

Scientific data visualisation using techniques normally reserved for more frivolous activities

Paul Bourke
University of New South Wales
Sydney, Australia
paul.bourke@unsw.edu.au

Abstract—Here we present technologies one normally encounter being used for frivolous merchandising, to the presentation and visualisation of scientific data and mathematics. Three particular print technologies will be presented, they are glasses free lenticular prints, 3D printing (also known as rapid prototyping) and crystal engraving. The relative merits of each technology will be presented, what forms of data they are most suited to, the challenges and limitations, and some details of how required data formats are created. The discussion will be illustrated by examples from the authors application of the technologies in the context of creating engaging data presentations and visualisation for research, public outreach, education, museum and art gallery exhibition.

Keywords—component; Visualisation, data representation lenticular, crystal engraving, 3D printing, rapid prototype, hologram, glasses free.

I. INTRODUCTION

The visualisation of scientific data and mathematics often employs high end graphics hardware and advanced human computer interfaces. For example stereoscopic displays are used to present geometrically complicated structures, the additional sense of depth provided by these displays can greatly enhance both the researchers understanding of the structure but also convey that same insight to their peers and members of the public. High resolution displays can also be used to assist with the resolving of data structures resulting in a reduced need to continually zoom in and out, the process of which is a tradeoff between seeing the detail and understanding the global context. Surround displays are sometimes used to immerse the viewer in the data, to give a sense of "being there" with the resulting heightened engagement. Finally, displays designed for visualisation in the sciences and mathematics can employ novel human computer interface devices to provide richer multidimensional control or, in the case of haptic devices, an interaction that leverages the additional sense of touch through our hands and fingers.

There are a number of characteristics the above laboratory based and bespoke visualisation displays and devices possess that can limit their use outside the research organization where they are hosted. Some of these are:

- **Cost.** Expensive equipment, while safe in the research laboratory and in the hands of the researchers, may not suit a more public and unsupervised environment. Even the cost of insurance to move the hardware can in some cases be prohibitive.

- **Delicacy.** Some devices may not be suitably robust for use by the public. Care for devices can be taken by researchers in the laboratory who more fully understand the limitations and the consequences of damage.
- **Specificity.** Some devices, particular human computer interfaces, may require some training in order to gain optimal benefit. They may additionally need to be fitted or adapted to the individual, again acceptable for a small groups of researchers but not the wider public.
- **Portability.** Specialized displays are often fixed in their location or may be heavy and thus not readily relocated.
- **Health and safety.** Devices such as head mounted displays (HMDs) and even glasses for stereoscopic displays have a means of transmitting medical conditions due to skin-to-skin contact. Immersive displays can result in balance issues in some people, stereoscopy can induce eye strain and headaches in approximately the 10% of the population who are susceptible.

These factors can make it difficult to present data in rich and engaging ways outside the research laboratory. This might include for example at conferences, in museums, as part of art gallery exhibitions, to the public and in the classroom environment. On the other hand one only needs to glance at touristic shops in cities or airports around the world to see a number of novel technologies that are being used to present and sell various forms of merchandise. Can these technologies be used to present data in science and mathematics in equally engaging ways and at more appropriate price points. In the remainder of this paper three possibilities that have been explored by the author will be discussed. The choice has been limited to the creation of physical objects rather than the many purely digital technologies that may be employed such as through novel browser based interfaces, on smart phones, through social media, augmented reality, gaming engines and so on.

The three technologies to be discussed are lenticular prints, rapid prototyping, and crystal engraving. The technology for each will be introduced and the opportunities it offers discussed along with the limitations. This will be in the context of actual practice by the author and in each case one or two examples will be presented along with technical details on how the data needs to be prepared and pre-processed in order to make it realizable with the technology.

II. LENTICULAR PRINTS

Providing depth perception by exploiting stereopsis (consequence of human binocular vision) has been a standard visualisation tool for over 30 years, becoming mainstream with the invention of the CRT display. The benefits for the understanding of geometrically complicated datasets by supporting depth perception is intuitively obvious and this was born out in practice. The change in the last 5 years has been the increased commoditization of the technology at least for 3D television size displays, once only the domain of well funded laboratories. There is however a gap between the digital stereo3D capable display and the more traditional printed means of presenting data that is still relevant even today. Researchers use stereoscopic 3D displays in the laboratory but are reduced to flat 2 dimensional representations in print, posters and at conferences.

There are a few options for representing depth in printed form. As with any 3D presentation there fundamentally needs to be a left and right eye channel which are presented independently to each eye. At one end of the spectrum are anaglyphic prints, these separate the channels by using two colours and the viewer to required to wear matching glasses. While anaglyph prints have many desirable properties, for example they can included in standard print material and on unmodified digital displays, they generally have poor colour fidelity. Another option are full synthetic holograms [1], while these are glasses free they currently have a high associated cost and have an uncertain future since they rely on film as the high resolution medium for capturing the light field. The method presented here uses the familiar lenticular [2] sheets that are traditionally used for gimmicky cards of sports people distributed in cereal packets, cheap touristic postcards, and 3D views of spiritual icons. These can be employed to show multiple objects and even simple animations but the interest here will be limited to multiple views of a single object resulting in the sensation of depth.

Lenticular technology is a convenient way to induce depth perception, it is a print from a normal high resolution printer subsequently mounted onto the lens layer. It exists as a physical object that can be hung on a wall or passed around. It does not require any form of eyewear to separate the channels, as such it is referred to as autostereoscopic and incurs minimal eye strain. Finally, they can be produced for a modest cost, and replicated for even less.

The detailed theory behind how lenticular or the simpler barrier strip displays or prints work will be left to the reader [2]. In the case of the barrier strip method a left and right pair of images are multiplexed vertically, the barrier strip ensures the left eye sees only the vertical strips from the original left eye image, and the right eye only sees the strips from the original right eye image. In the case of lenticular prints the barriers are replaced by small lenses and instead of just two, a number of images are multiplexed vertically. The lenticular lenses present these vertical strips each in their own viewing zones, if the images are created correctly each pair of viewing zones presents a pair of images to the viewers eyes and every pair is a valid stereo pair. The parallax effect of looking around the object is supported when moving ones head horizontally if each

image is computed along a path in front of the object. While the source images for lenticular prints can be created photographically, we are only interested here in synthetic data representations.

The image creation then involves creating multiple renditions of the 3D model, each pair of rendered images being a valid stereo pair, see figure 1. The range over which the camera is translated and the field of view of the camera are determined directly from the final viewing conditions (size of the print and the distance it will be from the intended viewing position), as is the case with all correct stereoscopic content creation. An example of images rendered from a model of Australian indigenous rock art is shown in figure 2. The number of images required depends on a few factors, the lenticular resolution here uses about 25 images. The different parallax can be seen in the final frame compared to the first few frames. Two views of the resulting lenticular print can be seen in the photographs in figure 3, while the 3D nature cannot be captured here the difference in parallax is visible.

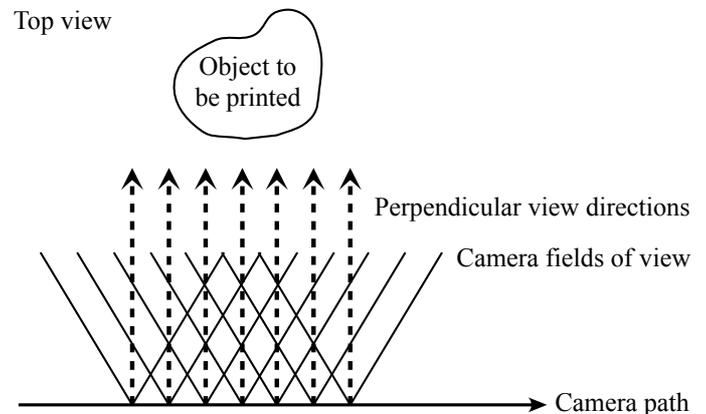


Figure 1. Straight camera path. Field of view and distance of the virtual camera from the object determined by the geometry of the lenticular sheet and planned viewing distance.

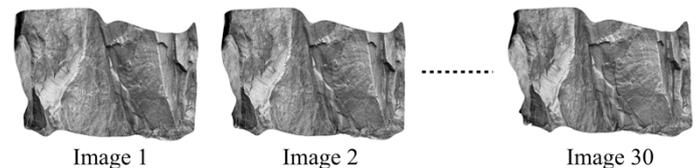


Figure 2. Sample images from a camera moving along linear path, note the parallax between the earlier and later image in the sequence.

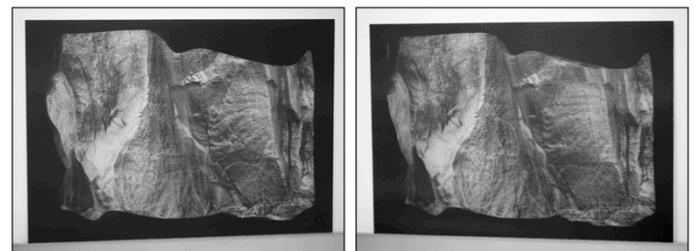


Figure 3. Two photographs of the final lenticular print from two view points of an example from Australian indigenous rock art.

III. RAPID PROTOTYPING

Visualisation, as the name implies, generally uses our sense of vision in order to convey information to the human brain. Sonification is the term given to the use of our sense of hearing in the visualisation process. While there are obvious examples of this such as the Geiger counter or the hospital "machine that goes ping", the conversion of data into audio or music is generally in support of visuals rather than instead of. Haptics is the use of our sense of touch for visualisation but this is generally mediated through the use of a mechanical and imperfect force feedback system. An alternative to haptics is to create a physical model of data, this allows the object to be explored using exactly the same natural eye-hand coordination we use to explore and understand 3D objects in our real world.

While rapid prototyping has been around for at least 20 years, it was rarely used outside the realm of industrial product design. Recent advances in rapid prototyping (RP), more commonly now known as simply 3D printing, has enabled a much wider range of objects to be realizable than was possible 10 years ago as well as printing in a wider range of materials. Today the techniques are increasingly becoming commoditized including online bureau services [3] both high end machines capable of very intricate and general designs [4] to low cost machines [5] that may not have the same finesse of design capability but are accessible at a much lower price point. In order for a particular structure to be realizable it must meet a number of conditions.

- RP machine can only create solid objects, that is, they cannot print the idealized points, lines, and planes often used within computer graphics based visualisations. Such ideal mathematically building blocks need to be given thickness to become "physical".
- The digital representation of the model surface must be watertight, that is, it needs to be a closed mesh without gaps or cracks arising from numerical imperfections.
- A particular 3D printing method and material imposes limits on the finest structures that can be built. The limit may be due to cleaning or other post processing stages, or it may simply be related to the strength of the underlying material. Current top of the range 3D printing in monochrome (any single colour) can resolve filaments down to about 0.75mm whereas full colour printers [6] are generally limited to at least 5mm.
- For practical reasons a model needs to be in one piece otherwise each piece should be considered as separate models.
- There is a maximum size that can be created given a particular printing technology.
- The cost, except for some more exotic coatings in precious metals, is a function of the volume of material used. In some cases the way the model is represented may influence the cost, for example large portions may be created hollow.

All the above generally mean that software written to convert data into 3D printable geometry needs to have a certain

rigor. Two particular commonly encountered requirements are how to represent curves (infinitely thin lines) or planar surfaces (also infinitely thin). It should be noted that all geometry will be transferred to the RP machine as planar approximations of higher level curves or surfaces, otherwise known in computer graphics as a triangulated (3 vertex bounded region of an infinitely thin plane) mesh.

The simplest method of creating lines with thickness is to replicate spheres along the curve or its straight line approximations, see figure 4. While simple, because each sphere is represented by a large number of triangular faces, it is a very inefficient method. Doubly inefficient since a large number of spheres may be required for a smooth representation of the line or curve. Slightly better is to remove all triangular faces internal to the spheres, this is a very fast test that reduces the triangle count significantly. A more efficient method is to sweep cylinders along each line segment of the curve, see figure 5. There are a number of ways the transition between two line segments can be handled, they all work tolerably well for slowly changing angles but some methods work better than other for sharp angles. It should be noted that invisible geometry due to overlapping spheres or cylinders is not an issue except that it may be less than optimal efficiency.

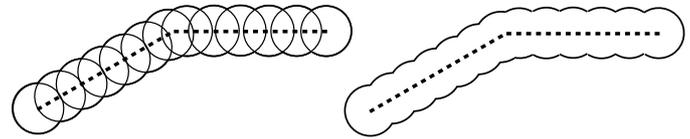


Figure 4. Distributing closed spheres along the curve (left). Optionally removing internal triangular faces making up each sphere (right).

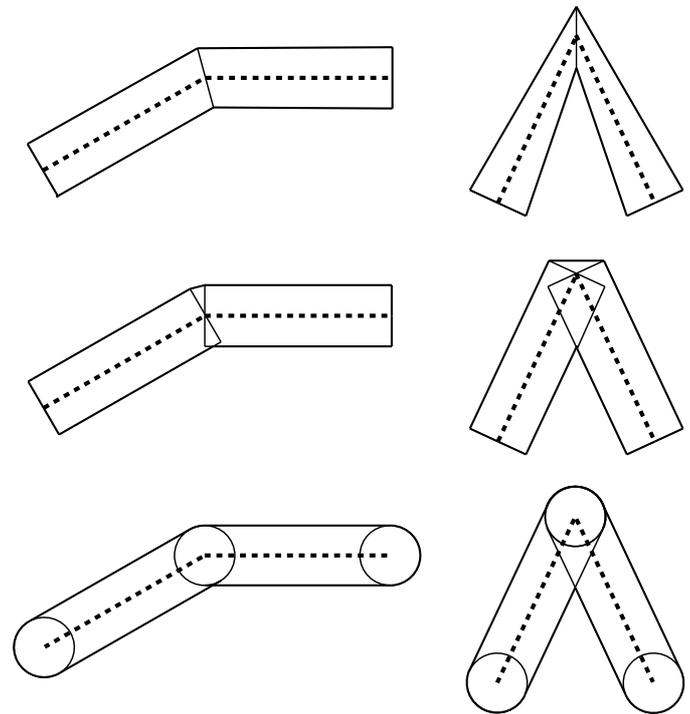


Figure 5. Each row illustrates various methods of thickening a curve using cylinders and spheres. Most work acceptably for low curvature line sequences (left column) but are differentiated by sharp angles (right column).

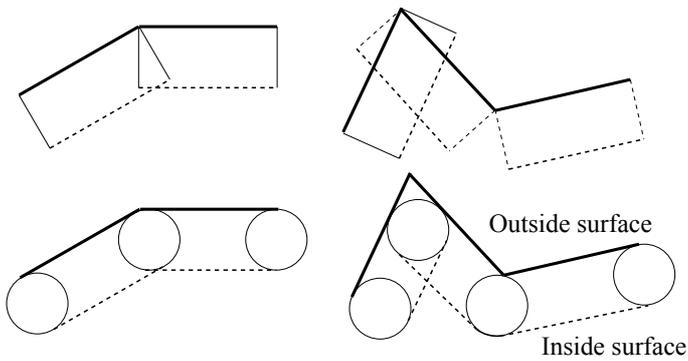


Figure 6. Problem with simple extrusion of surfaces for sharp angles (top right), rolling ball algorithm proposed by the author (bottom right).

The representation of spheres or cylinders by a triangular mesh involves a discrete approximation to the surface and as such a sphere is much more expensive than a cylinder for the same fidelity. For example if a cylinder is approximated by 10 degree steps, there are a total of 108 triangular faces (36 at each end and 2 for each extruded face). A sphere represented in the normal polar coordinates parameterisation (not the most efficient) would require 1296 triangular faces (36 lines of longitude, 18 lines of latitude, and 2 triangles per face). As such it is common to use more efficient representations of spheres such as recursive tessellations of a platonic solid followed by vertex radius normalisation, but even there the triangle cost of spheres for the same degree of surface smoothness is generally significantly higher than for cylinders.

As with the thickening of lines, 2D surfaces also need to be given some thickness. Intuitively one might imagine extruding each planar section perpendicular to the normal of the plane, see top row of figure 6. Note that this introduces the notion of an "inside" and "outside" to the surface, the inside being the direction the planar section is extruded in. The main problem with this approach is the punch-through for corners with greater than a right angle turn. The solution proposed by the author is called the "rolling ball" algorithm. Imagine a ball rolling on the inside of the surface. Create the extruded surface by tracing out the position of the point on the ball most distant to the surface. The resulting thickened surface has an embedded notion of the inside and outside surface in so far as the inside surface is smoother than the outside surface since concave transitions are replaced by smoothed areas of the rolling sphere.

Models such as those discussed here are conveyed to the RP machine by means of a data file that describes the geometry in terms of triangular faces. The exact file format used may depend on the machine, the vendors software, and whether colour [6] is supported. The most basic file format for monochrome models is the STL format, an extremely simple file format to create and one that heralds from the early days of STereoLithography, hence its name. Simple colour models are most simply conveyed as WRL or X3D formats, both of which originated from the VRML standard of the 1980 intended to present 3D models in the context of the web browser. Textured models can be described in a number of formats but the simplest is as Wavefront OBJ files, these are all plain text formats and are thus relatively straightforward to create from ones own software.

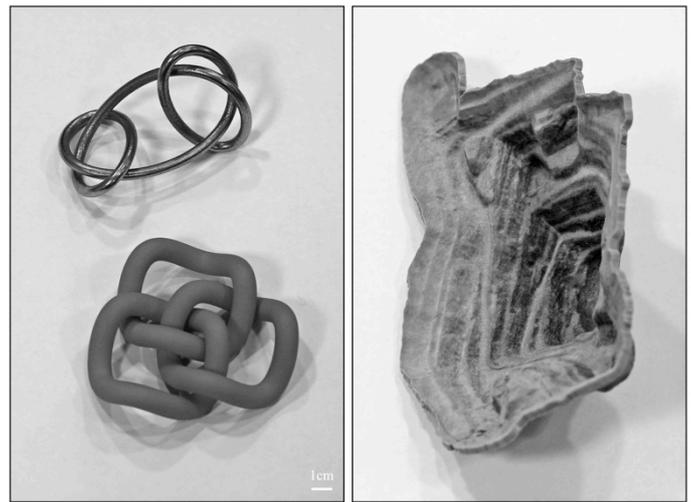


Figure 7. Mathematical knot illustrates circle extruded curves (left), extruded mine model surface (right).

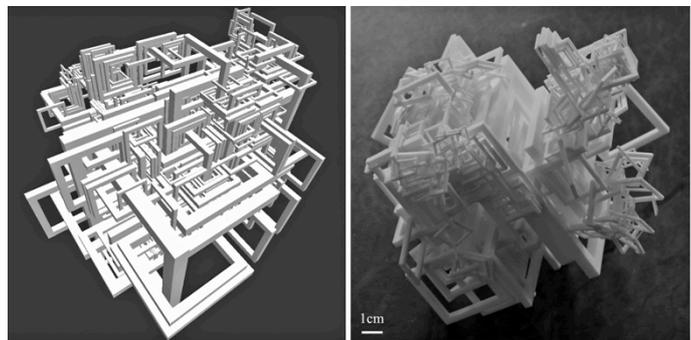


Figure 8. Computer rendering of linked toroidal space filling [7] (left), 3D printed version (right) also illustrating information carried through real world lighting of the printed models.

IV. CRYSTAL ENGRAVING

As discussed above, RP models cannot represent points and are only appropriate for single connected models. There are many imaginable examples where RP processes would be unsuitable as a means of physically representing data. One such example would be representing the points from a galaxy survey, which, while simple to draw on a computer display would be totally unsuitable for 3D printing since it would result in simply a collection of unattached points, see figure 9.

There is however a 3D printing technology ideally suited to such non-connected data, indeed ideally suited to astronomy data. It can often be found in tourist shops around the world, stores dealing in glass and crystal works, and there is even a world wide franchise where a persons bust or favorite pet can be photographed and printed. The technique is laser based printing [8] within solid glass blocks, although other shapes are possible. The basic technology is called Sub Surface Laser Engraving (SSLE) and involves focusing a laser beam at precise locations within the crystal block, at each position a small bubble forms representing one data point. The accuracy at which these bubbles can be formed is very high and they can be smaller than 1/10mm in diameter.

There are clearly a number of limitations of this technology that are fundamental in the process.

- Colour is not possible, the bubbles only appear as white dots due to scattering of incident light through the crystal block. It is however possible to illuminate the crystal block with coloured light (usually a coloured LED light base) and some practitioners have used this with some success.
- The dots are of fixed size so they cannot be used directly to represent grey scale information. However the dot density can be varied within a region to convey a linear scale and it should be noted this is not so much a limitation of the technology as much as the limitation of the current implementations. In theory variable powered lasers could produce a range of bubbles sizes.
- There is a limit on the bubble density, localized defects and even cracks will appear if the bubble density is too high, for example, if the bubbles overlap. Equally, if the point density is too low then the object appears too faint.

In the process of representing data using SSLE, the author has developed algorithms to extract suitable point clouds from various types of dataset, namely volumetric and polygonal. In both cases the key is creating an appropriate density of dots, dense enough so that the surfaces are clearly visible and not too dense that cracking occurs. Perhaps surprisingly these models can be created very quickly, within minutes, despite there being perhaps millions of points. As such the author need now overly concern themselves with the number of points but rather their optimal deployment.

Two approaches have been explored for creating suitable point clouds from volumetric data, the first is to create a point at positions in the volume if the voxel value at that position lies within some range, see figure 10 (left). Another approach is to first create an isosurface, using marching cubes say, and then polygonise that polygon model, see figure 10 (right). The direct volume sampling is the easiest and works best when the resolution of the volumetric data is a reasonable match to the final point density in the crystal.

A number of options exist for creating a point cloud from a polygonal dataset. Whichever is used one generally needs to only add a point to the final point cloud if the point is at least some minimum distance from any other point, this is to avoid large connected bubbles and in the worse case internal cracking. Candidate points can be added by considering the vertices, the center of triangles, by sampling of the edges, or even sampling the faces, see figure 11. The best method depends on the characteristics of the triangles within the model. For example, the triangles in meshes derived from marching cubes [10] tend to be approximately equal size and high density in which case sufficient point density may be achieved by simple considering the vertices themselves. For other models there may be large polygons which will need their surfaces subsampled. Another approach is to contour the model along one or two axis planes and sample points along those contour lines. As with surfaces displayed on a computer display this often leads to natural illustration of the local surface curvature.

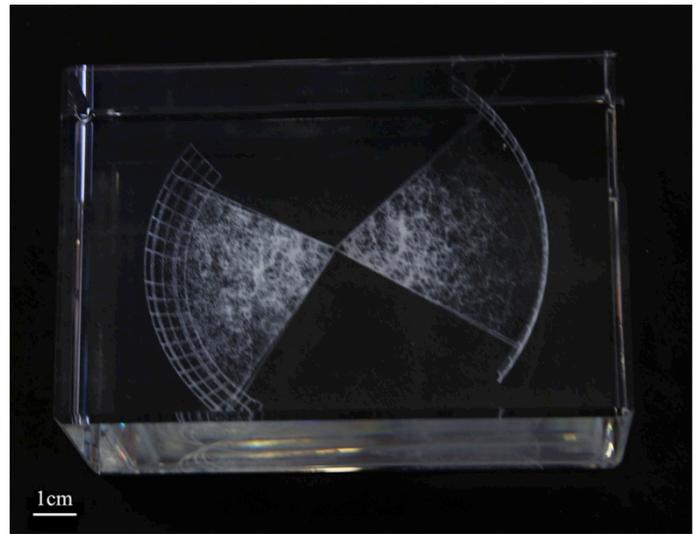


Figure 9. Crystal block of the 2dF galaxy redshift survey [9], each point is a 3D galaxy position and as such the data preparation is trivial.

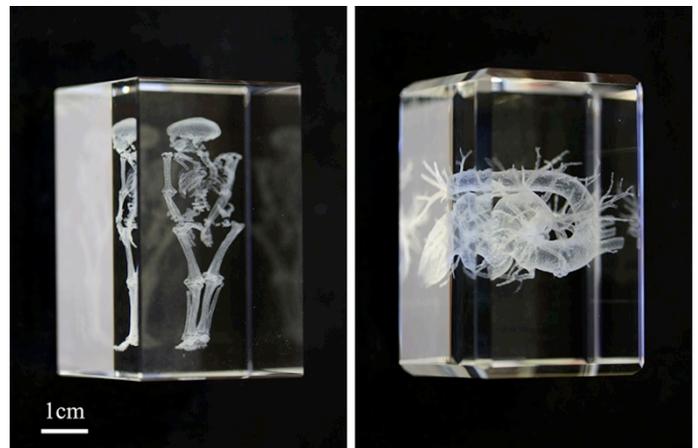


Figure 10. Direct volumetric point cloud generation of an Egyptian mummy (left), isosurface sampling of a CT scan of a human heart (right).

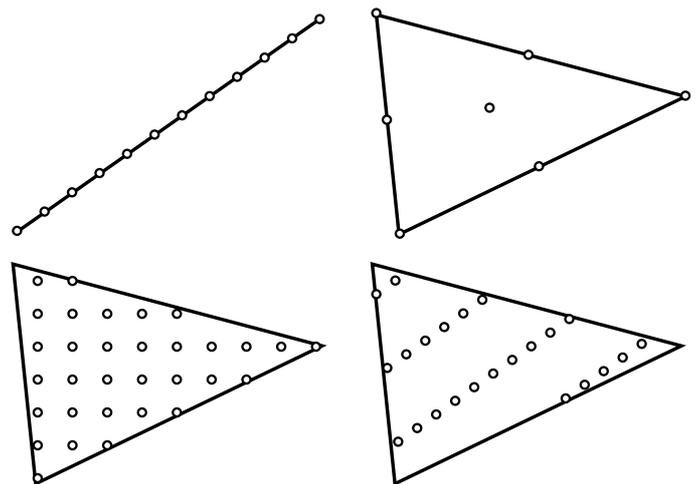


Figure 11. Sampling lines with defined dot spacing (top left), sampling planar faces at vertices and optionally midpoints (top right), sampling planar faces at predefined dot spacing (bottom left), sampling planar faces along contours (bottom right).

V. CONCLUSION

Presented here are three technologies that are more commonly encountered in areas of merchandise and marketing but which are suited to more serious applications, in particular, the representation and visualisation of data from science and mathematics. There are parallels here with the use of so called game engines that can be used to create virtual environments and deployed for equally serious applications.

All three technologies result in actual physical objects and can be produced at relatively low cost making them suited to use where more specialized digital displays or other hardware would not be possible or would be prohibitively expensive. The original application of these technologies for frivolous merchandising can now be applied to more meaningful forms of educational merchandise.

The applications, while using technologies many of the public will have encountered, have not generally been applied to the visualisation of scientific data and can therefore be viewed as novel and thus give rise to an increased engagement. This increased engagement has been shown to provide a heightened learning experience for school age education.

Visualisation is generally performed using only the sense of vision, the tactile aspect of the rapid prototypes and crystal blocks utilizes another of the human senses. The multimodel aspects of printed 3D models has already proved to have benefit in computational chemistry and microscopy where researchers can explore the molecular structures with their joint visual/tactile sense.

VI. ACKNOWLEDGEMENTS

The work was supported by iVEC through the use of advanced computing resources located at The University of Western Australia.

REFERENCES

- [1] Synthetic dynamic holograms. Web reference (Access 1 Feb 2015). <http://paulbourke.net/miscellaneous/hologram/>
- [2] Oster, Gerald: Optical Art, Applied Optics (1985) 4 (11): 1359–69. doi:10.1364/AO.4.001359
- [3] Shapeways: Online 3D printing bureau service. Web reference (Access 1 Feb 2015). <http://shapeways.com/>
- [4] ObJet 3D printer. Web Reference (Access 1 Feb 2015). <http://www.stratasys.com/>
- [5] MakerBot desktop 3D printing. Web Reference (Access 1 Feb 2015). <http://www.makerbot.com/>
- [6] Zcorp (3dsystems) colour 3D printer. Web Reference (Access 1 Feb 2015). <http://www.3dsystems.com/>
- [7] J. Shier and P.D.Bourke: An Algorithm for Random Fractal Filling of Space, Computer Graphics Forum, The Eurographics Association and John Wiley Sons Ltd. Vol. 32, Issue 8, pp. 89-97, December 2013.
- [8] Crystal engraving hardware. Web reference (Access 1 Feb 2015). <http://www.hanslaser.net/product/>
- [9] D. S. Madgwick, E. Hawkins, O. Lahav, S. Maddox, P. Norberg, J. Peacock, I. K. Baldry, C. M. Baugh, J. Bland-Hawthorn, T. Bridges, R. Cannon, S. Cole, M. Colless, C. Collins, W. Couch, G. Dalton, R. De Propris, S. P. Driver, G. Efstathiou, R. S. Ellis, C. S. Frenk, K. Glazebrook, C. Jackson, I. Lewis, S. Lumsden, B. A. Peterson, W. Sutherland, K. Taylor. The 2dF Galaxy Redshift Survey: galaxy clustering per spectral type. Mon. Not. R. Astron. Soc. 000, 1-10 (2003).
- [10] W. E. Lorensen, H. E. Cline: Marching Cubes: A high resolution 3D surface construction algorithm. Computer Graphics, Vol. 21, Nr. 4, July 1987.