


3DATAPRINT

On the use of 3D printing for scientific visualization

June 15th, 2018



P.E.C. Melis¹

(Visualization group, SURFsara)

S.J. Marrink², S. Thallmair

(Molecular Dynamics group, University of Groningen)

A. Blass³, V. Spandan Arza

(Physics of Fluids group, University of Twente)

J. Nijkamp⁴, N. Janssen

(Surgical Oncology department, Antoni van Leeuwenhoek Hospital - Netherlands Cancer Institute)

G. Závodszky⁵

(Computational Science group, University of Amsterdam)



¹paul.melis@surfsara.nl,
⁵g.zavodszky@uva.nl

²s.j.marrink@rug.nl,

³a.blass@utwente.nl,

⁴j.nijkamp@nki.nl,

Abstract

We present the results of a case-based study in using 3D printing for applications in scientific visualization. The goal was to get insight into the practical sides of producing 3D prints from 3D scientific data and make a (qualitative) evaluation of the added values of using 3D printing for visualization purposes.

A set of different use cases, from varying scientific domains, is presented. The 3D prints produced for these cases were deemed a success, with more detail in the 3D prints possible than expected, while providing alternative interpretation and communication options compared to on-screen 2D/3D visualizations. For example, in providing more natural interaction with the 3D model and easier communication by physically pointing out elements of interest, or just handing over a 3D print. For the cases involving complex 3D surfaces the advantages appeared to be in a better sense and understanding of spatial structures compared to on-screen 3D visualization and more natural interaction.

Unforeseen effects that were observed were that a high amount of detail in a 3D print can look impressive (more so than seeing the same data on a screen) and inspired reactions when confronting people with 3D printed versions of data they were familiar with. In cases where the size of 3D models being visualized approaches the limits of graphical hardware used for interactive visualization a 3D print can be comparatively “low-end” alternative, as no hardware is needed to view a 3D print, the view is always naturally stereoscopic, etc.

Although 3D printing has been commercially available for some time now and prices are no longer prohibitive for small- to medium-sized scientific models, several challenges for regular application to scientific visualization remain, depending on the actual use case and goal of a 3D print. The limitations mostly have to do with the amount of detail than can be printed related to physical size, limitations in using color and the effort needed to turn a 3D dataset into a printable 3D model. The latter involves design and creative aspects that in most cases involves manual effort which is hard to automate and requires some experience and skill.

3D printing for visualization will most likely not be applied as commonly as on-screen visualization and will stay a fairly unique way of visualizing, compared to current on-screen visualization techniques that are flexible and ubiquitous. It should therefore not be viewed as an alternative to on-screen visualization, but as an addition to it that triggers different senses, speaks more to intuition and so provides different opportunities.

Contents

1	Introduction	4
2	Related work	6
3	3D printing workflow	8
3.1	Step 1: Preprocess	8
3.2	Step 2: Design	9
3.3	Step 3: Print	11
3.3.1	Automated 3D model inspection	11
3.3.2	Manual 3D model inspection	11
3.3.3	Input preparation	11
3.3.4	3D printing	12
3.3.5	Postprocessing	12
3.3.6	Output inspection	13
3.4	Step 4: Evaluate	13
4	Printing and design limitations	13
4.1	Upper limit in the amount of detail per print	13
4.2	Limited use of color	14
4.3	Minimal feature size	15
4.4	Maximum printable size	16
4.5	Budgetary constraints	17
4.6	Miscellaneous issues	17
5	Use cases	19
5.1	Turbulent structures - University of Twente	19
5.1.1	Taylor-Couette flow	19
5.1.2	Meandering features in sheared thermal convection	21
5.2	Synthetic cell model - University of Groningen	24
5.3	Oncological surgery support - Netherlands Cancer Institute	26
5.4	Flow properties around aneurysms - University of Amsterdam	30
5.4.1	Virtual stenting	32
6	Conclusions	33
	References	36
	Appendices	38
A	Upper bound on the amount of printable polygons	38
B	Links to online background material	39
C	Tools	39

List of Figures

1	3D printing workflow	8
2	Turning a social network into a 3D printed object [20]	9
3	The influence of material choice: the exact same 3D model printed in two different materials (and with different printing techniques)	10
4	Sub-steps of the “print” step	11
5	Visible staircase effect due to layered printing, which varying orientation within the 3D prints	12
6	Remains of support structures added during input preparation	12
7	Red blood cell 3D model, with heatmap showing issues with thickness (model after [22])	15
8	Packed 3D models for a 3D printing batch, shown in Autodesk Netfabb	16
9	Visualization of a Taylor-Couette simulation	19
10	Taylor-Couette 3D print	20
11	Meandering structures for different simulation parameters (left-to-right: increasing Reynolds number, top-to-bottom: increasing Rayleigh number	21
12	Full print of the 8 parts (left); Detail view (right)	22
13	Large print in 8 parts (keyboard for scale)	23
14	SynthetiCell model, artist impression by Graham Johnson	24
15	Final model (16x15x21 cm)	26
16	3D printed scan container with “phantom” tissue sample used for testing	26
17	CT-scan of a tissue sample inside a container. Note that the container is clearly visible	27
18	Location marked on the 3D scan (in 3 view directions) and corresponding hole on the container to use for inking that location	27
19	Evolution of the container design	29
20	Part of a simulated blood vessel with an aneurysm in blue. Cross-section at the red circle for different flow parameters	30
21	Aneurysm print	31
22	Color difference in the discs between on-screen visualization (background) and 3D print	31
23	Cubes	32
24	Imaginary model with “cubes” and “ridges”	38

1 Introduction

This report presents and discusses the results of the 3DATAPRINT project, which was an activity performed within the SURF Open Innovation Lab at SURFsara in the period July 2017 to March 2018. Within the project a number of research partners from Dutch research institutes worked together with SURFsara to investigate the use and applicability of 3D printing for scientific visualization purposes.

3D printing has become accessible to anyone as the cost of creating 3D prints has dropped significantly. Consumer-quality 3D printers cost only a few hundred euros these days and the most common printing material (plastic) is very cost-effective in use. And for those not wishing to own or operate a 3D printer themselves several commercial 3D printing service providers have entered the market.

Although there definitely have been numerous applications of 3D printing within science (see the Related work section below) a 3D printed dataset still appears to be somewhat of an exotic phenomenon in a scientific context. This might be because it is *perceived* to be of little value, involve too much effort to create or to be too expensive. In general, the major applications of 3D printing appear to be in manufacturing, prototyping, art and design. Hence, applying it to scientific visualization is a bit of a niche. But it is therefore an interesting topic in itself, as one encounters different limitations and challenges.

The focus in the project was therefore on gaining insight in the practical sides of producing 3D prints. What is needed to turn a scientific dataset into a 3D print? How much effort (and cost) is involved? What possibilities and limitations need to be taken into account? Are the results useful and do they add something compared to current on-screen and on-paper visualizations? Are there shared elements in these areas between the different use cases?

All 3D printing use cases shown are based on existing scientific research by the relevant project partner:

- Turbulent structures - Physics of Fluids group, University of Twente
- Synthetic cell model - Molecular Dynamics group, University of Groningen
- Oncological surgery support - Surgical Oncology department, Antoni van Leeuwenhoek Hospital - Netherlands Cancer Institute
- Flow properties around aneurysms - Computational Science group, University of Amsterdam

For some use cases 3D printing had been applied earlier on a similar use case, in a pilot study.

The already mentioned Related Work section discusses previous applications of 3D printing to science and visualization. But 3D printing is still a developing field and it is therefore useful to continuously check what is possible in terms of technology and software tooling, but also commercial services.

Some topics are out of the scope of this report:

- Describe the history of 3D printing
- Providing an extensive overview of all 3D printing techniques and materials. Where relevant we provide some background in the text.
- Describe how to set up and operate your own 3D printer. In contrast, we used a commercial service provider (Shapeways) for all the 3D prints shown in this report.
- Teach how to create 3D printable designs from scientific data. This is a topic in-itself that we could devote a lot of time to. It is part experience, part creativity, part experimenting, etc. We do detail some of the design choices and experiences in handling the use cases.

For readers that want more background in some of these areas, please check out the excellent online references listed in Appendix B.

As mentioned above a commercial service provider was used to produce the 3D prints, based on 3D models prepared within the project. Operating a 3D printer can be a time-consuming task: print jobs need to be prepared, started and monitored, printing material needs to be refilled, the printer cleaned and maintained, etc. Furthermore, the actual printing is only part of the full process to go from 3D dataset to finished 3D print. 3D models need to be prepared in the software controlling the 3D printer, which is a separate step from producing the 3D models in the first place. And once the printer has done its job the resulting model usually needs to be checked, cleaned, polished or any other number of finishing steps applied. Finally, to be able to print in different materials in general implies using different 3D printers and 3D printing techniques, adding to the burden of maintaining a 3D printing pipeline.

Different service providers will have a different selection of available materials, printing options and pricing. For this project Shapeways was chosen as the printing service provider, as it provides a wide range of materials, has a nice online workflow and we had prior experiences with their services. Obviously, there are more providers on the market and this report does not attempt to promote one over the other. Where applicable we do describe differences that we noted. One consequence of choosing a specific service provider is that some of the material names used in this report denote specific products of Shapeways. Usually, the same material and/or printing technique is available from other providers as well, but under a different name.

One note about the Shapeways material names used is in order. Recently (May 2018) Shapeways phased out a number of materials and renamed almost all others. Although maybe somewhat confusing to readers of this report we have chosen to use the original material names, as both the old and new names are currently used in the material descriptions on the Shapeways website. Furthermore, we used at least one material (High Definition Acrylate) that was recently phased out, but for which we have not evaluated available alternatives.

We hope that the results presented in this report can serve as inspiration for further applications of 3D printing in science, in particular for visualization purposes.

2 Related work

Since 3D printing has become available numerous applications of the technique to scientific subjects have been published. Here, we list a number of them, covering different fields of science. Some of these applications are straightforward “data to 3D print”, others are closer to art and design. Some of these publications mention that the 3D models used for the prints are publicly available. This is an interesting addition to making the underlying *data* for a publication available.

Within astronomy and astrophysics Clements *et al.* [1] visualized the Cosmic Microwave Background (CMB) radiation, as measured by the Planck mission, in a sphere-like 3D print the size of an orange. The printed objects “capture the essence of the CMB anisotropies [...] and can provide a non-visual appreciation of the CMB for both the sighted and unsighted.” Interestingly, the focus appears to be more on *feeling* structures than viewing them. They mention having to scale up the bumps in the 3D model for the anisotropies to become large enough to be felt and seen, as the differences are actually on the order of 1 part in 100,000 [2]. In [3] 3D prints are created of density structures attained from N-body simulations. In the 3D prints they discover “wall-like features that are invisible in two-dimensional projections”, thereby providing a benefit of the 3D prints over on-screen/on-paper visualizations. An example of 3D printing applied to stellar winds is given in [4].

Somewhat more abstract results are produced with 3D prints of mathematical concepts. Knill and Slavkovsky [5] use 3D printing to visualize proofs in mathematics. They mention that “[as was] already known to educators in ancient Greece, models allow to bring mathematics closer to the public. The new 3D printing technology makes the realization of such tools more accessible than ever.” and that even with modern display technologies on the web like WebGL and SVG “the possibility to manipulate an object with our bare hands is still unmatched.” An example of better interpretation of mathematical structures is [6], where hyperbolic honeycombs are 3D printed as “representations of only edges are difficult to interpret from a two-dimensional picture, but a 3D printed sculpture is much more useful.” In [7] 3D printing is applied to several different mathematical concepts, such as projections, tilings and 4D hypercubes. For the latter example a “zoetrope” is created, which consists of a large rotating disc holding a set of 3D printed models, lit by a strobe light. The rotation of the disc together with the correct strobing frequency gives the illusion of the models staying in place, while changing over time. Over 200 mathematically-themed prints by Henri Segermans are listed in [8], copies of which can even be directly ordered. Notable examples are shown in more detail in [9] where different types of curved surfaces are 3D printed using dozens of interlocking, but movable, parts.

Weiss *et al.* [10] create and analyze a 3D print of a small gene-gene interaction network, suggesting that “[...] there may be advantages unique to experiencing this data through a physical medium that should not be ignored”.

In [11] a 3D printed shape of the MIT campus, based on LIDAR data, was used as a base object for projection of color-coded “big data” related to social media, such as geo-tagged Twitter messages in the area. Another proposed application is to show traffic

hot spots. The benefits listed are that the 3D printed model provides a “shared and tangible 3D visualization” for a group of people. In [12] an installation is shown that allows a user to interactively show different statistics of regions of Italy, in effect providing an articulated bar graph shaped as a country. More generally, representing abstract data in physical form is also known as “data physicalization”, “physical visualization” or a “data sculpture”. In [13] quite a number of 3D printed examples of this are shown. But there’s a larger collection of other examples, some over a hundred years old, that were not 3D printed and constructed by other means. These can serve as inspiration for producing 3D printed models.

Most of the above examples need effort in producing printable 3D models, as the underlying data isn’t directly usable. But when data already has a 2D or 3D structure making a physical print is a natural extension to existing visualization techniques. In [14] 3D printing is applied to (amongst others) urban planning and geology. They note an advantage of 3D printing in the longevity of the results, as “printing can take a few minutes to few days depending on the level of details and scale of application. But the models can stay for years.” They also provide a systematic comparison of 3D digital visualization versus 3D printed models, noting that 3D printed models are easier to handle and explore by inexperienced users compared to 3D digital visualizations. The possibility of having a larger group of people around a physical model is another benefit listed. In [15] (published in Dutch only) 3D terrain elevation models based on LiDAR scans are printed, including with color overlays. Among the advantages of 3D printing listed are being able to use 3D prints for easier interpretation compared to a 2D visualization and being able to use the resulting 3D prints without any dependence on IT facilities.

Structures in fluid flow are printed in [16], which notes the challenges to print small-scale features and the need for extra construction steps, such as gluing printed pieces together and adding wires to suspend 3D printed elements.

Interesting research in the area of 3D scanning and then printing faces is shown in [17], where they note that “3D prints of scanned faces often appear flat and lifeless.” They make the connection with sculptors “that often deviate from geometric accuracy in order to enhance the appearance of their sculpture.” The described work in effect enhances physical prints of 3D scans by operating on the underlying 3D models.

In medicine 3D printing has found many applications, from producing instructional models and medical equipment to more patient-oriented applications, such as producing well-fitting prosthetic parts and creating cranium replacement in neurosurgery. Printing biological material, such bone and ear cartilage, is under active research. An example of a more visualization-oriented application is described in [18], where 3D printing is applied to diffusion-tensor MRI imaging. They contrast the use of 3D virtual reality (VR) applications for showing brain data with “review[ing] the data with a very robust, natural, and fast haptic interface: their own hands”. Furthermore, they note initial results are encouraging and end users are excited about the possibilities of the technique, realizing that “the data visualization is not dependent on computing power [and] the frame rate does not drop for more complex models.” However, limitations and challenges are also clear in that pieces need to be structurally sound and extra supports

added might obstruct the view of the data being visualized. In [19] pre-operative imaging of kidney tumors was used to produce a 3D printed model of both the tumor and the surrounding healthy kidney. This “allowed surgeons and patients to pre-operatively visualize both the pre- and post-tumor-resection kidney status”, thereby helping both groups: surgeons could better plan complex operations, while patients could be better informed about such an operation and the underlying disease. Post-operative analysis was done to compare the 3D printed models with the actual surgical specimens removed and were found to be “nearly identical”.

3 3D printing workflow




Figure 1: 3D printing workflow

A general iterative workflow to go from (scientific) data to a corresponding 3D print is shown in figure 1. Most of this workflow is straightforward, but we describe some details below for each of the steps.

One thing to note in this workflow is the distinction between *3D model* and *3D model (printable)*. A printable 3D model will need to satisfy several limitations inherent to the chosen printing technique and material. Overcoming these limitations was a major part of the effort in producing 3D prints for the use cases described in this report.

3.1 Step 1: Preprocess

In principle, any type of input data can be used to create a 3D print, but certain types of data obviously lend themselves better than others. Data that is already in the form of a 3D model requires minimal preprocessing. More specifically, a model consisting of a *3D polygonal model* is ideal, as that is the de-facto input required for making a 3D print. Such a 3D model consists of polygons: small planar facets connected together to form a surface. 3D polygonal models are widely used in 3D modeling, gaming, but also scientific research. More details on the 3D model are discussed below with the *design* step.

Another type of often used 3D scientific data is volumetric data. This type of data somehow needs to be converted to a 3D surface, for example, by extracting an isosurface. Point-based 3D datasets (say, an N-body simulation of a galaxy) can be turned into a polygonal model by, for example, placing small 3D spheres at the point locations, although this would produce unconnected spheres which would need to be connected in some way in a second step. See Figure 2 for an example (from [20]).




Figure 2: Turning a social network into a 3D printed object [20]

Geographic data from a GIS system might only consist of 2D outlines and therefore need to be solidified and extruded to 3D shapes. Other GIS data might have associated height values and therefore already form a 3D surface.

Abstract data, like a network of nodes and edges, can be turned into a 3D model by placing 3D objects (like spheres) at node positions and representing edges with, for example, 3D cylinders, as in Figure 2.

All in all, depending on the type of data and goal of the 3D print, this step might require a somewhat creative approach to produce an appropriate 3D model. Conversely, this step (together with the *design* step) might also highlight that there are too many hurdles to go from data to printable 3D model. In practice the *preprocess* step will tightly interact/overlap with the next step, *design*.

3.2 Step 2: Design

The rationale for having a separate *design* step has to do with the specific limitations of 3D printing and the restrictions it places on the input 3D model. Even though the data to print might already consist of a 3D model, it might not be printable in that form.

First, the 3D model needs to be *watertight*. Put simply, this means that the model must consist of a closed surface with a hollow interior, without any holes giving access to the interior space. Another restriction is that the 3D model must have a certain minimal “thickness” everywhere (determined by the chosen material and printing technique). More formal checks on watertightness and other properties of the 3D model can be performed by the user in tools like MeshLab, Autodesk NetFabb or Blender (see Appendix C). When using a printing service provider model checks will also be performed in the *print* step. In some cases the online processing tools of the provider used can automatically fix certain deficiencies of the 3D model, such as minor thickness issues.

Challenges that need to be solved in this step are:

- Producing a 3D model that matches the goal for the 3D printed version
- Overcoming printing, material and budgetary limits, for example
 - creating a watertight 3D model
 - handling features in the 3D model that are too small or fragile to print
 - balancing target size of the print versus available printing budget
- Making design decisions, such as
 - the exact material to use
 - the best way to incorporate color
 - where to add supporting geometry, with minimal occlusion of the model being shown and the right center of gravity for models that need to stand by themselves
 - if (and where) to add escape holes to reduce the volume of printing material used
 - designing interlocking parts for 3D prints that needs to be separable



Figure 3: The influence of material choice: the exact same 3D model printed in two different materials (and with different printing techniques)

In Figure 3 the importance of choosing a material with the appropriate properties is highlighted. It shows two prints of the same 3D model, but in different materials: High-definition Acrylate (left) versus PLA. The contrast in surface precision, cleanliness and layer visibility is enormous. Note that the left model was printed using a process where a photo-reactive resin is cured with light, while the right model uses Fused Deposition Modeling, in which layers of molten plastic are deposited.

The “design” step is the most crucial step in successfully producing a 3D print. In many cases this step will highlight the (im)possibility of producing a 3D print with the needed features. In this step creativity is usually needed, to work around some of the design issues. In the project we frequently found in this step that we had to be less ambitious in the 3D prints we wanted to make.




Figure 4: Sub-steps of the “print” step

3.3 Step 3: Print

When using a printing service provider the *print* step can be treated mostly as a black box: the provider takes the printable 3D model created by the user as input and produces the final 3D print as output. However, the step can be broken down into a number of sub-steps, see figure 4. The actual actions in these steps will vary by printing material and technique, so some knowledge of these steps is useful. And when not using a printing service provider, but using a 3D printer one owns, most of these steps will need to be done anyway, increasing the total effort to produce a print.

In these sub-steps the print can get rejected by the automated and manual checks, or fail during printing, hence the arrow from *print* to *evaluate* in the workflow (Figure 1) above.

3.3.1 Automated 3D model inspection

This step is usually already performed when uploading a 3D model to the service provider. It includes checks on watertightness, minimal feature size, maximum printable size and other material and printing limits. These are all based on the geometric features of the 3D model and can take a substantial time for models with high detail and complex features (15-60 minutes was not uncommon for complex models used in the project).

3.3.2 Manual 3D model inspection

After a model has been accepted by the automated checks and an order is placed there’s usually a manual inspection done on the 3D model at the provider’s side. One experience with this step in the project is that the operator used different software for checking feature size of the 3D model which had stricter limits. So even though the automatic checks succeeded, a print would occasionally get rejected in this step. Furthermore, based on experience the operator performing the inspection might provide useful feedback if he/she deems the 3D print unlikely to succeed.

3.3.3 Input preparation

In most cases a user’s 3D model is added to a printing batch of multiple (tens) of other 3D prints and assigned to a 3D printer, as part of the necessary steps to prepare for printing. Here, specific factors related to printing technique, printer model and material

are taken into account. Most of these are only of interest to the printer operator, but one notable choice in this step is the orientation in which the 3D model is produced in the 3D printer.

A number of printing techniques produce a visible staircase effect on the printed surface due to the layered printing approach used, shown in Figure 5. The user might have a preference in which direction these layers occur and orient the input 3D model appropriately. If no orientation is explicitly set by the user the operator might change it based on experience or to use less space in the printing batch.



Figure 5: Visible staircase effect due to layered printing, which varying orientation within the 3D prints

Another preparation step for certain printing techniques is the addition of support structures, for example, for resin-based printing. Here, complex (overhanging) structures need to be reinforced to counteract warping. The added supports will be removed manually after printing by the service provider.

3.3.4 3D printing

Actually producing the 3D print on a 3D printer, including a cooling down period for techniques that use high temperatures during printing.

3.3.5 Postprocessing

After a print batch is finished the 3D print needs to be retrieved. For powder-based materials this involves taking the print out of the bed of powder, doing a rough cleanup, etc. For some materials the prints are cleaned in an alcohol bath. Prints to be made in color might have been printed in white and need to be dyed. For polished materials a polishing step will be performed. Supports added in the input preparation step are removed. Certain materials might need further curing before further postprocessing. All in all, the steps here are very specific to the used material and printing technique.



Figure 6: Remains of support structures added during input preparation

Figure 6 shows a mistake than can happen when supports need to be added: too much material was removed during cutting away the supports. The two hemispherical models are basically different sizes of the same model (although the number of holes in the hemisphere differs). Note that for the left model remains of the supports are visible on the top rim, while for the right model they are visible on the outside of the hemisphere, indicating a different orientation of the model during 3D printing. The four pins on the rim of the right model have been inadvertently cut away from the left model.

3.3.6 Output inspection

Some 3D prints might have obvious structural failures, but other times the printed model needs to be compared to the input 3D model to see if the print was a success. This is a manual process in most cases, but might get (more) automated in future using 3D scans of the printed model and doing shape-matching against the input 3D model.

3.4 Step 4: Evaluate

The output from the whole workflow is one or more 3D prints. These will need to be evaluated by the end-user for printing defects, correctness, usability, etc. Apart from evaluating the 3D print results themselves the “visual impact” of the model can be judged. If needed, the model design, printing parameters (such as material and size), or even the input data might have to be altered to improve the 3D print.

In the project, most use cases needed more than one iteration, but sometimes the first print was satisfactory.

In some cases a part of a 3D model was printed as a test, to see how a specific material looked in terms of color and/or surface smoothness. Test pieces where also printed to judge material strength and precision. These test prints can be seen as a first iteration of the workflow.

4 Printing and design limitations

To zoom in on the most often encountered limitations in the project we describe a number of them in detail, roughly in descending order of importance.

4.1 Upper limit in the amount of detail per print

The printing provider used in the project advertises a limit of 1,000,000 polygons per printable 3D model². Other providers use similar limits. Furthermore, there is also a limit on the file size of the uploaded 3D model of 64MB. However, a model file can be compressed in a ZIP file to effectively increase this limit.

²Although the actually enforced limit didn't appear to be exactly this number.

For design-oriented prints these limits are reasonable, but many scientific datasets and corresponding 3D models can easily have orders of magnitude more detail. One rationale given for a detail limit is that the printing precision and maximum printable size (see section below) together result in an upper bound on the maximum size of the input 3D model in terms of polygons. There *is*, of course, an upper limit, but it would appear to be much higher than 1 million (see Appendix A).

It is unclear if the limit is an actual limitation in the used 3D printing equipment and/or software at the provider's end or that this was merely chosen as a conservative limit for increasing the chances of successfully making a 3D print. It might also be used to limit the processing burden on the provider's back-end systems, as each uploaded model is automatically checked for deficiencies and printability, which can take quite a bit of compute power.

In practice, this limit implies either having to reduce the surface complexity of the 3D model to print, or to cut up the 3D model in smaller pieces that are individually printable and can be used to reconstruct the larger object. In the use cases described below we applied both approaches.

4.2 Limited use of color

For many materials 3D prints can only be made in a single color, in most cases from a limited set of color choices. Combining parts printed in different colors is an option, but in practice this means designing parts to interlock in some way, so they can be combined after being printed separately.

For scientific visualization having only a few colors, even just one per separately printed part, is very limited. In visualization color is used freely and plays an important role in data interpretation, with detailed color maps being used to represent values such as pressure or velocity. Making 3D prints with similar color fidelity, so-called *full-color* prints, is possible, but this is limited in practice to sandstone-like materials. These materials are more brittle and not as detailed as many of the often used (single color) plastics. Larger models also tend to become quite heavy. Similarly, printing only those parts that need color in a full-color material, with other parts printed in regular plastics is possible, but will take extra effort. Especially as the different types of parts will need to be physically combined in some way.

And even among color choices for a single-color material there can be differences. For example, Shapeways' "strong and flexible plastic" comes in roughly a dozen color options, some of which only come in a polished variant (e.g. there's Red Polished, but not Red). The colored options (apart from white and black) are dyed after printing and apparently the dyeing process leaves a residue that needs to be removed, hence the need for polishing. This polishing step is done with a fairly rough procedure involving a large vibrating bucket filled with small ceramic pellets that smoothen the surface of the 3D print [21]. But these pellets can get stuck in small holes or cause small features of a print to get damaged. There is, however, no option to print in these colors *without* applying the polishing step. So for fragile models to be printed in plastic the number of

colors might effectively be limited to the two non-polished options: white and black.

Finally, the colors used in a 3D model when viewer on a screen will be somewhat different from the actual colors in the corresponding 3D print (the “Flow properties around aneurysms” use case in the next section shows an example of this). This is due to differences in the range of possible colors between screen and 3D print and the way former color space is converted to the latter. The magnitude of the difference is, of course, material dependent and might be acceptable in a lot of applications. But somewhat more worrying is that a printing provider might optimize the colors per printed 3D model, in effect altering the mapping from colors in the 3D model to the 3D print, leading to slightly unpredictable and unreproducible results. For example, for Shapeways’s Sandstone material the following remark is made about the material’s color³:

Q: How well will the colors of the product match the colors of the file?

A: The colors of the product will not perfectly match the colors of digital file or render. When printing in this material, some variation in color shade is unavoidable, especially with difficult-to-produce colors such as browns and skin tones. *We optimize each print run to make individual models look great, [so you] may see differences in color when looking at multiple prints of the same model side by side.*

4.3 Minimal feature size

The material chosen for a print, together with the associated printing technique, imposes inherent limitations on, for example, print precision, the minimal size of small features, and the minimum thickness of unsupported structures.

There are limits to the smallest detail visible by the human eye and the current layer thickness for high-resolution 3D printing techniques is on the order of the width of a human hair. Printing in thinner layers might be interesting for certain manufacturing applications, but won’t add much visually as the extra detail can’t be perceived in most cases. However, the layer thickness isn’t the whole story: small protruding features and thin walls need to be much thicker than a single layer to become physically strong enough, in most cases several to dozens of material layers.

In practice, print precision is on the order of 0.1-0.5 mm for ridges on a surface, while the minimal thickness for an unsupported protruding wire is in the range 1-4 mm. For each material and printing technique detailed instructions on this subject and other limitations is made available by printing providers. In general these

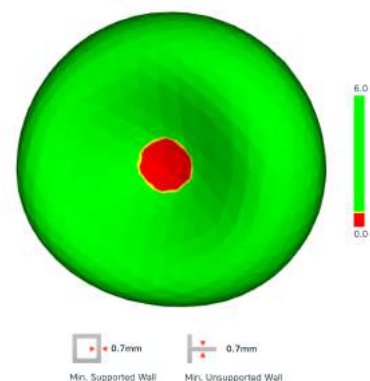


Figure 7: Red blood cell 3D model, with heatmap showing issues with thickness (model after [22])

³<https://www.shapeways.com/materials/sandstone>

instructions contain hard technical and material limitations, but also design guidelines to prevent unwanted bending or breaking of a 3D print or its small features.

A 3D model can be checked against the minimal feature sizes using software tools, either one the user has installed themselves or one provided by the printing provider as part of the online workflow. In most cases areas in the 3D model where small or thin features are problematic are fairly predictable. But an example where these limitations can lead to unexpected issues is shown in Figure 7. A 3D model of a simulated red blood cell⁴ has a center part that is too thin: around 0.15 mm, where the recommended minimal wall thickness is 0.7 mm. Even though the full model is large and thick enough to be printed, the center part might not come out as intended.

Another effect of printing details that are close to the minimal feature and layer size for a material is that the print will not perfectly match the 3D model. In the process of preparing the 3D printer input data the 3D model is discretized into elements that are on the order of the minimal feature size that the printer can produce. For small features this discretization will have a larger relative error than for medium to large features and can lead to a noticeable difference in shape or smoothness of small features compared to the 3D model.

The bottom line is that the smallest details and thicknesses in the 3D model to print must be matched against the detail limitations of the material and technique. Reducing detail and/or scaling up the 3D model is often necessary, but this might not always be acceptable for detailed scientific models.

4.4 Maximum printable size

On the macro-level a 3D print will have a maximum size, dictated by the physical limits imposed by the 3D printer used. 3D models that would be too large physically need to be scaled down. But if a large print is really desirable a 3D model must be split into parts by the user, printed separately and then later re-assembled (see the “Turbulent structures” use case in the next section).

The maximum size limitation also has an important interaction with limitations on the minimum feature size mentioned above, as a model that is scaled up to make small features printable can become too large to print.




Figure 8: Packed 3D models for a 3D printing batch, shown in Autodesk Netfabb

⁴This is one cell out of many hundreds in a full simulation

4.5 Budgetary constraints

The available budget for 3D printing will limit the physical size of a 3D print, as the cost of a print is in principle based on the *printing volume* (the amount of material used) and *machine space*, apart from the material choice and printing technique. The machine space cost is due to several printing orders, from different customers, usually being printed in a single batch, where the 3D models are packed together as shown in figure 8. The more space a model takes in this packing, the higher the cost. So even a model that uses very little printing material, for example a hollow tube with a thin wall, might still use a lot of material space.

As an indication: the cost for each of the separate parts of the printed models in this report was in the range 50 to 300 Euros.

4.6 Miscellaneous issues

Here we describe some elements not directly related to the design of a 3D print, but that might be of interest when printing scientific models.

Material properties

Depending on the goal for a 3D print material properties other than those related to the actual printing might be of importance. For example, some materials cannot withstand high temperatures. Colors might fade over time under the influence of UV light. Other materials are brittle and therefore not applicable for prints that are to be handled or inspected often. The color and handling properties are probably most important for visualization prints. In general, the material information provided by the printing provider contains all the details.

Service provider workflow

Printing providers usually have an online workflow that, after uploading a 3D model, performs a number of checks on the 3D model. These checks can take a very long time (sometimes 15-60 minutes), especially for complex models with a high number of polygons and/or intersecting parts. The existence of a queue of models to be checked (from other customers) on the providers system might also cause delays.

Frequently, issues are detected by the online checks and the user has to improve the 3D model locally and re-upload, leading to a slow iterative process. Rescaling a model in the online workflow usually also causes the model checks to be re-executed, taking up time.

By performing a number of 3D model checks and optimizations locally the user can increase the success rate of the online checks:

- Make sure the geometry is watertight, fits within the maximum printable size, doesn't have features smaller than the minimal feature size, etc. See Appendix C for software that can perform these kinds of checks.

- Give the model the correct physical scale, so rescaling isn't necessary in the online workflow.

One limitation in the workflow of the provider used was that for uploaded 3D model files only 3 units could be picked: millimeters, inches and meters. So for a 3D model which used centimeter as units it needed to be rescaled after being uploaded. Directly designing the model in of the available units saves having to do the rescaling step.

- For models that consist of lots of separate (but overlapping) 3D parts merging all the parts into a single mesh with a boolean join might be faster than letting the online tools perform the same operation.

It would be interesting if service providers could provide an offline tool to let users locally perform the model checks that are currently being done online. This should save the providers processing power, while allowing users to get their models correct before uploading. If this tool could be integrated in 3D modeling packages and tools that users use for designing their models this would save time spent on exporting, compressing and uploading.

Data storage and ownership

One downside of using a service provider, as in this project, is that the 3D data to print needs to be transferred to such a provider. As most of the providers use a web-based workflow, where 3D models are uploaded as part of placing an order, one particular aspect is of interest: what happens with the 3D data (and derived content, like preview images) after it has been uploaded? And who owns the rights to the design once uploaded?

Provider Shapeways is fairly clear in this respect ([23], section 7): “[we] shall not use, modify or display your 3D Models or derivatives thereof” except for manufacturing to fulfill orders or when using the option to sell 3D prints through the Shapeways website. Intellectual property rights on 3D models are also retained by the user. The user has the option to delete their account, but no explicit mention is made what happens with the uploaded 3D content in this case.

3D Hubs, a different service provider, appears to have similar terms in [24], but phrases them somewhat unclear, by using the overlapping terms “3D Content” (3D printable files supplied by a user) and “Content” (3D Content, images, videos, text, or any other content submitted by a User on the Website). Section 6.3 of [24] reads - emphasis added - “By submitting Content (*except for the User’s 3D Content*) [...] the User grants 3D Hubs a [liberal] license to use, edit, modify, reproduce, distribute, prepare derivative works of, display, perform, and otherwise fully exploit the Content in connection with the Website and/or the Services [...]”. What exactly this means in practice isn't clear, nor whether, for example, preview images derived from 3D models fall under “Content” or “3D Content”.

Other providers might or might not apply similar terms and conditions. In case tight control over uploaded 3D data and derived elements like preview images is of importance (which is not unlikely in scientific research) thoroughly checking the terms and

conditions of any 3D printing provider used is recommended.

5 Use cases

Within the project four use cases from different scientific fields and with a different focus for the 3D prints were chosen. Each use case is discussed separately below, together with the resulting 3D prints. Where necessary the research context for the use case is described in some detail.

5.1 Turbulent structures - University of Twente

Fully developed turbulence is one of the big unsolved problems in fluid dynamics. The main question is the distribution of rare events, which has important implications for, for example, flight safety. We approach this problem from a fundamental point of view using numerical simulations.

In particular, we focus on generating high levels of turbulence in canonical systems such as Rayleigh-Benard (RB) convection and Taylor-Couette (TC) flows. In RB convection, a fluid in a closed box is heated from below and cooled from above, while in TC flows, the flow is confined between two independently rotating coaxial cylinders. These two systems are paradigmatic systems in fluid dynamics to test various new concepts. These systems are interesting because (i) they are mathematically well defined by the extended Navier-Stokes equations with their respective boundary conditions; (ii) exact global balance relations between the respective driving and dissipation can be derived and (iii) they are experimentally accessible with high precision, thanks to simple geometries and high symmetries.

These systems encompass a wide a range of flow structures for different drivings. These are typically visualized on a 2D screen through iso-surfaces or volume rendering in visualization software. A clear and more concise picture of the multi-scale turbulent structures in the flow can be better explained to the 'non-expert' by visualizing these structures with the help of 3D printing.

5.1.1 Taylor-Couette flow

Figure 9 shows a visualization of the Taylor-Couette simulation data in a visualization application (VisIt)⁵. Two isosurfaces are extracted from the data, for specific values. This is the usual way in which the data is visualized.

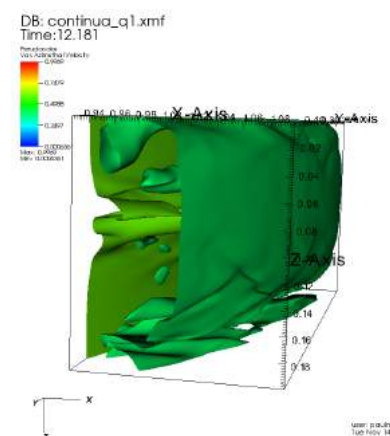


Figure 9: Visualization of a Taylor-Couette simulation

⁵Note that the view is vertically flipped compared to the 3D prints shown later

The 3D models of the two isosurfaces were processed in Blender in order to make them printable:

- Reduced the number of polygons in the surfaces, by about 90%. As the surfaces have large areas with low curvature and detail the reduction did not significantly alter the overall shape.
- Added thickness to the surfaces, to make the infinitely thin isosurfaces watertight
- Added the cylinder walls (the domains from the simulation) at the correct positions
- Removed loose parts that did not contribute significantly to the flow structure
- Added supports and “foot” planes, so the models can support themselves standing up. A somewhat difficult issue was making sure that the center of mass of the models was not pulling them over, given the limited options for connecting the feet parts.
- Split the foot plane in two parts so the two isosurface models can be separated, while still allowing them to be placed against each other in the correct relative position

As no color is needed printing in a strong plastic material is the most cost-effective and durable solution. The 3D prints came out quite nicely and with a good fit:




Figure 10: Taylor-Couette 3D print

Being able to view the prints as a single domain as well as separate parts is a nice feature. This leads to quite natural interaction with the parts compared to on-screen visualization, where showing side-by-side views of the two parts would be a similar method. The printed 3D model is, of course, fixed in terms of the data that it shows, while an interactive on-screen visualization can usually be changed as part of analyzing the data.

All in all producing the 3D prints for this use case was fairly straightforward without too many surprises. The addition and design of the foot plates took a bit of effort, but could perhaps be automated in a 3D modeling package, given input 3D model parts.

5.1.2 Meandering features in sheared thermal convection

A second simulation shows shear-driven flow in a heat-driven Rayleigh-Benard simulation. In this simulation a fluid is present between a cooled top plate and a heater bottom plate. The plates also move in relative opposite directions, causing the fluid to shear. The shearing motion together with the temperature difference results in stable thermal structures in the fluid. Depending on the chosen simulation parameters different meandering structures emerge, as shown in Figure 11.




Figure 11: Meandering structures for different simulation parameters (left-to-right: increasing Reynolds number, top-to-bottom: increasing Rayleigh number)

The center case in the picture above, with two meandering structures, was chosen for printing as it has a good balance between amount of detail and number of structures.

The challenges in 3D printing this use case where the amount of detail in the isosurface model and small features that extend without support. Preparation steps for 3D printing were:

1. Reduction of the isosurface complexity from 33 million polygons to 8 million
2. Adding thickness to the (infinitely thin) isosurface
3. Splitting the model into 8 parts, that each were printed separately. Each part had roughly 1 million polygons and could therefore be printed directly.
4. Cutting the full model in individual parts resulted in small pieces of 3D geometry protruding from one part over a neighbouring part without any connection. These, dozen or so, small features were re-attached manually in Blender.

Printing was done in strong and flexible plastic. Currently, the different parts do not have extra geometry for lining up and/or connecting them into a single 3D model. The “printing in parts” approach is a first step in trying to experiment with 3D printing models with high surface complexity, in order to get around the upper limit in model complexity. Further reduction of the complexity might be a different way around the limitation, but would lead to loss of detail in the printed model. Some of the small overhanging features are very fragile. Scaling up the model might be wise if it is intended to be handled a lot (as opposed to being placed only on display).

The full model, consisting of the 8 parts, is around 40x18cm. The parts don't perfectly



Figure 12: Full print of the 8 parts (left); Detail view (right)

align, perhaps due to printing inaccuracies and/or slight bending of small protruding features during printing and cooling. That such an amount of detail was 3D printable at all was a nice surprise. The 3D model contains lots of small and thin structures that protrude without support, but they are easily expressed in the material used. A few really thin wire-like parts will probably break off if the print is handled a lot, though.

We do consider the full print to be a success, as it shows the structures in the flow at different levels without too much compromise. The periodic nature of the simulation domain is also present in the printed version, meaning the tiles can be shifted one position along an axis and still match up. The detail is challenging to take in, which might be due to the white color which reduces the contrast somewhat (less shadowing). One possible visual improvement we came up with would be to give the flatter areas between the meandering features a darker color, to make the meandering features stand out better. This would, however, need to be done manually after production of the print, due to the limitations of printing in color described earlier.

One nice feature of the 3D print is that it is very easy and natural to pick up to study in detail or to point out something, compared to performing the same actions with an on-screen 3D visualization. The size of the 3D model in this use case, in terms of number of polygons, also approaches what can currently be visualized interactively on a PC with a single high-end GPU, especially when using advanced lighting effects, such as shadowing. Enabling such lighting effects is usually desirable in on-screen visualization as they provide visual clues for understanding the spatial structure of the data. But enabling these effects costs graphics performance, reducing the smoothness of the interactive view.

Another benefit of a 3D print is that it naturally provides a stereoscopic view, which helps in understanding the spatial structure of the data. On-screen stereoscopic visualizations are feasible, but require dedicated display hardware. Using stereoscopic rendering for increased spatial perception, or even fully immersive Virtual Reality with a headset like the HTC Vive, will further decrease rendering performance, at some point

leading to an unusable viewing experience.⁶ So for highly detailed 3D datasets a “low-tech” 3D printed model has some clear advantages over on-screen interactive 3D visualization: no PC hardware and/or GPU is needed for viewing, “rendering performance” is infinite, the view is perfectly stereoscopic, lighting and shadowing are realistic, etc.

As a test both fluid use cases were printed a second time in a different size: the Taylor-Couette case was printed 50% smaller and the meandering features case almost 200% larger, see Figure 13 for the latter.




Figure 13: Large print in 8 parts (keyboard for scale)

For the Taylor-Couette case the difference in scale did not noticeably influence the value of the 3D print with respect to data interpretation, probably due to the low surface complexity. So a smaller (and therefore cheaper) print can be used for similar datasets in future.

The very large meandering features print of Figure 13 was mostly aimed at producing an impressive display of the simulation results. The turbulent structures are, unsurprisingly, easier to see and acknowledge in the larger print. The larger print did not contain more detail, it was merely a larger version of the print shown in Figure 12. Producing a new print with higher detail by splitting the original unreduced isosurface into more than 8 parts would have been an option, but was not chosen due to time constraints. One effect of the large print size is bending in some of the printed parts, most noticeable in the second part from the left on the bottom row. This might be caused by temperature effects after the print was delivered or already have been present after printing.

One conclusion from the creation of these prints is that printing in separate parts is a

⁶A quick test in Unity3D with the 3D model in VR using a high-end GPU and HTC Vive quickly showed the limits: the full model of 8 million polygons could be displayed with a bit of stuttering in the frame rate, but only with very simple lighting (no shadowing). The display resolution with this VR device is also fairly low, so small details in the 3D model were hard to observe.

feasible approach to overcome the physical size limits in the 3D printing process. Splitting input 3D models into separate parts is feasible in a semi-automated way in 3D modeling packages and so does not involve too much manual effort. And for repeated printing of the same type of data, for example the isosurface models shown here, the processing steps from input data to printable parts can probably be completely automated with only a final visual human check on the output. Whether such a 3D printing workflow provides added value is obviously another matter and case-dependent.

5.2 Synthetic cell model - University of Groningen

This use case involved visualization of the structure of a biological cell using a 3D printed model. The input 3D model was produced for the SyntheticCell project, a 10-year FET Flagship initiative to produce a computational synthetic model. In Figure 14 an impression of a synthetic cell is shown, representing the three basic processes taking place in a living cell: cell fueling (green), DNA processing (orange/red) and cell division (blue). The model is useful both for education and instruction.




Figure 14: SynthetiCell model, artist impression by Graham Johnson

The input 3D model to print was commissioned from a 3D artist. The model was quite well constructed in terms of detail and makeup of individual parts. But it still took quite a while before it was acceptable for 3D printing. This was due to non-manifold and duplicate geometry being present in the model (perhaps originating from further 3D editing after the model was delivered by the artist). This is where a bottleneck in the online workflow becomes apparent: after uploading a new version of the 3D model the automatic checks done by the printing provider could take 10-15 minutes or sometimes would not complete at all. Secondly, the upload and initial view of the model worked, but subsequent scaling of the model (from millimeter units to centimeter units) would lead to the model getting stuck in the provider's processing pipeline. As mentioned earlier, having a workflow where the automatic checks can be performed locally before uploading, including providing clear error feedback, would save quite a bit of time. The 3D model used here was actually somewhat "exotic", as it consisted of a large number of separate, but intersecting, pieces of geometry totaling around 700,000 triangles. We suspect the automatic online checks attempt to merge all pieces into a single 3D mesh as a first step for easier processing. Doing this merging step locally helped in substantially reducing the online check time, at the cost of extra local geometry preprocessing.

As use of color was essential in this model and the properties of the full-color sandstone wasn't exactly known several test prints were made to evaluate material strength and printability. The sandstone material was stronger than expected, although breakable by hand when applying a medium amount of force. The alternative, printing in strong white plastic, would have implied adding color to the model after 3D printing as a post-process, which would involve quite a bit of manual work.

The 3D geometry in the model representing DNA-processing parts, specifically the DNA strands, comes with several challenges for 3D printing:

- The DNA strands are very thin (2-3mm thickness when printed), increasing the chance they will break, either immediately or over time. A strong-and-flexible-plastic print of these strands, was partly successful, in that the thicker strands stay intact, while the very thin strands broke on delivery. The same parts in sandstone were delivered without the strands being included, these probably already broke at the service provider during printing.
- As can be seen in the illustration above these DNA parts are floating freely in the cell interior. In case of a 3D print they will need to be attached to some part of the cell wall, but when using supporting geometry (e.g. straight cylinders) these will quickly need to be as thick (or thicker) as the wires themselves. Adding these supports also leads to features being included in the 3D print that are not part of the cell, so they should be either clearly recognizable as such (e.g. different color) and/or placed in such a way that they don't obstruct the view too much. How to include loose parts is obviously a general challenge when 3D printing a model.

A surprising feature when 3D printing the thin DNA strands in plastic was the flexibility and strength of the material, even at such elongated and thin shapes. This opens the door to including such shapes for flexible parts and features in 3D prints. One element of the original input model that was simply left out are the free-floating small yellow parts visible in Figure 14. We currently do not have a good way to add these features to the 3D print without adding a lot of supports that occlude the actual geometry of interest.

Based on properties of the test pieces, full-color sandstone was chosen to print the "outer shell" part of the complete model shown in Figure 14, as it allows both printing in color while providing a relatively sturdy model that can be picked up and inspected. The free-floating DNA parts were printed in white strong-and-flexible plastic and colored by hand, as this involved only a small amount of work.

The first printed model arrived with some of the larger protruding pieces broken off, so improvements in the 3D model were made to strengthen these parts. The loose DNA strands were suspended within the cell's interior using nylon strands. The final model is shown in Figure 15.

Apart from the long cycle time of the online workflow the amount of manual geometry preparation to make the model printable was quite substantial and involved the most effort of all the use cases presented. Again, this model here is somewhat different from the other ones, due to the large number of individual and intersecting parts mentioned above.

All in all quite a bit of manual post-processing of the 3D print was involved, but this is defensible for a print that is to be used for communication and education. Here, the ultimate goal and use of the 3D print comes into play: it is expected to be shown and handled many times, so spending more effort on creating it is worth it.




Figure 15: Final model (16x15x21 cm)

5.3 Oncological surgery support - Netherlands Cancer Institute

The context of this use case is a surgical procedure called “breast lumpectomy”, which is used in the treatment of breast cancer. In this type of surgery a breast tumour is removed, together with some of the surrounding tissue. After surgery the removed specimen is microscopically inspected by a pathologist to see if tumor tissue is present on the borders of the specimen. If no tumour cells are found then the surgery was successful. But the presence of tumour cells on the border implies the tumour was not fully removed. In that case, the patient needs follow-up surgery or radiotherapy.

It takes roughly a week from removal of the specimen during surgery to the result of the microscopic inspection. One method for substantially reducing this time is to perform a micro-CT scan of the removed specimen *during* surgery. If this scan shows that tumour cells are present on the border of the specimen the surgeon can be notified for removal of extra tissue without needing a second operation.

The CT scan will also show where on the specimen tumour cells are present at the surface. This location is then marked with surgical ink, so further pathological inspection can be performed on the region. Currently, the CT scan process is still too time-consuming for direct feedback to the surgeon during the surgery, but this should improve in future.

One challenge is to relate the tumour surface location as seen in CT images to the



Figure 16: 3D printed scan container with “phantom” tissue sample used for testing




Figure 17: CT-scan of a tissue sample inside a container. Note that the container is clearly visible

corresponding physical location on the specimen. For this, a 3D printed container was developed. The specimen is placed in the container, which is then placed in the CT scanner. See figures 16 and 17.

The container is marked with orientation labels as part of the 3D print, so the tissue can be placed and oriented correctly in the container.⁷ Because the tissue is correctly oriented in the scanning container any tumour tissue located on the scan images can be related to the physical position on the specimen and therefore to a location in the patient.

The container also contains a regular pattern of holes all around. The orientation labels on the container are readable in the CT-scan visualization. This allows choosing the best hole to use for marking the tumour location on the specimen surface with ink. See figure 18. The 3D print is thereby helping to make the connection between the virtual world of the CT-scan visualization and the physical world of the specimen.




Figure 18: Location marked on the 3D scan (in 3 view directions) and corresponding hole on the container to use for inking that location

⁷The removed tissue is marked with small pieces of suture thread by the surgeon (two long ones for the feet direction, etc). These indicate the orientation that the specimen had in the patient.

The container and its design had a number of criteria to satisfy:

- The used 3D printed material must have an adequate CT density: it must be visible on the CT scan images, as otherwise it cannot be related to the tissue being scanned. But its X-ray density must not be too high, as that would occlude the tissue inside the container.
- It must be possible to sterilize the container after each use. In the present use this means placing the container in a washer that uses liquid of 80°C. The washing procedure also places quite some stress on the container.
- Containers of different sizes are needed, to match the size of the tissue specimen. Using a single large container is not preferable as that leads to less precise marking with ink for small tissue specimens.

In figure 19 successive designs of the container are shown, highlighting the iterative design process. Changes were made to strengthen the container due to stresses imposed during the sterilization procedure, for easier access to the sample inside the container through the holes and adding a third piece on which the container is placed. This piece serves as a tight connection between the container and the rotating platform of the CT scanner, which is necessary for generating a high-precision CT scan. The piece also captures any liquid material from the specimen. The latest evolution of the design was programmatically generated in a 3D modeling package, while the earlier designs were hand-modeled. Being able to programmatically generate new container designs (with varying sizes and/or number of holes) is a nice improvement over the earlier manual designs.

The chosen material, “High Definition Acrylate”, can withstand the needed high temperatures and stresses during sterilization and has low CT density⁸. Note that the high-definition nature of the material (which is printed in layers of 50 microns) isn’t very important here. The material was chosen as it was the only one to satisfy the earlier listed criteria. One disadvantage of the material is that it is roughly 6 times more expensive than the simple strong-and-flexible-plastic in terms of material volume. This currently isn’t prohibitive, but will be more noticeable with larger models, as volume (and therefore price) scales with N^3 when scaling the model uniformly by a factor of N .

A downside of this specific material with the service provider used is that the maximum printable size is quite limited: 8.4 x 6.3 x 16 cm. This is much smaller than for strong-and-flexible plastic (at 65 x 35 x 55 cm). The main reason for this is the use of a 3D printer with less machine space. This has limited the container size used so far. One way around the limitation would be to split the container parts in two halves, which would need to be glued together after printing. As this is not very attractive, we have not attempted this yet. An easier alternative would be to use a different material that is suitable, but this would involve evaluating multiple material choices for CT density.

The different printed containers are currently being used in a clinical study at the Netherlands Cancer Institute (to be reported on separately) and so far perform their intended function well. They allow matching up a 3D on-screen visualization (the CT scan data) with a physical location on the specimen. The concept of creating a 3D printed

⁸This material has since been phased out at Shapeways



Figure 19: Evolution of the container design

container that shows visibly in a CT scan could be applied to non-medical contents as well, for example in cultural heritage objects.

5.4 Flow properties around aneurysms - University of Amsterdam

This use case deals with cerebral aneurysms: lesions on brain vessel walls. By the age of 65 most of the population harbors at least one aneurysm, however luckily, most of them are symptomless and never rupture. However, if they rupture they are often fatal as they cause brain hemorrhage (one of the main categories of strokes).

Current methods for risk assessment have a prediction rate of around 70%. Seemingly, it is not possible to increase this number based on the geometry or flow conditions of the aneurysm, unraveling the pathogenesis is necessary to identify weaker segments in the vessel wall. In current research the University of Amsterdam targets the very first steps of the aneurysmal formation process. They argue, that (at least in the specific section investigated) aneurysms are caused by mechanical effects on the innermost cell layer of the blood vessel, the endothelium. This is caused by the inhomogeneous, periodically changing loads of the pulsating blood flow.

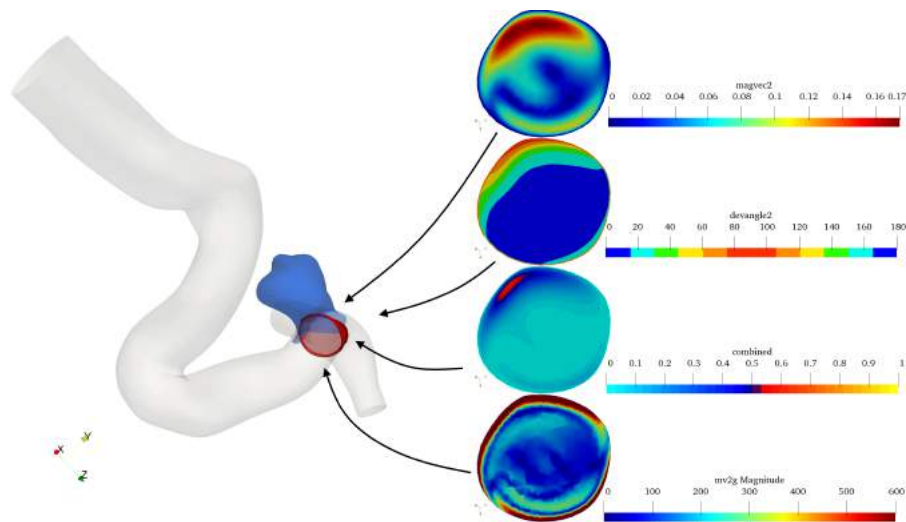


Figure 20: Part of a simulated blood vessel with an aneurysm in blue. Cross-section at the red circle for different flow parameters

Flow simulations that take these mechanical effects into account have been developed and Figure 20 shows one of the results. The blood vessel is shown in light grey, with the aneurysm in blue. The vessel geometry is based on 3D rotational angiography of an actual vessel containing an aneurysm. The geometry of the vessel before the aneurysm formed was derived from the angiographic data using a digital vessel reconstruction technique. The reconstructed geometry was then used for flow simulations.

The flow in these simulations is pulsating, which causes the flow direction to vary over time. At the red circular cross-section of the vessel different flow parameters, for a single timestep in the simulation (at the systolic peak), are shown on the right. The top

one, “magvec2”, is proportional to the magnitudes of the secondary velocity components (those that lie in the plane of the cross-section). The second cross-section shows the relative magnitudes of primary and secondary velocities. Technically, it shows the deviation of the flow direction from the primary direction, pointing along the centerline.

This deviation changes during a heartbeat, exerting a varying load on the vessel wall. In yellow-to-red areas the flow changes substantially over time, in blue areas much less. The third cross-section shows the combination of the first two. It shows those areas, where the velocity is high, and its deviation from the primary direction is also large. As can be seen there is such an area near the top of the vessel precisely at the location where the aneurysm formed.

The initial goal of making a 3D print of this simulation result was to have the base vessel geometry as one part and the aneurysm as a second, detachable, part. Furthermore, for each of the parameter cross-sections a disc would be printed that could be slotted in the correct location in the vessel. In this way different parameters could be inspected in-situ. By using multiple cross-sections along the vessel the behaviour of the flow can be followed and inspected.

However, designing this “slot-and-disc” setup turned out to be substantially more work than anticipated. As the discs need to have color these would have to be printed in a full-color (sandstone) material,⁹ while the vessel itself needs to be strong due to its relative size and therefore is best printed in a plastic material. Furthermore, some arrangement would be needed to keep an inserted disc immovably in place while inspecting the full model, while still allowing to easily interchange a disc for another. A simpler to realize alternative was to combine one set of discs with an outline of the vessel geometry in a single 3D model and print it in full-color. See Figure 21.

Analogous to regular color printing on paper the color space available on-screen versus 3d-print isn't equal. The 3D printed colors are more washed out and slightly different in tint than on screen, see Figure 22. Here, the use of the well-known “rainbow” color scale helps in the interpretation of the data, as it provides regions with highly perceptible contrast. The specific parameter value is not absolute, but relative, so exact values are not really important, only differences between regions.

Note that the right-most discs show some stripe-like artifacts in the coloring. This is

⁹Or print the discs in plastic and glue (paper) printed visualizations on them



Figure 21: Aneurysm print




Figure 22: Color difference in the discs between on-screen visualization (background) and 3D print

due to the underlying discretization of the disc 3D model and the use of vertex colors, not due to the 3D printing.

5.4.1 Virtual stenting

A related use case, only partially concerning visualization, is the virtual placement of a flow diverter (stent) in the blood vessel: a metal cylindrical mesh. Such a stent can be surgically placed to enforce the area of the aneurysm and positively influence the local flow. The stent is placed using a catheter inserted in the blood vessel.

At the UvA software has been created that simulates placing and unfolding a stent inside vessel geometry. Different stent geometries and materials can be virtually tested this way on actual vessel geometry acquired from medical scans. However, verification of the virtual stenting with actual physical placement of a stent is needed. Here, the 3D model of the vessel can be 3D printed, followed by performing an actual physical stent placement. Both the virtual and physically placed stents can then be compared. This has added benefits compared to validating stent positioning in real, living patients. The 3D printed model can receive much larger doses of x-ray radiation which leads to higher resolution and higher contrast scans.

Challenges for 3D printing here are the small size of the blood vessels (diameter around 3-5 mm) and the needed smoothness of the 3D printed surface. The latter will influence the physical placement of the stent. Several test prints were made in two different materials and with optional surface smoothing and optional clipping. See Figure 23. Here, the vessel shape was subtracted from a cubical region.

The yellowish transparent cube is made of “transparent acrylic”. This material is described as becoming yellowish transparent after polishing, which wasn’t applied yet to the cube in the picture. The walls of the vessel are somewhat visible through the transparent cube, but doing a negative print (i.e. print the vessel surface itself) might be more instructive for showing the vessel geometry. Visually observing the physically placed stent might be difficult with the current quality of the transparent print. However, a CT-scan is planned for actual inspection of the stent placement. One noticeable print artifact is the amount of “rubble” inside the blood vessel, which is left-over print residue that should not be there. When the print was cleaned the actual vessel turned out to not be open all the way through. The print is therefore not of a very high quality.




Figure 23: Cubes

The two right-most cubes differ only in the amount of surface smoothing applied on the 3D mesh. But the resulting prints are visually identical. There are also small grooves visible due to the printing process, which should be below the manufacturer specified

threshold, but might be too large for actual use in verifying stent placement. An interesting extension could be to print the results of the 3D virtual stent placement, i.e. vessel wall with the placed stent.

The 3D printed cubes shown here were prototypes and, separately from this project, more work has been done on verifying physical stent placement using further developed 3D prints.

6 Conclusions

From realizing the different use cases discussed a number of clear conclusions and lessons learned can be described, most notably in the area of possibilities, limitations and effort involved. However, in some other areas clear and objective conclusions are harder to attain.

First, the 3D prints produced for the different use cases were deemed a success. In general it was possible to produce 3D prints with more detail than was expected, even when taking feature size limitations into account. The 3D prints provided inspiration, as well as reactions of joy - for lack of a better word - when confronting people with 3D printed versions of data they were familiar with (even with researchers that worked on the described use cases). The physical form of the data, or the specific way in which it was turned into a 3D print, was apparently something of an eye-opener. An unforeseen effect that was observed was that a high amount of detail in a 3D print can look impressive, more so than seeing the same data on a screen. Again, the physical manifestation of complex data normally only seen and operated on “virtually” on a 2D screen provides a different sensation. Perhaps this is due to being able to observe a 3D print using stereoscopic vision, versus monoscopic vision for a 3D rendering on a 2D display. Apart from being able to physically touch and feel a 3D print, interactions of the 3D print with the physical environment, for example due to lighting and shadowing, might also be a factor.

Focusing more on the use of 3D printing for visualization the (qualitative) benefits found of the 3D prints over on-screen 2D/3D visualizations were that they allowed more intuitive and natural communication and discussion of the underlying (scientific) material, both with peers, but also for lay audiences. Being able to simply physically point at a location on a 3D print is much more intuitive than having to orient the corresponding 3D model in visualization software followed by indicating that location on a screen that does not display a 3D structure very well. For cases involving complex 3D surfaces the lack of depth perception and stereoscopic view on a screen hinder spatial understanding. Complex 3D models, especially those resulting from scientific simulations, can also stretch the limits of graphical hardware used for their visualization. In this respect a 3D print is a low-end visualization, in that once the 3D print has been created no hardware is needed to view it and it provides a naturally stereoscopic and intuitive view that is easy to manipulate. Being able to hand over a 3D print related to a complex research topic might make that topic more accessible and easier to grasp¹⁰

¹⁰Pun intended

for a lay-person than by looking at an on-screen 3D visualization that stays abstract and distant.

But 3D printing also has quite a number of downsides. First and foremost are the limitations in amount of detail and use of color. The 1 million polygon limit per printable part enforced by the printing provider used (and similar limits at other providers) seems fairly arbitrary and in practice was often encountered. Raising this limit would enable printing of larger high-complexity models, as it is clear that 3D models with many more polygons would still fit into the maximum printable size. So physically there should not be too much of a problem to increase the limit. However, unknown factors related to processing power and/or 3D printer hardware limitations might be involved and difficult to overcome. And as mentioned earlier printing *precision*, i.e. layer thickness, is already quite sufficient, but the minimal size of small (protruding) features is much higher and even then results in quite fragile physical details for certain materials. On the opposite side of the spectrum was the need to scale down 3D models to overcome limitations in maximum printable size, but this could introduce features that became too small to print. Balancing these two aspects was definitely something that was a recurring challenge in the project.

The limitations on the use of color are another serious drawback in current 3D printing options. The most widely available full-color (sandstone) material is fragile, relatively imprecise and heavy, which leads to a very limited area of application. For some use cases it might be an alternative to physically combine multiple prints of different single-color materials with higher strength and detail, but it adds to the manual effort involved. For producing transparent parts there currently are not very many options that are of high quality it seems.

The sweet spot of a lightweight and strong full-color material combined with a printing technique that can produce prints with small sturdy details (while being cheap to produce) is not yet filled. But several developments in this area might come to market soon [25].

We only briefly experimented with different printed sizes of the same 3D model and cannot really draw firm conclusions from this. One limitation that we noted in these tests was that the used printing provider did not offer the option to directly print a differently sized version of a 3D model. The uploaded model needs to be rescaled first (overwriting the 3D model already present), followed by performing the automatic checks. So testing with different sizes of a 3D model involves uploading it multiple times and applying different scales by hand. The provider *did* offer easier variation of printable sizes, but only for 3D prints that were offered for sale (selling 3D printed models is an option for Shapeways customers). Such a commercial approach is probably not of interest in most scientific use cases.

Not related to 3D printing per se, is the need to turn a 3D model into some appropriate physical form. Here, several design challenges need to be solved, which most of the time involve manual (and creative) input, making automation harder. The iterative process of designing, uploading and finally producing 3D prints can take quite some effort. Results were not always predictable and more than once a 3D print arrived that had some defect that was missed during design. The use of supporting software helps, but

can't always indicate or even predict the issues. Here, experience with all facets of the process is very valuable. Working with highly detailed 3D models can be a challenge in general, especially when needing to edit small surface features or perform compute-intensive operations like a boolean intersection of two 3D models. Experience with 3D geometry processing and optimization is usually needed for turning complex 3D models into feasible 3D prints, but can be supported with appropriate software that automates (part) of this process.

The price for producing (large) prints can be prohibitive, but was not a major factor in the use cases here. However, material and machine-space costs increase quickly when scaling up 3D prints. Also, the iterative design process involved some use cases can be costly and time-consuming.

All in all the added value of 3D printing for visualization is not easily summarized, even though certain benefits are listed above. Regardless of all the limitations mentioned above the project results definitely motivate us to apply 3D printing more often. Not only to produce actual 3D prints that can be used as part of research, but also because it provides inspiration and helps in communication. For example, the option to make available 3D-printable models of datasets, as part of making research data freely available, is an interesting addition to the concept of open science.

But 3D printing for visualization will most likely not be applied as commonly as on-screen visualization and will stay a fairly unique way of visualizing. Obviously, on-screen interactive visualizations provide infinite flexibility and changes can be made instantly, something that 3D printing will never match. It should therefore not be viewed as an alternative to on-screen visualization, but as an addition to it. One that triggers different senses, speaks more to intuition and so provides different opportunities. In this respect there appears to be room for more regular application of 3D printing for visualization.

The oncological surgery support case was special in that it made a two-way connection between regular on-screen visualization (of a CT-scan) and the corresponding object in the physical world. Whereas in other use cases information flowed from 3D visualization/model to 3D print the surgery case reversed this direction: the 3D print was turned back into a 3D (volumetric) model through the CT-scan. This principle could be applied more generally, for example, by allowing users to add physical markers (or even scribbles with a pen) onto a 3D print that are then recognized and tracked using computer vision techniques in order to augment the original 3D model that was printed. In this respect future developments in Augmented and Virtual Reality are of interest, as these techniques deal with complementary goals and challenges. For example, the surface of a 3D print cannot easily be altered once printed, but visually overlaying a differently colored surface onto a white print is relatively easy with AR techniques. This could be used to show different parameters sets on a 3D print of, say, a simulated surface. Or a combination of 3D printed scene elements and AR could be used to let a user interactively construct an AR-based visualization that provides all the intuitive benefits of manipulating physical objects.

Time will tell what great applications of 3D printing to visualization are going to be created.

References

1. Clements, D. L., Sato, S., & Fonseca, A. P. (2017). Cosmic Sculpture: A new way to visualise the Cosmic Microwave Background. *Eur. J. Phys.*, *38*, 015601. doi:10.1088/0143-0807/38/1/015601
2. 3D print your own baby universe. (2016, October). Imperial College London; <https://www.youtube.com/watch?v=3GSwl21-Zwl>.
3. Diemer, B., & Facio, I. (2017). The fabric of the universe: Exploring the cosmic web in 3D prints and woven textiles. *Publications of the Astronomical Society of the Pacific*, *129*(975), 058013. Retrieved from <http://stacks.iop.org/1538-3873/129/i=975/a=058013>
4. Madura, T. I., Clementel, N., Gull, T. R., Kruip, C. J. H., & Paardekooper, J.-P. (2015). 3D printing meets computational astrophysics: deciphering the structure of η Carinae's inner colliding winds, *449*, 3780–3794. doi:10.1093/mnras/stv422
5. Knill, O., & Slavkovsky, E. (2013). Illustrating mathematics using 3D printers.
6. Nelson, R., & Segerman, H. (2015). Visualizing Hyperbolic Honeycombs. *ArXiv e-prints*.
7. Brilliant geometry: A physical 3D zoetrope of a 4D cube. (2017). <https://www.youtube.com/watch?v=LOVzytir7bM>.
8. Mathematical art by henry segerman. (n.d.-). <https://www.shapeways.com/shops/henryseg>.
9. Segerman, H., & Irving, G. (2014). Hinged negatively curved surfaces. <https://www.youtube.com/watch?v=RnKulbKauXk>.
10. Weiss, T., Zieselman, A., Hill, D., Diamond, S., Shen, L., Saykin, A., & Moore, J. (2015). The role of visualization and 3-d printing in biological data mining. *BioData Mining*, *8*. doi:10.1186/s13040-015-0056-2
11. Weber, Z., & Gadepally, V. (2014). Using 3D printing to visualize social media big data. *CoRR*.
12. Masserdotti, A. (2016). Touch that data - actuated prism map of italy. Retrieved from <http://dataphys.org/list/actuated-prism-map-of-italy/>
13. List of physical visualizations and related artifacts (3D printed). (n.d.-). <http://dataphys.org/list/tag/3d-printing>.
14. Ghawana, T., & Zlatanova, S. (2013, May). 3D printing for urban planning a physical enhancement of spatial perspective.
15. Magnússon, R., Boer, W. de, & Seijmonsbergen, A. (2016). Landschappen in 3D. - 3D-printen van hoge resolutie terreinmodellen met overlays van geologie en orthofoto's met arcgis en blender. *ESRI Magazine*. <http://dare.uva.nl/document/2/175447>.
16. Taira, K., Sun, Y., & Canuto, D. (2017). 3D printing of fluid flow structures.

17. Jachnik, J., Goldman, D. B., Luo, L., & Davison, A. J. (2015). Interactive 3D face stylization using sculptural abstraction. *CoRR*, *abs/1502.01954*. Retrieved from <http://arxiv.org/abs/1502.01954>
18. Acevedo, D., Zhang, S., Laidlaw, D. H., & Bull, C. (2004). Color rapid prototyping for diffusion tensor MRI visualization. In *Proceedings of miccai 2004 short papers*.
19. Komai, Y., Sugimoto, M., Gotohda, N., Matsubara, N., Kobayashi, T., Sakai, Y., ... Saito, N. (2016). Patient-specific 3-dimensional printed kidney designed for "4D" surgical navigation: A novel aid to facilitate minimally invasive off-clamp partial nephrectomy in complex tumor cases. *Urology*. doi:10.1016/j.urology.2015.11.060
20. Create and print out your social network in 3D. (2013, October). <https://www.3ders.org/articles/20131011-create-and-print-out-your-social-network-in-3d.html>.
21. (2013). <https://www.youtube.com/watch?v=qecZEwyN3ak>.
22. Zavodszky, G., Rooij, B. van, Azizi, V., & Hoekstra, A. (2017). Cellular level in-silico modeling of blood rheology with an improved material model for red blood cells. *Frontiers in Physiology*, 8. doi:10.3389/fphys.2017.00563
23. Shapeways terms and conditions. (2017). Shapeways; https://www.shapeways.com/terms_and_conditions.
24. Terms of use, 3D hubs - terms and conditions. (2014). 3D Hubs; <https://www.3dhubs.com/terms-of-use>.
25. Peels, J. (2017, September). How close are we to full color 3D printing? 3Dprint.com; <https://3dprint.com/187404/full-color-3d-printing/>.

Appendices

A Upper bound on the amount of printable polygons

What follows is a back-of-the-envelope calculation of a (theoretical) upper bound on the number of polygons that can be printed. In this case for a specific material and maximum print size: Shapeway's "Strong & flexible plastic" material, with which models can be printed up to 650x350x550 mm.¹¹ Imagine a very simple 3D model consisting of a thin rectangular part of 2x2mm, containing a cube of 1x1x1 mm at one of its corners, as shown in the left part of figure 24. This model consist of 6 polygons (quads), marked with black dots. If we repeat this simple model side-by-side in the plane we get the larger model shown on the right.

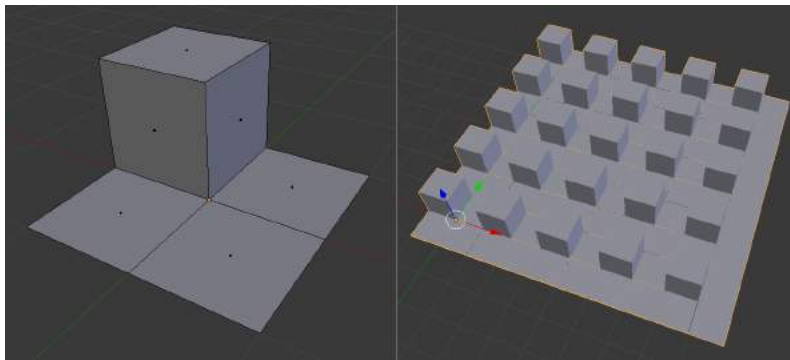


Figure 24: Imaginary model with "cubes" and "ridges"

At a maximum print size (in one plane) of 650x550m we could fit our simple 3D model $325 \times 275 = 89,375$ times in the available space, using 536,250 polygons. If we repeat the larger model in the 3rd dimension (of 350mm) as well, say with 1 mm spacing in between, we could fit it 175 times. This increases the total amount of polygons used to 93,843,750.

We are ignoring the fact that our models are not watertight and would need more polygons to make them so. Also, the feature size of 1 mm for the cube and ridges used here could be further decreased for this material.¹² So almost two orders of magnitude more detail, in terms of number of polygons, can be placed within the available space than the actual printing limit of 1,000,000 polygons.

A second example: 3D print platform 3D Hubs does not explicitly enforce a polygon limit. The only explicit limit is a maximum file size of 128 MB. For an STL file using 50 bytes per triangle this would give an upper limit of around 2.65 million polygons.¹³ If the STL file is compressed with a (conservative) factor of 2 this would allow at least 5 million polygons. This is already better than the 1M limit above.¹⁴

¹¹<https://www.shapeways.com/materials/strong-and-flexible-plastic>

¹²The material documentation lists 0.2 mm as the "minimum embossed detail" size

¹³Twelve 32-bit floating-point values (3 vertices + normal vector) plus a 16-bit attribute value

¹⁴However, uploading a model of 4M polygons (size 20x12x14 cm) resulted in *printability* and *mesh integrity* feedback being unavailable, leading to the impression that large models are problematic.

B Links to online background material

- General 3D printing overview
 - <https://www.3dhubs.com/knowledge-base/additive-manufacturing-technologies-overview>
 - <https://www.3dhubs.com/what-is-3d-printing>
 - <https://www.shapeways.com/tutorials/how-your-3d-printed-product-is-made>
- Designing 3D printable models
 - <https://www.3dhubs.com/3d-printing-handbook>
 - <https://www.shapeways.com/tutorials>
- Materials
 - <https://www.shapeways.com/materials>
 - <https://www.3dhubs.com/material-index?action#?page=1&sortBy=popularity&order=desc>
- Commercial 3D printing service providers (selection)
 - <https://www.shapeways.com/>
 - <https://www.3dhubs.com/>
 - * An online platform where owners of 3D printers can provide printing services.

C Tools

- Meshlab
 - <http://www.meshlab.net/>
 - Open-source general geometry processing and inspection tool
 - Beside being able to export to STL (one of the most common formats for 3D printing), MeshLab can be used to prepare 3D models for printing by creating inner shells, resampling/remeshing the 3D model to make slicing easier, closing small holes to obtain watertight meshes, and flattening the bottom area to have a better platform adherence.
- Blender
 - <https://www.blender.org>
 - Open-source general purpose 3D modeling, animation and rendering package
 - Includes “3D Print Toolbox”
 - Model volume and area, checking manifoldness, wall thickness, intersections, ...
 - Scripting geometry creation and operations
- Autodesk Netfabb
 - <https://www.autodesk.eu/products/netfabb/overview>
 - Commercial additive manufacturing and design software (Microsoft Windows only)
 - Analyze and repair models, support creation, topology optimization, ...
- Ultimaker Cura

- <https://ultimaker.com/en/products/ultimaker-cura-software>
- Open-source software for preparing 3D models for printing on a wide range of 3D printers
- Provided by Ultimaker, a 3D printer manufacturing company
- Targeted at owners of 3D printers
- Provides lots of control over the actual input sent to the 3D printer, in terms of slicing, infill, printing quality, etc.