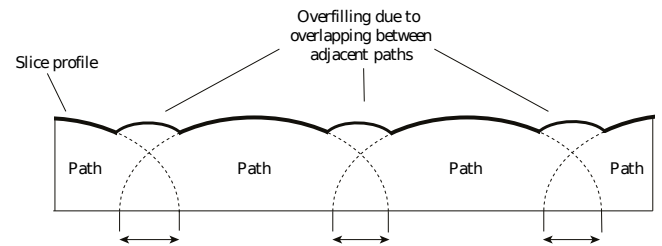

Figure 17: RP forming principle.

Continuity. In FDM and other material deposition processes it is important to keep the amount of material deposited along the path constant. To this end, many authors observed that controlling the amount of material being deposited when a path begins or ends is very difficult. Recent works aim to reduce as much as possible the number of disconnected paths which are necessary to cover the slice area, so as to minimize discontinuities in the deposition process [ZCG*16].

Geometry. Not only the endpoints of a machine path but also the geometry of the path itself may affect the quantity of material locally deposited in FDM. Long and low curvature paths are to be preferred to short paths with sharp turns. In the latter case the speed of the nozzle would decrease, thus increasing the time necessary to complete the deposition and possibly triggering under or overfilling of the filament [JHF*14].

Patterns. Internal volumes of parts are a major factor in time and material consumption, in particular for FDM where the motion of the print head – a relatively heavy mechanical device – is slow. To save time and material, several infill patterns have been proposed. These patterns are usually specific to FDM due to the aforementioned continuity and geometry requirement, but also due to the

overhang constraints. See Section 3.3.2 for an in-depth discussion of sparse infill patterns.

Domain split. In order to be able to process arbitrarily complex slices, most of the approaches proposed in the literature use a divide-and-conquer strategy, partitioning the slice into a set of pockets to operate on, rather than trying to cover the whole area with a single connected curve. To this end, many different strategies have been proposed in the literature. In [DK04] the slice is decomposed into a set of monotone polygons; [DPCL14] uses convex decomposition; [ZCG*16] uses the iso-contours of an inward distance field from the border; [HLA94] uses a Voronoi-based approach and [KP98] uses a medial axis based approach.

Performances. Besides the method used to decompose the domain and the particular curves used to fill each pocket, another important factor in the definition of an efficient toolpath is the reduction of the so called machine *airtime*, that is, the time necessary to move the nozzle (or the laser) from the end of a curve to the beginning of the subsequent one. Given a toolpath composed of a set of disconnected curves, the machine airtime can be minimized by acting on to three separate variables: the order in which the curves are processed, the path that takes the nozzle from the end of a curve to the beginning of the next one, and the orientation of each curve (from begin to end, or from end to begin). For the latter, notice that for closed loops, the machine may choose any point within the loop as a starting point. The path planning optimization problem has been shown to be related to the Travelling Salesman Problem (TSP), which is NP-complete. Many methods have been proposed for its approximate solution, using genetic algorithms [Wei09, WMJC02], the Christofides algorithm [FGCC16, GCFC16] or other heuristics [CDW03].

3.5.2. Raster case

Recent printing technologies such as Z-Corp. and the DLP/SLA printers (i.e. stereolithography printers that exploit the Digital Light Processing technology) treat each slice as an array of pixels. For the Z-Corp. case the machine toolpath consists in a trivial visit of a regular 2D grid. From a machine usage point of view an important parameter is the binder saturation, which affects both the strength and the accuracy of a printed part [VC11]. Typically, a binder is first applied with a higher saturation to the edges of the part, creating a strong *shell* for the exterior. Next, an infrastructure is created for the part walls, which are also built with a higher saturation. The remaining interior areas are printed with a lower binder saturation, which gives the part its stability [ZCo07]. For the DLP-based printers there is no notion of machine toolpath, because the whole image/slice is projected onto the first layer of the resin tank. In this case it is important to project each slice for the proper amount of time (exposure time), in order to balance surface accuracy with part strength. Exposure time depends on material, layer thickness and slice resolution. Notice that commercial slicers allow to export vector slices (e.g. in SVG format) even for raster printers [Sli11]. This is because slice rescaling is usually performed internally prior to projection, and scaling a raster image may introduce unnecessary artifacts such as image blurring.

3.6. Multiple components / Batch printing

In principle, a 3D model made of multiple parts can be printed all at once, that is, pre-assembled. However, with current technologies this approach might still have several drawbacks:

- The pre-assembled model can be too large to fit the printing chamber, whereas each single constituting part would fit;
- The model is small enough, but it contains tangency parts which would be fused together during the printing process (e.g. this might happen for bearing spheres);
- The model has none of the aforementioned issues, but printing the assembly would require the insertion of support structures that could not be removed;
- The model has none of the aforementioned issues, but printing the assembly would require the insertion of support structures that, after removal, would produce rough surfaces (above the tolerances set in the design).

For all these reasons, the current practice for the manufacturing of components made of multiple parts consists in building each part separately, and thus requires a reassembly of the physical objects to form the ultimate component. Some elements of the design (e.g. screws, bolts, bearings, ...) are bought from external producers that offer a catalogue of standardized components, whereas all the remaining parts are produced by additive manufacturing. Typically, each part undergoes most of the process planning steps independently. Each so-prepared part (that is, checked, oriented, and supported) is sent to the software that drives the printer and, through this software, an operator places the part in a free portion of the building plate. When the plate is full, or when there are no more parts to print, the operator runs the actual printing process. In this case, the only process planning steps which are common to all the parts are the slicing and the creation of the toolpath. It is also possible to slice one part at a time while using the same layer thickness: this is done, for example, on EOS machines which use the PSW software to place pre-sliced parts on the building platform. Algorithmic approaches exist that try to optimize for both quality and packing efficiency by properly orienting and placing the various parts on the platform [CDMS06, ZBHK15].

3.7. Model partitioning for 3D printing

3.7.1. Printing big objects

When even a single object is larger than the printing chamber, solutions exist to split it into parts to be printed independently and reassembled afterwards. A noticeable example is given in [LBRM12], where the algorithm decomposes a 3D object into reassemblable parts, each contained within a given printing volume. Besides reassemblability, this algorithm strives to avoid the production of too small parts, and considers the structural soundness by avoiding to put seams in areas of high mechanical stress (see Figure 18). In a similar work [HFW11], the seams are placed along lines of high curvature so as to minimize their impact on the aesthetics. In both [LBRM12] and [HFW11], connectors are created to ease the actual reassembly. Conversely, instead of using connectors, glue, or screws, in [SFLF15] the subdivision strives to create a part configuration that allows self-interlocking, such that the assembled

	Orientation	Supports	Slicing	Toolpath
Cost				
Pre-build	-	-	-	Sec. 3.5.1
Build	Sec. 3.2.1	Sec. 3.3.2	Sec. 3.4.1	Sec. 3.5.1
Post-processing	Sec. 3.2.2	Sec. 3.3.1	-	-
Fidelity				
Form	Sec. 3.2.2	-	Sec. 3.4.1	-
Texture	Sec. 3.2.2	Sec. 3.3.1	-	-
Design compliancy	-	-	-	-
Functionality				
Robustness	Sec. 3.2.3	-	-	-
Mass distribution	-	Sec. 3.3.2	-	-
Thermal/Mechanical prop.	-	Sec. 3.3.2	-	-

Table 4: *Relations between process planning steps and quality metrics.*

Figure 18: *A 3D object is partitioned into reassemblabe parts that can be printed. Image courtesy of [LBRM12].*

object can be not only repeatedly disassembled and reassembled, but also strongly connected by the parts' own geometry.

3.7.2. Print, pack and ship paradigm

Delivering a printed object (e.g. to a customer) has a cost that grows as the size of the *pack* grows. Thus, it is important to investigate how a 3D model can be split into easily printable parts that can eventually be tightly packed in a box and reassembled at the destination. In [Att15], an algorithm that performs this kind of split is proposed based on a user-controllable tradeoff between the packing efficiency and the number of parts to be produced. In [VGB*14b], the focus is mainly on the reduction of the material usage, for which hollowing and orientation optimization are used: however, the eventual packing efficiency is one of the parameters that drives the overall optimization process. In [YCL*15], the structural stresses are also considered to perform the subdivision prior to the packing. In [ZSMS14], a generic shape is converted into a foldable set of nearly-cubical parts. The printer can produce the object in its folded configuration which occupies a smaller volume, and thus might fit the printing chamber. Afterwards, the printed prototype can be unfolded into the desired shape.

3.7.3. Strive for quality prints through divide-and-conquer

After their removal, support structures easily leave tiny defects on the surface. To avoid this drawback, in [HLZCO14] the 3D model is split into so-called *approximate pyramidal shapes* that can be printed without supports. A similar approach is employed

in [HMA15], where a limited distortion is allowed to minimize the number of parts to be produced. In [HBA13] the surface quality is improved by splitting the model into few pieces so that each piece can be consistently sliced with a small geometric error along one of three orthogonal slicing directions.

3.8. Relations between quality metrics and basic steps

Each of the process planning stages discussed so far has an impact on some of the quality metrics. Since the relative importance of these metrics depends on the application, herewith we provide an overview of these process-quality relations that may help in the process of tuning the whole process planning. Table 4 summarizes these relations and, for each of the specific PP steps, refers to the corresponding section where the causes and effects are described in detail.

Note that in many cases several quality metrics should be optimized at the same time, and that is why multiple decision systems and genetic algorithms have been proposed to support the user in the difficult task of finding a tradeoff [INRT17, AAR01].

4. Open challenges

In this final section we discuss some of the open challenges in the process planning pipeline for additive manufacturing.

4.1. Challenges in an industrial perspective

AM has a huge potential for industry, thanks to the many advantages introduced by the technology itself and also the extreme simplification of the process planning compared to traditional approaches. Nevertheless, several challenges still have to be solved to enable a large adoption of AM by the industry.

Industry typically seeks integrated solutions for PP, i.e. algorithms for model repair, part orientation, support creation, slicing, toolpath calculation etc. shall be preferably integrated in a single, comprehensive tool. In this regard, one of the key challenges is to keep such tools up to date and allow a suitable degree of customization. At the academic level, the development of a common software framework for AM would help towards this goal: by developing their research prototypes within a common framework,

academics could reach a larger audience and ease transfer towards the industry.

An interesting trend is to attempt to automatically optimize most steps of the process (part orientation, supports, batch printing, etc.). However, it is important to give enough controls to users. For instance some users will prioritize the performance of the part (structural resistance, surface finish, etc.), while other users would be interested in optimizing the process (production time, material waste, etc.). To this end, the parameter(s) to be optimized by PP algorithms must be selected by the user. Unlike for non-professional users, tools where the parameters cannot be tuned would not be acceptable.

Another area of improvement of PP is the process simulation. As we have seen in this survey, some key characteristics of the final result can already be reliably forecast and some of the features can now be subjected to an optimization loop. The future development will hopefully produce more and more accurate tools that can be integrated in the PP, allowing the users to improve desiderata and minimize the uncertainties related to the AM technology.

4.2. Complementarity with subtractive technologies

In the majority of the mechanically demanding applications, AM is nowadays coupled to subtractive manufacturing. In fact, AM technologies are not capable yet of achieving the precision often required in mechanical parts, especially in terms features positioning, shape and surface finish, as illustrated in Figure 19. In order to address this issue, the common practice is to modify the CAD model leaving some extra material on the features that require high precision (e.g. couplings, holes, pins, shoulders for bearings, etc.). This extra material will be then milled until the desired shape is obtained with the required surface finish. Thanks to this approach, the parts manufactured by means of AM can be used effectively and coupled to traditional parts and commercial components, with a suitable level of precision. Some recent works in process planning integrate both technologies together, e.g. [NZDS15].

Nevertheless, the current limitation in the precision, stability and roughness of the features produced by AM is expected to be solved in the next years of technology development. Similarly to 5-axes milling, which was considerably improved with the development of CNC technologies and with the evolution of CAM tools, it is ex-

pected that soon AM will be able to produce parts perfectly meeting the stringent requirements of mechanical couplings.

This would be a double advantage, because not only it will simplify the production of parts and reduce the manufacture cost, but it would allow the users to fully exploit the freedom of AM, overcoming any constraint imposed by the finishing of features by subtractive manufacturing.

4.3. Design for additive manufacturing

In Section 3 we have presented a number of techniques that, given an input shape, aim to find the best possible way to 3D print it so as to match one or a combination of the targets listed in Section 2. In this section we approach the same problem from a different perspective: how can I design a shape that, if 3D printed, will satisfy the quality, cost or mechanical requirements that I would like to satisfy? The key observation is that, in many cases, changes on the design itself would make a big difference in terms of AM. Although shape design is out of the scope of this survey, we would like to give the reader some examples of works in this topic, to emphasize that most of the problems that the process plan aims to solve can also be tackled at design time. We believe that future software will couple modeling and process planning more tightly in order to guide the user towards shape the can be fabricated more reliably.

Metamaterials. We have discussed in Section 3.3.2 how filling objects with microstructures impact their physical properties (e.g. elasticity, porosity). Metamaterials give designers a higher level control by linking desired macroscale properties to specific microstructures, and are thus a key ingredient in design for additive manufacturing [Ros07]. In terms of design, this implies that the software should be able to automatically choose the metamaterial and its parameters from a desired macroscale behavior specified by the user [BBO*10, CLD*13, XLCB15]. Recent works explore user interfaces that let the user paint properties to create, for instance, flexible objects [SBR*15, PZM*15] and even mechanisms [IFW*16].

Design to reduce supports. Support structures impact on AM in many ways (material, removal effort, surface artifacts,...). Two interesting recent works attempt to reduce supports at the design stage. In [HJW15] the shape is modified to make it more self-sustainable and, thus, require a minor amount of external supports. In [RL16] the authors propose three basic 3D sculpting operations (trim, preserve and grow) to produce shapes that can be printed without supports structures.

In general, supports are one of the most difficult and penalizing issues for AM, and more research effort is required at all stages (modeling, planning, materials and machines) to alleviate their detrimental impact.

Design to ensure structural soundness. Often, it is discovered only after fabrication that a printed shape will not be able to sustain its own weight or will be too fragile. Stava et al. [SVB*12] proposed not only to predict fragilities in a part, but also proposed an algorithm to reinforce the shape by means of local operations such as hollowing, thickening ans strut insertion. Although not in
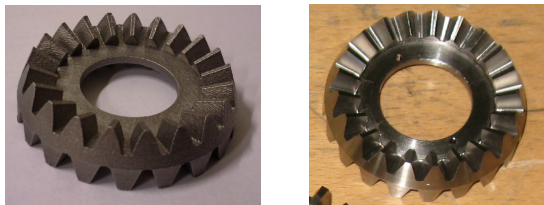

Figure 19: Example of mechanical component (a bevel gear) manufactured with AM (left) and with subtractive technique (right). The surface finish achieved by AM is not compatible with the specifications of the gear, so a finishing operation (e.g., by milling) is required.

the context of AM, Umetani et al. [UIM12] proposed a design tool for furniture that suggests changes based on structural strength prediction. This type of feedback and automated suggestions during design could enable a widespread adoption of AM, allowing anyone to model a shape while the algorithm verifies the design and suggests fixes.

Industry. People operating in industrial design have been creating shapes to be fabricated with subtractive techniques for decades. Nowadays AM poses new challenges for designer, which have to think of new good principles that fit a different manufacturing paradigm. *Printing-aware design* is often used in literature to refer to attempts to go in this direction. We point the reader to [HPE16, TMV*16, SAZ15] and references therein for inspiring discussions on this topic.

4.4. Embedding devices into prints

A fascinating and rapidly developing topic is the ability to insert sensors or devices into a print. For instance, Savage et al. proposed a design tool to augment objects with specially fabricated touch sensitive areas [SZH12], and to embed a camera into 3D printed objects to turn them into custom controllers [SCH13]. The RevoMaker [GZN*15] fabricates an object *around* a core containing sensors and electronic devices. The *Voxel8* printer goes beyond this by proposing to integrate electronic components connected together by conductive ink deposited within the layers. Similarly, [PGMH16] describes a prototype printer that, besides depositing fused plastic, can also roll copper coils that produce magnetic fields when traversed by a current.

Such conductive materials open a wide range of applications, in particular for antenna or battery designs [RL15].

Such applications open a whole new set of challenges overlapping between material science, geometry, 3D modeling, circuit routing, and process planning for additive manufacturing.

4.5. Self-assembling structures

Another promising direction of research is the study of materials that are able to morph into a target shape after being printed, the so-called 4D printing [Tib14]. For instance, to print plant-inspired structures that change their shape after immersion in water [GMN*16]. Self-assembly structures have applications in autonomous robotics or tissue engineering, among others [KTL*15].

4.6. Discretization-free pipeline

As discussed earlier in this report, the usual entry point for Process Planning consists of a tessellation that approximates the smoothly curved CAD model with a collection of triangles. This step simplifies geometric operations and exchanges between software to the point that the triangle-based STL format has become a de-facto standard for process planning. Nonetheless, since the tessellation step is often the cause of many issues with the geometry (see Section 3) and introduces unnecessary approximations, many attempts have been made during the years to avoid this conversion and perform all the calculations based on the original CAD geometry (see

Section 3.4.2). In spite of these efforts, STL remains an undiscussed standard and CAD-compliant geometric processing for PP is still a very limited practice.

A main obstacle to a widespread adoption of such pipelines is the intrinsic hardware limitations of current printing devices. 3D printers, indeed, can just execute a finite set of commands that typically makes the tool move from one point to another along a straight line segment or, in the most sophisticated models, to follow circular arcs while extruding or solidifying material.

Nevertheless a discretization-free pipeline would allow a form of device independent process planning, where discretization would only occur at the very last moment, on the machine. Thus, the approximations would be optimally decided depending upon machine capabilities and resolution.

Acknowledgements

References

[AAD98] ALEXANDER P., ALLEN S., DUTTA D.: Part orientation and build cost determination in layered manufacturing. *Computer-Aided Design 30*, 5 (1998), 343–356. 5, 6

[AAR01] ANITHA R., ARUNACHALAM S., RADHAKRISHNAN P.: Critical parameters influencing the quality of prototypes in fused deposition modelling. *J. Mater. Process. Tech. 118*, 1 (2001), 385–388. 19

[ACC*88] ALLISON J. W., CHEN T. P., COHEN A. L., SMALLEY D. R., SNEAD D. E., VORGITCH T. J.: Boolean layer comparison slice, 1988. US Patent 5854748, 3D Systems Inc. 10

[ACK13] ATTENE M., CAMPEN M., KOBBELT L.: Polygon mesh repairing: an application perspective. *ACM Computing Surveys 45*, 2 (2013), 15:1–15:33. 8

[AD94] ALLEN S., DUTTA D.: On the computation of part orientation using support structures in layered manufacturing. In *Proceedings of Solid Freeform Fabrication Symposium* (1994), pp. 259–269. 8

[AD95] ALLEN S., DUTTA D.: Determination and evaluation of support structures in layered manufacturing. *Journal of Design and Manufacturing 5* (1995), 153–162. 10

[AKL07] AHN D., KIM H., LEE S.: Fabrication direction optimization to minimize post-machining in layered manufacturing. *Int. J. Mach. Tool. Manu. 47*, 3 (2007), 593–606. 8, 10

[ALS14] ANDREASSEN E., LAZAROV B. S., SIGMUND O.: Design of manufacturable 3D extremal elastic microstructure. *Mechanics of Materials 69*, 1 (2014), 1–10. 14

[AT15] ALLEN R. J. A., TRASK R. S.: An experimental demonstration of effective curved layer fused filament fabrication utilising a parallel deposition robot. *Additive Manufacturing 8* (2015), 78–87. 16

[Att10] ATTENE M.: A lightweight approach to repair polygon meshes. *The Visual Computer* (2010), 1393–1406. 8

[Att14] ATTENE M.: Direct repair of self-intersecting meshes. *Graphical Models 76* (2014), 658–668. 8

[Att15] ATTENE M.: Shapes in a box: Disassembling 3D objects for efficient packing and fabrication. *Comp. Graph. Forum 34*, 8 (2015), 64–76. 19

[Aut11] AUTODESK: Meshmixer. http://www.meshmixer.com, 2011. 8

[BAH11] BRACKETT D., ASHCROFT I., HAGUE R.: Topology optimization for additive manufacturing. In *Proceedings of the Solid Freeform Fabrication Symposium* (2011), pp. 348–362. 7

[BAU15] BRUNTON A., ARIKAN C. A., URBAN P.: Pushing the limits of 3D color printing: Error diffusion with translucent materials. *ACM Trans. Graph.*, 1 (2015), 1–13. 17

[BBO*10] BICKEL B., BÄCHER M., OTADUY M. A., LEE H. R., PFISTER H., GROSS M., MATUSIK W.: Design and fabrication of materials with desired deformation behavior. *ACM Tr. Graph. 29*, 4 (2010), 1. 20

[BDDH11] BRETEAU T., DAMAY T., DUC E., HASCOËT J.-Y.: Design for manufacturing with tool paths adapted to marine propeller. *Int. J. on Interactive Design and Manufacturing 5*, 4 (2011), 271–275. 1

[BL06a] BYUN H.-S., LEE K. H.: Determination of optimal build direction in rapid prototyping with variable slicing. *Int. J. Adv. Manuf. Tech. 28*, 3-4 (2006), 307–313. 5, 9, 10

[BL06b] BYUN H.-S., LEE K. H.: Determination of the optimal build direction for different rapid prototyping processes using multi-criterion decision making. *Robotics and Computer-Integrated Manufacturing 22*, 1 (2006), 69–80. 5, 9, 10

[BM97] BENOUAMER M. O., MICHELUCCI D.: Bridging the gap between CSG and Brep via a triple ray representation. In *Procs. of the fourth ACM symposium on Solid modeling and applications* (1997). 15

[BPH13] BROCKMEYER E., POUPYREV I., HUDSON S.: PAPILLON: designing curved display surfaces with printed optics. In *Proceedings of UIST* (2013), pp. 457–462. 7

[BPK05] BISCHOFF S., PAVIC D., KOBBELT L.: Automatic restoration of polygon models. *ACM Trans. Graph. 24*, 4 (2005), 1332–1352. 7

[BWBSH14] BÄCHER M., WHITING E., BICKEL B., SORKINE-HORNUNG O.: Spin-It: Optimizing moment of inertia for spinnable objects. *ACM Trans. Graph. 33*, 4 (2014), 96:1–96:10. 7

[Cal14] CALIGNANO F.: Design optimization of supports for overhanging structures in aluminum and titanium alloys by selective laser melting. *Materials & Design 64* (2014), 203–213. 11

[Car84] CARPENTER L.: The A-buffer, an antialiased hidden surface method. In *Proceedings of SIGGRAPH* (1984), pp. 103–108. 15

[CB14] CALDERON S., BOUBEKEUR T.: Point morphology. *ACM Trans. Graph. 33*, 4 (2014), 45:1–45:13. 11, 12

[CC05] CHOI S., CHEUNG H.: A multi-material virtual prototyping system. *Computer-Aided Design 37*, 1 (2005), 123–136. 16

[CC06] CHOI S., CHEUNG H.: A topological hierarchy-based approach to toolpath planning for multi-material layered manufacturing. *Computer-Aided Design 38*, 2 (2006), 143–156. 16

[CDMS06] CANELLIDIS V., DEDOUSSIS V., MANTZOURATOS N., SOFIANOPOULOU S.: Pre-processing methodology for optimizing stereolithography apparatus build performance. *Computers in industry 57*, 5 (2006), 424–436. 18

[CDW03] CASTELINO K., D'SOUZA R., WRIGHT P. K.: Toolpath optimization for minimizing airtime during machining. *J. Manuf. Syst. 22*, 3 (2003), 173–180. 18

[CFN*95] CHENG W., FUH J., NEE A., WONG Y., LOH H., MIYAZAWA T.: Multi-objective optimization of part-building orientation in stereolithography. *Rapid Prototyping J. 1*, 4 (1995), 12–23. 8

[Che07] CHEN Y.: 3D texture mapping for rapid manufacturing. *Comput. Aided Des. Appl. 4*, 6 (2007), 761–771. 13

[CJR95] CHALASANI K., JONES L., ROSCOE L.: Support generation for fused deposition modeling. In *Solid Freeform Fabrication Symposium* (1995), pp. 229–241. 10

[CK10a] CAMPEN M., KOBBELT L.: Exact and robust (self-)intersections for polygonal meshes. *Comp. Graph. Forum 29*, 2 (2010), 397–406. 8

[CK10b] CAMPEN M., KOBBELT L.: Polygonal boundary evaluation of Minkowski sums and swept volumes. *Comp. Graph. Forum 29*, 5 (2010), 1613–1622. 15

[CLD*13] CHEN D., LEVIN D. I. W., DIDYK P., SITTHI-AMORN P., MATUSIK W.: Spec2fab: A reducer-tuner model for translating specifications to 3D prints. *ACM Trans. Graph. 32*, 4 (2013), 1. 20

[CLQ13] CHEN Y., LI K., QIAN X.: Direct geometry processing for telefabrication. *J. Comput. Inf. Sci. Eng. 13*, 4 (2013), 041002. 10, 16

[CRC08] CHAKRABORTY D., REDDY B. A., CHOUDHURY A. R.: Extruder path generation for curved layer fused deposition modeling. *Computer-Aided Design 40*, 2 (2008), 235–243. 16

[Cru89] CRUMP S. S.: Apparatus and method for creating three-dimensional objects - U.S. patent 5.121.329, 1989. 3

[CSP*01] CHO W., SACHS E., PATRIKALAKIS N., CIMA M., JACKSON T., LIU H., SERDY J., STRATTON C., WU H., RESNICK R.: Methods for distributed design and fabrication of parts with local composition control. In *Proceedings of the 2001 NSF Design and Manufacturing Grantees Conference* (2001). 17

[CT98] CHIU W., TAN S.: Using dexels to make hollow models for rapid prototyping. *Computer-Aided Design 30*, 7 (1998), 539–547. 12

[CTF*94] COX J. J., TAKEZAKI Y., FERGUSON H. R., KOHKONEN K. E., MULKAY E. L.: Space-filling curves in tool-path applications. *Computer-Aided Design 26*, 3 (1994), 215–224. 12

[CW08] CHEN Y., WANG C. C.: Layer depth-normal images for complex geometries: Part one – accurate modeling and adaptive sampling. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (2008). 11, 15

[CW11] CHEN Y., WANG C. C.: Uniform offsetting of polygonal model based on layered depth-normal images. *Comp.-Aided Design 43*, 1 (2011), 31–46. 12

[CW13] CHEN Y., WANG C. C. L.: Regulating complex geometries using layered depth–normal images for rapid prototyping and manufacturing. *Rapid Prototyping J. 19*, 4 (2013), 253–268. 7, 15

[CZ10] CHOI S., ZHU W.: A dynamic priority-based approach to concurrent toolpath planning for multi-material layered manufacturing. *Computer-Aided Design 42*, 12 (2010), 1095–1107. 16

[CZCL13] CADMAN J. E., ZHOU S., CHEN Y., LI Q.: On design of multi-functional microstructural materials. *Journal of Materials Science 48*, 1 (2013), 51–66. 13

[DGLC15] DINH H. Q., GELMAN F., LEFEBVRE S., CLAUX F.: Modeling and toolpath generation for consumer-level 3D printing. In *ACM SIGGRAPH 2015 Courses* (2015), pp. 17:1–17:273. 13

[DHL14] DUMAS J., HERGEL J., LEFEBVRE S.: Bridging the gap: automated steady scaffoldings for 3D printing. *ACM Trans. Graph. 33*, 4 (2014), 98:1–98:10. 9, 10, 11

[DJZ15] DEDE E. M., JOSHI S. N., ZHOU F.: Topology optimization, additive layer manufacturing, and experimental testing of an air-cooled heat sink. *Journal of Mechanical Design 137*, 11 (2015), 111403. 7

[DK04] DWIVEDI R., KOVACEVIC R.: Automated torch path planning using polygon subdivision for solid freeform fabrication based on welding. *J. Manuf. Syst. 23*, 4 (2004), 278–291. 18

[DM94] DOLENC A., MÄKELÄ I.: Slicing procedures for layered manufacturing techniques. *Comp.-Aided Design 26*, 2 (1994), 119–126. 2, 6, 14

[DPCL14] DING D., PAN Z. S., CUIURI D., LI H.: A tool-path generation strategy for wire and arc additive manufacturing. *Int. J. Adv. Manuf. Tech. 73*, 1-4 (2014), 173–183. 12, 17, 18

[DTS16] DELFS P., ŤOWS M., SCHMID H.-J.: Optimized build orientation of additive manufactured parts for improved surface quality and build time. *Additive Manufacturing* (2016). 6, 8, 10

[DWP*10] DONG Y., WANG J., PELLACINI F., TONG X., GUO B.: Fabricating spatially-varying subsurface scattering. *ACM Trans. Graph. 29*, 4 (2010), 1. 17

[EME15] Ezair B., Massarwi F., Elber G.: Orientation analysis of 3D objects toward minimal support volume in 3D-printing. *Computers & Graphics 51* (2015), 117–124. 8, 10

[EMET06] El-Midany T. T., Elkeran A., Tawfik H.: Toolpath pattern comparison: Contour-parallel with direction-parallel. In *Geometric Modeling and Imaging–New Trends* (2006), pp. 77–82. 12

[ER07] Eggers G., Renap K.: Method and apparatus for automatic support generation for an object made by means of a rapid prototype production method, 2007. US Patent 20100228369, Materialize. 10, 11

[Far85] Farouki R.: Exact offset procedures for simple solids. *Comput. Aided Geom. D. 2*, 4 (1985), 257–279. 11

[FF95] Frank D., Fadel G.: Expert system-based selection of the preferred direction of build for rapid prototyping processes. *J. Intell. Manuf. 6*, 5 (1995), 339–345. 8

[FGCC16] Fok K.-Y., Ganganath N., Cheng C.-T., Chi K. T.: A 3D printing path optimizer based on Christofides algorithm. In *IEEE Int. Conf. on Consumer Electronics-Taiwan* (2016), pp. 1–2. 18

[For95] Forsyth M.: Shelling and offsetting bodies. In *Procs of the 3rd ACM Symp. on Solid Modeling and Applications* (1995), pp. 373–381. 11

[FPRJ00] Frisken S. F., Perry R. N., Rockwood A. P., Jones T. R.: Adaptively sampled distance fields: A general representation of shape for computer graphics. In *Procs of SIGGRAPH* (2000), pp. 249–254. 11

[FVP13] Fryazinov O., Vilbrandt T., Pasko A. A.: Multi-scale space-variant FRep cellular structures. *Computer-Aided Design 45*, 1 (2013), 26–34. 13

[GCFC16] Ganganath N., Cheng C.-T., Fok K.-Y., Chi K. T.: Trajectory planning for 3D printing: A revisit to traveling salesman problem. In *2nd Int. Conf. on Control, Automation and Robotics* (2016), pp. 287–290. 18

[GMN*16] Gladman A. S., Matsumoto E. A., Nuzzo R. G., Mahadevan L., Lewis J. A.: Biomimetic 4d printing. *Nature materials 15* (2016), 413–418. 21

[Gri94] Griffiths J. G.: Toolpath based on Hilbert's curve. *Computer-Aided Design 26*, 11 (1994), 839–844. 12

[GZN*15] Gao W., Zhang Y., Nazzetta D. C., Ramani K., Cipra R. J.: Revomaker: Enabling multi-directional and functionally-embedded 3d printing using a rotational cuboidal platform. In *Proceedings of UIST* (2015), pp. 437–446. 21

[GZR*15] Gao W., Zhang Y., Ramanujan D., Ramani K., Chen Y., Williams C. B., Wang C. C., Shin Y. C., Zhang S., Zavattieri P. D.: The status, challenges, and future of additive manufacturing in engineering. *Computer-Aided Design 69* (2015), 65–89. 1

[HA13] Hayasi M. T., Asiabanpour B.: A new adaptive slicing approach for the fully dense freeform fabrication (FDFF) process. *J. Intell. Manuf. 24*, 4 (2013), 683–694. 14

[HBA13] Hildebrand K., Bickel B., Alexa M.: Orthogonal slicing for additive manufacturing. *Computers & Graphics 37*, 6 (2013), 669–675. 8, 10, 14, 19

[Hei10] Heide E.: Method for generating and building support structures with deposition-based digital manufacturing systems, 2010. US Patent 20110178621 A1, Stratasys. 10

[Hel91] Held M.: A geometry-based investigation of the tool path generation for zigzag pocket machining. *The Visual Computer 7*, 5-6 (1991), 296–308. 17

[HFM*10] Hašan M., Fuchs M., Matusik W., Pfister H., Rusinkiewicz S.: Physical reproduction of materials with specified subsurface scattering. *ACM Trans. Graph. 29*, 3 (2010). 17

[HFW11] Hao J., Fang L., Williams R. E.: An efficient curvature-based partitioning of large–scale STL models. *Rapid Prototyping J. 17*, 2 (2011), 116–127. 18

[HJS03] Han W., Jafari M. A., Seyed K.: Process speeding up via deposition planning in fused deposition-based layered manufacturing processes. *Rapid Prototyping J. 9*, 4 (2003), 212–218. 17

[HJW15] Hu K., Jin S., Wang C. C.: Support slimming for single material based additive manufacturing. *Comp.-Aided Design 65* (2015), 1–10. 9, 20

[HKM07] Hachenberger P., Kettner L., Mehlhorn K.: Boolean operations on 3D selective Nef complexes: Data structure, algorithms, optimized implementation and experiments. *Computational Geometry 38*, 1–2 (2007), 64–99. 8

[HL98] Hur J., Lee K.: The development of a CAD environment to determine the preferred build-up direction for layered manufacturing. *Int. J. Adv. Manuf. Tech. 14*, 4 (1998), 247–254. 8, 9, 10

[HL14] Hergel J., Lefebvre S.: Clean color: Improving multi-filament 3d prints. *Comp. Graph. Forum 33*, 2 (2014), 469–478. 16

[HLA94] Held M., Lukács G., Andor L.: Pocket machining based on contour-parallel tool paths generated by means of proximity maps. *Computer-Aided Design 26*, 3 (1994), 189–203. 12, 18

[HLDC16] Hornus S., Lefebvre S., Dumas J., Claux F.: Tight printable enclosures and support structures for additive manufacturing. In *EG Workshop on Graphics for Digital Fabrication* (2016). 13

[HLZCO14] Hu R., Li H., Zhang H., Cohen-Or D.: Approximate pyramidal shape decomposition. *ACM Trans. Graph. 33*, 6 (2014), 213:1–213:12. 19

[HMA15] Herholz P., Matusik W., Alexa M.: Approximating freeform geometry with height fields for manufacturing. *Computer Graphics Forum 34*, 2 (2015), 239–251. 19

[HMS*99] Hartquist E., Menon J., Suresh K., Voelcker H., Zagajac J.: A computing strategy for applications involving offsets, sweeps, and Minkowski operations. *Computer-Aided Design 31*, 3 (1999), 175–183. 12

[HNCS16] Hildreth O. J., Nassar A. R., Chasse K. R., Simpson T. W.: Dissolvable metal supports for 3D direct metal printing. *3D Printing and Additive Manufacturing 3*, 2 (2016), 90–97. 11

[HNX*14] Huang Q., Nouri H., Xu K., Chen Y., Sosina S., Dasgupta T.: Predictive modeling of geometric deviations of 3D printed products-a unified modeling approach for cylindrical and polygon shapes. In *IEEE International Conference on Automation Science and Engineering* (2014), pp. 25–30. 6

[Hol05] Hollister S. J.: Porous scaffold design for tissue engineering. *Nature materials 4*, 7 (2005), 518–524. 13

[Hoo86] Hook T. V.: Real-time shaded NC milling display. In *Proceedings of SIGGRAPH* (1986), pp. 15–20. 11, 15

[HPE16] Hällgren S., Pejryd L., Ekengren J.: (Re) design for additive manufacturing. *Procedia CIRP 50* (2016), 246–251. 1, 21

[HRJ97] Hope R. L., Roth R. N., Jacobs P. A.: Adaptive slicing with sloping layer surfaces. *Rapid Prototyping J. 3*, 3 (1997), 89–98. 14

[HS12] Huang B., Singamneni S.: Alternate slicing and deposition strategies for fused deposition modelling of light curved parts. *J. of Achievem in mat and manuf 55*, 2 (2012), 511–517. 16

[Hui94] Hui K.: Solid sweeping in image space – application in NC simulation. *The Visual Computer 10*, 6 (1994), 306–316. 11

[HWC13] Huang P., Wang C. C., Chen Y.: Intersection-free and topologically faithful slicing of implicit solid. *J. Comput. Inf. Sci. Eng. 13*, 2 (2013), 021009. 7, 15

[HWC14] Huang P., Wang C. C., Chen Y.: Algorithms for layered manufacturing in image space. *ASME Advances in Computers and Information in Engineering Research 1* (2014), 377–410. 10, 11, 15

[HWH15] Huang T., Wang S., He K.: Quality control for fused deposition modeling based additive manufacturing: Current research and future trends. In *First International Conference on Reliability Systems Engineering* (2015), pp. 1–6. 6

[HYML09] HUANG X., YE C., MO J., LIU H.: Slice data based support generation algorithm for fused deposition modeling. *Tsinghua Science and Technology 14*, S1 (2009), 223–228. 10

[HYW*09] HUANG X., YE C., WU S., GUO K., MO J.: Sloping wall structure support generation for fused deposition modeling. *Int. J. Adv. Manuf. Tech. 42*, 11-12 (2009), 1074–1081. 10

[HZH*16] HUANG Y., ZHANG J., HU X., SONG G., LIU Z., YU L., LIU L.: Framefab: Robotic fabrication of frame shapes. *ACM Trans. Graph. 35*, 6 (2016). 16

[ice] IceSL. http://www.loria.fr/~slefebvr/icesl/. 7

[IFW*16] ION A., FROHNHOFEN J., WALL L., KOVACS R., ALISTAR M., LINDSAY J., LOPES P., CHEN H.-T., BAUDISCH P.: Metamaterial mechanisms. In *Proceedings of UIST* (2016). 20

[INRT17] INGRASSIA T., NIGRELLI V., RICOTTA V., TARTAMELLA C.: Process parameters influence in additive manufacturing. In *Adv. on Mechanics, Design Eng. and Manuf.* Springer, 2017, pp. 261–270. 19

[ISO97] ISO: ISO 4287: 1997: Geometrical product specifications (GPS)–surface texture: Profile method–terms, definitions and surface texture parameters, 1997. 6

[JBAG12] JHABVALA J., BOILLAT E., ANDRÉ C., GLARDON R.: An innovative method to build support structures with a pulsed laser in the selective laser melting process. *Int. J. Adv. Manuf. Tech. 59*, 1-4 (2012), 137–142. 11

[JH95] JAMIESON R., HACKER H.: Direct slicing of CAD models for rapid prototyping. *Rapid Prototyping J. 1*, 2 (1995), 4–12. 15

[JHF13] JIN Y. A., HE Y., FU J. Z.: An adaptive tool path generation for fused deposition modeling. In *Advanced Materials Research* (2013), vol. 819, Trans Tech Publ, pp. 7–12. 12

[JHF*14] JIN Y.-A., HE Y., FU J.-Z., GAN W.-F., LIN Z.-W.: Optimization of tool-path generation for material extrusion-based additive manufacturing technology. *Additive Manufacturing 1* (2014), 32–47. 17

[JKSH13] JACOBSON A., KAVAN L., SORKINE-HORNUNG O.: Robust inside-outside segmentation using generalized winding numbers. *ACM Trans. Graph. 32*, 4 (2013), 33:1–33:12. 8

[JLP*99] JACKSON T., LIU H., PATRIKALAKIS N., SACHS E., CIMA M.: Modeling and designing functionally graded material components for fabrication with local composition control. *Materials & Design 20*, 2–3 (1999), 63–75. 13

[JQH15] JIN Y., QIN S. J., HUANG Q.: Out-of-plane geometric error prediction for additive manufacturing. In *IEEE International Conference on Automation Science and Engineering* (2015), pp. 918–923. 6

[Ju04] JU T.: Robust repair of polygonal models. *ACM Trans. Graph. 23*, 3 (2004), 888–895. 7

[Ju09] JU T.: Fixing geometric errors on polygonal models: A survey. *Computer Science and Technology 24*, 1 (2009), 19–29. 8

[KC02] KIM B. H., CHOI B. K.: Machining efficiency comparison direction-parallel tool path with contour-parallel tool path. *Computer-Aided Design 34*, 2 (2002), 89–95. 12

[KD96] KULKARNI P., DUTTA D.: An accurate slicing procedure for layered manufacturing. *Comp.-Aided Design 28*, 9 (1996), 683–697. 14

[KD98] KUMAR V., DUTTA D.: An approach to modeling & representation of heterogeneous objects. *J. Mechanical Design 120*, 4 (1998), 659. 16

[KJA92] KIRSCHMAN C., JARA-ALMONTE C.: A parallel slicing algorithm for solid freeform fabrication processes. In *Solid Freeform Fabrication Proceedings* (1992), pp. 26–33. 15

[KJAB*91] KIRSCHMAN C., JARA-ALMONTE C., BAGCHI A., DOOLEY R., OGALE A.: Computer aided design of support structures for stereolithographic components. In *Proceedings of the 1991 ASME Computers in Engineering Conference* (1991), pp. 443–451. 10

[KKD98] KUMAR V., KULKARNI P., DUTTA D.: Adaptive slicing of heterogeneous solid models for layered manufacturing. In *Proceedings of the ASME Design Technical Conferences* (1998). 16

[KM06] KHARDEKAR R., MCMAINS S.: Fast layered manufacturing support volume computation on GPUs. In *ASME 2006 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (2006), pp. 993–1002. 8

[KMD00] KULKARNI P., MARSAN A., DUTTA D.: A review of process planning techniques in layered manufacturing. *Rapid Prototyping J. 6*, 1 (2000), 18–35. 2

[KP98] KAO J.-H., PRINZ F. B.: Optimal motion planning for deposition in layered manufacturing. In *Procs of DETC* (1998), vol. 98, pp. 13–16. 12, 18

[KPA09] KUMAR G. S., PANDITHEVAN P., AMBATTI A. R.: Fractal raster tool paths for layered manufacturing of porous objects. *Virtual and Physical Prototyping 4*, 2 (2009), 91–104. 13

[KT07] KOU X., TAN S.: Heterogeneous object modeling: A review. *Computer-Aided Design 39*, 4 (2007), 284–301. 13

[KTL*15] KHOO Z. X., TEOH J. E. M., LIU Y., CHUA C. K., YANG S., AN J., LEONG K. F., YEONG W. Y.: 3d printing of smart materials: A review on recent progresses in 4d printing. *Virtual and Physical Prototyping 10*, 3 (2015), 103–122. 21

[LBRM12] LUO L., BARAN I., RUSINKIEWICZ S., MATUSIK W.: Chopper: partitioning models into 3D-printable parts. *ACM Trans. Graph. 31*, 6 (2012), 129. 7, 18, 19

[LCCG97] LAN P.-T., CHOU S.-Y., CHEN L.-L., GEMMILL D.: Determining fabrication orientations for rapid prototyping with stereolithography apparatus. *Computer-Aided Design 29*, 1 (1997), 53–62. 8

[LCH15] LAPUT G., CHEN X. A., HARRISON C.: 3D printed hair: Fused deposition modeling of soft strands, fibers, and bristles. In *Proceedings of UIST* (2015), pp. 593–597. 16

[LDJC15] LI D., DAI N., JIANG X., CHEN X.: Interior structural optimization based on the density-variable shape modeling of 3D printed objects. *Int. J. Adv. Manuf. Tech. 83*, 9 (2015), 1627–1635. 7, 13

[Lef13] LEFEBVRE S.: IceSL : A GPU accelerated modeler and slicer. In *18th European Forum on Additive Manufacturing* (2013). 15

[Lef15] LEFEBVRE S.: 3d infilling: faster, stronger, simpler, 2015. http://sylefeb.blogspot.fr/2015/07/3dprint-3d-infilling-faster-stronger.html. 13

[LHL14] LEFEBVRE S., HORNUS S., LASRAM A.: Per-pixel lists for single pass a-buffer. In *GPU Pro5*. Informa UK Ltd, 2014, pp. 3–23. 15

[Lie08] LIEN J.-M.: Covering Minkowski sum boundary using points with applications. *Comput. Aided Geom. D. 25*, 8 (2008), 652–666. 12

[LL16] LEE J., LEE K.: Block-based inner support structure generation algorithm for 3D printing using fused deposition modeling. *Int. J. Adv. Manuf. Tech.* (2016). 13

[LLMZ16] LI D., LEVIN D. I., MATUSIK W., ZHENG C.: Acoustic voxels: Computational optimization of modular acoustic filters. *ACM Trans. Graph. 35*, 4 (2016), 88. 7

[LM14] LI W., MCMAINS S.: A sweep and translate algorithm for computing voxelized 3D Minkowski sums on the GPU. *Computer-Aided Design 46*, 0 (2014), 90–100. 12

[LSD*16] LANGLOIS T., SHAMIR A., DROR D., MATUSIK W., LEVIN D. I.: Stochastic structural analysis for context-aware design and fabrication. *ACM Trans. Graph. 35*, 6 (2016). 7

[LSZ*14] LU L., SHARF A., ZHAO H., WEI Y., FAN Q., CHEN X., SAVOYE Y., TU C., COHEN-OR D., CHEN B.: Build-to-last: Strength to weight 3D printed objects. *ACM Trans. Graph. 33*, 4 (2014), 97. 5, 13

[LW11] LIU S., WANG C. C. L.: Fast intersection-free offset surface generation from freeform models with triangular meshes. *IEEE Trans. Autom. Sci. Eng 8*, 2 (2011), 347–360. 11

[LW13] LEUNG Y.-S., WANG C. C. L.: Conservative sampling of solids in image space. *IEEE Comp. Graph. & Appl. 33*, 1 (2013), 32–43. 15

[LWZL03] LIU G. H., WONG Y. S., ZHANG Y. F., LOH H. T.: Error-based segmentation of cloud data for direct rapid prototyping. *Computer-Aided Design 35*, 7 (2003), 633–645. 16

[MA07] MAKHANOV S. S., ANOTAIPAIBOON W.: *Advanced numerical methods to optimize cutting operations of five axis milling machines*. Springer Science & Business Media, 2007. 12, 17

[MAB*15] MUSIALSKI P., AUZINGER T., BIRSAK M., WIMMER M., KOBBELT L.: Reduced-order shape optimization using offset surfaces. *ACM Trans. Graph. 34*, 4 (2015), 102. 7

[MBH04] MA W., BUT W.-C., HE P.: NURBS-based adaptive slicing for efficient rapid prototyping. *Comp.-Aided Design 36*, 13 (2004), 1309–1325. 15

[MCJ*16] MORGAN H., CHERRY J., JONNALAGANNA S., EWING D., SIENZ J.: Part orientation optimisation for the additive layer manufacture of metal components. *Int. J. Adv. Manuf. Tech. 86*, 5 (2016), 1679–1687. 8, 10

[Mco05] MCOR: Mcor technologies, 2005. http://mcortechnologies.com/. 3, 17

[MCTB11] MAULE M., COMBA J. L. D., TORCHELSEN R. P., BASTOS R.: A survey of raster-based transparency techniques. *Computers & Graphics 35*, 6 (2011), 1023–1034. 15

[MDL16] MARTÍNEZ J., DUMAS J., LEFEBVRE S.: Procedural Voronoi foams for additive manufacturing. *ACM Trans. Graph. 35*, 4 (2016), 44:1–44:12. 14

[MeSEPC16] MEDEIROS E SÁ A., ECHAVARRIA K. R., PIETRONI N., CIGNONI P.: State of the art on functional fabrication. In *Eurographics Workshop on Graphics for Digital Fabrication* (2016). 1

[MeSMEC15] MEDEIROS E SÁ A., MELLO V. M., ECHAVARRIA K. R., COVILL D.: Adaptive voids. *Vis. Computer 31*, 6-8 (2015), 799–808. 13

[MHCL15] MARTÍNEZ J., HORNUS S., CLAUX F., LEFEBVRE S.: Chained segment offsetting for ray-based solid representations. *Computers & Graphics 46* (2015), 36–47. 12

[MHR*16] MUSIALSKI P., HAFNER C., RIST F., BIRSAK M., WIMMER M., KOBBELT L.: Non-linear shape optimization using local subspace projections. *ACM Trans. Graph. 35*, 4 (2016), 87:1–87:13. 7

[MIG*14] MUELLER S., IM S., GUREVICH S., TEIBRICH A., PFISTERER L., GUIMBRETIÃLRE F., BAUDISCH P.: Wireprint: 3D printed previews for fast prototyping. In *Proceedings of UIST* (2014). 3, 5, 16

[Mit09] MITUTOYO: Surface roughness measurement. *Mitutoyo America Corporation Buletin n. 1984* (2009). 6

[MKD99] MANI K., KULKARNI P., DUTTA D.: Region-based adaptive slicing. *Computer-Aided Design 31*, 5 (1999), 317–333. 14

[MN13] MICROSOFT, NETFABB: 3D model repair service. https://modelrepair.azurewebsites.net/, 2013. 8

[MRI00] MASOOD S., RATTANAWONG W., IOVENITTI P.: Part build orientations based on volumetric error in fused deposition modelling. *Int. J. Adv. Manuf. Tech. 16*, 3 (2000), 162–168. 6, 8, 10

[MRI03] MASOOD S., RATTANAWONG W., IOVENITTI P.: A generic algorithm for a best part orientation system for complex parts in rapid prototyping. *J. Mater. Process. Tech. 139*, 1 (2003), 110–116. 6, 8, 10

[MS99] MCMAINS S., SÉQUIN C.: A coherent sweep plane slicer for layered manufacturing. In *Proceedings of the fifth ACM symposium on Solid modeling and applications* (1999). 15

[MSWS00] MCMAINS S., SMITH J., WANG J., SÉQUIN C.: Layered manufacturing of thin-walled parts. In *ASME Design Engineering Technical Conference* (2000). 12

[Nic11] NICOLL G.: Rhombic dodecahedron infill, 2011. http://www.thingiverse.com/thing:12535. 13

[NZDS15] NEWMAN S. T., ZHU Z., DHOKIA V., SHOKRANI A.: Process planning for additive and subtractive manufacturing technologies. *CIRP Annals-Manufacturing Technology 64*, 1 (2015), 467–470. 20

[Oxm11] OXMAN N.: Variable property rapid prototyping. *Virtual and Physical Prototyping 6*, 1 (2011), 3–31. 13

[Par05] PARK S. C.: Hollowing objects with uniform wall thickness. *Computer-Aided Design 37*, 4 (2005), 451–460. 12

[PBJSH16] PRÉVOST R., BÄCHER M., JAROSZ W., SORKINE-HORNUNG O.: Balancing 3D models with movable masses. In *Proceedings of the Vision, Modeling and Visualization Workshop* (2016). 7

[PC15] PAN Y., CHEN Y.: Smooth surface fabrication based on controlled meniscus and cure depth in microstereolithography. *Journal of Micro and Nano-Manufacturing 3*, 3 (2015), 031001. 16

[PCB01] PARK S.-M., CRAWFORD R. H., BEAMAN J. J.: Volumetric multi-texturing for functionally gradient material representation. In *ACM symposium on Solid modeling and applications* (2001). 14, 16

[PDG99] PHAM D., DIMOV S., GAULT R.: Part orientation in stereolithography. *Int. J. Adv. Manuf. Tech. 15*, 9 (1999), 674–682. 5, 9, 10

[PFV*11] PASKO A., FRYAZINOV O., VILBRANDT T., FAYOLLE P.-A., ADZHIEV V.: Procedural function-based modelling of volumetric microstructures. *Graphical Models 73*, 5 (2011), 165–181. 13

[PGMH16] PENG H., GUIMBRETIÃLRE F., MCCANN J., HUDSON S.: A 3d printer for interactive electromagnetic devices. In *Proceedings of UIST* (2016), pp. 553–562. 21

[PHL11] PARK I. B., HA Y. M., LEE S. H.: Dithering method for improving the surface quality of a microstructure in projection microstereolithography. *Int. J. Adv. Manuf. Tech. 52*, 5-8 (2011), 545–553. 16

[PJB04] PRIEDEMAN JR W. R., BROSCH A. L.: Soluble material and process for three-dimensional modeling, September 2004. US Patent 6,790,403. 11

[PK08] PAVIĆ D., KOBBELT L.: High-resolution volumetric computation of offset surfaces with feature preservation. *Computer Graphics Forum 27*, 2 (2008), 165–174. 11

[Pol98] POLYJET: Polyjet technology, 1998. https://en.wikipedia.org/wiki/Objet_Geometries. 3

[PP12] PHATAK A. M., PANDE S.: Optimum part orientation in rapid prototyping using genetic algorithm. *J. Manuf. Syst. 31*, 4 (2012), 395–402. 9, 10

[PRM14] PEREIRA T., RUSINKIEWICZ S., MATUSIK W.: Computational light routing: 3d printed optical fibers for sensing and display. *ACM Trans. Graph. 33*, 3 (2014), 24. 7

[PVRD03] PANDEY P., VENKATA REDDY N., DHANDE S. G.: Slicing procedures in layered manufacturing: a review. *Rapid Prototyping J. 9*, 5 (2003), 274–288. 2

[PWLSH13] PRÉVOST R., WHITING E., LEFEBVRE S., SORKINE-HORNUNG O.: Make It Stand: Balancing shapes for 3D fabrication. *ACM Trans. Graph. 32*, 4 (2013), 81:1–81:10. 7

[PZM*15] PANETTA J., ZHOU Q., MALOMO L., PIETRONI N., CIGNONI P., ZORIN D.: Elastic textures for additive fabrication. *ACM Trans. Graph. 34*, 4 (2015), 135:1–135:12. 14, 20

[PZZC12] PAN Y., ZHAO X., ZHOU C., CHEN Y.: Smooth surface fabrication in mask projection based stereolithography. *Journal of Manufacturing Processes 14*, 4 (2012), 460–470. 16

[QS03] QU X., STUCKER B.: A 3D surface offset method for STL format models. *Rapid Prototyping J. 9*, 3 (2003), 133–141. 11

[QZQ11] QIU Y., ZHOU X., QIAN X.: Direct slicing of cloud data with guaranteed topology for rapid prototyping. *Int. J. Adv. Manuf. Tech. 53*, 1 (2011), 255–265. 16

[QZY13] QI D., ZENG L., YUEN M. M. F.: Robust slicing procedure based on surfel-grid. *Comp. Aided Des. Appl. 10*, 6 (2013), 965–981. 15

[RC95] RICHARD D. T., CRAWFORD R. H.: Optimizing part quality with orientation. The University of Texas at Austin, 1995. 8

[RC97] REEVES P. E., COBB R. C.: Reducing the surface deviation of stereolithography using in-process techniques. *Rapid Prototyping J. 3*, 1 (1997), 20–31. 6

[RCM*14] REINER T., CARR N., MĚCH R., ŠTAVA O., DACHSBACHER C., MILLER G.: Dual-color mixing for fused deposition modeling printers. *Computer Graphics Forum 33*, 2 (2014), 479–486. 17

[RL15] RANEY J. R., LEWIS J. A.: Printing mesoscale architectures. *MRS Bulletin 40*, 11 (2015), 943–950. 14, 21

[RL16] REINER T., LEFEBVRE S.: Interactive modeling of support-free shapes for fabrication. In *Eurographics - Short Papers* (2016). 9, 20

[RNDA13] ROLLAND-NEVIERE X., DOERR G., ALLIEZ P.: Robust diameter-based thickness estimation of 3D objects. *Graphical Models 75* (2013), 279–296. 7

[Ros07] ROSEN D. W.: Computer-aided design for additive manufacturing of cellular structures. *Comp. Aided Des. Appl. 4*, 5 (2007). 1, 13, 16, 20

[RR86] ROSSIGNAC J. R., REQUICHA A. A.: Offsetting operations in solid modelling. *Comput. Aided Geom. D. 3*, 2 (1986), 129–148. 11

[SARW*15] SITTHI-AMORN P., RAMOS J. E., WANGY Y., KWAN J., LAN J., WANG W., MATUSIK W.: MultiFab: A machine vision assisted platform for multi-material 3D printing. *ACM Trans. Graph. 34*, 4 (2015), 129:1–129:11. 3

[SAZ15] SALONITIS K., AL ZARBAN S.: Redesign optimization for manufacturing using additive layer techniques. *Procedia CIRP 36* (2015), 193–198. 1, 21

[SBK15] SIKDER S., BARARI A., KISHAWY H. A.: Global adaptive slicing of NURBS based sculptured surface for minimum texture error in rapid prototyping. *Rapid Prototyping J. 21*, 6 (2015), 649–661. 14

[SBR*15] SCHUMACHER C., BICKEL B., RYS J., MARSCHNER S., DARAIO C., GROSS M.: Microstructures to control elasticity in 3D printing. *ACM Trans. Graph. 34*, 4 (2015), 136:1–136:13. 14, 20

[SCH13] SAVAGE V., CHANG C., HARTMANN B.: Sauron: Embedded single-camera sensing of printed physical user interfaces. In *Proceedings of UIST* (2013), ACM, pp. 447–456. 21

[SDW*16] SONG P., DENG B., WANG Z., DONG Z., LI W., FU C.-W., LIU L.: CofiFab: Coarse-to-fine fabrication of large 3D objects. *ACM Trans. Graph. 35*, 4 (2016), 45:1–45:11. 5

[SFLF15] SONG P., FU Z., LIU L., FU C.-W.: Printing 3D objects with interlocking parts. *Comp. Aided Geom. D. 35–36* (2015), 137 – 148. 18

[SGHS98] SHADE J., GORTLER S., HE L.-W., SZELISKI R.: Layered depth images. In *Proceedings of SIGGRAPH* (1998), pp. 231–242. 15

[SHB96] SABOURIN E., HOUSER S. A., BØHN J. H.: Adaptive slicing using stepwise uniform refinement. *Rapid Prototyping J. 2*, 4 (1996), 20–26. 14

[SHB97] SABOURIN E., HOUSER S. A., BØHN J. H.: Accurate exterior, fast interior layered manufacturing. *Rapid Prototyping J. 3*, 2 (1997), 44–52. 14

[SHCW94] SACHS E., HAGGERTY J., CIMA M., WILLIAMS P.: Three-dimensional printing techniques, Aug 1994. US Patent 5,340,656. 3, 17

[Sig95] SIGMUND O.: Tailoring materials with prescribed elastic properties. *Mechanics of Materials 20*, 4 (1995), 351–368. 13

[SIM16] STEUBEN J. C., ILIOPOULOS A. P., MICHOPOULOS J. G.: Implicit slicing for functionally tailored additive manufacturing. *Computer-Aided Design 77* (2016), 107–119. 13

[Sli11] SLICER: Slicer. http://slic3r.org/, 2011. 18

[SLS*05] STARLY B., LAU A., SUN W., LAU W., BRADBURY T.: Direct slicing of STEP based NURBS models for layered manufacturing. *Computer-Aided Design 37*, 4 (2005), 387–397. 15

[SNDM03] SHIN K.-H., NATU H., DUTTA D., MAZUMDER J.: A method for the design and fabrication of heterogeneous objects. *Materials & Design 24*, 5 (2003), 339–353. 16

[SPP04] SHIN H., PARK S., PARK E.: Direct slicing of a point set model for rapid prototyping. *Comp. Aided Des. Appl. 1*, 1-4 (2004), 109–115. 16

[SRDH12] SINGAMNENI S., ROYCHOUDHURY A., DIEGEL O., HUANG B.: Modeling and evaluation of curved layer fused deposition. *J. Mater. Process. Tech. 212*, 1 (2012), 27–35. 16

[SU14] SCHMIDT R., UMETANI N.: Branching support structures for 3D printing. In *ACM SIGGRAPH 2014 Studio* (2014), ACM, p. 9. 11

[SVB*12] STAVA O., VANEK J., BENES B., CARR N., MĚCH R.: Stress relief: improving structural strength of 3D printable objects. *ACM Trans. Graph. 31*, 4 (2012), 48:1–48:11. 7, 20

[SW*94] SUH Y. S., WOZNY M. J., ET AL.: Adaptive slicing of solid freeform fabrication processes. In *Solid Freeform Fabrication Symposium* (1994), pp. 404–411. 14

[SZH12] SAVAGE V., ZHANG X., HARTMANN B.: Midas: Fabricating custom capacitive touch sensors to prototype interactive objects. In *Proceedings of UIST* (2012), pp. 579–588. 21

[TB98] TYBERG J., BØHN J. H.: Local adaptive slicing. *Rapid Prototyping J. 4*, 3 (1998), 118–127. 14

[Tib14] TIBBITS S.: 4d printing: Multi-material shape change. *Architectural Design 84*, 1 (2014), 116–121. 21

[TJ11] TELEA A., JALBA A.: Voxel-based assessment of printability of 3D shapes. In *Proceedings of ISMM* (2011), Springer, pp. 393–404. 7

[TJ13] TAUFIK M., JAIN P. K.: Role of build orientation in layered manufacturing: a review. *International Journal of Manufacturing Technology and Management 27*, 1-3 (2013), 47–73. 2

[TJ14] TAUFIK M., JAIN P. K.: Volumetric error control in layered manufacturing. In *ASME Int. Design Engineering Technical Confs and Computers and Information in Engineering Conf* (2014). 6

[TMV*16] THOMPSON M. K., MORONI G., VANEKER T., FADEL G., CAMPBELL R. I., GIBSON I., BERNARD A., SCHULZ J., GRAF P., AHUJA B., ET AL.: Design for additive manufacturing: Trends, opportunities, considerations, and constraints. *CIRP Annals-Manufacturing Technology 65*, 2 (2016), 737–760. 1, 21

[TPR04] THRIMURTHULU K., PANDEY P. M., REDDY N. V.: Optimum part deposition orientation in fused deposition modeling. *Int. J. Mach. Tool. Manu. 44*, 6 (2004), 585–594. 5, 9, 10

[TSB*16] TOWNSEND A., SENIN N., BLUNT L., LEACH R., TAYLOR J.: Surface texture metrology for metal additive manufacturing: a review. *Precision Engineering 46* (2016), 34–47. 6

[UIM12] UMETANI N., IGARASHI T., MITRA N. J.: Guided exploration of physically valid shapes for furniture design. *ACM Trans. Graph. 31*, 4 (2012), 86:1–86:11. 21

[UKY*15] ULU E., KORKMAZ E., YAY K., OZDOGANLAR O. B., KARA L. B.: Enhancing the structural performance of additively manufactured objects through build orientation optimization. *Journal of Mechanical Design 137*, 11 (2015), 111410. 9, 10

[UPSW16] UMETANI N., PANOTOPOULOU A., SCHMIDT R., WHITING E.: Printone: interactive resonance simulation for free-form print-wind instrument design. *ACM Trans. Graph. 35*, 6 (2016), 184. 7

[US13] UMETANI N., SCHMIDT R.: Cross-sectional structural analysis for 3D printing optimization. In *SIGGRAPH Asia 2013 Technical Briefs* (2013), pp. 5:1–5:4. 9, 10

[VC11] VAEZI M., CHUA C. K.: Effects of layer thickness and binder saturation level parameters on 3D printing process. *Int. J. Adv. Manuf. Tech. 53*, 1-4 (2011), 275–284. 18

[VFLS13] VOLPATO N., FRANZONI A., LUVIZON D. C., SCHRAMM J. M.: Identifying the directions of a set of 2D contours for additive manufacturing process planning. *Int. J. Adv. Manuf. Tech. 68*, 1-4 (2013), 33–43. 17

[VGB14a] VANEK J., GALICIA J., BENES B.: Clever support: efficient support structure generation for digital fabrication. *Computer Graphics Forum 33*, 5 (2014), 117–125. 11

[VGB*14b] VANEK J., GALICIA J., BENES B., MĚCH R., CARR N., STAVA O., MILLER G.: Packmerger: A 3D print volume optimizer. *Computer Graphics Forum 33*, 6 (2014), 322–332. 19

[Vil99] VILAR R.: Laser cladding. *J Laser Applicat. 11* (1999), 64–79. 3

[VKF∗92] VUYYURU P., KIRSCHMAN C., FADEL G., BAGCHI A., JARA-ALMONTE C.: A NURBS-based approach for rapid product realization. In *5th Int. Conf. on Rapid Prototyping* (1992). 15

[VKWM16] VIDIMČE K., KASPAR A., WANG Y., MATUSIK W.: Foundry: Hierarchical material design for multi-material fabrication. In *ACM User Interface Software and Technology Symposium* (2016). 17

[VM06] VARADHAN G., MANOCHA D.: Accurate Minkowski sum approximation of polyhedral models. *Graphical Models 68*, 4 (2006), 343–355. 11

[VWRKM13] VIDIMČE K., WANG S.-P., RAGAN-KELLEY J., MATUSIK W.: OpenFab: A programmable pipeline for multi-material fabrication. *ACM Trans. Graph. 32*, 4 (2013), 136:1–136:12. 13, 17

[Wan11] WANG C. C. L.: Approximate boolean operations on large polyhedral solids with partial mesh reconstruction. *IEEE Trans. Vis. Comput. Graphics 17*, 6 (2011), 836–849. 15

[WBHP12] WILLIS K., BROCKMEYER E., HUDSON S., POUPYREV I.: Printed optics: 3d printing of embedded optical elements for interactive devices. In *Proceedings of UIST* (2012), pp. 589–598. 7

[WC13] WANG C. C., CHEN Y.: Thickening freeform surfaces for solid fabrication. *Rapid Prototyping J. 19*, 6 (2013), 395–406. 7

[WCT∗15] WANG W., CHAO H., TONG J., YANG Z., TONG X., LI H., LIU X., LIU L.: Saliency-preserving slicing optimization for effective 3D printing. *Computer Graphics Forum 34*, 6 (2015), 148–160. 14

[Wei09] WEIDONG Y.: Optimal path planning in rapid prototyping based on genetic algorithm. In *Chinese Control and Decision Conference* (2009), pp. 5068–5072. 18

[WKW16] WU J., KRAMER L., WESTERMANN R.: Shape interior modeling and mass property optimization using ray-reps. *Computers & Graphics 58* (2016), 66–72. 7

[WLC10] WANG C. C. L., LEUNG Y.-S., CHEN Y.: Solid modeling of polyhedral objects by layered depth-normal images on the GPU. *Computer-Aided Design 42*, 6 (2010), 535–544. 15

[WM13a] WANG C., MANOCHA D.: Efficient boundary extraction of BSP solids based on clipping operations. *IEEE Trans. Vis. Comput. Graphics 19*, 1 (2013), 16–29. 8

[WM13b] WANG C. C., MANOCHA D.: GPU-based offset surface computation using point samples. *Comp.-Aided Design 45*, 2 (2013), 321–330. 12

[WMJC02] WAH P. K., MURTY K. G., JONEJA A., CHIU L. C.: Tool path optimization in layered manufacturing. *Iie Transactions 34*, 4 (2002), 335–347. 18

[WMP∗97] WEISS L. E., MERZ R., PRINZ F. B., NEPLOTNIK G., ET AL.: Shape deposition manufacturing of heterogeneous structures. *J. Manuf. Syst. 16*, 4 (1997), 239. 16

[WPGM16] WU R., PENG H., GUIMBRETIÈRE F., MARSCHNER S.: Printing arbitrary meshes with a 5DOF wireframe printer. *ACM Trans. Graph. 35*, 4 (2016), 1–9. 16

[WSP∗00] WU H., SACHS E., PATRIKALAKIS N., BRANCAZIO D., SERDY J., JACKSON T.: Distributed design and fabrication of parts with local composition control. In *Proceedings of the 2000 NSF Design and Manufacturing Grantees Conference* (2000). 17

[WW16] WANG L., WHITING E.: Buoyancy optimization for computational fabrication. *Computer Graphics Forum 35*, 2 (2016), 49–58. 7

[WWLZ04] WU Y. F., WONG Y. S., LOH H. T., ZHANG Y. F.: Modelling cloud data using an adaptive slicing approach. *Computer-Aided Design 36*, 3 (2004), 231–240. 16

[WWY∗13] WANG W., WANG T. Y., YANG Z., LIU L., TONG X., TONG W., DENG J., CHEN F., LIU X.: Cost-effective printing of 3D objects with skin-frame structures. *ACM Trans. Graph. 32*, 6 (2013), 177. 5, 11, 13

[WWZW16] WU J., WANG C. C. L., ZHANG X., WESTERMANN R.: Self-supporting rhombic infill structures for additive manufacturing. *Computer-Aided Design* (2016). 13

[WZK16] WANG W., ZANNI C., KOBBELT L.: Improved surface quality in 3D printing by optimizing the printing direction. *Computer Graphics Forum 35*, 2 (2016), 59–70. 6, 8, 10, 14

[XLCB15] XU H., LI Y., CHEN Y., BARBIČ J.: Interactive material design using model reduction. *ACM Trans. Graph. 34*, 2 (2015), 18. 20

[XLL16] XU W.-P., LI W., LIU L.-G.: Skeleton-sectional structural analysis for 3D printing. *J Comp. Science and Tech. 31*, 3 (2016), 439–449. 7

[XLW99] XU F., LOH H., WONG Y.: Considerations and selection of optimal orientation for different rapid prototyping systems. *Rapid Prototyping J. 5*, 2 (1999), 54–60. 5

[YBSH16] YAMAN U., BUTT N., SACKS E., HOFFMANN C.: Slice coherence in a query-based architecture for 3D heterogeneous printing. *Computer-Aided Design 75-76* (2016), 27–38. 13

[YCL∗15] YAO M., CHEN Z., LUO L., WANG R., WANG H.: Level-set-based partitioning and packing optimization of a printable model. *ACM Trans. Graph. 34*, 6 (2015), 214:1–214:11. 19

[YLFW02] YANG Y., LOH H., FUH J., WANG Y.: Equidistant path generation for improving scanning efficiency in layered manufacturing. *Rapid Prototyping J. 8*, 1 (2002), 30–37. 12

[YLQ10] YANG P., LI K., QIAN X.: Topologically enhanced slicing of MLS surfaces. In *30th Computers and Information in Engineering Conference, Parts A and B* (2010), vol. 3. 16

[YQ08] YANG P., QIAN X.: Adaptive slicing of moving least squares surfaces: Toward direct manufacturing of point set surfaces. *J. Comput. Inf. Sci. Eng. 8*, 3 (2008), 031003. 16

[YZRL08] YUKSEK K., ZHANG W., RIDZALSKI B. I., LEU M. C.: A new contour reconstruction approach from dexel data in virtual sculpting. In *3rd Int Conf on Geometric Modeling and Imaging* (2008). 15

[ZBHK15] ZHANG Y., BERNARD A., HARIK R., KARUNAKARAN K. P.: Build orientation optimization for multi-part production in additive manufacturing. *J. Intell. Manuf.* (2015), 1–15. 18

[ZCG∗16] ZHAO H., CHEN B., GU F., HUANG Q.-X., GARCIA J., CHEN Y., TU C., BENES B., ZHANG H., COHEN-OR D.: Connected Fermat spirals for layered fabrication. *ACM Trans. Graph. 35*, 4 (2016), 1–10. 12, 17, 18

[ZCo07] ZCORPORATION: Zprinter 450 hardware manual. http://www.3dsystems.com/, 2007. 18

[ZCYK11] ZHOU C., CHEN Y., YANG Z., KHOSHNEVIS B.: Development of multi-material mask-image-projection-based stereolithography for the fabrication of digital materials. In *Annual solid freeform fabrication symposium* (2011). 3

[ZJ15] ZHANG Z., JOSHI S.: An improved slicing algorithm with efficient contour construction using STL files. *Int. J. Adv. Manuf. Tech. 80*, 5-8 (2015), 1347–1362. 15

[ZL08] ZHOU S., LI Q.: Microstructural design of connective base cells for functionally graded materials. *Materials Letters 62*, 24 (2008), 4022–4024. 14

[ZL09] ZHANG W., LEU M. C.: Surface reconstruction using dexel data from three sets of orthogonal rays. *J. Comput. Inf. Sci. Eng. 9*, 1 (2009), 011008. 15

[ZLP∗15] ZHANG X., LE X., PANOTOPOULOU A., WHITING E., WANG C. C. L.: Perceptual models of preference in 3D printing direction. *ACM Trans. Graph. 34*, 6 (2015), 215:1–215:12. 6, 8, 10, 11

[ZLQ∗11] ZENG L., LAI L. M.-L., QI D., LAI Y.-H., YUEN M. M.-F.: Efficient slicing procedure based on adaptive layer depth normal image. *Computer-Aided Design 43*, 12 (2011), 1577–1586. 15

[ZPLZ07] ZHANG W., PENG X., LEU M. C., ZHANG W.: A novel contour generation algorithm for surface reconstruction from dexel data. *J. Comput. Inf. Sci. Eng. 7*, 3 (2007), 203–210. 15

[ZPZ13] ZHOU Q., PANETTA J., ZORIN D.: Worst-case structural analysis. *ACM Trans. Graph. 32*, 4 (2013), 137:1–137:12. 7

[ZSMS14] ZHOU Y., SUEDA S., MATUSIK W., SHAMIR A.: Boxelization: Folding 3D objects into boxes. *ACM Trans. Graph. 33*, 4 (2014), 71:1–71:8. 19

[ZWCJ11] ZHAO H., WANG C. C. L., CHEN Y., JIN X.: Parallel and efficient boolean on polygonal solids. *Vis. Computer 27*, 6-8 (2011), 507–517. 15

[ZXW*15] ZHANG X., XIA Y., WANG J., YANG Z., TU C., WANG W.: Medial axis tree – an internal supporting structure for 3D printing. *Comput. Aided Geom. D. 35-36* (2015), 149–162. 13

[ZY01] ZHU W. M., YU K. M.: Dexel-based direct slicing of multi-material assemblies. *Int. J. Adv. Manuf. Tech. 18*, 4 (2001), 285–302. 15, 16