

Op Art rendering with lines and curves

Tiffany C. Inglis (Advisor: Craig S. Kaplan)
University of Waterloo
piffany@gmail.com

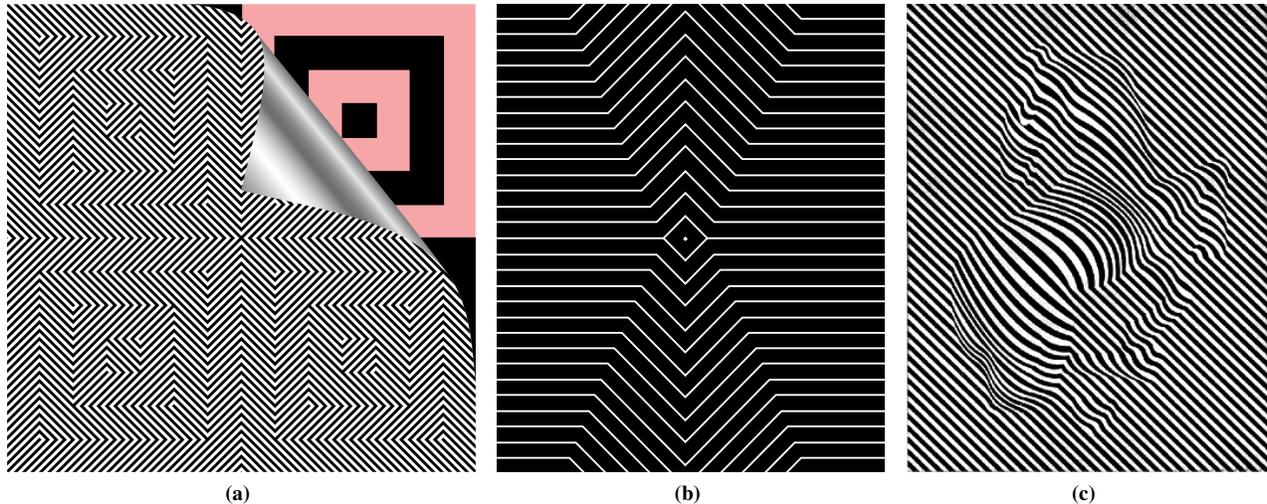


Figure 1: (a) *Square of Two* by Neal (1965) and (b) *Gavotte*: from *Black Series II* by Stella (1967) are examples of 2- and 3-colour Op Art respectively, in which each line direction represents a colour in the underlying image. (c) Vasarely's *Zèbre* (1944) is an example of curve-based Op Art.

1 Problem and Motivation

Non-photorealistic rendering (NPR) is an area of computer graphics that studies artistic expressions for digital art. Unlike traditional computer graphics that emphasizes photorealism, NPR explores a wide variety of artistic styles including painting, sketching, and animated cartoons, and develops algorithms for recreating them. One art form that is largely unexplored in computer graphics is Op Art. Short for Optical Art, Op Art is a style of abstract art associated with optical illusions that gained popularity in 1965 with an exhibition called *The Responsive Eye* [Seitz 1965]. It explores many fundamental artistic and scientific concepts, from geometric shapes and perspectives to colour theory and the psychology of visual perception. To exaggerate certain visual effects, Op artists often use simple shape primitives such as circles and lines in highly contrasting colours such as black and white. We will focus on two types of Op Art—line-based and curve-based—that make use of densely packed lines and curves to convey a hidden underlying image.

In line-based Op Art, line directions are used in place of colours to depict salient regions. Reginald H. Neal's *Square of Two* (Figure 1a), for example, uses two orthogonal line directions to represent an underlying 2-colour planar map. This particular arrangement of lines creates an optical effect where the different regions seem to scintillate as the eye moves across the image. Moreover, between two regions containing lines of different directions, there is a series of collinear line bends that creates the illusion of a boundary. A contour that is perceived without a change in luminance from one side to another is called an *illusory contour* [Petry and Meyer 1987]. In *Square of Two*, all the concentric squares are outlined with illusory contours. Notice that line breaks and T-junctions (Figure 2) are found on the vertices of these squares. Through experimentation, we found that features other than line bends interfere with the

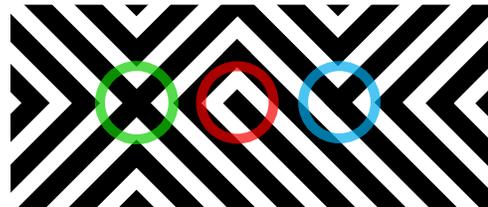


Figure 2: Examples of artifacts in line-based Op Art include: (circled, from left to right) a cross, a line break, and a T-junction.

perception of illusory contours and should therefore be avoided.

Our goal in this project was to develop an algorithm that converts a coloured planar map into a line-based Op Art composition subject to the following constraints:

1. Lines in adjacent regions travel in different directions.
2. As much as possible, boundaries between regions appear as illusory contours formed by line bends only. In other words, artifacts such as line breaks, T-junctions and crosses should be minimized, if not completely eliminated.

Square of Two is a 2-colour lined-based Op Art composition because it uses two line directions to represent a 2-colour map. However, line-based Op Art can have more than two different line directions, as demonstrated by Frank Stella's *Gavotte* (Figure 1b), which uses three line directions. To be as general as possible, we want our algorithm to operate on inputs with any number of colours.

For curve-based Op Art, the idea is to take a set of parallel lines and curve them in certain regions to create a bas-relief effect. For

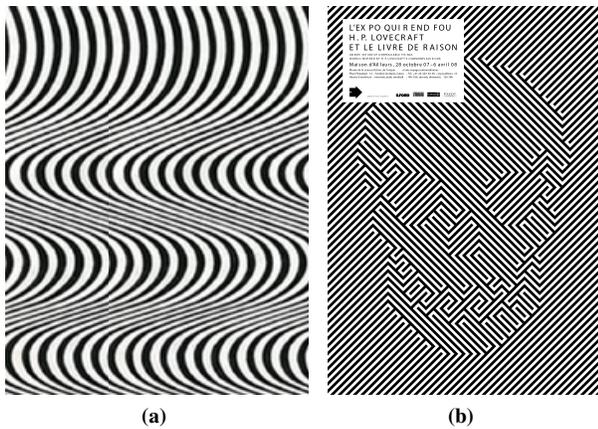


Figure 3: (a) *Current* by Riley (1964) and (b) a Lovecraft exhibition poster by Notter and Vigne (2007).

example, Victor Vasarely *Zèbre* (see Figure 1c) contains a flat background (depicted by straight lines) and a slightly elevated foreground (depicted by curves). When developing an algorithm to generate curve-based Op Art, we want to avoid two things: intersecting curves, and the use of 3D inputs. Having to specify a 3D surface places an enormous burden on the artist. We want to generate curve-based Op Art given only a decomposition of the plane into regions. The algorithm should be able to infer plausible depth information from the shapes of the regions, and then the artist can fine-tune the curves as necessary.

In this paper, we present three novel algorithms for automatically generating Op Art. The first two algorithms address different types of input for line-based Op Art. For a 2-colour input, there is an efficient algorithm that always finds artifact-free Op Art compositions. For an input with more colours, it is in general not possible to remove all artifacts. Therefore we use an optimization algorithm to minimize the number of artifacts. The third algorithm generates curve-based Op Art based on a physical simulation that curves a set of parallel lines towards the target configuration while ensuring that intersections do not occur.

2 Background and Related Work

Many Op artists have experimented with line-based and curve-based Op Art because the amount of contrast between the densely packed lines and curves creates a variety of optical effects. Bridget Riley’s *Current* (Figure 3a) uses curves to express current flow, while a Lovecraft exhibition poster by Julien Notter and Sébastien Vigne (Figure 3b) uses line-based Op Art to depict a skull.

Op Art also has a significant influence on graphic and fashion design. Dioptical and OPTICA Normal (Figure 4a and 4b) are two typefaces designed using 2-colour line-based Op Art. A Brazilian anti-smoking ad campaign used Op Art to show the nauseating effect of smoking. In fashion design, Op Art patterns were prominently featured in a recent collection by Alexander McQueen.

Despite the prevalence of line-based and curve-based Op Art, there are no general approaches to creating them. The letters in the Dioptical font appears to be assembled through a tiling method. However, this approach only creates blocky designs and cannot convey curvy illusory contours unless the line density is much higher. For curve-based Op Art, there are related techniques developed for digital micrography [Maharik et al. 2011] and facial engraving [Ostro-

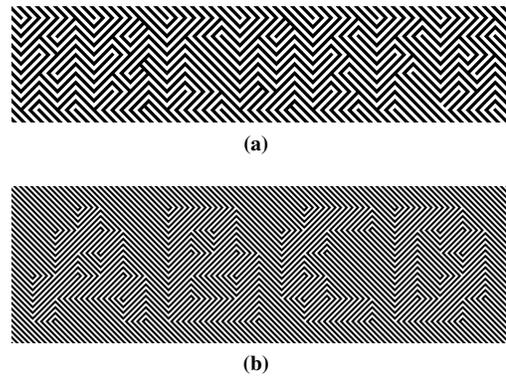


Figure 4: Two typefaces inspired by Op Art: (a) *OPTICA Normal* by Manuel Guerrero, and (b) *Dioptical* by House42.

moukhov 1999]. Some scientific analyses have also been performed on well-known Op Art compositions. Dodgson studied the regularity and randomness of circular patterns in Riley’s works [Dodgson 2008], while Troncoso et al. analyzed the perceived corner salience in Vasarely’s works [Troncoso et al. 2009] from a neuroscience perspective.

3 Approach and Uniqueness

We developed two algorithms for creating line-based Op Art, and one for curve-based Op Art. After evaluating the outputs, we describe the types of inputs that are best suited for creating Op Art. Finally, we introduce a hybrid Op Art that combines the strengths of both styles.

3.1 2-colour algorithm for line-based Op Art

Given a 2-colour planar map, we want to find the corresponding 2-colour line-based Op Art composition that uses two different line directions to represent the colours. To start, we will use horizontal and vertical lines. The goal is to create clearly defined illusory contours as region boundaries, which means the artifacts should be entirely, or at least mostly, removed.

Given a 2-colour map and a desired line spacing s for the Op Art, our algorithm proceeds as follow. First, create a square grid with $s \times s$ squares and resample the input map on it so that each square contains one of two colours (see Figure 5a). Then replace each colour with a line direction by filling in the interior of each region (see Figure 5b). Along the region boundaries, there are loose ends waiting to be joined to remove artifacts. To do so, draw alternate edges along the boundaries (Figure 5c). We proved that this method is guaranteed to remove all the artifacts along the region boundaries [Inglis and Kaplan 2011]. Lastly, to maximize the contrast of the image, the black lines are thickened until they are as wide as the white gaps (see Figure 5d).

Some line-based Op Art compositions such as Neal’s *Square of Two* use diagonal lines, which is equivalent to rotating the horizontal and vertical lines by 45° . This effect can be achieved with our algorithm by first rotating the input by -45° (in other words, apply the inverse transformation), applying the above algorithm, and finally rotating the output by 45° . In general, any invertible transformation can be applied to the lines.

The 2-colour algorithm scales linearly with the size of the underlying grid, which depends on the input size and the line spacing.

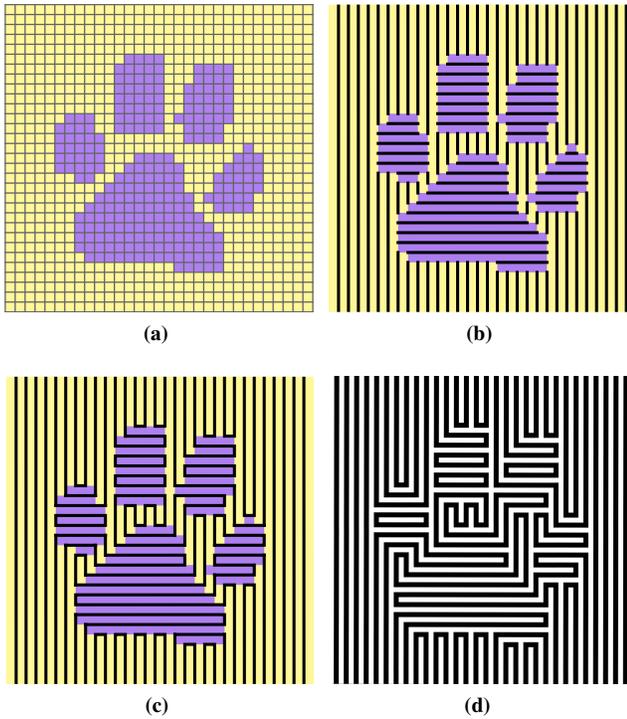


Figure 5: The 2-colour algorithm for creating line-based Op Art is as follows: (a) put the input 2-colour map on a square grid, (b) fill each region's interior with lines corresponding to the region colour, (c) draw alternate edges along region boundaries, and (d) render the lines with the appropriate thickness.

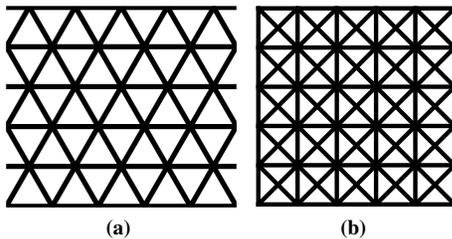


Figure 6: (a) 3-colour equilateral triangular grid, and (b) 4-colour grid with horizontal, vertical, and 45° diagonal lines.

For any reasonable-sized input, the corresponding Op Art composition can be generated in real-time, which means we can create an interactive drawing tool that allows the user to paint 2-colour Op Art. An interesting application of this algorithm would be to create animated Op Art.

3.2 3- and 4-colour algorithm for line-based Op Art

To generalize the algorithm further, we would like to create Op Art compositions from n -colour maps, where $n > 2$. If n is large, using one line direction per colour would result in illusory contours that are difficult to read because the lines may bend at shallow angles. To solve this problem, we want to reduce the number of colours by recolouring. The Four Colour Theorem states that any planar map can be coloured using at most four colours such that neighbouring regions have different colours. Thus any n -colour input can be re-

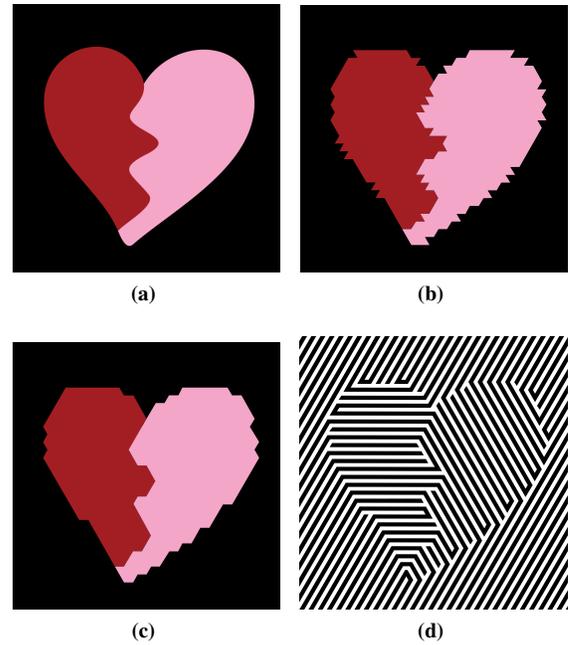


Figure 7: Preprocessing can help create smoother Op Art results. (a) First, choose a 3-colour input. (b) When resampling it on a triangular grid, the region boundaries may become jagged. (c) We apply a smoothing operation to remove acute angles. (d) Then we use simulated annealing to find an optimal Op Art composition.

duced to up to four colours, which means we only need to extend the line-based Op Art algorithm to handle 3- and 4-colour inputs.

For 3-colour Op Art, we use the three line directions found on a grid of equilateral triangles (Figure 6a) so that the angles formed are always multiples of 60° and that the three families of lines have the same equal spacing. The obvious approach for a 3-colour input is to extend the 2-colour algorithm where region interiors are first filled in, and then alternate boundary edges are drawn. However, this method does not remove all artifacts. In fact, there are simple cases in which it is impossible to remove all the artifacts along region boundaries. Therefore, instead of attempting to remove all artifacts, we will focus on minimizing them.

We perform an optimization based on simulated annealing [Press et al. 1992]. Consider each Op Art composition as a configuration of the triangular grid edges; that is, an edge is either drawn or not. Suppose we have chosen a line direction to represent each colour. Then it remains to find the set of edges along the region boundaries to draw such that the number of artifacts is minimized. Simulated annealing is a technique for searching this space effectively. We define the cost of a configuration as the number of artifacts it contains. Beginning with a random configuration, we decide at each iteration whether to move from the current configuration to a neighbouring configuration. If the neighbour has a lower cost, then we always accept the move. Otherwise, the decision is based on an acceptance probability. This probability diminishes over time, making it less and less likely to accept a less optimal configuration. Eventually, the annealing process settles on one configuration.

Since there are six ways to assign three line directions to three colours, we apply simulated annealing six times to get the optimal configuration in each case. The best of the six is then returned as the final Op Art composition. While the output is not guaranteed to be

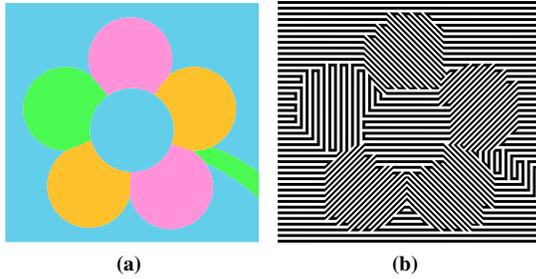


Figure 8: (a) Suppose the input has four colours. (b) Simulated annealing can be applied to find an optimal Op Art composition, but we use two different line thicknesses to ensure all regions are approximately 50% grey.

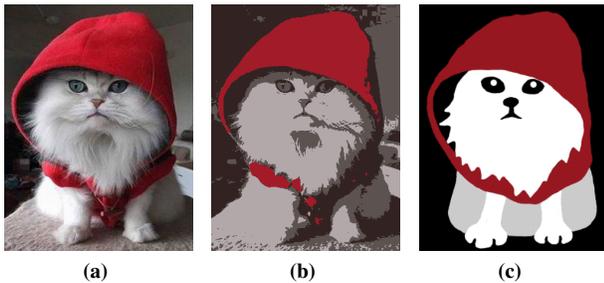


Figure 9: To convert (a) a photograph into a suitable input for the Op Art algorithm, we can apply (b) colour quantization to reduce the palette size. However, (c) simple images such as cartoons typically generate better Op Art results.

artifact-free, the Op Art compositions generated using this method contain relatively few artifacts that do not significantly affect the illusory contours. If necessary, we can run the simulation for a longer period of time to find configurations with fewer artifacts.

Unlike the square grid, working with a triangular grid has the disadvantage that smooth boundaries may become jagged when resampled on the grid (Figure 7a and 7b), resulting in a less accurate Op Art depiction of the input. To fix this problem, we apply a morphological smoothing operation prior to optimization that removes all acute angles along region boundaries. Figure 7 illustrates the steps in the 3-colour algorithm.

For 4-colour inputs, we use the same optimization approach. The only difference is that the underlying grid now contains four different line directions (Figure 6b). Notice that in this grid, the spacing between the horizontal and vertical lines is different from the spacing between diagonal lines. This discrepancy implies that if the same line thickness is used, some regions will look darker than others. To remove the tonal difference, we will adjust the line thickness in each region so that the black lines are as wide as the white gaps. Figure 8 shows one input and result of the 4-colour algorithm.

3.3 Suitable inputs for line-based Op Art

The examples shown so far use simple inputs, but what if we wanted to convert a more complex image, such as the photograph in Figure 9a, into line-based Op Art? We will need some way of processing it to make it a suitable input. There are many colour quantization algorithms that can reduce the colour palette of an image (Figure 9b), but they do not provide large regions with smooth

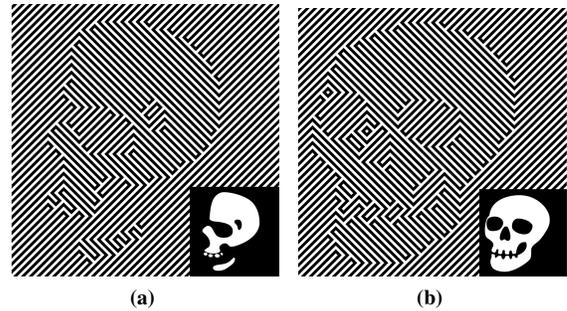


Figure 10: The type of input image can affect the Op Art compositions produced. (a) An input image that focuses on tonal variation translates into a less readable Op Art composition. (b) An input image that emphasizes shape over tones generates a more readable Op Art composition.

boundaries, leading to Op Art compositions that are difficult to decipher. For greyscale images, algorithms such as artistic thresholding [Xu and Kaplan 2008] can generate 2-colour maps with smoother boundaries. For coloured images, image segmentation can be used to create maps with large salient regions. In general, simpler inputs such as cartoon images (see Figure 9c) are better suited for creating Op Art than complex inputs generated from photographs and paintings.

Even among simple input images, some produce better Op Art compositions than others. Consider the two skull images in Figure 10: the left one provides information about the lighting while the right one provides information about the shapes. Line-based Op Art is based entirely on communicating shape, because the lines always approximate 50% grey. Thus, the Op Art composition in Figure 10a is difficult to decipher because the tonal variation was lost when the input was converted into Op Art. In summary, the best type of input for generating line-based Op Art should contain large regions with smooth boundaries, and it should also not rely on tonal information for readability.

3.4 Algorithm for curve-based Op Art

Curve-based Op Art is composed of a series of densely packed lines that are curved in a foreground region to create the effect of a bas-relief (Figure 1c). Our goal is to develop an algorithm that converts a 2D input—a silhouette defining the background and foreground regions—into a curve-based Op Art composition in which the curves do not touch or intersect. Working from a 2D input is easier because the artist does not have to create a 3D model, although our algorithm supports 3D inputs as well.

Our algorithm starts with a set of vertical lines representing a flat background. From the input silhouette (Figure 11a), we can approximate the surface normals of the intended 3D shape. These normal vectors are then mapped to a heat field which is used to curve the vertical lines. The surface normals are computed by applying Sobel edge detection to the input (see Figure 11b), then interpolating the gradient values (see Figure 11c). The interpolated gradient values are mapped to the range of -1 to 1 , representing heat sinks (in blue) and heat sources (in red). The idea is to curve the lines left towards heat sinks and right towards heat sources until equilibrium is reached (see Figure 11d). In other words, given the heat sinks and sources, we allow the system to reach equilibrium and use the final temperature to displace the lines.

To simulate heat flow, we use the finite difference method [Thomas

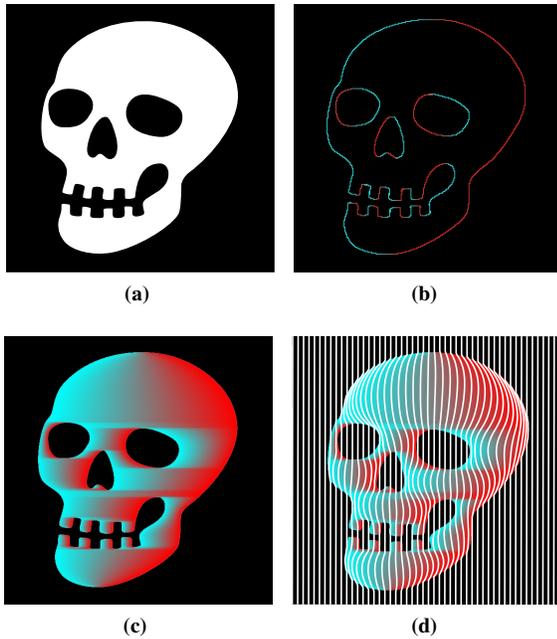


Figure 11: (a) Starting with the input silhouette, (b) apply Sobel edge detection and (c) calculate the heat sources (red) and sinks (blue) by interpolating the gradient values. (d) Allow the temperature to reach equilibrium, then curve the lines by map the temperature to the displacement.

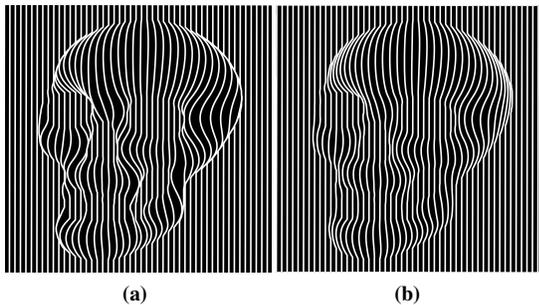


Figure 12: (a) The initial simulation may produce uneven spacing among the curves. (b) We can fix it by adding repulsive forces between neighbouring curves.

1995]. For small enough time steps, we can ensure that the paths do not touch or intersect one another. However, the spacing between neighbouring curves can still be too small or too large (see Figure 12a). To obtain a more even spacing, we add a small repulsive force between every pair of neighbouring paths (see Figure 12b).

3.5 Suitable inputs for curve-based Op Art

The inputs suitable for 2-colour line-based Op Art are also good candidates for curve-based Op Art, but unlike line-based Op Art, the readability of illusory contours is a big concern in curve-based Op Art. An illusory contour is defined by a series of the line bends, and in curve-based Op Art, the line bends are sometimes too shallow. Also, when a region is small, the lines inside can only curve by a small amount. As a result, the region looks flat and nearly

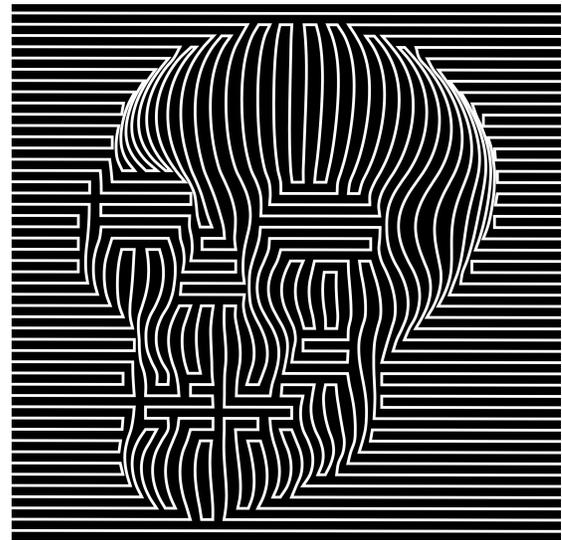


Figure 13: A skull rendered in the hybrid Op Art style.

indistinguishable from the background. Therefore, a suitable input for curve-based Op Art should contain foreground regions that are large relative to the line spacing.

3.6 Hybrid Op Art

In line-based Op Art, the regions are more clearly defined but they all appear flat. In contrast, curve-based Op Art creates the illusion of depth through curves but the illusory contours are less readable. To combine the advantages of both styles, we propose a hybrid form of Op Art that uses 2-colour line-based Op Art as the base configuration, with the foreground lines curved to depict a 3D surface. Figure 13 shows an example of a skull rendered in this style, where the skull is depicted using curves, and between the foreground and background, there are sharp line bends that clearly indicate where the illusory contours are.

4 Results and Contributions

We described three novel algorithms for automatically generating line-based and curve-based Op Art. We analyzed the results and determined the types of inputs that produce successful Op Art compositions. Finally, we demonstrated a hybrid art form that combines the strengths of line-based and curve-based Op Art.

There are several avenues for future work. For the line-based Op Art algorithm that uses simulated annealing to find the optimal Op Art configuration, we would like to improve its efficiency by exploring a large set of moves to search the configuration space. For curve-based Op Art, we want to improve the readability further by finding the best line direction to use for a given input and exaggerating certain features by increasing the amount of curving.

In addition to creating Op Art, our algorithms could also be used for algorithms in recreational computational geometry, such as designing mazes [Xu and Kaplan 2007] or puzzles containing hidden images. By computing suitable mesh parameterizations, Op Art designs could also be created for 3D surfaces to be used for fashion design or decorating sculptures.

References

- DODGSON, N. A. 2008. Regularity and randomness in Bridget Riley's early Op Art. In *Proceedings of Eurographics Workshop on Computational Aesthetics in Graphics, Visualization, and Imaging*, Eurographics Association, 107–114.
- INGLIS, T. C., AND KAPLAN, C. S. 2011. Generating Op Art lines. In *Proceedings of the International Symposium on Computational Aesthetics in Graphics, Visualization, and Imaging*, ACM, CAe '11.
- MAHARIK, R., BESSMELTSEV, M., SHEFFER, A., SHAMIR, A., AND CARR, N. 2011. Digital micrography. *Transactions on Graphics (Proc. SIGGRAPH 2011)*, 100:1–100:12.
- OSTROMOUKHOV, V. 1999. Digital facial engraving. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, SIGGRAPH '99, 417–424.
- PETRY, S., AND MEYER, G. E. 1987. *The Perception of Illusory Contours*. Springer.
- PRESS, W. H., FLANNERY, B. P., TEUKOLSKY, S. A., AND VETTERLING, W. T. 1992. Simulated annealing methods. In *Numerical Recipes in C: The Art of Scientific Computing*, second ed. Cambridge University Press, ch. 10.9, 444–455.
- SEITZ, W. C. 1965. *The Responsive Eye*. Museum of Modern Art, New York, N.Y.
- THOMAS, J. 1995. *Numerical Partial Differential Equation*. Springer.
- TRONCOSO, X. G., MACKNIK, S. L., AND MARTINEZ-CONDE, S. 2009. Corner salience varies linearly with corner angle during flicker-augmented contrast: a general principle of corner perception based on vasarely's artworks. *Spatial Vision* 22, 3, 211 – 224.
- XU, J., AND KAPLAN, C. S. 2007. Image-guided maze construction. *ACM Transactions on Graphics* 26 (July), 29–1–29–9.
- XU, J., AND KAPLAN, C. S. 2008. Artistic thresholding. In *Proceedings of the 6th international symposium on Non-photorealistic animation and rendering*, ACM, NPAR '08, 39–47.