

Chapter 10

Computational Aesthetic Evaluation: Past and Future

Philip Galanter

Abstract Human creativity typically includes a self-critical aspect that guides innovation towards a productive end. This chapter offers a brief history of, and outlook for, computational aesthetic evaluation by digital systems as a contribution towards potential machine creativity. First, computational aesthetic evaluation is defined and the difficult nature of the problem is outlined. Next, a brief history of computational aesthetic evaluation is offered, including the use of formulaic and geometric theories; design principles; evolutionary systems including extensions such as coevolution, niche construction, agent swarm behaviour and curiosity; artificial neural networks and connectionist models; and complexity models. Following this historical review, a number of possible contributions towards future computational aesthetic evaluation methods are noted. Included are insights from evolutionary psychology; models of human aesthetics from psychologists such as Arnheim, Berlyne, and Martindale; a quick look at empirical studies of human aesthetics; the nascent field of neuroaesthetics; new connectionist computing models such as hierarchical temporal memory; and computer architectures for evolvable hardware. Finally, it is suggested that the effective complexity paradigm is more useful than information or algorithmic complexity when thinking about aesthetics.

a63c5a3857db9a4d63ca53b55e68664f ebrary 10.1 Introduction

This chapter looks at computers and aesthetic evaluation. In common usage the word *creativity* is associated with bringing the new and innovative into being. The term, whether used in reference to the arts or more generally, connotes a sort of self-directedness and internal drive. Evaluation or criticism is by its very nature reactive. Something is first created and only then can it be evaluated. Evaluation and creativity at first seem to be two different kinds of activity performed at different times.

But almost any exploration of creativity will quickly reveal evaluation threaded throughout the entire process. For accomplished artists there are usually at least

P. Galanter (✉)

Department of Visualization, Texas A&M University, College Station, Texas, USA
e-mail: galanter@viz.tamu.edu

J. McCormack, M. d'Inverno (eds.), *Computers and Creativity*,
DOI 10.1007/978-3-642-31727-9_10, © Springer-Verlag Berlin Heidelberg 2012

255

three ways evaluation becomes an intrinsic part of the creative process. First, artists typically exercise evaluation as they experience, study, and find inspiration in the work of other artists. In practice artists will execute countless micro-evaluations as part of making aesthetic decisions for works-in-progress. Once completed, artists evaluate the final product, gaining new insights for the making of the next piece.

If computers are to become artistically creative their need for an evaluative function will be no less acute. Computer artists have invented a great variety of fecund computational methods for generating aesthetic possibilities and variations. But computational methods for making aesthetically sound choices among them have lagged far behind.

This chapter provides specific examples of computational methods for making aesthetic choices. Longer examples have been selected as good illustrations of a particular approach, with shorter examples providing variations. Some examples show where a path is already known to lead, while others are provided as trail heads worthy of further exploration.

10.1.1 What Do We Mean by Computational Aesthetic Evaluation?

The word *evaluation* is sometimes prone to ambiguous use due to the multiple meanings of the word *value*. For example, a mathematician can be said to evaluate an expression or formula. An art expert might evaluate a given object for market value or authenticity. Part of that might involve an evaluation of style and provenance.

For this discussion *aesthetic evaluation* refers to making normative judgements related to questions of beauty and taste in the arts. It's worth noting that the word "aesthetics" alone can imply a broader critical contemplation regarding art, nature, and culture. The topic of aesthetics, including evaluation, goes back at least to Plato and Aristotle in the West (for a good overview of philosophical aesthetics see Carroll

1999). The term *computational aesthetics* has been somewhat instable over time. For some the term includes both generative and analytic modes, i.e. both the creation and evaluation of art using a computer. For others it purely refers to the use of computers in making aesthetic judgements. This chapter concentrates on systems for making normative judgements, and to emphasise this I've used the terms "computational aesthetic evaluation", "machine evaluation", and "computational evaluation" as synonyms (Hoenig 2005, Greenfield 2005b).

Computational aesthetic evaluation includes two related but distinct application modes. In one mode aesthetic evaluations are expected to simulate, predict, or cater to human notions of beauty and taste. In the other mode machine evaluation is an aspect of a meta-aesthetic exploration and usually involves aesthetic standards created by software agents in artificial worlds. Such aesthetics typically feel alien and disconnected from human experience, but can provide insight into all possible aesthetics including our own.

Finally, it's worth noting that *aesthetic evaluation* and the *evaluation of creativity* are somewhat related but quite distinct. For example, accomplishments in non-artistic fields such as science and mathematics can also be evaluated as to their degree of creativity. And in the arts it's possible to have an artwork of high aesthetic value but without much creativity, or a highly creative artwork where the aesthetics are poor or even irrelevant.

10.1.2 Why Is Computational Aesthetic Evaluation so Difficult if not Impossible?

It should be noted at the outset that computational aesthetic evaluation is an extremely difficult problem. In the abstract, notions of computational aesthetic evaluation and computational creativity lead to deep philosophical waters regarding phenomenology and consciousness. Let's assume a computational evaluation system is created that appears to duplicate human aesthetic judgement. Would such a machine actually experience a sense of redness, brightness or other qualia? How would we know? Can machine evaluation be successful without such experience? If such a machine isn't conscious does that mean human aesthetic judgement and computational aesthetic evaluation are quite different? Or could it be that they aren't so different after all because the brain is itself a machine? All of these interesting questions are outside of the scope of this chapter.

Some feel that effective practical computational evaluation will remain out of reach in our lifetime and perhaps forever. The complications begin with the likely fact that the human aesthetic response is formed by a combination of genetic predisposition, cultural assimilation, and unique individual experience. Despite a growing research literature, the psychology of aesthetics is a mostly incomplete science, and our understanding of each component is limited.

Even if we had a full understanding of aesthetics' genetic, cultural, developmental, and psychological modalities, the creation of comparable computational functionality would remain a daunting task. It would probably require the resolution of a number of standing hard problems in artificial intelligence. A model of human aesthetics, or human intelligence in general, has to represent more than a hypothetical brain-in-a-jar. Our aesthetic sense and psychological makeup are in part the result of embodied experience situated in a specific environment. Machine evaluation will have to account for perception not as a passive mental process, but rather as a dynamic interaction between our bodies and the world (Davis and Rebelo 2007, McCormack 2008). Additionally, it will have to allow for emotions and the irrational Dionysian element in the arts.

10.2 A Brief History of Computational Aesthetic Evaluation

Any suggested computational aesthetic evaluation mechanism is going to contain, at least implicitly, a theory of aesthetics. Most theories from the history of aesthetics do

not immediately suggest algorithms, quantifiable properties, or objective formulas. But some do and it is with those that our discussion begins.

10.2.1 *Formulaic and Geometric Theories*

The mathematician George David Birkhoff published a mostly speculative book in 1933 titled "Aesthetic Measure". Birkhoff limits his theory to aspects of pure form (the "formal") and doesn't address symbolic meaning (the "connotative"). He then proposes the formula $M = O/C$ where M is the measure of aesthetic effectiveness, O is the degree of order, and C is the degree of complexity. Birkhoff (1933) notes, "The well known aesthetic demand for 'unity in variety' is evidently closely connected with this formula."

Birkhoff warns that his measure can only be applied within a group of similar objects and not across types such as a mix of oil and watercolour paintings. He also finesses variation in experience and taste intending M to be a measure for an "idealised 'normal observer'" as a sort of mean of the population.

While most of the book is presented from a mathematical point of view, it is sometimes forgotten that Birkhoff begins with an explicit psychoneurological hypothesis. He describes complexity (C) as the degree to which unconscious psychological and physiological effort must be made in perceiving the object. Order (O) is the degree of unconscious tension released as the perception is realised. This release mostly comes from the consonance of perceived features such as "repetition, similarity, contrast, equality, symmetry, balance, and sequence." While Birkhoff views complexity and order as ultimately psychological phenomena, for analysis he operationalises those concepts using mathematical representations. He then goes on to analyse examples such as polygons, vases, and harmonic structures in music to illustrate his theory.

Birkhoff's theory has been disputed from its first publication. For example, in 1939 Wilson published experimental results showing that Birkhoff's measure did not correlate with actual subjects' stated aesthetic preferences regarding polygons (Wilson 1939). Alternate formulas have been offered that seem to correlate more closely with the judgements of subjects (Boselie and Leeuwenberg 1985, Staudek 1999). And for some, Birkhoff's formula seems to measure orderliness rather than beauty, and penalises complexity in a rather unqualified way (Scha and Bod 1993).

But there are at least two aspects of Birkhoff's work that remain in legitimate play today. First is the intuition that aesthetic value has something to do with complexity and order relationships. Second is the idea that modelling brain function can illuminate discussions of aesthetics. Indeed, both of these reappear as themes throughout this chapter.

The positing of mathematical bases for aesthetics long predate Birkhoff. Pythagoras is traditionally credited with the discovery that dividing a vibrating string following simple consecutive integer ratios such as 1:2, 2:3, and 3:4 yields pleasing harmony relationships. The Golden Ratio ϕ , an irrational constant approximately equal to 1.618, and the related Fibonacci series have been said to generate

proportions of optimal aesthetic value. It is claimed they are embedded in great works of art, architecture, and music.

Psychologist Gustav Fechner is credited with conducting the first empirical studies of human aesthetic response in the 1860s. His experiments seemed to show that golden rectangles had the greatest appeal relative to other aspect ratios. But subsequent studies have cast strong doubt on those results. As noted in a special issue of the journal *Empirical Studies of the Arts*, there were methodological flaws and cultural bias in previous confirmatory studies (McCormack 2008, Holger 1997).

In addition, Livio has credibly debunked supposed Golden Ratio use in works including the Great Pyramids, the Parthenon, the Mona Lisa, compositions by Mozart, and Mondrian's late paintings. However, he notes that use of the Golden Ratio as an aesthetic guide has become something of a self-fulfilling myth. For example, Le Corbusier's Modulor, a design aid for proportions, was consciously based on the Golden Ratio (Livio 2003).

On a bit firmer ground is a principle credited to linguist George Kingsley Zipf commonly referred to as *Zipf's law*. As first applied to natural language, one can begin with a large body of text and tally every word counting each occurrence. Then list each word from the most to the least frequent. The observed result is that for the frequency P_i of a given word with a given rank i :

$$P_i \approx \frac{1}{i^a} \tag{10.1}$$

where the exponent a is near 1 (Zipf 1949).

Manaris et al. (2005; 2003) note that this power law relationship has not only been verified in various bodies of musical composition, but also "colours in images, city sizes, incomes, music, earthquake magnitudes, thickness of sediment depositions, extinctions of species, traffic jams, and visits of websites, among others." They go on to show how Zipf metrics can be used to classify specific works as to composer, style, and an aesthetic sense of "pleasantness". In addition Machado et al. (2007) apply Zipf's law in the creation of artificial art critics. Much earlier work showed that both frequency and loudness in music and speech conform to a $1/f$ statistical power law. The authors suggest using $1/f$ distributions in generative music (Voss and Clarke 1975).

Studies by Taylor have shown that late period "drip" paintings by Jackson Pollock are fractal-like. He has also suggested that the fractal dimension of a given Pollock painting is correlated with its aesthetic quality. Fractals are mathematical objects that exhibit self-similarity at all scales. Examples of real world objects that are fractal-like in form include clouds, mountains, trees, and rivers. In the case of Pollock's paintings the fractal dimension is a measure of the degree to which the canvas is filled with finely detailed complex structures. A paint mark with a fractal dimension of 1 will no more fill the canvas with detailed structures than a typical straight line. A paint mark with a fractal dimension of 2 will entirely fill the canvas with fine detail. These correspond well with our everyday topological sense of one and two dimensional spaces (Peitgen et al. 1992).

Pollock's paint marks exhibit detail between these two extremes, and have a non-integer dimension somewhere between 1 and 2. When measured empirically the

fractal dimension of his paintings increases over time from 1.12 in 1945 to 1.72 in 1952. Presumably Pollock's innovative "dripping" technique improved over time and in this very limited realm the fractal dimension can be used for aesthetic evaluation (Taylor 2006). Use of a related measure applied to non-fractal two-dimensional patterns correlates well with beauty and complexity as reported by human subjects (Mori et al. 1996).

Work has been done in the fields of medical reconstructive and cosmetic surgery to quantify facial and bodily beauty as an objective basis for evaluating the results of medical procedures. Hönn and Göz (2007) in the field of orofacial orthopaedics cite studies indicating that infants preferentially select for facial attractiveness, and that such judgements by adults are consistent across cultures. Atiyeh and Hayek (2008) provide a survey for general plastic surgery, indicating a likely genetic basis for the perception of both facial and bodily attractiveness. Touching on rules of proportion used by artists through the centuries they seem ambivalent or even supportive of the Golden Ratio standard. However, in conclusion they write, "The golden section phenomenon may be unreliable and probably is artifactual".

To date when it comes to quantifying human facial and bodily beauty there is no medical consensus or standardised measure. More broadly, many now feel that any simple formulaic approach to aesthetic evaluation will be inadequate. Beauty seems to be too multidimensional and too complex to pin down that easily.

10.2.2 *Design Principles*

Another source of aesthetic insight is the set of basic principles taught in typical design foundations courses. A standard text in American classrooms includes considerations such as: value and distribution; contrast; colour theory and harmony; colour interaction; weight and balance; distribution and proportion; and symmetrical balance. Also included are Gestalt-derived concepts like grouping, containment, repetition, proximity, continuity, and closure (Stewart 2008).

However, to date there is very little in the way of software that can extract these features and then apply rule-of-thumb evaluations. Among the few is a system that makes aesthetic judgements about arbitrary photographs. Datta et al. (2006; 2007) began with a set of photos from a photography oriented social networking site. Each photo was rated by the membership. Image processing extracted 56 simple measures related to exposure, colour distribution and saturation, adherence to the "rule of thirds," size and aspect ratio, depth of field, and so on. The ratings and extracted features were then processed using both regression analysis and classifier software. This resulted in a computational model using 15 key features. A software system was then able to classify photo quality in a way that correlated well with the human ratings.

Some work has been done using colour theory as a basis for machine evaluation. Tsai et al. (2007) created a colour design system using genetic searching and noted, "... auto-searching schemes for optimal colour combinations must be supervised

by appropriate colour harmony theories since if such supervision is not applied, the search results are liable to be dull and uncoordinated. . .” Others have applied a variation of Birkhoff’s aesthetic measure for colour harmony attempting to better define order in colour schemes (Li and Zhang 2004).

But overall there has been little progress in automating design principles for aesthetic evaluation. Feature extraction figures heavily in this problem, so perhaps future computer vision researchers will take on this problem.

10.2.3 Artificial Neural Networks and Connectionist Models

Artificial neural networks are software systems with designs inspired by the way neurones in the brain are thought to work. In the brain neurone structures called *axons* act as outputs and *dendrites* act as inputs. An axon to dendrite junction is called a *synapse*. In the brain, electrical impulses travel from neurone to neurone where the synaptic connections are strong. Synapse connections are strengthened when activation patterns reoccur over time. Learning occurs when experience leads to the coherent formation of synapse connections.

In artificial neural networks virtual neurones are called *nodes*. Nodes have multiple inputs and outputs that connect to other nearby nodes similar to the way synapses connect *axons* and *dendrites* in the brain. Like synapses these connections are of variable strength, and this is often represented by a floating point number. Nodes are typically organised in layers, with an input layer, one or more hidden layers, and finally an output layer. Connection strengths are not manually assigned, but rather “learned” by the artificial neural network as the result of its exposure to input data.

For example, a scanner that can identify printed numbers might be created by first feeding pixel images to the input layer of an artificial neural network. The data then flows through the hidden layer connections according to the strength of each connection. Finally, one of ten output nodes is activated corresponding to one of the digits from “0” to “9”. Before being put into production the scanner would be trained using known images of digits.

Some of the earliest applications of neural network technology in the arts consisted of freestanding systems used to compose music (Todd 1989). Later in this chapter artificial neural networks will be described as providing a component in evolutionary visual art systems (Baluja et al. 1994).

A significant challenge in using artificial neural networks is the selection, conditioning, and normalisation of data presented to the first layer of nodes. It was noted in Sect. 10.2.1 that ranked music information following Zipf’s law can be used to identify composers and evaluate aesthetics. Manaris et al. (2005; 2003) reported an impressive success rate of 98.41 % in attempting to compute aesthetic ratings within one standard deviation of the mean from human judges.

A similar effort was made to evaluate a mix of famous paintings and images from a system of evolved expressions. The machine evaluation used Zipfian rank-frequency measures as well as compression measures as proxies for image complexity. The authors reported a success rate of 89 % when discriminating between human

and system-produced images. Using famous paintings in the training set provided stability and human-like standards of evaluation. Using system produced images allowed the evolution of more discerning classifiers (Machado et al. 2008). In a related paper the authors demonstrate artificial neural networks that can discriminate works between: Chopin and Debussy; Scarlatti and Purcell; Purcell, Chopin, and Debussy; and other more complicated combinations. In another demonstration, a neural network was able to discriminate works between Gauguin, Van Gogh, Monet, Picasso, Kandinsky, and Goya (Machado et al. 2004, Romero et al. 2003).

Without explicit programming, artificial neural networks can learn and apply domain knowledge that may be fuzzy, ill defined, or simply not understood. Phon-Amnuaisuk (2007) has used a type of artificial neural network called *self-organising maps* to extract musical structure from existing human music, and then shape music created by an evolutionary system by acting as a critic. Self-organising map-based music systems sometimes produce reasonable sequences of notes within a measure or two, but lack the kind of global structure we expect music to have. In an attempt to address this problem self-organising maps have been organised in hierarchies so that higher-level maps can learn higher levels of abstraction (Law and Phon-Amnuaisuk 2008). In another experiment, artificial neural networks were able to learn viewer preferences among Mondrian-like images and accurately predict preferences when viewing new images (Gedeon 2008).

10.2.4 Evolutionary Systems

The evolutionary approach to exploring solution spaces for optimal results has had great success in a diverse set of industries and disciplines (Fogel 1999). Across a broad range of approaches some kind of evaluation is typically needed to steer evolution towards a goal. Much of our discussion about computational aesthetic evaluation will be in the context of evolutionary systems. But first consider the following simplified industrial application.

Assume the problem at hand is the design of an electronic circuit. First, chromosome-inspired data structures are created and initially filled with random values. Each chromosome is a collection of simulated genes. Here each gene describes an electronic component or a connection, and each chromosome represents a circuit that is a potential solution to the design problem. The genetic information is referred to as the *genotype*, and the objects and behaviours they ultimately produce are collectively called the *phenotype*. The process of genotype-creating-phenotype is called *gene expression*. A chromosome can reproduce with one or more of its genes randomly mutated. This creates a variation of the parent circuit. Or two chromosomes can recombine creating a new circuit that includes aspects of both parents.

In practice, a subset of chromosomes is selected for variation and reproduction, and the system evaluates the children as possible solutions. In the case of circuit design a chromosome will be expressed as a virtual circuit and then tested with a software-based simulator. Each circuit design chromosome is assigned a score based

on not only how well its input and output match the target specification, but perhaps other factors such as the cost and number of parts, energy efficiency, and ease of construction.

The formula that weights and combines these factors into a single score is called a *fitness function*. Chromosomes with higher fitness scores are allowed to further reproduce. Chromosomes with lower fitness scores are not selected for reproduction and are removed from evolutionary competition. Using a computer this cycle of selection, reproduction, variation, and fitness evaluation can be repeated hundreds of times with large populations of potential circuits. Most initial circuits will be quite dysfunctional, but fortuitous random variations will be retained in the population, and eventually a highly optimised “fit” circuit will evolve. For an excellent introduction to evolutionary systems in computer art see Bentley and Corne (2002). In that same volume, Koza et al. (2002) illustrate the application of genetic programming in real world evolutionary circuit design.

Evolutionary systems have been used to create art for more than 20 years (Todd and Latham 1992). But an evolutionary approach to art is particularly challenging because it is not at all clear how aesthetic judgement can be automated for use as a fitness function. Nevertheless, evolution remains a popular generative art technique despite this fundamental problem (for an overview of current issues in evolutionary art see McCormack 2005 and Galanter 2010).

From the outset there have been two popular responses to the fitness function problem. The first has been to put the artist in the loop and assign fitness scores manually. The second has been to use computational aesthetic evaluation and generate fitness scores computationally. More recently there have been efforts to create systems with fitness functions that are emergent rather than externally determined.

10.2.5 Interactive Evolutionary Computation

From the earliest efforts interactive (i.e. manual) assignment of fitness scores has dominated evolutionary art practice (Todd and Latham 1992, Sims 1991). There was also early recognition that the human operator creates a “fitness bottleneck” (Todd and Werner 1998). This labour-intensive bottleneck forces the use of fewer generations and smaller populations than in other applications (for a comprehensive overview of interactive evolutionary computing across a number of industries, including media production, see Takagi 2001).

There are additional problems associated with the interactive approach. For example, human judges become fatigued, less consistent, and prone to skew towards short term novelty at the expense of aesthetic quality (Takagi 2001, Yuan 2008). One suggested remedy for such fatigue problems has been to *crowd-source* evaluation. This involves recruiting large numbers of people for short periods of time to render judgements. In Sim’s *Galapagos*, choices viewers make as to which of a number of monitors to watch are used as implicit fitness measures (Sims 1997). The *Electric Sheep* project provides evolutionary fractal flame art as a screen saver on thousands

of systems around the world. Users are invited to provide online feedback regarding their preferences (Draves 2005).

But the crowd-sourcing solution is not without its own potential problems. Artists Komar and Melamid executed a project called *The People's Choice* that began by polling the public about their preferences in paintings. Based on the results regarding subject matter, colour, and so on they created a painting titled *America's Most Wanted*. The result is a bland landscape that would be entirely unmemorable if it were not for the underlying method and perhaps the figure of George Washington and a hippopotamus appearing as dada-like out-of-context features. As should be expected the mean of public opinion doesn't seem to generate the unique vision most expect of contemporary artists. Komar and Melamid's critique in this project was directed at the politics of public relations and institutions that wield statistics as a weapon. But the aesthetic results advise caution to those who would harness crowd-sourced aesthetic evaluation in their art practice (Komar et al. 1997, Ross 1995). It's also worth noting that Melamid observed that some aesthetic preferences are culturally based but others seemed to be universal. The evolutionary implications of this will be discussed later in the section on Denis Dutton and his notion of the "art instinct", Sect. 10.3.1.

Another approach has been to manually score a subset, and then leverage that information across the entire population. Typically this involves clustering the population into similarity groups, and then only manually scoring a few representatives from each (Yuan 2008, Machado et al. 2005). Machwe (2007) has suggested that artificial neural networks can generalise with significantly fewer scored works than the interactive approach requires.

10.2.6 Automated Fitness Functions Based on Performance Goals

The Mechanical Turk was a purported mechanical chess-playing machine created in the late 18th century by Wolfgang von Kempelen. But it was really more a feat of stage magic than computation. Exhibitors would make a great show of opening various doors revealing clockwork-like mechanisms. Despite appearances, a human operator was hidden inside the cabinet, so the chess game was won or lost based on the decisions the operator made (Aldiss 2002, Standage 2002).

To some extent using interactive evolutionary computing for art is a similar trick. These systems can generate and display a variety of options at every step, but ultimately the aesthetic challenge is won or lost based on the decisions made by the artist-operator.

Fully automated evolutionary art systems call for, rather than offer, a solution to the challenge of computational aesthetic evaluation. Machine evaluation can be relatively simple when the aesthetic is Louis H. Sullivan's principle that "form follows function" (Sullivan 1896). Computational evaluation here is tractable to the extent the needed functionality can be objectively evaluated via computation. For example, Gregory Hornby and Jordan Pollack created an evolutionary system for

designing furniture (tables). Their fitness function sought to maximise height, surface structure, and stability while minimising the amount of materials required. This approach is similar to the optimisation-oriented evolutionary systems found in industry (Hornby and Pollack 2001).

Similarly, specific performance goals can provide a fitness function in a straightforward way in art applications. Sims' *Evolved Virtual Creatures* is an early example. His evolutionary system bred virtual organisms with simple "neuron" circuitry and actuators situated in a world of simulated physics. The initial creatures, seeded with random genes, would typically just twitch in an uncoordinated way. But then selection pressure was applied to the evolving population using a simple fitness function that might reward jumping height, walking speed, or swimming mobility. As a result, the evolved creatures exhibited very competent locomotion behaviour. Some seemed to rediscover movement found in the natural world, while others exhibited strange and completely novel solutions (Sims 1994).

Performance goals can also be useful in the development of characters for computer games through evolution. For example, the amount of time a character survives can be used as a fitness function yielding incrementally stronger play (Wu and Chien 2005).

Diffusion limited aggregation (DLA) systems can be used to create growing frost- or fern-like patterns, and have been studied using evolutionary performance goals. They grow as particles in random Brownian motion adhere to an initial seed particle. To study optimal seed placement, Greenfield (2008a) applied an evolutionary system where the size of the resulting pattern served as an effective fitness measure. In another project he used an evolutionary system to explore the effect of transcription factors on morphology. Each transcription factor was assigned a different colour. The performance and aesthetics of the result were improved by using a fitness function that rewarded transcription factor diversity (Greenfield 2004). Similarly, an evolutionary sculpture system using cubic blocks as modules has produced useful emergent forms simply by rewarding height or length (Tufte and Gangvik 2008).

In their project "Breed" Driessens and Verstappen created a subtractive sculpture system. Each sculpture is started as a single cube treated as a cell. This cell is subdivided into eight smaller sub-cells, one for each corner. Rules driven by the state of neighbouring cells determine whether a sub-cell is kept or carved away. Then each of the remaining cells has the subdivision rules applied to them. And so on. The final form is then evaluated for conformance to goals for properties such as volume, surface area and connectivity. In "Breed" the rule-set is the genotype, the final sculpture is the phenotype, and evaluation relative to performance goals is used as a fitness function. Unlike most other evolutionary systems there is a population size of just one. A single mutation is produced and given an opportunity to unseat the previous result. At some point the gene, i.e. rule set, ceases to improve by mutation and the corresponding sculpture is kept as the result.

Whitelaw (2003) points out that unlike industrial applications where getting stuck on a local maximum is seen as an impediment to global optimisation, this project uses local maxima to generate a family of forms (differing solutions) related

by their shared fitness function. Also Whitelaw points out that unlike some generative systems that reflect human selection and intent, Driessens and Verstappen have no particular result in mind other than allowing the system to play itself out to a final self-directed result. In this case performance goals play quite a different role than those used in optimisation-oriented industrial systems.

10.2.7 Evolutionary Fitness Measured as Error Relative to Exemplars

Representationalism in visual art began diminishing in status with the advent of photographic technologies. Other than use as an ironic or conceptual gesture, mimesis is no longer a highly valued pursuit in contemporary visual art. Similarly a difference or error measure comparing a phenotype to a real-world example is not typically useful as an aesthetic fitness function. In the best case such a system would merely produce copies. What have proven interesting, however, are the less mimetic intermediate generations where error measures can be reinterpreted as the degree of abstraction in the image.

For example, Aguilar and Lipson (2008) constructed a physical painting machine driven by an evolutionary system. A scanned photograph serves as the target and each chromosome in the population is a set of paint stroke instructions. A model of pigment reflectance is used to create digital simulations of the prospective painting in software. A software comparison of pixel values from the simulated painting and the original image generates a fitness score. When a sufficient fitness score is achieved the chromosome is used to drive a physical painting machine that renders the brush strokes on canvas with acrylic paint.

Error measurement makes particularly good sense when programming music synthesisers to mimic other sound sources. Comparisons with recordings of traditional acoustic instruments can be used as a fitness function. And before the evolutionary system converges on an optimal mimesis interesting timbres may be discovered along the way (McDermott et al. 2005, Mitchell and Pipe 2005).

Musique concrete is music constructed by manipulating sound samples. For evolutionary musique concrete short audio files can be subjected to operations similar to mutation and crossover. They are then combined and scored relative to a second target recording. Again mimesis is not the intent. What the audience hears is the evolving sound as it approaches but does not reach the target recording (Magnus 2006, Fornari 2007). Gartland-Jones (2002) has used a similar target tracking approach with the addition of music theory constraints for evolutionary music composition.

In a different music application Hazan et al. (2006) have used evolutionary methods to develop regression trees for expressive musical performance. Focusing on note duration only, and using recordings of jazz standards as a training set, the resulting regression trees can be used to transform arbitrary flat performances into expressive ones.

There are numerous other examples of error measures used as fitness functions. For example, animated tile mosaics have been created that approach a reference portrait over time (Ciesielski 2007). The fitness of shape recognition modules have been based on their ability to reproduce shapes in hand drawn samples (Jaskowski 2007). An automated music improviser has been demonstrated that proceeds by error minimisation of both frequency and timbre information (Yee-King 2007). Alsing (2008) helped to popularise the error minimisation approach to mimetic rendering with a project that evolved a version of the “Mona Lisa” using overlapping semi-transparent polygons.

10.2.8 Automated Fitness Functions Based on Complexity Measures

a63c5a3857db9a4d63ca53b55e68664f
ebrary

Fitness scores based on aesthetic quality rather than simple performance or mimetic goals are much harder to come by. Machado and Cardoso’s *NEvAr* system uses computational aesthetic evaluation methods that attempt to meet this challenge. They generate images using an approach first introduced by (Sims 1991) called *evolving expressions*. It uses three mathematical expressions to calculate pixel values for the red, blue, and green image channels. The set of math expressions operates as a genotype that can reproduce with mutation and crossover operations.

Machado and Cardoso take a position related to Birkhoff’s aesthetic measure. The degree to which an image resists JPEG compression is considered an “image complexity” measure. The degree it resists fractal compression is considered to be proportional to the “processing complexity” that will tax an observer’s perceptual resources. Image complexity is then essentially divided by processing complexity to calculate a single fitness value.

Machado and Cardoso reported surprisingly good imaging results using evolving expressions with their complexity-based fitness function. But the authors were also careful to note that their fitness function only considers one formulaic aspect of aesthetic value. They posit that cultural factors ignored by *NEvAr* are critical to aesthetics. In later versions of *NEvAr* a user guided interactive mode was added (Machado and Cardoso 2002; 2003, Machado et al. 2005, see also Chap. 11 in this volume for their extended work in this vein).

a63c5a3857db9a4d63ca53b55e68664f
ebrary

10.2.9 Automated Fitness Functions in Evolutionary Music Systems

For evolutionary music composition some have calculated fitness scores using only evaluative rules regarding intervals, tonal centres, and compliance to key and meter. Others, like *GenOrchestra*, are hybrid systems that also include some form of listener evaluation. The *GenOrchestra* authors note that unfortunately without human

a63c5a3857db9a4d63ca53b55e68664f
ebrary

evaluation “the produced tunes do not yet correspond to a really human-like musical composition” (Khalifa and Foster 2006, De Felice and Fabio Abbattista 2002).

Others have used music theory-based fitness functions for evolutionary bass harmonisation (De Prisco and Zaccagnino 2009), or to evolve generative grammar expressions for music composition (Reddin et al. 2009). For mimetic evolutionary music synthesiser programming McDermott et al. (2005) used a combination of perceptual measures, spectral analysis, and sample-level comparison as a fitness function to match a known timbre.

Weinberg et al. (2009) have created a genetically based robotic percussionist named *Haile* that can “listen” and trade parts in the call and response tradition. Rather than starting with a randomised population of musical gestures *Haile* begins with a pool of pre-composed phrases. This allows *Haile* to immediately produce musically useful responses. As *Haile* runs, however, the evolutionary system will create variations in real time. The fitness function used for selection uses an algorithm called *dynamic time warping*.

Dynamic time warping here provides a way to measure the similarity between two sequences that may differ in length or tempo. In response to a short rhythmic phrase played by a human performer, *Haile* applies the dynamic time warping-based fitness function to its population of responses and then plays back the closest match. The goal is not to duplicate what the human player has performed, but simply to craft a response that is aesthetically related and thus will contribute to a well-integrated performance.

10.2.10 Multi-objective Aesthetic Fitness Functions in Evolutionary Systems

Aesthetic judgements are typically multidimensional. For example, evaluating a traditional painting involves formal issues regarding colour, line, volume, balance, and so on. A fitness function that has to include multiple objectives like these will typically have a sub-score for each. Each sub-score will be multiplied by its own coefficient that serves as a weight indicating its relative importance. The weighted sub-scores are then summed for a final total score.

However, the weights are typically set in an ad hoc manner, and resulting evaluations may not push the best work to the front. And there is no reason to assume that the weights should maintain a static linear relationship regardless of the sub-score values. For example, various aspects of composition may influence the relative importance of colour.

Pareto ranking can address some of these concerns as an alternative to simple weights. In Pareto ranking one set of scores is said to *dominate* another if it is at least as good in all component sub-scores, and better in at least one. A *rank 1* set of scores is one that isn't dominated. When there are multiple objectives there will typically be multiple rank 1 sets of scores. The dimension of the problem they dominate is what differentiates rank 1 genotypes, and all can be considered viable.

Genotypes of less than rank 1 can be considered redundant. Note, however, that some redundancy in the gene pool is usually considered a good thing. In situations where a single genotype must be selected, a rank 1 genotype is sometimes selected based on its uniqueness relative to the current population (Neufeld et al. 2008, Ross and Zhu 2004, Greenfeld 2003).

Both weighting and Pareto ranking are approaches to the more general problem of multi-objective optimisation. For multidimensional aesthetics a computational evaluation system will have to deal with multi-objective optimisation either explicitly as above, or implicitly as is done in the extensions to evolutionary computation noted below.

10.2.11 *Biologically Inspired Extensions to Simple Evolutionary Computation*

Evolutionary art faces significant challenges beyond machine evaluation-based fitness functions. For example, the expression of genes in nature doesn't happen in a single step. There is a cascading sequence of emergence across a number of scales from DNA, to proteins, organelles, cells, tissues, and finally organs resulting in an individual. Life's capacity for complexification is unmatched in the known universe. By comparison evolutionary computing systems are simple in that they typically only support a single level of emergence, i.e. the genotype directly generates the phenotype (Casti 1994, Galanter 2010).

And so current evolutionary computing technologies have a very limited capacity for the creation of complexity. This isn't a problem in most industrial applications because their solution spaces are well explored by the search and optimisation strategies evolutionary computing offers. But art is one of the most complex activities of the arguably most complex unitary system known, the human mind.

A number of nature-inspired extensions for evolutionary art have been explored in part to meet this need for increased complexity. Each suggests new perspectives regarding computational aesthetic evaluation. For example, with the addition of co-evolution two or more species may compete for survival. This can create an evolutionary "arms race" making fitness a moving target for all. But it is also possible that species will coevolve to fill mutually beneficial symbiotic roles, and possibly exhibit convergent aesthetics. In such systems the ecology is a dynamic system offering increased complexity. Some species will specialise and dominate an ecological niche while others remain flexible generalists. And some species may in fact alter the ecology creating a niche for itself. Meanwhile, within a species individuals may interact via social transactions further modulating what constitutes fitness. These extensions are explored in the following sections.

10.2.11.1 Coevolution

Coevolution in evolutionary art and design has been investigated since at least 1995. Poon and Maher (1997) note that in design a fixed solution space is undesirable because the problem itself is often reformed based on interim discoveries. They suggest that both the problem space and solution space evolve with each providing feedback to the other. Each genotype in the population can combine a problem model and a solution in a single chromosome. Or there can be two populations, one for problems and one for solutions. Then current best solutions are used to select problem formulations, and current best problem formulations are used to select solutions. Both methods allow a form of multi-objective optimisation where the problem emphasis can shift and suggest multiple solutions, and well-matched problem formulations and solutions will evolve.

One challenge with coevolutionary systems is deciding when to stop the iterative process and accept a solution. The authors note termination can be based on satisfactory resolution of the initial problem, but that such an approach loses the benefit of the coevolved problem space. Other termination conditions can include the amount of execution time allowed, equilibrium where both the solution and problem spaces no longer exhibit significant change, or where a set of solutions cycle. The last case can indicate the formation of a Pareto-optimal surface of viable solutions (Poon and Maher 1997).

Todd and Werner were early adopters of a coevolutionary approach to music composition. Prior to their work there had been attempts to create fitness functions based on rule-based or learning-based *critics*. But such critics typically encouraged compositions that were too random, too static, or otherwise quite inferior to most human composition. It's worth remembering that genes in evolutionary systems seek high fitness scores and only secondarily produce desirable compositions. Sometimes trivial or degenerate compositions will exploit brittle models or faulty simulations, thereby "cheating" to gain a high score without providing a useful result.

Based on the evolution of bird songs through sexual selection, the system devised by Todd and Werner consists of virtual male composers that produce songs and virtual female critics that judge the songs for the purpose of mate selection. Each female maintains an expectation table of probabilities for every possible note-to-note transition. This table is used to judge males' songs in three ways. The first two methods reward males the more they match the female's expectations. In the third method males are rewarded for surprising females. And for each of these three methods transition tables can be static, or they can coevolve and slowly vary with each new generation of females.

The first two matching methods quickly suffered from a lack of both short term and long term variety. However, rewarding surprise lead to greater variety. One might expect that rewarding surprise would encourage random songs. But this didn't happen because random songs accidentally contain more non-surprise elements than songs specifically structured to set up expectations and then defy them.

Initially the females were created with transition tables derived from folk songs. At first this resulted in human-like songs. But the authors note:

One of the biggest problems with our coevolutionary approach is that, by removing the human influence from the critics (aside from those in the initial generation of folk-song derived transition tables), the system can rapidly evolve its own unconstrained aesthetics. After a few generations of coevolving songs and preferences, the female critics may be pleased only by musical sequences that the human user would find worthless.

Todd and Werner suggest that adding some basic musical rules might encourage diversity while also encouraging songs that are human-like. Additionally a learning and cultural aspect could be added by allowing individual females to change their transition tables based on the songs they hear (Todd and Werner 1998).

Greenfield (2008b) has presented an overview of coevolutionary methods used in evolutionary art including some unpublished systems made by Steven Rooke. Rooke first evolved critics by training them to match his manually given scores for a training set of images. The critics then coevolve with new images. Individual critics are scored by comparing their evaluations to those of previous critics. Critics are maintained over time in a sliding window of 20 previous generations. Rooke found that while the coevolved critics duplicated his taste, the overall system didn't innovate by exploring new forms.

Greenfield then describes his own system where images and 10×10 convolution filters are coevolved. Parasite filters survive by generating result images similar to original. Images survive by making the parasite filter results visible. A number of subtleties require attention such as setting thresholds that define similarity, the elimination of do-nothing filters, adjusting the evolutionary rates of parasites versus images, and the balancing of unary and binary operators to control high frequency banding. He cites Ficici and Pollack (1998) and confirms observing *evolutionary cycling*, where genotypes are rediscovered again and again, and *mediocre stable* states where the coevolving populations exhibit constant change with little improvement. Greenfield notes:

In all of the examples we have seen: (1) it required an extraordinary effort to design a population to coevolve in conjunction with the population of visual art works being produced by an underlying image generation system, and (2) it was difficult to find an evaluation scheme that made artistic sense. Much of the problem with the latter arises as a consequence of the fact that there is very little data available to suggest algorithms for evaluating aesthetic fitness. . . It would be desirable to have better cognitive science arguments for justifying measurements of aesthetic content.

In later sections we will survey some of the work in psychology and the nascent field of neuroaesthetics that may contribute to computational aesthetic evaluation as Greenfield suggests.

10.2.11.2 Niche Construction by Agents

As discussed in McCormack and Bown (2009) an environment can be thought of as having both *properties* and *resources*. Properties are environmental conditions such as temperature or pH, and resources are available consumables required by

organisms such as water and specific kinds of food. Each organism will have specific needs as to the properties and resources it requires of its environment. A given organism's preferred properties and resources define its ecological niche.

In typical "artificial life" systems evolutionary computing is implemented within the context of a simulated ecosystem. In those systems adaptation to ecological niches can increase diversity and enhance multi-objective optimisation. But beyond simple adaptation genotypes within a species can actively construct niches to their own advantage. McCormack and Bown have demonstrated both a drawing system and a music system that exploit niche construction.

In the first system drawing agents move leaving marks, are stopped when they intersect already existing marks, and sense the local density of already existing marks. Each agent also has a genetic preference for a given density. Initially agents that prefer low density will succeed in dividing large open sections of the canvas. Over time some agents will create higher densities of marks, which in turn act as constructed niches for progeny with a predisposition for high density. As a result some, but not all, sections of the canvas become increasingly dense and provide niches for high-density genotypes. The visual result exhibits a wide range of densities. Similar agent-based systems without niche construction tend to create drawings with homogeneous density. This system is further discussed in Chap. 2.

In the second system a single row of cells is connected head-to-tail as a toroid. Each cell generates a sine wave creating a single frequency tone. A line runs through all of the cells, and at each cell the line height is mapped into the loudness of its sine wave. Agents inhabit the cells, and each has a genetic preference for line height and slope. Each agent applies these preferences as pressure to the line in its cell as well as the cell to its left. Depending on the local state of their niche, i.e. the line height and slope in their cell, agents will stay alive and reproduce or die and not pass on their genotype. This sets up a dynamic system with localities that benefit certain genotypes. Those genotypes then modify the ecosystem, i.e. the line, to the benefit of their progeny. The resulting sound exhibits a surprising diversity of dynamics even though it is initialised at zero. As with many evolutionary and generative systems, this is due to the random variation in the initial population of agents.

10.2.11.3 Agent Swarm Behaviour

In most of the evolutionary systems discussed so far there is no interaction between phenotypes. Each is independently evaluated via user selection or fitness function. Other than this comparison a given phenotype has no impact on another. When phenotypes begin to interact in other ways, typically in the context of a simulated ecosystem, they can be thought of as simulated organisms or *agents* that exhibit behaviours. With niche creation agents modify their ecology establishing a mediated form of agent interaction. But agents can also interact directly creating an emergent group behaviour or *swarm behaviour*.

The canonical natural example of such an agent is the ant. An ant colony uses swarm intelligence to optimise the gathering and retrieval of food. As an ant finds

food and brings it back to the nest it selectively leaves a chemical pheromone trail. Other ants happening upon the chemical trail will follow it, in effect joining a food retrieval swarm. Each ant adds more pheromone as they retrieve food. Because the pheromone spreads as it dissipates ants will discover short cuts if the initial path has excessive winding. In turn those short cuts will become reinforced with additional pheromone. Once the food is gone the ants stop laying down pheromone as they leave the now depleted site, and soon the pheromone trail will disappear. This behaviour can be simulated in software agents (Resnick 1994).

Artists have simulated this behaviour in software using agents that lay down permanent virtual pigment as well as temporary virtual pheromone trails. Variation and some degree of aesthetic control can be gained by breeding the ant-agents using an interactive evolutionary system (Monmarché et al. 2003).

Greenfield (2005a) automates the fitness function based on a performance metric regarding the number of cells visited randomly or due to pheromone following behaviour. Measuring fitness based only on the number of unique cells visited results in “monochromatic degeneracies”. Rewarding only pheromone following creates a slightly more attractive blotchy style. Various weightings of both behaviours produce the best aesthetic results exhibiting organic and layered forms.

Urbano (2006) has produced striking colourful patterns using virtual micro-painters he calls “Gaugants”. In the course of one-to-one transactions his agents exert force, form consensus, or exhibit dissidence regarding paint colour. The dynamics are somewhat reminiscent of scenarios studied in game theory. Elzenga’s agents are called “Arties”. They exhibit mutual attraction/repulsion behaviour based on multiple sensing channels and genetic predisposition. The exhibited emergence is difficult to anticipate, but the artist can influence the outcome by making manual selections from within the gene pool (Elzenga and Pontecorvo 1999).

10.2.11.4 Curious Agents

Saunders and Gero (2004), and Saunders (2002) have extended swarming agents to create what they have called *curious agents*. They first note that agents in swarm simulations such as the above are mostly reactive. Flocking was originally developed by Reynolds (1987) and then extended by Helbing and Molnar (1995; 1997) to add *social forces* such as goals, drives to maximise efficiency and minimise discomfort, and so on. Social forces have been shown, for example, to create advantages in foot traffic simulation.

Sanders and Gero expand the dynamics of aesthetic evaluation behaviour by adding curiosity as a new social force. Their implementation uses a pipeline of six primary modules for sensing, learning, detecting novelty, calculating interest, planning, and acting. Sensing provides a way to sample the world for stimulus patterns. Learning involves classifying a pattern and updating prototypes kept in long term memory. Novelty is assessed as the degree to which error or divergence from previous prototypes is detected. Based on novelty a measure of interest is calculated.

Changes in interest result in goals being updated, and the current ever-changing goals determine movement.

Unsupervised artificial neural networks are used for classification, and classification error for new inputs is interpreted as novelty. But greater novelty doesn't necessarily result in greater interest. The psychologist Daniel Berlyne proposed that piquing interest requires a balance of similarity to previous experience and novelty. So, as suggested by Berlyne (1960; 1971), a *Wundt curve* is used to provide the metric for this balance and produces an appropriate interest measure. More about Berlyne's work follows in Sect. 10.3.2.

Based on this model Sanders created an experimental simulation where agents enter a gallery, can sense other agents, and can also view the colours of monochrome paintings hanging on nearby walls. There are also unseen monochrome paintings with new colours in other rooms. Along with other social behaviours agents learn the colours presented in one room, and then are potentially curious about new colours in other rooms. Depending on the sequence of colour exposure and the related Wundt-curve mapping, agents may or may not develop an interest and move to other areas.

10.2.11.5 Human Aesthetics, Meta-aesthetics, and Alternatives to Fitness Functions

Commenting on systems like those above using coevolution, niche creation, swarms, and curiosity Dorin (2005) notes:

... the "ecosystemic" approach permits simultaneous, multidirectional and automatic exploration of a space of virtual agent traits without any need for a pre-specified fitness function. Instead, the fitness function is implicit in the design of the agents, their virtual environment, and its physics and chemistry.

This avoids the problem of creating a computational aesthetic evaluation system by hand, and allows for the creation of evolutionary systems that generate surprising diversity and increased dynamics. Thus, if the goal is the creation of robust systems for meta-aesthetic exploration these evolutionary system extensions seem to be quite beneficial.

However, if the goal is to evolve results that appeal to our human sense of aesthetics there is no reason to think that will happen. Recall the earlier differentiation between human aesthetic evaluation and meta-aesthetic explorations. Creating evolutionary diversity and dynamics via artificial aesthetics foreign to our human sensibility is one thing. Appealing to human aesthetics is quite another. As observed by Todd and others, to date extensions and emergent aesthetics like those above do not provide machine evaluation that mirrors human aesthetic perception.

10.2.12 Complexity Based Models of Aesthetics

One of the recurring themes in computational aesthetics is the notion that aesthetic value has something to do with a balance of complexity and order. Birkhoff's

aesthetic measure proposed the simple ratio $M = O/C$ where M is the measure of aesthetic effectiveness, O is the degree of order, and C is the degree of complexity.

But what is complexity? And what is order? Birkhoff suggested that these are proxies for the effort required (complexity) and the tension released (order) as perceptual cognition does its work. As a practical matter Birkhoff quantified complexity and order using counting operations appropriate to the type of work in question. For example, in his study of polygonal compositions complexity was determined by counting the number of edges and corners. His formula for order was:

$$O = V + E + R + HV - F \tag{10.2}$$

Here he sums the vertical symmetry (V), equilibrium (E), rotational symmetry (R), horizontal-vertical relation (HV), and unsatisfactory or ambiguous form (F). These notions of complexity and order at first appear to be formulaic and objective, but they nevertheless require subjective decisions when quantified.


In an attempt to add conceptual and quantitative rigour, Bense (1965) and Moles (1966) restated Birkhoff's general concept in the context of Shannon (1948)'s information theory creating the study of *information aesthetics*. Shannon was interested in communication channels and the quantification of information capacity and signal redundancy. From this point of view an entirely unpredictable random signal maximises information and complexity, and offers no redundancy or opportunity for lossless compression. In this context disorder or randomness is also called *entropy*. Extending this, Moles equated low entropy with order, redundancy, compressibility, and predictability. High entropy was equated with disorder, complexity, incompressibility, and surprise (see Chap. 3 for further discussion of information aesthetics).

As previously noted, Machado (1998) has updated this approach by calculating aesthetic value as the ratio of *image complexity* to *processing complexity*. Processing complexity refers to the amount of cognitive effort that is required to take in the image. Image complexity is intrinsic to the structure of the image. This led them to propose functional measures where image complexity is inversely proportional to JPEG compressibility and processing complexity is directly proportional to fractal compressibility.

With the advent of complexity science as a discipline defining order and complexity has become much more problematic. This account begins with *algorithmic complexity* or *algorithmic information content* as independently developed by Kolmogorov (1965), Solomonoff (1964), Chaitin (1966). In this paradigm the complexity of an object or event is proportional to the size of the shortest program on a universal computer that can duplicate it. From this point of view the most complex music would be white noise and the most complex digital image would be random pixels. Like information complexity, algorithmic complexity is inversely proportional to order and compressibility.

For physicist Murray Gell-Mann the information and algorithmic notions of complexity don't square with our experience. When we encounter complex objects or situations they aren't random. Despite being difficult to predict they also have some degree of order maintaining integrity and persistence.

Fig. 10.1 Information and algorithmic complexity increase monotonically with increasing disorder. Effective complexity peaks where there is a mix of order and disorder such as is found in biological life



a63c5a3857db9a4d63ca53b55e68664f
ebrary

Consider two situations, one where there is a living frog and another where there is a long dead and decaying frog. The decaying frog has greater entropy because relative to the living frog it is more disordered, and over time it will become more even more disordered to the point where it will no longer be identifiable as a frog at all. Intuitively we would identify the living frog as being more complex. It displays a repertoire of behaviours, operates a complex system of biochemistry to process food, water, and oxygen to generate energy and restore tissues, maintains and exchanges large amounts of genetic information in the course of reproduction, and so on. Along with these orderly processes the frog remains flexible and unpredictable enough to be adaptive and avoid becoming easy prey. In terms of entropy our highly complex living frog is somewhere between simple highly ordered crystals and simple highly disordered atmospheric gases.

To better capture our intuitive sense of complexity Gell-Mann has proposed the notion of *effective complexity*, a quantity that is greatest when there is a balance of order and disorder such as that found in the biological world (Gell-Mann 1995). Unlike information and algorithmic complexity, effective complexity is not inversely proportional to order and compressibility. Rather both order and disorder contribute to complexity (Fig. 10.1, please note that this graph is only meant as a qualitative illustration with somewhat arbitrary contours).

a63c5a3857db9a4d63ca53b55e68664f
ebrary

Complexity science continues to offer new paradigms and definitions of complexity. In a 1998 lecture by Feldman and Crutchfield at the Santa Fe Institute well over a dozen competing theories were presented (Feldman and Crutchfield 1998)—the debate over complexity paradigms continues. Measuring aesthetic value as a relationship between complexity and order is no longer the simple proposition it once seemed to be. (For an alternate view of complexity and aesthetics see Chap. 12.)

Artists working in any media constantly seek a balance between order and disorder, i.e. between fulfilling expectations and providing surprises. Too much of the former leads to boredom, but too much of the latter loses the audience. It is a dynamic that applies to visual art, music, and the performing arts alike. And it helps differentiate genres in that styles that cater to established expectations are considered

a63c5a3857db9a4d63ca53b55e68664f
ebrary

to be more “traditional” while styles that serve up more unorthodox surprises are considered to be “cutting edge.”

Notions of Shannon information and algorithmic complexity have their place. But in aesthetics it is misleading to treat order and complexity as if they are polar opposites. My suggestion is that the notion of effective complexity better captures the balance of order and disorder, of expectation and surprise, so important in the arts. This offers the challenge and potential benefit that effective complexity can serve as a measure of quality in computational aesthetic evaluation.

10.3 The Future of Computational Aesthetic Evaluation

As should be obvious by now, computational aesthetic evaluation is a very difficult and fundamentally unsolved problem. To date any marginal successes have tended towards narrow application niches using methods that do not generalise very well.

The irony is that aesthetic evaluation is something we all do quite naturally. Could it be that the solution to the computational aesthetic evaluation problem is within us and just not yet understood?

Artists and engineers have always learned from nature. There is a significant and growing literature around the psychology and neurology of aesthetics. But this challenge to understanding seems no less daunting than the difficulty of machine evaluation. The human brain that gives rise to the human mind is arguably the most complex unitary system known. The brain includes approximately 10^{15} neural connections. In addition, recent research regarding the brain’s glial cells reveals that they contribute to active information processing rather than, as previously thought, merely providing mechanical support and insulation for neurones. Glial cells make up 90 % of the brain and some scientists speculate that they are specifically engaged in creative thought (Koob 2009). Computing hardware can only make up for part of this gap by exploiting electronic switching speeds that are about 10^7 times faster than human neurones.

Nevertheless, it seems reasonable that an improved understanding of natural aesthetic perception will contribute to computational aesthetic evaluation efforts, and science has made some significant progress in this regard. Perhaps a good place to start is recent scientific thinking as to the origins of human aesthetics.

10.3.1 The Origins of Art and the Art Instinct

Denis Dutton notes that evolutionary scientist Stephen Jay Gould claims that art is essentially a nonadaptive side effect, what Gould calls a *spandrel*, resulting from an excess of brain capacity brought about by unrelated adaptations. Dutton (2009) argues that the universality of both art making behaviour and some aesthetic preferences imply a more direct genetic linkage and something he calls the *art instinct*.

Dutton points out that like language every culture has art. And both language and art have developed far beyond what would be required for mere survival. The proposed explanation for the runaway development of language is that initially language provided a tool for cooperation and survival. Once language skills became important for survival language, fluency became a mate selection marker. The genetic feedback loop due to mate selection then generated ever-increasing language ability in the population leading to a corresponding *language instinct* (Pinker 1994).

Additionally, Dutton posits that early human mate selection was, in part, based on the demonstration of the ability to provide for material needs. Like language, this ability then became a survival marker in mate selection subject to increasing development. Just as a peacock's feather display marks a desirable surplus of health, works of art became status symbols demonstrating an excess of material means. It is not by coincidence then that art tends to require rare or expensive materials, significant time for learning and making, as well as intelligence and creativity. And typically art has a lack of utility, and sometimes an ephemeral nature. All of these require a material surplus.

One could argue that even if art making has a genetic basis it may be that our sense of aesthetics does not. In this regard, Dutton notes the universal appeal, regardless of the individual's local environment, for landscape scenes involving open green spaces trees and ample bodies of water near by, an unimpeded view of the horizon, animal life, and a diversity of flowering and fruiting plants. This scene resembles the African savannah where early man's evolution split off from other primate lines. It also includes numerous positive cues for survivability. Along with related psychological scholarship Dutton quotes the previously noted Alexander Melamid:

... I'm thinking that this blue landscape is more serious than we first believed... almost everyone you talk to directly—and we've already talked to hundreds of people—they have this blue landscape in their head... So I'm wondering, maybe the blue landscape is genetically imprinted in us, that it's the paradise within, that we came from the blue landscape and we want it... We now completed polls in many countries—China, Kenya, Iceland, and so on—and the results are strikingly similar.

That our aesthetic capacity evolved in support of mate selection has parallels in other animals. This provides some hope for those who would follow a psychological path to computational aesthetic evaluation, because creatures with simpler brains than man practice mate selection. In other words perhaps the computational equivalent of a bird or an insect is "all" that is required for computational aesthetic evaluation. But does mate selection behaviour in other animals really imply brain activity similar to human aesthetic judgement? One suggestive study by Watanabe (2009) began with a set of children's paintings. Adult humans judged each to be "good" or "bad". Pigeons were then trained through operant conditioning to only peck at good paintings. The pigeons were then exposed for the first time to a new set of already judged children's paintings. The pigeons were quite able to correctly classify the previously unseen paintings as "good" or "bad".

10.3.2 Psychological Models of Human Aesthetics

Conspicuously missing from most work by those pursuing machine evaluation that mimics human aesthetics are models of how natural aesthetic evaluation occurs. Rudolf Arnheim, Daniel Berlyne, and Colin Martindale are three researchers who stand out for their attempts to shape the findings of empirical aesthetics into general aesthetic models that predict and explain. Each has left a legacy of significant breadth and depth that may inform computational aesthetic evaluation research. The following sections provide an introduction to their contributions.

10.3.2.1 Arnheim—Gestalt and Aesthetics

If one had to identify a single unifying theme for Arnheim it would have to be the notion of perception as cognition. Perception isn't something that happens to the brain when events in the world are passively received through the senses. Perception is an activity of the brain and nothing short of a form of cognition. And it is this perceptual cognition that serves as the engine for gestalt phenomena.

First written in 1954 and then completely revised in 1974, Arnheim's book *Art and Visual Perception: A Psychology of the Creative Eye* established the relevance of gestalt phenomena as art and design principles (Arnheim 1974). *The law of prägnanz* in gestalt states that the process of perceptual cognition endeavours to order experience into wholes that maximise clarity of structure. From this law come the notions of closure, proximity, containment, grouping, and so on now taught as design principles (Wertheimer 2007).

The neurological mechanisms behind these principles were not, and still are not, well understood. Arnheim wrote of *forces* and *fields* as existing both as psychological and physical entities; the physical aspects being neurological phenomenon in the brain itself. Some have suggested it is more useful to take these terms metaphorically to describe the dynamic tensions that art exercises (Cupchik 2007).

Arnheim's theory of aesthetics is much more descriptive than normative. Nevertheless, those interested in computational aesthetic evaluation have much to take away with them. That perception is an active cognitive process, and that the gestalt whole is something more than the sum of the parts, is now taken by most as a given. And the difference between maximising clarity of structure and maximising simplicity of structure is a nuance worthy of attention (Verstegen 2007).

10.3.2.2 Berlyne—Arousal Potential and Preferences

Daniel E. Berlyne published broadly in psychology, but his work of note here regards physiological arousal and aesthetic experience as a neurological process (Konečni 1978). One of Berlyne's significant contributions is the concept of *arousal potential* and its relationship to *hedonic response*.

Arousal potential is a property of stimulus patterns and a measure of the capability of that stimulus to arouse the nervous system. Arousal potential has three sources. First, there are *psychophysical* properties such as very bright light, very loud sounds, sensations with an abrupt onset, very low or high frequency sounds, and so on. Second, there are *ecological* stimuli such as survival threats like pain or predator sightings, or cues associated with the availability of food. But the third and strongest according to Berlyne are referred to as *collative* effects. These are combined and comparative experiences that present arousal potential in a context dependent and relative manner. Examples include “novelty, surprisingness, complexity, ambiguity, and puzzlingness.” Berlyne (1971) explicitly notes the correspondence between many of these collative effects and concepts from Shannon’s information theory.

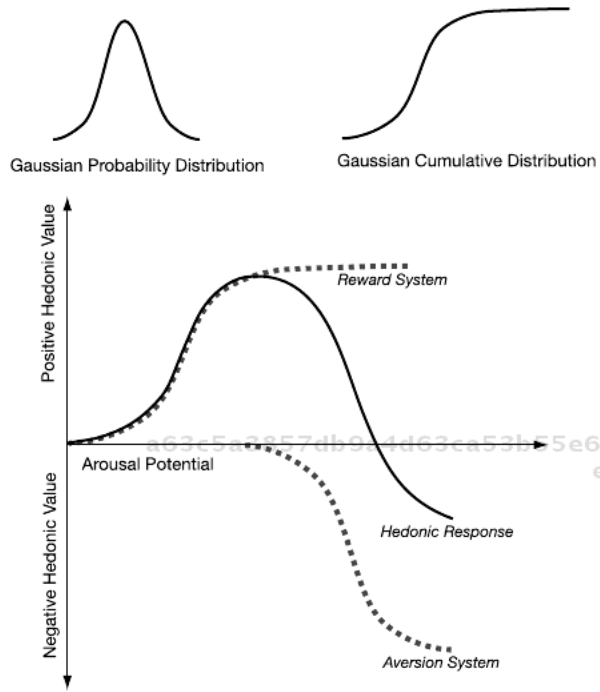
The hedonic response to sources of arousal potential refers to the spectrum of pleasure and pain we experience. Berlyne proposes that the hedonic response is the result of separate and distinct reward and aversion systems. Each of these systems is made up of neurones. The firing thresholds of individual neurones will vary according to the *normal* or *Gaussian probability* distribution as is typical in nature (see Fig. 10.2). Therefore the strength of the arousal potential will determine the number of neurones that fire in response. The number of neurones responding will increase as a *Gaussian cumulative distribution*, i.e. the area under the Gaussian probability distribution as the threshold moves from left to right. Berlyne further proposes that the reward system requires less arousal potential exposure to activate, but that when activated the aversion system will produce a larger response.

The result is the hedonic response as a summation of the positive reward system and the negative aversion system. With no arousal potential there is a hedonic response of indifference. As more arousal potential is presented the hedonic response increases manifesting itself as a pleasurable experience. Beyond a certain point, however, two things happen. First, the reward system reaches maximum activation and plateaus. Second, the aversion system begins to activate. As the aversion system reaches higher levels of activation the hedonic response will lessen and eventually cross into increasing levels of pain.

Berlyne notes that this function is usually called *the Wundt curve*, as it was first presented by the “father of experimental psychology” Wilhelm Wundt in 1874. But in Wundt’s model the *x*-axis represents low-level neural intensity. Berlyne’s arousal potential on the *x*-axis includes psychophysical intensity, but it also includes ecological stimuli and most importantly collative effects. For Berlyne increasing collative effects such as novelty and surprise also represent increasing complexity in the information theory sense. From this point of view works of only moderate information complexity maximise the hedonic response. This resonates well with the intuitive artistic notion that audiences respond best to works that are not so static as to be boring, and yet also operate within learned conventions so as to not be experienced as chaotic.

There is, however, another interpretation. The notion of Gell-Mann’s effective complexity was previously mentioned. From that point of view complexity is a balance of order and disorder, and biological life presents complexity at its peak.

Fig. 10.2 Wundt curve as applied by Berlyne. Redrawn from Berlyne (1971)



The Wundt and effective complexity curves both peak in the middle suggesting that positive hedonic response may be proportional to effective complexity. Effective complexity has, in a sense, the balance of order and disorder “built in.” One might hypothesise that the most important and challenging survival transactions for humans have to do with other living things and especially fellow humans. Perhaps that created evolutionary pressure leading to the optimisation of the human nervous system for effective complexity, and human aesthetics and related neurological reward/aversion systems reflect that optimisation.

10.3.2.3 Martindale—Prototypicality and Neural Networks

Colin Martindale was an active empiricist and in 1990 he published a series of articles documenting experiments intended to verify the arousal potential model of Berlyne. Martindale et al. (1990) notes:

Berlyne... developed an influential theory that has dominated the field of experimental aesthetics for the past several decades... Berlyne is often cited in an uncritical manner. That is, he is taken as having set forth a theory based upon well-established facts rather than, as he actually did, as having proposed tentative hypotheses in need of further testing. The result has been a stifling of research on basic questions concerning preference, because these questions are considered to have been already answered. In this article, we report a series of experiments that test obvious predictions drawn from Berlyne’s theory. It was in the firm

expectation of easily confirming these predictions that we undertook the experiments. The results are clear-cut. They do not support the theory.

The debate pitting collative effects versus prototypicality would dominate experimental aesthetics for almost 20 years (North and Hargreaves 2000). For some Berlyne's notion of collative effects was especially problematic. First it was odd for a behaviourist like Berlyne to make an appeal to a concept so much about the inner state of the individual. Additionally, terms like *novelty* and *complexity* were problematic both in specification and mechanism.

However, Martindale's primary critique was empirical. For example, contrary to Berlyne's model he found that psychophysical, ecological, and collative properties are not additive, nor can they be traded off. Significantly more often than not empirically measured responses do not follow the inverted-U of the Wundt curve, but are monotonically increasing. Finally, a number of studies showed that meaning rather than pure sensory stimulation is the primary determinant of aesthetic preference (Martindale et al. 1990; 2005, Martindale 1988b).

In a series of publications Martindale (1981; 1984; 1988a; 1991) developed a natural neural network model of aesthetic perception that is much more consistent with experimental observation. Martindale first posits that neurones form nodes that accept, process, and pass on stimulation from lower to higher levels of cognition. Shallow sensory and perceptual processing tends to be ignored. It is the higher semantic nodes, the nodes that encode for meaning, that have the greatest strength in determining preference. Should the work carry significant emotive impact the limbic system can become engaged and dominate the subjective aesthetic experience.

Nodes are described as specialised recognition units connected in an excitatory manner to nodes corresponding to superordinate categories. So, for example, while one is reading nodes that extract features will excite nodes for letters, and they will in turn excite nodes for syllables or letter groupings, leading to the excitation of nodes for words, and so on. Nodes at the same level, however, will have a lateral inhibitory effect. Nodes encoding for similar stimuli will be physically closer together than unrelated nodes. So nodes encoding similar and related exemplars will tend towards the centre of a *semantic field*. The result is that the overall nervous system will be optimally activated when presented an unambiguous stimulus that matches a prototypically specific and strong path up the neural hierarchy (Martindale 1988b).

Commenting on prototypicality North and Hargreaves (2000) explain:

... preference is determined by the extent to which a particular stimulus is typical of its class, and explanations of this have tended to invoke neural network models of human cognition: this approach claims that preference is positively related to prototypicality because typical stimuli give rise to stronger activation of the salient cognitive categories.

Martindale's neural network prototypicality model carries with it great explanatory and predictive power. Towards the end of his life he penned a chapter describing the results of 25 widely disparate empirical studies, and how his single model can provide a foundation for understanding all of them (Martindale 2007).

While most in the field agree that Martindale's prototypicality model explains more of the empirical data than Berlyne's collative effect model, some cases remain where prototypicality is the weaker explanation. Some have suggested ways

to reconcile the two models to provide more cover than either can alone (North and Hargreaves 2000, Whitfield 2000).

10.3.3 Empirical Studies of Human Aesthetics

Along with unifying theories such as those offered by Arnheim, Berlyne, and Martindale, the field of psychology offers a vast catalogue of very specific findings from experimental aesthetics. It is difficult in aesthetics research to identify and control the myriad factors that may influence hedonic response. And because human subjects are typically required it is difficult to achieve large sample sizes. Nevertheless empirical studies of human aesthetics seem to be on the increase, and many are highly suggestive and worth consideration by those interested in computational aesthetic evaluation.

Empirical studies of human aesthetics usually focus on viewers, artists, or objects. Studies of viewers have to account for audiences that are expert and not. Some experiments focus on the impact setting has on aesthetic perception. Others are attempts to correlate aesthetic response with social or personality factors. Studies of artists usually focus on aspects of divergent thinking, creativity, and self-critical abilities. Studies of objects typically include some form of analysis relative to a hypothesised aesthetic mechanism.

A full or even representative cataloguing of these studies is unfortunately well outside of the scope of this chapter. What stands out in reading the literature though is the large number of variables that determine or shade human aesthetic experience. For example:

- Subjects first asked to think about the distant future are more likely to accept unconventional works as art than those who first think about their near future (Schimmel and Forster 2008).
- A hedonic contrast effect has been established in music listening. In absolute terms the same music will be evaluated more positively if preceded by bad music, and less positively if preceded by good music (Parker et al. 2008).
- Not all emotions lend themselves to musical expression. Those that do tend to be general, mood based, and don't require causal understanding (Collier 2002).
- Individual preference differences can form on the basis of experience. Relative to non-professionals, photo professionals exhibit a greater ability to process photographic information, and show a relative preference for photographs that are uncertain and unfamiliar (Axelsson 2007).
- Artists and non-artists were presented with a sequence of 22 work-in-process images leading to Matisse's 1935 painting, *Large Reclining Nude*. Non-artists judged the painting as getting generally worse over time consistent with the increasing abstraction of the image. In contrast, art students' judgements showed a jagged trajectory with several peaks suggesting an interactive hypothesis-testing process (Kozbelt 2006).

- Whether isolated or within a larger composition, note intervals in music carry significant and consistent emotional meaning. There is also softer evidence that these interval-emotional relationships are universal across different times, cultures, and musical traditions. Speculation is that this is related to universal aspects of vocal expression (Oelmann and Laeng 2009).

10.3.4 *Neuroaesthetics*

Beginning with Birkhoff, and throughout this chapter, neurology has frequently been the backstory for aesthetic and computational aesthetic evaluation models described at higher levels of abstraction. To some extent Arnheim, and certainly Berlyne and Martindale, all had in mind neurological models as the engines of aesthetic perception. In no small part due to new imaging technologies such as functional magnetic resonance imaging (fMRI), positron emission tomography (PET) scanning, and functional near-infrared imaging (fNIR), science seems to be preparing to take on perhaps the deepest mystery we face everyday, our own minds.

It is in this context that the relatively new field of neuroaesthetics has come into being (Skov and Vartanian 2009a). Neuroaesthetics is the study of the neurological bases for all aesthetic behaviour including the arts. A fundamental issue in neuroaesthetics is fixing the appropriate level of inspection for a given question. It may be that the study of individual neurones will illuminate certain aspects of aesthetics. Other cases may require a systems view of various brain centres and their respective interoperation.

A better understanding of representation in the brain could illuminate not only issues in human aesthetics but more generally all cognition. This in turn may find application not only in computational aesthetic evaluation, but also broadly across various artificial intelligence challenges. And finally, a better understanding of neurology will likely suggest new models explaining human emotion in aesthetic experience. If we better understand the aesthetic contributions of both the cortex and the limbic system, we will be better prepared to create machine evaluation systems that can address both the Dionysian and the Apollonian in art (Skov and Vartanian 2009b).

10.3.5 *Computing Inspired by Neurology*

Computer science has felt the influence of biology and brain science from its earliest days. The theoretical work of Von Neumann and Burks (1966) towards a universal constructor was an exploration of computational reproduction and evolution. Turing (1950) proposed a test essentially offering an operational definition for machine intelligence. Turing also invented the reaction diffusion model of biological morphogenesis, and towards the end of that article he discusses implementing a computer simulation of it (Turing 1952). Computing models inspired by neurology have

fallen in and out of fashion, from Rosenblatt's early work on the perceptron (Rosenblatt 1962), to Minsky and Papert's critique (Minsky and Papert 1969), and to the later successful development of non-linear models using backpropagation and self-organisation.

A number of artificial neural network applications already noted showed only limited success as either a fitness function or a standalone machine evaluation system. It would be premature to conclude such use in has hit a permanent plateau. But it would be glib to suggest that since the brain is a neural network that the successful use of artificial neural networks for computational aesthetic evaluation is inevitable. The brain's 10^{15} neural connections and presently unknown glial cell capacity presents a daunting quantitative advantage artificial systems will not match any time soon.

Perhaps a better understanding of natural neurology and subsequent application to connectionist technologies can help overcome what present artificial systems lack in quantity. This is the approach Jeff Hawkins has taken in the development of hierarchical temporal memory.

10.3.6 The Neocortex and Hierarchical Temporal Memory

Hawkins has proposed the *hierarchical temporal memory model* for the functionality found in the neocortex of the brain. He proposes that this single mechanism is used for all manner of higher brain function including perception, language, creativity, memory, cognition, association, and so on. He begins with a typical hierarchical model where lower cortical levels aggregate inputs and pass the results up to higher levels corresponding to increasing degrees of abstraction (Hawkins and Blakeslee 2004).

Neurologists know that the neocortex consists of a repeating structure of six layers of cells. Hawkins has assigned each layer with functionality consistent with the noted multi-level hierarchical structure. What Hawkins has added is that within a given level higher layers constantly make local predictions as to what the next signals passed upward will be. This prediction is based on recent signals and local synapse strength. Correct predictions strengthen connections within that level. Thus the neocortex operates as a type of hierarchical associative memory system, and it exploits the passage of time to create local feedback loops for constant training.

Artificial hierarchical temporal memory has been implemented as software called *NuPIC*. It has been successfully demonstrated in a number of computer vision applications where it can robustly identify and track moving objects, as well as extract patterns in both physical transportation and website traffic (Numenta 2008). To date NuPIC seems to work best when applied to computer vision problems, but others have adapted the hierarchical temporal memory model in software for temporal patterns in music (Maxwell et al. 2009).

10.3.7 Computer Architectures for Evolvable Hardware

Another promising technology is reconfigurable hardware that evolves in a way to best solve the problem at hand. Evolvable hardware exploits programmable circuit devices such as *field programmable gate arrays* (FPGAs). These are integrated circuit chips with a large number of simple logic units or *gates*. Settable switches called *architecture bits* or *configuration memory* program the logical function and interconnection of these gates. Field programmable gate arrays allow the mass manufacture of standardised silicon that has its circuit-level functionality postponed for later definition. This circuit-level functionality is lower and faster than that achieved by executing machine language code (Yao and Higuchi 1997).

By treating the architecture bits as a chromosome the configuration of field programmable gate arrays can be determined using evolutionary methods. Evolution in this case doesn't design the gate array configuration so much as it designs the chip's behaviour relative to some fitness function defined need. In this some see a parallel to the way neurones exhibit emergent learning. And because these chips can be reprogrammed on the fly there is the possibility of learning adaptation to changing conditions.

It's worth noting that a proposed evolvable hardware system has been simulated in software, and used as a pattern recognition system for facial recognition with an experimental accuracy of 96.25 % (Glette et al. 2007).

10.4 Conclusion

Computational aesthetic evaluation victories have been few and far between. The successful applications have mostly been narrowly focused point solutions. Negative experience to date with low dimensional models such as formulaic and geometric theories makes success with similar approaches in the future quite unlikely.

Evolutionary methods, including those with extensions such as coevolution, niche construction, and agent swarm behaviour and curiosity, have had some circumscribed success. The noted extensions have allowed evolutionary art to iterate many generations quickly by eliminating the need for interactive fitness evaluation. They have also allowed researchers to gain insight into how aesthetic values can be created as emergent properties. In such explorations, however, the emergent artificial aesthetics themselves seem alien and unrelated to human notions of beauty. They have not yet provided practical leverage when the goal is to model, simulate, or predict human aesthetics via machine evaluation.

I've suggested that a paradigm like effective complexity may be more useful than information or algorithmic complexity when thinking about aesthetics. Effective complexity comes with the notion of balancing order and disorder "built in", and that balance is critical in all forms of aesthetic perception and the arts.

There is also a plausible evolutionary hypothesis for suggesting that effective complexity correlates well with aesthetic value. Effective complexity is maximised

in the very biological systems that present us with our greatest opportunities and challenges. Hence there is great survival value in having a sensory system optimised for the processing of such complexity. There is also additional survival value in our experiencing such processing as being pleasurable. As in other neurological reward systems such pleasure directs our attention to where it is needed most.

The fields of psychology and neurology have been noted as possible sources of help for future work in computational aesthetic evaluation. Models of aesthetic perception such as those from Arnheim, Berlyne, and especially Martindale invite computational adaptation. Results from empirical studies of human aesthetics can stimulate our thinking about computational evaluation. At the same time they warn us that aesthetic evaluation in humans is highly variable depending on setting, context, training, expectations, presentation, and likely dozens of other factors.

Will robust human-like computational aesthetic evaluation be possible someday? There is currently no deductive proof that machine evaluation either is or isn't possible in principle. Presumably an argument for impossibility would have to establish as key an aspect of the brain or human experience that goes beyond mechanical cause and effect. Others might argue that because the brain itself is a machine our aesthetic experience is proof enough that computational aesthetic evaluation is possible. These in-principle arguments parallel philosophical issues regarding phenomenology and consciousness that are still in dispute and far from settled.

As a practical matter, what is currently possible is quite limited. The one consistent thread that for some will suggest a future direction relates to connectionist approaches. The current leading psychological model, Martindale's prototypicality, presents natural aesthetic evaluation as a neural network phenomenon. We know that animals with natural neural systems much simpler than those in the human brain are capable of some forms of aesthetic evaluation. In software, new connectionist computing paradigms such as hierarchical temporal memory show promise for both higher performance and closer functional equivalency with natural neural systems. In hardware we are beginning to see systems that can dynamically adapt to problem domains at the lowest gate level. Perhaps this will all someday lead to a synergy of hardware, software, and conceptual models yielding success in computational aesthetic evaluation.

Acknowledgements My interest in writing this chapter began at the "Computational Creativity: An Interdisciplinary Approach" seminar in July of 2009 at the Schloss Dagstuhl—Leibniz Center for Informatics. I would like to thank Margaret Boden, Mark d'Inverno and Jon McCormack for organising the seminar. In addition my thanks go to my fellow members of the "Evaluation" discussion group at the seminar including Margaret Boden, David Brown, Paul Brown, Harold Cohen, and Oliver Deussen. Finally I enjoyed and appreciated the lively post-seminar e-mail discussion of related topics with David Brown, Paul Brown, Harold Cohen, Jon McCormack, and Frieder Nake. Please note, however, that any matters of opinion or error in this chapter are purely my own.

References

Aguilar, C., & Lipson, H. (2008). A robotic system for interpreting images into painted artwork.

- In C. Soddu (Ed.), *International conference on generative art* (Vol. 11). Generative Design Lab, Milan Polytechnic.
- Aldiss, B. (2002). The mechanical turk—the true story of the chess-playing machine that changed the world. *TLS-the Times Literary Supplement*, 5170, 33.
- Alsing, R. (2008). *Genetic programming: evolution of Mona Lisa*. <http://rogeralsing.com/2008/12/07/genetic-programming-evolution-of-mona-lisa/>. Accessed 7/21/2011.
- Arnheim, R. (1974). *Art and visual perception: a psychology of the creative eye* (new, expanded and revised ed.) Berkeley: University of California Press.
- Atiyeh, B., & Hayek, S. (2008). Numeric expression of aesthetics and beauty. *Aesthetic Plastic Surgery*, 32(2), 209–216.
- Axelsson, O. (2007). Individual differences in preferences to photographs. *Psychology of Aesthetics, Creativity, and the Arts*, 1(2), 61–72.
- Baluja, S., Pomerleau, D., & Jochem, T. (1994). Towards automated artificial evolution for computer-generated images. *Connection Science*, 6(1), 325–354.
- Bense, M. (1965). *Aesthetica; Einführung in die neue Aesthetik*. Baden-Baden: Agis-Verlag.
- Bentley, P., & Corne, D. (2002). An introduction to creative evolutionary systems. In P. Bentley & D. Corne (Eds.), *Creative evolutionary systems* (pp. 1–75). San Francisco/San Diego: Morgan Kaufmann/Academic Press.
- Berlyne, D. E. (1960). *Conflict, arousal, and curiosity*. New York: McGraw-Hill.
- Berlyne, D. E. (1971). *Aesthetics and psychobiology*. New York: Appleton-Century-Crofts.
- Birkhoff, G. D. (1933). *Aesthetic measure*. Cambridge: Harvard University Press.
- Boselie, F., & Leeuwenberg, E. (1985). Birkhoff revisited: beauty as a function of effect and means. *The American Journal of Psychology*, 98(1), 1–39.
- Carroll, N. (1999). *Philosophy of art: a contemporary introduction, Routledge contemporary introductions to philosophy*. London: Routledge.
- Casti, J. L. (1994). *Complexification: explaining a paradoxical world through the science of surprise* (1st ed.). New York: HarperCollins.
- Chaitin, G. J. (1966). On the length of programs for computing finite binary sequences. *Journal of the ACM*, 13(4), 547–569.
- Ciesielski, V. (2007). Evolution of animated photomosaics. In *Lecture notes in computer science* (vol. 4448, pp. 498–507).
- Collier, G. L. (2002). Why does music express only some emotions? A test of a philosophical theory. *Empirical Studies of the Arts*, 20(1), 21–31.
- Cupchik, G. C. (2007). A critical reflection on Arnheim's gestalt theory of aesthetics. *Psychology of Aesthetics, Creativity, and the Arts*, 1(1), 16–24.
- Datta, R., Joshi, D., Li, J., & Wang, J. Z. (2006). Studying aesthetics in photographic images using a computational approach. In *Proceedings: Vol. 3953. ECCV 2006* (Pt. 3, pp. 288–301).
- Datta, R., Li, J., & Wang, J. Z. (2007). Learning the consensus on visual quality for next-generation image management. In *Proceedings of the ACM multimedia conference* (pp. 533–536). New York: ACM.
- Davis, T., & Rebelo, P. (2007). Environments for sonic ecologies. In *Applications of evolutionary computing* (pp. 508–516). Berlin: Springer.
- De Prisco, R., & Zaccagnino, R. (2009). An evolutionary music composer algorithm for bass harmonization. In *Applications of evolutionary computing* (Vol. 5484, pp. 567–572). Berlin: Springer.
- Dorin, A. (2005). Enriching aesthetics with artificial life. In A. Adamatzky & M. Komosinski (Eds.), *Artificial life models in software* (pp. 415–431). London: Springer. Chap. 14.
- Draves, S. (2005). The electric sheep screen-saver: A case study in aesthetic evolution. In *Lecture notes in computer science: Vol. 3449. Evo workshops* (pp. 458–467).
- Dutton, D. (2009). *The art instinct: beauty, pleasure, and human evolution* (1st U.S. ed.). New York: Bloomsbury Press.
- Elzenga, R. N., & Pontecorvo, M. S. (1999). *Arties: meta-design as evolving colonies of artistic agents*. Generative Design Lab.

- De Felice, F., & Fabio Abbattista, F. S. (2002). Genorchestra: an interactive evolutionary agent for musical composition. In C. Soddu (Ed.), *International conference on generative art* (Vol. 5). Generative Design Lab, Milan Polytechnic.
- Feldman, D. P., & Crutchfield, J. (1998). *A survey of complexity measures*. Santa Fe Institute.
- Ficici, S., & Pollack, J. (1998). Challenges in co-evolutionary learning; arms-race dynamics, open-endedness, and mediocre stable states. In C. Adami (Ed.), *Artificial life VI: proceedings of the sixth international conference on artificial life* (pp. 238–247). Cambridge: MIT Press.
- Fogel, L. J. (1999). *Intelligence through simulated evolution: forty years of evolutionary programming*. Wiley series on intelligent systems. New York: Wiley.
- Fornari, J. (2007). Creating soundscapes using evolutionary spatial control. In *Lecture notes in computer science* (Vol. 4448, pp. 517–526).
- Galanter, P. (2010). The problem with evolutionary art is. In C. DiChio, A. Brabazon, G. A. DiCaro, M. Ebner, M. Farooq, A. Fink, J. Grahl, G. Greenfield, P. Machado, M. O'Neill, E. Tarantino, & N. Urquhart (Eds.), *Lecture notes in computer science: Vol. 6025. Applications of evolutionary computation, pt. II, proceedings* (pp. 321–330). Berlin: Springer.
- Gartland-Jones, A. (2002). Can a genetic algorithm think like a composer? In C. Soddu (Ed.), *International conference on generative art* (Vol. 5). Generative Design Lab, Milan Polytechnic.
- Gedeon, T. (2008). Neural network for modeling esthetic selection. In *Lecture notes in computer science* (Vol. 4985(2), pp. 666–674).
- Gell-Mann, M. (1995). What is complexity? *Complexity*, 1(1), 16–19.
- Glente, K., Torresen, J., & Yasunaga, M. (2007). An online EHW pattern recognition system applied to face image recognition. In *Applications of evolutionary computing* (pp. 271–280). Berlin: Springer.
- Greenfield, G. R. (2003). Evolving aesthetic images using multiobjective optimization. In *CEC: 2003 congress on evolutionary computation* (pp. 1903–1909).
- Greenfield, G. (2005a). Evolutionary methods for ant colony paintings. In *Lecture notes in computer science: Vol. 3449. Evo workshops* (pp. 478–487).
- Greenfield, G. (2005b). On the origins of the term computational aesthetics. In *Computational aesthetics 2005: Eurographics workshop on computational aesthetics in graphics, visualization and imaging*, Girona, Spain, 18–20 May, 2005. Eurographics.
- Greenfield, G. (2008a). Evolved diffusion limited aggregation compositions. In *Applications of evolutionary computing* (pp. 402–411). New York: Springer.
- Greenfield, G. R. (2004). The void series—generative art using regulatory genes. In C. Soddu (Ed.), *International conference on generative art* (Vol. 7). Generative Design Lab, Milan Polytechnic.
- Greenfield, G. R. (2008b). Co-evolutionary methods in evolutionary art. In J. Romero & P. Machado (Eds.), *Natural computing series. The art of artificial evolution* (pp. 357–380). Berlin: Springer.
- Hawkins, J., & Blakeslee, S. (2004). *On intelligence* (1st ed.). New York: Times Books.
- Hazan, A., Ramirez, R., Maestre, E., Perez, A., & Pertusa, A. (2006). Modelling expressive performance: a regression tree approach based on strongly typed genetic programming. In *Applications of evolutionary computing* (pp. 676–687). Berlin: Springer.
- Helbing, D., & Molnar, P. (1995). Social force model for pedestrian dynamics. *Physical Review*, E(51), 4282–4286.
- Helbing, D., & Molnar, P. (1997). Self-organization phenomena in pedestrian crowds. In F. Schweitzer (Ed.), *Self-organization of complex structures: from individual to collective dynamics* (pp. 569–577). London: Gordon and Breach.
- Hoenig, F. (2005). Defining computational aesthetics. In L. Neumann, M. Sbert & B. Gooch (Eds.), *Computational aesthetics in graphics, visualization and imaging*, Girona, Spain.
- Holger, H. (1997). Why a special issue on the golden section hypothesis? An introduction. *Empirical Studies of the Arts*, 15.
- Hönn, M., & Göz, G. (2007). The ideal of facial beauty: a review. *Journal of Orofacial Orthopedics/Fortschritte der Kieferorthopdie*, 68(1), 6–16.

- Hornby, G. S., & Pollack, J. B. (2001). The advantages of generative grammatical encodings for physical design. In *Proceedings of the 2001 congress on evolutionary computation* (Vol. 601, pp. 600–607).
- Jaskowski, W. (2007). Learning and recognition of hand-drawn shapes using generative genetic programming. In *Lecture notes in computer science* (Vol. 4448, pp. 281–290).
- Khalifa, Y., & Foster, R. (2006). A two-stage autonomous evolutionary music composer. In *Lecture notes in computer science: Vol. 3907. Evo workshops* (pp. 717–721).
- Kolmogorov, A. N. (1965). Three approaches to the quantitative definition of information. *Problems in Information Transmission*, 1, 1–7.
- Komar, V., Melamid, A., & Wypijewski, J. (1997). *Painting by numbers: Komar and Melamid's scientific guide to art* (1st ed.). New York: Farrar Straus Giroux.
- Konečni, V. J. (1978). Daniel E. Berlyne: 1924–1976. *The American Journal of Psychology*, 91(1), 133–137.
- Koob, A. (2009). The root of thought: what do glial cells do? <http://www.scientificamerican.com/article.cfm?id=the-root-of-thought-what>. Accessed 11/29/09.
- Koza, J. R., Bennett, F. H. I., Andre, D., & Keane, M. A. (2002). Genetic programming: biologically inspired computation that exhibits creativity in producing human-competitive results. In P. Bentley & D. Corne (Eds.), *Creative evolutionary systems* (pp. 275–298). San Francisco/San Diego: Morgan Kaufmann/Academic Press.
- Kozbelt, A. (2006). Dynamic evaluation of Matisse's 1935 large reclining nude. *Empirical Studies of the Arts*, 24(2), 119–137.
- Law, E., & Phon-Amnuaisuk, S. (2008). Towards music fitness evaluation with the hierarchical SOM. In *Applications of evolutionary computing* (pp. 443–452). Berlin: Springer.
- Li, Y.-F., & Zhang, X.-R. (2004). Quantitative and rational research for the sense quantum—research of the order factors for color harmony aesthetic. *Journal of Shanghai University (English Edition)*, 8(2), 203–207.
- Livio, M. (2003). *The golden ratio: the story of phi, the world's most astonishing number* (1st ed.). New York: Broadway Books.
- Machado, P. (1998) Computing aesthetics. In *Lecture notes in artificial intelligence: Vol. 1515*.
- Machado, P., & Cardoso, A. (2002). All the truth about NEvAr. *Applied Intelligence*, 16(2), 101–118.
- Machado, P., & Cardoso, A. (2003). *NEvAr system overview*. Generative design lab, Milan Polytechnic.
- Machado, P., Romero, J., Cardoso, A., & Santos, A. (2005). Partially interactive evolutionary artists. *New Generation Computing*, 23(2), 143–155.
- Machado, P., Romero, J., & Manaris, B. (2008). Experiments in computational aesthetics—an iterative approach to stylistic change in evolutionary art. In J. Romero & P. Machado (Eds.), *The art of artificial evolution: a handbook on evolutionary art and music* (pp. 311–332). Berlin: Springer.
- Machado, P., Romero, J., Santos, A., Cardoso, A., & Pazos, A. (2007). On the development of evolutionary artificial artists. *Computers and Graphics*, 31(6), 818–826.
- Machado, P., Romero, J., Santos, M. L., Cardoso, A., & Manaris, B. (2004). Adaptive critics for evolutionary artists. In *Lecture notes in computer science. Applications of evolutionary computing* (pp. 437–446). Berlin: Springer.
- Machwe, A. T. (2007). Towards an interactive, generative design system: integrating a 'build and evolve' approach with machine learning for complex freeform design. In *Lecture notes in computer science* (Vol. 4448, pp. 449–458).
- Magnus, C. (2006). Evolutionary musique concrete. In F. Rothlauf & J. Branke (Eds.), *Applications of evolutionary computing, EvoWorkshops 2006* (pp. 688–695). Berlin: Springer.
- Manaris, B., Machado, P., McCauley, C., Romero, J., & Krehbiel, D. (2005). Developing fitness functions for pleasant music: Zipf's law and interactive evolution systems. In *Lecture notes in computer science: Vol. 3449. Evo workshops* (pp. 498–507).
- Manaris, B., Vaughan, D., Wagner, C., Romero, J., & Davis, R. B. (2003). Evolutionary music and the Zipf-Mandelbrot law: developing fitness functions for pleasant music. *Applications of*

Evolutionary Computing, 2611, 522–534.

- Martindale, C. (1981). *Cognition and consciousness. The Dorsey series in psychology*. Homewood: Dorsey Press.
- Martindale, C. (1984). The pleasures of thought: a theory of cognitive hedonics. *Journal of Mind and Behavior*, 5(1), 49–80.
- Martindale, C. (1988a). Cognition, psychobiology, and aesthetics. In F. H. Farley & R. W. Neperud (Eds.), *The foundations of aesthetics, art, and art education* (pp. 7–42). New York: Praeger Publishers.
- Martindale, C. (1988b). Relationship of preference judgements to typicality, novelty, and mere exposure. *Empirical Studies of the Arts*, 6(1), 79–96.
- Martindale, C. (1991). *Cognitive psychology: a neural-network approach*. Pacific Grove: Brooks/Cole Publishing Company.
- Martindale, C. (2007). A neural-network theory of beauty. In C. Martindale, P. Locher & V. Petrov (Eds.), *Evolutionary and neurocognitive approaches to aesthetics, creativity, and the arts* (pp. 181–194). Amityville: Baywood.
- Martindale, C., Moore, K., & Anderson, K. (2005). The effect of extraneous stimulation on aesthetic preference. *Empirical Studies of the Arts*, 23(2), 83–91.
- Martindale, C., Moore, K., & Borkum, J. (1990). Aesthetic preference: anomalous findings for Berlyne's psychobiological theory. *The American Journal of Psychology*, 103(1), 53–80.
- Maxwell, J. B., Pasquier, P., & Eigenfeldt, A. (2009). Hierarchical sequential memory for music: a cognitive model. In *International society for music information retrieval*.
- McCormack, J. (2005) Open problems in evolutionary music and art. In *Lecture notes in computer science: Vol. 3449. Evo workshops* (pp. 428–436).
- McCormack, J. (2008). Facing the future: evolutionary possibilities for human-machine creativity. In J. Romero & P. Machado (Eds.), *The art of artificial evolution: a handbook on evolutionary art and music* (pp. 417–451). Berlin: Springer.
- McCormack, J., & Bown, O. (2009) Life's what you make: Niche construction and evolutionary art. In *Lecture notes in computer science: Vol. 5484. Evo workshops* (pp. 528–537).
- McDermott, J., Griffith, N. J. L., & O'Neill, M. (2005). Toward user-directed evolution of sound synthesis parameters. In *Lecture notes in computer science: Vol. 3449. Evo workshops* (pp. 517–526).
- Minsky, M. L., & Papert, S. (1969). *Perceptrons; an introduction to computational geometry*. Cambridge: MIT Press.
- Mitchell, T. J., & Pipe, A. G. (2005). Convergence synthesis of dynamic frequency modulation tones using an evolution strategy. In *Applications on evolutionary computing* (pp. 533–538). Berlin: Springer.
- Moles, A. A. (1966). *Information theory and esthetic perception*. Urbana: University of Illinois Press.
- Monmarché, N., Aupetit, S., Bordeau, V., Slimane, M., & Venturini, G. (2003). Interactive evolution of ant paintings. In B. McKay et al. (Eds.), *Congress on evolutionary computation* (Vol. 2, pp. 1376–1383). New York: IEEE Press.
- Mori, T., Endou, Y., & Nakayama, A. (1996). Fractal analysis and aesthetic evaluation of geometrically overlapping patterns. *Textile Research Journal*, 66(9), 581–586.
- Neufeld, C., Ross, B. J., & Ralph, W. (2008). The evolution of artistic filters. In J. Romero & P. Machado (Eds.), *The art of artificial evolution: a handbook on evolutionary art and music* (pp. 335–356). Berlin: Springer.
- North, A. C., & Hargreaves, D. J. (2000). Collative variables versus prototypically. *Empirical Studies of the Arts*, 18(1), 13–17.
- Numenta (2008). Advanced nupic programming. http://www.numenta.com/for-developers/software/pdf/nupic_prog_guide.pdf. Accessed 16/04/10.
- Oelmann, H., & Laeng, B. (2009). The emotional meaning of harmonic intervals. *Cognitive Processing*, 10(2), 113–131.
- Parker, S., Bascom, J., Rabinovitz, B., & Zellner, D. (2008). Positive and negative hedonic contrast with musical stimuli. *Psychology of Aesthetics, Creativity, and the Arts*, 2(3), 171–174.

- Peitgen, H.-O., Jürgens, H., & Saupe, D. (1992). *Chaos and fractals: new frontiers of science*. New York: Springer.
- Phon-Amnuaisuk, S. (2007). Evolving music generation with SOM-fitness genetic programming. In *Lecture notes in computer science* (Vol. 4448, pp. 557–566).
- Pinker, S. (1994). *The language instinct* (1st ed.). New York: Morrow.
- Poon, J., & Maher, M. L. (1997). Co-evolution and emergence in design. *Artificial Intelligence in Engineering*, 11(3), 319–327.
- Reddin, J., McDermott, J., & O'Neill, M. (2009). Elevated pitch: automated grammatical evolution of short compositions. In *Lecture notes in computer science: Vol. 5484. EvoWorkshops 2009* (pp. 579–584).
- Resnick, M. (1994). *Complex adaptive systems. Turtles, termites, and traffic jams: explorations in massively parallel microworlds*. Cambridge: MIT Press.
- Reynolds, C. (1987). Flocks, herds, and schools: a distributed behavioural model. *Computer Graphics*, 21(4), 25–34.
- Romero, J., Machado, P., & Santos, M. L. (2003). *Artificial music critics*. Generative Design Lab, Milan Polytechnic.
- Rosenblatt, F. (1962). *Principles of neurodynamics; perceptrons and the theory of brain mechanisms*. Washington: Spartan Books.
- Ross, A. (1995). Poll stars. *ArtForum*, 33(5), 72–77.
- Ross, B. J., & Zhu, H. (2004). Procedural texture evolution using multi-objective optimization. *New Generation Computing*, 22(3), 271–293.
- Saunders, R. (2002). *Curious design agents and artificial creativity*. PhD thesis, University of Sydney.
- Saunders, R., & Gero, J. S. (2004). Curious agents and situated design evaluations. *AI Edam-Artificial Intelligence for Engineering Design Analysis and Manufacturing*, 18(2), 153–161.
- Scha, R., & Bod, R. (1993). Computationale esthetica. *Informatie en Informatiebeleid*, 11(1), 54–63.
- Schimmel, K., & Forster, J. (2008). How temporal distance changes novices' attitudes towards unconventional arts. *Psychology of Aesthetics, Creativity, and the Arts*, 2(1), 53–60.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(3), 379–423.
- Sims, K. (1991). Artificial evolution for computer-graphics. *Siggraph '91 Proceedings* 25, 319–328.
- Sims, K. (1994). Evolving virtual creatures. *Siggraph '94 Proceedings*, 28, 15–22.
- Sims, K. (1997). Galapagos interactive exhibit. <http://www.karlsims.com/galapagos/index.html>. Accessed 11/16/2010.
- Skov, M., & Vartanian, O. (2009a). Introduction—what is neuroaesthetics? In M. Skov & O. Vartanian (Eds.), *Neuroaesthetics—foundations and frontiers in aesthetics* (pp. iv, 302 p.). Amityville: Baywood.
- Skov, M., & Vartanian, O. (2009b). *Neuroaesthetics, foundations and frontiers in aesthetics*. Amityville: Baywood.
- Solomonoff, R. J. (1964). A formal theory of inductive inference, part I and part II. *Information and Control*, 7, 1–22. 224–254.
- Standage, T. (2002). *The mechanical turk: the true story of the chess-playing machine that fooled the world*. London: Allen Lane.
- Stauderk, T. (1999). *On Birkhoff's aesthetic measure of vases* (Vol. 2009). Faculty of Informatics, Masaryk University.
- Stewart, M. (2008). *Launching the imagination: a comprehensive guide to basic design* (3rd ed.). Boston: McGraw-Hill Higher Education.
- Sullivan, L. H. (1896). The tall office building artistically considered. *Lippincott's Magazine*, 57, 403–409.
- Takagi, H. (2001). Interactive evolutionary computation: fusion of the capabilities of EC optimization and human evaluation. *Proceedings of the IEEE*, 89(9), 1275–1296.