

The Monumental Geometry of E Pluribus Unum

Craig S. Kaplan*
Cheriton School of Computer Science
University of Waterloo

Chris Jordan
Artist
Seattle, Washington

Abstract

We describe the conception and construction of *E Pluribus Unum*, a large-scale artwork that combines text and geometry to depict over a million names of global organizations dedicated to social and environmental causes. The names are rendered in small type along lines arranged symmetrically in a disc. We examine the aesthetic choices involved in designing the artwork, and the technical challenges we faced in laying out and rendering a monumental amount of text.

CR Categories: I.3.8 [Computer Graphics]: Applications J.5 [Computer Applications]: Arts and Humanities—Fine arts I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

Keywords: typography, geometry, mandala, art

1 Introduction

The scale and interconnectedness of contemporary global culture are too vast, too complex, and too rapidly changing for the human mind to comprehend. We are steadily bombarded by quantities measured in terabytes, tonnes, and trillions of dollars, but these abstractions rob the real world of its substance and provide only a facile summary.

Art is a medium in which to confront these questions of scale, a means of building a metaphorical lens through which an individual might glimpse humanity as a whole.

The second author is a digital artist who has been tackling this problem for 15 years. He has created a number of works that present the viewer with a concrete manifestation of an otherwise incomprehensible datum. His images typically consist of simple arrangements of small elements—plastic cups, or cigarette packages, or credit cards—by the thousands or millions. Often the small elements are arranged by colour or tone so that they collectively give rise to a single larger image, in the manner of an image mosaic [Rosin and Collomosse 2013, Chapter 10]. These works are gigapixel images intended to be presented as large-scale digital prints, ranging from two to twelve square meters in surface area. The imposing physical size of these artworks turns the blandness of a number into a visceral experience. They also afford a unique opportunity to experience complex information at multiple scales; a viewer can stand back to take in the entire composition at a glance, but then approach continuously until they can resolve each of the individual elements. When a large-scale physical artwork cannot be displayed at its native size, an interactive visualization that permits exploration via deep zooming can serve as a substitute.

*e-mail: csk@uwaterloo.ca

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CAe 2014, August 08 – 10, 2014, Vancouver, British Columbia, Canada.
2014 Copyright held by the Owner/Author. Publication rights licensed to ACM.
ACM 978-1-4503-3019-0/14/08 \$15.00

Most of the artist's works depict aspects of modern society that can be considered negative: the deterioration of the natural world, over-reliance on disposable goods, and human tendencies towards violence, vice, and vanity. The goal is not to confront the viewer with purely negative imagery, but to depict these aspects of society as they are; the viewer is invited to contemplate the consequences of humanity's collective action, and to consider change if they find the result objectionable.

In 2008, however, he conceived of a new work that would present a positive aspect of our interconnectedness. He envisioned a large-scale composition made up of a symmetric arrangement of lines criss-crossing a circle. The design would evoke the symbolism of the circle in the world's cultural traditions—most immediately the mandala in Hinduism and Buddhism, but also the Gothic rose window, the compass rose, Indra's net, and the human eye. At a distance, the lines would represent links between people and groups all over the world. Close up, it would become clear that each line is built from text, showing the names of a million worldwide organizations devoted to peace, environmental stewardship, social justice, and the preservation of diverse and indigenous culture (the true number of such organizations is estimated to be much higher, and growing). The intention is to depict the ability of these organizations to draw humanity closer together, into a consistent force for good on earth. The piece was to be called *E Pluribus Unum* (from the many, one).

The creation of such an artwork embodies several novel technical challenges, particularly concerning the layout and rendering of such an enormous amount of text in a single composition (the final piece contains nearly 33,000,000 characters, equivalent to more than 25 copies of *Moby-Dick* on a single page). The first author, a computer graphics researcher, collaborated remotely with the artist by creating a suite of interactive and offline software tools. The tools allowed us to work together on decisions related to aesthetic goals (such as the geometry of the mandala and choice of typeface) and engineering goals (such as the print size and resolution).

This paper describes the artistic and technical collaboration that led to the production of *E Pluribus Unum*. We will discuss the collection of a database of organization names (Section 2), the exploration of the space of mandala designs (Section 3), and the layout (Section 4) and rendering (Section 5) of the text.

2 Taking names

In 2007, author and entrepreneur Paul Hawken published *Blessed Unrest* [Hawken 2008]. This book explores the global rise and impact of environmental and social justice organizations, viewing them collectively as a single movement. The book's development led to the creation of the WiserEarth project (www.wiser.org), a virtual community in which individuals and organizations can exchange information. The WiserEarth website includes a directory of organizations from around the world, whose names are an ideal list for *E Pluribus Unum*. From WiserEarth we obtained a list of 113,037 names, capturing a snapshot of their directory as of early 2008 (their directory has since grown to 114,985). The list is perforce incomplete, and biased towards English- and Spanish-speaking parts of the world.

In such a list, errors must be accepted as inevitable. The list would consume more than 900 printed pages, making it impossible to proofread every entry for spelling mistakes. A subtler technical complication is that while most names were provided in a standard UTF-8 encoding, a few used alternate encodings within the same file. Because of the frequency of accented characters, this inconsistency cannot be ignored. We developed a Python script that sanitized the input. It would naively attempt to convert the string from UTF-8 to ASCII, replacing all accented characters with equivalent XML character entity references (for example, “ë” would be replaced with “é”). If the conversion failed, it indicated the use of an alternate non-ASCII character encoding. We repaired these failure cases by manually constructing an auxiliary table mapping byte sequences to XML entity references; fortunately, only twenty or so of these fixups were needed to process the entire file. The process leaves us hoping that digital text will continue to converge on a single global standard.

3 Mandala design

The artistic goal for this project was to depict global connections between people symbolically via lines crossing a disc. It is natural to gravitate towards a symmetric layout, with people represented by N evenly distributed points (v_1, \dots, v_N) on the disc’s boundary. We will connect these points with line segments in such a way that the resulting design has d_N symmetry (i.e., the symmetries of a regular N -sided polygon). The lines we choose to draw will then be a subset of the edges of the complete graph on these N vertices. We can create different designs by assigning different tones or colours to sets of edges, including the possibility of omitting them from the drawing altogether. We refer to any design created this way as a “mandala”. A mandala will form the geometric scaffolding upon which we will eventually render text.

In order to proceed, we needed an effective and collaborative way to explore the space of mandala designs. The first author created an interactive Java application to facilitate this exploration. After a value of N is chosen on the command line, the interface displays three windows, as shown in Figure 1. The first window is a colour palette consisting of eight shades of grey and a “disabled” colour. The second window is a single vertex view showing v_1 at the “north pole” of the mandala, together with its $N - 1$ connections to all other vertices. The third window is a zoomable preview of the complete design. Each edge in the single vertex view terminates in a small disc; clicking on that disc applies the currently selected colour to the corresponding edge. Clicking on v_k applies the same change to v_{N+2-k} , the edge’s reflection across the diameter containing v_1 , thus preserving global d_N symmetry.

The application renders the mandala by drawing the packet of coloured edges, as shown in the single vertex view, at every vertex in the mandala. We draw the edge from v_i to v_j only if $i < j$, which avoids drawing every edge twice. In our greyscale model, edges are drawn strictly from lightest to darkest (and disabled edges are skipped). This drawing order can lead to a feeling of depth, with fainter lines receding into the background.

The mandala can be saved as an Encapsulated Postscript file for easy viewing. The file contains an embedded comment with high-level state information, allowing mandalas to be read back into the interface and used by subsequent algorithms in the construction of *E Pluribus Unum*. This simple expedient avoids the need for a separate save file, and guarantees that the symbolic and graphical views of the same information remain synchronized.

This application allowed us to explore possible mandalas quickly and productively. Figure 2 gives several examples. We concentrated on mandalas with 108 and 144 points, with a strong bias

towards 108 because of its spiritual significance in Hinduism and Buddhism.¹ We find these mandalas to be simple, compelling works of abstract line art. For large values of N , a variety of interesting textures emerge as one travels from the centre of the design out to the boundary. Note that when N is even, the diameters of the disc are part of the complete graph. We consistently disabled these to keep the centre of the disc open. We favour the aesthetic qualities of an open centre: it creates a space of relative calm in the centre of a chaotic design.

Once we began generating complete prototypes of *E Pluribus Unum*, we further iterated the mandala design, driven by the goal of rendering at least a million names at the chosen mandala size and text size. The threshold of one million was chosen as a means of estimating the true number of service organizations, of which our list of names represented only a small fraction. In the end, we were able to achieve our goal with a relatively simple design with maximal impact (Figure 2d): a 108-pointed mandala with diameters disabled and all other edges coloured black.

4 Layout

The previous section yields a set of lines in the plane. Each line must now be packed with text, specifically a sequence of names taken from the list collected in Section 2. Given our intention to place at least a million names from a list of around 100,000, we must expect names to repeat, and take measures to ensure that the repetitions do not produce any discernible patterns in the layout. Ideally, it would be impossible to find any repetitions upon a cursory inspection of the finished piece.

We proceed under the assumption that the name sequence for each line can be computed in isolation (i.e., that sufficient per-line randomness will lead to sufficient randomness across the complete design). Let the complete list of K names be denoted by (s_1, \dots, s_K) , where each s_i is a unicode string. Each s_i has some length w_i when rendered in a given typeface at a given size. We also define a separator string s_{sep} that will be placed between adjacent names, with its associated length w_{sep} . For a line segment L of length d , we must then find a sequence of M indices (i_1, \dots, i_M) such that $d - \sum_{j=1}^M w_{i_j} - (M - 1)w_{\text{sep}}$ is small but positive. That is, the total length of all text in the line, including separators, should come close to d without exceeding it.

We can immediately recognize this question as a variant of the SUBSET SUM problem [Cormen et al. 2009, Section 35.5], for which an optimal solution is intractable. Fortunately, there are two important reasons why optimality is not crucial. First, at our planned mandala and text sizes, lines will be much longer than names. Because we have a large set of names evenly distributed over a range of lengths, we can expect to pack lines greedily, never leaving behind space longer than the shortest name. Second, each vertex will be the meeting place of a large number of overlapping lines (107 in our final design). The termination point of incomplete lines will still be well within the region of heavy overlap. In any case, true optimality would be deterministic, undermining our quest for randomness.

In practice, we use a simple greedy algorithm. Assume the names are sorted by length, so that $w_1 < w_2 < \dots < w_K$. Let d represent the amount of empty space remaining on a line. Using binary search, we can find i_{min} , the smallest index for which $w_{i_{\text{min}}} > d$. If $i_{\text{min}} = 0$ then no names fit in the remaining space and the algorithm terminates. Otherwise, we add $s_{\lfloor \sqrt{u} i_{\text{min}} \rfloor}$ to the line, where u is a uniformly distributed random number in $[0, 1]$. The square

¹See, for example, the Wikipedia page for the number 108.

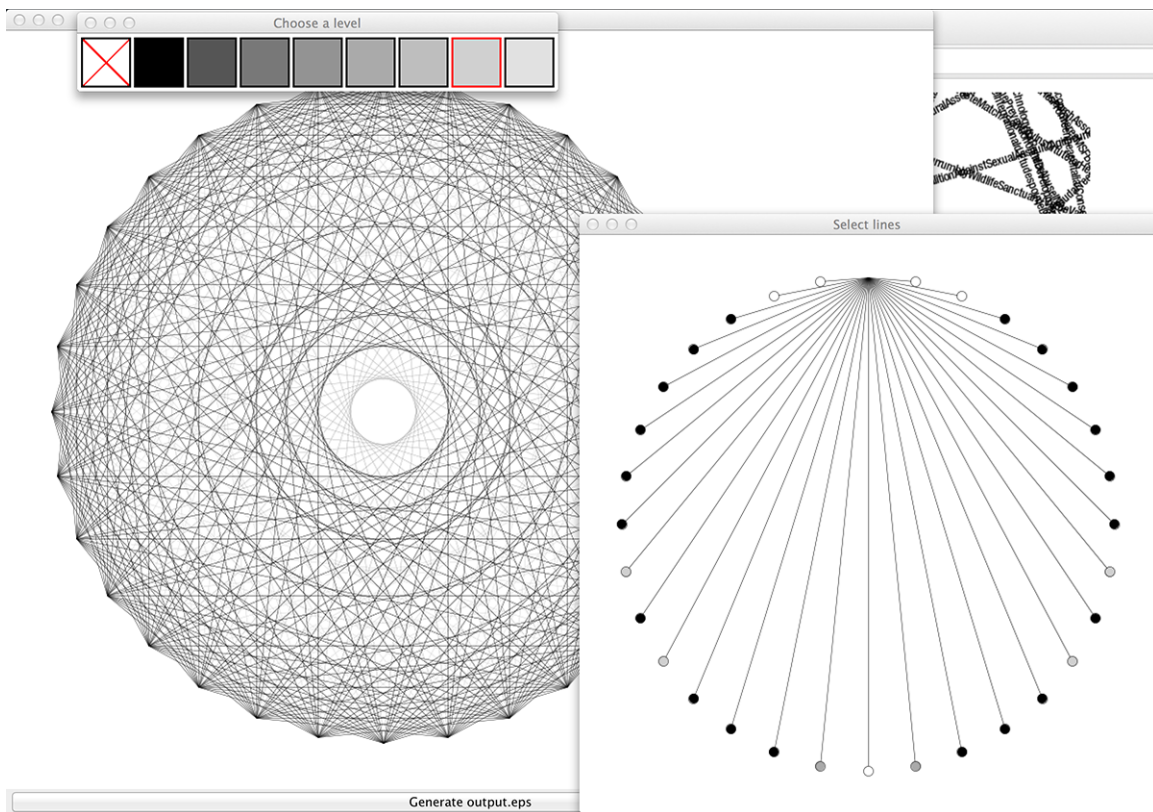


Figure 1: A screenshot of the interactive mandala design software. The small window in the top left displays a palette of eight shades of grey and a “disabled” colour. The lower right window shows the edges emanating from a single vertex; each edge terminates in a small indication of that edge’s current colour. The large window on the left shows a preview of the complete mandala.

root operation biases the selection towards larger indices, avoiding an unnecessary proliferation of short names. We subtract the length of this string, plus the length of a separator if necessary, from d . We then repeat this process, iterating until d is so small that no new names can be added. We can further reduce the residual line length remaining simply by running this algorithm multiple times and choosing the best solution. We also post-process the name sequence by randomizing its order (avoiding the risk of many lines ending with short names), after which we interleave copies of the separator string.

5 Rendering

The final step in generating *E Pluribus Unum* is to render the strings generated above along the lines in the mandala. Although text rendering might first appear to be a thoroughly solved problem, we faced several significant challenges related to the sheer scale of the design.

In our first prototype implementation, we assembled transformed lines of text into an SVG file. The SVG standard is well designed to handle type, supports unicode, and includes advanced features such as kerning and letterspacing. However, SVG rendering engines are simply not built to handle the torrent of text we were producing. We found that our attempts to display the finished composition would crash or cripple SVG viewing software.

A subtler complication arises in the computation of the geometric length w_i of each name. SVG rendering engines do not all measure text the same way, owing to differences in the treatment of hinting,

kerning, letterspacing, and ligatures or other substitutions. The layout algorithm of the previous section might incorporate a particular choice of text measurement, only to run afoul of a renderer with a different opinion. When measuring strings consisting of thousands of characters, minute discrepancies can accumulate, causing a line to overshoot its intended termination point by a wide margin.

Our solution to these problems is to move to the FreeType library (www.freetype.org), a free, open-source, cross-platform library for digital typography. With FreeType we can measure and render text under a consistent set of measurement assumptions, and we can render transformed text directly to a raster image, ensuring each character is placed precisely at its intended location.

Even with FreeType, we still did not want to render the whole design in a single image. Our planned canvas size was over 20 feet on a side, or at least 86400 pixels at 360 DPI, for a total of more than seven gigapixels. Moreover, FreeType addresses pixels using 16-bit integer coordinates, limiting image resolution to 32768 pixels on a side. Instead we implemented a tiled renderer. Given a tile, the renderer iterates over all translated and rotated names. It performs an initial clipping operation to determine if the name intersects the tile’s bounding box; if so, it uses FreeType to render antialiased greyscale pixels.

By rendering a few sample tiles at a time, we were able to ensure the correctness of the renderer and agree on the stylistic and technical parameters we wanted to use for the final design. At this point, the computer scientist needed to send a complete set of tiles to the artist. This task was itself somewhat daunting; even though the tiles permit a great deal of standard image compression, the complete

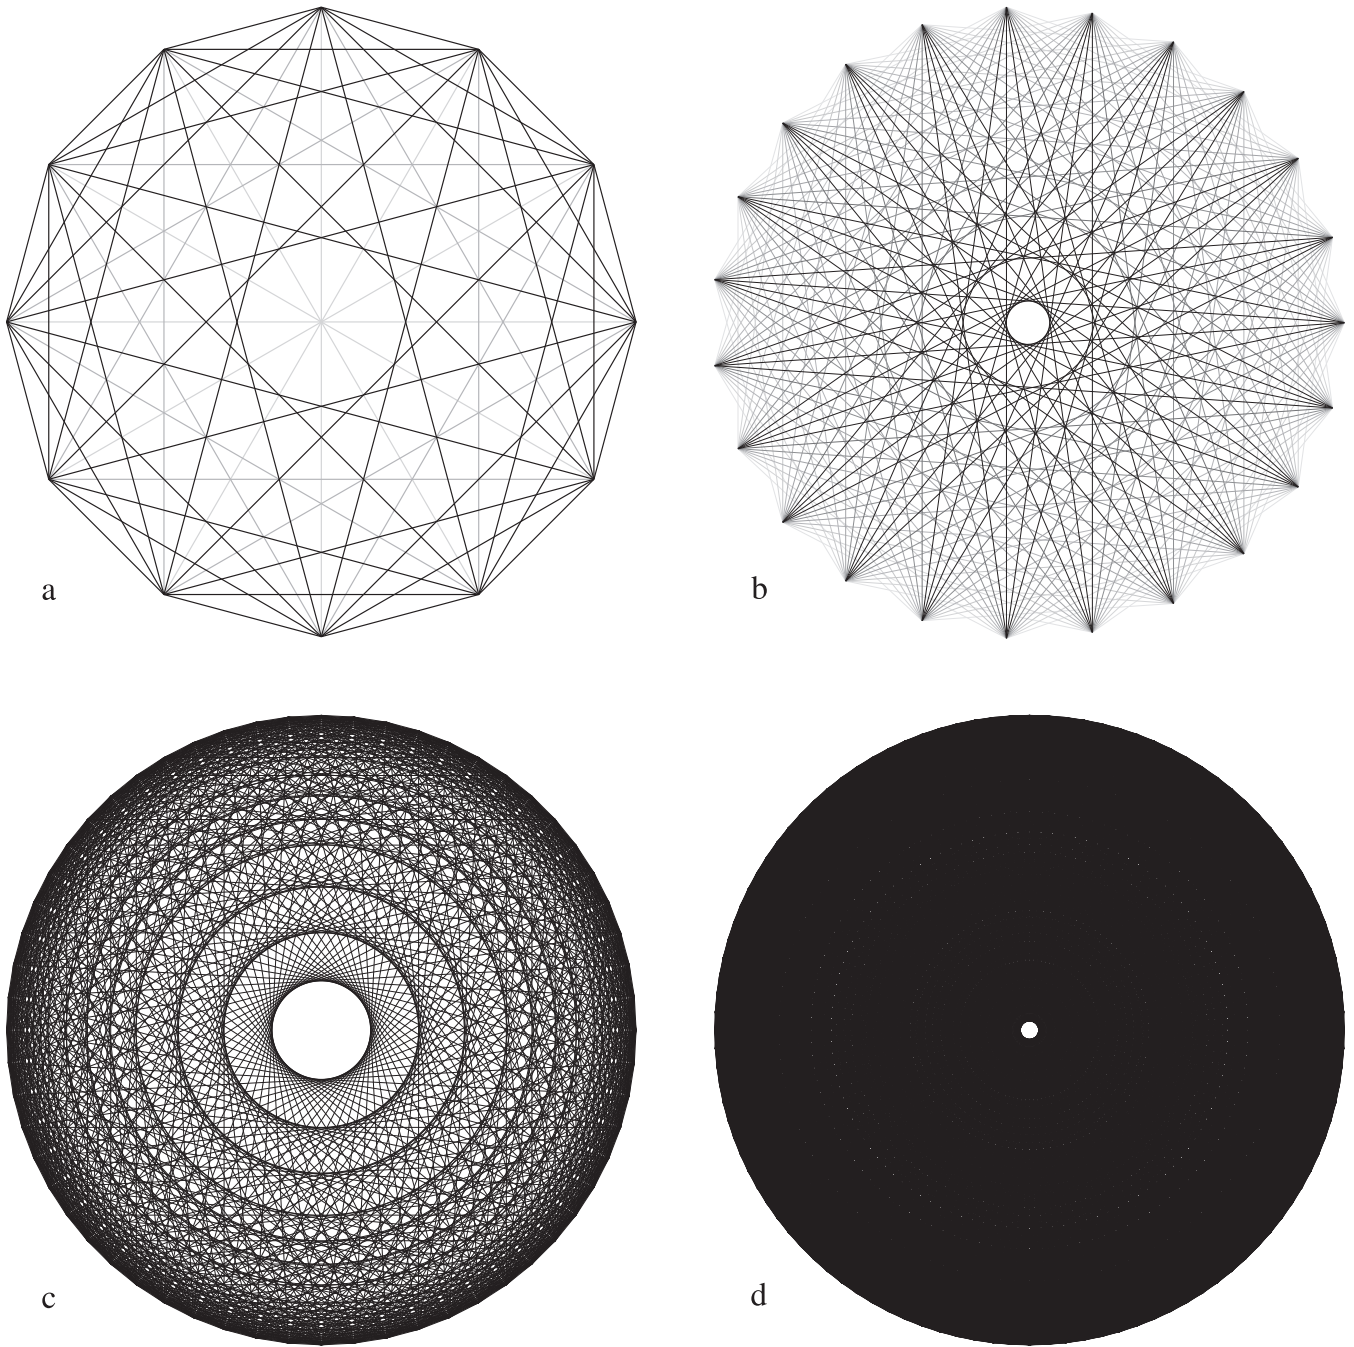


Figure 2: Four sample mandalas created using the software described in Section 3. The designs have 12, 23, 60 and 108 points. The 108-pointed mandala (d) is the design used in the final composition of *E Pluribus Unum*. At standard screen resolutions, the geometric structure the 108-pointed mandala is nearly invisible; the reader is encouraged to zoom in to see the fine details.

set would still be a heavy amount of data to transmit over the internet. In the end, the artist received the tiles in the most compressed possible form: the renderer executable, together with the ASCII-encoded list of names and a shell script that invoked the renderer to regenerate the tiles on the artist's computer.

As much as possible, text is oriented to be read left-to-right. That is, for an edge connecting two vertices, we begin the text at the leftmost of the two. Before rotating, we shift every line of text down vertically so that the centre of the line's bounding box coincides with the mandala edge, rather than the typeface's baseline.

6 The final design

With the entire system in place, we explored mandala configurations and typographic variations until we arrived at a satisfactory design. The final version is intended to be rendered at 360 DPI on a canvas 24 feet on a side, containing a mandala 22 feet in diameter. The text is set in 4 point Helvetica Neue Medium Condensed, very small but legible when viewed from a position close to the design. The small scale of the text is consistent with previous compositions by the artist, in which individual primitives sometimes hover on the edge of visibility.

In addition to using a condensed typeface, we took several other measures to increase the perception of solid lines at a distance. First, we removed all internal spaces from organization names, to avoid leaving gaps. Second, we incorporated negative letterspacing to pack characters snugly together. Third, we chose a simple slash character as the separator. Naturally, the first two changes also allowed us to pack more names into the same mandala.

The structure of *E Pluribus Unum* is itself a lesson in numbers that are difficult to comprehend. The 5724 lines of the mandala run for a total of almost 16.3 miles, and contain 1,005,714 names. Different names repeat different numbers of times, in an uneven distribution: 87% of names repeat fewer than 20 times, another 5.4% repeat more, up to a maximum of 544 times (still fewer than one in 1800 names, ensuring that repetitions are difficult to find). By luck of the draw 7.6% of names in the original list do not appear in the final design. It would be interesting to revise the layout algorithm to minimize the set of unused names, though doing so is primarily an algorithmic challenge—it is unlikely to affect the aesthetics of the design.

Figure 3 shows a sequence of images of *E Pluribus Unum*, beginning with the whole design and zooming in on a small section. At typical screen and document sizes, no single image can hope to convey the evolving forms and textures contained in the design over its full range of magnifications. By taking advantage of Deep Zoom Images in Microsoft Silverlight, we also created a web page where the design can be explored interactively (see <http://www.chrisjordan.com/gallery/epu/>).

We built a prototype physical print by mounting large-format inkjet prints on foamcore; this prototype was shown as part of an exhibition of the artist's work. Our long-term goal, however, is to fabricate *E Pluribus Unum* from laser-etched panels of anodized aluminium.

7 Conclusions

One goal of *E Pluribus Unum* was to reflect on the positive aspects of interconnectedness in the modern world. We find it fitting, then, that this work was undertaken by two collaborators who have never met face-to-face, using email to bridge a continent-sized geographical separation.

In the end, the technical challenges we faced were fairly straightforward, and did not require radical new graphics techniques. Nevertheless, we feel that the story of our collaboration, and the work that emerged from it, may be of interest to researchers interested in computational aesthetics, geometric art, and the visualization of large amounts of text.

References

- CORMEN, T. H., LEISERSON, C. E., RIVEST, R. L., AND STEIN, C. 2009. *Introduction to Algorithms*, third ed. MIT Press.
- HAWKEN, P. 2008. *Blessed Unrest: How the Largest Social Movement in History Is Restoring Grace, Justice, and Beauty to the World*. Penguin Books.
- ROSIN, P., AND COLLOMOSSE, J., Eds. 2013. *Image and Video-Based Artistic Stylisation*, vol. 42 of *Computational Imaging and Vision*. Springer.

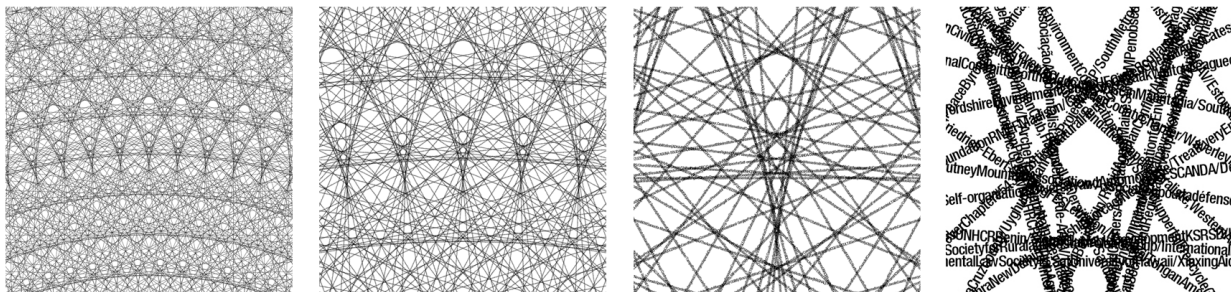



Figure 3: A low-resolution rendering of *E Pluribus Unum*. The row of smaller images at the bottom zooms in progressively on a point about one third of the way from the centre to the top of the disc, showing the progression of textures at different scales. Close inspection of the third image in the row suggests the presence of text; the text becomes readable in the fourth, which is shown slightly larger than actual size.