Musings about beauty

Walter Kintsch

University of Colorado

In this essay I explore how cognitive science could illuminate the concept of beauty. Two results from the extensive literature on aesthetics guide my discussion. As the term "beauty" is overextended in general usage, I choose as my starting point the notion of "perfect form." Aesthetic theorists are in reasonable agreement about the criteria for perfect form. What do these criteria imply for mental representations that are experienced as beautiful? Complexity theory can be used to specify constraints on mental representations abstractly formulated as vectors in a high-dimensional space. A central feature of the proposed model is that perfect form depends both on features of the objects or events perceived and on the nature of the encoding strategies or model of the observer. A simple example illustrates the proposed calculations. A number of interesting implications that arise as a consequence of reformulating beauty in this way are noted.

»Können Sie mir sagen, was Schönheit sei? « rief er aus.

»Vielleicht nicht!« versetzte ich, »aber ich kann es Ihnen zeigen.«

»…….daß die Schönheit etwas Unbegreifliches….sei.
Was man nicht begreifen kann, das ist nicht,
was man mit Worten nicht klarmachen kann, das ist Unsinn.«[1]

Johann Wolfgang Goethe
Der Sammler und die Seinigen

Aesthetics has been a human concern throughout history. Cognitive science is a relatively new development and its implications for a theory of aesthetics have been largely unexplored. What would a cognitive theory of aesthetics look like? Can cognitive science provide a framework for the analysis of beauty, does our knowledge about cognition provide insights into the what and why of beauty? Or must it remain inconceivable?

Almost everything can be beautiful to someone, at some time. Kant noted that judgments of beauty are sensory, emotional, and intellectual all at once (Kant, 1781). He is surely right. Nevertheless, that is an unacceptable starting point for a theory of aesthetics, because it leads inexorably to this quote from Wikipedia: "Thus aesthetic judgments might be seen to be based on the senses, emotions, intellectual opinions, will, desires, culture, preferences, values, subconscious behavior, conscious decision, training, instinct, sociological institutions, or some

---

[1] »Can you tell me what is beauty?« he exclaimed. »Perhaps not!« I replied, »but I can show it to you. ……beauty is something inconceivable….what one cannot conceive of is nothing, what one cannot explain with words is nonsense.«

complex combination of these, depending on exactly which theory one employs". In other words, beauty is hopelessly complex. If one accepts this conclusion, a theory of aesthetics is impossible. The meaning of the word "beautiful" has become so over-extended in everyday discourse that it is no longer useful as a theoretic construct. Therefore, the scope of a theory of aesthetics must be restricted somehow.

Schopenhauer did so by drawing a strict line, perhaps an illusory one but a useful one, between the world of the intellect and the world of the will (Schopenhauer, 1818). Aesthetics is located in the world of the intellect, far removed from the world of the will. Beauty is the contemplation of *perfect form* that has nothing to do with utility or personal and societal values. The contemplation of perfect form is the purest activity in the world of the intellect. Thus, following Schopenhauer, one might suggest that a theory of aesthetics should be concerned with explicating the notion of perfect form. Of course, such a theory by itself could not provide a complete account of aesthetic judgments, which are influenced by just about everything and, as I have argued above, represent a too difficult and messy problem for theoretical analysis. Thus, my goal here will be to clarify the notion of perfect form. But as I shall show below, this will also help us to understand the role of observer related factors in the perception of beauty, and thus bring us closer to a fuller understanding of the features of mental representations related to the experience of beauty. These then would constitute the universal basis that Kant argues underlies judgments of beauty – including the beauty of a rose in the garden, a towering mountain, a piano sonata, a tennis serve, and of a mathematical proof. What is it that can make these objects and events beautiful?

Before continuing, a remark on art and beauty seems in order. Much of art, especially modern art, strenuously avoids being beautiful. To quote a recent newspaper article:

> *Is beauty dead? The answer that springs from much of contemporary art is an unapologetic "yes." … Today, beauty is no longer about what's pretty, symmetrical, or harmonious. It's about what stirs the viewer to grapple with the world as it really is.*[2]

Emotional and intellectual impact are the criteria by which modern art is judged, rather than beauty. The article goes on to argue for a new conception of beauty, one that allows us to appreciate the "beauty of ugliness." Instead, I want to retain the traditional meaning of beauty, that is, to reserve the concept of beauty for "what is pretty, symmetrical, and harmonious." I shall not be concerned with the questions what the proper role of beauty is or should be in art.

Mental Representations

Beauty is in the eye of the beholder. Another way of saying that is that beauty is not a property of objects but of their mental representations created by the mind/brain. To experience an object as beautiful, it must possess certain affordances and it must be appreciated properly by someone. The objective and the subjective must mesh in just the right way to yield the experience of beauty. Only some objects afford mental representations that some of us experience as beautiful. What are the properties of an object so that it can be perceived as beautiful? What allows one person to experience beauty where another remains unmoved?

---

[2] C. Strickland, *Christian Science Monitor*, 12-20-2007

The mind/brain represents objects as activity patterns in an (intricately structured) cloud of neurons, or, more abstractly, as vectors in a high-dimensional space. Mental processes are manipulations of these representations, namely, computation. Thus, an object D is represented by a vector in n-dimensions, $D = \{d_1, d_2, d_3,.....d_n\}$. One could think of d_i as the activation value of a specific neuron, in which case the dimensionality of D would be extremely large; millions of neurons may be involved in mental representations. However, since each neuron is attuned to only a few aspects of the input it receives (neurons perform dimension reduction), the actual dimensionality of mental representations is much lower. It can be thought of as the set of independent quantities needed to specify D. To describe D, we can either enumerate all its elements or specify a computer program that generates the set of elements.

Specifying D via the computer program that generates it makes it possible to measure its complexity – specifically, its *Kolmogorov complexity*. Kolmogorov (1965) defined complexity as the shortest code in a universal programming language. He showed that the length of the shortest description of an object x is invariant up to a constant between different universal languages, which makes Kolmogorov complexity a very useful construct. The Kolmogorov complexity of an object x is denoted as K(x). The conditional Kolmogorov complexity K(y|x) is the length of the shortest program that yields y, given x as an input.

Mental representations are not static: they can suffer interference, they may decay in time, and they are subject to consolidation as well as active, controlled re-coding under conscious control. Thus, an object such as a painting may have an embodied, perceptual representation, which in itself may be complex, consisting of a perceptual record

combined with later experiences and memory images. Abstractions of this embodied representation may be created, often involving verbal and symbolic recoding, compressing the original vector representation and thereby reducing its information content. In reading a text, we distinguish the mental representation of the text itself – the textbase – and its knowledge and interest-based elaboration – the situation model (Kintsch, 1998). This elaboration can be automatic, as when experts fill in the gaps and links in a text within their domain of expertise, but it also is often conscious and indeed effortful, as when we labor to understand a difficult text in an unfamiliar domain. This labor is not confined to the initial reading but may continue as we ruminate about the implications of what we have read. Thus, memory traces are dynamic, with important consequences for the perception of beauty.

It is important, especially when we are concerned with the subjective experience of beauty, to emphasize that the processes that contribute to the vector D comprise the total human response to a beautiful object – sensory, conceptual, emotional. Imagine standing on a mountain top, contemplating the beautiful scenery below. The mental representation of that event is not just visual, but emotional (the euphoria of having reached the summit), and embodied (the fatigue of the ascent). Cognition does not occur in a disembodied mind, but in an acting, feeling, perceiving person. The awe with which we respond to a certain piece of music is as much a part of our mental representation as the pattern of sounds we perceive. The happiness we experience looking at some picture becomes an integral part of its mental representation. Emotions and bodily events are represented in our brains as patterns of neural activity, as are visual scenes or the semantic content of a story. The vector D is all brain activity, but this activity is the result of multimodal processes, ranging from the embodied to the symbolic.

Mental representations are physical symbol systems, realized as patterns of activation over sets of neurons in the brain. Describing a mental representation is to specify the probability distribution of activation values. Hence, for present purposes, Kolmogorov complexity can be given an information theoretic interpretation. Let $\mathbf{X}$ be a set of neurons $\{x\}$ with a probability distribution $p(x)$. The Kolmogorov complexity $K(\mathbf{X})$ can be estimated by its entropy

$$H(X) = E_x[I(x)] = - \sum_{x \varepsilon X} p(x) \log p(x).$$

Let $q(x)$ be a different probability distribution over $\mathbf{X}$. The conditional Kolmogorov complexity $K(\mathbf{X}_p|\mathbf{X}_q)$ can be equated with the Kullback–Leibler divergence – the average number of additional bits per datum necessary to transform $\mathbf{X}_p$ into $\mathbf{X}_q$:

$$Div_{KL}(p(\mathbf{X})||q(\mathbf{X}) = \sum_{x \varepsilon \mathbf{X}} p(x) \log(p(x)/q(x)) \qquad (1)$$

The vector D is the result of perceptual and conceptual processes in the human mind/brain. In the same situation, different brains will construct different representations. (Trivially, the representation of a red apple will be different for a color-blind person and one with normal color perception). Moreover, in the same situation, the same person at different times may construct different representations, depending on the person's momentary goals, motivation, available knowledge, and so on. (The tiger admired in the zoo and the one ready to attack in the jungle.) Hence D, for a given object, depends on the conditions of encoding, or the model M used for encoding.

Thus, I take mental representation to consist of an object (event, performance) with properties D (for data) and an encoding procedure M (for model, or method), {DM}.[3] What is there about D so that, for an appropriate M, it is perceived as beautiful? What sort of models M can result in encodings that we experience as beautiful?

All mental representations, beautiful or not, are subject to a general constraint that has been called the *simplicity principle* (Chater, 1999; Chater & Brown, 2008). The simplicity principle was first espoused by Mach (1886), who suggested that the cognitive system prefers patterns that provide simple descriptions of the data. A number of psychological observations supported Mach's conjecture, most notably the Gestalt theorists' Law of Prägnanz (Koffka, 1935; other examples are cited in Chater & Vitanyi, 2003). Simplicity is related to redundancy, as Simon pointed out: if no aspect of a complex system "can be inferred from any other then it is its own simplest description. We can exhibit it, but we cannot describe it by a simpler structure" (Simon, 1969, p. 221). In a loose way, the idea that the cognitive system prefers a simple, non-redundant code also relates to the observation that people tend to remember the gist of a text, rather than all its detail. The macropropositions in a text provide an efficient code from which some of the rest of the text can be reconstructed by a reader familiar with the domain.

In its modern form the simplicity principle states that the cognitive system tends to minimize the Kolmogorov complexity $K(x)$. The simplicity principle is closely related to a probabilistic Bayesian analysis. There are two results that are particularly useful for present purposes (more detail

---

[3] See also Chater & Brown (2008) and Schmidhuber (1997)

and proofs are given by Vitanyi and Li, 2000, and Chater and Brown, 2008).

First, a probabilistic analysis suggests that the cognitive system favors an M such that P(M|D) is maximized. Thus, it can be shown that there exists a close relationship between complexity and probability

max P(M|D) = min {K(M) + K(D|M)}.

The model M that has the highest a priori probability is the model that has the shortest code for the data. That code is a two-part description, involving a specification of the model M and an encoding of the data, given M.

Secondly, complexity theory can be extended to the situation where data accumulate over the course of time. Predictions of the next item based on simplicity and predictions based on the item's probability are equivalent.

The relevance of Kolmogorov complexity for the analysis of beauty was first recognized by Schmidhuber (1997) in his discussion of low-complexity art. He notes that the observer's model (Schmidhuber uses the term coding algorithm) embodies the observer's subjectivity, and argues that, given the model, the most beautiful drawing among a set of drawings classified as comparable by a given subjective observer is the one for which the information to compute the model from the data is minimized.

The Kolmogorov complexity theory is defined for universal programming languages. The human mind is no Turing machine, however. In terms of the Chomsky hierarchy of formal languages, the mind is at the level of a context sensitive language (Miller & Chomsky,

1964). That means that we cannot compute Kolmogorov complexity unless we know just what the constraints are that are imposed by the human cognitive system. In general, we do not know what form mental representations take, what sort of coding mechanisms are possible. The problem, then, is that the Kolmogorov complexity and the complexity that humans experience do not always correspond. For instance, the transformation 1 5 7 0 4 8 2 8 $\rightarrow$ 3 1 4 0 9 6 5 6 is exceedingly simple, but it may not be immediately obvious. The pattern in TTHTTTHHTH… appears complex to us, but, as the parity of the digits of π, it is trivial to generate for a computer (Griffiths & Tenenbaum, 2003). Such problems limit the applicability of complexity theory, but as a general framework it can still be useful, and, as demonstrated below, when we can make precise assumptions about the form of mental representations, complexity computations may be computationally tractable.

To use complexity theory as a basis for a cognitive theory we must, therefore, identify the constraints imposed by the human cognitive system with respect to the perception of beauty, or perfect form. Psychology today is far from being able to do so in a comprehensive and definitive manner. Nevertheless, there are some results that permit us to formulate a provisional answer. I have already mentioned a general constraint that all mental representations are subject to: the tendency to prefer the simplest possible code for the data. This suggests the hypothesis that beauty is related to the notion of a good gestalt: if an object can be coded so as to result in a good gestalt, we may perceive it as beautiful. Gestalt psychologists (e.g. Koffka, 1935) have specified the notion of "good gestalt" in some detail. Basically, it amounts to a whole whose various parts or aspects fit together in a harmonious way. Below, I shall try to specify further the notion of harmony in terms of complexity theory.

Harmony, however, is not enough. There are two more classic results from psychology that need to be taken into account. One says that we need variety, the other that there can be too much variety. Any stimulus that is presented for a long time or repeatedly will evoke a progressively weaker neural response. Importantly, this is not only true for perception, but also for emotion: affective habituation is very much like sensory adaptation, as Titchner (1908) had already noted. Thus, adaptation is a basic property of the cognitive system that must be counteracted. To be perceived as perfect form, a stimulus must afford the possibility for varied and dynamic encodings. A too simple object quickly becomes boring. The same activity, however blissful it may be initially, loses its ability to please when mindlessly repeated – variety, as they say, is the spice of life, and of beauty, too.

Variety is needed, but within limits. What the limits on complexity are varies greatly with the nature of the perceptual stimulus (through chunking) and the level of domain expertise of the perceiver. Human consciousness or working memory has a limited capacity. What can be apperceived is limited in complexity. Wundt (1887) and Miller (1956) have shown that the number of elementary stimuli that can be simultaneously perceived is limited, although that limit can be expanded through the process of chunking. Chunking is a process of grouping whereby the group functions as a unit. Furthermore, with training, people can form retrieval structures that make available large amounts of relevant information in working memory (the long-term working memory of Ericsson & Kintsch, 1995).

The relationship between a whole and its parts is fundamental to the notion of beauty discussed below. Concept such as part and aspect

are, of course, perfectly good commonsense terms, but we must take note of some important research results. Parts (nose and eyes are parts of the face) or aspects (the color, texture, or orientation of an object) can be defined physically as well as psychologically, but the two do not always correspond. This is shown most clearly by research that has used discrete static stimuli that vary along several dimensions. When subjects are supposed to pay attention to one dimension or aspect, they sometimes experience interference from variations in an irrelevant dimension. This happens, for instance, when subjects are asked to evaluate the brightness of a stimulus and an irrelevant aspect such as saturation is varied at the same time. Aspects that interfere with each other are called "integral dimension" whereas aspects that do not interfere with each other (the size of an object does not interfere with brightness judgments) are called "separable dimensions" (Garner, 1974). Integral dimensions are perceived holistically, that is, people cannot selectively attend to one dimension. Thus, a physical analysis of a stimulus and a psychological analysis do not always coincide: physically there are two aspects, but psychologically there is only one.

The parts of objects can be defined in different ways, for example by means of the geons of Biederman (1987), which are a set of basic shapes that can be used to represent a large number of objects. Objects, of course, are rarely perceived alone but as part of a visual scene. Visual scene analysis (e.g. Rensink, 2000) distinguishes between the gist of a scene (*a mountain landscape*), its parts (*a mountain peak, a lake*) and the relation among the parts (*the lake is below mountain and to the left*). The process is bidirectional and iterative. Gist is recognized by invoking one of the scene schemata stored in long-term memory on the basis of just a few recognized objects. The gist then participates in identifying further objects and modifying existing ones. The spatial layout is used to check

the emerging perceptual representation. In turn, the new data refine or change the gist representation, with the interplay of bottom-up and top-down processes continuing until a stable perceptual representation results.

Stimuli that evolve over time, as in listening to music or reading a book, contain cues that alert the perceiver to their part-whole structure. The chapter – and paragraph – organization of a book serves as an obvious guide to the reader, but there are many other cues to help the reader (Kintsch, 1998). The reader forms a macrostructure (the gist) on the basis of some key propositions, then uses it to construct a detailed representation of the text, including the intricate relationships among propositions. Text representations, both macro- and microstructure, evolve over time and, if the reader expends the required cognitive effort, are constantly being modified as new aspects are being noted.

Two important points have to be made about the recognition of whole objects and their parts. First, it is hierarchical. What is a part at one level of analysis may be the whole at a more fine grained level. This is the case for reading or listening to music as well as visual scene analysis. Second, less obviously, the process is dynamic. The limited capacity of working memory prohibits the simultaneous recognition of all parts of an object. Only a limited amount of information can be held in working memory when reading a book or listening to music. Similarly, it is in general impossible to perceive more than two or three parts or aspects of an object or scene at a time, and a shift of attention is required to focus on other parts, or for a more detailed analysis. Thus, a whole series of ephemeral structures must be built up in working memory before a detailed, stable percept is obtained. Similarly for reading: text comprehension and a stable text memory are generated via temporary

structures in working memory, roughly corresponding to sentence interpretations. This process can be automatic, as when experts are working in their domain of expertise, but in general requires a certain amount of cognitive effort. If that effort is not forthcoming, perception and comprehension may remain deficient.

Perfect form

Historically, the role of harmony and the prevention of habituation effects has been recognized by writers who sought to define beauty. Theorists of aesthetics have provided a number of feature lists that they claim characterize beauty. One such list by the painter William Hogarth I find particularly instructive. Hogarth (1753) lists six principles, as cited in Wikepedia:

(1) harmony, or fitness of the parts to the whole;

(2) variety;

(3) uniformity, regularity or symmetry, which must be in the service of fitness;

(4) simplicity or distinctness, which enables the viewer to appreciate the variety;

(5) intricacy, which allows for active processing, leading the eye "a wanton kind of chase";

(6) quantity or magnitude, which attracts attention and produces admiration and awe.

Of these (3) is really subordinate to (1) – fitness can be based on uniformity, regularity or symmetry – and (4) is subordinate to (2) – variety must be perceptible, simple and clear, rather than chaotic or random; variety must not interfere with fitness. (6) seems rather like a special case – significant as far as beauty is concerned, but not a necessary characteristic of harmony; I shall disregard it.

That leaves us with just three criteria: (a) fitness, (b) variety and (c) active energy. Harmony is required of the various aspects of a beautiful object or event. But by itself it does not guarantee beauty: the four sides of a square score high on harmony, but that does not make the square beautiful. The parts of a beautiful object, though harmonious, should also be varied and distinct. Finally, beauty requires an object that affords the active participation of an observer. A beautiful object is interesting, in the sense that an observer can find ever new ways to admire it; it "allows for active energies," Hogarth said.

For comparison, here is a modern list of the "path to beauty," by Ramachandran and Hirstein (1999):
1. Extremes: exaggeration along some dimension relative to mean
2. Structural goodness: good Gestalt
3. Unidimensionality: standing out in a single dimension
4. Contrast: strong features emerge prior to grouping
5. Challenge: requirement to engage in perceptual problem solving
6. Genericity: avoids statistically suspicious coincidences
7. Metaphoricity: hidden connection between different parts or aspects
8. Symmetry

This list maps quite well into Hogarth's list: Structural goodness and symmetry are what he called "fitness"; Challenge and metaphoricity belong to the "active energy" category; and the other criteria have to do with "variety." However, the Ramachandran and Hirstein list is not much of an improvement over Hogarth: symmetry (8) is clearly redundant, since it is part of a good Gestalt (2); metaphoricity (7) is equally redundant, since it is just a particular kind of challenge (5) – finding hidden

connections is perceptual problem solving. The other criteria specify how variety can be achieved: (1) and (2) imply that the most successful variety involves exaggeration along a single perceptual dimension; (4) and (6) are related to the simplicity and distinctiveness principle of Hogarth: (6) excludes randomness and arbitrariness, and (4) says there should be distinct parts (which then can be grouped into a whole).

Edelman (2008) has a chapter on "Beauty" where he discusses the Ramachandran and Hirstein principles (8), symmetry, and (1), exaggeration. He credits the English painter Reynolds (1723–1792), the first president of the Royal Academy, with an algorithm for distilling "perfect beauty": averaging over many instances that may not be very beautiful in themselves, which irons out the imperfections of particular instances and yields an idealized image. Data from face perception experiments are cited in which synthetic faces are judged beautiful the more common a face is (the larger the number of instances blended into it). The most common face – the average – is the most harmonious one, because, by definition, it retains the essential features of what makes a face, but has lost individual imperfections.[4] Thus, the Reynolds/Edelman algorithm can be viewed as an application of Hogarth's harmony criterion.

There are innumerable other lists in the literature which seek to specify the conditions for beauty, but to examine them would be a waste of time. Making more lists and citing more examples is not going to get us anywhere. We know, roughly, what stimuli are perceived as beautiful. We need to specify what that means for their mental representations – how beauty is computed by the mind. The Kolmogorov theory of complexity can help us arrive at a more principled analysis of beauty.

---

[4] However, for a critical review of this work, see Schmidhuber (1997)

As has already been mentioned, Schmidhuber (1997) pioneered this approach. He distinguished between an observer's subjective encoding scheme, the model, and the data, and argued that beauty in a certain type of art can be identified with low Kolmogorov complexity. Here, a more differentiated analysis will be presented, involving multiple criteria and framed within the part-whole distinction. The discussion will be motivated both by results from the literature on aesthetics and considerations about the nature of mental representations from cognitive science.

According to the literature reviewed above, beauty has to do with the relation of an object to its parts and the relations among its parts, as encoded by a human observer. Formally, we have represented the mental representation of an object by a vector D of n dimensions, $D = \{d_1, d_2, d_3, \ldots d_n\}$. An aspect A_i is similarly represented by a vector $A_i = \{a_1^i, a_2^i, a_3^i, \ldots a_{ni}^i\}$. If D has k aspects, its representation is the kth-rank tensor, $D = A_1 \otimes A_2 \otimes \ldots \otimes A_k$. For computational simplicity, this could be approximated in a variety of ways, for example, through circular convolution or simply by taking the average, $D = \Sigma_k A_i$. I assume that the parts (or aspects) of an object D are a property of the object D as it exists independently of human observers. What concerns us here is how these aspects are encoded by a human observer, that is, not the aspect A_i itself but how it is encoded by some model, $A_i M_j$. In general, an object may have aspects that a given observer does not encode at all, or that different observers (or the same observer at different times) encode in different ways, depending on the model M_j that is being used. Thus, as far as mental representations are concerned, there are two things to consider: a real-world object D with several aspects, and a human observer who uses a particular encoding procedure, the model M_j, to

encode the object and its aspects; in general, an object D can be interpreted with several alternative models.

The distinction between an object and its aspects allows us to state three specific hypotheses about the conditions for perfect beauty.

(1) <u>Harmony</u>. Let DM be a mental representation of the data D given the model M. Let D be composed of the aspects $A_1,...A_i,...A_k$. According to the simplicity principle, the cognitive system will select a model M that minimizes K(DM), and if the aspects A_i are perceived in isolation $K(A_i M_i)$ will also be minimized. Harmony requires that $M = M_i$ for all aspects of D. That is, the same model that is appropriate for D is also appropriate for its aspects A_i. Hence, the model M_1 that is selected by the cognitive system among all possible models M_i because it minimizes the Kolmogorov complexity of D is also the model that is selected for each of its aspects. In other words, harmony implies that D as well as its aspects {A} can be encoded with the same procedure, the model M_1, and that the parts should fit the whole – that the cost of deriving a part from the whole is small:

$$\Sigma_k K(A_i M_1 | DM_1) \text{ is small} \tag{2}$$

Another way to describe the harmony criterion is to make use of the relation between probability and complexity:

$$\max P(M|D) = \min \{K(M) + \Sigma_k K(A_i | M\}$$

That is, to minimize complexity, only a single model needs to be specified, not one for each aspect A_i. Edelman's criterion of familiarity yields a method for finding minima; thus, for a set of faces D_i,

$$\min_i K(D_i, M_i) = \text{average } K(D_i, M_j).$$

(2) <u>Variety</u>. While harmony is a necessary condition, it is not a sufficient condition for beauty, because very boring objects could be

generated in this way. A second requirement for a harmonious object to be perceived as beautiful is that its aspects are varied. While all aspects must fit the whole, the parts must be distinct from each other, which means that

$$K(A_iM \mid A_jM) > K(A_iM, A_jM \mid DM). \qquad (3)$$

In words, the similarity between two aspects is smaller than the similarity between each aspect and the whole. Thus, perfect form requires that the model that the cognitive system arrives at (the one that minimizes complexity) is the same for the whole and its aspects, and the aspects are distinct. That is, under the model M, it is more difficult to transform the aspects into each other than to transform each aspect into the whole.

The obvious application of these ideas is to a painting or a sculpture, which is fully present as a physical object at any moment in time. Yet it also applies to music and literature, which develop over time. A sentence or a musical phrase exists within the context of the whole work and must be interpreted taking into account what has come before as well as what still is to come. Indeed, the difference between the visual arts and writing and music is more apparent than real, as the next criterion for perfect form will make obvious: a visual scene is present physically all at once, but its reception is anything but instantaneous.

(3) <u>Compression</u>. The third requirement for perfect form requires viewing cognition as a dynamic system. Mental representations are constructions, and constructions can change over time. As time passes, the model that an observer uses to encode a datum D changes. Perfect form implies that the model changes in a particular way: it becomes ever better, in the sense that the complexity at time $t+1$ is less than at time t, that is, harmony – how well parts fit the whole, given a particular model – increases:

$$0 > \Sigma_k K_t(A_i M \mid DM) - \Sigma_{k'} K_{t+1}(A_i M \mid DM) \qquad (4)$$

A beautiful object allows for the compression of complexity; it affords the discovery of new harmonies.

To make this constructive activity possible, both D and M must fulfill certain conditions. D must have a sufficient number of potential aspects, so that it is possible to discover new ones even after prolonged study. Beautiful objects may have regularities, symmetries, fractal properties and the like that are not always discovered immediately. Objects that do not have this potential, even though they may appear beautiful at first have no lasting appeal. D thus must have the potential for the discovery of novel aspects and harmonies. But equally important, the observer must have the capability and desire to develop new, simpler, compressed models, to take advantage of the affordances the object offers. That ability comes with experience and training; to really appreciate beauty one needs to become a connoisseur. Connoisseurship is a kind of expertise. Like all expertise it requires an intimate, long-term familiarity with a particular domain. To the connoisseur, things that only elicit a cursory look from most people become fascinating objects of contemplation and reanalysis. New aspects are discovered; familiar ones are seen in new ways. What the connoisseur does is best described as a continual process of complexity reduction: contemplation reveals new harmonies, new aspects.

Schmidhuber (1991, 2006) has identified the compression of complexity with interest: objects are interesting when they afford compression with further study. Interestingness, in this sense as

complexity reduction[5], is seen as another factor in beauty.[6] When we can always perceive something new in an object, when objects afford the discovery of new harmonies, we appreciate them most. When is complexity compression possible? If an object is so complex and unfamiliar that it is impossible for an observer to discover regularities in it, there will be no complexity reduction. A random display (or an object whose structure is so complex that it appears random) will remain random, and quickly becomes boring. At the other extreme, a totally familiar object that has been exhausted does not afford further possibilities for comprehension reduction – it too is boring. Objects that are complex and unfamiliar enough afford complexity reduction. Thus, Eq. (4) implies the well known U-shaped relation between interest and familiarity (Wundt, 1887; Berlyne, 1960): the very familiar and the totally unfamiliar are not interesting; what is just familiar enough tends to be interesting.

The more one learns about something, the more one is able to appreciate its complexity. A good example is twelve-tone music (Schmidhuber, 1996). To audiences a hundred years ago, as well as neophytes today, it appears as random noise. Modern ears, however, have become used to it, are hearing something in it, can find it interesting and even beautiful – not as harmonious as Mozart, but far from the aggressive cacophony that it once appeared to be. Learning to see something as interesting and beautiful occurs not only with respect to art, but also nature. An experienced naturalist can make us see and appreciate things in nature, large and small, that we would have walked

---

[5] There are other sources of interest, too, besides compressibility: some things are intrinsically interesting – roughly, sex and violence.
[6] Schmidhuber distinguishes between beauty and interest; here, interest is considered a criterion for beauty. This seems primarily a question of terminological preference.

right past on our own. It was always there – the D is the same, but we learn to use a new model M that encodes D in a way that makes it interesting and, if the conditions of harmony and variety are also met, beautiful.

Complexity theory relates probability and complexity. Predictions based on a model that yields minimum complexity and predictions based on maximum probability are the same. Prediction plays an important role in the experience of beauty, not only with events that unfold in time, as in literature, music, dance, or ceremonies, but, as we have just argued, also with static objects such as a painting, a flower, or a landscape. Beauty in music implies a model that maximizes the probability of the next element as we listen, but also that can be compressed, so that when we listen the next time, a new, more sophisticated model will guide our perception. Why is it that we can read a great book many times and it becomes more interesting with each reading? Because it affords us the opportunity to fine-tune our model, to construct a novel interpretation every time. The book remains the same, but we – our model – change.

Perfect form, thus, is a construction of the mind, involving the active participation of an informed observer. Piaget described the cognitive development of children as an interplay between the processes as assimilation and accommodation (Piaget, 1955; Schmidhuber, 2009). Assimilation means that old schemas are used to encode new information; in accommodation, new encoding schemas are generated. Mental representations that we characterize as perfect form also involve assimilation in terms of an existing model, and accommodation through the evolution of a new model; however, not all forms of assimilation and accommodation result in the perception of beauty. The perception of beauty requires certain conditions to be satisfied – harmony and variety

for assimilation, and compression for accommodation. With respect to the emphasis on the active contribution of the observer, the perception of beauty resembles other cognitive processes that have been traditional objects of study within cognitive science. There are, for example, numerous studies of the generation effect (e.g. McNamara & Healy, 2000) that show better memory for information that has been generated rather than information passively absorbed. Similarly, the importance of actively constructing situation models in text comprehension and memory is widely attested (e.g., Kintsch, 2009). The central role of background knowledge (e.g., Kintsch, 1998) in the construction of situation models also has its parallel in the perception of beauty: the connoisseur with a lifetime of experience perceives things that escape the untrained eye or ear. Thus, the study of aesthetic experience appears to be more closely related to the mainstream of cognitive science than is apparent at first.

A simplified example of complexity computations

In general, we are unable to specify either the mental representation of an object D or of the model M, which makes actual computations of complexity impossible. However, it is possible to illustrate the computations involved within specific models. For instance, current models of discourse comprehension make specific assumptions about the mental representation of a text and the model used to encode it. Specifically, within the framework of the Topic Model (Griffiths, Steyvers, & Tenenbaum, 2007) a text (or document) is represented as a multinomial probability distribution over topics. Topics are semantic elements that are inferred from a linguistic corpus and are represented as probability distributions over the words of the corpus. Thus, a model M may be defined as a set of n topics $\{T_i\}$ that are used to construct the mental representation of a document D,

$$D = \Sigma_n p_i T_i \ ,$$

Where p_i is the probability of topic T_i.

Consider a story A that has aspects (e.g., paragraphs) A_1, A_2,..., A_n and a set of topics $\{T_i\}$. The aspects are the property of the data; the way they are encoded in terms of the available topics constitutes the model; the whole story as well as each aspect is to be represented with the same model, the *n* topics T. Harmony requires that it is indeed possible to encode the whole as well as the parts with the same model, that is, the cost of transforming the aspects A_1, A_2,..., A_n into the whole must be small (Eq. 2). Variety requires that the parts are all different, that is, that the cost of transforming one aspect into another is large (Eq.3). Dynamics (Eq. 4) requires that there be another set of topics $\{T^*\}$ that yields less complex encodings than $\{T\}$. Since we are comparing probability distributions, cost measures are simply the KL–divergence between them (Eq. 1).

A serious test of such a model would be a major undertaking, requiring a careful selection of representative texts. At this point, all I can offer is a greatly oversimplified and highly artificial example – intended merely to illustrate the computations involved.

Consider two texts A and B, each with three sections, or aspects. Our goal is to estimate how beautiful A and B are, taking into account only their content, not their stylistic and rhetorical properties. Suppose that both texts as well as all their sections can be encoded with a model that has only two topics, T_1 and T_2, with probabilities shown in Table 1a and 1b. That is, both texts A and B load approximately equally on T_1 and T_2, but A_1 loads heavily on T_1 while A_2 loads mostly on T_2, and so on. To compare probability distributions we use the Kullback–Leibler

Divergence, Eq. 1. Harmony requires that the sections of each text as well as the text as a whole can be encoded with the two-topic model and that the divergence between the codes for the sections and the code for the text as a whole is small. The codes in this case are the probability distributions shown in Table 1. Using Eq. 2, DIV $(A_1, A_2, A_3 || A) = .24$ and DIV $(B_1, B_2, B_3 || B) = .04$. Both of these numbers are small,[7] so the parts fit the whole harmoniously in both cases. However, the three sections of A are more varied than the three sections of B. Therefore, A but not B fulfills the variety criterion. Specifically, ave DIV $(A_i || A_j) = .93$, while ave DIV $(B_i || B_j) = .14$. That is, the sections of A are dissimilar, while those of B are not.

Text C, shown in Table 1c, is an example of disharmony: a different model is required to encode each of its sections: C_1 requires mostly topics T_1 and T_2, C_2 requires mostly T_3 and T_4, and C_3 requires mostly T_5 and T_6. For this text, DIV $(C_1, C_2, C_3 || C) = 1.20$, a higher value than was obtained for A and B. Hence C fails the harmony criterion.

Finally, we illustrate the dynamic aspects of complexity reduction with the example shown in Table 1d. What we have assumed here is that upon reflection, a new dimension is discovered for encoding Text A. This discovery somewhat changes the old topic assignments, as we illustrate with the example in Table 1d. The important point is that the additional dimension actually reduces the complexity of the code for A, DIV $(A_1^*, A_2^*, A_3^* || A^*) = .16$. Thus, the parts fit the whole even better than before as required by Eq. 4. Hence of the three "texts" in Table 1, only A fulfills all the criteria for perfect form. C scores too low on harmony; B lacks variety; but A is harmonious, varied, and as its code changes from A to A*,

---

[7] The KL-divergence has no upper bound.

dynamic. Needless to say that these examples are intended merely to show how the complexity computations proposed here might be realized.

Insert Table 1 here

To see how far the above calculations are from a processing model of aesthetic perception, it is useful to consider Solso's distinction between levels of representation in art (Solso, 2003): the surface information (e.g., visual properties of an artwork), the conceptual level (the concepts and ideas represented), and the interpretational level (including emotional responses), in analogy with the surface level–textbase–situation model representations in discourse comprehension (Kintsch, 1998). The calculations presented here concern purportedly necessary conditions at Solso's second level. One could also regard them as formalizations of Berlyne's notions about the role of symmetry and balance in aesthetic perception (Berlyne, 1974). They do not constitute a processing model, however. Millis and Larson (2008) have made a promising start in this respect, and the approach taken here might be compatible with their work. Level-1 effects have also been studied, for example rhyme and alliteration in poetry (Kintsch, 1994; Lea et al., 2009). Level-3 effects are easier to talk about than to model, but they too are open to experimental investigation, as shown, for example, by Millis (2001), who manipulated inferences about paintings by the kind of titles he provided.

The focus of the work discussed here has been on the nature of the mental representations that are experienced as beautiful and their dependence on both the features of a stimulus and the perceiver's model that encodes that stimulus. Alternatively, research on the perceptual fluency hypothesis (Reber, Schwarz, & Winkielman, 2004) emphasizes the subjective experience of encoding fluency and the resulting positive

affect as the source of beauty. It, too, sees beauty as a result of how people and objects relate. The two approaches make similar predictions about the characteristics of objects (e.g., symmetry, good gestalt) and perceivers (e.g., expertise), probably because perceptual fluency depends on low complexity and harmony, as defined here by Eq. (2). The variety condition (Eq. 3) that is proposed here as a criterion for perfect form is not accounted for by the perceptual fluency hypothesis, and neither is compression (Eq. 4), in part, because the experiments about perceptual fluency mostly employ rather simple stimuli and are concerned primarily with first impressions rather than the extended consideration of an object.

Perfect form and beauty

To summarize, I have argued that the essential characteristic of perfect form is harmony. Harmony is a property of mental representations. The mind constructs mental representations that have minimal complexity, that is, for a given datum D an encoding procedure or model M is selected that minimizes $\{K(M) + K(D|M)\}$. Harmony is a relation between a mental representation $\{DM\}$ and the mental representation of its aspects $\{A_1M,...,A_kM\}$: the whole D and its parts, the aspects, $A_1,...,A_k$, can all be encoded with the same model M in such a way that the aspects fit the whole, that is, the aspects can be derived from the whole: the cost of generating an aspect from the whole is low. Harmony, however, is a necessary but not a sufficient condition for perfect form. An object must have aspects that are varied and distinct, and it must be interesting. Distinctiveness can be defined in terms of conditional complexity: the complexity of going from the whole to the aspect, $K(A|D)$ is low, while the complexity of transforming one aspect into another $K(A_i|A_j)$ is high. Interestingness here means the possibility of

complexity reduction: contemplation of an object leads to the selection of a new model M* that yields a lower complexity encoding. In other words, an interesting object allows for active re-coding that improves predictability on the basis of newly detected harmonies and structures.

The framework explored here emphasizes the crucial role of the perceiver, his or her model M, in the experience of beauty. Indeed, it makes perceiving beauty a creative act. When it comes to the enjoyment of beauty in nature, it may be surmised that our native endowment allows every one to appreciate it, modulo some cultural biases, though here too exposure and experience play a significant role. When it comes to the perception of beauty in art, connoisseurship matters even more. To perceive beauty in art requires a well developed, sophisticated model that enables the mental construction of the artwork in the perceiver's mind, to recreate it mentally in its ideal form. It certainly matters what the orchestra plays, or how the pianist plays, but what matters even more is what we hear, the heavenly music constructed in the mind of the listener or player.

I started this essay with an argument to focus on perfect form instead of beauty. While understanding beauty is our goal, the concept of beauty is so overextended as to make an analysis extremely difficult. Perfect form, instead, turned out to be much more tractable: we could analyze perfect form in terms of some kind of model selection. Note that this move has allowed us to reconsider some of the factors that that made the concept of beauty so messy. A person's model depends on his or her experience, knowledge, social conventions, taste. All the factors that we had initially excluded in favor of pure perfect form now find a new place within a novel framework: the question now becomes how mental representations are constructed, a promising research question

for cognitive science, though as yet a poorly understood one. Beauty, thus, comes in by the back door.

How is perfect form to be achieved? How can we have both harmony and variety, and change as well? Surely, there is no single, general answer to these questions, but rule systems must play an important role. Constructing an object according to a system of rules is one way to achieve harmony: all of its aspects obey the same rules. Of course, the rule system has to be complex enough to allow for diversity and dynamic reinterpretations. An example might be the rules that govern the composition of a fugue: start with a theme of several bars; replicate it with varied onset times; then stretch it, compress it, reverse it, mirror it, start changing it in subtle but surprising ways – in the hands of a skilled composer such a rule system can generate beauty. Or take a more complex rule system, a twelve-tone series, and play the same sort of games with it to create endlessly varied structures. The danger here is that the rules generate such complex structures that they can no longer be perceived as harmonious. In the extreme, only the composer knows, our ears do not. Another example from music is tonal ambiguity. A given note may have a place in more than one structure, it may, for instance, belong to one key, but can also be heard as part of a new modulation. Similarly for visual ambiguity: the makers of village and nomadic rugs often play with figure-ground reversals, inviting the viewer to "a wanton kind of chase," as Hogarth put it. Harmony is achieved here by adhering to strict rules in the construction of these textiles: the weave, the patterns, the colors, the use that a textile is intended for, are all conventionally determined. But these conventions leave the weavers enough freedom to generate a pleasing variety of objects. Indeed, the physical constraints imposed by the very materials of construction may contribute to harmony: given the weft-and warp foundation and a

decorative weft, there is only so much that is possible, which enforces coherence in the eventual product. Lack of physical constraints in an artwork can be problematic: in a wood or marble sculpture, the material limits what can be done and imparts a degree of unity; in a plastic sculpture or video installation everything is possible – harmony and hence beauty becomes harder to achieve. This is not to say that beauty is not possible in the absence of physical constraints, nor that such constraints guarantee success. It simply means that technical limitations are not necessarily the enemy of beauty, and complete freedom not necessarily fruitful.

<u>Truth and beauty</u>

> The true and the beautiful are akin. Truth is beheld by the intellect which is appeased by the most satisfying relations of the intelligible: beauty is beheld by the imagination which is appeased by the most satisfying relations of the sensible.
>
> James Joyce,
> *A Portrait of the Artist as a Young Man*

The concept of beauty, viewed in this way, is closely linked to that of truth. Beauty is a property of certain kinds of mental representations, and the same properties may play a role in truth. In particular, the fitness criterion may be as relevant for truth as it is for beauty. One condition for truth may be that a model M is true if it makes possible an efficient encoding of a datum D and all its aspects. Beauty may thus light the way to truth, as has indeed happened repeatedly in the history of physics. Freeman Dyson "points to several famous examples from the history of

physics when theories designed to be beautiful turned out to be true. The best-known examples are the Dirac wave-equation for the electron and the Einstein theory of General Relativity for gravity."[8] Thus, there may be deep relations between the concepts of beauty and truth that remain to be explored.

<u>Conclusions</u>

So what has been achieved with these musings about beauty? They are at a far too general and abstract level to be useful for a cognitive model. Nevertheless, they may be of some small value by focusing on the role of mental representation in aesthetic experience. Restating our problem in terms of complexity theory does not solve that problem, but it may nevertheless be illuminating. The simplicity principle makes some strong, if general, claims about the nature of mental representations. Within that framework, complexity theory allows us to formulate a coherent account of the notion of perfect form.

---

[8] New York Review of Books, 4-9-2009

References

Berlyne, D. E. (1960). *Conflict, arousal, and curiosity*. New York: McGraw-Hill.

Berlyne, D. E. (1974). *Studies in the new experimental aesthetics*. Washington D. C.: Hemisphere.

Biederman, I. (1987). Recognition by components: A theory of human image understanding. *Psychological Review*, *94*, 115-147.

Chater, N. (1999) The search for simplicity: A fundamental cognitive principle? *Quarterly Journal of Experimental Psychology*, *52A*, 273-302.

Chater, N. & Brown, G. D. A. (2008) From universal laws of cognition to specific cognitive models. *Cognitive Science*, 32, 36-67

Chater, N. & Vitanyi, P. (2003) Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Science*, *7*, 19-22.

Edelman, S. (2008). *Computing the mind*. Oxford: Oxford University Press.

Garner, W. R. (1962). *Uncertainty and structure as psychological concepts*. New York: Wiley.

Griffiths, T. L. & Tenenbaum, J. B. (2003). Probability, algorithmic complexity, and subjective randomness. Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society.

Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review, 114*, 211-244.

Hogarth, William (1753) *The Analysis of Beauty*. London. Paul Mellon Center for Studies in British Art, 1998.

Kant, I. (1781) *Kritik der reinen Vernunft*. English translation *The critique of pure reason*, Encyclopedia Britannica, 1952.

Kintsch, W. (1994). Kognitionspsychologische Modelle des Textverstehens: Literarische Texte. In K. Reusser & M. Reusser-Weyeneth (Eds.) *Verstehen: Psychologischer Process und didaktische Aufgabe*. Bern: Huber. Pp. 39–53.

Kintsch, W. (1998). *Comprehension: A paradigm for cognition*. New York: Cambridge University Press.

Kintsch, W. (2009). Learning and constructivism. In S. Tobias & T. M. Duffy (Eds.) *Constructivist Instruction: Success or failure?* Pp. 223–241. New York: Taylor & Francis.

Koffka, K. (1935) *Principles of Gestalt Psychology*. Harcourt, Brace & CO., NY.

Lea, R. B., Rapp, D. N., Elfenbein, A., Mitchel, A. D., & Romine, R. S. (2008). Sweet Silent Thought: Alliteration and Resonance in Poetry Comprehension. *Psychological Science, 19*, 709–716.

Mach, E. (1886) *Analyse der Empfindungen.* Jena: Fischer. English translation *The analysis of sensation*, Chicago: Open Court, 1897.

McNamara, D. S. & Healy, A. F. (2000). A procedural explanation of the generation effect for simple and difficult multiplication problems and answers. *Journal of Memory and Language, 43,* 652–679.

Miller, G. A., & Chomsky, N. (1963). Finitary models of language users. In R. D. Luce, R. R. Bush & E. Galanter (Eds.), *Handbook of mathematical psychology*. New York: Wiley.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits of our capacity for processing information. *Psychological Review, 63*, 81–97.

Millis, K., & Larson, M. (2008). Applying the Construction-Integration framework to aesthetic responses to representational artworks. *Discourse Process, 45*, 263–287.

Millis, K. K. (2001). Making meaning brings pleasure: The influence of titles on aesthetic experiences. *Emotion, 1*, 320–329.

Ramachandran, V. S., & Hirstein, W. (1999). The science of art: A neurological theory of aesthetic experience. *Journal of Consciousness Studies*, *6*, 15–51.

Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? *Personality and Social Psychology Review, 8*, 364–382.

Rensink R. A. (2000). Scene Perception. In A. E. Kazdin (Ed.), *Encyclopedia of Psychology*. Vol. 7. (pp. 151–155). New York: Oxford University Press.

Schopenhauer, A. (1818/19). *Die Welt als Wille und Vorstellung*. English translation *The world as will and representation*, Dover Editions 1966.

Schmidhuber, J. (1991). Curious model-building control systems. Proc. International Joint Conference on Neural Networks, Singapore. *IEEE*, *2*, 1458–1463

Schmidhuber, J. (1997). Low-complexity art. *Leonardo, Journal of the International Society for the Arts, Sciences, and Technology*. *30*, 97–103.

Schmidhuber, J. (2006). Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Connection Science*, *18*, 173–187.

Schmidhuber, J. (2009). Driven by compression processes: A simple principle explains essential aspects of subjective beauty, novelty, surprise, interestingness, attention, curiosity, creativity, art, science, music, jokes. *Journal of SICE*, *48*, 21–32.

Simon, H. A. (1969) *The sciences of the artificial*. Cambridge MA: MIT Press.

Solso, R. (2003). *The psychology of art and the evolution of the conscious brain.* Cambridge, MA: MIT Press.

Titchner, E. B. (1908). *Lectures on the elementary psychology of feeling*

and attention. New York: MacMillan.

Vitanyi, P. M. B. & Li, M. (2000). Minimum description length induction, Bayesianism, and Kolmogorov complexity. *IIIE Transactions on Information Theory*, *IT–46*, 446–464.

Wundt, W. (1887) *Grundzüge der physiologischen Psychologie*. Leipzig: Engelmann.

Table 1. Examples of hypothetical documents A, B, and C, and their representations as probability distributions over topics.

1a. Topic probabilities for text A and its three sections analyzed with a two-topic model.

	A_1	A_2	A_3	A
T_1	.8	.1	.5	.48
T_2	.2	.9	.5	.52

$$DIV(A_1,A_2,A_3 || A) = .24$$

1b. Topic probabilities for text B and its three sections analyzed with a two-topic model.

	B_1	B_2	B_3	B
T_1	.3	.6	.5	.48
T_2	.7	.4	.5	.52

$$DIV(B_1,B_2,B_3 || B) = .04$$

1c Topic probabilities for text C and its three sections analyzed with a six-topic model.

	C_1	C_2	C_3	C
T_1	.8	.01	.01	.27
T_2	.16	.01	.01	.06
T_3	.01	.06	.01	.03
T_4	.01	.9	.01	.30
T_5	.01	.01	.16	.06
T_6	.01	.01	.8	.27

$$\text{DIV}(C_1, C_2, C_3 \,||\, C) = 1.20$$

1d Topic probabilities for text A^* and its three sections analyzed with a three-topic model.

	A_1^*	A_2^*	A_3^*	A^*
T_1	.6	.1	.4	.34
T_2	.2	.7	.5	.48
T_3	.2	.2	.1	.18

$$\text{DIV}(A^*_1, A^*_2, A^*_3 \,||\, A^*) = .16$$