IMPROVING AESTHETIC MEASURES FOR EVOLUTIONARY VASE DESIGN

by

KATE REED

A thesis submitted to the
University of Birmingham
for the degree of
MASTER OF RESEARCH

School of Computer Science
University of Birmingham
September 2012

Contents

Abstract

In order to avoid the expense of interactive evolution some researchers have begun using aesthetic measures as automatic fitness functions. This thesis explores the potential of one of the earliest aesthetic measures by George Birkhoff as a fitness function in vase design after suitable modifications.

Results suggest that Bezier curves provide a suitable representation for rotationally symmetric vases. Initial testing of vases of this form suggested that Birkhoff's original measure had some limited potential as a fitness function with some adaptations. However it had several flaws which could be exploited in the evolutionary process, most obviously it did not determine the orientation of the vase. These results have also revealed several further properties with a positive correlation with human–awarded scores.

Initial testing had suggested that Birkhoff's measure had little correlation with modern tastes but further investigation revealed that it could produce excellent results once other criteria such as orientation had been considered.

A suitable balance of these new measures along with Birkhoff's measure was found and vases evolved by the measure were assessed for their aesthetic potential by a group of volunteers. Although the initial designs suffered from lack of diversity some modifications led to a measure that enabled the evolution of a range of vases which were liked by many of the volunteers.

The final range of vases included many shapes similar to those developed by human designers. They show that at least within this test case it is possible to replicate the human

design process. Coupled with 3D printing techniques they allow automation of the whole process from conception to product. It is hoped that this demonstration of the theory will enable further work on other aesthetic products.

Acknowledgements

I wish to thank my supervisor Xin Yao and co–supervisor Shan He for their considerable help and support during my MRes course. It is a daunting task moving into a new subject at postgraduate level but they have kept me moving in the right direction, encouraging me at every step. I also wish to thank the final member of my RSMG group Peter Tino for his support as well as my other lecturers at Birmingham for their excellent introductions to this fascinating subject.

I am also grateful to my Mother for accommodating me for a year and my Dad for proof–reading my work. As well this direct support of my thesis I am also extremely grateful for all the support they have given me to enable me to get this far.

Finally I wish to thank the Royal Commission for the Exhibition of 1851 who sponsored my work. They saw the potential of my research, believing that it fitted with their charter of 'increasing the means of industrial education and extending the influence of science and art upon productive industry'. Their support has allowed me free rein to follow my research wherever it led, I only hope I can live up to the amazing achievements of my fellow recipients.

List of Figures

List of Tables

Chapter 1

Introduction

In recent years there has been a growing interest in using aesthetic measures as a fitness function in evolutionary design. Much of this work has concerned 2D applications such as art works, screen layouts and graphic design, although recent attempts have been made to extend these theories to 3D design. This work will continue to look at the application of aesthetic measures as a fitness function in the design of 3D objects, namely vases.

Vases were chosen as a subject as they have clear design constraints (they have to stand up, hold water etc.) but they are essentially a decorative object with aesthetic considerations.

The aim of this project was to develop an aesthetic measure of vases that is capable of being used a fitness function in an evolutionary algorithm. The basis for this measure was the work of George Birkhoff who developed an aesthetic measure for vases in his book *Aesthetic Measure*[6]. However the measure needed to be adapted to make it viable as a fitness function as it was never designed for the purpose and makes many assumptions. Evolutionary algorithms can take advantage of bad fitness functions to create unsuitable but high scoring designs. Also the original measure was poorly tested; a good measure should be more rigorously tested to ensure its preferences reflect the opinion of a large sample of people.

Fitness functions based on aesthetic measures have been used before but this project differs as it focuses on a 3D object; a vase which is also functional. The majority of the previous work had been 2D with only a few example of 3D work. This project will help to expand on this line of research by applying new measures as well as providing further testing for some previous measures.

A rotated Bezier curve was found to be an excellent representation which produced a wide range of interesting designs but also responded well to evolutionary functions. A simple embryology was used to ensure every design was viable. These settings allowed the design to be manipulated without compromising the practical considerations.

Although initial results were discouraging the final experiment demonstrated a good correlation for Birkhoff's measure with human–awarded scores. However the most critical factor in ensuring a good aesthetic design was correctly orienting the design. Two new measures, one based on Birkhoff's further requirements (as stated in his book) and one related to aesthetic measure work done on graphic design provide this orientation measure.

These measures allow a computer to evolve original, attractive vases unaided by a human. Although the design of vases is not in itself a major achievement it is hoped that by demonstrating that it is possible to analyse the aesthetics of this one object, it will enable similar work on other items[40]. This will allow an element of automation in the evolution of aesthetic objects, reducing the reliance on expensive interactive evolution. This could be combined with the emerging 3D printing technology to automatically design and make attractive consumer products.

Aesthetic Measures for Evolutionary Vase Design, a paper based on this work, was presented at the EvoMusArt 2013 conference and included in the proceedings.

Chapter 2

Literature Review

2.1 Evolutionary Approaches to 3D Form Design

2.1.1 Representations

The representation of a 3D object is critical to the form it will finally take. Certain forms can never be created in some representations but will arise naturally in others.

2.1.1.1 Design using Primitives

Both 2D and 3D forms can be created by building a complex form out of several simple shapes. These shapes are usually predefined and the overall structure is described as a list of these shapes, each with a shape descriptor (e.g. sphere), size, proportions and location.

2.1.1.1.1 GADES This program, developed by Bentley[4] is capable of developing a wide range of 3D items and has successfully designed items as diverse as tables, steps, prisms, streamlined vehicles and floor plans.

In his PhD thesis[3], Bentley discusses the importance of choosing a representation. In particular he highlights the importance of ensuring similar phenotypes have similar genotypes. He also argues that reducing the number of parameters improves the program as it reduces the design space and therefore simplifies the problem.

His representation allows for a vast range of designs. Some of the designs (such as the table) required straight lines and right angled corners and here GADES produces final

Figure 2.1: Example of a form evolved by Mutator[23]

designs capable of being constructed from scratch in only a few hours. However some items (such as the racing car) would benefit from having curves in the design. In these instances GADES can be regarded as more of a sketching tool; a human designer would still need to come in and finalise the design. Further information on this representation can be found in appendix A.1.1.

2.1.1.1.2 Mutator This 3D evolutionary art program was developed by Latham and Todd from an earlier art program called FormSynth[24]. Mutator evolves horn-like shapes such as the one in figure 2.1. These are formed of spheres and other primitives with genes controlling the number of horns and primitives as well as the curve and twist of the horns. The predefined layout of the spheres into 'horns' and 'ribs' gives the resulting artwork a very particular style.

Designs were chosen by the user and both mutation and breeding were used to generate the next generation. Within this design space the designer was given free rein with the only restrictions being placed on the ability of the computer to render the image.

Figure 2.2: Example of a form evolved by BlobTrees[19]

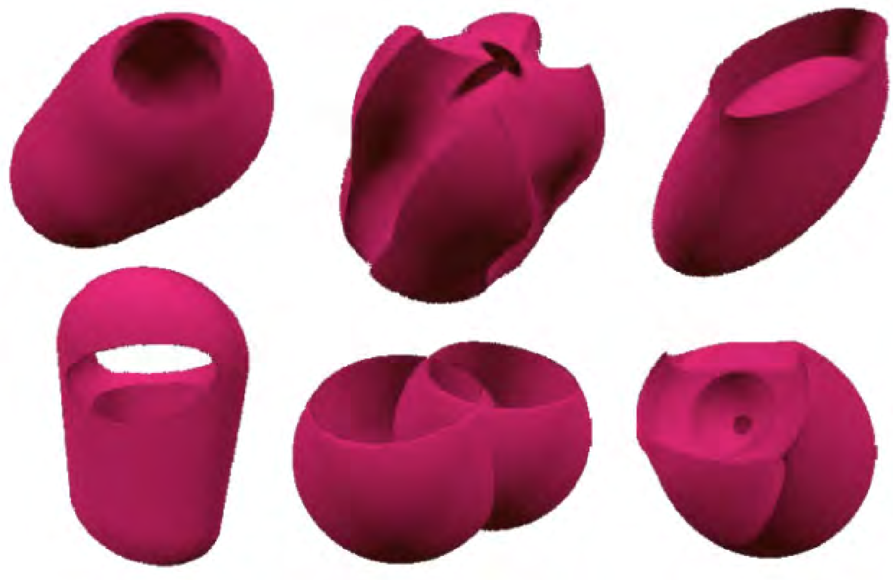

2.1.1.1.3 Implicit Surfaces A few researchers have combined the implicit surface method (see appendix A.1.2) with evolutionary algorithms; Jacob, Kwong and Wyvill developed a system that evolved implicit surface *BlobTrees* using genetic programming, to produce items with a distinctive 'blobby' form (figure 2.2). Each terminal node would be one of the primitive points and would have location coordinates and colour as well as a scaling factor along each axis. The interactions such as blend were described in the function nodes to build the complete form.

They themselves admit in [20] that their program is not intended to produce final designs but is just there to act as inspiration for a human designer. The limited functions they have used with their implicit surfaces gives all the designs a very specific look and like GADES this may only be appropriate for a limited range of designs.

Another implicit surface approach by Nishino, Takagi, Cho and Utsumiya[33] uses a far wider range of functions to allow a greater diversity of forms which include tapering, shearing, twisting and bending. They also use more flexible primitive shapes (superellipsoids and supertoroids) to build their forms. These are capable of creating blobby forms as well as straight lines and sharp angles. Here the genotype is a binary string with each gene represented by 8 bits. These genes correspond to parameters such as primitive type, position and scale and more complex ones such as tapers, shearing and bending. The interactions between primitives are also in the string and include field strength that determines how much the primitives blend together.

This program has the widest diversity of any of the methods looked at and in theory is capable of creating almost any shape. However with this vast design space to explore it is very user intensive. When the user started with a random object it took over 50 generations to produce an object that even remotely resembles the desired object. Starting with a prototype does reduce this to just 10 generations but then a user is required to create the prototype and must have a clear design goal.

2.1.1.1.4 Summary of Primitives Representation Design using primitives allows the designer to explore a vast design space and as a result it is capable of designing a wide selection of objects. This makes it ideal for multi–purpose design programs. One of the consequences of this increased design space is the increased time taken to achieve a suitable design and a prototype may need to be used to begin the process in a suitable area of the design space.

Stylistically the design is completely defined by the choice of primitives that are available. A representation such as the one used in GADES that uses only angular primitives is incapable of designing the blobby forms of the BlobTrees representation and vice–versa. If the program is being used to design a final product, rather than just a sketch design, then great care must be taken when choosing the primitives to allow the desirable forms to be evolved.

2.1.1.2 Cross-sections

3D forms are often described by 2D images. Cross–sections are used to give more detail of a 3D form than would be achievable with just elevations in many 3D design disciplines. Certain forms can be completely described by their cross–sections. Rotationally symmetrical objects are one such example.

2.1.1.2.1 Tuboid This is an evolutionary art project by artist couple Driessens and Verstappen[48]. Unlike other projects mentioned so far, Tuboid is a sculpture of the evolutionary process itself; every layer of the sculpture is one generation of the cross–section that the process evolves.

The cross–sections are formed of 32 spokes that mutate their length and rotation as the form grows. The only selection criteria is the requirement of the shape the spokes form to not intersect itself. This means that Tuboid is actually the evolution of a 2D shape with the 3D form being the representation of the process in time.

Figure 2.3: Example of forms evolved by Tuboid[10]

2.1.1.2.2 Engineering Problems The evolution of crushers by Hingston et al [17] is a good example of a 3D design optimisation problem, simplified to 2D while still being an accurate representation of the object. The 2D representation allowed the evolution of a design that had a simulated capacity of 140% of the original. Details of the model are in appendix A.1.3

A similar simplification to a 2D cross–section in an engineering problem was the evolution of flywheels by Eby et al [11]. Here the optimisation program looked for the solution with the maximum rotational energy given a maximum rotational velocity. It was also refined to reduce the stress in the flywheel as many of the highest energy candidates would have a dangerously high stress level in the wheel.

2.1.1.2.3 Shells This representation by Prusinkiewicz and Fowler is not used as the basis of an evolutionary algorithm but as part of a study into simulating the patterns and forms of seashells[30]. However it is a simple representation that could potentially create beautiful forms when combined with an evolutionary algorithm. To create the seashell form they use a generating curve that is rotated around a helico–spiral and steadily reduced in size. The style of the shell is determined by both the shape of the generating curve and proportions of the helico–spiral. The generating curve itself is represented by a Bezier curve that is described by a list of coordinates.

2.1.1.2.4 Summary of Cross–Sections Cross–sections allow the 3D form to be simplified to a 2D problem. With this simplification the design space is reduced to forms such as rotational objects. Stacked layers such as the ones seen in Tuboid do in theory have the ability to describe any form, but the number of layers required and the complexity of the required shape in each layer prohibits the use of cross–sections to explore this extended design space except in special cases such as Tuboid, where only a single 2D shape is changed at each evolutionary step.

The remaining design space is explored using the 2D shapes in the cross–section. As 2D shapes are usually easier to describe than 3D this shape can be very complex, leading to the organic forms seen in Tuboid and the shells. But equally the shape could have sharp corners and straight lines as in the crusher design. Although this representation does limit the design space it is ideal for engineering design and other areas where the limited design space may be the only possible design space.

2.1.2 Algorithms

2.1.2.1 Genetic Programming

Genetic programming has been widely used in 2D art. Its ability to evolve complex functions allows works such as 'In the Beginning' by Rooke[37]. This piece has 427 nodes in its tree and is very unlikely to have ever been manually designed due to its complexity. However attractive images of this complexity are hard to come by beginning with entirely random variables and for this reason Rooke prefers to seed his new images with previous successes from a compiled library.

A certain amount of strategy appears to be required in the evolutionary process. Some stages in the evolution of 'In the Beginning' are shown on the artist's website[36]. The breeding of a great–great–grandparent (shown in figure 2.4) is particularly interesting as one of its parents is just plain blue, not in itself interesting but provides a dramatic change to the colour of the child while leaving the form mostly unchanged from its second parent. This demonstrates one of the features of genetic programming; a unfit parent may still have valuable traits that can be passed down. This makes it harder to select parents but encourages greater genetic diversity. If you are picking by hand and have the necessary experience to see the potential of an unfit parent, then this trait can lead to very successful designs.

Another demonstration of this trait is shown by BlobTrees (see subsection 2.1.1.1.3). During the evolution of chairs it was found that two chairs that were unfit (two legs were

Figure 2.4: Two Parents (left and centre) and Child (right) in Rooke's Art[36]

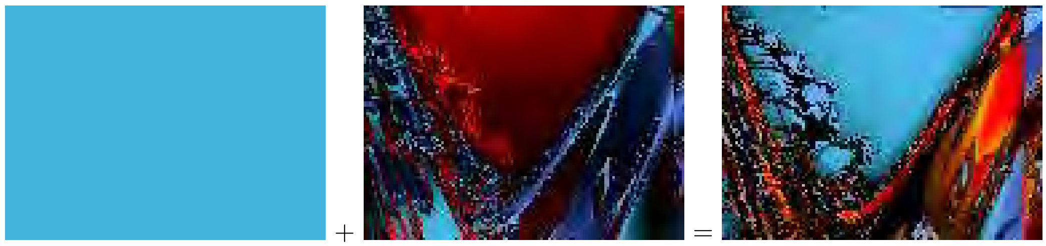

not attached to the chair) could produce an offspring that not only had all its legs attached but was also interesting and attractive.

However the flexibility of BlobTrees and their tree genotypes are not without problems. Bentley[3] deliberately rejected the use of constructive solid geometry with multiple types of primitives due to the dramatic changes that could be caused by a small mutation. He gives the example of a mutation from a cylinder to a plane to demonstrate this.

2.1.2.2 Genetic Algorithms

Genetic programming has found many uses in exploring large design spaces with abstract forms. However genetic algorithms are more suitable when greater control of the form is required. For example a 2D image may become unattractive if a large step is taken within the design space, but it will still be an image. If a table or any other object with a clearly defined purpose makes a dramatic change then it may no longer function as required. These objects benefit from the stability that a genetic algorithm offers.

2.1.2.3 Summary

Both algorithms considered here have been used in art and design projects to great effect. However they are each suited to a different task; genetic programming is more suited to abstract forms that will not be destabilised by dramatic changes, while genetic algorithms are more suited to items with a purpose and allow a more methodical exploration of the design space.

2.1.3 Variation Operators

2.1.3.1 Recombination

Recombination or reproduction is one of the most common ways of producing the next generation, although not all algorithms use it.

Mutator[24] (see section 2.1.1.1.2) is a good example of different recombination options, as in Mutator the breeding of two or more parents could be done in a number of ways. These include common methods such as random selection, weighted average and splicing.

A particularly interesting method that was used is dominant and recessive genes. This method determines if a gene is recessive or dominant by studying how fast the gene has been changing. A gene that has remained relatively unchanged for several generations must have been selected by the user at every step and therefore is probably desirable; it is set to be dominant. A gene that has changed several times in the last few steps is likely to be undesirable and is set to be recessive.

The wide range of options available to the user allowed a much more varied evolutionary experience but would increase the time taken to move onto the next generation. However when Mutator was being developed the time taken to run each generation dwarfed the time taken to set up the following step and therefore it was worth spending extra time on this variety of options.

One of the potential problems of recombination is if very different 'species' have been evolved with equally good fitness there is the possibility that the offspring will have very poor fitness. This is particularly a problem with functional designs where one part of the design may have a large impact on another, making the fitness dependent on the whole form.

2.1.3.2 Mutation

Even with a large initial population, recombination is limited to the design space covered by the initial genes. To introduce further variation and allow the design to evolve significantly from the first designs mutation is required. This could either be changing a value by adding or subtracting another value within a set range or changing an abstract representation, such as a primitive shape or function, to another. However mutation must be within a strict range otherwise children may bear no resemblance to their parents, particularly if it is a primitive that is to be changed.

2.1.3.3 Replacing the Previous Generation

Some algorithms such as GADES[4] (section 2.1.1.1.1) allow successful designs from earlier generations to survive in the population for many generations. This ensures that good genes have a chance to spread through the population before that parent dies out. Others such as Picbreeder [34] will evolve an entirely new population each time. Picbreeder is an online 2D evolutionary art tool that uses the evolution of neural networks to create images. In each generation the user selects one or more designs and then selects ‘evolve’. The new generation is evolved from the selection only.

How the next generation replaces the previous therefore depends on the way that fitnesses are assessed. GADES gave a numerical fitness value based on a series of physical properties that were measured in the computer while Picbreeder has only two fitness options, selected or rejected. The rejected options have no opportunity to pass on their genes and are dropped from the population immediately. In GADES and other programs with a spectrum of fitness values it is possible (although unlikely) for even the worst designs to pass on their genes and if the new generation is poor they can also survive into the next generation.

2.1.3.4 Summary

Both recombination and mutation have useful traits for the design of vases. As vases have a simple concept and an easily defined purpose, recombination could be used without the design becoming unsuitable. It is important to ensure good inheritance as the child must have traits of the parents, particularly aesthetic traits. The ability of mutation to introduce new genetic material would also be desirable.

2.2 Fitness Evaluation

2.2.1 Interactive Evolution

Interactive aesthetic design is a very simple idea and has been used successfully in a wide range of applications in both 2D and 3D. 2D examples include evolutionary art such as the work of Rooke[36] and Picbreeder[34] as well as more design–based work such as batik design[25] where a specific style and use is required. 3D examples include Mutator[24] and the implicit surface methods.

In the implicit surface method by Nishino et al[33] (see section 2.1.1.1.3) the selection process relies only on human evaluation. This allows the user to evolve almost anything from the randomised objects although beginning with a prototype form increased the speed at which a suitable design was reached. The examples given in the paper demonstrate the program's flexibility as they evolved both organic objects (vegetables) and man–made objects. This highlights one of the benefits of interactive evolutionary algorithms: there are no physical fitness functions that a computer could test in the evolution of a vegetable but most users would have a good mental image of what they were trying to achieve.

One of the disadvantages of interactive design is user fatigue. This is due to the number of generations that may be required to find the ideal design and the potential problem of the design becoming stagnated, causing the user to lose interest. The batik design system[25] prevents stagnation by using an *out-breeding mechanism*. This maintains two populations; the main population evolved using interactive selection and a second population that is evolved to be as different as possible. This increases the genetic diversity and allows the user to bring in a new individual if the main population begins to stagnate.

2.2.2 Automatic/semi-automatic Fitness Evaluation

2.2.2.1 Automatic Measures

Aesthetic measures attempt to assign an aesthetic fitness based on some predetermined rules. One of the earliest and most famous is Birkhoff's measure (see section 2.3.1):

$$M = \frac{O}{C}. \tag{2.1}$$

Where O is the order of the object and C is the complexity. This has formed the basis of many subsequent measures that have been used as fitness functions in the evolution of aesthetic objects.

In 2D evolutionary art they have already been used by many practitioners with varying degrees of success. Den Heijer and Eiben[16] compared four different measures including the ratio of image complexity and processing complexity by Machado and Cardoso[26]. This ratio was based on Birkhoff's order/complexity ratio as the amount that an image can be compressed by indicates the amount of patterns present in the image or its order and the full size of the image could be considered as its complexity. The other measures were the normal distribution of colours by Ross and Ralph[38], the fractal dimension by Spehar et al[44] and a weighted sum of the previous measures.

They conclude that a particular measure will evolve a certain style of image. They also observe that the images evolved using any of the other measures rate poorly using the fractal measure. They observe that; 'many people like fractal images, but in reality, not many images actually have fractal properties'. This highlights one of the dangers of using aesthetic measures - they may describe a property that is perceived as beautiful but it is unlikely to be the only such property and designs may exist that are attractive only in a cultural context.

In the evolution of 3D forms they have been used in a range of tasks including computer generated terrains[47]. Here a Kolmogorov complexity measure was used to analyse images of the evolved terrains. The Kolmogorov measure also uses the ratio between the amount of space saved by compressing the image and the original image size similar to the first of the measures used in Den Heijer and Eiben's work.

True 3D aesthetic measure have been less widely used, one recent example is by Bergen[5]. This measure uses several properties of 3D graphics objects to calculate a fitness measure. These include the angles between adjacent face normals as well as the differences in these face areas. These different properties are then processed in several ways to find different measures, including comparing the distribution of the angles with the normal distribution with the same mean and standard deviation as well as similar comparisons of the logarithm of the results and the entropy. He also studies the symmetry of the object. This is found by dividing the volume the object is contained within into cubes and then comparing the number of vertices within opposite cubes along each axis.

A range of good values was found by studying a set of existing graphics objects. These were from an online database and although they were not necessarily created to be beautiful they were considered by Bergen to cover a range of good aesthetic traits. These allowed him to find 'sweet spots' in the aesthetic measure values. These measures were then used to evolve a range of items including organic forms and city layouts. The measures were tested by evolving a series of forms that were then judged in pairs by volunteers. The higher scoring individual using the deviation from normal fitness function were chosen 95% of the time, supporting the measure's validity. However Bergen concludes that the area is in its early stages and requires further research.

2.2.2.2 Learning from Users

Some researchers have attempted to develop automatic measures by learning from interactive processes. One of these by Ekárt, Sharma and Chalakov[14] for abstract 2D art tested several previous measures to determine if any were reflected in an interactive evolutionary process. The measures that they used were Machado and Cardoso's measure[26] and Ross and Ralph's measure[38] as used by Den Heijer and Eiben[16] above. The other measures are based on physical entropy and Shannon entropy.

They created an interactive evolutionary process and recorded the value of each measure for each iteration, testing the processes of both computer scientists and graphic designers. Their results suggest that there was little evidence of continuous increase of a single measure during the interactive evolutionary process. However it was noted that the pareto front of two of the measures did improve suggesting that the user does not select for a single trait but for a combination of the traits.

Other researchers such as Gu, Tang and Frazer[18] develop neural networks to mimic human preferences. This was done by developing a fitness landscape from a series of choices by a user. Unlike the other automatic fitness functions, where the fitness is derived from a set of features of the phenotype, this fitness landscape is based on genotype variables.

2.3 Aesthetic Measures

2.3.1 Birkhoff's Original Measure

In his book *Aesthetic Measure*[6], first published in 1933, the American mathematician George Birkhoff described his formal definition of a measure of beauty.

$$Measure = \frac{Order}{Complexity} \tag{2.2}$$

He then applied this to a number of areas, including art, polygonal forms, music, poetry and vase designs. Although his overall measure was applied to all of the areas the definitions of 'Order' and 'Complexity' varied between the tasks.

He applied his polygonal form measure (appendix A.2.1) to 90 polygons and verified his results by seeking the opinion of members of two classes of students. He does not go into detail as to the extent of this verification but states 'The results so obtained were found to be in substantial agreement with the arrangement obtained by the formula'

The measure that is obviously most applicable is Birkhoff's measure of vases. Birkhoff had some misgivings about applying his measure to vases. He admitted that 'I should expect that any vase form ... would appeal to some intelligent person, if only by virtue of its novelty'.

All the vases studied by Birkhoff were rotationally symmetric and therefore he only looked at the silhouette of the vase. As they do not have well defined edges and angles the measure is based on a set of 'characteristic points'. The order is defined by the ratios of 1:1

and 1:2 between them as well as tangent properties at these points. The complexity is the number of points. These are all defined in detail in appendix A.2.2. However, this formal measure was not the only consideration that Birkhoff discussed; he also noted the value of the 'Regularity of Contour' and 'Conventional and Utilitarian Requirements'.

For regularity of contour Birkhoff studied the properties of the curvature of the curve. Given the points and tangents of a vase he states that the curvature between these points should vary continuously and should oscillate as little as possible and it should not oscillate more than once between points of inflection. He also requires the rate of change of curvature to be as small as possible. He notes that the application of these requirements would be difficult and this perhaps explains why he does not formally include them in his aesthetic measure.

Finally Birkhoff looked at the conventional and utilitarian requirements, or the features that need to be present for the vase to function as such. These are listed in full in appendix A.2.2. They include recommended sizes and locations for maximum and minimum dimensions and rules regarding angles.

2.3.2 Testing the Measure

Birkhoff both developed and tested the vases using the examples from China. He used these as Chinese ceramics were highly regarded during that period. They also did not routinely embellish the vase using handles or other features, allowing the form to be studied on its own.

He applied the measure to eight vases and although he does not attempt to compare their measured aesthetics to any human comparison he states that:

> 'In my opinion the large number of these elements of order H, V, HV and T which are present cannot be explained as accidental, because vases taken at random contain very few such relationships. It seems to me probable that, more or less intuitively, forms of vases were gradually adopted which involved the types of relation recognized by our theory'.

He further tested his theory by designing 3 vases of his own that all scored very highly using his measure. Although this writer would agree that they are aesthetically pleasing he did not attempt to validate these designs.

In his conclusion he admits that this measure is 'more uncertain' than his other measures but does still have some value. He adds that it only has value if the other two requirements, 'regularity of contour' and the 'utilitarian and conventional' list are satisfied.

2.3.3 Analysing and Improving Birkhoff's Measure

2.3.3.1 Psychological Analysis of Birkhoff's Measures

2.3.3.1.1 Experimental Results Since the publication of his work there have been several attempts to test Birkhoff's claims. Most of these have focussed on his work with geometric shapes.

Several tests were carried out in the decade after the publication of *Aesthetic Measure*. Among these was a paper by Davis in 1936[9]. Davis chose 10 of Birkhoff's shapes and asked two groups of students put them in their preferred order. These two groups were 162 psychology students and 55 art students. The correlations between their average ordering and Birkhoff's measure was poor with a correlation of only 0.11 for the psychology students and 0.05 for the art students. The two groups had a correlation of 0.9 with each other.

The next tests were carried out by Beebe–Center and Pratt in 1937[2] and Harsh, Beebe–Center and Beebe–Center[15] in 1939. The first of these used results from another researcher, Smith which used 15 polygons and 6 students. This time much better correlations were found, suggesting that 'Professor Birkhoff's formula for polygons has a considerable degree of validity'. The second test involved new data collected by the authors, this time 26 shapes were used and 30 psychology students. The range of correlations between the measure and the students ranged from -0.47 to 0.92 with an average of 0.34. This was similar to the range of the correlations between the students, suggesting that the 'formula predicts an aesthetic order about as typical of the group as the average observer'.

In the first of his papers Beebe–Center and his co–authors also studied a set of data that had been gathered on vases by Schnittkind. This time there were 3 groups, 5 psychology students, 8 'naive observers' and seven fine art students and instructors. Here Birkhoff's measure correlated at least as well as the the other individuals in the first two groups (0.23 in a range 0.20–0.34 for psychology students and 0.16 in a range -0.21–0.32), but it correlated very poorly with the art group (0.09 in a range 0.02–0.51), suggesting that it reflects the opinions of untrained individuals but not those with more expert opinions.

A fourth set of geometrical shape results by Wilson[49] also did not support the measure and he suggests that weighting for features that are found to be desirable may be a better approach to a measure for geometric shapes.

In conclusion the validity of Birkhoff's geometric shape measure seems to vary depending on the group sampled and with the method used. This suggests that although it may be applicable under certain circumstances, as a general measure in is not sufficient.

2.3.3.1.2 Problems with the Methods One of the problems with validating measures is the inconsistency of the human subjects. Ecker et al[13] describe many potential inconsistencies that have been noticed over the years. These include 'Circular Triads' (noted by Pratt[35]) where 'if a subject states his preference for A over B and of B over C, then presuming preference to be logically transitive, the same subject ought to declare a preference for A over C'. However in practice this is not always the case; further experimentation by Pratt suggests that these triads exist but are rarer than chance would allow, suggesting a degree of logical preference.

The subject's training also has an impact on their choice of the 'best' design. Child[8] carried out a series of experiments to compare the preferences of experts with the preferences of school children. After teaching the children the reasons why types of work were preferred the children could guess the 'best' work. However their preferences did not necessarily change. A similar phenomenon was observed by Ecker[12] where an expert would identify the best painting of two but still express a personal preference for the other.

This unqualified preference could be for any number of reasons. Barnhardt[1] carried out a further experiment with Birkhoff's shapes that aimed to identify the reasons for a particular preference. Alongside preferences based on measurable criteria such as complexity and symmetry some participants mentioned associations that biased their preferences. These included resemblance to familiar objects such as stars, books and hats. One participant was particularly attached to a shape that reminded her of a diamond ring due to the fact she had recently got engaged.

Some of these biases can be removed from the sample set, for example in the second paper Beebe–Center et al attempted to limit the impact by removing shapes with obvious connotations such as stars and crosses. But less obvious and more personal relations such as the diamond ring are difficult to remove. The search for an overall aesthetic will also be affected by changing influences. The inclusion of a swastika in the shapes is an obvious example; during Birkhoff's research in the early 1930s it would have had little meaning to most people but by 1939, when Beebe–Center et al were working it had very strong connotations and was therefore removed from the sample.

Given these difficulties it is unsurprising that the many tests of Birkhoff's shape measure have been inconclusive. These problems will also be apparent in any subsequent research although association problems may be less applicable to a vase.

2.3.3.2 Developing Birkhoff's Vase Measure

Staudek[45] developed Birkhoff's method further. He morphed one design into another, suggesting that the aesthetic measure should vary continuously over this process. However the aesthetic measure was not continuous as the equations rely on perfect ratios. To amend this he added a tolerance margin so that any ratio falling within the margin would be counted as a ratio (or suitable tangent property) but the value of this would be reduced proportionally with the difference between the actual ratio and the desired ratio. His method is discussed in more detail in appendix B.2.2.1.

This produced a graph that was much more continuous in appearance than the original so that two very similar vases now had very similar scores. In some ways this is faithful to the original measures which were done by eye on real vases which had some degree of inaccuracy which the eye could not perceive. Modern methods using computers can work with a much higher degree of accuracy and in a randomly generated form it is likely that no ratios will be found due to this unnecessary level of precision. Adding a tolerance margin allows the computer to recognise the ratios that the eyes can see. One noticeable difference with his method however is his lack of caps and independence requirements for the order components.

2.3.4 Other Measures

Since Birkhoff there have been other attempts to develop aesthetic measures for other applications. One of the areas that has been well documented is the layout of computer screens. This is an area of great interest as improving the aesthetics of a layout can improve its usability[32]. They are of interest here as they use properties that are easily transferred to vase design.

Ngo and Byrne[32] define five different measures to describe the aesthetics of a computer screen. These are: balance, equilibrium, symmetry, sequence and order and complexity, these are described in detail in appendix A.3. The first 4 order measures are then summed and divided by the complexity; the number of features. This is once again based on Birkhoff's $M = \frac{O}{C}$.

Their measure was tested by applying it to seven mock layouts. These were then scored by six experienced designers and the average results compared. The scoring was done in the range 0–3. This was equivalent to the measure as it was shown that this would also assign scores between 0–3. The designers often disagreed, with one design receiving a wide range of scores between 0–2. However the average scores had a good correlation with the aesthetic measure and the authors determined that this correlation was significant despite the low number of layouts and volunteers. A further test was carried out by redesigning some existing screens and receiving positive feedback on the changes.

These types of measures, that concern themselves with the balance of a design, are not currently included in Birkhoff's measure, however they may provide new areas of investigation to improve upon it.

2.4 Gaps in the Literature

Many of the new measures developed since Birkhoff have used his fundamental idea of order/complexity but all have developed new definitions of order and complexity for their specific tasks such as screen layout and 2D art. Some of these have then been combined with evolutionary algorithms to allow automatic evolution. On the whole the tested measures have shown positive if rather limited results.

The use of 3D aesthetic measures in evolution is small and limited to a few abstract designs evolved by a limited number of measures. These and other 3D measures should be extended to a wider range of familiar objects to further test their validity.

Little effort has been made to test the claims made by Birkhoff regarding his vase measure or to compare it with modern tastes. As some of the work testing his theories on geometric shapes has been disputed it is concerning that many of the subsequent measures are based on his work. Further study of the vase application will provide further evidence of the validity (or lack of validity) of his work.

Chapter 3

Preliminary Studies

A series of preliminary studies was carried out with the aim of being able to evolve vases using an aesthetic measure based on Birkhoff's. A suitable representation and preliminary tests of the measures were required.

3.1 Representation

The decision was made to reduce the design space to rotationally symmetric vases. This greatly reduced the number of non–viable designs while still allowing for a wide range of styles. It was consistent with Birkhoff's measure and also allowed the representation to be reduced to just a 2D curve.

3.1.1 Parametric Curves

The first method used to create the 2D curve was parametric equations. Although they were shown to be capable of achieving a wide selection of forms they did not perform well when applying genetic operators, therefore a new representation was sought. Information on these trials can be found in appendix B.1.

During this investigation an embryology was introduced to prevent the problems caused by the 2D curve intersecting the axis of rotation. The minimum radius value was deducted from the entire curve and then a positive value (here 1) was added. This ensured the vase was feasible.

3.1.2 Bezier Curves

The second representation of a 2D curve that was tested was Bezier curves as used in the shell forms of Prusinkiewicz and Fowler[30]. These are represented by a series of (x, z) coordinates. The tests are described here in brief, further details are in appendix B.1.2.

The effect of varying the number of coordinates was considered and once again forms with both fixed z and variable z were studied to assess the potential design space. When fixed the z coordinates was set to 0, 2, 4 etc, increasing by 2 each time. It was found that 2 or 3 pairs of coordinates were too simple to produce interesting variations but higher numbers all gave good results. The number chosen for further tests was 5 as this gave a good variation in designs but was not too large. Examples are shown in appendix B.1.2.1. As before the curve was adjusted to give a positive minimum value to prevent the vase from crossing the rotational axis.

By fixing the z coordinates it was found that all the random designs would be viable. Varying z randomly produced many non–viable designs where the curve would intersect itself. This was prevented by setting the first coordinate as 0 and then adding a random number to the previous coordinate to generate each new coordinate. Gaussian random numbers were used with standard deviation 1 and variable mean. A positive mean ensured most of the added numbers were positive but negative values were also possible. This also gave a good range of viable vase designs. Examples are shown in appendix B.1.2.2.

The effect of genetic algorithms on the Bezier curves was studied. They were chosen as they are ideal for manipulating a matrix of data such as the set of coordinates used for the Bezier curves. As the coordinates have potentially continuous values, floating point numbers were used as opposed to binary numbers or integers.

Two different forms of recombination were studied; these were splicing and averaging. Both were found to produce designs that bore a good resemblance to both of their parents as seen in examples in appendix B.1.2.3.

Figure 3.1: Parent and Offspring by Mutating Bezier Curves

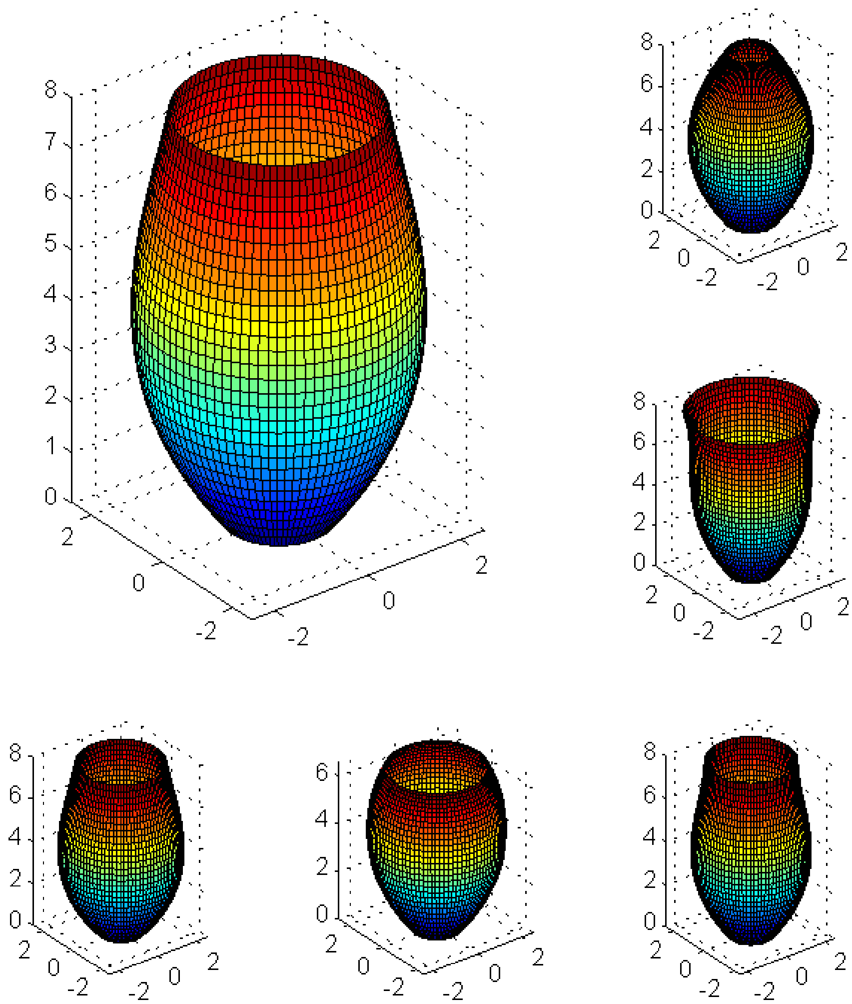

To test the effect of mutation a single point, either in the x or z coordinates was selected and had a Gaussian random number added with mean 0 and standard deviation 1. Some examples are given in figure 3.1. All of the offspring resemble their parent and appear to be simple changes to the form such as adjusting the width or the addition of new detail to the rim.

3.1.3 Conclusion

The use of the 2D rotational representation appears to work well once a suitable embryology is used to prevent the curve from intersecting the rotational axis. Given a suitable curve this then creates a viable vase with a stable base and open top. A suitable curve is any curve that does not intersect itself and increases in the z direction overall.

The parametric equations with fixed z would occasionally produce non–viable vases but with a suitable selection of primitive functions the majority were viable and there was a good variation of designs. The variable z version produced less consistent results and would require constraints to the primitive functions if it were to be used.

Every vase found using Bezier curves with fixed z during these initial tests produced a viable vase. Changing the number of coordinates did not affect the viability of the vase and attractive shapes were formed by all. Lower numbers would have limited design space but would be suitable if the number of coordinates were to change during the evolutionary process. Allowing z to vary also produced interesting forms but the z coordinates would require a greater level of control to ensure a viable vase.

The genetic programming of the parametric equations was both difficult to implement and gave unsatisfactory results. This may be a problem with the algorithm as the significant changes seen have been reported before by Bentley. However other researchers have been able to use this algorithm to their advantage and produce good results. This suggests that the problem here may be due to the implementation as Matlab is not suited to variable length trees.

The other problem could be the representation; in order to get a suitable inheritance of features, similar phenotypes must have similar genotypes to allow a smooth transition between forms. The primitive functions that the parametric equation are built from all produce distinctly different curves and therefore changing a function through mutation or crossover could have a huge effect on the form.

By contrast the genetic algorithm used on the Bezier curves performed as expected. As the numbers in the genotype represented coordinates, similar numbers produced similar vases. Mutation would change a coordinate by a small amount, deforming the original shape into a subtly different form. Splicing would take the first coordinates from one design and the last coordinates from another and produced a vase that looked as you would expect, with the base of one and the top of the other. Averaging also worked as expected with the child having the average form of its parents.

3.2 Initial Aesthetic Measure Assessment

3.2.1 Collection of Data

To test the aesthetic measure a random selection of Bezier curve vases was generated using fixed z with maximum height 4 and minimum radius 0.5. These were then reviewed by a group of volunteers. The sample size was small (30 designs and 8 reviewers) but it was hoped that it would give an idea of the effectiveness of the measure. They were asked to give each design a mark out of 10 and assess the design in both orientations. The designs and results are shown in appendix B.2.1. The information regarding the difference in regular and inverted scores was gathered as the lack of orientation requirements was one of the noticeable flaws in Birkhoff's measure. It was hoped that this information would allow the development of an orientation measure.

The scores were analysed to see if there were noticeable aesthetic trends that could be mirrored by an aesthetic measure. One problem that was initially noticed was the range of values used; some users would award scores from 1–9 while others would use a smaller range. It was not clear whether those with larger range had stronger feelings regarding the designs or a lower threshold for awarding these scores. However it was hoped that the mean of the scores would reduce this problem.

The data was treated as 60 separate vases, 30 of each orientation. The Pearson correlation of the mean and individual scores were found. These were in the range 0.7031–0.8431 which suggest that the average scores are a reasonable representation of the data trends for all of the individuals.

Uncapped, Dependent	-0.2795
Capped, Dependent	-0.1815
Uncapped, Independent	-0.1808
Capped, Independent	-0.1240

Table 3.1: Spearman Correlation Coefficients for Aesthetic Measure

The choice of orientation was also investigated. For every vase the number of people who disagreed with the average orientation was found. This ranged between 0–3 with an average of just 0.9, showing that the majority of the volunteers agree with the orientation given by the mean. Therefore it is believed that the mean accurately reflects the overall opinions of the group. These average values can then be used to analyse potential measures without requiring the individual assessment of every volunteer's scores.

3.2.2 Original Measure

The implementation of Birkhoff's method was based on the algorithm used by Staudek[45] (appendix B.2.2.1) with a few changes, including the addition of a measure of independence and capped values.

The impact of the independence and the capping was investigated. Figure 3.2 shows the four combinations of results for capped/uncapped and dependent/independent. The aesthetic measures were plotted against the average human results for the preferred orientation as the original measure assumed the vases were correctly orientated. The sizes of the points indicate the complexity of the vases with a larger point showing a higher complexity.

The correlation coefficient was also calculated for each one (table 3.1), this revealed very poor correlations for all four options as they were all very low and negative, where a positive was expected. Spearman correlation was used as it was not known if any correlation would be linear.

Closer analysis shows that the vases that are scoring too high are the more complex vases while it is the simple vases that are scoring low. This suggests that the division by the complexity is not having sufficient impact on the measure. Many of the highest scoring vases

Figure 3.2: Variations on Birkhoff's Aesthetic Measure

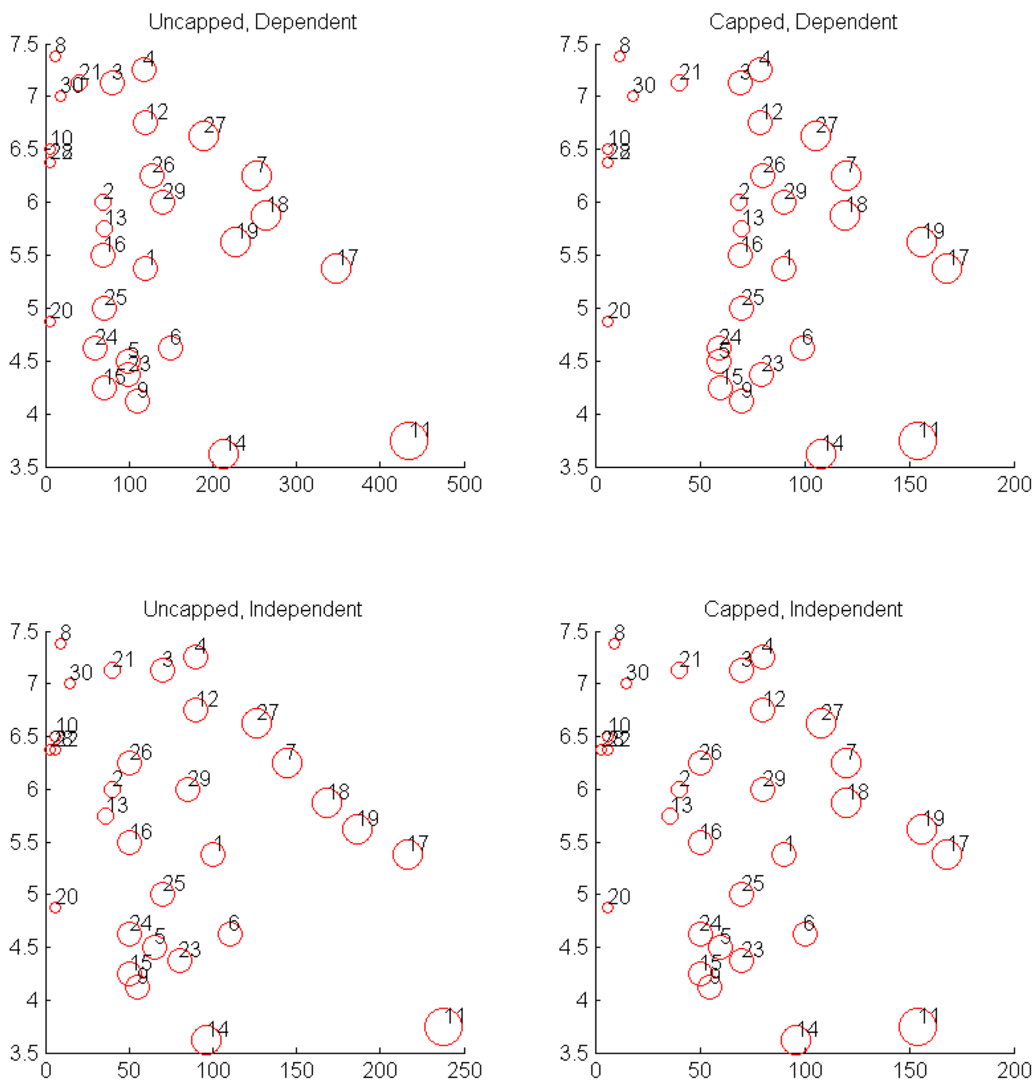

Uncapped, Dependent	0.3628
Capped, Dependent	0.3811
Uncapped, Independent	0.4114
Capped, Independent	0.4084

Table 3.2: Spearman Correlation Coefficients for New Complexity Value

have a low complexity but due to this low number of points they are unable to generate order scores that would allow them to have a high measure.

To increase the impact of a high complexity a new function for complexity was used:

$$C = e^{N-3} \tag{3.1}$$

Where N is the number of significant points on one side of the vase. As the range of N in these examples is 3–7 this gives a range for C between 1–54.6 instead of the original 6–14. Applying this new function we get the new plots shown in figure 3.3 and the new correlation values given in table 3.2.

By suppressing the complex vases the correlations have improved, in particular the independent scores now show a good positive correlation for both and all the confidences are over 97%. Looking at the plots we see that the measure has not awarded any high scores to the vases with a low human score. This is a trait that will be very useful in a fitness function to avoid the evolution of bad designs.

This new complexity function demonstrates that a modified version of Birkhoff's measure with both the caps and independence requirement could be useful in the development of a fitness function. However the correlations are still lower than would be desired and these results presume that the vase is correctly orientated. Further traits will be studied to improve the correlation of the aesthetic measure and to introduce a new orientation measure to ensure the vases are correctly orientated.

To explore the effect of different ratios a second set of ratios $1:1$ and $1:\phi$, where ϕ is the golden ratio were also tested. These produced the results shown in table 3.3, all with confidence over 90%.

Figure 3.3: Variations on Birkhoff's Aesthetic Measure with modified order function

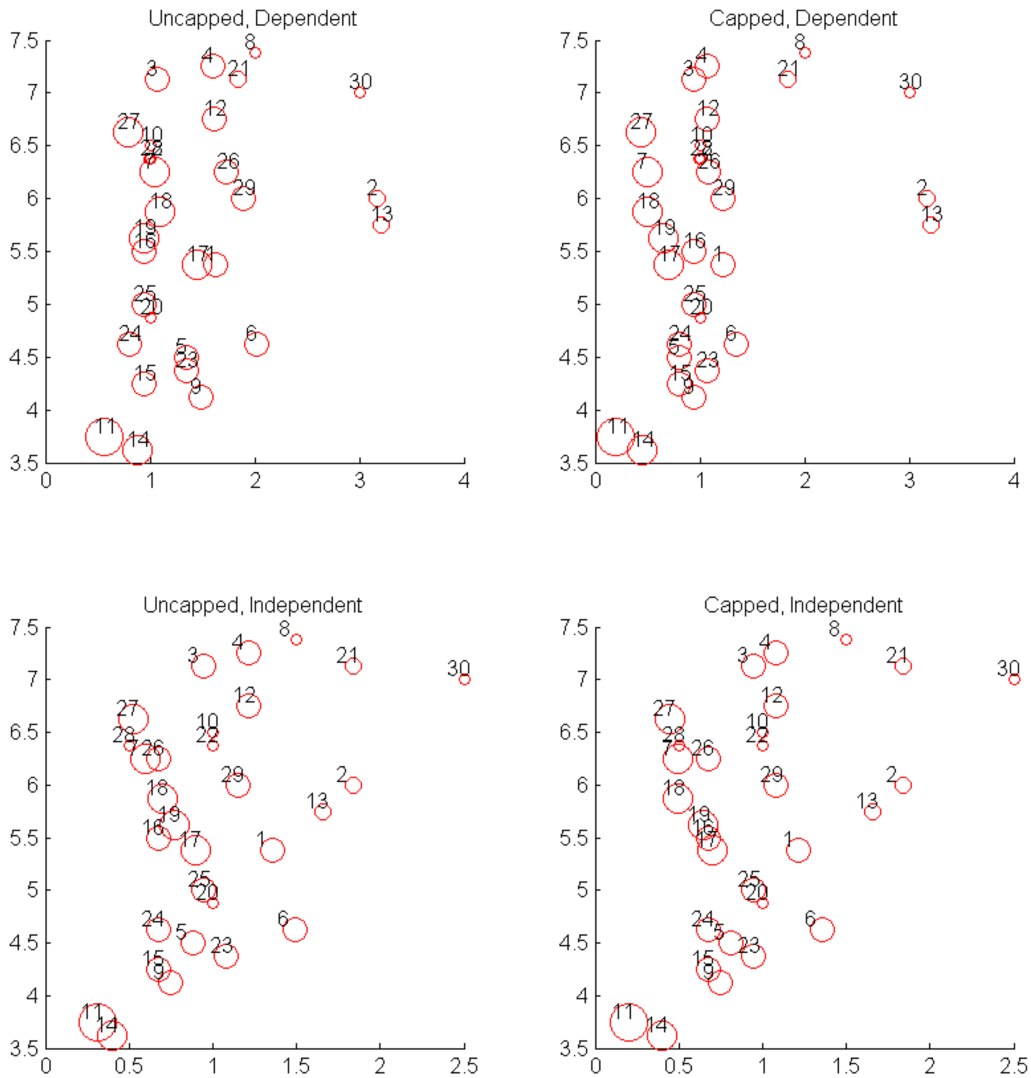

Uncapped, Dependent	0.3139
Capped, Dependent	0.3794
Uncapped, Independent	0.2723
Capped, Independent	0.2657

Table 3.3: Spearman Correlation Coefficients for Φ and New Complexity Value

3.2.3 Further Style Analysis

3.2.3.1 Curvature

The next property that was considered by Birkhoff was the curvature of the outline of the vase. He specifies that given the points and tangents the curvature should change as little as possible. He believed that this would be a hard trait to analyse but using discrete differential geometry it is possible to form an estimate of the curvature (see appendix B.2.2.2). The difficulty in this case comes from the inability to compare different fitted curves given a set of points and tangents. Other curves could be fitted to the existing points and tangents but further human testing would be required to find the comparison.

Despite the problems, curvature was still analysed to see if it could provide a suitable measure, this is shown in appendix B.2.2.3. This suggested that a small gap between points of inflection was undesirable and formal plotting gave a very good Spearman correlation of 0.5197 with a confidence of over 99%. As these points of inflection would exist even if the curve was adjusted to satisfy Birkhoff's suggested requirement, it could be more accurate to fix the maxima, minima and end points along with their tangents and allow the change of curvature points to move or even disappear when adjusting the curve. However again this is not testable with the current set of results. This observation is supported by the literature; in his thesis on the regularity of curvature Moreton[31] discuss the definition of a fair curve. As well as Birkhoff's description he also includes a requirement by Bu–Qing and Ding–Yuan[7] paraphrased as 'there are no unwanted inflection points on the curve'. A similar definition was also used by Roulier, Rando and Piper [39], also mentioned by Moreton.

Other curvature measures that were studied included the spacing of the significant points, which had a correlation of 0.7467 (confidence >99%) and the total signed curvature (0.4203 with over 98% confidence). The unsigned curvature had a very poor correlation of 0.0554 and a confidence of 61.44%. These are all discussed in greater detail in appendix B.2.2.3.

3.2.3.2 Birkhoff's Further Requirements

This set of requirements gives guidelines that (according to Birkhoff, detailed in appendix A.2.2) should be satisfied to ensure the vase is suitable for purpose. The requirements that have an impact on the style are the proportion–based ones that determine the maximum and minimum widths of the form and their location. The rectilinear requirement also has aesthetic considerations.

The requirements state that the breadth at the base and the minimum width should not be less than $\frac{1}{8}$ of the height. This is satisfied in all of the example vases as the minimum width anywhere on the vase is 1 while the height of all the vases is 4. This is by construction and will not change throughout the evolutionary process. Likewise the requirement that the maximum breadth should be at least $\frac{1}{4}$ of the height was similarly satisfied.

The requirement that the maximum breadth should be less than the height and should not be at either end of the vase was not predetermined. This requirement was partly to ensure that the form was vase–like and not pot–like, however the designs were judged only on their aesthetics and not their type. Therefore it is expected that a preference for a maximum value would not be apparent in the results. The details are in appendix B.2.2.4.

As expected there was no significant correlation (value 0.1320, confidence 75%) for the size of the maximum breadth as it has little impact on the aesthetics of the design. However this could still be used if greater control over the type of evolved pot is required. There was a greater correlation with the position of the maximums (value 0.4198, confidence >98%) and therefore having the maximum at one end does indeed appear to be an undesirable trait.

The rectilinear requirement, that is there should be no straight edges except at the ends, was covered during the curvature analysis. Actual straight lines appear very rarely in a design produced using Bezier curves but visually straight lines do appear. These are characterised by the curvature remaining close to 0 resulting in points of inflection appearing close together. This trait has already been identified as being unattractive and provides a suitable component of a measure.

3.2.3.3 Applicable Measures from Other Work

One style measure that may be applicable from other work is Bergen's deviation from normal (DFN)[5] measure used on 3D graphics. This was the difference between a normalised histogram distribution of the angles between faces and a normal distribution with the same mean and standard deviation. The details are in appendix B.2.2.4 and the correlation was found to be 0.3010 with confidence roughly 95%.

3.2.4 Orientation Analysis

Unlike the original trials the random nature of the test designs meant that many of the designs appeared to be upside down with a higher score for the inverted view. Therefore an essential part of the analysis is to find a measure for the orientation to allow any subsequent fitness function to only evolve vases that are the correct way up. Inspiration was taken from Birkhoff's other requirements as well as the work on screen design.

3.2.4.1 Birkhoff's Further Requirements

Of the six further requirements only two are applicable to the orientation; the requirement for the minimum width to be in the top half of the vase and the restriction on the angles at the base.

Analysis of these two properties (details in appendix B.2.3) show that the minimum width has no notable correlation but the angles at the base have a correlation of 0.8983 with a confidence of over 99%. This could provide a valuable fitness function which uses the angle at the base to judge the design's orientation awarding a high fitness to those with a mid to high angle and a low fitness to those with a low angle.

3.2.4.2 Applicable Measures from other Work

One of the aesthetic measures from screen design that may be useful for orientation is the balance of the screen. This is based on the balance of features either side of the two central axes. As the vases are symmetrical by design there is no information to be gained from investigating the horizontal balance but the vertical balance was investigated. The desired trait from screen design was an even distribution of features either side of the horizontal axis, however it was apparent early on that very few vases of any aesthetic score exhibit this trait. Instead it was observed that 'top heavy' vases were preferred.

The reason for this trait is not obvious, especially considering it contradicts the screen design work. However it has already been observed that measures applicable to one area do not always translate; for example in Barnhardt's study into geometric shapes[1] he found that many people prefer complex geometric shapes while Birkhoff's theory prefers simple forms. The results here for vases support Birkhoff's theory. I would venture the opinion that a top heavy vase is more desirable as it gives the impression of elevating the weight, therefore creating a more elegant solution. This trait is not required in a screen layout.

The weight was measured in two ways (detailed in appendix B.2.3); first using the centre of mass, giving a correlation of 0.5064 (confidence over 99%) and second by finding the balance of the significant points, giving a correlation of of 0.8266 (confidence over 99%).

3.2.5 Derivation of Weights

As several measures were to be used, suitable weights would be needed to find a final fitness function. These will need to be derived automatically to enable the program to vary the weights as it receives further feedback.

A range of methods were tried, including just summing the measures. However the most successful was weighting by the correlations. Better weights were discovered but no simple automatic for finding them was discovered. Full details of this trial are in appendix B.2.4.

3.2.6 Representation Compatibility with Measures

The measures were also tested for their compatibility with the chosen representation. A random vase was generated and mutated several times (shown in figure 3.4), the measures were found for each vase and plotted to see how they changed as the design changed (figure 3.5).

Figure 3.4: Sequence of Mutating Vases

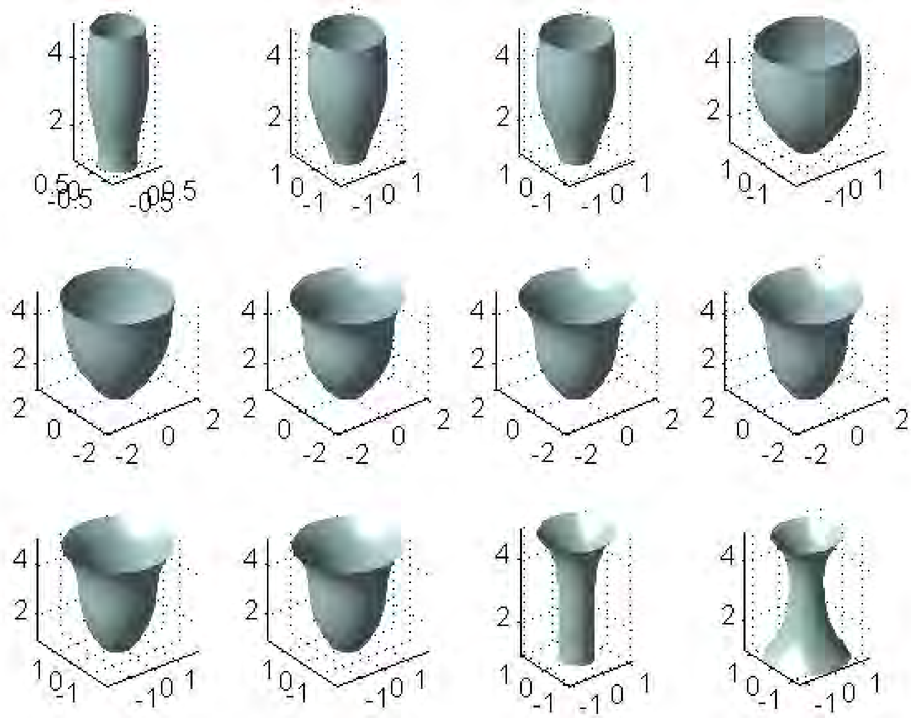

Figure 3.5: Measure Scores for Mutating Vases

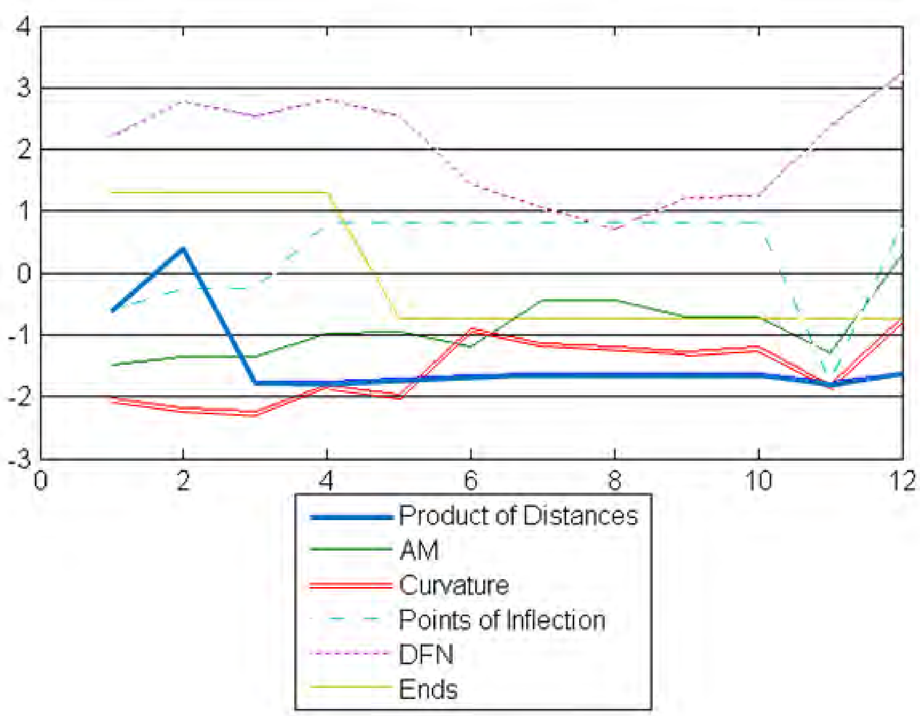

Figure 3.6: Measure Scores for Crossover Vases

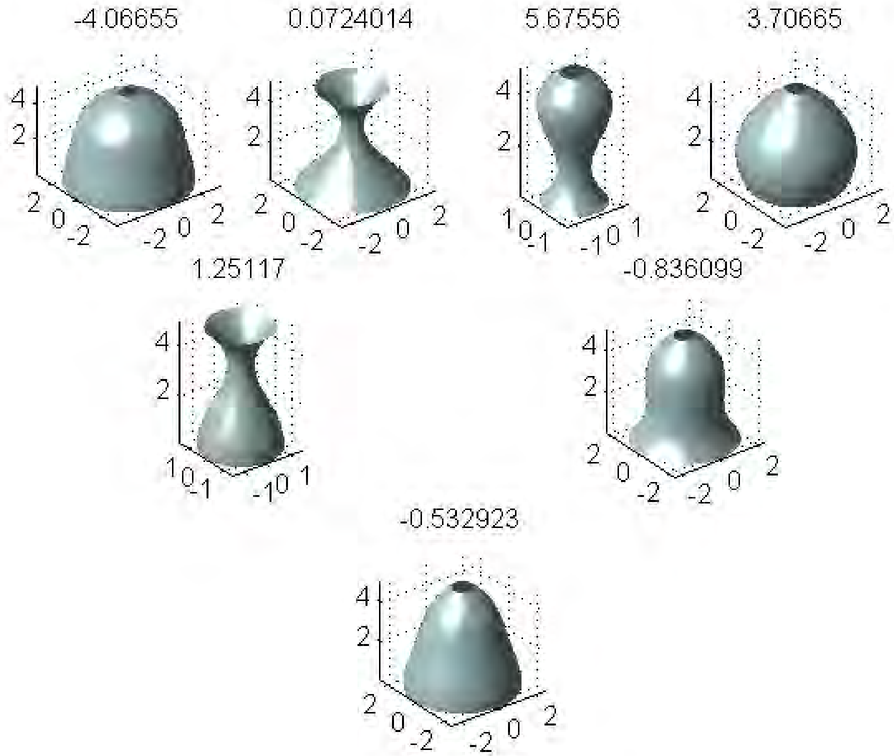

A compatible measure would change very little if there was no visual change and would change much more for a large visual change. Overall there appears to be a good correlation between visual changes and the scale of changes to the measure scores, there are some anomalies such as the large change to the product of distances scores between 2–3 despite no visual change but most are as expected. Repeating the test on other mutated vases produces similar results.

However crossover had a much larger effect on the score, with the offspring's score often very different to the parents. An example of this is shown in figure 3.6, which shows a 'family tree' of vases with the sum of their scores. This tree shows that it is possible for a child to have a very different score to its parents but often the child's score is closer to the parent that it more closely resembles. For example the first child on the left has a closer resemblance to its right hand parent and its score is closer to that one although higher than both. This is good as it demonstrates the ability of a child to be an improvement on its parents.

3.3 Summary

The preliminary studies within this project aimed to find a suitable representation for a vase as well as an initial selection of measures for further testing.

The representation found was a rotational shape formed using a 2D Bezier curve. Parametric equations were also considered but as discussed in detail in section 3.1.3 they proved to be erratic when genetic functions were applied. The Bezier curves behaved very well when using genetic functions; mutation would create a small but noticeable change to the form and the various methods of crossover used had recognisable traits from both parents.

By using rotational forms and a simple embryology that moved the curve a suitable distance from the rotational axis, all of the vases produced by this method were viable vases, capable of standing and holding flowers and water.

Forms randomly generated using this representation were judged to have aesthetic scores ranging between 1–9. This is good for the later experiments as it allows a wide range to be considered and indicates that the representation is capable of producing designs that will be liked.

Further analysis of the preliminary results also indicated 6 possible style measures and 3 orientation measures. Birkhoff's measure required some modification, possibly as a result of changing tastes since his ideals of aesthetics were formed. His further requirements also required some modification. Some, such as the angles at the base provided a very good measure of orientation which was formalised to enable it to be used within the evolutionary process. Others, such as the minimum width being towards the top of the vase seem to be contradicted entirely by these results. Again possibly the result of changing tastes but perhaps also the effect of analysing random forms as opposed to designed vases. Other measures were also found, inspired by a range of other sources.

Although the measures noted here appear to have a very strong correlation, the limited numbers of vases and volunteers meant that further work was required to test the validity of the measures.

Chapter 4

Experimental Setup

The main experiment was designed to automatically find a weighted sum of the previously described measures with the aim of finding a final fitness function capable of evolving a range of interesting and attractive vase designs. As such it had several requirements:

- To gain human–awarded scores for a range of designs from a large sample of volunteers.
- To automatically generate weights in response to those scores.
- To evolve designs using this weighted fitness function to test for exploitable loopholes.

The relative weights of the different measures should reveal which measures are universal and by collecting data on the volunteers it may also reveal some that depend on the individual. It is hoped that the collection of this data will lead to a measure that can:

- Evolve a wide range of interesting and attractive designs.
- Generate good designs more consistently than random evolution or those generated by the original Birkhoff measure.
- Create designs that will be suitable for manufacture into physical products.

4.1 Parameter Choices

4.1.1 Representation and Genotype

The final representation was a Bezier curve based form, with a genotype similar to the original list of x, z coordinates used in the initial trials. It was decided that the z coordinates would be fixed to ensure consistent viable designs. This also meant the genotype could be reduced to a single list of x coordinates. As variable length lists are difficult to handle in Matlab a fixed list of 5 coordinates was chosen.

In initial tests a 5–point Bezier curve was shown to be capable of representing a wide variety of designs. However during program development it was observed that in practice many of the possible designs seem similar. Therefore some forms seemed to repeat frequently. One possible method of increasing the diversity was to increase the number of coordinates but this increased the diversity by adding more complex forms to the existing range. As initial tests showed a preference for simple forms it was felt these would not be a valuable addition.

The method chosen to create a greater diversity of forms modified the representation to allow sharp angles to exist between two or more Bezier curves. These forms were created by breaking the curve at a set coordinate point or points, with the coordinates either side forming 2 or more Bezier curves. This introduced a wide variety of new designs, some that potentially could be considered attractive and a valuable addition. Further information is given in appendix C.1.1.

The height was increased from 4 to 5 to help restrict the designs to vases as opposed to bowls. It meant that it was possible for a design to have a maximum width below $\frac{1}{4}$ of the height which was one of Birkhoff's further requirements but it is unlikely as the total difference between maximum and minimum radius would have to be less than 0.25.

Finally the vases were now shown in a neutral blue colour with natural shading rather than the multicoloured Matlab figures used in the initial trials.

4.1.2 Measures

Initial work suggested 6 possible style measures and 3 orientation measures. To recap these were:

Style:

- Birkhoff's Aesthetic Measure with new complexity function
- Sum of the signed curvature
- Number and separation of points of inflection
- Maximum width between ends
- Difference from Mean
- Product of the Distances

Orientation:

- Difference between angles at base and rim
- Balance of features
- Location of physical center of mass

Some modification were made, mainly to cope with the new addition to the representation. These modifications are detailed in appendix C.1.2. A new measure was also added to judge the effect of the new angles. The m files that were used to calculate the measures are also found in appendix C.3.

4.1.3 Scoring

The scoring system was chosen to generate a wide distribution of scores. The initial tests used a 1-10 scoring system but it was felt that this was not intuitive and the generosity of the volunteer would vary.

The scoring system used in the final program used the comparison of 4 designs at a time with the volunteer asked to select their favourite. This one was then awarded 3 points and the others had one point deducted from an initial score of 0. To highlight the subtle differences in designs the 4 were deliberately selected so that designs with similar scores were directly compared.

The final scores used in the weights were then divided by the number of times that they were tested plus 1. The division was to prevent artificial increases in scores from the random nature of the selection. For example if there were two designs that were both loved and always selected then if one was tested twice and the other 3 times then the score for the first would be 6 and the other 9, which may not necessarily reflect their relative attractiveness. The plus 1 was to prevent dividing by zero.

This method was tested with the initial results to ensure it accurately represented the preferences of the individuals. A similar set up was created where instead of a volunteer choosing their preference the vase with the highest score from the initial test was chosen instead. The test was run 10 times for each initial test subject giving a total of 80 tests and the correlation was consistently above 0.8% with confidence above 99%. The method therefore gave a very good approximation of the preferences of the individual.

4.1.4 Weights

The m file *WeightFinder.m* finds the weights by finding the correlation between the scores and measures. The correlation used is Spearman's rank correlation as it considers non–linear relationships.

The first 5 measures are correlated directly with the highest score out of the two orientations. The confidence of the positive correlation is also found and if it is below 80% then the weight is set to 0, above this it is given the correlation value. This also sets any negative correlation to 0.

The 6th measure was introduced to discover if there is any relationship regarding the new sharp angles. As it is not known if there is a relationship and whether it may be positive or negative this is not tested for confidence. As the correlation would be low if the confidence was low then this would not make a significant difference. It also allows a negative value if the relationship suggests it.

The weights of the 3 orientation measures are found by comparing the differences between the scores of the two orientations and the differences between the measures. Again these are tested for their positive confidence and changed to 0 if the correlation is too small.

The weights would change at various intervals in the program using all of the previous results. They were initialised using a sample set of 4 vases which gave the weights for all the measures as 1.

4.1.5 Evolutionary Parameters

4.1.5.1 Mutation

The mutation method used was the addition of a Gaussian random number to one of the x coordinate points. This was done with a probability of 20%. In addition there was a possibility of the sharp angle cut–off points appearing or disappearing with a probability of 5% or moving with a probability of 20%. Finally there was a 20% chance of a design being flipped.

4.1.5.2 Recombination

Recombination was done using the splicing method, using the values of two different vases and the splicing point was random. The cut–off points were split at the same point.

4.1.5.3 Population Size and Number of Generations

The population size and number of generations was restricted by the requirement that the evolution takes place between user selections. If the population size was too large very few generations could take place between each selection and the program would feel unresponsive as a generation must finish before it would respond. The final number chosen was 20 as this allowed the program to respond quickly but allowed for a good range of designs.

To increase the range of designs further a much larger population of 200 was initially generated and all the measures applied during the initial set–up while the instructions were being read by the volunteer. The initial population was selected from this group and new individuals were added to the main population from this group after each change to the weight. This helped maintain genetic diversity in this small population without increasing the generation time.

The actual number of generations depended on the individual as someone who took a long time to choose would allow more generations to take place than someone who responded quickly.

4.1.5.4 Selection

The selection was carried out after the genetic operators as they could be disruptive to some measures. The selection of the parents for the genetic operators was evenly distributed through the entire population. The next generation was then selected from the group of both offspring and parents.

Although several methods of random selection were tried they usually showed little increase in best score during the evolutionary process. This was possibly due to the small population size. To improve this elitist selection was used.

4.1.5.5 Constraints

One of the early problems was the preference for very wide vases. To prevent this a separate test that set the score of every vase of radius > 2 to -40. This appeared to be sufficient to prevent the extra wide designs.

As the final evolutionary step called for 10 vases to be included at one time there was a need to create diversity within the population. To do this a similarity measure was used; the differences between all radius values was found and summed together, a larger value therefore creating a larger difference. It was not perfect as visually similar vases could have a large calculated difference but this method provided a better range of designs in tests than other similarity measures such as the difference in the derivatives.

After a vase had been selected for the next generation the vases with a similarity score of less than 10 were found and their measure scores reduced proportionally with a reduction of 10 for identical vases and 0 for vases with a similarity score of 10. The next vase was chosen based on these new scores.

4.1.6 Program Choice and Machine Details

Matlab[29] was chosen for this project as it allowed the 3D vase renderings to be produced from the genotype quickly and easily. It also supported the computational requirements for the evolutionary process and allowed a simple GUI to be produced for the experiment. Matlab version R2007a was used and the program was built and run on a laptop running Windows 7 64-bit with an Intel 2.4GHz Core i5 CPU and 4GB RAM.

4.1.7 Vase Types

To allow comparison with Birkhoff's original measure and randomly generated designs it was decided that the program should include vases of these types to allow the collection of data. A set of 200 designs were evolved using only Birkhoff's formal measure (including caps and independence but no further requirements). These were then available for the program to randomly select at the beginning of the test.

In addition 10 set designs were manually chosen from a large set of random designs to include a wide range of every measure, including extremes, as well as being visually different. These would be used for every volunteer to see the extent of individual taste.

4.2 Program Layout

The program is split into two main sections that run alternately to enable the results to be both collected and used in the evolution of new designs. The first part is the human interface, which displays 4 vases to the volunteer, prompting them to chose their favourite. The score of the favourite then increases and the other 3 have theirs reduced.

The second is the behind–the–scenes part which evolves vases in the time it takes for the volunteer to make their choice. When appropriate it adds new vases to the selection and updates the weights.

A full description and screenshots are included in appendix C.2.

4.3 Summary

The initial work was refined; the representation was modified to allow angles within the curves and this necessitated changing some of the measures. One measure (the difference from mean) was dropped due to its long computation time and the problems this would cause during evolution.

The new representation was able to describe a wide variety of vases with various aesthetic qualities. The measures consider many different properties of a vase and with further testing they have the potential to evolve attractive forms.

To test the measures and representation a program was written to allow vases of various types to be considered by volunteers. These included vases evolved using the latest weighted measure. This weighted measure was found automatically during the process, responding to the preferences of new users.

A test program of this type was essential to the development of a good measure as it allowed the volunteer to test the measure directly so that any loopholes (giving good scores to bad designs) found by the evolutionary process could be removed. Tests only involving set or random designs, as per the initial tests, could easily develop loopholes of this kind.

The program was designed to be used consecutively by the volunteers, thus the later volunteers would be shown examples of evolved designs using measures developed from the results of all previous volunteers. This should allow the program to converge onto global weights from all users. If successful these weights will then provide a suitable measure for the evolution of attractive vases.

Chapter 5

Analysis and Discussions of Experimental Results

A total 22 volunteers completed the tests. This led to the development of a global measure using weighted values of the individual measures. In turn this allowed the evolution of new designs.

5.1 Developed Global Measure

The weights of the final measure are given in table 5.1. The totals found by weighting the individual measures by these weights gives an overall correlation of 0.3271 (confidence over 99%).

The progression of the weights though the test period is shown in appendix D.1.1. It was observed that as the weights stabilised the evolved vases converged to a particular type.

Measure	Weight
Birkhoff	0.0486
Curvature	0
Points of Inflection	0.1628
Width at Ends	0
Product of Distances	0.0594
Existence of Sharp Points	0.3035
Angles at Ends	0.3970
Balance of Features	0.0859
Centre of Mass	0

Table 5.1: Final Weights

This convergence of vase designs had an undesirable effect on the range of vases to be judged; nearly all of the 20 evolved designs would be similar. The test program had been designed to remove loopholes where a bad design was awarded a good score but as these designs were liked further testing only reinforced this particular design.

Although this design was popular it was rarely able to achieve the highest scores. This was possibly due to the way the scores were allocated whereby designs with similar scores would be shown together. This meant that many volunteers would be faced with a selection of 4 nearly identical vases and would have to chose one at random. A unique design could score much higher as if it was suitably liked it could be picked every time. Alternatively, the design may not be popular enough to achieve the highest score, but with the effect of the score sharing it is impossible to tell.

It was hoped that the final measure would be capable of designing a range of forms and this does not achieve this, however it provides a starting point for a more flexible measure.

5.2 Comparison of Evolved, Random and Birkhoff vases

It was desirable that the new measure should be capable of producing good designs more consistently than either random generation or Birkhoff's original measure. The sets of each type were divided up and the average scores found. The preselected vases were not included in this analysis. Further analysis was also carried out to find the effect of the new sharp angles (further details are in appendix D.1.2).

As expected the evolved designs had a much higher average than the random or Birkhoff evolved designs. Closer examination revealed that this was mainly due to the ability of the fitness function to remove sharp edges and correctly orientate the designs. Manually selecting for these properties gives the Birkhoff evolved vases a huge average score of 0.7922. It should be noted, however, that the evolved designs may be victim to the effect of the recurring 'ideal' design. If 40% of the designs are very similar then their scores will be divided between them while unique designs are able to achieve much higher values. This is also shown in the distribution of scores. There is a higher incidence of evolved designs in the higher scores but the very highest has a much wider mix of types.

5.3 Test Group

The test group consisted of 22 individuals with no overlap between them and the preliminary group of 8. The group was roughly divided into two distinct sets of people, scientists (mainly physics and maths) with little or no design experience and designers, all of whom were architects.

During the testing several observations were made:

- Some individuals would choose intuitively while others would take a very long time. As the tests were not timed there is no way of comparing the results of any individuals with their time taken.

- Many volunteers, in particular the designers, would comment on the practical aspects of the vase such as the stability and suitability for flowers and selected for these traits as well as aesthetics. One of the designers pointed out that as architects they are encouraged to consider holistic design and more art–based designers may have different preferences.

- Some of the volunteers, all designers, commented on the colour of the vase. The colour had been chosen to be neutral and not a factor but they suggested that some designs with a different colour or pattern may be much more attractive than the current version.

- Several confessed to changing their minds half way through and it was suggested by some that comparing all the vases simultaneously would have made it easier for them to chose.

Measure	Designer	Non–Designer
Birkhoff	0	0.0469
Curvature	0	0
Points of Inflection	0.0526	0.2052
Width at Ends	0	0
Product of Distances	0	0.0709
Existence of Sharp Points	0.175	0.3500
Angles at Ends	0.5005	0.3579
Balance of Features	0.1000	0.0889
Centre of Mass	0	0

Table 5.2: Designer Vs. Non–Designer Comparison.

Measures for both groups were also found to see if there were any differences. The measures are shown in table 5.2. Previous studies have often found that trained volunteers have a lower correlation with aesthetic measures than untrained volunteers. Here the style results seem to agree with this but the orientation measures contradict it. As the number of designers was low (7) it is unknown whether this is a general trend or the influence of a few individuals. Further testing would be required to test the potential differences.

The pre–set designs were used to provide a direct comparison between the volunteers. The correlations between them were found. These ranged between 0.4621 and -0.3977. Looking at the plots in appendix D.1.3 we see that some are generally liked while others are disliked. Some also appear to have a preferred orientation. However with all of these there are exceptions suggesting a strong element of personal taste.

5.4 Other Measures and Possible Adaptations

5.4.1 Modifications

The surprisingly high scores for the Birkhoff designs suggested that it was worth revisiting the original measure to see if the correlation was an improvement over the initial tests.

The Spearman correlation for the original Birkhoff measure, with the original ratios had a value of 0.2082 with a confidence of over 99%. This is significantly higher than the modified version with a value of 0.0330 and confidence about 86%. This may be partially due to the high occurrence of vases with 5 significant points (as most of the evolved have this property). These will all have low scores from the modified Birkhoff measure but are generally liked by the volunteers.

One of the main causes of the convergence was the strong weight of the edge angles measure which encouraged the base to curve outwards and the rim to flare. This combination is easily satisfied by the vase that is evolved. However it is valuable as an orientation measure and this is one of the main strengths of the new measure. Therefore the overall measure was changed so that its value was dependent only on the style measures. The weighted sum of the orientation measures was also found and if the sign of this measure was negative then the design was flipped with a probability of 50%. This biased the design towards correctly oriented vases but did not affect the style.

This also allowed the reintroduction of the radius component of the Balance of Features orientation measure. This was originally removed to prevent encouraging the evolution of very wide designs, but also reduced its correlation. Reintroducing it increases the correlation from 0.0859 to 0.1891.

Another cause of the convergence was the lack of conflict between the measures. All the significantly weighted measures could achieve high values on this design with none achieving only at the expense of another. To combat this problem another measure was introduced deliberately designed to conflict with the existing measures. As the preferred design had 5 significant points the new measure awarded those with 4 or less. The weighting was varied, not by any correlation value but randomly with a slight bias toward positive weightings. Along with the other modifications this gave a much wider variety of designs.

The new weightings are given in table 5.3. Using these with the weight for number of significant points as 0.1, the new overall correlation for the style scores is 0.2064 with >99% confidence. The orientation scores have a correlation of 0.2643, again with a confidence of over 99%.

Measure	Weights
Original Birkhoff	0.2082
Curvature	0.0916
Points of Inflection	0.1628
Number of Significant Points	Variable
Product of Distances	0.0594
Existence of Sharp Points	0.3035
Angles at Ends	0.3970
Balance of Features	0.1891

Table 5.3: New Weights

5.4.2 Testing the Final Designs

To test the modifications 10 designs for each volunteer were found; 6 designs were generated using the new measures with a wide variety of weights for the variable measure. This was evolved using *EvolveExamples.m* in appendix D.2. One was an example of the convergent designs using the unmodified measure and two used weightings derived from each individual's results. The final vase was the highest–scoring design from the previous work that could have been any type used. Each volunteer was asked to put these in order of preference, where 1 was their favourite and 10 was their least favourite. Not all of the volunteers were available for the follow–up but 13 of the original 22 responded. They were also asked to give marks out of 10 for their favourite and least favourite to see the range. The favourites ranged in scores from 6–9 and the least favourites from 0–6.

Overall the results suggest that the new measure (with a positive weight for the variable measure) is capable of evolving a range of designs in which one or more will be liked by most people. A breakdown of these results is given in appendix D.1.4. This range could be further extended by relaxing the strict maximum height and minimum width currently used in the program.

5.5 Final Designs

The final weights are capable of evolving a range of vases, although some repetition still occurs. Figure 5.1 shows 9 independently evolved vases using this fitness measure. It also shows the value of the random weight for each vase, with low values giving complex vases and high giving simple designs.

The designs evolved here are still only sketches. They have no base and have no thickness to their walls. In order to translate these into viable vases a final program *FullVase.m* (appendix D.2) was written. This used the evolved form as the outer skin, but created an inner skin as well as a base. Figure 5.2 shows an example evolved vase after this process. The base was designed with a lip that would be necessary if these were to be produced in ceramic to allow application of glaze. The program also used a function called *surf2stl.m* written by Bill McDonald from the Matlab file exchange[46] to change the Matlab surface into a .stl file that would be usable by 3D printing firms.

Companies such as *Shapeways*[43] and *Sculpteo*[41] offer ceramic printing. Example models were ordered from *Sculpteo*. For reasons of cost the vases were small and in order to satisfy their printing requirements it was necessary to increase the minimum radius from 0.5 to 0.75. New vases were evolved with this width and some selected to be printed (figure 5.3).

None of the vases are particularly unusual in design; on the contrary many are unintentionally similar to designs available on the high street (figures 5.4 and 5.5). However rotational vases with no patterns do not offer great potential for original new forms as other evolutionary art such as fractal patterns have done. The achievement here is the ability for the computer to consistently produce designs, equivalent to those designed by humans, which are attractive and potentially marketable. It is hoped that this case study will allow the extension of this theory to other, more complex products, which would take a human much longer to design. Combined with 3D printing it could allow the automatic production of large selections of products, allowing consumers to chose a unique item without a designer having to design each and every one.

Figure 5.1: 9 Independently evolved vases using final selection. Figure above indicates value of random weight

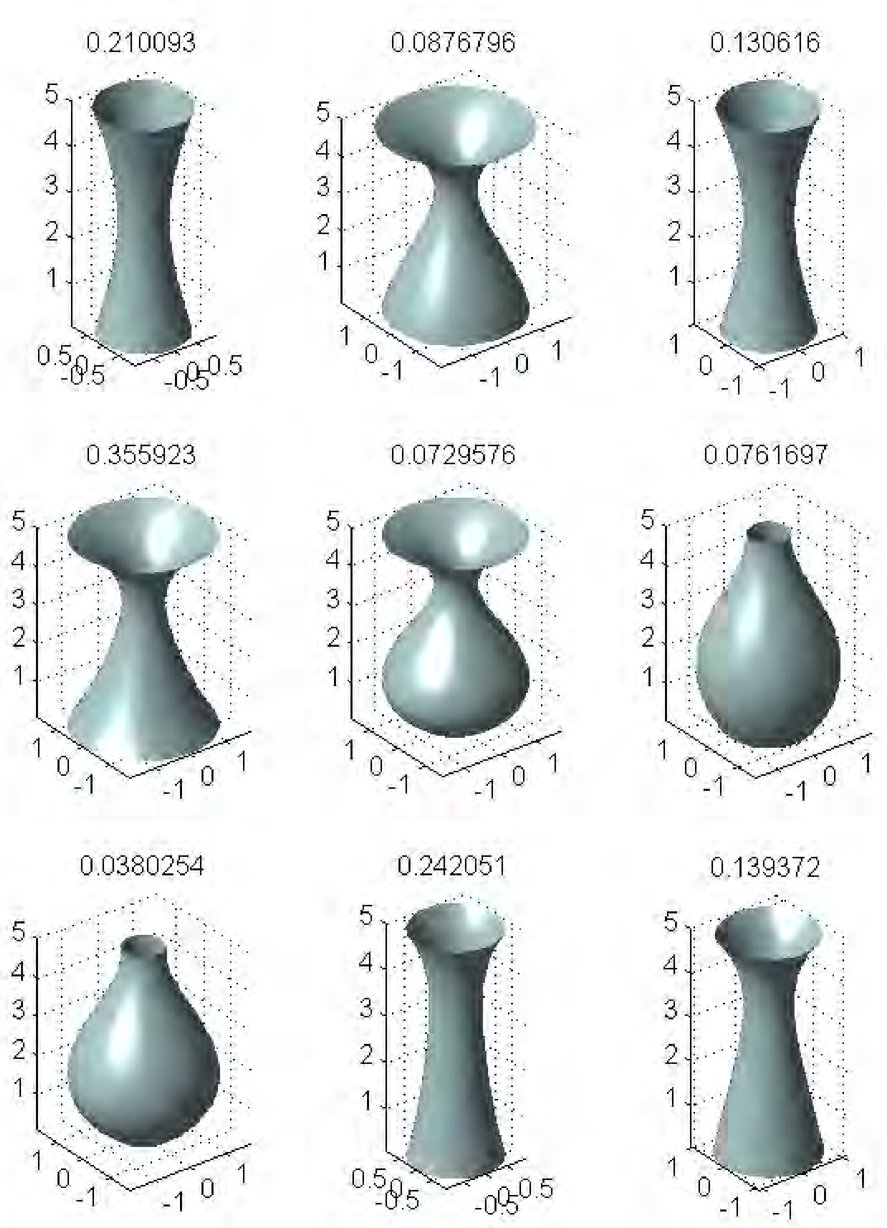

Figure 5.2: Top and Base of Full Vase

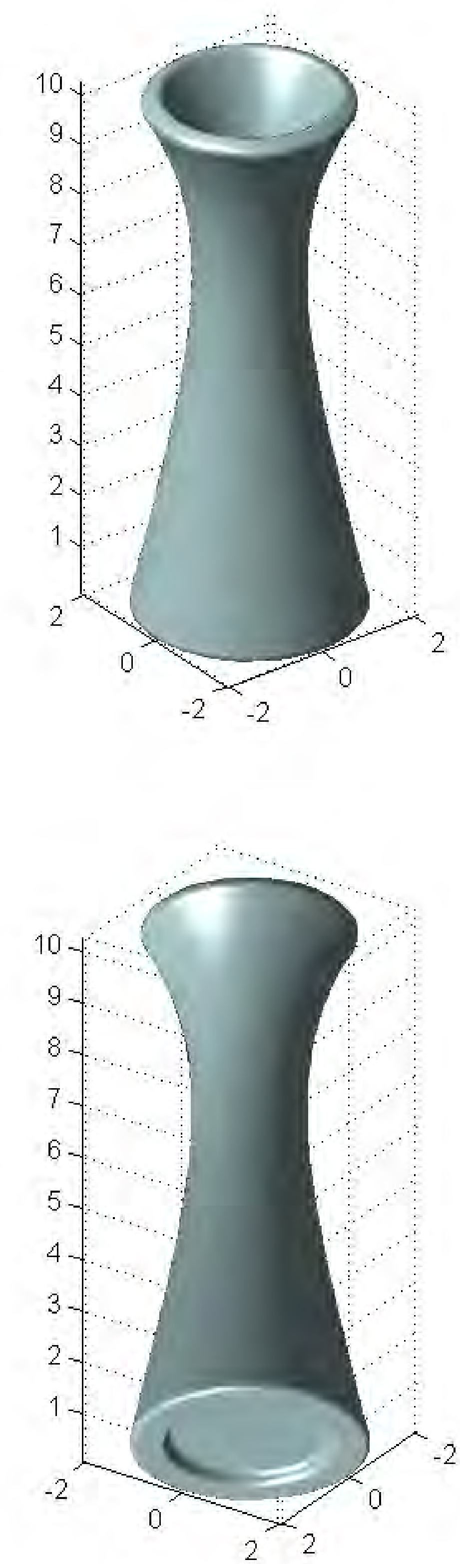

Figure 5.3: Printed Vases in White Ceramic

Figure 5.4: Example of Vases on sale at John Lewis [21][22]

[redacted]

Figure 5.5: Example of Vases on sale at Marks and Spencer [27][28]

5.6 Summary

The results of the experiment showed that many of the measures did have a positive correlation with the opinions of the volunteers. However in the original form the resulting fitness function would always evolve the same design. Further analysis of the results and re–examination of some of the earlier forms of the measures suggested a new method of combining the measures into a fitness function.

One of the vital steps was the separation of style measures and orientation measures, using the orientation to reorient upside down vases but not to influence the style. The other was to introduce a conflicting measure as it was observed that the recurring design was successful in all of the measures. This new measure awarded designs with less than 5 significant points, conflicting with the recurring design. The weight of this was randomised to allow for a range of designs. Follow–up questions confirm that many of these were found to be attractive by the original test subjects.

Finally the possibility of creating a physical version of the form was explored. A program was created which turned the designs into full vase forms with a thickness to the walls and a functioning base. An example of one of these vases was printed in ceramic, creating a fully functional desirable product entirely designed and built by computers.

Chapter 6

Conclusion

Previous research has begun to demonstrate the potential of using aesthetic measures as fitness functions in evolutionary design. To develop suitable measures we need to understand what aspects of a design influence human preferences. The previous research, supported by the results of this project, suggest that there is unlikely to be a single measure of aesthetics, yet by studying human responses to different designs it is possible to determine some guidelines for different objects. These can then be used as the basis for a fitness function to evolve the objects independently.

The results show that despite studies that have contradicted Birkhoff's measures, including the initial work in this study, they do have some correlation with human preferences in the field of vase design and can form the base of an overall fitness function. However as predicted it requires additional measures to ensure the correct orientation. Other measures, judging the curvature as well as the number and location of significant points also have value. Many of these additions were mentioned by Birkhoff as additional requirements but here they have been formalised and confirmed.

The final evolutionary fitness function also required a random element to ensure a suitable range of designs. This was because all the measures agreed on one particular form that therefore dominated the evolutionary population. A final test demonstrated that the new function was capable of designing a range of vases that included designs liked by the majority of volunteers. Many named one of these their favourite of the entire test.

Finally it was shown that these designs could then be converted to full 3D vase models that were capable of being printed in ceramic, thus creating a final product from the process.

It is hoped that this demonstration of the potential of 3D aesthetic measures to design functional products will lead to work on other items, allowing the automatic design and manufacture of attractive products.

Bibliography

[1] Barnhardt, E.N.: The Criteria used in Preferential Judgments of Geometrical Forms. In: The American Journal of Psychology, Vol. 53, No. 3, pp. 354–370. University of Illinois Press. (1940)

[2] Beebe–Center J.G and Pratt, C.C.: A test of Birkhoff's Aesthetic Measure. In: Journal of General Psychology. Vol. 17, pp.339–354. (1937)

[3] Bentley, P. J.: Generic Evolutionary Design of Solid Objects using a Genetic Algorithm. Ph.D. Thesis, Division of Computing and Control Systems, Department of Engineering, University of Huddersfield. (1996)

[4] Bentley, P. J.: From Coffee Tables to Hospitals: Generic Evolutionary Design. In: Bentley, P. J. (ed): Evolutionary Design by Computers. pp.405–423, Morgan Kaufmann Publishers, San Francisco. (1999)

[5] Bergen, S.: Automatic Structure Generation using Genetic Programming and Fractal Geometry. Brock University, Ontario. (2011)

[6] Birkhoff, G.: Aesthetic Measure. Cambridge Massachusetts University Press. (1933)

[7] Bu-Qing, S. and Ding-Yuan, L.: Computational Geometry, Curve and Surface Modeling. Academic Press, San Diego. (1989)

[8] Child, I.L.: Development of Sensitivity to Aesthetic Values. In: Cooperative Research Project No. ED-003-483. New Haven, Connecticut. Yale University. (1964)

[9] Davis, R.C.: An Evaluation and Test of Birkhoff's Aesthetic Measure formula. In: Journal of General Psychology. Vol. 15, pp.231–240. (1936)

[10] Website of Driessens and Verstappen: `http://notnot.home.xs4all.nl/index.html`

[11] Eby, D., Averill, R. C., Punch, W. F. and Goodman, E. D.: The Optimization of Flywheels using an Injection Island Genetic Algorithm. In: Bentley, P. J. (ed): Evolutionary Design by Computers. pp.405–423, Morgan Kaufmann Publishers, San Francisco. (1999)

[12] Ecker, D.W.: Justifying Aesthetic Judgments. In: Art Education, Vol. 20, pp. 5-8. (1967)

[13] Ecker, D.W., Johnson, T.J. and Kaelin, E.F.: Aesthetic Inquiry. In: Review of Educational Research, Vol. 39, No. 5, pp. 577-592. (1969)

[14] Ekárt, A., Sharma, D. and Chalakov, S.: Modelling Human Preference in Evolutionary Art. In: Di Chio, C.(ed.): Applications of Evolutionary Computation 2011. LNCS, vol. 6625, pp. 303–312. Springer, Heidelberg (2011)

[15] Harsh, C.M, Beebe–Center, J.G. and Beebe–Center, R.: Further Evidence Regarding Preferential Judgment of Polygonal Forms. In: The Journal of Psychology, Vil. 7, pp. 343–350. Provincetown, Massachusetts. (1939)

[16] den Heijer, E. and Eiben, A. E.: Comparing aesthetic measures for evolutionary art. In: Di Chio, C. et al (eds.): Applications of Evolutionary Computation 2010. LNCS, vol. 6025, pp. 311–320. Springer, Heidelberg (2010)

[17] Hingston, P., Barone, L., While, L.: Evolving Crushers. In: Proceedings of the 2002 Congress on Evolutionary Computation. CEC '02. pp. 1109–1114. IEEE Press, New York (2002)

[18] Gu, Z., Tang, M.X. and Frazer, J.H.: Capturing Aesthetic Intention during Interactive Evolution. In: Computer–Aided Design, Vol. 38, pp. 224–237. Elsevier. (2006)

[19] Jacob, C., Kwong, H. and Wyvill, B.: Toward the Creation of an Evolutionary Design System for Implicit Surfaces. In: Western Computer Graphics Symposium 2001, Skigraph'01, Sun Peaks Resort, British Columbia. (2001)

[20] Jacob, C. and Hushlak, G.: Evolutionary and Swarm Design in Science, Art and Music. In: Romero, J. and Machado, P. (eds.) The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music. Springer, Berlin. (2008)

[21] John Lewis Vase: `http://www.johnlewis.com/231586111/Product.aspx`

[22] John Lewis Vase: `http://www.johnlewis.com/231586108/Product.aspx`

[23] William Latham's Website: `http://www.doc.gold.ac.uk/~mas01whl/`

[24] Latham, W. and Todd, S.: The Mutation and Growth of Art by Computers. In: Bentley, P. J. (ed): Evolutionary Design by Computers. pp.221–250, Morgan Kaufmann Publishers, San Francisco. (1999)

[25] Li, Y., Hu, C. J., Yao, X.: Innovative Batik Design with an Interactive Evolutionary Art System. In: Journal of Computer Science and Technology, 24(6), pp 1035–1047. (2009)

[26] Machado, P., and Cardoso, A.: Computing Aesthetics. In: Proceedings of the Brazilian Symposium on Artificial Intelligence, SBIA-98, pages 219–229. Springer, Heidelberg (1998)

[27] Marks and Spencer Vase: `http://www.marksandspencer.com/Marks-and-Spencer-Ceramic-Vase/dp/B003TM0U24?ie=UTF8&ref=sr_1_4&nodeId=48263031&sr=1-4&qid=1346339713&pf_rd_r=1W0JJKE6NDMGA53W1P02&pf_rd_m=A2BO0OYVBKIQJM&pf_rd_t=101&pf_rd_i=48263031&pf_rd_p=321381387&pf_rd_s=related-items-3`

[28] Marks and Spencer Vase: http://www.marksandspencer.com/Marks-and-Spencer-Floral-Small/dp/B007ZX6TJQ?ie=UTF8&ref=sr_1_5&nodeId=48263031&sr=1-5&qid=1346339713&pf_rd_r=1W0JJKE6NDMGA53W1P02&pf_rd_m=A2BO0OYVBKIQJM&pf_rd_t=101&pf_rd_i=48263031&pf_rd_p=321381387&pf_rd_s=related-items-3

[29] Matlab: http://www.mathworks.co.uk/products/matlab

[30] Meinhardt, H.: The Algorithmic Beauty of Seashells, Third Edition. Springer-Verlag, Berlin. (2003)

[31] Moreton, H.P.: Minimum Curvature Variation Curves, Networks and Surfaces for Fair Free-Form Shape Design. University of California at Berkeley. (1992)

[32] Ngo, D.C.L. and Byrne, J.G.: Aesthetic Measure from Screen Design. In: Computer Human Interaction Conference 1998. Proceedings. pp. 64–71. IEEE Press, New York (1998)

[33] Nishino, H., Takagi, H., Cho, S. B. and Utsumiya, K.: A 3D Modeling System for Creative Design. In: 15th International Conference on Information Networking, 2001. Proceedings. pp. 479–486. IEEE Press, New York (2001)

[34] Picbreeder Website http://picbreeder.org/

[35] Pratt, C.C.: The Stability of Aesthetic Judgements. In: Journal of Aesthetics and Art Criticism, Vol. 15, pp. 1-11. (1956)

[36] The Evolutionary Art of Steven Rooke Website http://srooke.com/

[37] Rooke, S.: Eons of Genetically Evolved Algorithmic Images. In: Bentley, P. J. and Corne, D. W. (eds): Creative Evolutionary Systems. pp. 129–145, Morgan Kaufmann Publishers, San Francisco. (2002)

[38] Ross, B., Ralph, W. and Zong, H.: Evolutionary image synthesis using a model of aesthetics. In: IEEE Congress on Evolutionary Computation (CEC) 2006, pp. 1087–1094, (2006)

[39] Roulier, J.A., Rando, T. and Piper, B.: Fairness and Monotone Curvature. In: C. K. Chui (ed.): Approximation Theory and Functional Analysis. Academic Press. (1991)

[40] Schnier, T. et al: Nature Inspired Creative Design – Bringing Together Ideas from Nature, Computer Science, Engineering, Art, Design. In: Designing for the 21st Century: Innes, T. (ed.): Interdisciplinary Questions and Insights, pp. 192–204, Gower Ashgate. (2007)

[41] Sculpteo Website: `http://www.sculpteo.com/en/`

[42] Serrano, A.: histnorm.m. From: `http://www.mathworks.com/matlabcentral/fileexchange/22802-normalized-histogram`, (2009)

[43] Shapeways Website: `http://www.shapeways.com/`

[44] Spehar, B., Clifford, C.W.G., Newell, B.R. and Taylor, R.P.: Universal aesthetic of fractals. Computers and Graphics, 27(5):813–820, (2003)

[45] Staudek, T.: On Birkhoff's Aesthetic Measure of Vases. Technical Report. Faculty of Informatics, Masaryk University. (1999)

[46] McDonald B.: surf2stl.m. From: `http://www.mathworks.com/matlabcentral/fileexchange/4512`, (2004)

[47] Walsh, P., Gade, P.: The Use of an Aesthetic Measure for the Evolution of Fractal Landscapes. In: IEEE Congress on Evolutionary Computation (CEC) 2011, pp.1613–1619, IEEE Press, New York. (2011)

[48] Whitelaw, M.: Breeding Artificial Objects: Art and Artificial Evolution. In: Bentley, P. J. and Corne, D. W. (eds): Creative Evolutionary Systems. pp. 129–145, Morgan Kaufmann Publishers, San Francisco. (2002)

[49] Wilson, D.J.: An experimental investigation of Birkhoff's Aesthetic Measure. In: The Journal of Abnormal and Social Psychology, Vol. 34(3), pp. 390–394. (1939)

Appendix A

Appendix: Literature Review Further Infomation

A.1 Further Details of Representations

A.1.1 GADES

GADES by Bentley [3] uses spatial partitioning with simple primitives called 'clipped stretched cubes'. These are described by a cuboid that has been 'clipped' by a plane running through the shape which allows the development of more complex shapes than would be possible with just rectangular forms. Each primitive has only 9 parameters; x,y,z coordinates, height, width and depth as well as the description of the plane: angle1, angle2 and distance from centre. Lists of these 9 values made up the genotype of the object.

Unclipped cubes were also used, particularly in the middle of a form. If two primitives intersected then the phenotype was adjusted by shrinking the cubes until they no longer overlapped. Primitives were also able to be added or deleted. A new primitive was created by splitting an existing cube in two. This would have no effect on the design at that stage but would subsequently allow it to evolve more complex forms. By contrast deleting a primitive could be a very large change but was demonstrated to be tolerable.

A.1.2 Implicit Surfaces

This is a graphics method where the object is represented by points in space and an equation determining the field potential around those points[19].

$$f(x, y, z) = \sum_{i=1}^{n} f_i(x, y, z) \tag{A.1}$$

The set of points in space with a certain potential value are the level set or implicit surface.

$$\Sigma = \{M(x, y, z) \in \mathbb{R}^3, f(x, y, z) = T\} \tag{A.2}$$

These can be added or subtracted to produce the object and further refinement can be done using functions such as blend or warp. As well as points other field sources such as lines and 3D forms can also be used. In this way it can be considered to be constructed from primitives although the way the primitives interact is much more complex than those previously mentioned.

A.1.3 Crusher Method

The crusher in this method was a cone style crusher designed for crushing rocks. This is a revolving head in a matching bowl that is inclined at an angle to produce the crushing effect.

The variables that the program evolved were; the shape of the bowl and crusher head, the rotational speed, the angle of the head and the closest point of the head and bowl. The program was able to simulate every design and the next generation were selected based on the size distribution of the output, the power requirements and the capacity of the crusher.

The shape of the bowl and head were modelled as a 2D cross–section to simplify the problem. This 2D line was represented by a vector of coordinates that were joined together by straight lines to create the final shape.

The initial population was filled using basic designs similar to those currently in use. After 200 generations the new crushers had a simulated capacity of 140% of the original.

A.2 Details of Birkhoff's Aesthetic Measures

A.2.1 Polygonal Form Aesthetic Measure

Here complexity is defined as the number of straight lines that form the shape. This is not the same as the number of sides as two sides that are part of the same line count only once.

The order is made of 5 parts:

- V - The vertical symmetry
- E - The optical equilibrium, or the impression that the shape is able to stand up
- R - The rotational symmetry
- HV - The observation that the lines are part of a regular grid
- F - A negative value that is used as a penalty if the angles or edges are too large or small or other irregularities exist

A.2.2 Vases

The characteristic points that describe a vase are defined as follows [6]:

- The points at the base and rim.
- Points where the tangent to the curve is vertical, e.g. points of local maximum and minimum width.
- Points of inflection where the curvature of the curve is 0, e.g. the points where the curve changes direction.
- Points where the curve is not continuous, e.g. where there is a sharp change of direction.

The sum of these points is the value of the complexity of the vase. The order is given by the sum of four components:

$$O = H + V + HV + T \tag{A.3}$$

Where H, V, HV and T are given by:

- H - The number of independent horizontal distances that are in ratios 1:1, and 1:2. Horizontal distances are given by the values of the widths at each characteristic point. This is capped at 4.

- V - The number of independent vertical distances that are in the ratios 1:1 and 1:2. Vertical distances are measured between two characteristic points and the two that are in a ratio must share one point. This is capped at 4.

- HV - The number of independent ratios between horizontal and vertical distances. The distances are as above and characteristic point of the horizontal distance must be the same as one of the points for the vertical distance. This is capped at 2.

- T - The number of independent tangents to the characteristic points that are parallel, perpendicular or cross the central axis level with one of the characteristic points. Vertical tangents at the end points are also counted. This is capped at 4.

By independent we mean that if distances A and B are in a ratio and B and C are in a ratio then if a ratio between A and C exists then it will not be counted as it is not independent.

The independence requirement reduces the amount of order that a new complexity point can add. For example in figure A.1 the top 3 points have one vertical ratio between them (first solid arrows), adding the lowest point would add one independent ratio (second solid arrows) and two further dependent ratios (dashed). If you allow dependent ratios then

Figure A.1: Independent Ratios (Solid) and Dependent Ratios (Dashed)

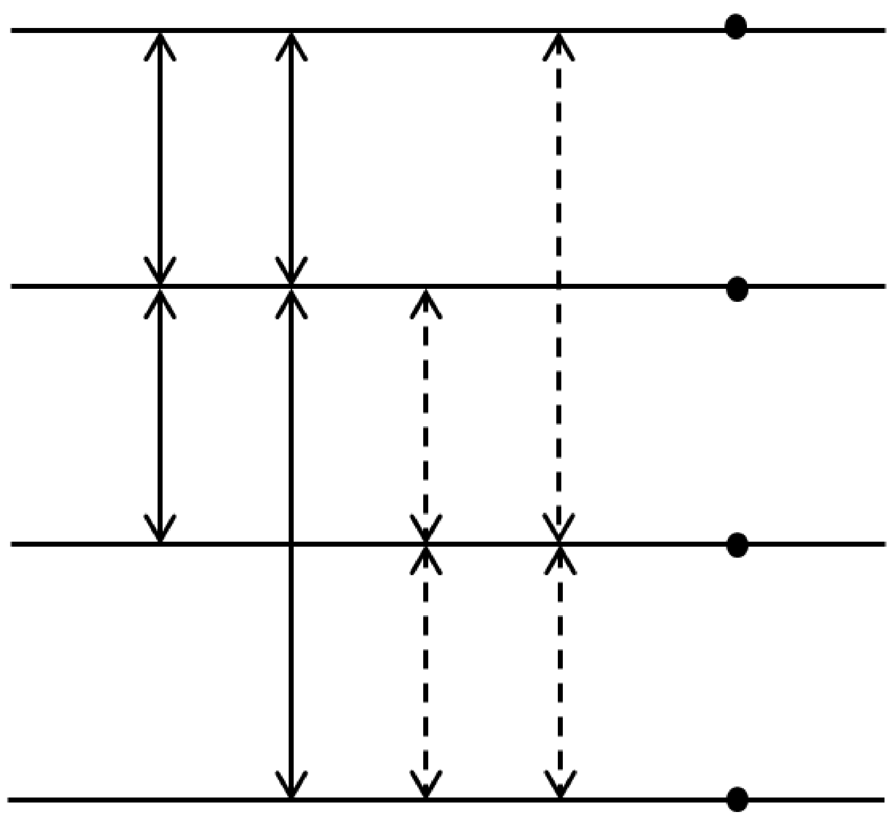

adding this point would vastly increase the score, strongly encouraging complex designs, using only the independent values slows this increase.

The capping ensures that eventually adding new points has no order value, but the complexity will increase. This reduces the value of over–complex designs.

Without the capping and independence it would beneficial to have a very high complexity as the new significant point would only require one new order point to achieve an overall increase in the measure.

Birkhoff also specified certain conditions that were to be fulfilled in order for the measure to be applicable. They are (as quoted from Birkhoff):

- *'The breadth at the base must be at least one eighth of the height.'*

  This was to ensure that the vase would be stable.

- *'The maximum breadth should be attained but once, between the ends of the vase. This should be at least one quarter of the height.'*

  This was to ensure a sufficient volume to the vase and to make it easier to carry.

- *'The minimum breadth should be attained in the upper half of the vase and should be at least one eighth the height.'*

  This was to give the vase a neck to allow it to be poured. He stated that it also suggested why a vase design could not be inverted and still remain functional.

- *'The height of the vase should be at least as great as its breadth'.*

  This was to define it as a vase rather than a pot.

- *'Rectilinear parts of the contour are only permitted at the ends, of lengths not more than one fourth of the breadth there, and directed so that the breadth is not decreasing towards the end in question.'*

  The reasoning behind preventing rectilinear parts was not elaborated on but they were allowed at the ends as a way to protect them.

- *'The characteristic tangents at the foot should be inclined at an angle of at least* $45°$*.'*

  The direction of the angle was not specified and although the justification states that the vase would 'bulge out' if this was not satisfied, the purpose was apparently to protect the base and therefore this rule is presumably to prevent the existence of sharp edges at the base.

A.3 Screen Layout Details

This gives the details of the four order measures used in the screen layout measure by [32].

- Balance - Describes how the features of the screen are distributed around the centre point. The area of each feature is multiplied by its distance from each axis and the overall balance is found. A high–scoring layout has an equal weight either side of both axes.

- Equilibrium - A similar measure that locates the centre of mass of the features and relates this to the actual centre, a high score is awarded if the two centres coincide.

- Symmetry - Found by assessing the weights of the four quadrants created by the vertical and horizontal axes. Each of these weights are then compared with the other three with the score increased if they are equal.

- Sequence - Also looks at quadrants. For each quadrant a value is given if the quadrant is not empty. This value is 4 for the top left, 3 for the top right, 2 for the bottom left and 1 for the bottom right. This reflects the way the eye scans the screen and favours information given in the top left.

Appendix B

Appendix: Preliminary Studies Further Information

B.1 Representation

This section contains plots to support the representation preliminary studies section.

B.1.1 Parametric Representation

Parametric equations are a very simple way of representing a 2D curve. The curve is described by two functions of a variable t. These represent the x and z coordinates of the curve along its length (z is used as it will represent the height of the vase).

Initial tests looked at a range of primitive functions and combinations of these. They also looked at the impact of fixing the z function compared with allowing this to vary. Fixing the z function reduced the design space which may prevent non–viable designs or restrict the space too far.

Initially only two primitives were used in each example e.g. $\cos(t + 2)$ used cos and $+$. These were chosen from the following functions; $+$, $\times$, cos and e. $-$, $/$, sin and log were not used as they could either be approximated by combinations of the others or they consistently produced invalid solutions (such as ∞ and complex numbers). The scalars used in $+$ and $\times$ were randomised. The two functions were described by a genotype of two integers. For example the function $x = \cos(e^t)$ was designated the genotype $[4, 3]$ as the exponential function is designated 4 and is applied first and cos is 3 and applied second. $+$ and $\times$ are designated 1 and 2 respectively.

During initial trials it was found that functions containing the exponential often failed. Therefore function 4 was changed to e^{-t}, this reduced the number of failures although they were not entirely prevented. It was also observed that many of the functions would give both positive and negative output which would cause the vase to cross over itself and if the 0 point was at the base the vase would not stand. For this reason the smallest value is deducted from every point and a positive number (here 1) is also added to define the narrowest point of the vase. This also ensured that the rim of the vase was open allowing the vase to hold water and flowers.

All combinations of the primitive functions were studied and were capable of producing viable vases. The only combination that was observed to produce a bad result was $e^{-t \cdot const}$ which occasionally produced a flat disk of radius of order 10^4. This was presumably when the constant was negative. Further adjustment to the e function may be required. Examples of these results can be seen in figures B.1 to B.4.

The second stage involved allowing different functions for the z coordinate as well. The same method was used as for the x coordinate above. Less viable vases were observed, but there were some interesting variations with bowl and plate–like forms appearing. This suggests that fixing z would be a safer option for evolving vases but allowing variable z could open up the program to the evolution of a wider range of ceramic forms (figure B.5).

To explore the effect of genetic programming on parametric equations, different programs were used to see the effects of both mutation and recombination. To get a clearer picture of the effects of the genetic operators a larger equation was used. To simplify the problem fixed z was used. Ideally variable–length trees should be considered for evolving the parametric equations. However Matlab is not suited to this task and to simplify the problem fixed–length trees were used with a matrix genotype.

Figure B.1: Results using $f(t + const)$

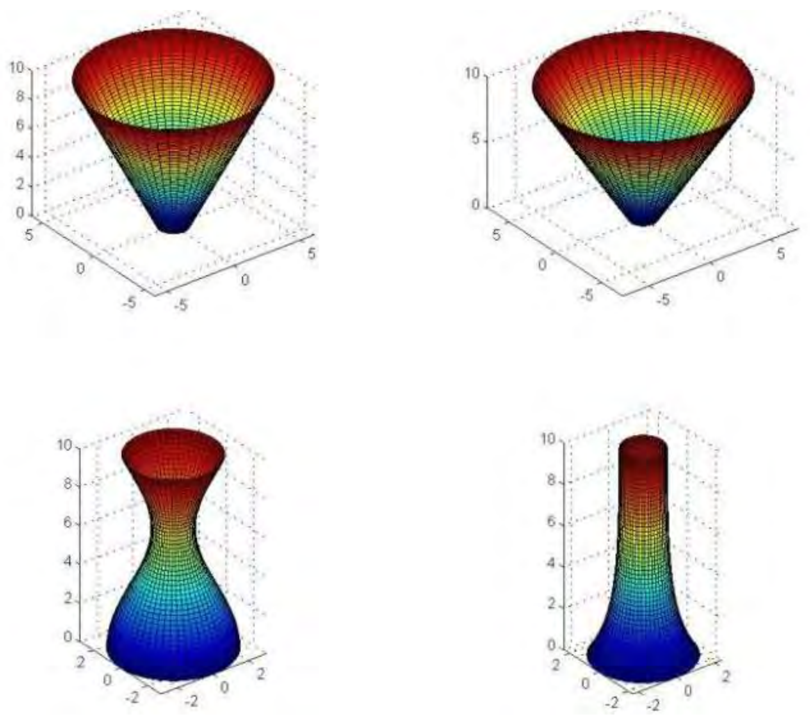

Figure B.2: Results using $f(t \cdot const)$

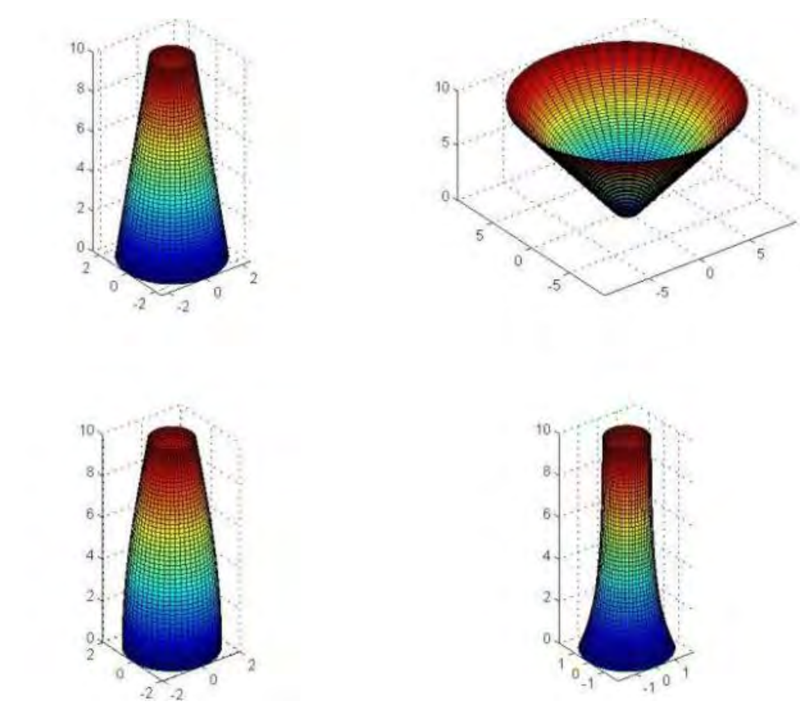

Figure B.3: Results using $f(\cos t)$

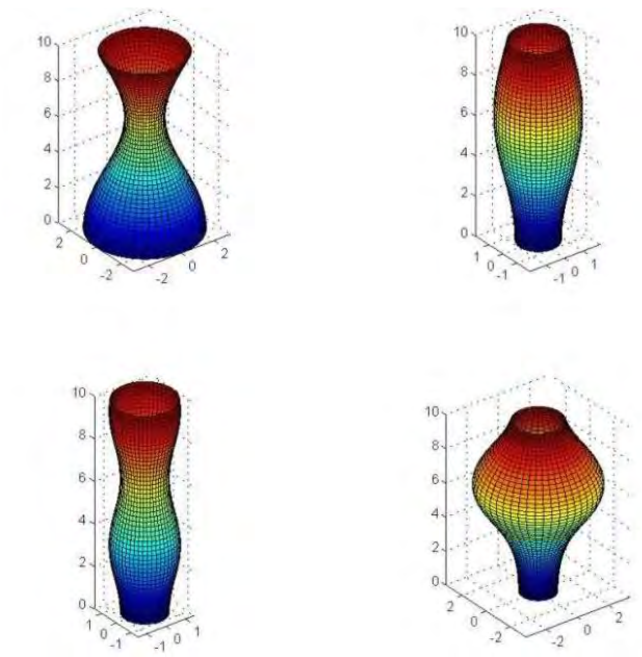

Figure B.4: Results using $f(e^t)$

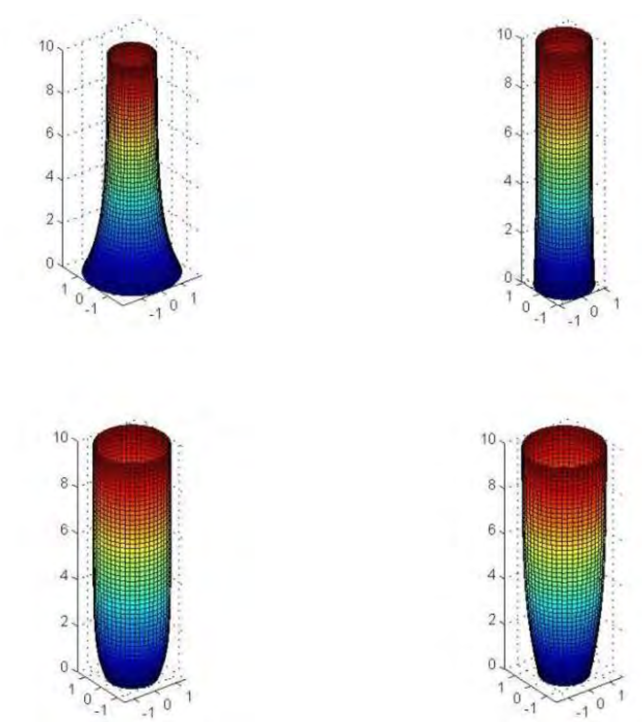

Figure B.5: Selection of Results using Variable z

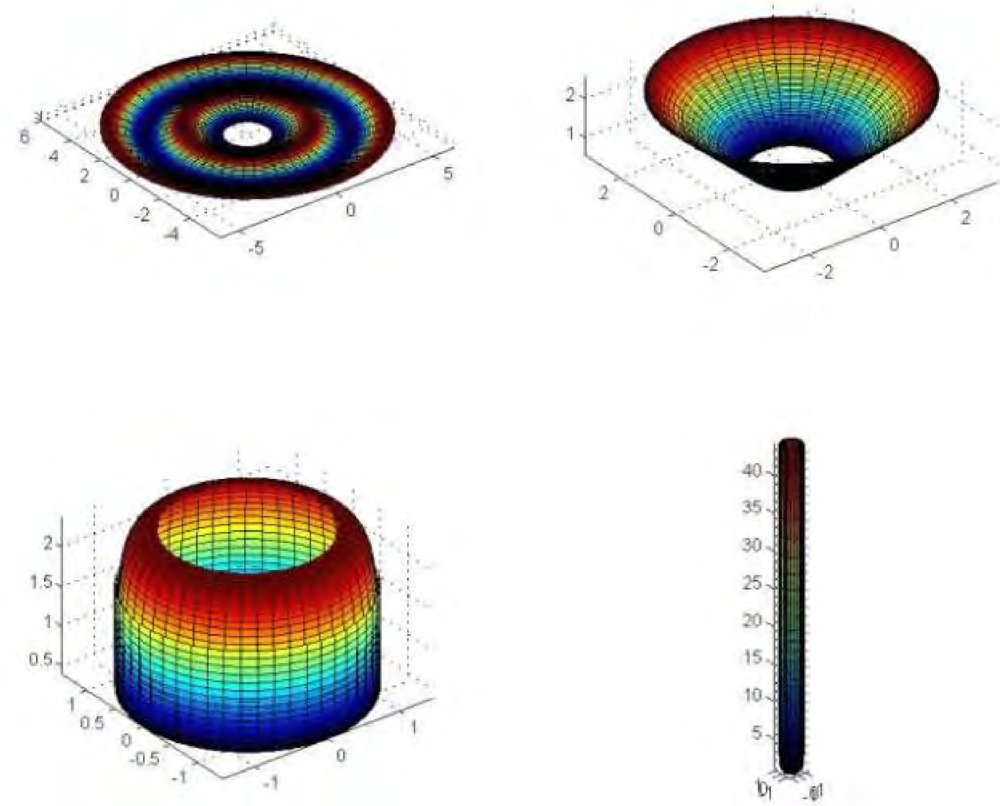

Figure B.6: Tree Structure of Equation: $\cos(4+t)+e^t\cdot(7+3)$

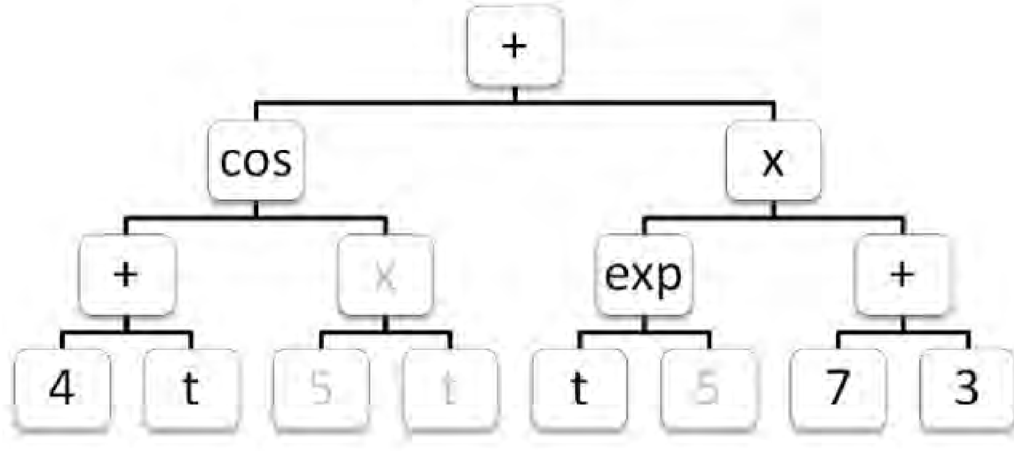

Table B.1: Matrix Structure of Equation: $\cos(4+t)+e^t\cdot(7+3)$

0	1	2	3	3	2	6	6	1	9	10	10	9	13	13
1	3	1	0	0	2	0	0	2	4	0	0	1	0	0
0	0	0	4	NaN	0	5	NaN	0	0	NaN	5	0	7	3

The length of the representation tree was fixed at 4 full levels. When using the exponential and cos functions only the first of the two variables that were in the lower level would be used. An example of a tree like this is shown in figure B.6 with the unused branches shaded out. The tree was represented in the computer by a matrix such as the one in table B.1, where the top row signifies the form of the tree, the second shows the type of functions used and the third row gives the variable where NaN is replaced by a vector of points (t).

Recombination was carried out by replacing part of one genotype matrix with the corresponding part of another. Figure B.7 shows an example output where a crossover is applied. The parents are the top three vases and the children are the lower three. The last child's parents are both 1 and as expected it is identical. Child 2 is the offspring of parents 1 and 3. The crossover point is 5, this is one of the terminal points on the first branch of the tree and therefore will only change that one value. Changing this one value may have very little impact unless it was a vector changing to a scalar and vis–versa, or a very significant change of scalar value. If the function on the node above was a *cos* or *exp* then it would not

Figure B.7: Parents and Offspring by crossover of Parametric Equations

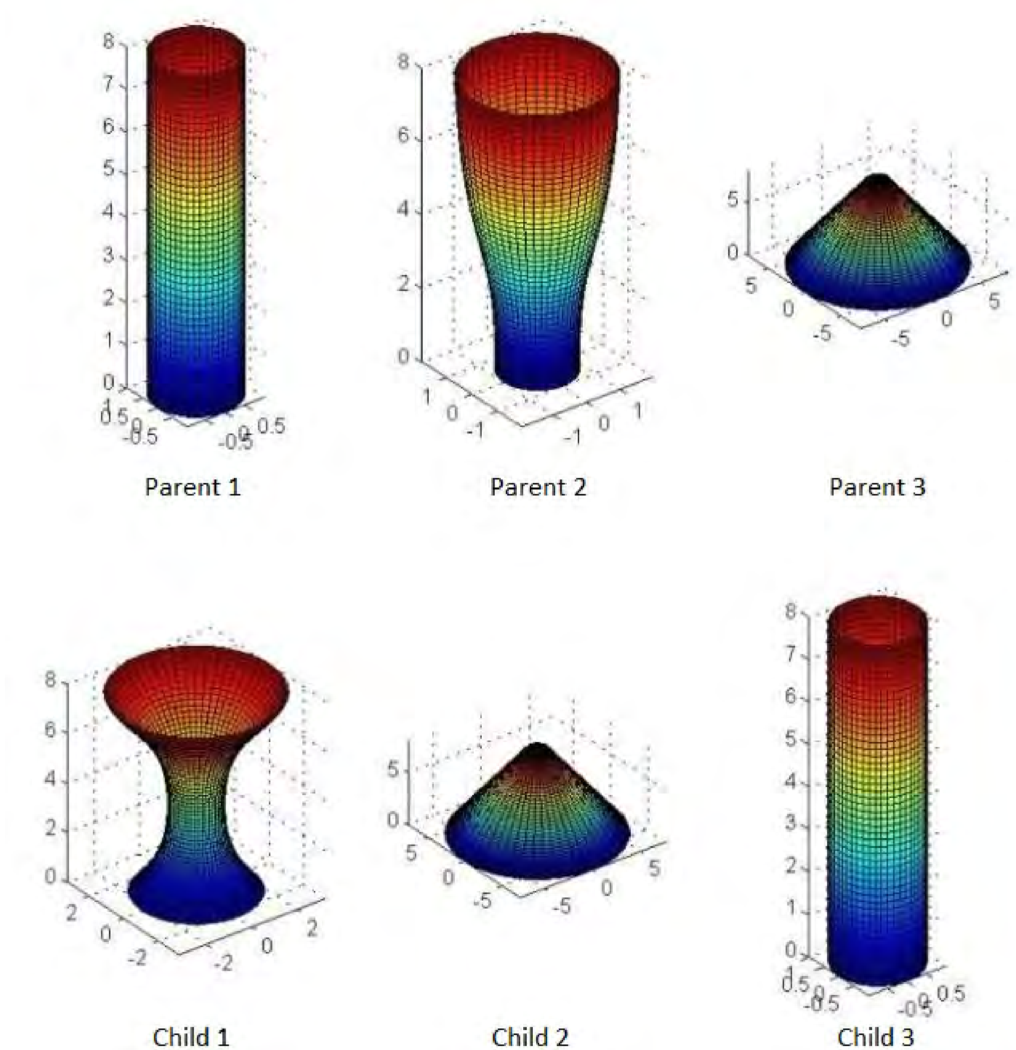

even be considered. Here it appears to have made no difference at all and the child looks identical to parent 3.

Child 1 is the offspring of parents 1 and 2 and has a crossover point of 2, meaning that the entire first branch of the tree will be replaced. This has had a very significant impact. The child bares little similarity to either parent. This potential level of variation could be too large.

Two different types of mutation were also tried, the mutation of one of the functions and the mutation of one of the scalar values. To test the two different mutations a single

Figure B.8: Parent and Offspring by Mutation of Parametric Equation Functions

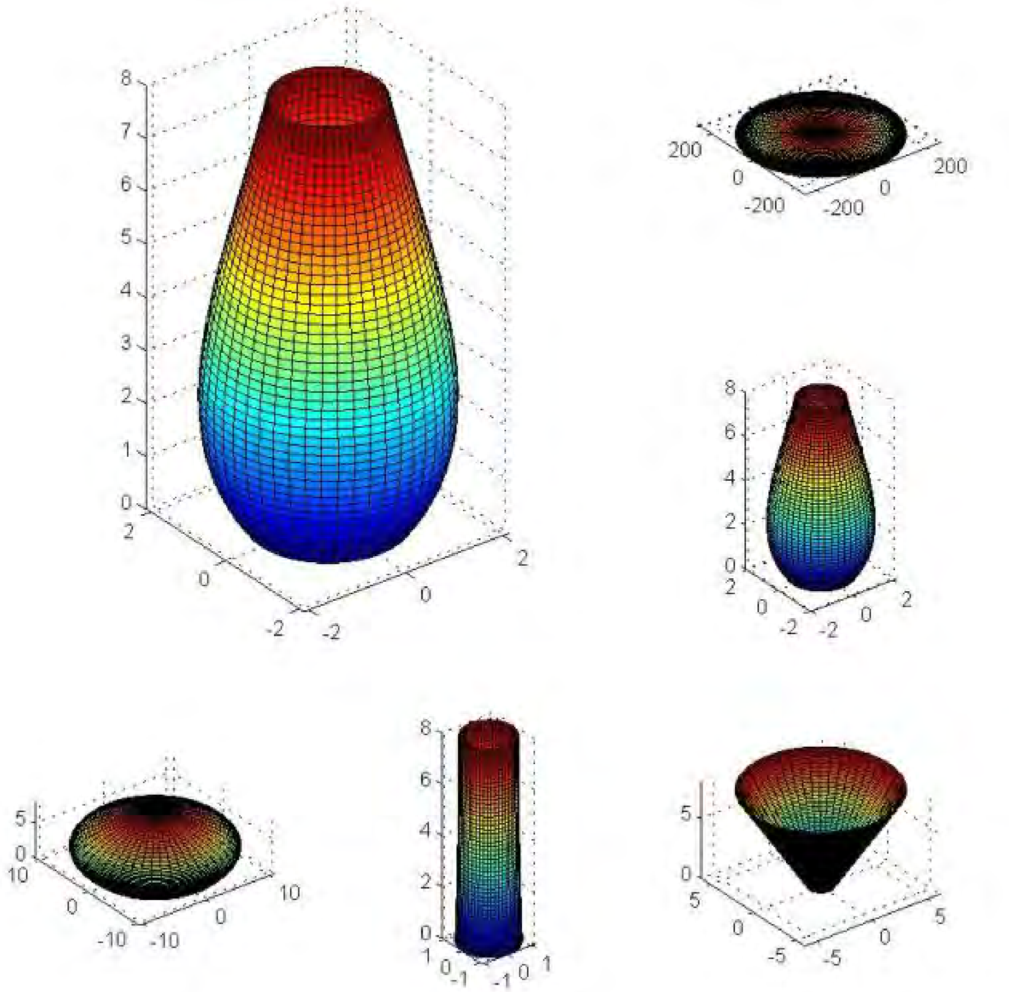

parent was used and five offspring produced each with a random mutation. Figure B.8 shows an example where only function mutation is used. The figure illustrates the severity of the changes a mutation of the function can create. Other results would produce little or no change but a range of severe changes would be more common. Very few results would provide a meaningful change where the child would resemble the parent but be visually different as well.

Figure B.9 gives a typical example of the effect of mutating a random value by addition of another random value. The results usually would be very similar to the parent with little or no visual change. More significant changes such as the middle child on the bottom row would appear rarely but would often be an interesting change.

It was observed that some parents seemed more stable than others, for example a vase with simple vertical sides would not produce any variation in its offspring in any of the trials that were seen but with other forms every child would be different, often in interesting ways yet with inherited traits from the parent. Therefore this type of mutation could find some uses in the evolutionary process with suitable constraints.

It was felt that the erratic behaviour of the crossover and mutation was too great for it to provide a viable representation. Parametric equations were rejected and another representation was found.

Figure B.9: Parent and Offspring by Mutation of Parametric Equation Values

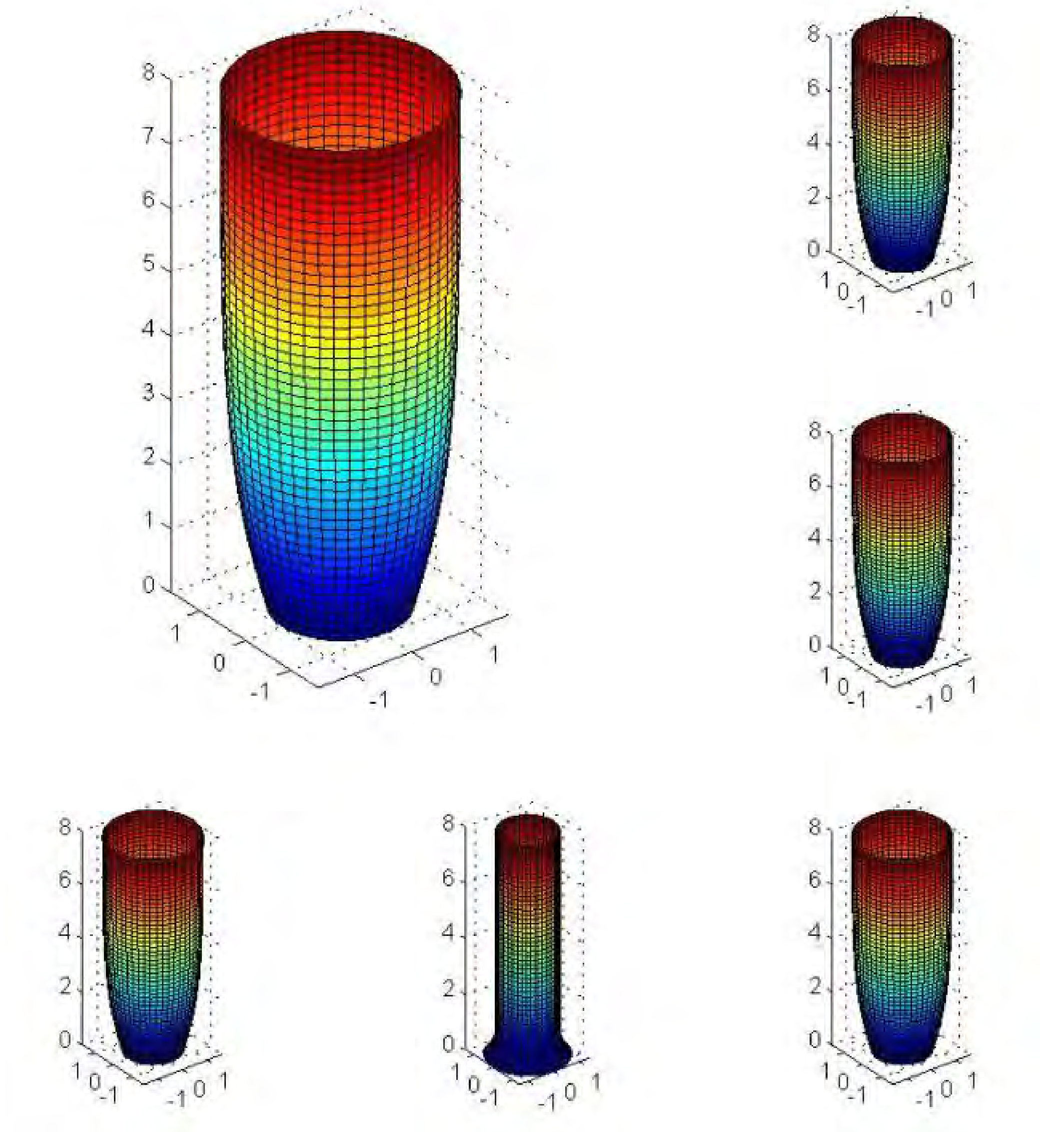

B.1.2 Bezier Curves

B.1.2.1 Demonstration of effect of number of Bezier coordinates, z is fixed:

Figure B.10: Increasing sets of Coordinates from 2–10

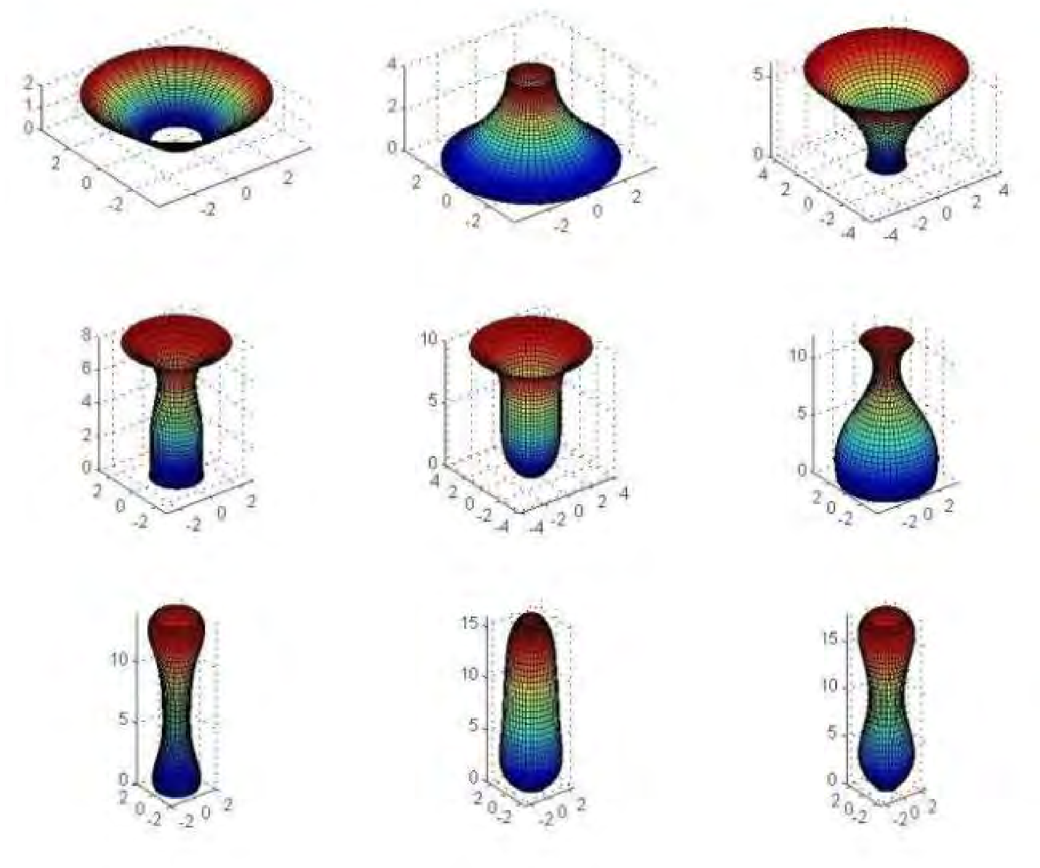

Figure B.11: Examples of Bezier Curve Vases with 5 sets of Coordinates

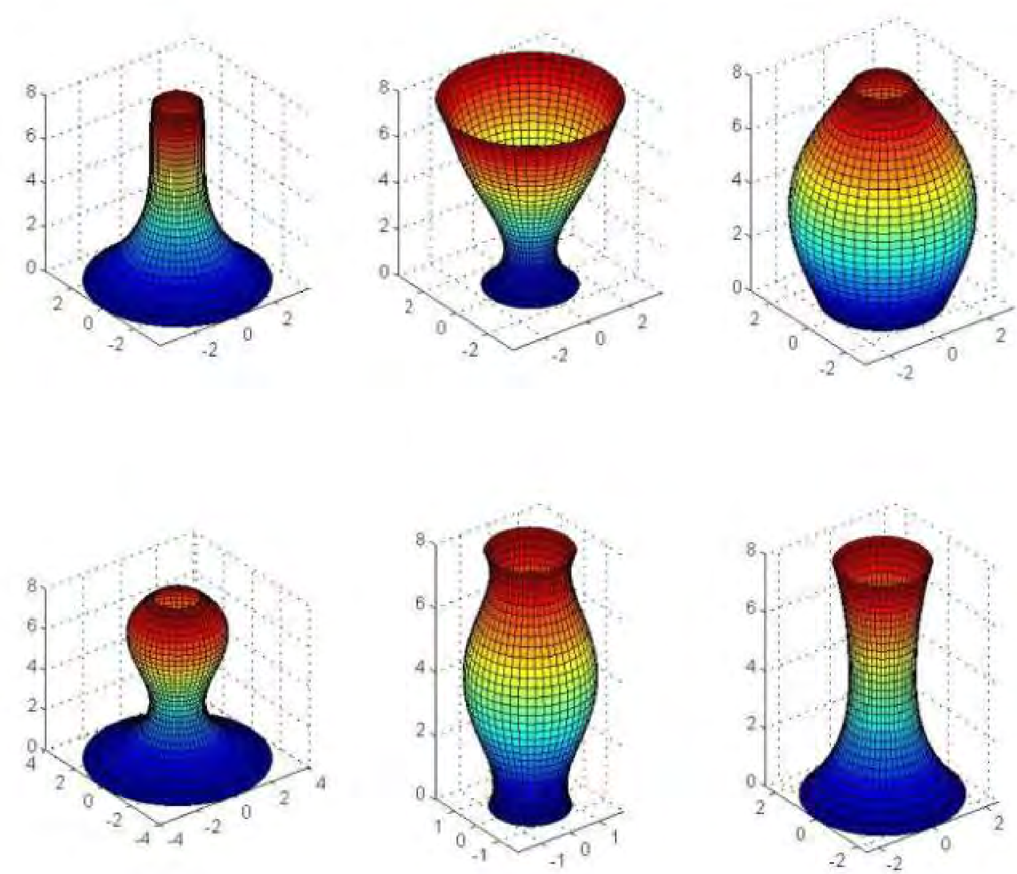

B.1.2.2 Allowing Variable z

An expectation of 2 seemed to be ideal for vases (figure B.13) but experimentation showed that reducing the mean to 1 would increase the number of bowl like shapes appearing (figure B.14). Using an expectation of 0 could create plates although many designs appeared to not be viable (figure B.15).

Figure B.12: Examples of Bezier Curve Vases with randomised z Coordinates

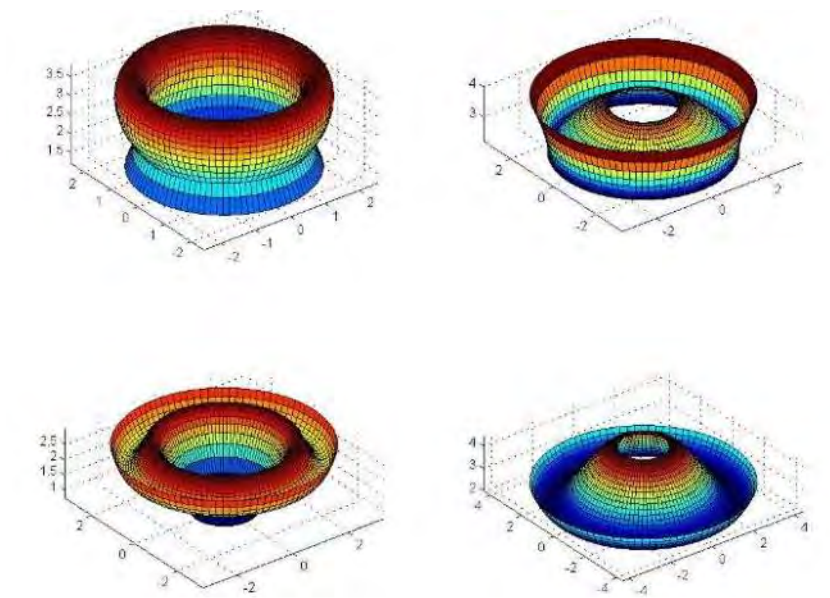

Figure B.13: Examples of z increasing with mean 2

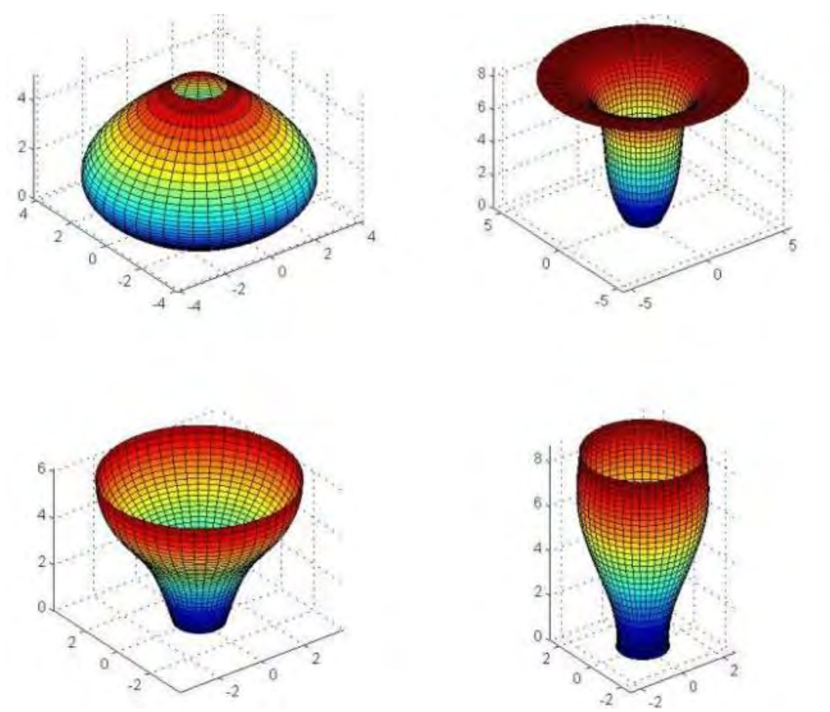

Figure B.14: Examples of z increasing with Mean 1

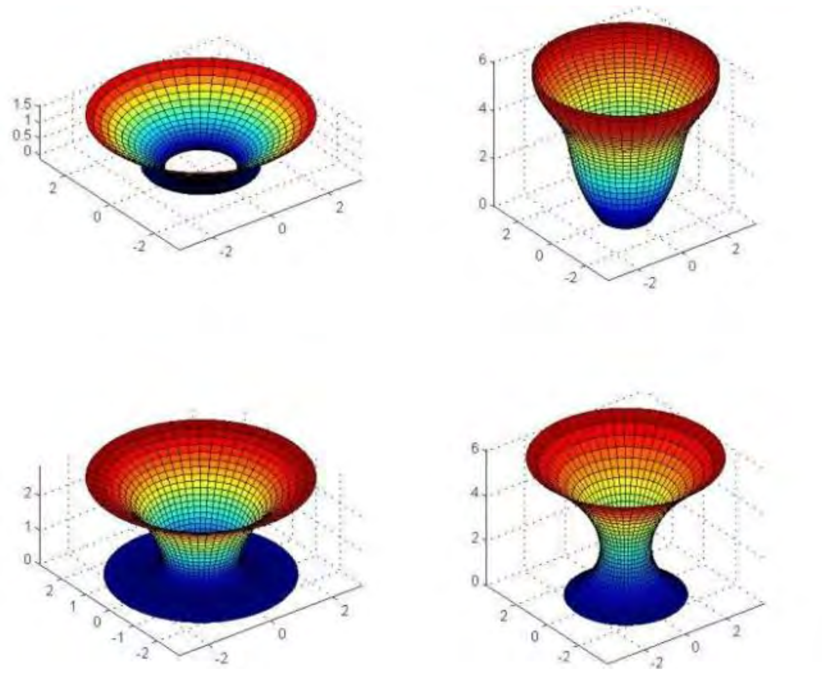

Figure B.15: Examples of z increasing with Mean 0

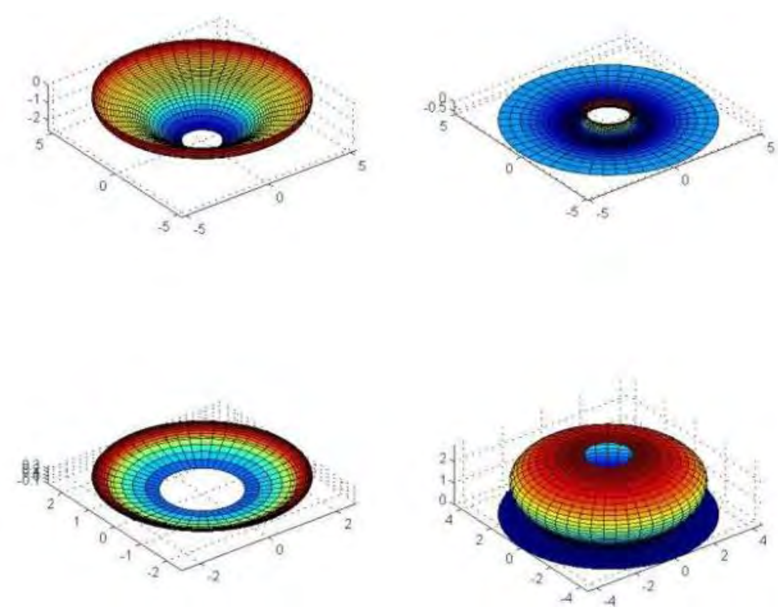

B.1.2.3 Effect of Recombination on Bezier Curves

Figure B.16: Parents and Offspring by Splicing Bezier Curves

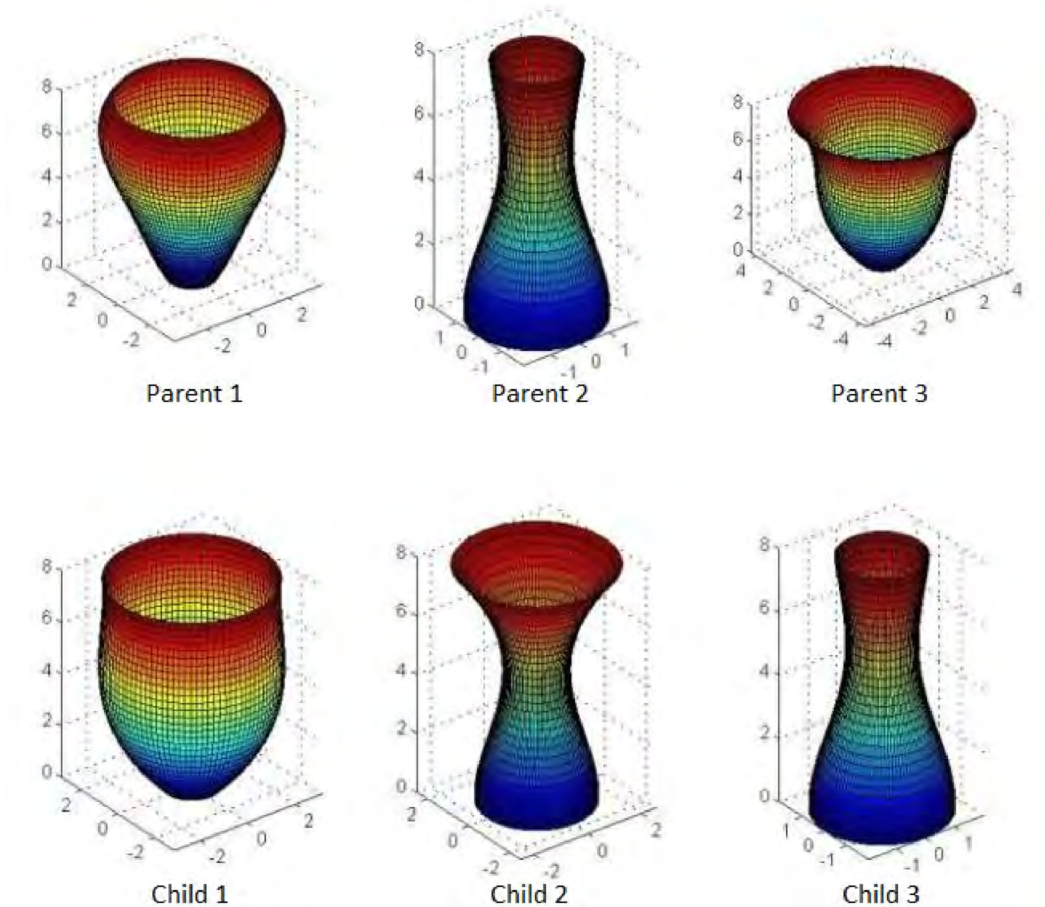

Figure B.16 shows the effect of splicing. This proved a very effective method as the child would have similarities to one or both of its parents. All the children in this figure were children of parents 2 and 3. Child 1 has the curved base of parent 3 and straight top of parent 2, while child 2 inherited the exact opposite traits from the two parents. Child 3 meanwhile appears identical to parent 2.

As genetic algorithms lend themselves to several different recombination options, an average was also considered. This also proved to be good at passing on phenotype properties as can be seen in figure B.17. For example the average of the tall thin parent 1 with the curved parent 2 gave the tall thin curved child 1. This allows the inheritance not of clear

Figure B.17: Parents and Offspring by averaging Bezier Curves

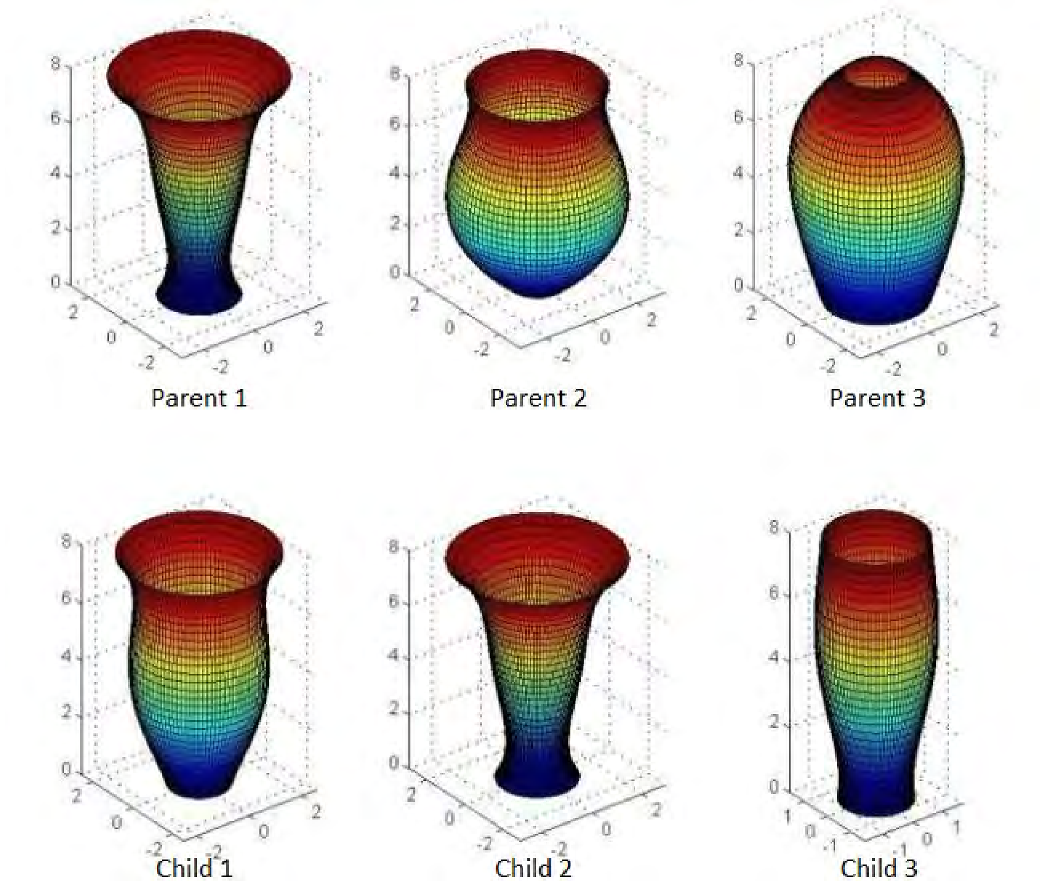

sections of the vase but of more general traits. The danger is that the combination of two very different interesting vases may produce a rather neutral child as has happened with child 3, its parents were numbers 1 and 3 which curve in almost the opposite directions and therefore the child has few distinct features. Child 2 had only one parent, parent 1 and as expected it is identical.

B.2 Aesthetic Measures

B.2.1 Initial Designs and Results

Form	Fig.	Uncapped	Capped	Significant Points
1	B.18	0.6091	0.5956	5
2	B.19	0.8126	0.8126	4
3	B.20	0.4521	0.4521	5
4	B.21	0.8523	0.6252	5
5	B.22	0.5131	0.4866	5
6	B.23	1.2010	0.8052	5
7	B.24	0.8976	0.5606	6
8	B.25	0.0603	0.0603	3
9	B.26	0.6532	0.4330	5
10	B.27	0.0295	0.0295	3
11	B.28	1.1780	0.7334	7
12	B.29	0.8664	0.7334	5
13	B.30	0.6470	0.6470	4
14	B.31	1.2720	0.6864	6
15	B.32	0.4801	0.4801	5
16	B.33	0.4683	0.4683	5
17	B.34	1.5111	0.9019	6
18	B.35	0.7005	0.5614	6
19	B.36	0.8703	0.8468	6
20	B.37	0.0705	0.0705	3
21	B.38	0.3586	0.3586	4
22	B.39	0.1666	0.1666	3
23	B.40	0.6237	0.5062	5
24	B.41	0.3885	0.3885	5
25	B.42	0.3460	0.3460	5
26	B.43	0.8351	0.6541	5
27	B.44	0.7885	0.4657	6
28	B.45	0.1158	0.1158	3
29	B.46	0.7782	0.6088	5
30	B.47	0.3579	0.3579	3

Table B.2: Aesthetic Measure Scores and the Number of Significant Points

Form	Fig.	Normal (N) Orientation Scores	Av.	Inverted (I) Orientation Scores	Av.	Max	N or I
1	B.18	3 2 3 2 1 1 1 1	1.750	6 7 5 3 4 7 6 5	5.375	5.375	I
2	B.19	5 5 6 7 6 6 5 8	6.000	4 7 7 3 5 6 4 7	5.375	6.000	N
3	B.20	8 8 8 6 7 5 6 9	7.125	6 7 5 3 4 7 6 5	5.000	7.125	N
4	B.21	5 3 8 5 4 2 1 2	3.750	4 7 7 3 5 6 4 7	7.250	7.250	I
5	B.22	4 6 3 5 6 4 2 6	4.500	2 2 4 2 4 6 5 2	3.375	4.500	N
6	B.23	8 5 6 7 7 1 1 2	4.625	3 2 4 2 2 1 2 7	2.875	4.625	N
7	B.24	4 2 5 2 3 2 1 2	2.625	6 8 7 5 7 5 6 6	6.250	6.250	I
8	B.25	7 6 9 7 6 8 8 8	7.375	7 6 8 5 6 8 8 1	6.125	7.375	N
9	B.26	5 5 3 5 5 4 3 3	4.125	2 1 2 1 2 1 1 1	1.375	4.125	N
10	B.27	7 7 5 5 4 8 8 8	6.500	3 3 3 1 1 1 1 2	1.875	6.500	N
11	B.28	1 2 7 2 2 5 2 3	3.000	4 4 4 1 3 7 3 4	3.750	3.750	I
12	B.29	7 8 6 4 7 7 8 7	6.750	5 5 5 2 2 4 1 2	3.250	6.750	N
13	B.30	4 4 3 3 2 4 3 3	3.250	6 6 5 7 6 7 7 2	5.750	5.750	I
14	B.31	4 4 4 3 4 5 2 3	3.625	4 3 3 2 4 1 2 5	3.000	3.625	N
15	B.32	6 3 3 1 1 1 1 4	2.250	6 5 3 8 5 1 1 5	4.250	4.250	I
16	B.33	7 8 6 2 3 5 7 6	5.500	2 4 3 1 1 1 1 2	1.875	5.500	N
17	B.34	3 6 7 4 6 6 5 6	5.375	2 2 3 1 4 6 2 2	2.750	5.375	N
18	B.35	3 7 8 5 6 8 5 3	5.625	5 7 4 5 6 7 7 6	5.875	5.875	I
19	B.36	4 6 6 6 7 6 3 6	5.625	3 4 2 1 2 2 3 2	2.375	5.625	N
20	B.37	8 5 5 5 5 5 1 5	4.875	5 2 1 1 2 1 1 6	2.375	4.875	N
21	B.38	6 7 7 7 7 7 8 8	7.125	7 6 6 5 7 8 7 8	6.750	7.125	N
22	B.39	2 2 3 2 1 1 2 2	1.875	6 7 5 6 7 6 7 7	6.375	6.375	I
23	B.40	4 5 6 4 3 4 3 6	4.375	2 3 3 1 2 1 1 2	1.875	4.375	N
24	B.41	3 2 5 5 6 2 6 3	4.125	4 6 7 3 6 7 2 2	4.625	4.625	I
25	B.42	4 3 3 1 3 1 2 4	2.625	4 4 4 6 7 8 5 2	5.000	5.000	I
26	B.43	6 6 6 4 5 5 9 9	6.250	6 5 5 6 5 8 5 6	5.750	6.250	N
27	B.44	5 4 4 3 3 4 3 2	3.500	5 7 7 5 7 8 7 7	6.625	6.625	I
28	B.45	4 5 5 4 3 2 3 7	4.125	6 7 4 8 7 5 8 6	6.375	6.375	I
29	B.46	2 1 4 2 2 1 1 1	1.750	7 7 8 4 7 5 5 5	6.000	6.000	I
30	B.47	6 4 7 6 7 8 8 6	6.500	6 6 7 7 6 8 8 8	7.000	7.000	I

Table B.3: Human Awarded Scores, Averages and Orientation

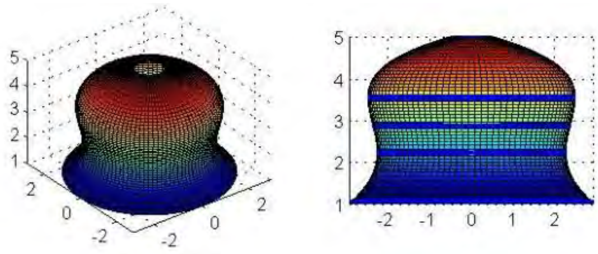

Figure B.18: Form 1

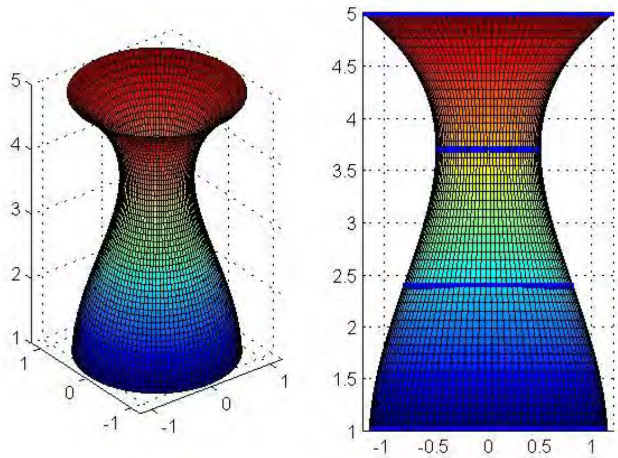

Figure B.19: Form 2

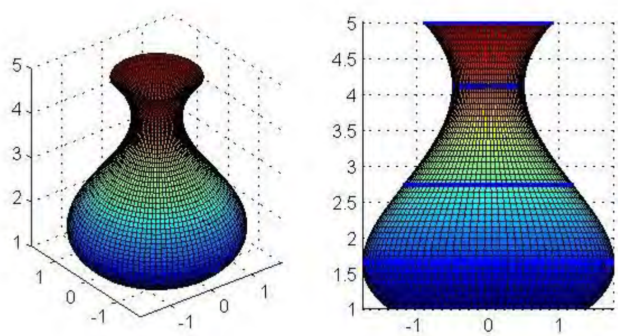

Figure B.20: Form 3

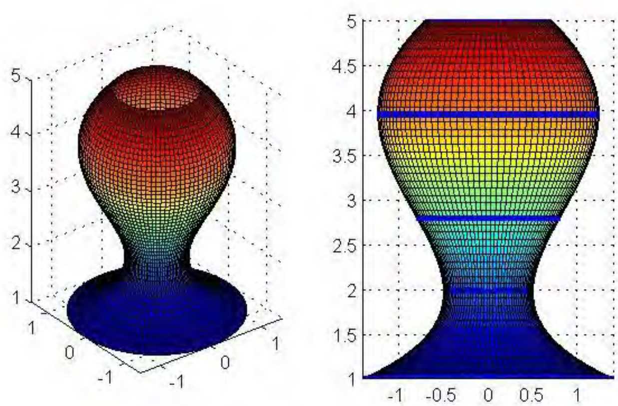

Figure B.21: Form 4

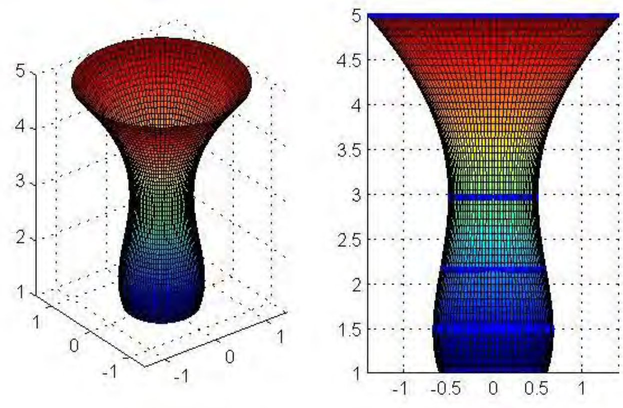

Figure B.22: Form 5

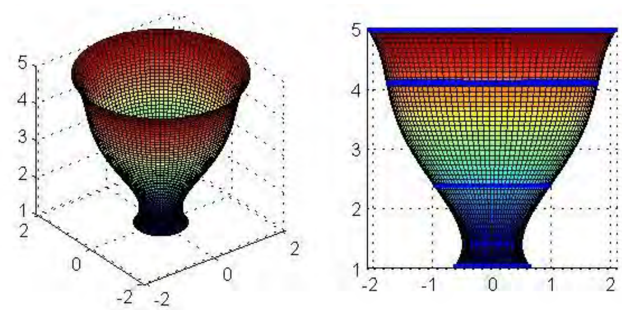

Figure B.23: Form 6

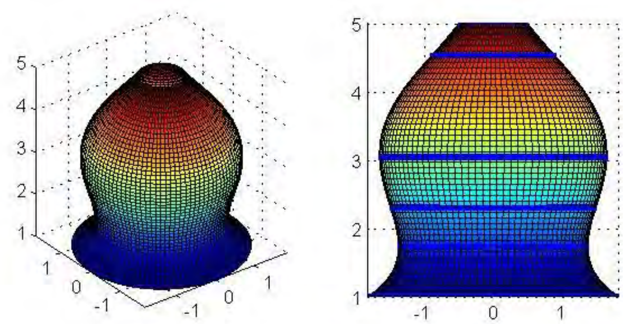

Figure B.24: Form 7

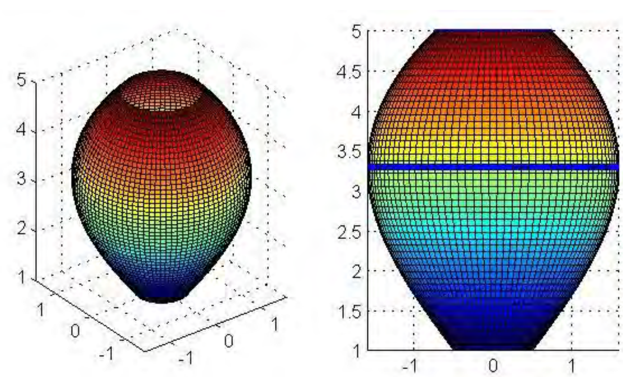

Figure B.25: Form 8

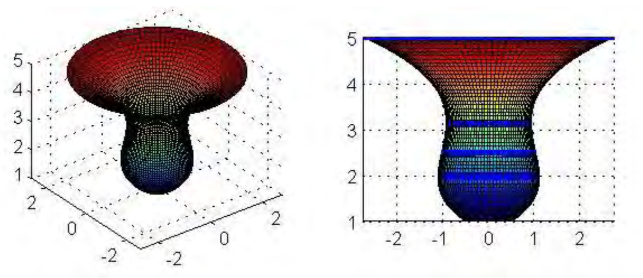

Figure B.26: Form 9

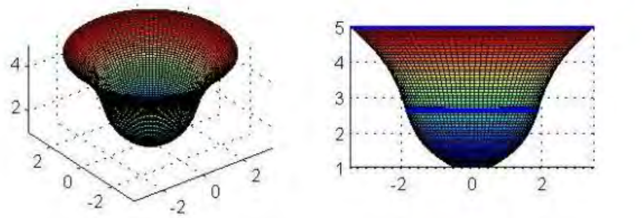

Figure B.27: Form 10

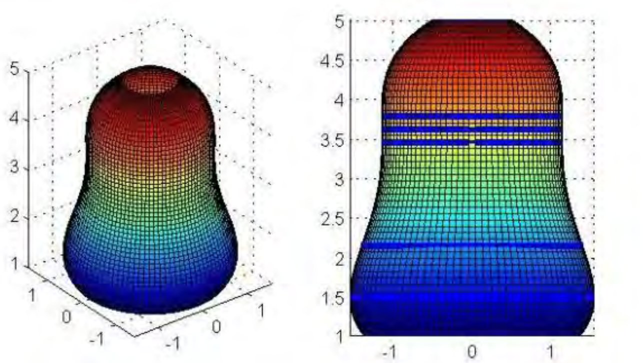

Figure B.28: Form 11

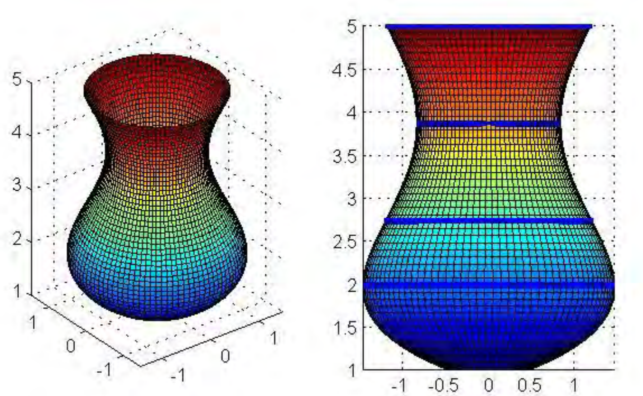

Figure B.29: Form 12

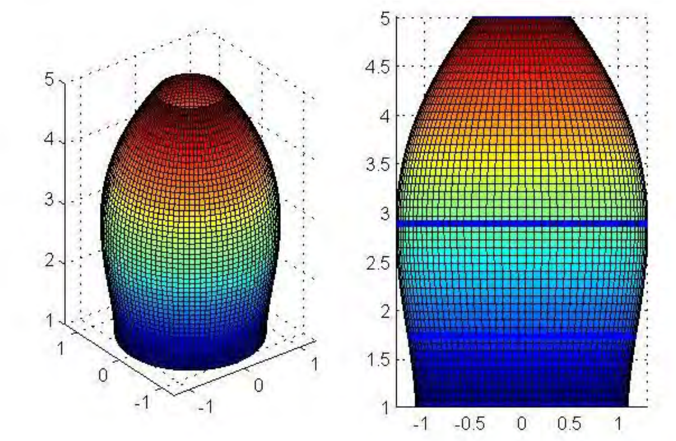

Figure B.30: Form 13

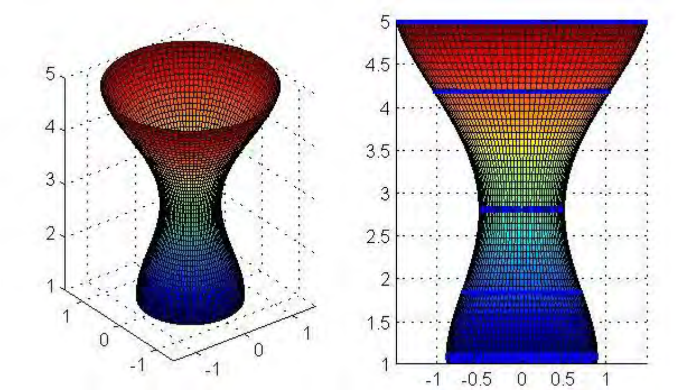

Figure B.31: Form 14

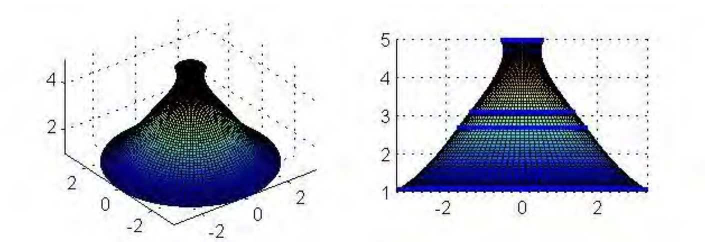

Figure B.32: Form 15

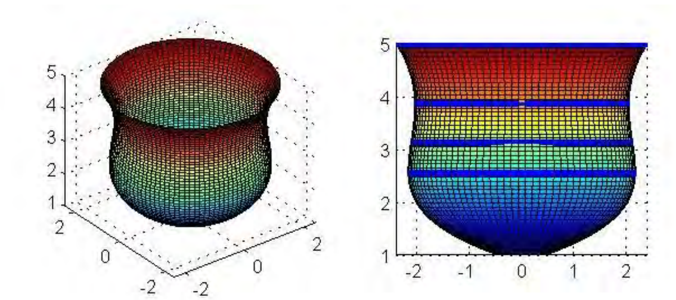

Figure B.33: Form 16

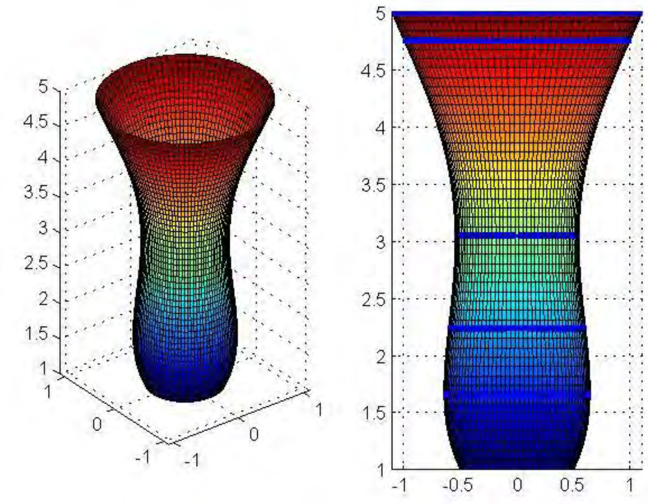

Figure B.34: Form 17

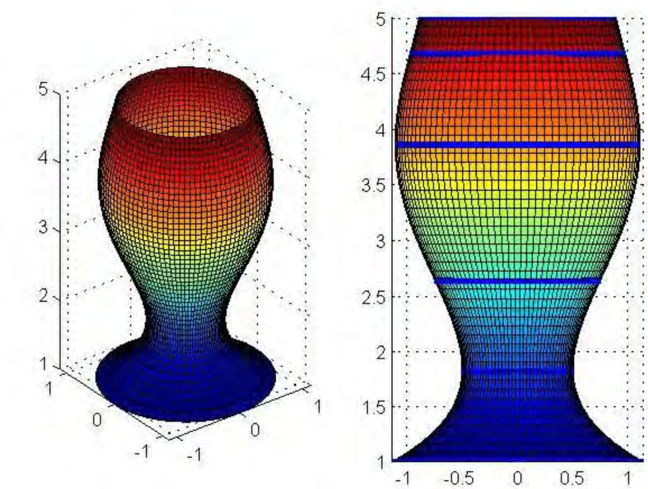

Figure B.35: Form 18

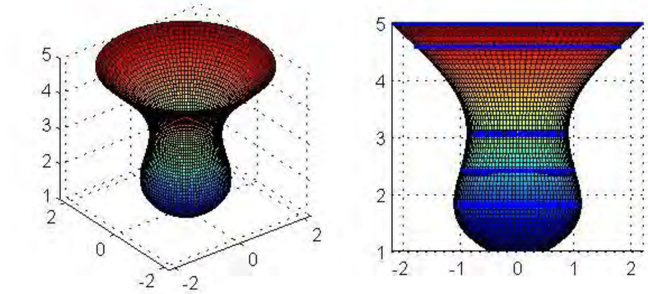

Figure B.36: Form 19

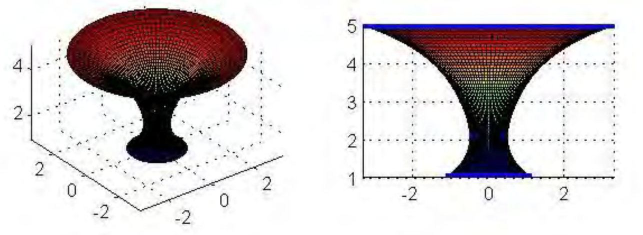

Figure B.37: Form 20

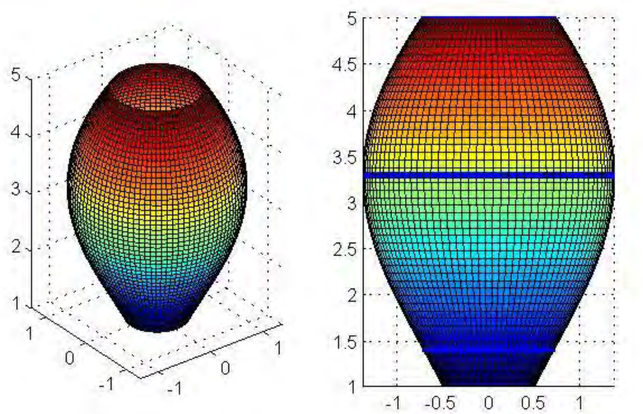

Figure B.38: Form 21

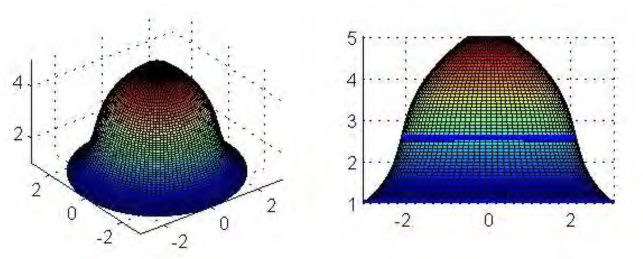

Figure B.39: Form 22

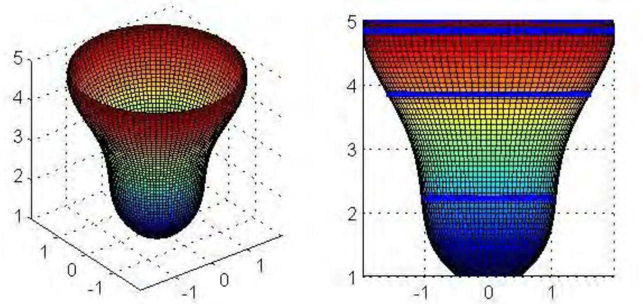

Figure B.40: Form 23

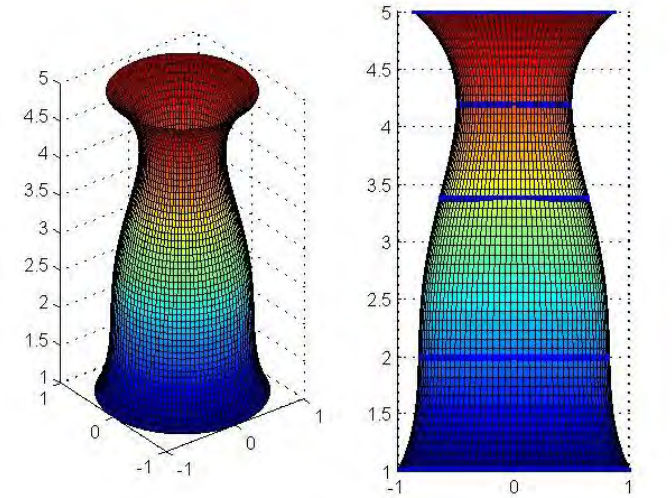

Figure B.41: Form 24

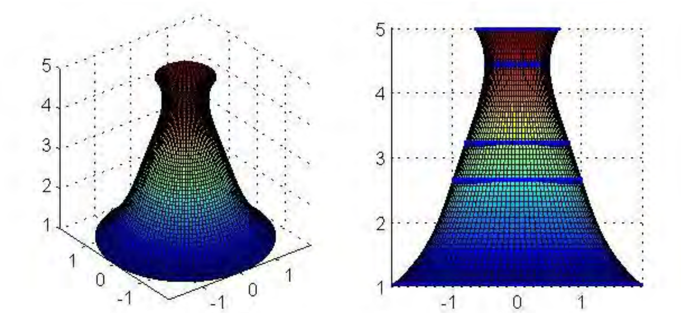

Figure B.42: Form 25

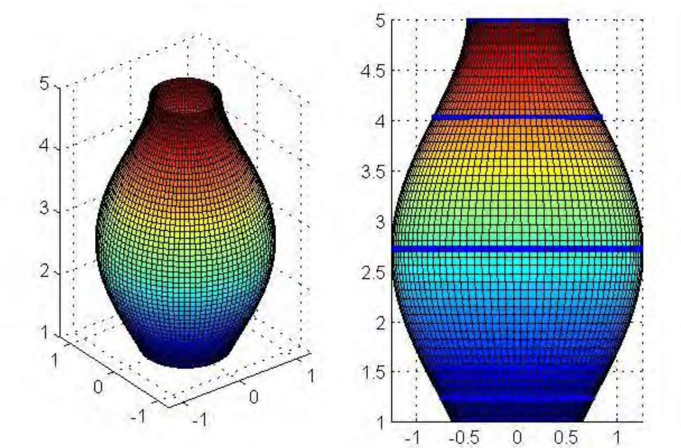

Figure B.43: Form 26

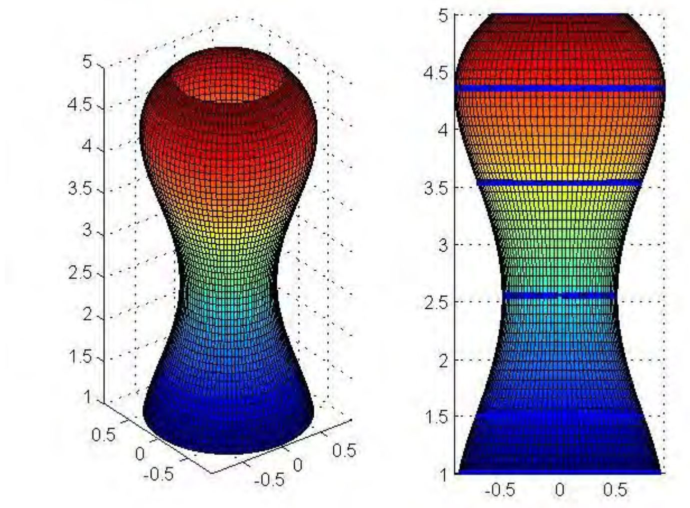

Figure B.44: Form 27

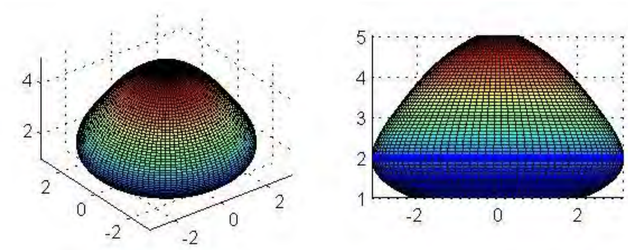

Figure B.45: Form 28

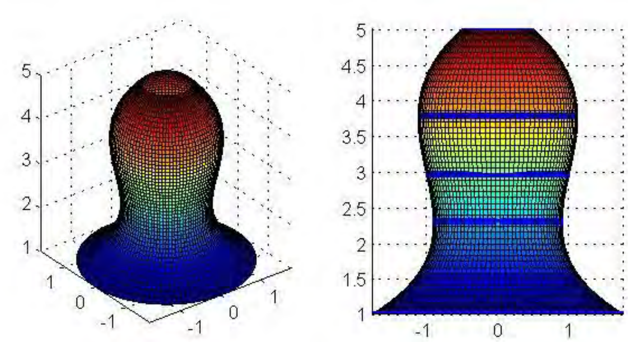

Figure B.46: Form 29

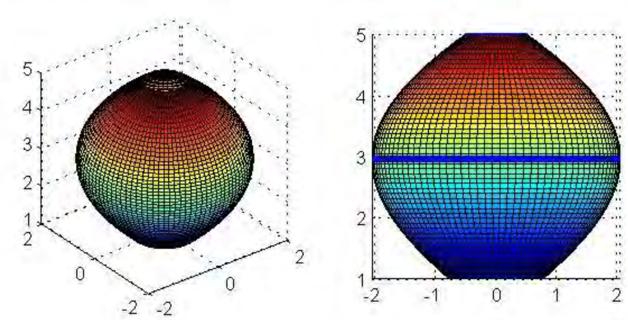

Figure B.47: Form 30

B.2.2 Aesthetic Measures

B.2.2.1 Staudek's Equations [45] and Subsequent Modifications

$$H(\epsilon) = \sum_{\substack{i,j=1..\frac{C}{2},\\ i<j}} R\left(h_i, h_j\right) \tag{B.1}$$

$$V(\epsilon) = \sum_{\substack{i=1..\frac{C}{2}-1,\\ j=1..i-1,\\ k=i..\frac{C}{2}-1}} R\left(\sum_{m=j}^{i-1} v_m, \sum_{n=i}^{k} v_n\right) \tag{B.2}$$

$$P(\epsilon) = \sum_{\substack{i=1..\frac{C}{2},\\ j=1..i-1,\\ k=i..\frac{C}{2}-1}} R\left(h_i, \sum_{n=i}^{k} v_n\right) + R\left(h_i, \sum_{n=i}^{k} v_n\right) \tag{B.3}$$

$$T(\zeta) = \sum_{\substack{t_i,t_j\\ tangents}} Q(t_i, t_j) \tag{B.4}$$

Where R and Q are given by

$$R(l_1, l_2) = \begin{cases} \max(0; 1 - 10(k\frac{l_1}{1_2} - 1)) & k\frac{l_1}{1_2} - 1 \in \langle 0; \epsilon\rangle \\ 0 & \text{Otherwise} \end{cases} \tag{B.5}$$

$$Q(t_1, t_2) = \begin{cases} \max(0; 1 - \eta/10) & t_1 \pm \eta, t_2 \pm \eta \in \tau \\ 0 & \text{Otherwise} \end{cases} \tag{B.6}$$

Where τ represents the properties of tangents as in the original measure and $k \in \{0.5, 1, 2\}$

The main change was to introduce a measure of independence when counting the number of ratios which made it more faithful to the original measure. The independence was found by adding a counter to each dimension that was being tested. When the counter reached 2 for a particular dimension any further ratios involving this dimension would not be counted. The original algorithm also included a tolerance distance for the measurement of the ratio, however as described it would only include instances where the ratio was slightly too large and not slightly too small so this was changed. The value would then be weighted by the difference between it and the exact ratio, so that the value would be 1 for the exact ratio reducing to 0 at the edge of the tolerance. This weighting was removed so that any value falling within the tolerance was given the value 1. This was due to the introduction of an independence measure as if the first ratios to be counted were inaccurate the score would be artificially low even if these distances were included in many other ratios.

Figure B.48: Edge Oscillating Circles

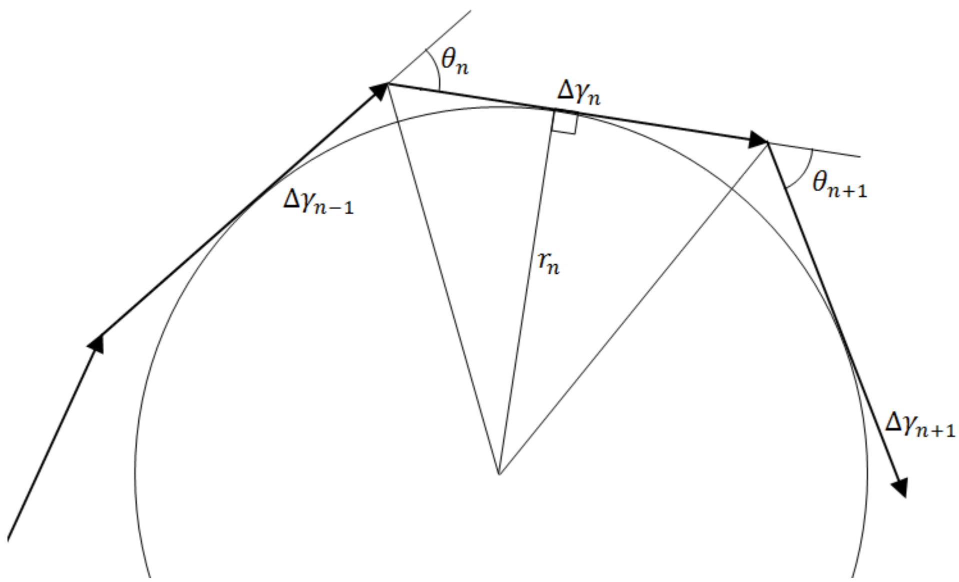

B.2.2.2 Discrete Curvature

A circle is fitted to the discrete curve and the inverse of the radius is found. Here we find the circle that is tangent to three successive edges as shown in figure B.48 and find the curvature using the equation:

$$\kappa = \frac{1}{r_n} = \frac{\tan(\frac{\theta_n}{2}) + \tan(\frac{\theta_{n+1}}{2})}{\parallel \Delta\gamma_n \parallel} \tag{B.7}$$

B.2.2.3 Curvature Analysis

The curvatures were plotted along with their corresponding designs (figure B.49) to see if there were any identifiable properties. They were ordered according to the average human score for the preferred orientation and plotted in that orientation. The curvature was signed to judge the effect of the direction of curvature.

The first observation was that the rate of change of the curvature does not seem to be as important as Birkhoff had suggested. Designs 3, 4 and 28 all have highly variable

Figure B.49: Vase Outlines with Curvature Plots

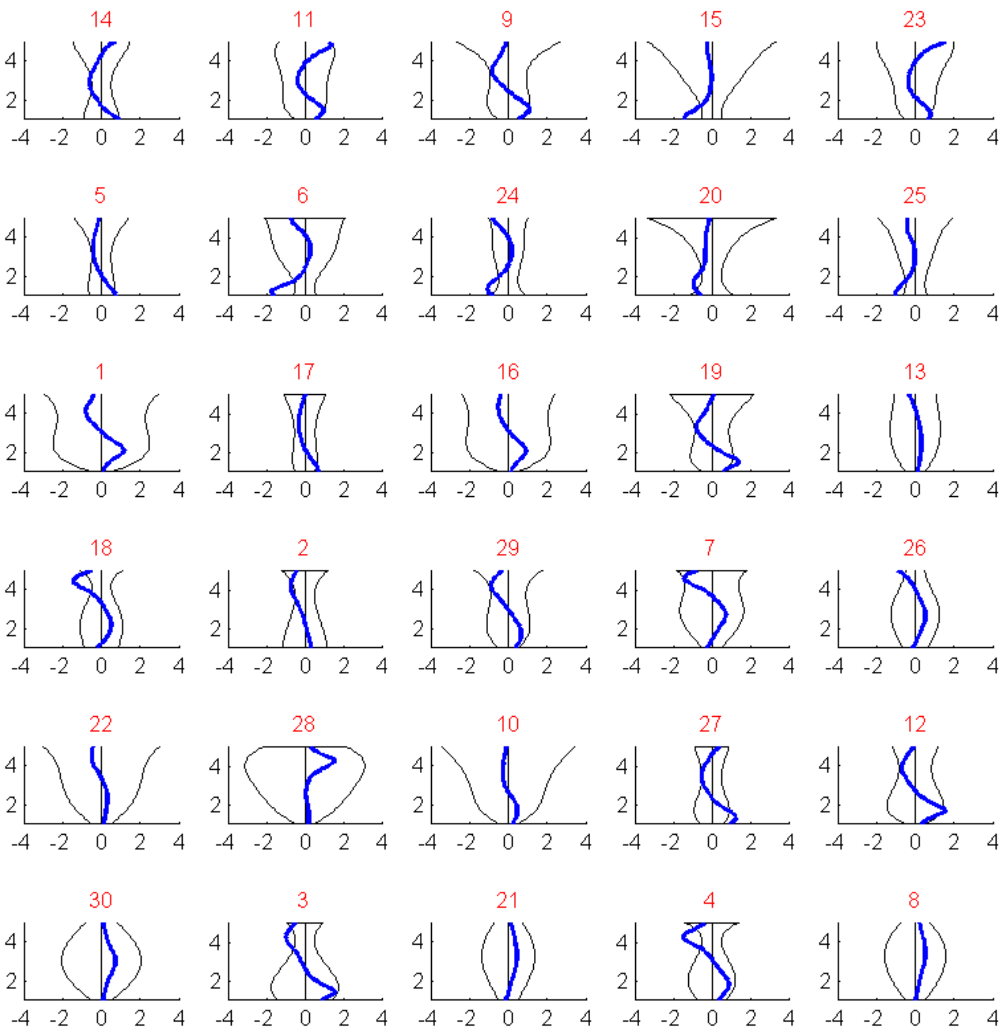

Figure B.50: Points of Inflection Measure

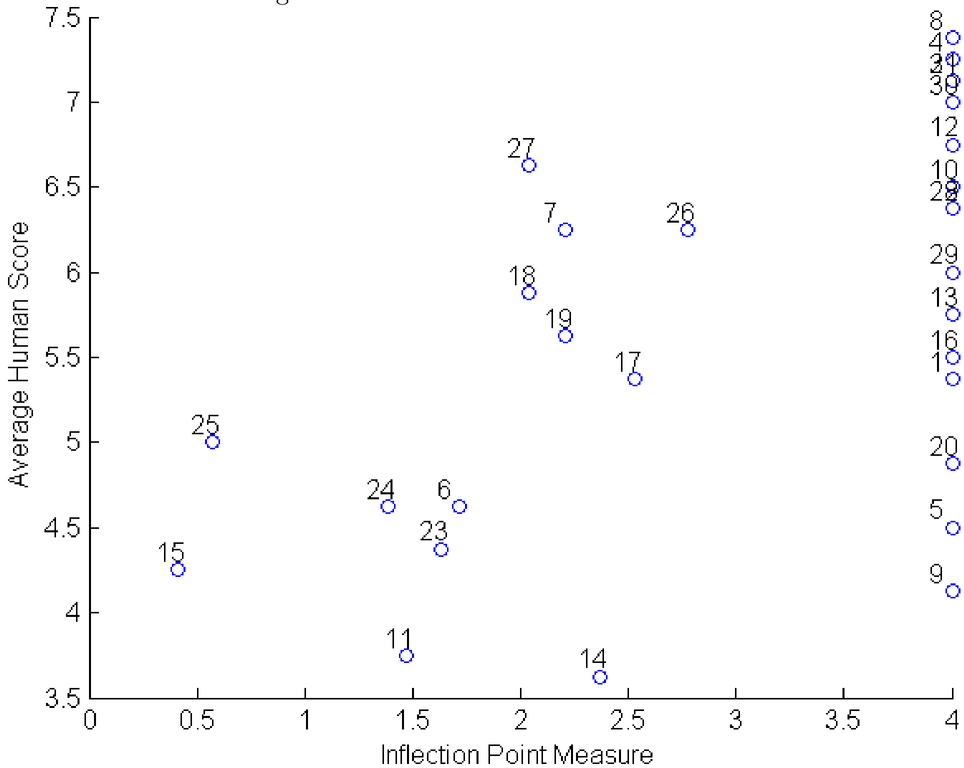

curvatures but are judged to be very good and designs 14, 11 and 5 have very regularly changing curvatures but are disliked.

One property that appears to be undesirable is the presence of several points where curvature is 0, in particular if these are close together or do not have a maximum or minimum in between them. The points are particularly close together on 11, 15 and 24 and are fairly close together on many others in the lowest scoring half. Most of these also do not have a maximum or minimum between the points.

By comparison the top scoring half only has 4 instances where there is more than one change of direction and all of these have them separated by at least half the height of the pot.

These were plotted to give a more formal assessment of this property (figure B.50). If 2 changes of curvature were present then the distance between them was plotted, otherwise a value of 4 (the full height of the vase) was used instead.

Figure B.51: Product of Distance Measure

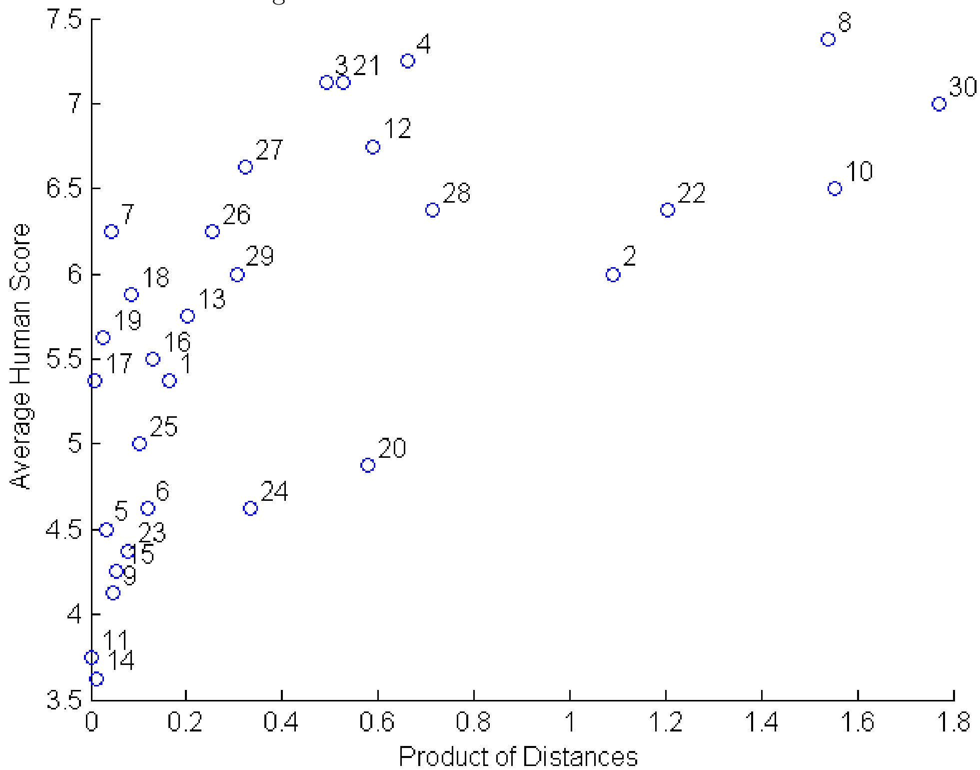

A more general measure looked at the spacings of all of the points. Observations suggested that the points in the preferred designs were quite evenly spaced. To measure this the distances between the points were multiplied together and the vertical product and the horizontal product were then summed. Before multiplying the distances they were divided by the largest distance. This prevented the domination of those with 3 significant points (which were guaranteed to have one or more vertical values greater than 1, giving them a large bias) and ensured all of the values were between 0 and 1. This is shown in figure B.51.

The overall curvature was also studied, first as a sum of the signed curvature and therefore favouring designs that are more convex and second as a sum of the unsigned curvature, favouring designs that had high curvatures in either direction. These are shown in figures B.52 and B.53.

Figure B.52: Signed Curvature Measure

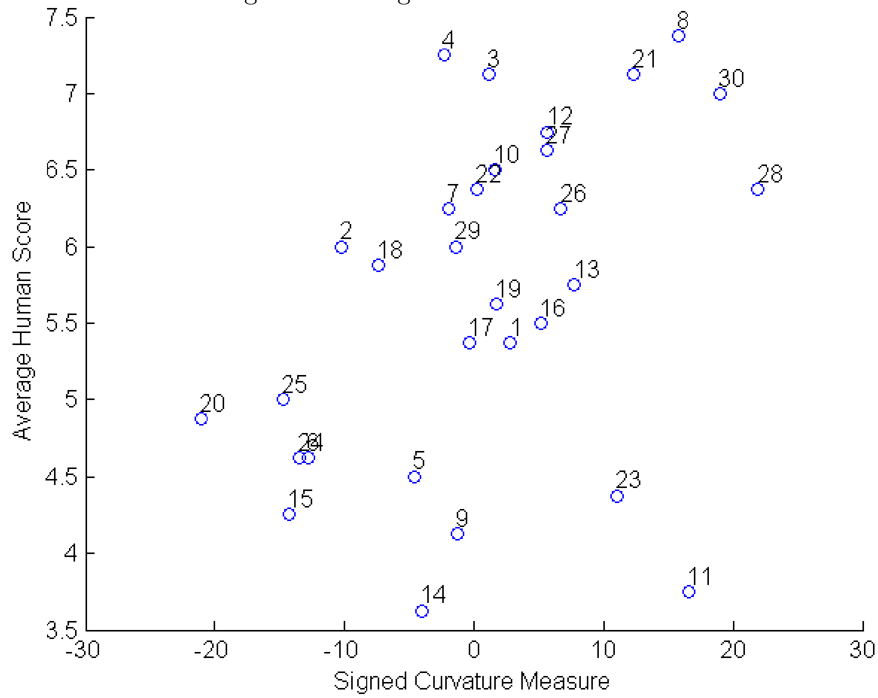

Figure B.53: Unsigned Curvature Measure

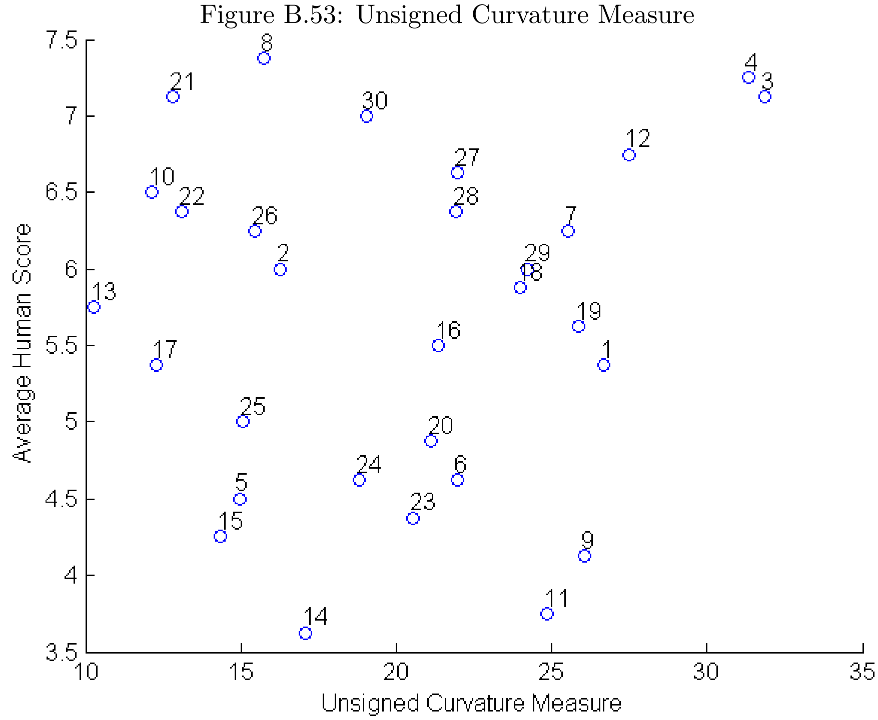

B.2.2.4 Other Style Properties Analysis

B.2.2.4.1 Width of Vase These two properties; the size of the maximum breadth and the position, produced the plots in figure B.54 where a value of 1 indicates the property was satisfied and 0 if not. A correlation value was also found for each plot.

B.2.2.4.2 Difference from Mean As in Bergen's measure the signed angles between neighbouring faces were found and a histogram produced. The normalised histogram was produced using a normalised function file called *histnorm.m* written by Arturo Serrano from the Matlab file exchange[42]. The difference between this and a normal distribution with the same mean and standard deviation was then found.

The inverse of this score (as a smaller score would be better) was plotted against the average human score as in the other measures and is shown in figure B.55.

Figure B.54: Maximum Width Requirements

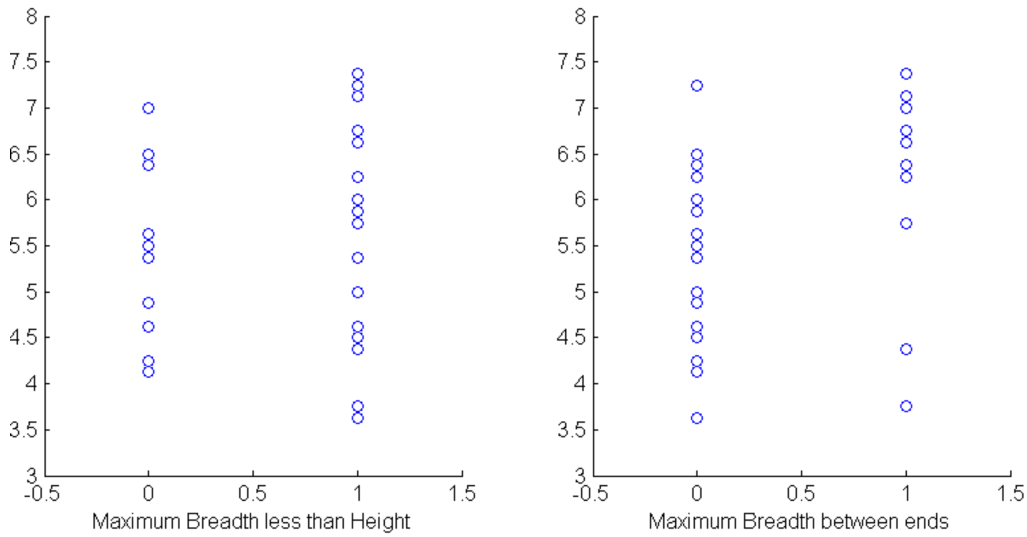

Figure B.55: Difference from Mean scores

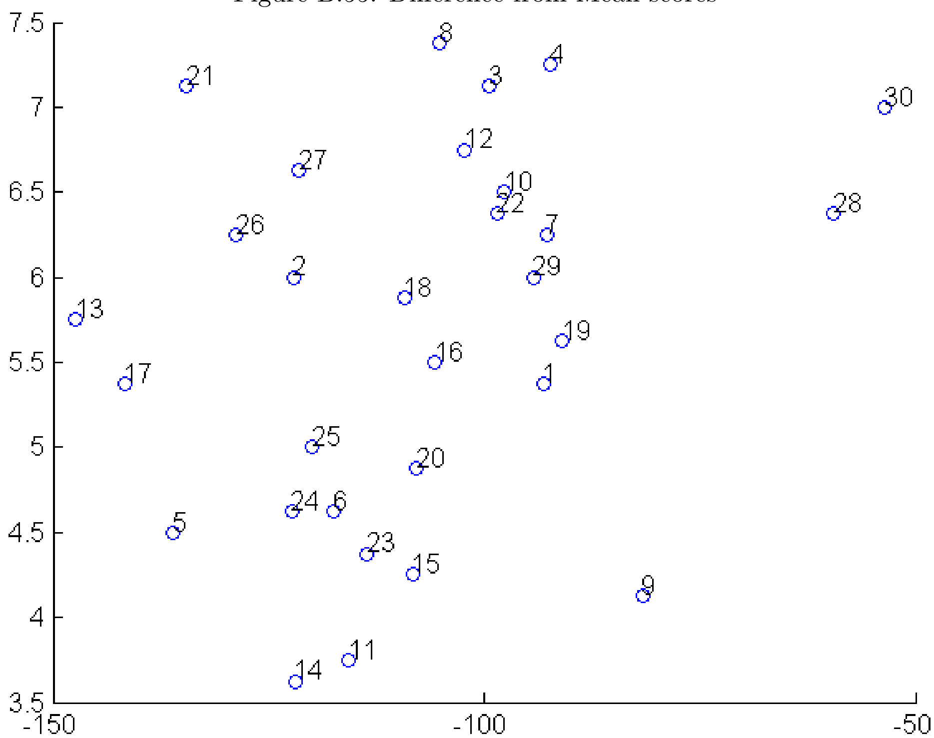

B.2.3 Orientation Analysis

The position of the minimum width was plotted against the average human score for both orientations and the two orientations visually compared using a connecting line either in blue (solid) or red (dashed). Blue indicates that the minimum point was in the upper half and red the lower half.

As you can see from figure B.56 the majority for the location of the minimum width were red, proving the opposite of Birkhoff's requirement. The correlation of the vertical and horizontal differences was also found, giving a value of -0.7805 with a positive confidence of 0. Looking closer we can see that most of the minima are positioned at the ends of the vase. In this case all but one favour the minimum at the base. However in the case of the minimum being between the ends the picture is less clear with five blue and seven red lines.

We can say, therefore, with some conviction that a minimum at the rim is undesirable but otherwise the desirable position of the minimum is undefined. This could be used in the fitness function by imposing a penalty on those with the minimum at the top but it would not be useful otherwise.

A similar graph was plotted for the angle at the base (figure B.57). As before both orientations were plotted. The angles from the horizontal were calculated and a low value indicates a sharp concave edge. If Birkhoff's statement is correct we would expect those that have a low angle to have a low score. This time the results support his observation. Most of those with a low angle at the base have low average scores. Those that have good scores are still improved by inverting the design so that the lowest angle becomes the rim. In fact nearly all of the designs prefer the largest angle at the base.

The weight balance was measured in two ways; first the centre of mass was found (using the trapezium rule to estimate the volume) and its position was plotted against the human score (figure B.58). This does show some preference for the top heavy orientation but has many exceptions. It is worth noting that all of the exceptions from the previous measure are

Figure B.56: Position of the Minimum Width

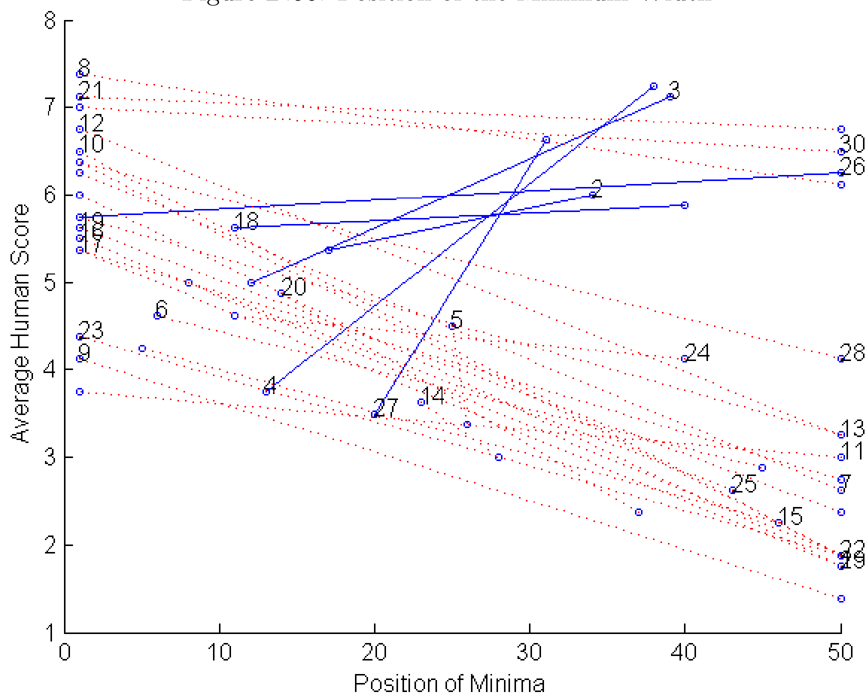

Figure B.57: Angles at Base

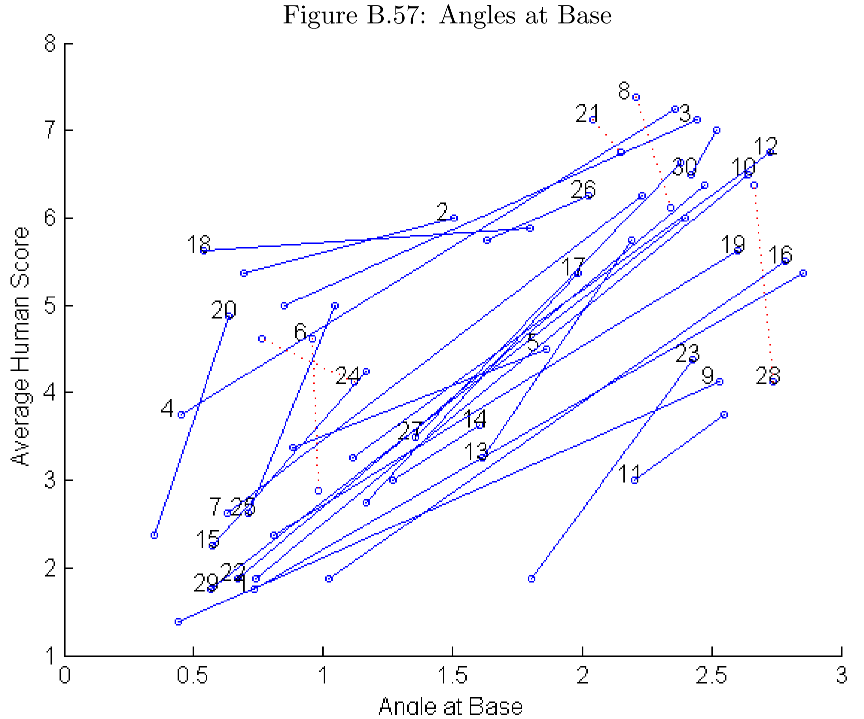

Figure B.58: Centre of Mass

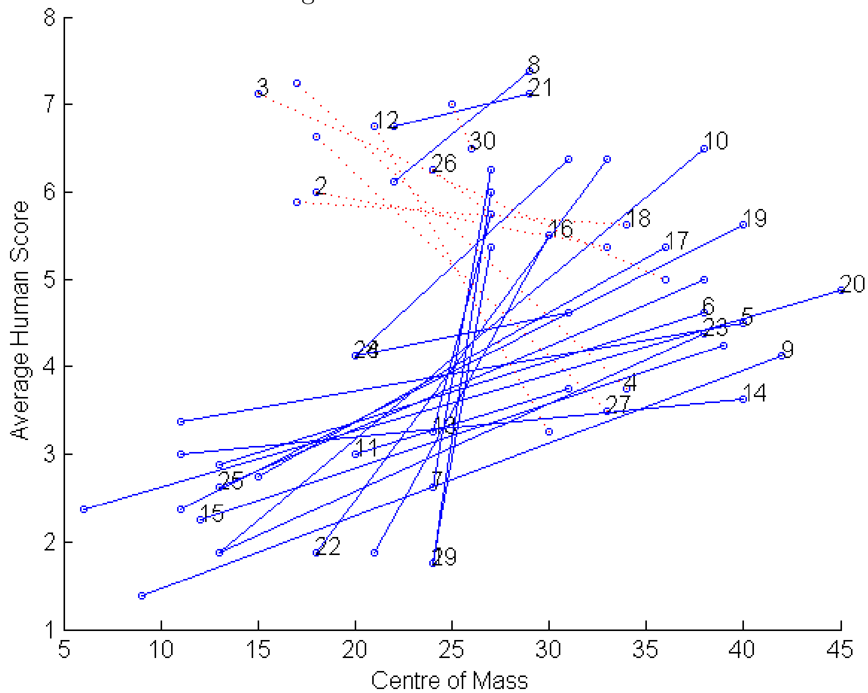

Figure B.59: Balance of Features

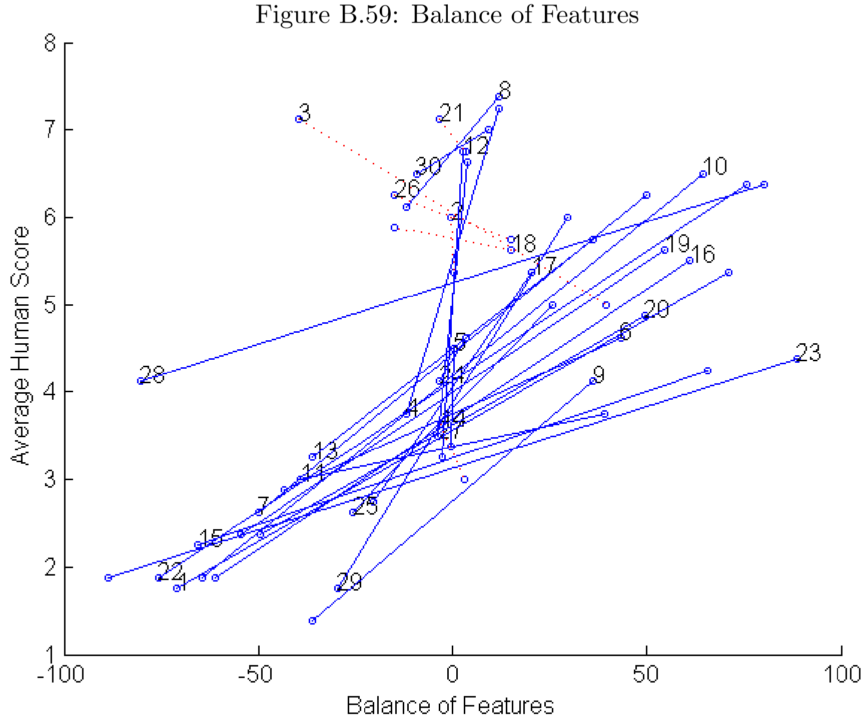

now clearly top heavy in their preferred orientation. Therefore it is likely to be a combination of these factors that contribute to the overall orientation.

The second weight measurement looked at the distribution of the significant points. Each point was assigned a score by finding the area of the rectangle between it and the centre of the vase. This centre was on the central axis and halfway up the height. The results are plotted in figure B.59.

B.2.4 Weights

In order to find suitable weights for the measures from the collected data in real time a simple, automatic method of calculating them was required.

The simplest method for this would be to add all the normalised component scores to find an overall score. The sum of the style scores have then been plotted against the average human scores (figure B.60). The sums were also normalised and then adjusted so their mean and standard deviation were the same as the human scores to aid comparison, the line shown in the figures is the 'ideal' correlation. The correlation of this measure was 0.6950, larger than most of the component scores and with no major anomalies in the lower right corner.

To attempt to improve the correlation further a second combination was tried where the component scores were weighted by their correlation values. The idea being a more significant trait would have a higher correlation value and would also be desirable as a large component in the full measure. The values weighted this way produced figure B.61 and had a correlation value of 0.7226. This is a good improvement over the unweighted version.

However a brief manual adjusting of the weights suggests that weighting by correlations does not produce the best fitness function. Using the weights 0.5, 0.2, 0.6, 0.8, 0.1, 0.5, for the components in the order above, a correlation value of 0.7331 was achieved. This is still less than the highest individual measure but further data will reveal if there is an advantage to combining the measures.

Figure B.60: Even Weighted Style Measure

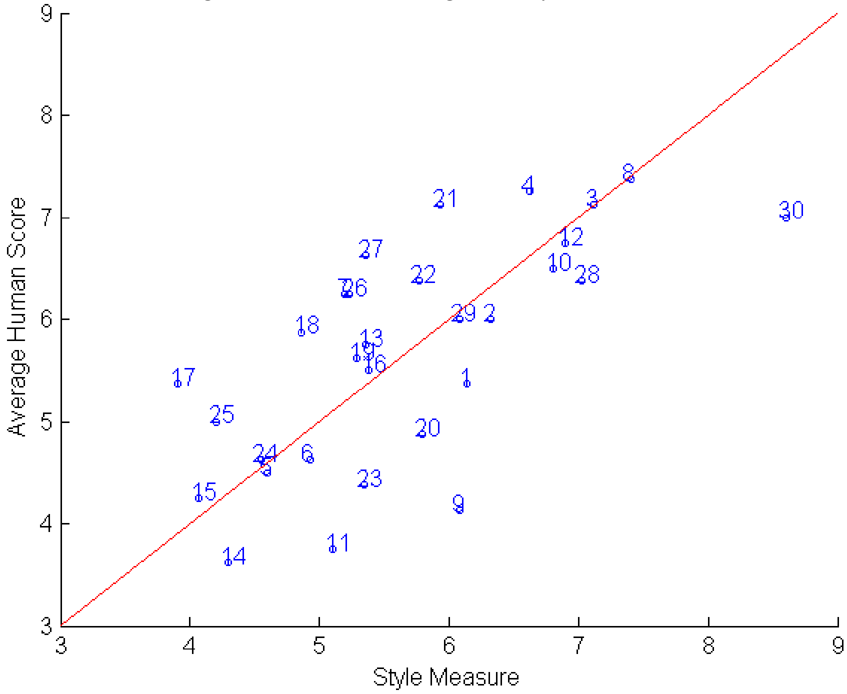

Figure B.61: Correlation Weighted Style Measure

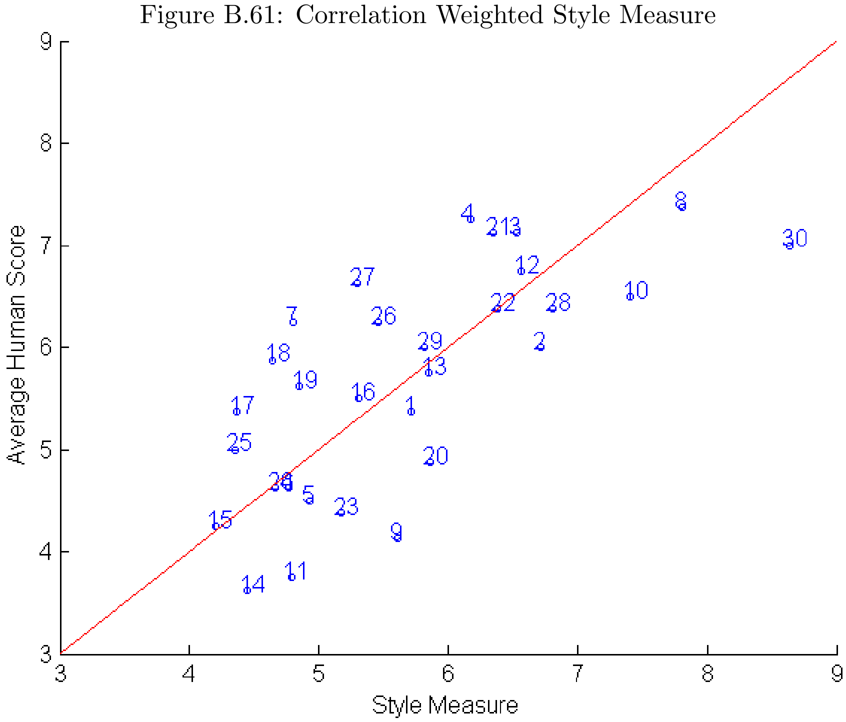

Figure B.62: Correlation Weighted Orientation Measures

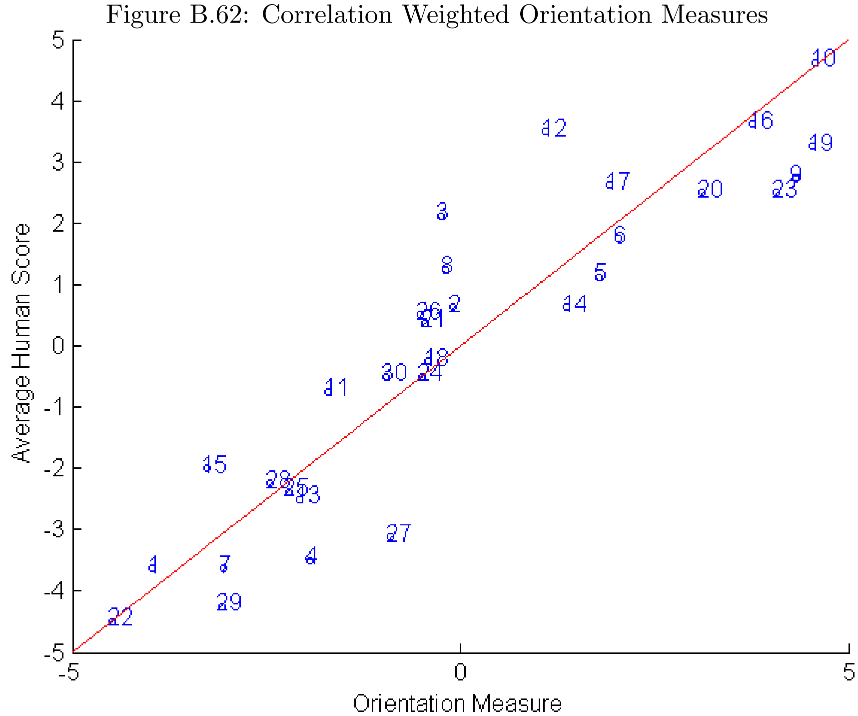

These correlations all used the original ratios for the Birkhoff measure, the ϕ measure was also studied. This gave a correlation weighted value of 0.7264, so although the ϕ ratio measure had a lower value than the original measure it provides a better combination with the other measures.

A similar process was used for the orientation measures; as before these were combined using three different weightings. These correlations were found using the differences between the two scores, which are also now used in the plots.

The even measure did not have any improvement over the best individual correlation with a correlation of 0.8858. The correlation–weighted measure performed much better, the correlation was now 0.9332 (see figure B.62).

The final weighting was manual and a brief search gave the weights 0.6, 0.1 and 0.3 to the different measures in the order above. This reflects the order of the correlations but the smaller orientations have been further reduced. This set of weights gave an correlation of 0.9680.

Other variations were tried such as squaring the correlation values or deducting a factor related to the presence of vases in the lower right hand corner. However these had mixed results and none showed improvement over the manual scores. Squaring the correlation values improved the orientation (0.9535) but not the style (0.7220). No simple automatic method was found for repeating the success of the manual weights. It would be possible to use a search method (possibly evolutionary) to find such weights but calculating this would be time consuming and unsuitable for real–time adjustments. As weighting by the correlation value appears to have good results this would be ideal for quick real–time weight adjustments.

Appendix C

Appendix: Experimental Setup Further Information

C.1 Changes

C.1.1 Representation

A break in the curve was added to increase the range of the representation. This new addition was represented as a second list of 3 points. These have an integer value and represent the state of the 3 internal coordinates. A value of 0 would mean that there was no break at that point, a non zero value would indicate a break.

As a full curve has 50 nodes where nodes 1 and 50 are the bottom and top points of the curve respectively, a break point can be positioned at any of the other 48 nodes. To ensure that multiple break points did not conflict, each of the three points could only occur within a set range of points. These were between 2–17 for point 1, 18–33 for point 2 and 34–49 for point 3.

Figure C.1 shows the impact of introducing a break point and the variation in break location, where the blue (dot and dash) line has no break point (value 0), red (solid) and green (dashed) have non–zero entries for value 3, red has 35 and green has 47.

Figure C.1: Effect of varying the break point value

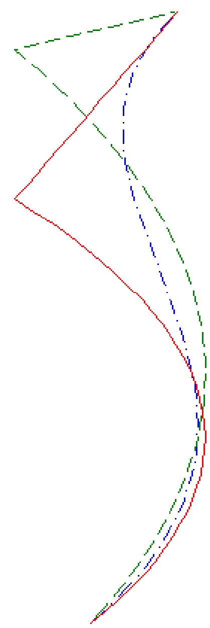

C.1.2 Measures

C.1.2.0.3 Modified Birkhoff's measure: The main adaptation was concerning the location of the significant points. Birkhoff had stated in his original measure that sharp corners should be considered significant points and this was added as a fourth type (after end points, points of vertical tangent and points of change of direction). The method of detecting change of direction points also presented problems as it detected points of zero curvature and all nodes on a curve generated by only 2 points (e.g. 2 consecutive break points) have this property, giving excess significant points. This was prevented by not searching for points within these types of curves as none of the significant point types would exist in these curves. Finally the ratios used were $1 : \phi$, $1 : 1$ and $1 : \frac{1}{\phi}$. Although the original ratios gave a better individual correlation the ϕ ratio had a better overall ratio with the other measures.

Therefore the full list of significant points (found using the m file *AComplexity.m*, shown in section C.3), and the methods used to locate them are as follows:

- End points: Nodes 1 and 50.

- Points of vertical tangent: The differences of consecutive radius values were found, these were called dr. A point of vertical tangent was found by finding the nodes where the preceding and following dr had different signs.

- Points of inflection: The perpendicular vector was found for each consecutive pair of dr values. These would point either in or out of the page and a change in sign indicated a point of inflection. The actual point of inflection was an edge but a node was required, therefore the program selected the node preceding the edge. This would mean that repeating the test with the inverse designs would produce slightly different values. This would usually make little difference but sometimes it was enough to vary the measures, particularly the Birkhoff measure as it has only small tolerances on its ratios.

  It was decided that although there may be some variation, to reduce calculation time it was preferable to use the same style values for both orientations as there was no visual change in the design.

- Sharp Points: Taken from the representation.

C.1.2.0.4 Sum of the signed curvature: The original measure used a 4 node version of discrete curvature, however it was found that the curvature values over the sharp angles would be very inaccurate. An alternative method, that finds the radius of a circle determined by used only 3 nodes (figure C.2) was found to be much more accurate at these points. It uses the equation:

Figure C.2: New Curvature

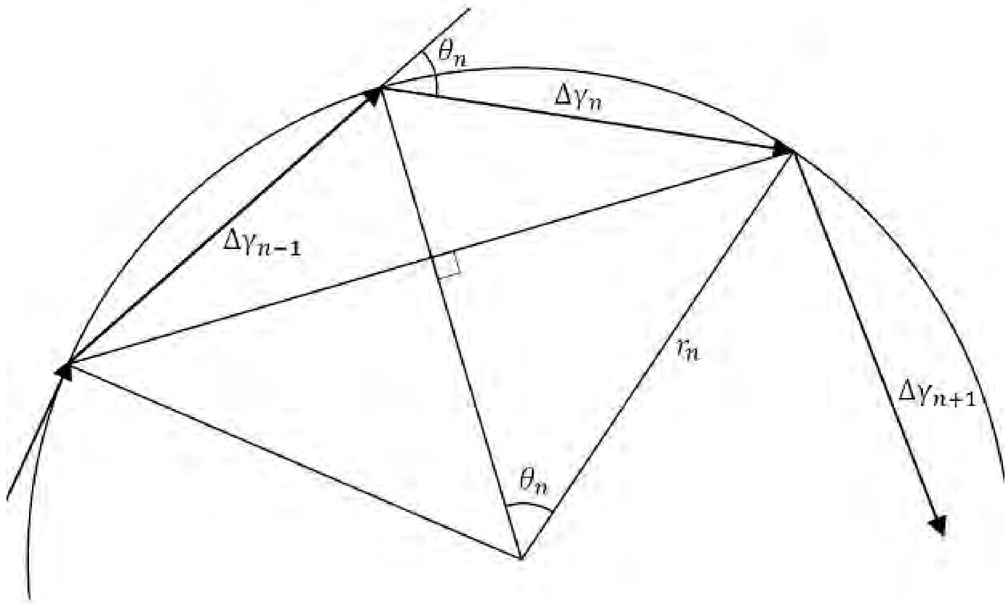

$$\kappa = \frac{2\sin\theta_n}{\|\Delta\gamma_{n-1} + \gamma_n\|} \tag{C.1}$$

This had a slight effect on the curvature values, changing the correlation to 0.4190 with the confidence over 98% which is still suitable.

C.1.2.0.5 Points of inflection: With the introduction of the sharp angles there were very few vases with 2 or more points of inflection, making this measure less useful. To amend this the new sharp points were added to the points of inflection, hoping that the dislike of complexity summarised by this measure would still be valid. However if there are other concerns that this measure observed then it may now be irrelevant. Also the value awarded if there is one or less points of inflection was increased to 5 in line with the increased maximum height.

C.1.2.0.6 Difference From Mean: This was not included in the final program. This was because it took considerably longer to calculate compared to the other measures due to its analysis of the 3D form rather than 2D. It may yet have potential as a measure but as the evolution runs in parallel with the data collection it made the program slow to respond, potentially annoying the users.

C.1.2.0.7 Existence of sharp points: As there is no initial data on the preference for the existence of these new sharp points an open measure was introduced, this simply awarded a point if the vase did not contain a sharp point and remained 0 if it did. The weights were allowed to be negative so that the preference could exist either way.

C.1.2.0.8 Balance of features: Early tests found that many of the evolved vases would become very wide. Although there are be several reasons why this may be, this measure was found to be a major cause. As it was looking for the differences between the areas of the top and bottom rectangles formed by the points it could be made to score very highly by making one rectangle very large, as the height is fixed the only option to do this was adjusting the width. To prevent this the width part was removed. Unfortunately this had the effect of removing the good correlation, giving a value of -0.0473. However as this would be ignored if the value remained this low it was still included to see if there was any potential for the modified version.

C.2 Full Program Description

This describes the full test program. Figure C.3 shows the overview diagram.

The human interface begins with a list of instructions and asks the volunteer for information concerning their design education (GCSE, A Level, Degree or Design Professional) (figure C.4). Once they have read the instructions they begin by clicking the 'Start' button. This chooses 4 designs and displays them to the volunteer (figure C.5). They are selected by first choosing the design with the lowest number of tests and then selecting 3 other designs with similar scores. This selection method ensures all the designs are studied and will

Figure C.3: Program Overview - User Interface part in Blue, Timer Section in White

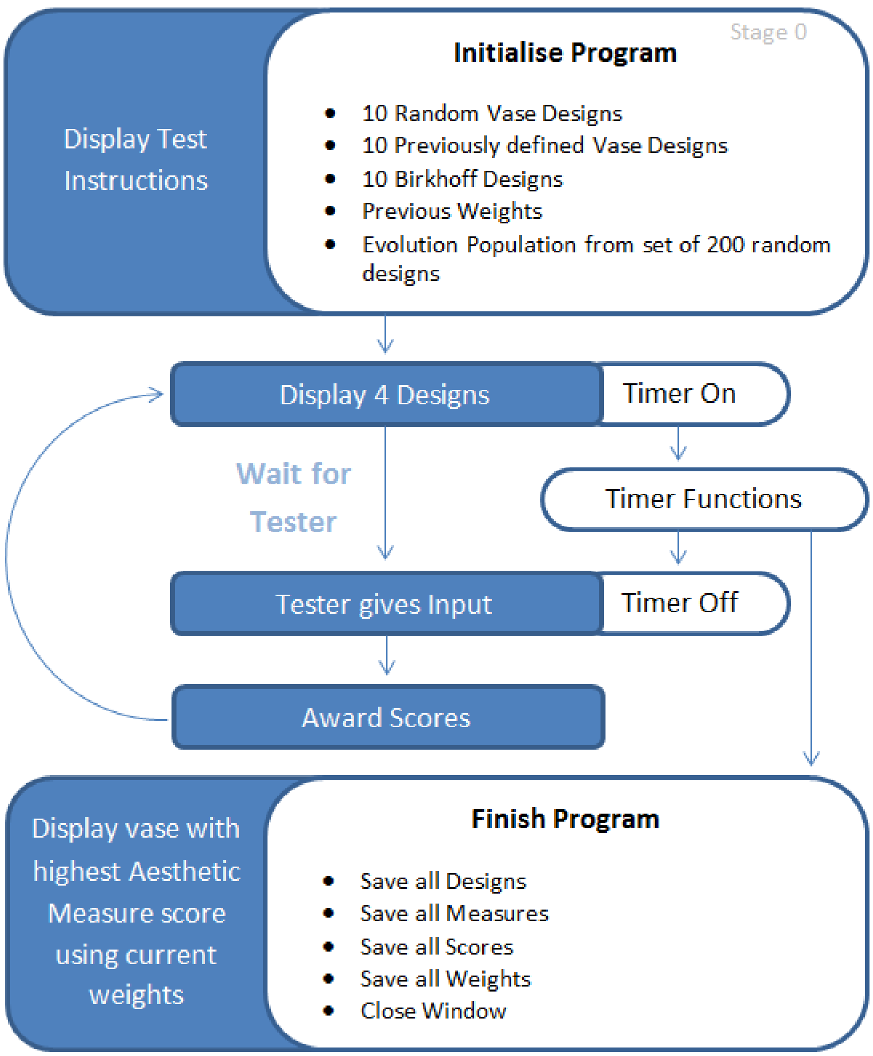

Figure C.4: Interface Initial Screen

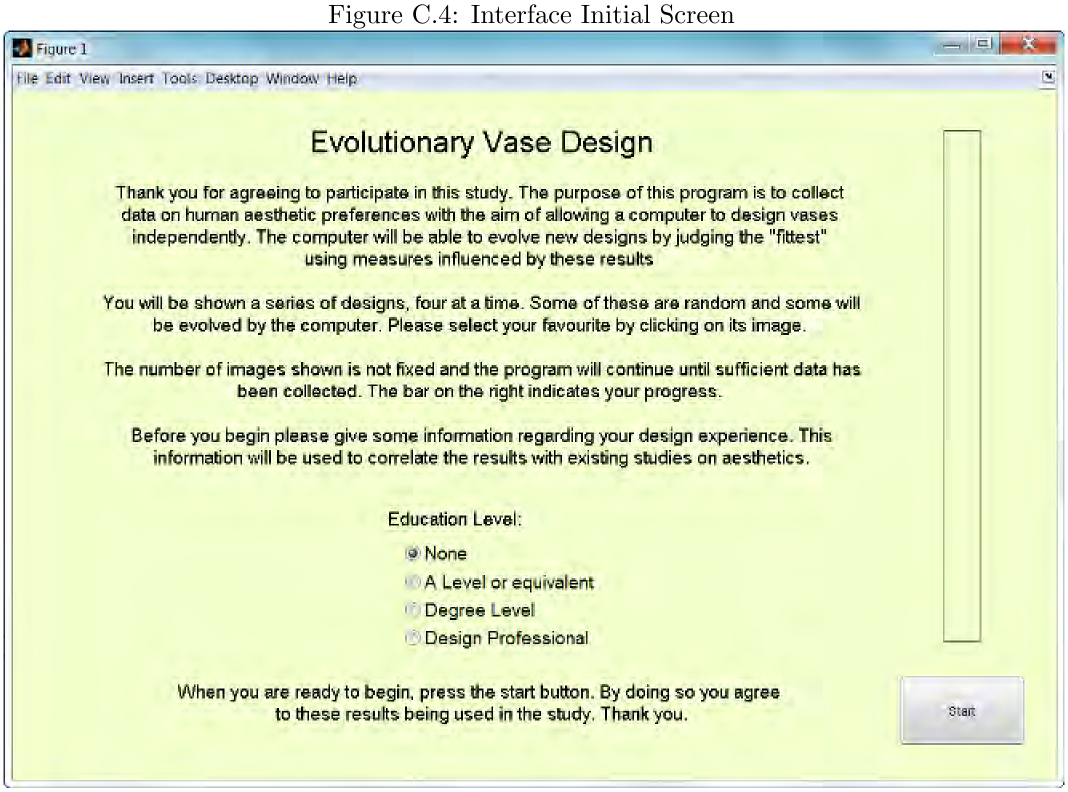

Figure C.5: Interface Selection Screen

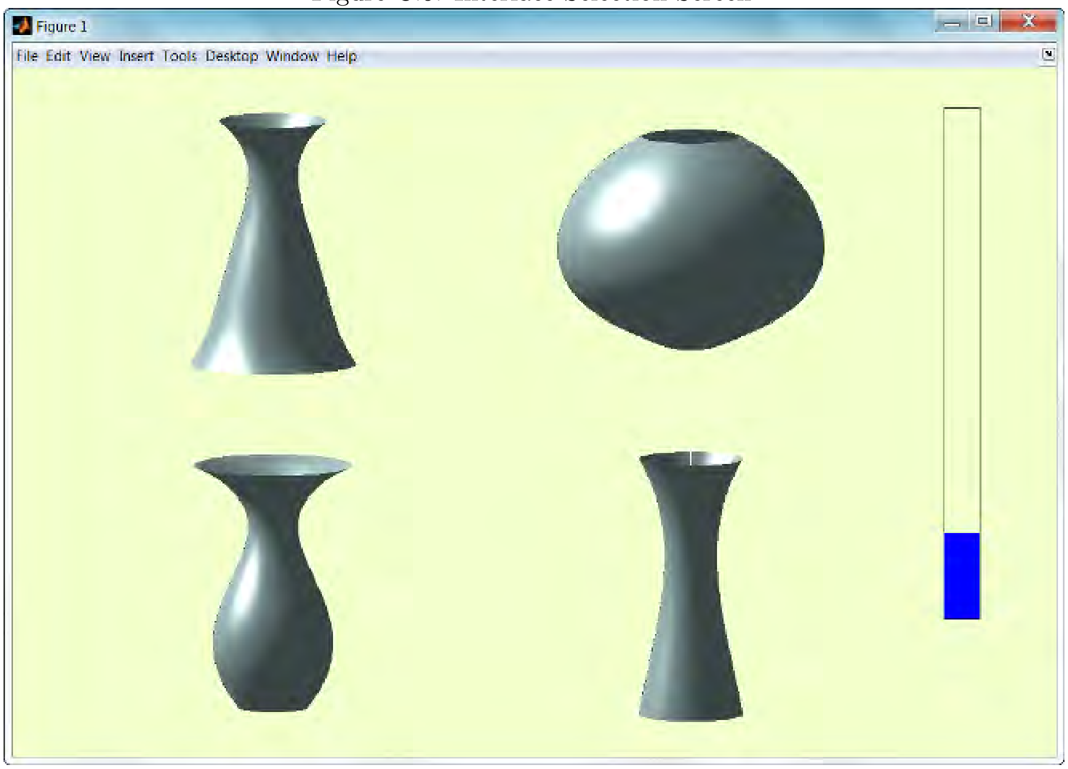

automatically favour the newest designs (as their number of tests will be small). By then selecting for similar scores it helps break up clumps of scores.

The volunteer selects their favourite by clicking on it. This design is then awarded 3 points, 1 is deducted from each of the others. The number of tests is updated by 1 for all 4. The program then select 4 new designs as before and displays them.

The timer driven section is responsible for the evolution, calculation of the weights and addition of new designs. The timer begins when the program is run, depending on different conditions the function that is triggered by the timer varies. Once running the timer continues to trigger new functions until the volunteer stops it by choosing a design, once the new designs have been displayed the timer begins again.

The first function (stage 0) sets up the initial set of 30 designs. These are made up of 10 random designs, 10 Birkhoff evolved designs (from a set of 200 previously generated with the

Figure C.6: Timer Function - Stage 1

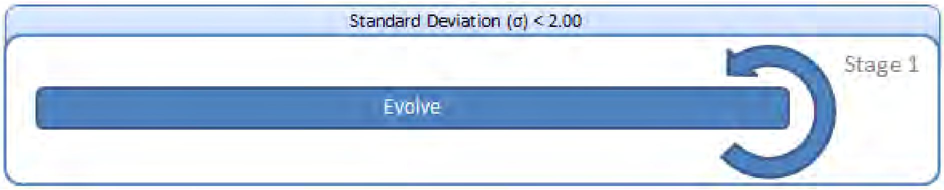

original Birkhoff function) and 10 set designs (chosen from a random set to display a wide range of measures, these are to allow direct comparison between volunteers). The inverse of these designs is also found so that the program can study the effects of orientation. This gives a total initial pool of 60 designs. The measures of all 60 are found and stored in a matrix along with the genotypes and scores (set to 0).

All of the stored designs are found in *.xlsx* files that accompany the program. Previous results are also stored in this way and the function also reads these to find the initial weights based on their scores.

The initial function also generates and finds the measures of a large pool of random designs, from these the top 20 using the previous weights and set as the initial population for the evolution. The rest are stored to enable some to be introduced to the evolution at later steps. The final action of this function is to make the 'Start' button appear.

The default function is the evolution function. This uses the current weights to assess the evolution and each cycle of the function is a single generation of the evolution. This function is called every time unless one of the other criteria is met, after the number of designs has reached 50 it is no longer called. This is stage 1 (see C.6).

If the standard deviation of the awarded scores has reached one of a several cut off points (number of vases/15, first is 2) then the function that adds a new design is triggered, this is stage 2 (see C.7). The top design from the evolved population is chosen and added to the matrix of designs, along with its inverse. At this point the weights are also re–evaluated. A new evolution population is then found by taking 15 of the original population (top 16 excluding the one added to the main pool) and then introducing new genes by selecting the

Figure C.7: Timer Function - Stage 2

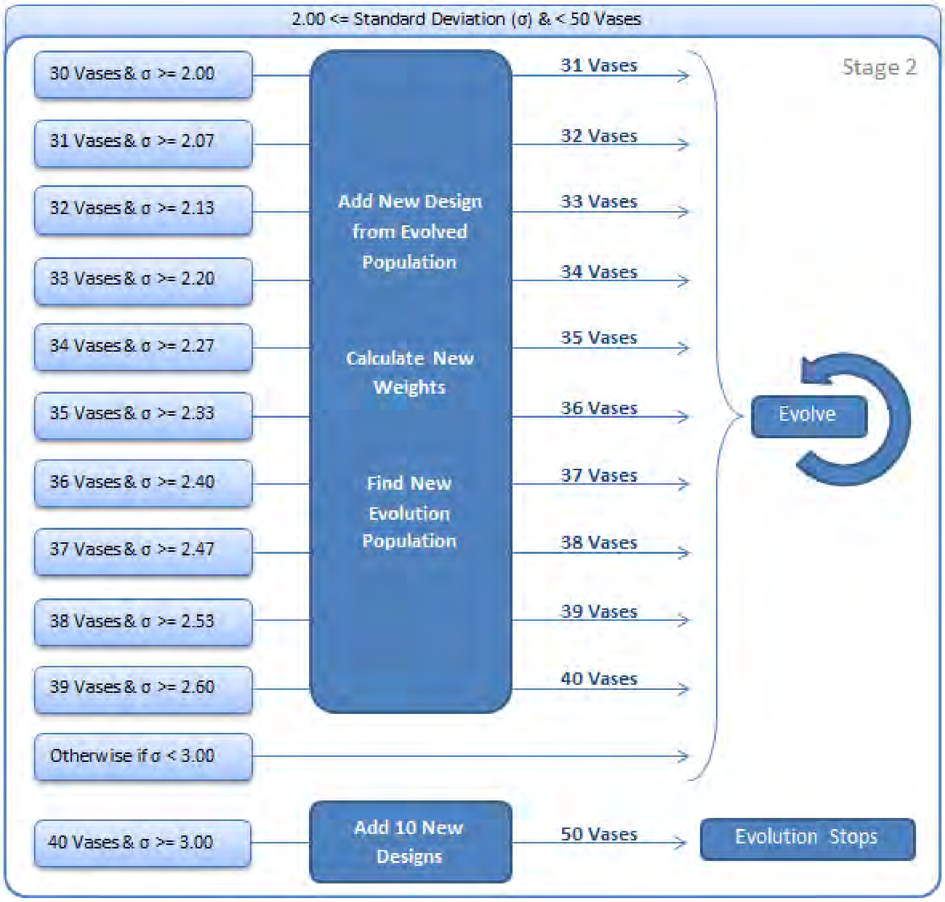

top 5 of the original set of 200 using the new weights. This function is triggered until there are 40 designs (or 80 if you include the inverses) in total. The standard deviation is used as the results need to have a suitable distribution to be useful.

When the standard deviation has reached 3.5 then the top 10 of the evolved designs are added. This larger value gives these vases longer to evolve with a fitness that now includes the preferences of the current volunteer. After these 10 have been added no further designs are evolved, this is stage 3 (see C.8).

The final function is called when the standard deviation has reached 4. This then stops the timer, finds the final weights and displays the vase judged the best by this measure. It also creates a 'Finish' button that stores all the new information in a master file when clicked.

Figure C.8: Timer Function - Stage 3

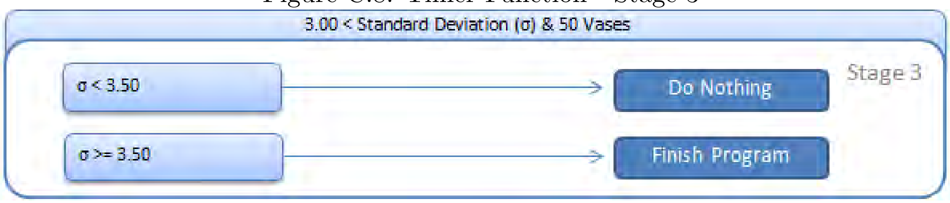

C.3 M Files

Vase.m

```
function Vase

f=figure('Visible','off','Units','normalized','Position',[0.1,0.1,0.8,0.8],'
    color',[0.95,1,0.8],'CloseRequestFcn',{@Close});

hnext    = uicontrol('Style','pushbutton','String','Start','Units','
    normalized','Position',[0.85 0.05 0.12 0.1],'Callback',{
    @randomizebutton_Callback},'Visible','off');

P1 = axes('Units','normalized','Position',[0.05,0.05,0.4,0.4],'Visible','off
    ');
P2 = axes('Units','normalized','Position',[0.45,0.05,0.4,0.4],'Visible','off
    ');
P3 = axes('Units','normalized','Position',[0.05,0.55,0.4,0.4],'Visible','off
    ');
P4 = axes('Units','normalized','Position',[0.45,0.55,0.4,0.4],'Visible','off
    ');

Bar= axes('Units','normalized','Position',[0.89,0.2,0.04,0.75],'Visible','off
    ');

string1='Thank you for agreeing to participate in this study. The purpose of
    this program is to collect data on human aesthetic preferences with the
    aim of allowing a computer to design vases independently.';
string1=sprintf('%s The computer will be able to evolve new designs by
    judging the "fittest" using measures influenced by these results', string1
    );
string1=sprintf('%s \n \n You will be shown a series of designs, four at a
    time. Some of these are random and some will be evolved by the computer.
    Please select your favourite by clicking on its image.', string1);
string1=sprintf('%s \n \n The number of images shown is not fixed and the
    program will continue until sufficient data has been collected. The bar on
     the right indicates your progress.', string1);
string1=sprintf('%s \n \n Before you begin please give some information
    regarding your design experience. This information will be used to
    correlate the results with existing studies on aesthetics.', string1);

Instructions1=uipanel('BackgroundColor',[0.95,1,0.8],'Units','normalized','
    Position',[0.45,0.45,0.01,0.01],'BorderType','none');

Title1=uicontrol('Parent',Instructions1,'Style','text','BackgroundColor
    ',[0.95,1,0.8],'Position',[-400 310 800 50],'FontSize',20,'FontName','
    Arial','String','Evolutionary Vase Design','Visible','on');

Instructions2=uicontrol('Parent',Instructions1,'Style','text','
    BackgroundColor',[0.95,1,0.8],'Position',[-400 -100 800 400],'FontSize
    ',12,'FontName','Arial','String',string1);
```

```
Education=uibuttongroup('Parent',Instructions1,'Units','pixels','
    BackgroundColor',[0.95,1,0.8],'BorderType','none','Position
    ',[-110,-190,300,150],'Visible','on','Title','Education Level:','FontSize
    ',12,'FontName','Arial');
Edu1=uicontrol('parent',Education,'Style','Radio','Tag','1','BackgroundColor
    ',[0.95,1,0.8],'String','None','Units','normalized','Position
    ',[0.1,0.71,0.8,0.2],'FontSize',12,'FontName','Arial');
Edu2=uicontrol('parent',Education,'Style','Radio','Tag','2','BackgroundColor
    ',[0.95,1,0.8],'String','A Level or equivalent','Units','normalized','
    Position',[0.1,0.48,0.8,0.2],'FontSize',12,'FontName','Arial');
Edu3=uicontrol('parent',Education,'Style','Radio','Tag','3','BackgroundColor
    ',[0.95,1,0.8],'String','Degree Level','Units','normalized','Position
    ',[0.1,0.25,0.8,0.2],'FontSize',12,'FontName','Arial');
Edu4=uicontrol('parent',Education,'Style','Radio','Tag','4','BackgroundColor
    ',[0.95,1,0.8],'String','Design Professional','Units','normalized','
    Position',[0.1,0.02,0.8,0.2],'FontSize',12,'FontName','Arial');

Instructions5=uicontrol('Parent',Instructions1,'Style','text','
    BackgroundColor',[0.95,1,0.8],'Position',[-320,-320,640,100],'FontSize
    ',12,'FontName','Arial','String',...
    'When you are ready to begin, press the start button. By doing so you
        agree to these results being used in the study. Thank you.');

Thanks=uicontrol('Style','text','Units','normalized','BackgroundColor
    ',[0.95,1,0.8],'Position',[0.35 0.9 0.2 0.05],'FontSize',20,'FontName','
    Arial','String','Thank you','Visible','off');
PFinal= axes('Position',[0.2 0.05 0.5 0.8],'Visible','off');

tim = timer('TimerFcn',@Time,'StartDelay',0.1,'Period',0.1,'ExecutionMode','
    fixedSpacing');
TimNo=uicontrol('Style','text','Value',1,'Visible','off');

axes(Bar)
rectangle;
set(Bar,'XLim',[-0.1 1.1],'YLim',[0 1.01])

set(f,'Visible','on')

start(tim)

    function Time(source,eventdata)
            global W
            global a
            global Ma
            global a1
            global MeanM
            global SDM
            global A
            global D
            global w
            global as1
            global as
            global A8a
            global A8b
            T=get(TimNo,'Value');
        if T==1
            N=10;
            a=NaN(5,20);
            as=NaN(3,20);
            V=xlsread('BirkhoffVase.xlsx',1);
            A8a=xlsread('Results.xlsx',1);
            A8b=xlsread('Results.xlsx',2);
            A1=[A8a,A8b];
```

```
    B=A1(2:6,:);
    D=size(B,2)/2;
    c=A1(10,:)./(A1(11,:)+1);
    C=[c(1:D)',c(D+1:2*D)'];
    M=A1(12:20,:);
    MeanM=mean(M,2);
    SDM=std(M,0,2);
    W=WeightFinder(C,M,D);
    w=W;
    A2=V(:,ceil(rand(1,N)*size(V,2)));
    A3=xlsread('Control.xlsx');
    D=2*N+10;
    A=linspace(1,2*D,2*D);
    A(2:6,1:N)=randn(5,N);
    A(2:9,N+1:2*N)=A2;
    A(2:9,2*N+1:D)=A3;
    A(2:6,D+1:2*D)=flipud(A(2:6,1:D));
    A(7:9,1:N)=(rand(3,N)<0.2).*(round(rand(3,N)*15)+repmat(linspace
        (2,34,3)',1,N));
    A(7:9,D+1:2*D)=flipud((51-A(7:9,1:D)).*(A(7:9,1:D)>0));
    M=MeasureFinder(A(2:6,:),D,A(7:9,:));
    A(10:11,1:2*D)=zeros(2,2*D);
    A(12:20,1:2*D)=M;
    a1=randn(5,200);
    as1=(rand(3,200)<0.2).*(round(rand(3,200)*15)+repmat(linspace
        (2,34,3)',1,200));
    Ma=MeasureFinder(a1,200,as1);
    AllR=NaN(50,200);
    for j=1:200
        [AllR(:,j),H]=Bezier2(a1(:,j),as1(:,j)');
    end
    r=min(AllR);
    AllR=AllR-repmat(r,50,1);
    Ind=Difference(AllR);
    Score=sum(Ma(:,1:200).*repmat(W,1,200));
    for j=1:20
        [S,I]=max(Score);
        a(:,j)=a1(:,I);
        as(:,j)=as1(:,I);
        Score(I)=NaN;
        Score=Score+((Ind(I,:)<10).*(Ind(I,:)-10));
    end
        set(hnext,'Visible','on')
    stop(tim)
else if T==2
        stop(tim)
        SetPlots(A)
        stop(tim)
    else if D>=50 && std(A(10,:))>=4
            set(hnext,'Visible','on')
            set(Thanks,'Visible','on')
            cla(P1)
            cla(P2)
            cla(P3)
            cla(P4)
            a=A(:,41:50);
            A=[A8a,A(:,1:D),A8b,A(:,D+1:2*D)];
            D=size(A,2)/2;
            g=A(11,:)+1;
            G=A(10,:)./g;
            W=WeightFinder([G(1:D)',G(D+1:2*D)'],A(12:20,:),D);
            w(:,12)=W;
            [Best,BestI]=max(sum(a(12:20,:).*repmat(W,1,10)));
            [R,H]=Bezier2(a(2:6,BestI),a(7:9,BestI)');
```

```
                R=R-min(R)+0.5;
                H=H-min(H);
                phi=linspace(0,2*pi,50);
                x=sin(phi);
                y=cos(phi);
                z=ones(1,50);
                C=ones(50,50,3)*0.9;
                C(:,:,1)=C(:,:,1)*0.85;
                Z=z'*H;
                X=x'*R;
                Y=y'*R;
                surf(PFinal,X,Y,Z,C,'FaceLighting','phong')
                set(PFinal,'DataAspectRatio',[1 1 1],'Visible','off');
                shading(PFinal,'interp')
                light('Parent',PFinal,'Position',[-5,0,5]);
                set(PFinal,'view',[0 8])
                stop(tim)

            else if D<40 && std(A(10,:))>=D/15;
                    addA
                    evolve(W,MeanM,SDM)
                else if D<50;
                        evolve(W,MeanM,SDM)
                        if std(A(10,:))>=3.5
                            AddAllA
                        end
                    end
                end
            end
        end
    end
    T=T+1;
    set(TimNo,'Value',T)
end

function randomizebutton_Callback(source,eventdata)
    global A
    global D
    global w
    switch get(hnext,'string')
        case 'Start'
            set(Education,'Visible','off')
            set(hnext,'String','Finish')
            set(hnext,'Visible','off')
            set(Instructions1,'Visible','off')
            set(Instructions2,'Visible','off')
            set(Instructions5,'Visible','off')
            set(Title1,'Visible','off')
            start(tim)
        case 'Finish'
            Ed=str2double(get(get(Education,'SelectedObject'),'Tag'));
            Ed1=xlsread('Results.xlsx',3);
            xlswrite('Results.xlsx',[Ed1,Ed],3);
            xlswrite('Results.xlsx',A(:,1:D),1);
            xlswrite('Results.xlsx',A(:,D+1:2*D),2);
            w1=xlsread('Results.xlsx',4);
            xlswrite('Results.xlsx',[w1,w],4);
            closereq
            clear A
            clear D
            clear p
            clear a
            clear MeanM
            clear SDM
```

```
            clear W
            clear Ma
            clear w
            clear as
            clear as1
            clear A8a
            clear A8b
            stop(tim)
            delete(tim)
    end
end

function Close(source,eventdata)
    closereq
    clear A
    clear D
    clear p
    clear a
    clear MeanM
    clear SDM
    clear W
    clear Ma
    clear w
    clear as
    clear as1
    clear A8a
    clear A8b
    stop(tim)
    delete(tim)
end

function evolve(W,MeanM,SDM)
    global a
    global as
    N=size(a,2);
    AllR=NaN(50,4*N);
    p1=ceil(rand(1,N*3)*N);
    p2=ceil(rand(1,N*3)*N);
    Xo=repmat(rand(1,3*N),5,1).*repmat(linspace(1,5,5)',1,3*N)>1.5;
    b=Xo.*a(:,p1)+(1-Xo).*a(:,p2)+(randn(5,3*N).*(rand(5,3*N)>0.8));
    c=Xo(2:4,:).*as(:,p1)+(1-Xo(2:4,:)).*as(:,p2);
    X1=(rand(1,3*N)>0.8);
    b(:,X1)=flipud(b(:,X1));
    C=c-repmat(linspace(2,34,3)',1,3*N).*(c>0);
    C=C+round(randn(3,3*N).*(rand(3,3*N)<0.05))+round(randn(3,3*N).*(rand
        (3,3*N)<0.2).*(C>0));
    c=mod(C,16)+repmat(linspace(2,34,3)',1,3*N).*(C~=0);
    c(:,X1)=flipud((51-c(:,X1)).*(c(:,X1)>0));
    b(:,3*N+1:4*N)=a;
    c(:,3*N+1:4*N)=as;
    for j=1:4*N
            [AllR(:,j),H]=Bezier2(b(:,j),c(:,j)');
    end
    r=min(AllR);
    AllR=AllR-repmat(r,50,1)+0.5;
    Ind=Difference(AllR);
    m=(MeasureFinder(b,4*N,c)-repmat(MeanM,1,8*N))./repmat(SDM,1,8*N);
    Score=sum(m(:,1:4*N).*repmat(W,1,4*N));
    MR=max(AllR);
    Score(MR>2)=-40;
    for j=1:N
            [S,I]=max(Score);
        a(:,j)=b(:,I);
        as(:,j)=c(:,I);
```

```
        Score(I)=NaN;
        Score=Score+((Ind(I,:)<10).*(Ind(I,:)-10));
    end
end

function addA
    global A
    global a
    global D
    global W
    global Ma
    global a1
    global MeanM
    global SDM
    global w
    global as1
    global as
    global A8a
    global A8b
    A1=[A8a,A(:,1:D),A8b,A(:,D+1:2*D)];
    g=A1(11,:)+1;
    G=A1(10,:)./g;
    d=size(A1,2)/2;
    W=WeightFinder([G(1:d)',G(d+1:2*d)'],A1(12:20,:),d);
    m=A1(12:20,:);
    MeanM=mean(m,2);
    SDM=std(m,0,2);
    M=MeasureFinder(a(:,1),1,as(:,1));
    B=A(:,1:D);
    B(1,D+1:2*D+2)=linspace(D+1,2*D+2,D+2);
    B(2:6,D+1)=a(:,1);
    B(7:9,D+1)=as(:,1);
    B(10:11,D+1)=0;
    B(12:20,D+1)=M(:,1);
    B(2:20,D+2:2*D+1)=A(2:20,D+1:2*D);
    B(2:6,2*D+2)=flipud(a(:,1));
    B(7:9,2*D+2)=flipud((51-as(:,1)).*(as(:,1)>0));
    B(10:11,2*D+2)=0;
    B(12:20,2*D+2)=M(:,2);
    A=B;
    D=D+1;
    w(:,D-30+1)=W;

    AllR=NaN(50,200);
    for j=1:200
        [AllR(:,j),H]=Bezier2(a1(:,j),as1(:,j)');
    end
    r=min(AllR);
    AllR=AllR-repmat(r,50,1);
    Ind=Difference(AllR);
    Score=sum(((Ma(:,1:200)-repmat(MeanM,1,200))./repmat(SDM,1,200)).*
        repmat(W,1,200));
    if D==30
        a2=NaN(5,20);
        as2=NaN(3,20);
        for j=1:20
            [S,I]=max(Score);
            a2(:,j)=a1(:,I);
            as2(:,j)=as1(:,I);
            Score(I)=NaN;
            Score=Score+((Ind(I,:)<10).*(Ind(I,:)-10));
        end
        a=a2;
        as=as2;
```

```
    else
        a2=NaN(5,5);
        as2=NaN(3,5);
        for j=1:5
            [S,I]=max(Score);
            a2(:,j)=a1(:,I);
            as2(:,j)=as1(:,I);
            Score(I)=NaN;
            Score=Score+((Ind(I,:)<10).*(Ind(I,:)-10));
        end
        a=[a(:,2:16),a2];
        as=[as(:,2:16),as2];
    end
end

function AddAllA
    global A
    global a
    global D
    global as
    M=MeasureFinder(a(:,1:10),10,as(:,1:10));
    B=A(:,1:D);
    B(1,D+1:2*D+20)=linspace(D+1,2*D+20,D+20);
    B(2:6,D+1:D+10)=a(:,1:10);
    B(7:9,D+1:D+10)=as(:,1:10);
    B(10:11,D+1:D+10)=0;
    B(12:20,D+1:D+10)=M(:,1:10);
    B(2:20,D+11:2*D+10)=A(2:20,D+1:2*D);
    B(2:6,2*D+11:2*D+20)=flipud(a(:,1:10));
    B(7:9,2*D+11:2*D+20)=flipud((51-as(:,1:10)).*(as(:,1:10)>0));
    B(10:11,2*D+11:2*D+20)=0;
    B(12:20,2*D+11:2*D+20)=M(:,11:20);
    A=B;
    D=D+10;
end

function Best1(source,eventdata)
    global A
    global p
    A(10,p(1))=A(10,p(1))+3;
    A(10,p(2))=A(10,p(2))-1;
    A(10,p(3))=A(10,p(3))-1;
    A(10,p(4))=A(10,p(4))-1;
    A(11,p)=A(11,p)+1;
    stop(tim)
    SetPlots(A)
end
function Best2(source,eventdata)
    global A
    global p
    A(10,p(1))=A(10,p(1))-1;
    A(10,p(2))=A(10,p(2))+3;
    A(10,p(3))=A(10,p(3))-1;
    A(10,p(4))=A(10,p(4))-1;
    A(11,p)=A(11,p)+1;
    stop(tim)
    SetPlots(A)
end
function Best3(source,eventdata)
    global A
    global p
    A(10,p(1))=A(10,p(1))-1;
    A(10,p(2))=A(10,p(2))-1;
    A(10,p(3))=A(10,p(3))+3;
```

```
    A(10,p(4))=A(10,p(4))-1;
    A(11,p)=A(11,p)+1;
    stop(tim)
    SetPlots(A)

end
function Best4(source,eventdata)
    global A
    global p
    A(10,p(1))=A(10,p(1))-1;
    A(10,p(2))=A(10,p(2))-1;
    A(10,p(3))=A(10,p(3))-1;
    A(10,p(4))=A(10,p(4))+3;
    A(11,p)=A(11,p)+1;
    stop(tim)
    SetPlots(A)
end
function SetPlots(A)
    global p
    global D
    [k,l]=min(A(11,:));
    j=A(10,l);
    s=0;
    q=0;
    while s<4
            I=A(10,:)<=j+q & A(10,:)>=j-q;
            B=A(1:9,I);
            s=size(B,2);
            q=q+1;
    end
    p=zeros(1,4);
    AllR=zeros(50,4);
    AllH=zeros(50,4);
    for i=1:4
            t=ceil(rand*(s+1-i));
            p(i)=B(1,t);
            [R,H]=Bezier2(B(2:6,t),B(7:9,t)');
            AllR(:,i)=R-min(R)+0.5;
            AllH(:,i)=H-min(H);
            B=[B(:,1:t-1),B(:,t+1:s+1-i)];
    end

    phi=linspace(0,2*pi,50);
    x=sin(phi);
    y=cos(phi);
    z=ones(1,50);
    C=ones(50,50,3)*0.9;
    C(:,:,1)=C(:,:,1)*0.85;

    Z=z'*AllH(:,1)';
    X=x'*AllR(:,1)';
    Y=y'*AllR(:,1)';
    surf(P1,X,Y,Z,C,'FaceLighting','phong','ButtonDownFcn',@Best1)
    set(P1,'DataAspectRatio',[1 1 1],'Visible','off');
    shading(P1,'interp')
    light('Parent',P1,'Position',[-5,0,5]);
    set(P1,'view',[0 8])

    Z=z'*AllH(:,2)';
    X=x'*AllR(:,2)';
    Y=y'*AllR(:,2)';
    surf(P2,X,Y,Z,C,'FaceLighting','phong','ButtonDownFcn',@Best2)
    set(P2,'DataAspectRatio',[1 1 1],'Visible','off');
    shading(P2,'interp')
```

```
        light('Parent',P2,'Position',[-5,0,5]);
        set(P2,'view',[0 8])

        Z=z'*AllH(:,3)';
        X=x'*AllR(:,3)';
        Y=y'*AllR(:,3)';
        surf(P3,X,Y,Z,C,'FaceLighting','phong','ButtonDownFcn',@Best3)
        set(P3,'DataAspectRatio',[1 1 1],'Visible','off');
        shading(P3,'interp')
        light('Parent',P3,'Position',[-5,0,5]);
        set(P3,'view',[0 8])

        Z=z'*AllH(:,4)';
        X=x'*AllR(:,4)';
        Y=y'*AllR(:,4)';
        surf(P4,X,Y,Z,C,'FaceLighting','phong','ButtonDownFcn',@Best4)
        set(P4,'DataAspectRatio',[1 1 1],'Visible','off');
        shading(P4,'interp')
        light('Parent',P4,'Position',[-5,0,5]);
        set(P4,'view',[0 8])

        k=std(A(10,:))*D/200;
        if k==0
            k=0.01;
        end
        if k>=1
                k=1;
        end
        axes(Bar)
        rectangle;
        rectangle('Position',[0 0 1 k],'FaceColor','b','EdgeColor','none');
        set(Bar,'XLim',[-0.1 1.1],'YLim',[0 1.01])

        start(tim)
    end

end
```

MeasureFinder.m

```
function M=MeasureFinder(A,D,B)
eps=0.1;
zeta=0.1;
M=zeros(9,D);
for k=1:D
    [R,H]=Bezier2(A(:,k),B(:,k)');
    r=min(R);
    R=R-r+0.5;
    b=[1 B(:,k)' 50];
    [c,n]=Acomplexity(R,H,b);
    N=sum(n);
    c1=c(:,1);
    Rc=R(c1);
    Hc=H(c1(2:N))-H(c1(1:N-1));

    d=exp(N-3);
    OHsI=min(4,HSumI(Rc,N,eps))/d;
    OVsI=min(4,VSumI(Hc,N,eps))/d;
    OPsI=min(2,PSumI2(Hc,Rc,N,eps))/d;
    OTsI=min(4,TSumI(R,H,c1,N,zeta))/d;
    M(1,k)=OHsI+OVsI+OPsI+OTsI;
```

```
    C=curvature(R,H);
    M(2,k)=sum(C);

    K=1;
    nn=n(3)+n(4);
    if (nn)>1
        for j=1:sum(n)
            if c(j,2)==3 || c(j,2)==4
               S(K)=c(j,1);
               K=K+1;
            end
        end
        M(3,k)=min(H(S(2:nn))-H(S(1:nn-1)));
    else
        M(3,k)=5;
    end

    [R1,I]=max(R);
    if I~=1 && I~=50
        M(4,k)=1;
    end

    dh=H(c(2:N,1))-H(c(1:N-1,1));
    dr=abs(R(c(2:N,1))-R(c(1:N-1,1)));
    M(5,k)=prod(dh/max(dh))+prod(dr/max(dr));

    if n(4)==0
        M(6,k)=1;
    end

    dH1=H(50)-H(49);
    dR1=R(50)-R(49);
    dH2=H(1)-H(2);
    dR2=R(1)-R(2);
    topangle=dR1/sqrt(dH1^2+dR1^2);
    bottomangle=dR2/sqrt(dH2^2+dR2^2);
    M(7,k)=acos(bottomangle)-acos(topangle);

    tw=sum((c(:,1)>25).*(c(:,1)-25));%.*(R(c(:,1))'));
    bw=sum((c(:,1)<=25).*(26-c(:,1)));%.*(R(c(:,1))'));
    M(8,k)=tw-bw;

    f=CoM(R,H);
    M(9,k)=f-25;
end
M(1:6,D+1:2*D)=M(1:6,1:D);
M(7:9,D+1:2*D)=-M(7:9,1:D);
```

Bezier2.m

```
function [R,H]=Bezier2(A,B)
B=[1,B,50];
b=linspace(1,5,5);
b=b(B~=0);
for i=1:size(b,2)-1;
    [R(B(b(i)):B(b(i+1))),H(B(b(i)):B(b(i+1)))]=BezierCurve([A(b(i):b(i+1)),
        linspace(B(b(i))/10,B(b(i+1))/10,b(i+1)-b(i)+1)'],B(b(i+1))-B(b(i))+1)
        ;
end
```

BezierCurve.m

```
function [X,Y]=BezierCurve(A,points)
a=size(A);
t=linspace(0,1,points);
X=zeros(1,points);
Y=X;
n=a(1)-1;
for i=0:n
    X=X+nchoosek(n,i).*(1-t).^(n-i).*t.^i.*A(i+1,1);
    Y=Y+nchoosek(n,i).*(1-t).^(n-i).*t.^i.*A(i+1,2);
end
```

AComplexity.m

```
function [a,n]=Acomplexity(X,Y,B)
derivX=X(2:50)-X(1:49);
derivY=Y(2:50)-Y(1:49);
n=zeros(4,1);
n(1)=1;
p=1;
t=1;
B=(B-1).*(B>0);
while p+t<=5
    while B(p+t)==0
        t=t+1;
    end
    StartX=B(p)+1;
    EndX=B(p+t)-2;
    a(sum(n),1)=B(p)+1;
    a(sum(n),2)=4;
    n(4)=n(4)+1;
    if t>1
        for k=StartX:EndX
            theta1=cross([derivX(k),derivY(k),0],[derivX(k+1),derivY(k+1),0])
                ;
            theta2=cross([derivX(k+1),derivY(k+1),0],[derivX(k+2),derivY(k+2)
                ,0]);
            if sign(derivX(k))~=sign(derivX(k+1))
                a(sum(n),1)=k+1;
                a(sum(n),2)=2;
                n(2)=n(2)+1;
            end
            if sign(theta1(3))~=sign(theta2(3))
                a(sum(n),1)=k+1;
                a(sum(n),2)=3;
                n(3)=n(3)+1;
            end
        end
    end
    p=p+t;
    t=1;
end
a(1,2)=1;
n(4)=n(4)-1;
n(1)=2;
a(sum(n),1)=50;
a(sum(n),2)=1;
```

HSumI.m

```
function h=HSumI(Ra,n,eps)
h=0;
Ind=zeros(n,1);
for i=1:n
    for j=i+1:n
        H=RFun(Ra(i),Ra(j),eps);
        if H>0
            Ind(i)=Ind(i)+1;
            Ind(j)=Ind(j)+1;
            if Ind(i)<3 && Ind(j)<3
                h=h+H;
            end
        end
    end
end
```

VSumI.m

```
function v=VSumI(Ha,n,eps)
v=0;
Ind=zeros(n,n);
for i=1:n-1
    for j=1:i-1
        for k=i:n-1
            l1=sum(Ha(j:i-1));
            l2=sum(Ha(i:k));
            V=RFun(l1,l2,eps);
            if V>0
                Ind(j,i)=Ind(j,i)+1;
                Ind(i,k)=Ind(i,k)+1;
                if Ind(j,i)<3 && Ind(i,k)<3
                    v=v+V;
                end
            end
        end
    end
end
```

PSumI.m

```
function p=PSum(Ha,Ra,n,eps)
p=0;
for i=1:n-1
    for j=1:i-1
        for k=i:n-1
            l1=sum(Ha(j:i-1));
            l2=sum(Ha(i:k));
            p=p+RFun(Ra(i),l1,eps)+RFun(Ra(i),l2,eps);
        end
    end
end
```

TSumI.m

```
function q=TSumI(R,H,a,n,zeta)
q=0;
derivR=R(2:50)-R(1:49);
derivH=H(2:50)-H(1:49);
da(1,1:n)=[derivR(1),(derivR(a(2:n-1))+derivR(a(2:n-1)-1))./2,derivR(49)];
da(2,1:n)=[derivH(1),(derivH(a(2:n-1))+derivH(a(2:n-1)-1))./2,derivH(49)];
Ind=zeros(n,1);
for i=1:n
    for j=i+1:n
        ang=real(acos(dot(da(:,i),da(:,j))/(norm(da(:,i))*norm(da(:,j)))));
        vert1=real(acos(dot(da(:,i),[0,1])/(norm(da(:,i)))));
        vert2=real(acos(dot(da(:,j),[0,1])/(norm(da(:,j)))));
        if ang<=zeta && vert1>=zeta && vert2>=zeta
            Q=1;
            Ind(i)=Ind(i)+1;
            Ind(j)=Ind(j)+1;
            if Ind(i)<3 && Ind(j)<3
               q=q+Q;
            end
        end
        if ang<=pi/2+zeta && ang>=pi/2-zeta
            Q=1;
            Ind(i)=Ind(i)+1;
            Ind(j)=Ind(j)+1;
            if Ind(i)<3 && Ind(j)<3
               q=q+Q;
            end
        end
    end
    for j=1:n
        x=[R(a(i));H(a(i))-H(a(j))];
        angx=acos(dot(da(:,i),x)/(norm(da(:,i))*norm(x)));
        if angx<=zeta || angx>=pi-zeta
            q=q+1;
        end
    end
end
```

RFun.m

```
function r=RFun(l1,l2,eps)
R=zeros(3,1);
phi=(1+sqrt(5))/2;
rat=[phi,1,1/phi];
for i=1:2
    K=abs(rat(i)*l1/l2-1);
    if K>=0 && K<=eps
        R(i+3)=1;
    end
end

r=max(R);
```

curvature.m

```
function C=curvature(R,H)
derivR=R(2:50)-R(1:49);
derivH=H(2:50)-H(1:49);
dR=derivR./sqrt(derivR.^2+derivH.^2);
```

```
dH=derivH./sqrt(derivR.^2+derivH.^2);
angle=real(acos(dot([dR(1:48);dH(1:48)],[dR(2:49);dH(2:49)])));
C=2*sin(angle(1:48))./sqrt((derivR(1:48)+derivR(2:49)).^2+(derivH(1:48)+
    derivH(2:49)).^2);
for k=1:48
    theta=cross([derivR(k),derivH(k),0],[derivR(k+1),derivH(k+1),0]);
    C(k)=C(k)*sign(theta(3));
end
```

CoM.m

```
function f=CoM(R,H)
n=size(R,2);
dh=H(2:n)-H(1:n-1);
F=sum(0.5*pi*R.^2.*([0,dh]+[dh,0]));
f=1;
t=0;
while t<F/2
    f=f+1;
    t=sum(0.5.*pi*R(1:f).^2.*([0,dh(1:f-1)]+[dh(1:f-1),0]));
end
f=f-1;
```

WeightFinder.m

```
function W=WeightFinder(DataY,M,D)
B1=NaN(1,D);
for i=1:D
        B1(i)=max(DataY(i,:));
end
A=DataY(:,1)-DataY(:,2);


[Spearman,Pval]=corr(M(1:5,1:D)',B1','type','spearman','tail','right');
Spearman(6)=corr(M(6,1:D)',B1','type','spearman','tail','both');
[Spearman(7:9),Pval(7:9)]=corr(M((7:9),1:D)',A,'type','spearman','tail','
    right');
Spearman=(Pval<0.2).*Spearman;

W=Spearman;
```

Appendix D

Appendix: Experimental Results Further Information

D.1 Results

D.1.1 Weights

The global measure was found as the program progressed from the correlation of all the previous vase measures with the scores. It was calculated at the beginning and end of each person's session as well as every time a new vase was added. Figure D.1 shows the changing weights over the experiment. The initial weights were all set to 1 and this drops off very quickly (1 indicates a perfect correlation). As expected the earlier weights vary considerably as there are less results to be correlated, these then stabilise later. Some exceptions to this are the late disappearance of the centre of mass measure and the late reappearance of the balance of features, modified Birkhoff and product of the distances. It is possible these all had small positive correlations for the entire experiment but the confidence was low and would only appear above 80% occasionally. This idea is reinforced for the Birkhoff measure as it appeared and disappeared several times before stabilising.

The dominant measures were the edge angles, presence of sharp angles and spacing of points of inflection measures. The presence of sharp angles was the only weight allowed to be negative due to the lack of previous evidence but this confirms that the correlation is strongly positive.

Figure D.1: Variation of weights over test period

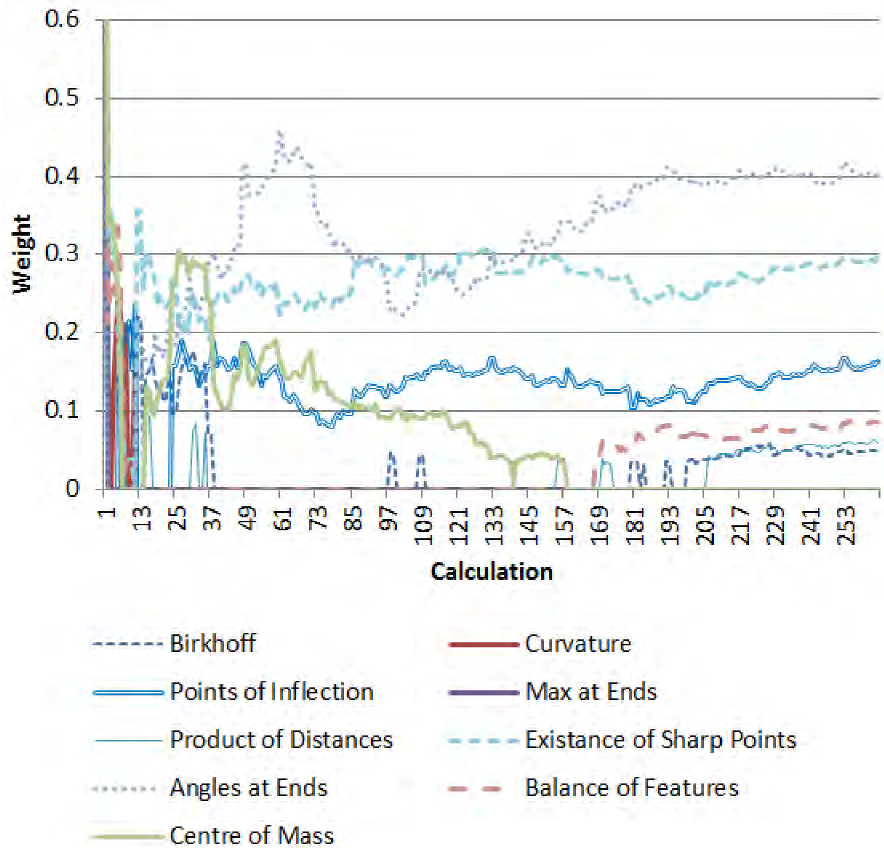

Figure D.2: First Convergent Vase Type

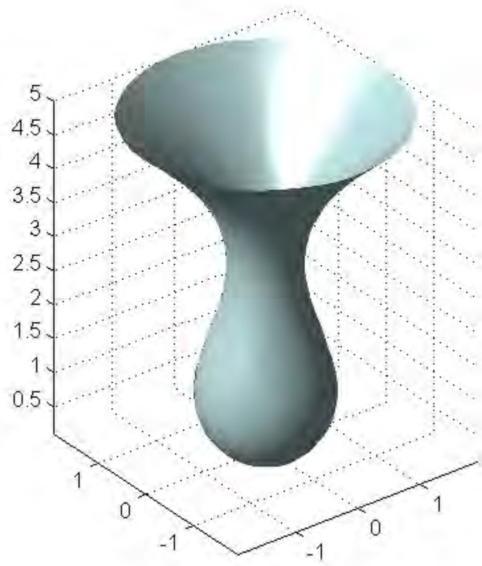

Figure D.3: Final Convergent Vase Type

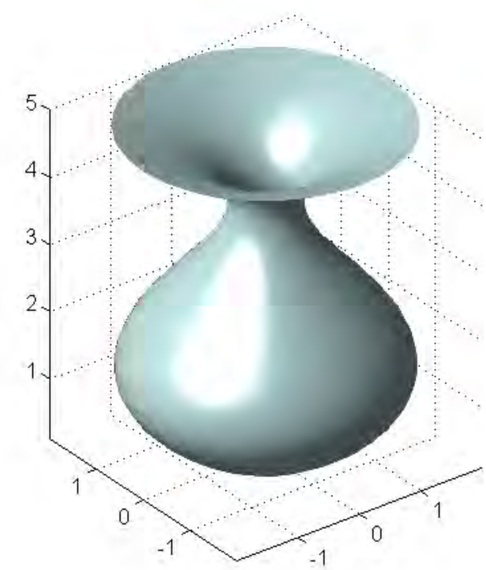

Looking at the vases evolved by the measures as they begin to stabilise it is obvious that this combination has an 'ideal' vase type. The evolved vases began to stabilise during the test period of volunteer 4 (weights 37 onwards). A vase of this type is shown in figure D.2 It is apparent that by this point there are 4 dominant measures, the 3 final dominant measures as well as the centre of mass. The only significant change to this ideal vase occurs during the test period of volunteer 11 (weights 121 onwards). It is during this period that the distribution of weight measure begins to reduce, this has the impact of allowing the base of the vase to increase in size, producing designs such as the one in figure D.3. The reintroduction of the other 3 measures later on has no visible effect on the 'ideal' vase as these were also all easily satisfied by variations of this vase.

D.1.2 Evolved vs. Random and Birkhoff Designs

To see the effect of the new fitness function the results were divided into their separate types and the average score found. The vases were supposed to be evolved to be the right way up and therefore every set was compared across the generated orientation and the inverted orientation separately as well as finding the maximum of these two scores.

The evolved vases ability tend to the correct orientation was obviously an advantage when the full sets were studied (scores are shown in table D.1). The evolved group had a high average score while its inverse group had a very low average score. As expected the scores for the Birkhoff and random designs had smaller variation between the orientations.

When the best score for each orientation is taken then the average scores have a clear pattern with random giving the lowest scores and the evolved the highest, indicating that the style component also has good influence. It is surprising how high the scores for the

Type	Original Orientation	Inverted Orientation	Preferred Orientation
Random	-0.0693	-0.1298	0.1829
Birkhoff Evolved	0.1894	0.0565	0.3999
New Evolved	0.3162	-0.2569	0.4415

Table D.1: Average Scores for different types of Design

Type	Original Orientation	Inverted Orientation	Preferred Orientation
Random	-0.2620	-0.2802	-0.0470
Birkhoff Evolved	-0.3047	-0.3490	-0.1068
New Evolved	-0.3282	-0.2549	0.0011

Table D.2: Average Scores for different types of Designs including Angles

Type	Original Orientation	Inverted Orientation	Preferred Orientation
Random	0.0913	-0.0045	0.3745
Birkhoff Evolved	0.5719	0.3705	0.7922
New Evolved	0.3359	-0.2569	0.4549

Table D.3: Average Scores for different types of Designs excluding Angles

Birkhoff designs are considering the apparent lack of correlation in the initial tests and the low weight assigned to the modified version in the full test.

To further investigate the reappearance of the Birkhoff measure the results were split further, separating those with sharp angles from those that did not. This was done for several reasons; first it was observed as being a property that was considered unattractive by the large weight assigned to that measure, it also directly affected the Birkhoff designs as evolving for Birkhoff's measure and allowing sharp angles produced designs that nearly always had sharp angles, to balance these a set of Birkhoff designs where angles were not allowed were also used, this divided the Birkhoff set in two. Finally the angles were the only major difference between the initial tests and the new designs, therefore their presence may contribute to the difference between the initial and new results.

The average scores for all those with one or more angles are shown in table D.2. As you can see they all have low scores, regardless of orientation or type. It should be noted that there were very few (13) evolved designs in this set as the measure preventing angles appeared very early.

However the results from those with no angles (table D.3) are very interesting. Once again we can see the clear advantage of the (almost) correct orientation provided by the evolved designs and the little change in orientation of the random designs but now the

Birkhoff scores have really changed. Just by cutting out the angles the score of the Birkhoff evolved designs has leapt to 0.572, and interestingly there is a definite preference for the original orientation. It is unknown why this should be, it is possible there is a bias within the evolution that is not apparent, potentially from the variation in the points of inflection location, but it is unknown why this should be bias towards evolving correctly oriented designs. It may also be possible that a bias has appeared due to some part of the testing process. This bias would be lost in the expected difference of the evolved set but should be apparent in the random set. All of the random sets do have slightly better scores for the original orientation but not high enough to be noticeable. For the time being this remains unexplained.

By cutting out the designs with angles we have now changed the order of the maximum scores. The random and Birkhoff scores have significantly increased so that now the random scores are only slightly worse than the evolved and the Birkhoff is much better than them both.

From this we can conclude that most of the benefit of the evolved designs comes from their orientation measures and the exclusion of sharp angles. The Birkhoff measure obviously provides an advantage over random designs and if correctly orientated and smooth provides a better measure than the evolved designs. This agrees with Birkhoff's original measure that assumed correct orientation and stipulated that the curve should only contain straight edges (a frequent side effect of the sharp angles) under certain circumstances.

The distribution of scores was also studied and is shown in figure D.4. This shows the percentage of each type of design within each section of the scores. The lowest scores have roughly equal numbers of Birkhoff, random and pre–set designs, low evolved designs and high inverted designs. As expected the percentage of evolved designs increases as the scores increase and the inverted decrease. The others are more constant but a slow increase of Birkhoff designs is apparent (these now include those with and without angles and both orientations).

Figure D.4: Percentage of each Type

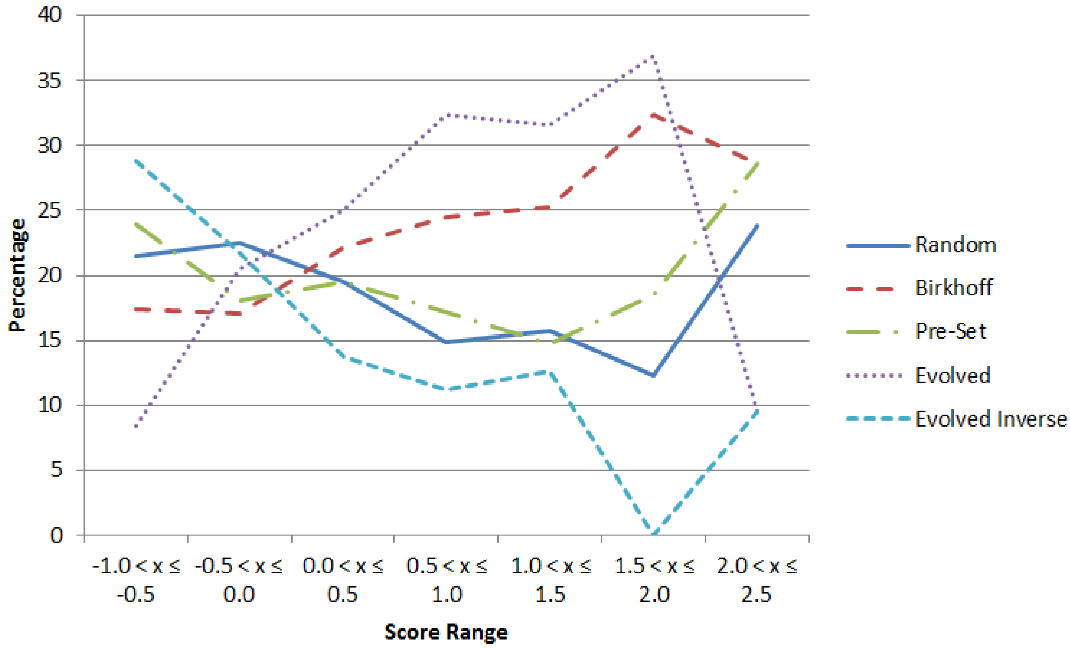

However this graph also reveals that the very highest scores (roughly the top 1%) have a much more varied distribution. This final set required that the volunteer would chose that particular design every time it appeared and therefore the large reduction of evolved is likely to be a result of the score sharing, where choosing the same design every time is nearly impossible due to the visual similarity. The impact of the score sharing is highlighted by the fact that 8 individuals gave their highest scores to designs very similar to this recurring vase, some of which were evolved and others that were random or Birkhoff. But these were often joint highest thus giving overall scores of 1.5 – 2, placing them in the next group down.

The reduction in evolved designs is balanced by the increase in the other types where their unique designs could easily be attractive enough to catch the volunteers eye and be selected consistently, thus allowing them to get perfect scores.

D.1.3 Set Design Results

The following figures show the 10 set designs along with their scores from all 22 volunteers. The blue line (with circles) shows the scores for the left hand orientation and the red

Figure D.5: Set Design 1

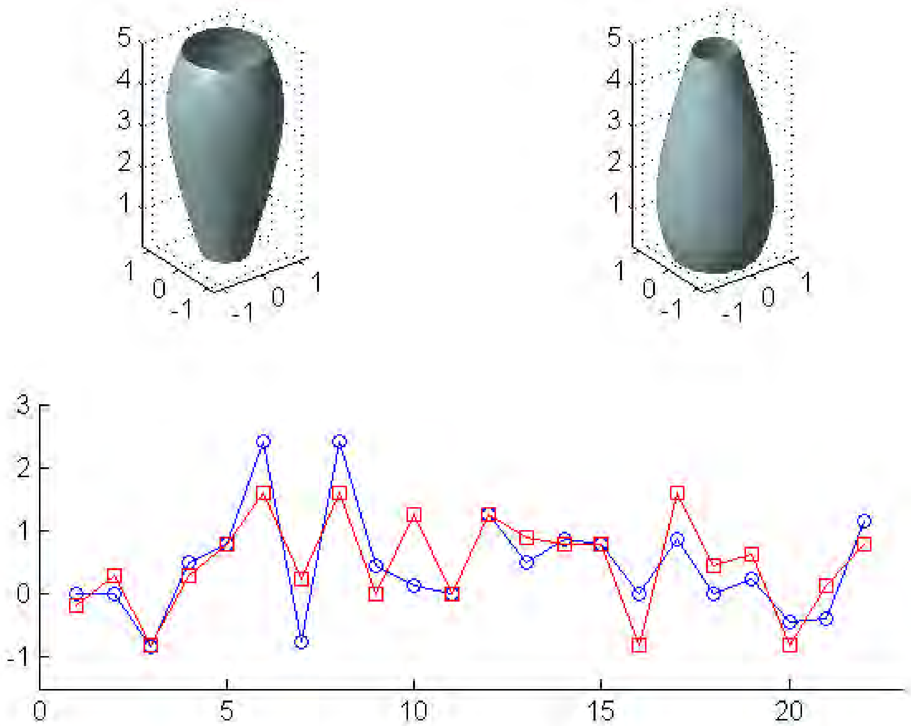

(squares) shows the right. These show that some of the vases were mostly liked (e.g set design 2) while others were disliked (7 and 10). However others had a greater range of scores and no design was liked or disliked by every volunteer. With regards to orientation it is apparent that the left hand design for 2 was generally preferred over the right, with the opposite being true for 4. Set design 1 appears to divide opinions as to its style but most people had consistent opinions in either orientation. Again there are always some that disagree and many have no apparent preferred orientation.

This indicates that there is a large range of personal tastes within the group. This would make any generalisations less clear but if there are patterns to be found they come from a wide variety of tastes and therefore will be more significant than if the test group had been chosen from individuals with similar tastes.

Figure D.6: Set Design 2

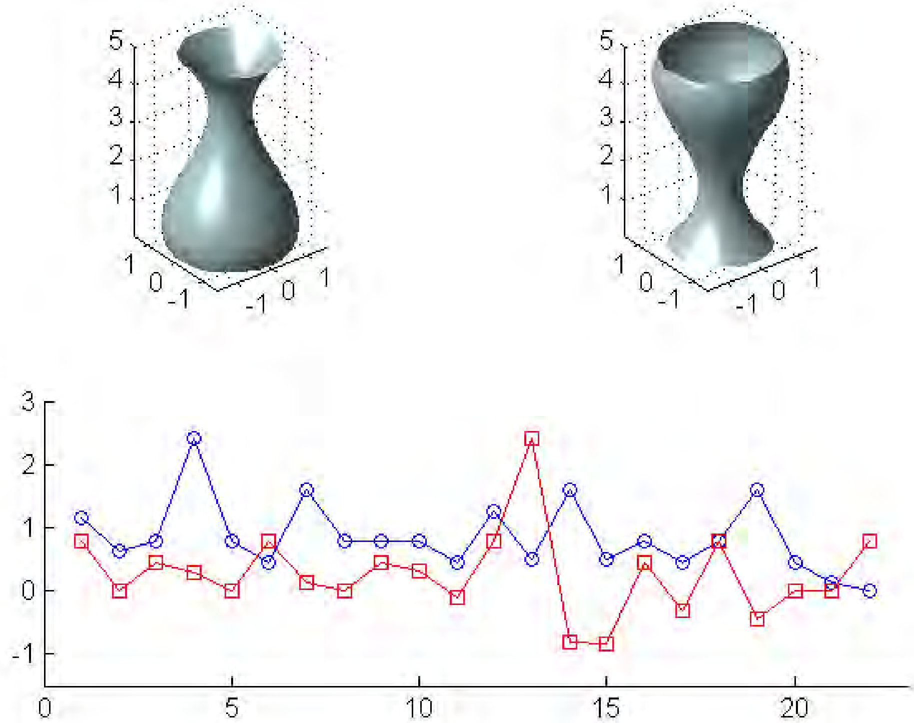

Figure D.7: Set Design 3

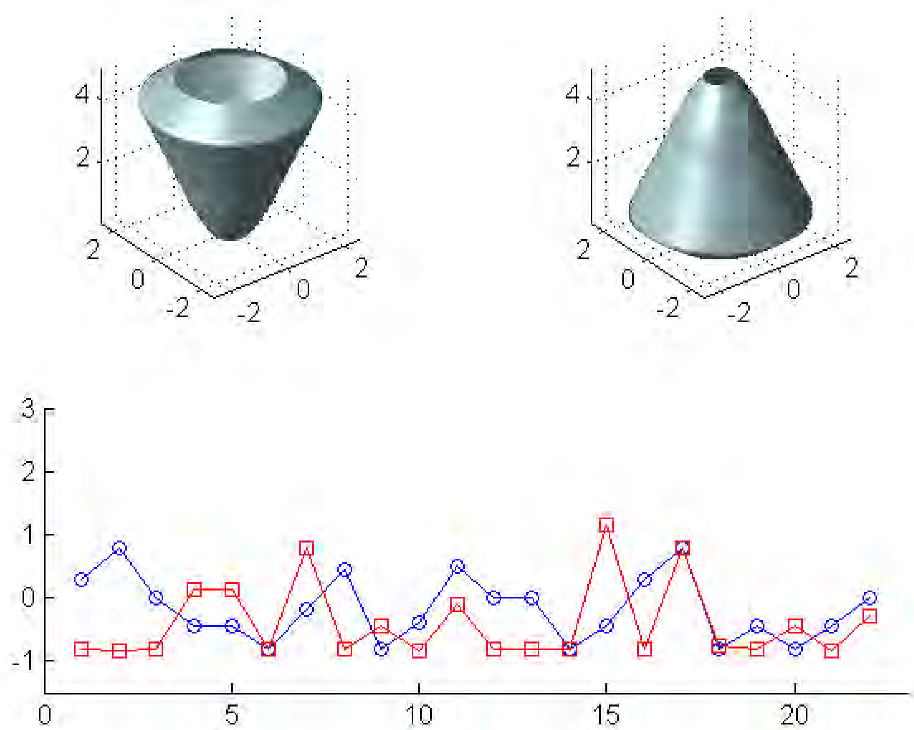

Figure D.8: Set Design 4

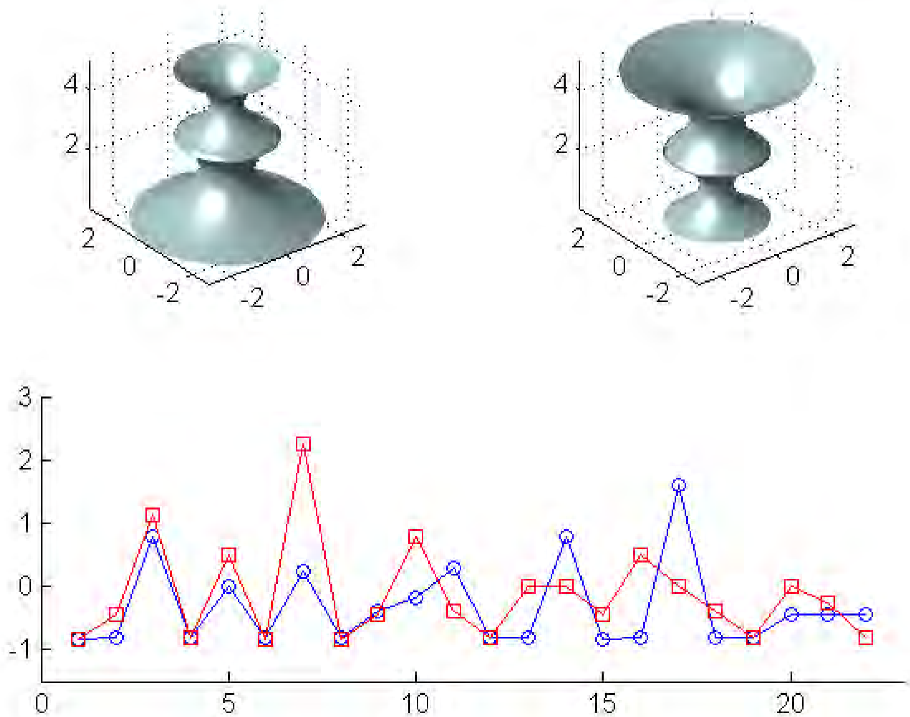

Figure D.9: Set Design 5

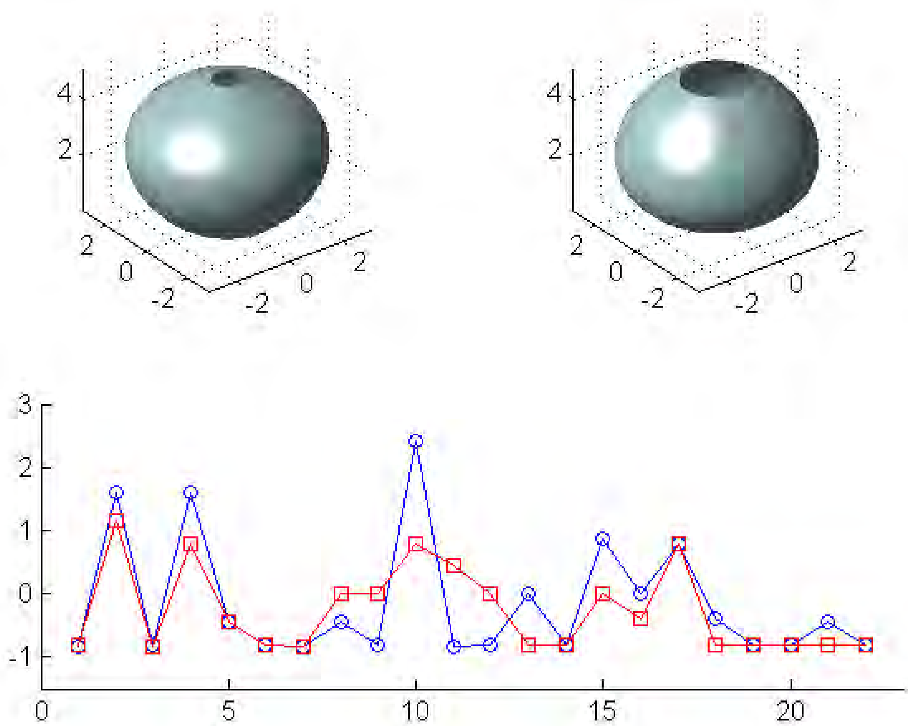

Figure D.10: Set Design 6

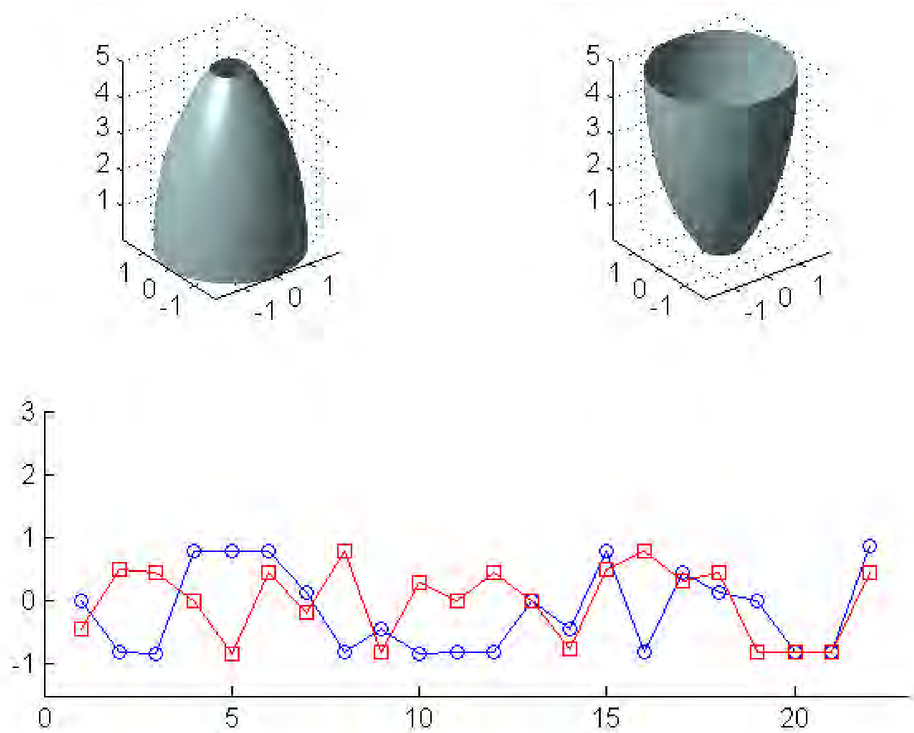

Figure D.11: Set Design 7

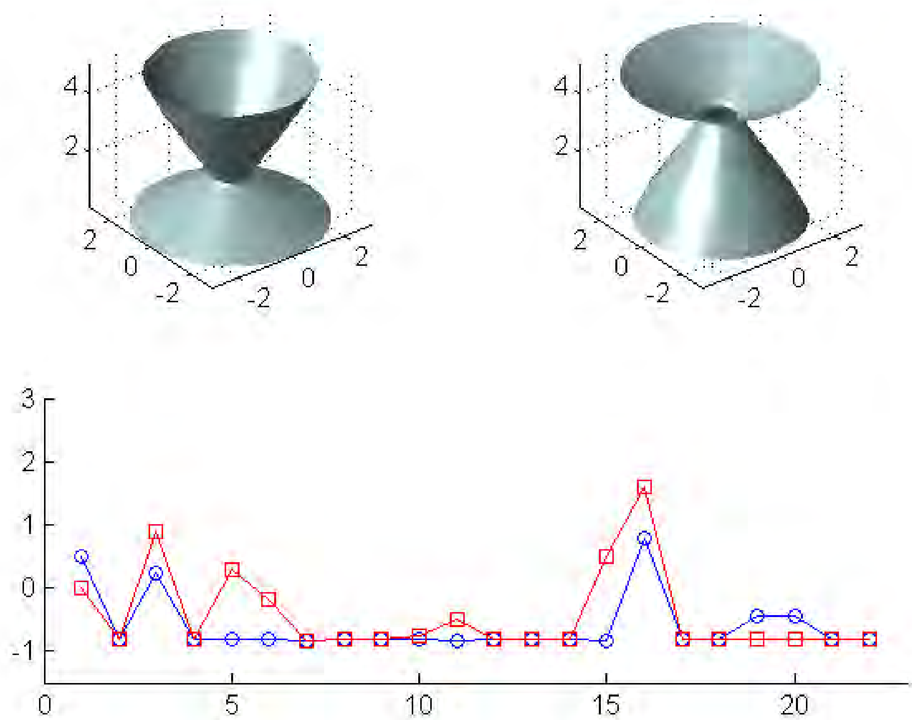

Figure D.12: Set Design 8

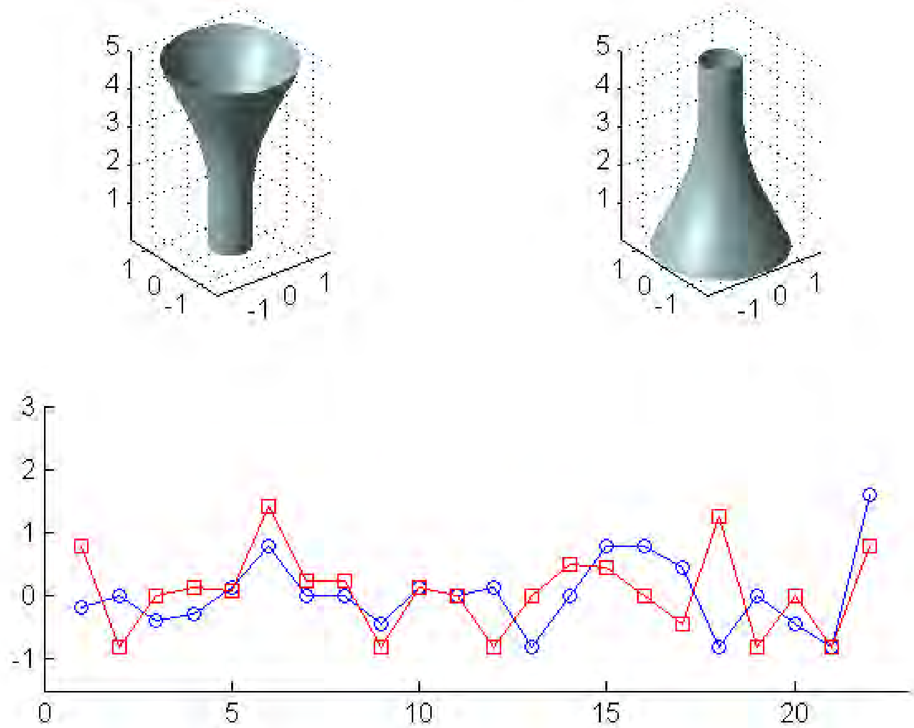

Figure D.13: Set Design 9

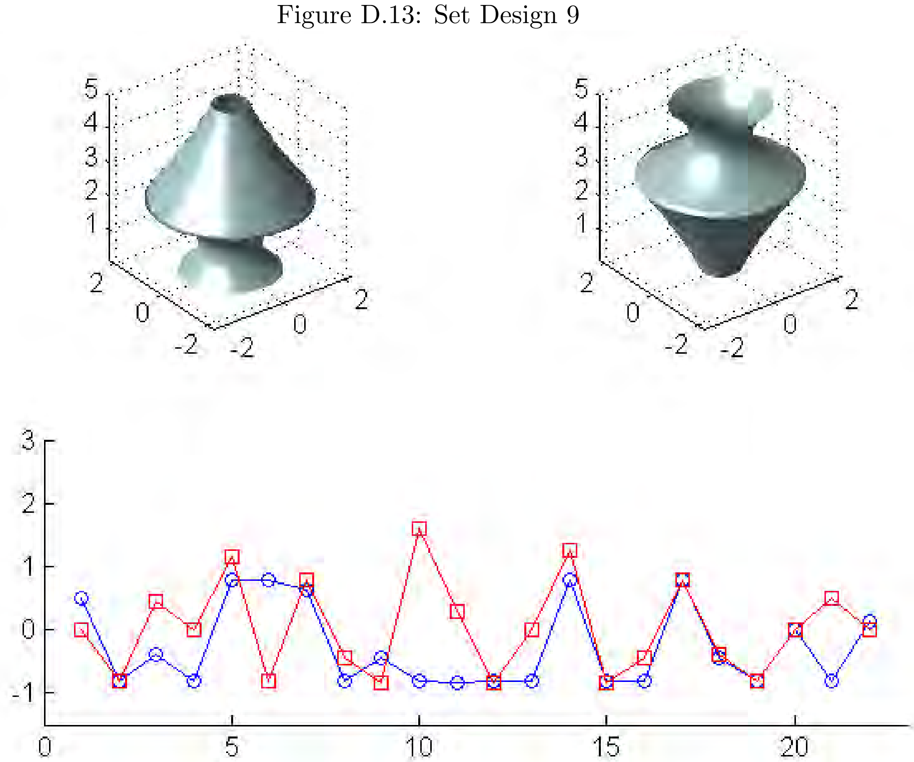

Figure D.14: Set Design 10

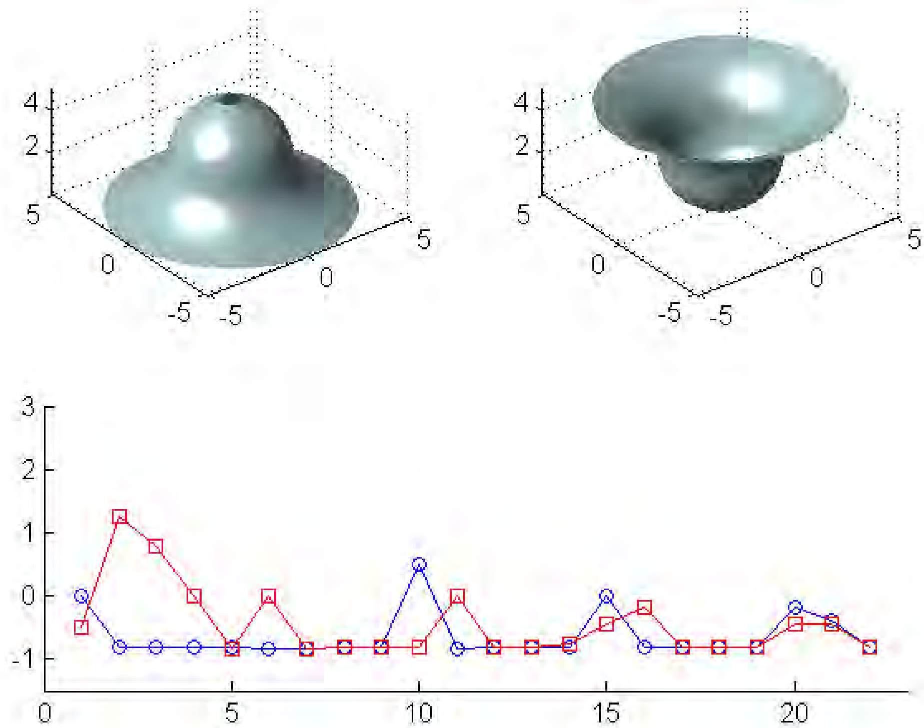

D.1.4 Extra Test Results

This shows the ranks given to the final set of vases tested as well as an average score derived from them. As the vases evolved by individual measures were different for each individual their results have been combined to produce a single score.

Vase Type	Rank													Score
New: Weight -0.9190	6	8	5	9	10	10	9	7	10	9	7	10	3	27
New: Weight -0.3578	5	6	10	10	9	4	3	10	5	10	5	8	10	35
Individual Measures	10	10	9	4	6	9	6	8	7	6	6	3	8	43
Individual Measures	9	1	7	3	7	8	4	9	8	7	10	2	7	
Original Evolved Design	4	4	3	7	8	6	1	6	9	8	4	4	9	57
New: Weight 0.7604	3	6	8	8	4	5	8	2	2	5	9	9	2	59
New: Weight 0.5045	8	1	6	6	5	7	10	5	4	4	8	6	1	59
New: Weight 0.2368	2	6	2	2	2	2	2	4	1	3	3	7	6	88
Previous Best	1	8	4	5	1	3	5	3	3	1	1	1	5	89
New: Weight 0.3103	7	1	1	1	3	1	7	1	6	2	2	5	4	89
Score of # 1	9	6	?	9	8	9	8	9	8	8	8	8	8	
Score of # 10	1	0	?	6	1	2	2	1	4	1	0	1	3	

Table D.4: Ranks and overall scores of the further test results

The three top scoring results were two of the new designs, both with positive weights for the variable, and the favourite design from the previous test. The next three consisted of two positive weighted new designs and the old evolved design, which was also potentially evolved when the variable weight near 0. The least liked four were the two negative weighted designs and the two individual evolved designs. As these were different for each individual the rankings were highly variable; Some made it into the top 3 while many others were really disliked. This suggests that any measure derived from a single set of results is not sufficient for evolving consistently. The favourite for each individual was an evolved one for seven of the respondents and the original top scorer for five, the last volunteer chose the original convergent design.

D.2 Final M Files

EvolveExamples.m

```
clear all
w=xlsread('AllWeights.xlsx',1);
MS=xlsread('AllWeights.xlsx',2);
N=20;
a=randn(5,N);
as=(rand(3,N)<0.2).*(round(rand(3,N)*15)+repmat(linspace(2,34,3)',1,N));
MeanM=MS(1:9);
SDM=MS(10:18);
W=w;
W(4)=W(4)+randn/12

for i=1:100
    AllR=NaN(50,4*N);
    p1=ceil(rand(1,N*3)*N);
    p2=ceil(rand(1,N*3)*N);
    Xo=repmat(rand(1,3*N),5,1).*repmat(linspace(1,5,5)',1,3*N)>1.5;
    b=Xo.*a(:,p1)+(1-Xo).*a(:,p2)+(randn(5,3*N).*(rand(5,3*N)>0.8));
    c=Xo(2:4,:).*as(:,p1)+(1-Xo(2:4,:)).*as(:,p2);
    X1=(rand(1,3*N)>0.8);
    b(:,X1)=flipud(b(:,X1));
    C=c-repmat(linspace(2,34,3)',1,3*N).*(c>0);
    C=C+round(randn(3,3*N).*(rand(3,3*N)<0.05))+round(randn(3,3*N).*(rand
        (3,3*N)<0.2).*(C>0));
    c=mod(C,16)+repmat(linspace(2,34,3)',1,3*N).*(C~=0);
    c(:,X1)=flipud((51-c(:,X1)).*(c(:,X1)>0));
    b(:,3*N+1:4*N)=a;
    c(:,3*N+1:4*N)=as;
    for j=1:4*N
        [AllR(:,j),H]=Bezier2(b(:,j),c(:,j)');
    end
    r=min(AllR);
    AllR=AllR-repmat(r,50,1)+0.5;
```

```
    Ind=Difference(AllR);
    m=(MeasureFinderMod(b,4*N,c)-repmat(MeanM,1,8*N))./repmat(SDM,1,8*N);
    Score=sum(m(1:6,1:4*N).*repmat(W(1:6),1,4*N));
    Orient=(sign(sum(m(7:9,1:4*N).*repmat(W(7:9),1,4*N)))==-1).*ceil(rand
        (1,4*N)-0.5);
    Orient=logical(Orient);
    b(:,Orient)=flipud(b(:,Orient));
    c(:,Orient)=flipud((51-c(:,Orient)).*(c(:,Orient)>0));
    MR=max(AllR);
    Score(MR>2)=-40;
    Best(i)=max(Score);

    for j=1:N
        [S,I]=max(Score);
        if j==1
            D=I;
        end
        a(:,j)=b(:,I);
        as(:,j)=c(:,I);
        Score(I)=NaN;
        Score=Score+((Ind(I,:)<10).*(Ind(I,:)-10)/10);
    end

end

figure(1)
p=linspace(1,22,22);
phi=linspace(0,2*pi,50);
x=sin(phi);
y=cos(phi);
z=ones(1,50);
C=ones(50,50,3)*0.9;
C(:,:,1)=C(:,:,1)*0.85;
[R,H]=Bezier2(a(1:5,1),as(1:3,1)');
r=min(R);
R=R-r+0.5;
Z=z'*H;
X=x'*R;
Y=y'*R;
surf(X,Y,Z,C,'FaceLighting','phong')
axis equal
shading('interp')
light('Position',[-5,0,5]);
figure(2)
plot(Best)
Best(100)
figure(3)
plot(linspace(1,9,9),W,linspace(1,9,9),w)
F=xlsread('Final.xlsx',1);
f=[a(:,1);as(:,1);W];
xlswrite('Final.xlsx',[F,f],1);
G=xlsread('Final.xlsx',2);
xlswrite('Final.xlsx',[G,Best'],2);
```

NewRFun.m

```
function r=NewRFun(l1,l2,eps)

R=zeros(3,1);
for i=-1:1
    k=2^i;
    K=abs(k*l1/l2-1);
```

```
    if K>=0 && K<=eps
        R(i+2)=1;
    end
end

r=max(R);
```

NewMeasureFinderMod.m

```
function M=NewMeasureFinderMod(A,D,B)
eps=0.1;
zeta=0.1;
M=zeros(9,D);
for k=1:D
    [R,H]=Bezier2(A(:,k),B(:,k)');
    r=min(R);
    R=R-r+0.5;
    b=[1 B(:,k)' 50];
    [c,n]=Acomplexity(R,H,b);
    N=sum(n);
    c1=c(:,1);
    Rc=R(c1);
    Hc=H(c1(2:N))-H(c1(1:N-1));

    d=2*N;
    OHsI=min(4,HSumI(Rc,N,eps))/d;
    OVsI=min(4,VSumI(Hc,N,eps))/d;
    OPsI=min(2,PSumI(Hc,Rc,N,eps))/d;
    OTsI=min(4,TSumI(R,H,c1,N,zeta))/d;
    M(1,k)=OHsI+OVsI+OPsI+OTsI;

    C=curvature(R,H);
    M(2,k)=sum(abs(C));

    K=1;
    nn=n(3)+n(4);
    if (nn)>1
        for j=1:sum(n)
            if c(j,2)==3 || c(j,2)==4
               S(K)=c(j,1);
               K=K+1;
            end
        end
        M(3,k)=min(H(S(2:nn))-H(S(1:nn-1)));
    else
        M(3,k)=5;
    end

    M(4,k)=12-d;

    dh=H(c(2:N,1))-H(c(1:N-1,1));
    dr=abs(R(c(2:N,1))-R(c(1:N-1,1)));
    M(5,k)=prod(dh/max(dh))+prod(dr/max(dr));

    if n(4)==0
        M(6,k)=1;
    end

    dH1=H(50)-H(49);
    dR1=R(50)-R(49);
    dH2=H(1)-H(2);
    dR2=R(1)-R(2);
```

```
    topangle=dR1/sqrt(dH1^2+dR1^2);
    bottomangle=dR2/sqrt(dH2^2+dR2^2);
    M(7,k)=acos(bottomangle)-acos(topangle);

    tw=sum((c(:,1)>25).*(c(:,1)-25).*(R(c(:,1))'));
    bw=sum((c(:,1)<=25).*(26-c(:,1)).*(R(c(:,1))'));
    M(8,k)=tw-bw;

    f=CoM(R,H);
    M(9,k)=f-25;
end
M(1:6,D+1:2*D)=M(1:6,1:D);
M(7:9,D+1:2*D)=-M(7:9,1:D);
```

All others remain the same as before with minor modifications to call the new functions in place of the old.

FullVase.m

```
clear all
P=1;
A=xlsread('Final.xlsx',1);
a=A(1:5,P);
as=A(6:8,P);
phi=linspace(0,2*pi,50);
x=sin(phi);
y=cos(phi);
z=ones(1,50);
N=25;

Scale=30;

[R1,H1]=Bezier2(a,as');
r=min(R1);
R1=R1-r+0.75;

dr=R1(1)-R1(2);
dh=H1(1)-H1(2);

dx=cross([dr/(sqrt(dr^2+dh^2)),dh/(sqrt(dr^2+dh^2)),0],[0,0,1]);
dphi=acos(dot([dr/(sqrt(dr^2+dh^2)),dh/(sqrt(dr^2+dh^2))],[0,1]));

if dr>0
    a=linspace(-dphi+pi,-pi/2,8);
else
    a=linspace(dphi-pi,-pi/2,8);
end
R=(cos(a)+dx(1))*0.8/Scale+R1(1);
H=(sin(a)+dx(2))*0.8/Scale+H1(1);
R(9)=R(8)-R(8)*3/Scale;
H(9)=H(8);
a=linspace(-pi/2,-pi,4);
R(9:12)=(cos(a))/Scale+R(9);
H(9:12)=(sin(a)+1)/Scale+H(9);
a=linspace(0,pi/2,4);
R(12:15)=(cos(a)-1)/Scale+R(12);
H(12:15)=(sin(a))/Scale+H(12);
R(16)=0;
H(16)=H(15);
```

```
R=[fliplr(R),R1];
H=[fliplr(H),H1];

dr=R1(1:49)-R1(2:50);
dh=H1(1:49)-H1(2:50);

X2=cross([dr./(sqrt(dr.^2+dh.^2));dh./(sqrt(dr.^2+dh.^2));zeros(1,49)],[zeros
    (1,49);zeros(1,49);ones(1,49)]);

R2=(R1(2:50)+R1(1:49))/2+X2(1,:)*3/Scale;
H2=(H1(2:50)+H1(1:49))/2+X2(2,:)*3/Scale;

S=0;
while S<3/Scale
    R2=R2(2:end);
    H2=H2(2:end);
    dr=R2(1)-R2(2);
    dh=H2(1)-H2(2);
    dx=cross([dr/(sqrt(dr^2+dh^2)),dh/(sqrt(dr^2+dh^2)),0],[0,0,1]);
    dphi=acos(dot([dr/(sqrt(dr^2+dh^2)),dh/(sqrt(dr^2+dh^2))],[0,1]));
    if dr>0
        a=linspace(-dphi+pi,-pi/2,8);
    else
        a=linspace(dphi-pi,-pi/2,8);
    end
    R3=(cos(a)+dx(1))*4/Scale+R2(1);
    H3=(sin(a)+dx(2))*4/Scale+H2(1);
    S=H3(8)-H(1);

end

R3(9)=0;
H3(9)=H3(8);

R2=[fliplr(R3),R2];
H2=[fliplr(H3),H2];

dr=R(66)-R(65);
dh=H(66)-H(65);

dx=cross([dr/(sqrt(dr^2+dh^2)),dh/(sqrt(dr^2+dh^2)),0],[0,0,1]);
dphi=acos(dot([dr/(sqrt(dr^2+dh^2)),dh/(sqrt(dr^2+dh^2))],[0,1]));

if dr>0
    a=linspace(-dphi+pi,-dphi,8);
else
    a=linspace(dphi+pi,dphi,8);
end
R1=(cos(a)-dx(1))/Scale+R(66);
H1=(sin(a)-dx(2))/Scale+H(66);
R1(1:4)=R1(1:4)-dx(1)/Scale;
H1(1:4)=H1(1:4)-dx(2)/Scale;

R=[R,fliplr(R1),fliplr(R2)];
H=[H,fliplr(H1),fliplr(H2)];
R=R*2;
H=H*2;

figure(1)
plot(R,H)
axis equal

Z=z'*H;
X=x'*R;
```

```
Y=y'*R;

C=ones(50,size(R,2),3)*0.9;
C(:,:,1)=C(:,:,1)*0.85;


surf(X,Y,Z,C,'FaceLighting','phong')
axis equal
shading('interp')
light('Position',[-5,0,-5]);
surf2stl('Vase2.stl',X,Y,Z)
```