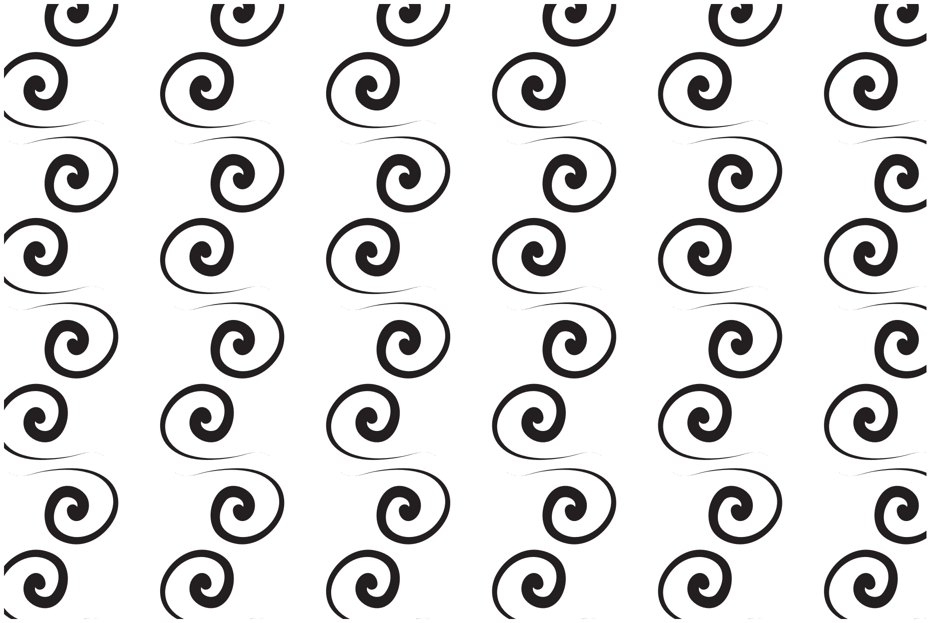

Figure 1: A symmetric pattern

Producing Symmetric Patterns with Contravarient Distributions

Chaim Goodman-Strauss
The Symmetry Group

cgstraus@comp.uark.edu

1 Introduction

Symmetric images have been common in virtually every culture on Earth at least since the Neolithic. Systematic methods for generating symmetric patterns have been developed for centuries and this has continued apace with the rise of computer technology. In essence, artists, artisans, designers and programers have always produced symmetric images by selecting some motif and replicating it in a regular manner. Here we give a new method for selecting motifs for replication using "contravarient distributions".

Here, we will first quickly outline the general methods used for producing patterns to date; all prior art will be seen as applications of these methods. We then will outline this new technique. We have suppressed the mathematics as much as possible and kept things informal.[1]

To begin, consider the image in Figure 2. The motif is a rectangle, marked with a swirl. It is replicated by rotating the motif by 180 degrees about the marked points on the rectangle. This procedure is repeated for each of the newly generated motifs, ad infinitum

[1] Accuracy requires, though, a few terms which the non-mathematician is free to ignore. We will try to confine this to footnotes. Strictly speaking, we are considering *discrete group actions* on the plane, sphere and other geometric spaces

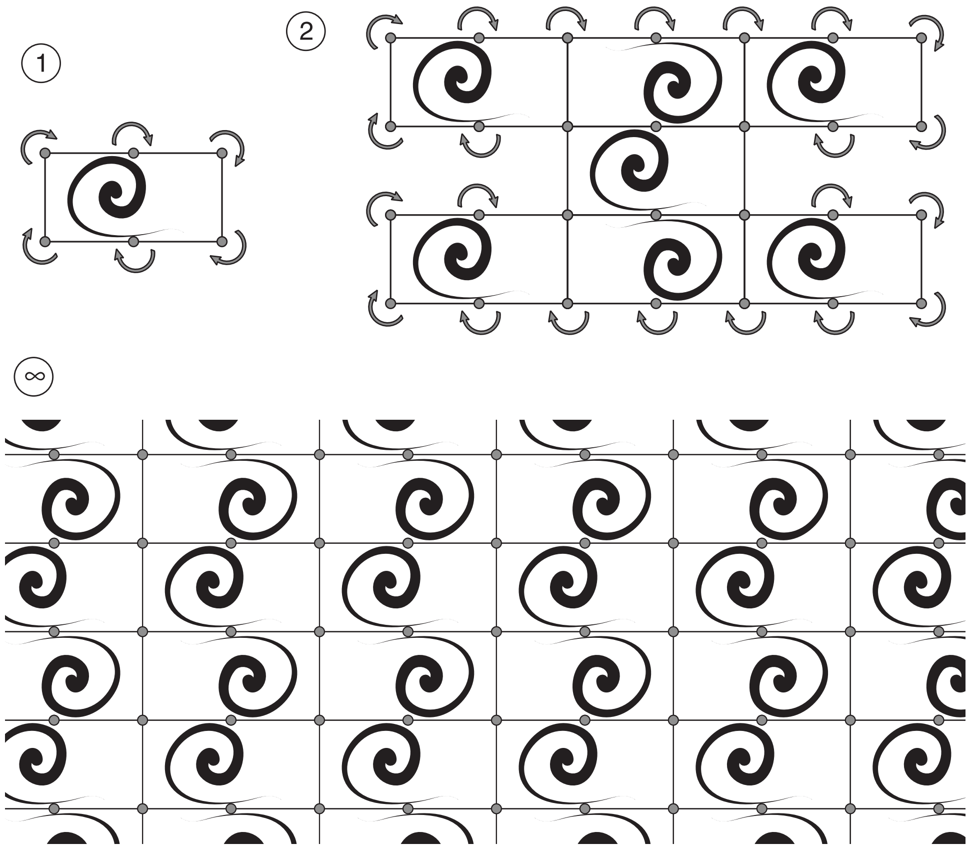

Figure 2: The regular repetition of a motif

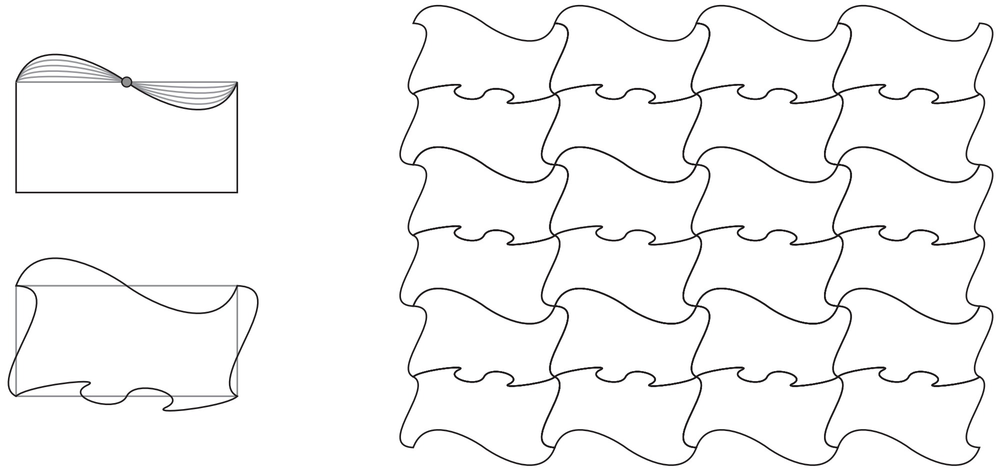

Figure 3: Adjusting the fundamental domain

(for the mathematician) or until a large enough piece of the pattern has been produced (for the designer).

There are only a few ways such a repetition can be carried out. First, there are only four kinds of transformation in the Euclidean plane: translations (a slide), rotations about point, reflections across a line, and glide reflections, which consist of a reflection and translation together. Up to trivial notions of equivalence, there are exactly 17 ways these operations can be combined to form a symmetrical pattern with a finite motif[2]

(There are a few more kinds of planar symmetry with infinite motifs or such that the pattern does not cover the whole plane. In addition, one can discuss symmetries of the sphere, or other geometric spaces. Our comments apply to all of these as well.)

These 17 types were first classified by Federov in 1885, then rediscovered by Polyá, Niggli and others in the 1920's. Grünbaum and Shephard gave a new derivation of this classification in the late 1970's. At about the same time W. Thurston gave an extremely elegant approach, which seems to give the "correct" mathematical perspective on this problem. The symmetry of Figure 2 is denoted $p2$ in the "international symbol" system widely used by crystalographers, and 2222 in the newer notation developed by J. Conway based on Thurston's classication.

However, it is quite interesting that this same classification was carried out, more or less naively, over and over again throughout the world over thousands of years. All 17 types of planar patterns appear, for example, in the pottery of the Pueblo Indians, the Alhambra, and elsewhere. Virtually every culture makes use of at least several of these types.

Implicit in the idea of generating a pattern using a motif is that of a "fundamental domain". A fundamental domain for a pattern is a tile that can be used as the motif; after replicating the tile, these fundamental domains exactly cover the pattern. For example, in Figure 2, the fundamental domain is a rectangle. For a given symmetry type, however, there may be infinitely many possible choices for fundamental domain.

For example, in the upper left of Figure 3, one edge of the original rectangle is being steadily modified, producing an infinite family of new fundamental domains. This can be carried out in general; the only constraints are that edges that correspond in the pattern must be changed together, and that the edges as a whole remain embedded (i.e. don't

[2] That is, in the Euclidean plane, up to affine conjugacy there are exactly 17 discrete group actions with compact fundamental domain.

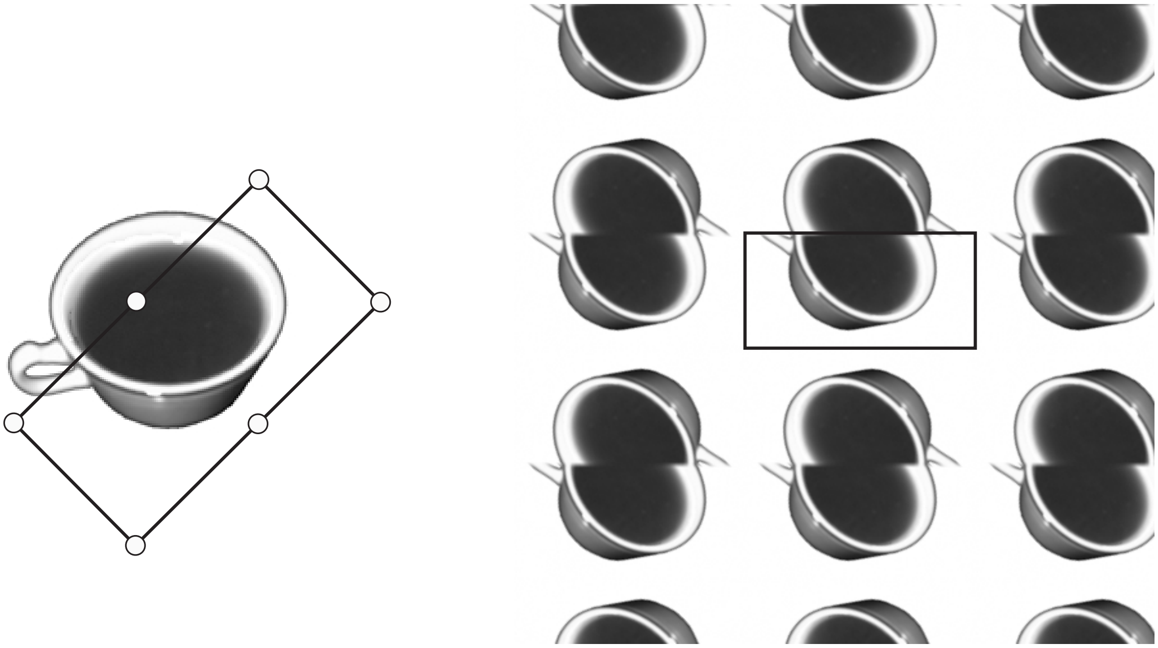

Figure 4: Selecting a portion of an image as the motif

criss-cross one another). But this kind of transformation does not change the particular way the fundamental domains meet one another— in this case, for example, they still meet four-to-a-corner. They remain of the same "Heesch type".

Up to a certain notion of equivalence, there are only a few distinct types of possible fundamental domain, first classified by Heesch in the 1930's and then independently by M.C. Escher about ten years later.[3] A given symmetry type may allow several distinct these "Heesch types" of fundamental domain— $p2$ for example admits four Heesch types of fundamental domain— but only finitely many, and in fact, all together, there are only 42 (?) Heesch types for planar symmetries.

It was precisely by analyzing the possible Heesch types of fundamental domains that allowed M.C. Escher to create his famous and remarkable prints. As is well known, he realized one could modify the edges of a given fundamental domain to resemble an image of any number of things. In fact, his basic technique is so useful that it is now used to excite and educate geometry students all over the world.

2 Software

With the advent of computer imaging techniques, various software packages have now been developed allow the user to select a portion of an image to use as the fundamental domain for a symmetric pattern.

This idea is illustrated in Figure 4; our original fundamental domain of Figure 2 has been placed over an image of a cup of coffee, and the resulting selection has been replicated to produce the pattern at right. The result is somewhat unsatisfying: one can clearly see the boundary of our original rectangle where it abrubtly truncates the image. To a limited

[3] Again, the new techniques of Thurston give a much more elegant classification of the types of fundamental domains.

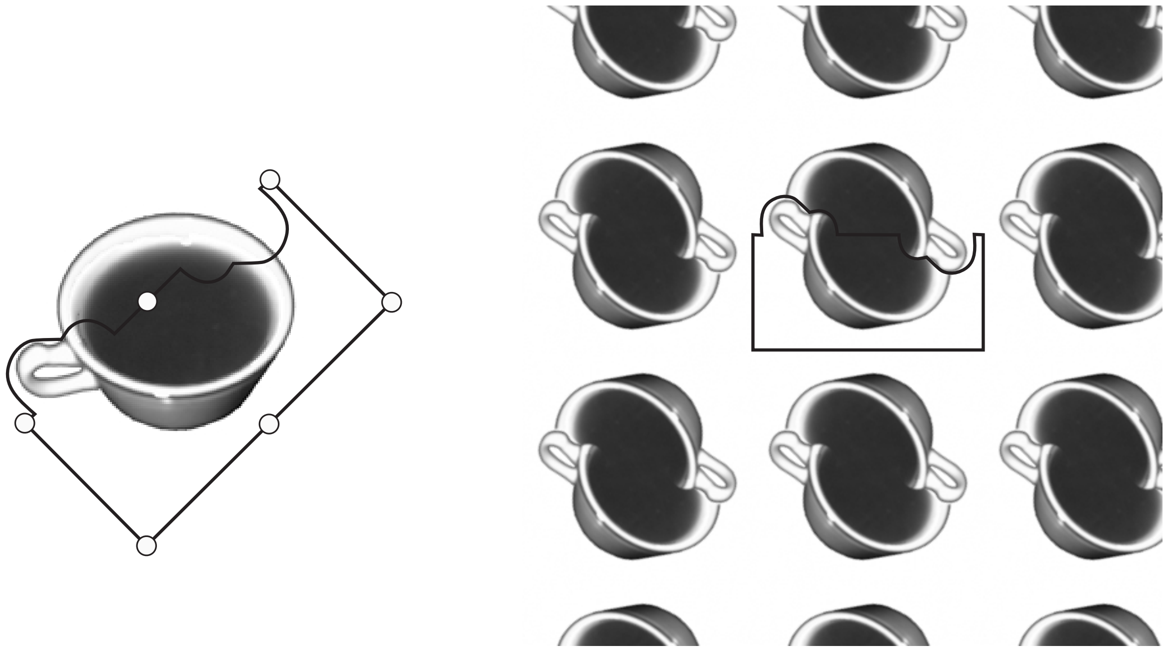

Figure 5: Adjusting the selection of an portion of an image as motif

degree, this technique has been implemented in *Tess*, *Terrazo*, and a small number of other pieces of software.

To date, only one commercial piece of software, KaleidoMania, has gone as far as the technique in Figure 5: in this figure, the edge of the fundamental domain has been modified, as in Figure 3, in order to follow the handle of the cup more closely (at the expense of the corresponding piece of rim).

There is one other technique that has appeared: "equivarient" drawing tools. As one brushes with an equivarient brush, for example, one sees copies of the brush at all corresponding points in the pattern.[4] With such a tool, one can draw designs such as that of figure 1. This has been implemented in *Lafite* and *Kali*.

All of this must be considered prior art; though no one single piece of software or *** encompasses all the ideas presented so far, each idea has appeared somewhere or another. In practice, however, there are technical difficulties in applying the idea of Figure 5 in full generality.

3 Contravarient distributions

We now introduce the idea of "orbits". In a symmetric pattern, there are many copies of each point in the original motif. At left in Figure 6, for example, the point labeled a lies at the tip of our swirl. The corresponding points in the pattern, at the tip of each copy of the swirl, form a single orbit. Similarly, each point in the pattern lies in exactly one orbit. And, in essence, a fundamental domain is a tile consisting of exactly one point from each orbit. The symmetric pattern is produced by replicating this tile. In the symmetric pattern, all of the points in each given orbit are colored precisely the same way.

[4] Hence the term "equivarient": each copy of the brush moves together

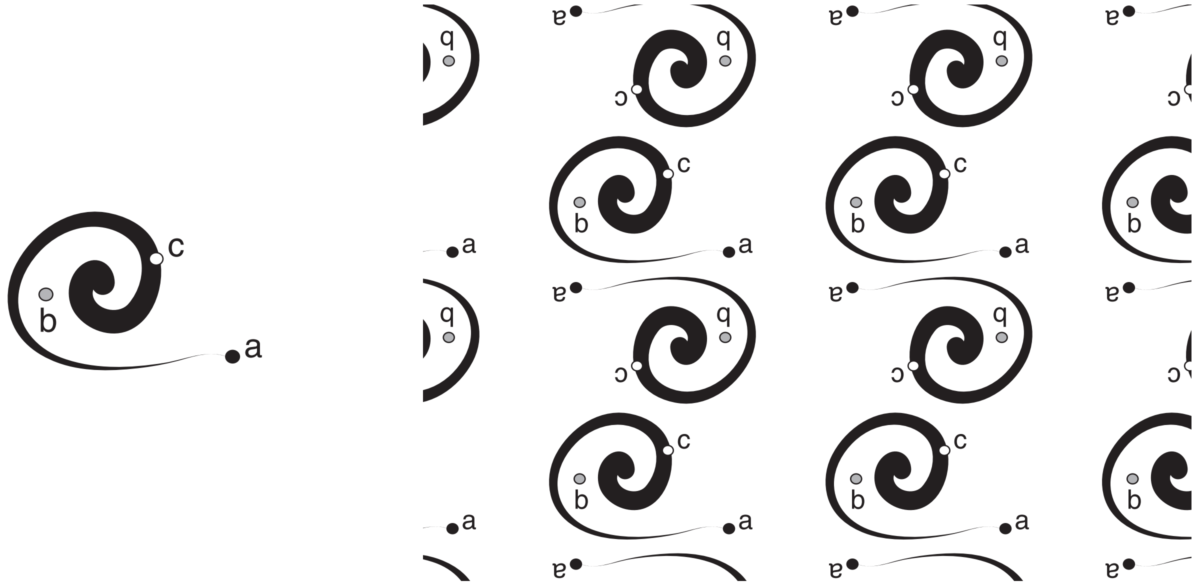

Figure 6: Orbits in a pattern

We now can discuss the new idea more fully. We provide an alternative method for selecting how each orbit should be marked in the pattern. Recall that *all* previous techniques of this sort have made this selection by first chosing a fundamental domain; each orbit corresponds to a point in the fundamental domain, and is colored in the same way as this point.

The essential idea here is to allow the fundamental domain[5] to consist of, say, half a point here, a tenth of a point there, etc, so long as the total value for any given orbit is 1.

More precisely: a "distribution" is simply a function assigning to each point in the plane some value. A distribution is "contravarient" under some symmetry if the values of the points in each orbit add up to 1. Given a contravarient distribution, we create a symmetric pattern by taking the weighted average of the points in the original image; the weights are the values of the distribution.

That is, an image is a map $I : X \to C$, where C is some space of colors or other quantities; we assume that we may take weighted averages of values in C. A distribution is a map $D : X \to \mathbf{R}$, where $\mathbf{R}$ is the set of real numbers. A contravarient distribution satisfies:

For all orbits O,

$$\Sigma_{x \in O} D(x) = 1$$

A symmetric image I_s is then defined by taking for each orbit:

$$I_s(O) = \Sigma_{x \in O} I(x) \cdot D(x)$$

and setting for all $x \in O$, $I_s(x) = I_s(O)$, producing the output image.

For example, in the following figures, we will illustrate distributions with shades of gray. If a given point x is shaded black, $D(x) = 0$; if x is white, $D(x) = 1$; if x is 50% gray, $D(x) = .5$, and so forth. We will mark the rotation points with white circles.

[5] The term is no longer accurate, technically, but we continue to use it informally.

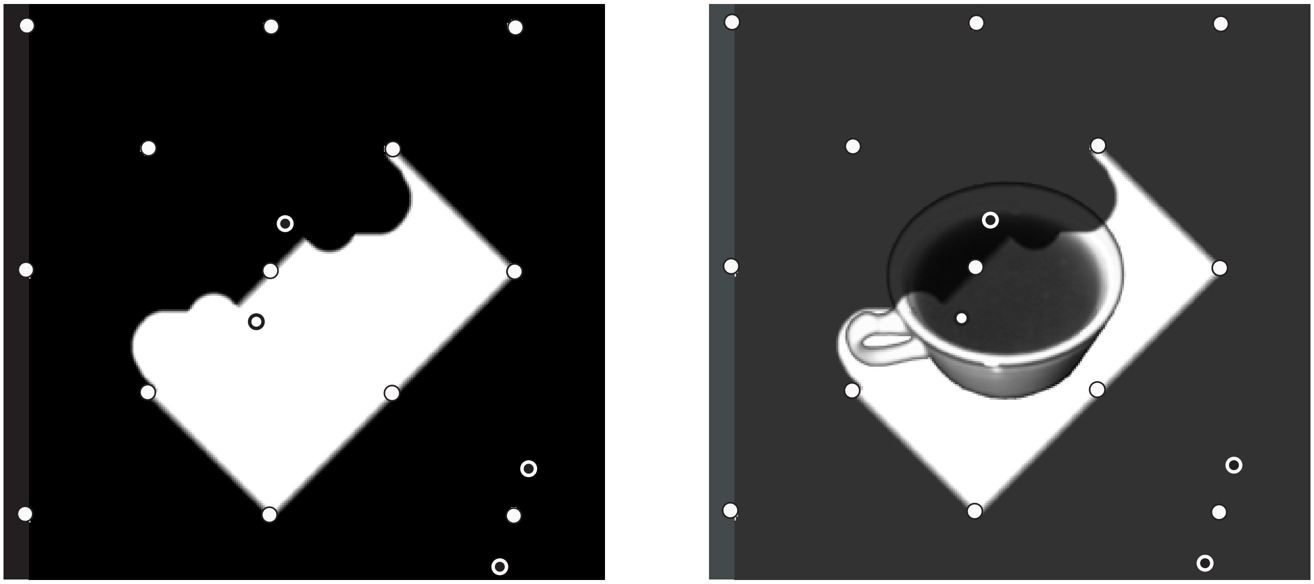

Figure 7: Using a distribution to define a fundamental domain

The images at left in Figure 7 and at top in Figure 9, and both images in Figure 8 represent contravarient distributions. The sum of values $D(x)$ in each orbit add to one, or to put it another way, the grayscale values over corresponding points in the symmetry sum to white.

At left in Figure 7, the contravarient distribution takes on only values 1 and 0; in other words, there are no shades of gray in the illustration of the distribution. At right in in Figure 7 we see the distribution overlaid on the original image. The distribution in this figure produces precisely the same image as that at right in Figure 5: in each orbit O, there is precisely one point x_O with value $D(x_O) = 1$, corresponding to a point in the fundamental domain of Figure 5. The remaining points x in each orbit have $D(x) = 0$. Consequently, we color the points in the orbit by

$$I_s(O) = \Sigma_{x \in O} I(x) \cdot D(x) = \Sigma_{x \in O, y \neq x_O} I(x) \cdot 0 + I(x_O) \cdot 1 = I(x_O)$$

So this method certainly encompasses earlier techniques. But allowing the distributions to take on a wider range of values is quite powerful. We can now anti-alias— or bleed out— the edges of our selection quite easily, by having the values of the distribution gently slope down from 1 to 0. We are no longer constrained by choice of Heesch type, and can freely select any sort of possible fundamental domain we please. In fact, we no longer need to have fundamental domains that are in a single connected piece. But most importantly, we can produce a completely new sort of symmetric images.

In Figure 9 we illustrate a more elaborate contravarient distribution. A portion of one orbit O is singled out specially with shaded circles; the shades correspond to the values of the distribution at these points. At lower left, we have illustrated this orbit O again, superimposed on the original image; here there is now room for labels $a_1, \ldots, a_5$ for the points in O that lie inside the bounds of the picture.

We have

$$D(a_1) = .54$$

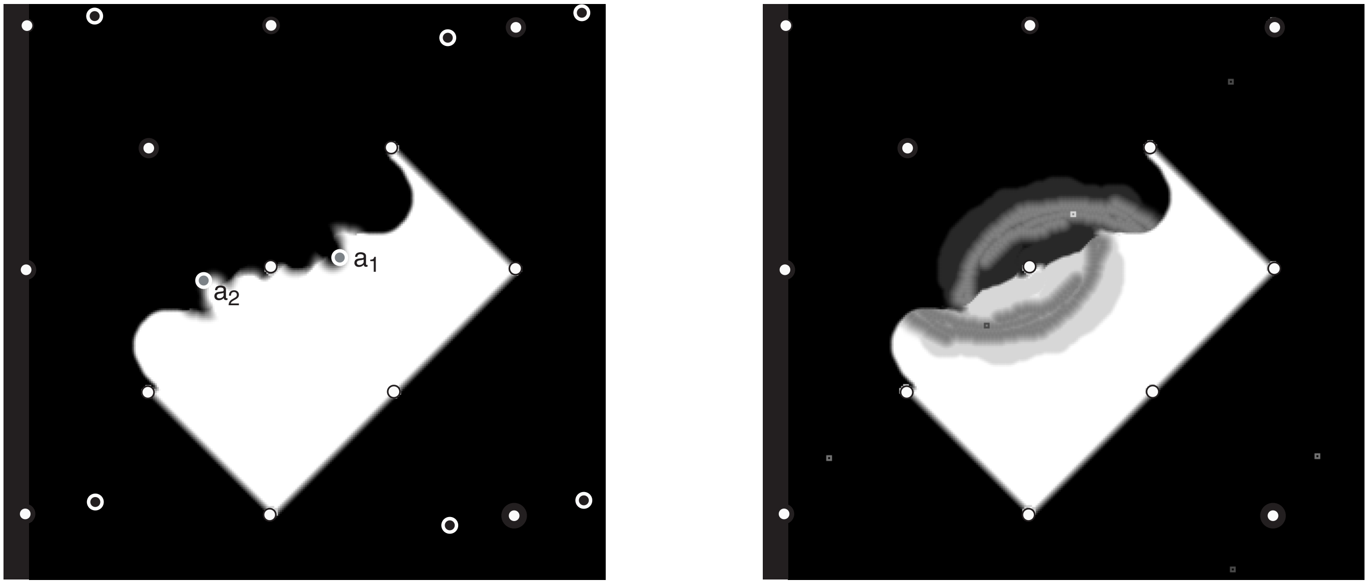

Figure 8: More contravarient distributions

$$D(a_2) = .05$$
$$D(a_3) = .18$$
$$D(a_4) = .04$$
$$D(a_5) = .19$$

There are actually infinitely more points a' in the orbit that lie outside the range of the picture; each of these has $D(a')$=0. Note then that

$$\Sigma_{a \in O} D(a) = 1$$

Now we color the corresponding points in the output image by

$$I_s(O) = \Sigma_{a \in O} I(a) \cdot D(a) = .54I(a_1) + .05I(a_2) + .18I(a_3) + .04I(a_4) + .19I(a_5) + 0 \ldots$$

Similarly, all points in all orbits in the output image are colored in a similar fashion, producing the image in the lower right of Figure 9.

4 The algorithm

First, one must construct and edit contravarient distributions. One may begin with any standard fundamental domain for a given symmetry and construct a distribution as we did in Figure 7. Then the user may use "contravarient" drawing tools (in contrast to the "equivarient" tools briefly discussed above): in essence, as one *adds* values to one point in an orbit, one must *subtract* values from all the other points in the orbit so that the sum of all the values in the orbit remains 1.[6].

For example, if the user wishes the value $D_{new}(a_1) = .68$, an increase of .14 from the value in Figure 9, a total of .14 must be subtracted from the other points in the orbit. One

[6]Hence "contravarient": the copies of the drawing tools act in precisely opposing ways

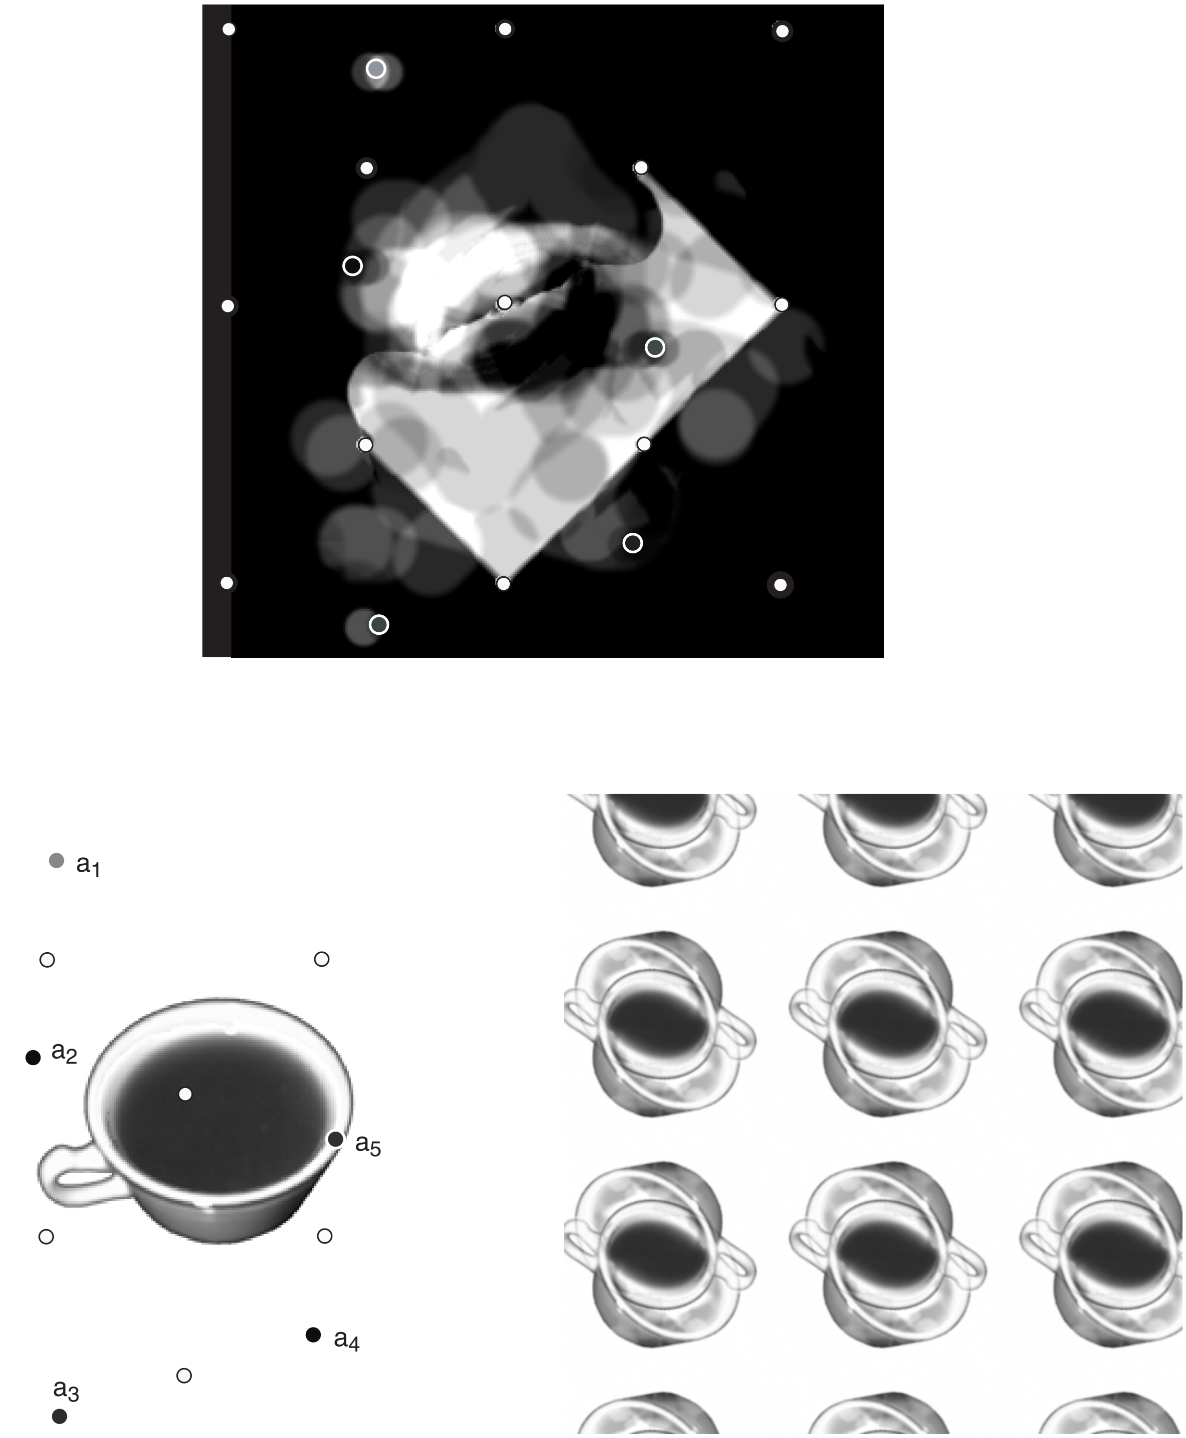

Figure 9: A more interesting example

must choose how this amount to be subtracted is spread across these points. This may be done in many ways, but we take the "fairness" approach and scale things proportionately. For example, $D(a_3) = .18$, which is $.18/.46$ of the total value among the points at which we will be subtracting. We will be subtracting a total of .14 so a_3's share is $.14\frac{.18}{.46}$ and $D_{new}(a_3) = D(a_3) - .14\frac{.18}{.46} \approx .13$

Thus, for each $i \neq 1$, we set

$$D_{new}(a_i) = D(a_i) - (D_{new}(a_1) - D(a_1)) \cdot \frac{D(a_i)}{1 - D(a_1)}$$

Once the contravaient distribution is chosen, one simply scans across the output image; at each point, one looks up the points in orbit of the corresponding points in the input image and the distribution, and then calculates the appropriate value for the point in the output image.

5 Extensions

In fact, one can define contravarient distributions in any setting for which one can define orbits for points. In particular, then, this algorithm applies to a much wider range of settings beyond creating regularly symmetric patterns (i.e. one does not need a discrete group action in the plane, sphere, or elsewhere). For example, the pattern in figure ?? is not symmetric in a classical sense; yet one can define orbits of points and so define a contravarient distribution. "Non-periodic" designs, such as that of Figure ??, have been widely studied and there are many techniques for their production. Many of these techniques produce what appear to be highly irregular patterns, yet have very clean underlying mathematical structure.

This algorithm, applied to such types of non-periodic designs, will be able to produce natural looking textures that are convincingly irregular. But the underlying mathematical regularity will allow the production of these textures to be very rapid.