Beyond the Creative Species

Making Machines That Make Art and Music

By: Oliver Bown

Citation:

Beyond the Creative Species: Making Machines That Make Art and Music
By: Oliver Bown
DOI: 10.7551/mitpress/10913.001.0001
ISBN (electronic): 9780262361750
Publisher: The MIT Press
Published: 2021

The open access edition of this book was made possible by generous funding from University of New South Wales.

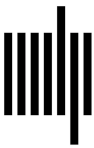

The MIT Press

4 Artistic Behavior

> It is self-evident that nothing concerning art is self-evident.
> —Adorno[1]

> A work of art is useless as a flower is useless. A flower blossoms for its own joy. We gain a moment of joy by looking at it. That is all that is to be said about our relations to flowers.
> —Oscar Wilde[2]

> Nothing is more necessary than the unnecessary.
> —*Life Is Beautiful*[3]

Art and Human Psychology

We shift our focus now specifically to artistic behavior. The last chapter's discussion of agency and function provides a framework for understanding creativity as an aspect of art practice, including what the underlying cultural and biological evolutionary foundations of art may be, how artistic creativity is distributed, and how function plays into this distribution. This ultimately leads to a better understanding of how creative art production and evaluation may be analyzed as sites of automation.

The term *art* is another bundle of concepts that will not be easily pinned down, and here it will be used in a particularly inclusive way to describe a myriad of activities including but not limited to visual art, music, decorative art, aspects of design, poetry, storytelling, fashion, and so on. Such a cluster can also be referred to as "cultural production." It is a given that bundling all of these activities together will gloss over the many critical

differences between these areas. While taking this liberty I will keep check of when this is being done and whether it is acceptable in each instance. The differences themselves will be of particular importance in this chapter, since a key concept is that each domain has distinct properties.

As we have experienced with creativity, the diverse and changeable concept of art eschews a strict categorical definition. One option is to take a family-resemblance approach. Each of us has a sense of what art is, based on a set of properties that may be more or less aligned with those concepts of other people. We may admit things to be art based on their association with the model we maintain of the concept "art." For example, Berys Gaut[4] provides the following list of family-resemblance features, emphasizing that his approach does not require that all criteria are satisfied by all artworks:

1. Possessing positive aesthetic properties
2. Being expressive of emotion
3. Being intellectually challenging
4. Being formally complex and coherent
5. Having the capacity to convey complex meanings
6. Exhibiting an individual point of view
7. Being original
8. Being an artifact or performance which is the product of a high degree of skill
9. Belonging to an established artistic form
10. Being the product of an intention to make a work of art[5]

Thus it may be that we have prototypical works of art—painting, music, poetry, film, and so on—that satisfy all of these things unambiguously, and are the things our minds naturally go to when asked to think of art, but we experience many other things that we tend to categorize or perceive as artistic in some less definitive way. This may include "bad" art or art that we don't like; such work satisfies criterion 10 and perhaps other criteria such as 8 and 9, but not criterion 1, as judged subjectively. We would of course accept that bad art is still art, even if we cannot use criterion 1 to directly evaluate it as such.

As with creativity, the family-resemblance approach then also permits us to include things that we may treat as works of art despite our various reservations: a beautiful technical diagram; a natural rock formation; the

Mandelbrot set; the junk shrines of a compulsive hoarder; art made by animals trained to use paintbrushes or drumsticks. These are things that we see as having artistic value or connotations even though something is not quite right—the work has not been created for aesthetic pleasure, or is not in a recognized field of art, or does not even have a human origin. As with creativity, this approach also doesn't fix requirements on other problematic cases that tend to lead to absurdities. For example, although we generally require originality in art, it seems pointless to insist that a forgery—or worse, something trivially novel but ultimately derivative, a slavish imitation—is not itself a piece of art.

Most importantly, this approach to defining art is in keeping with the complex tapestry of contexts in which it appears in human behavior. We expect that many aspects of human behavior consist of intertwined and overlapping activities that are often caught between categories: art or not-art. This approach relieves us of the problem of art in anthropology. Cross-cultural studies suggest that all cultures possess behaviors *resembling* aspects of Western art, and yet do not necessarily share the same conceptualization of art possessed by Western societies.[6]

A family-resemblance approach stands in contrast to the idea that art is defined by what art institutions or experts declare it to be, although this is one of the listed criteria. As Alfred Gell[7] notes, this is hopelessly limited to our specific cultural context, but even worse, even within such a context this would mean a circular and fickle uncertainty about what was in and what wasn't, depending on the ebbs and flows of a dynamic art world. A family-resemblance approach relieves us of having to draw a boundary between the world of art and the rest of our lives, accepting that categories are porous and permeable. Notions such as that art is lived experience or life is art are more easily constructed across this boundary. For other art-related words in other cultures there could be other associative lists like the one above. What ties artistic behavior together across cultures is not necessarily the conception of art itself but a set of underlying cognitive traits and social dynamics that are associated with humans.

Item 1 in Gaut's list, nevertheless, may appear to many the most familiar and possibly the strongest determining factor in deciding what is and isn't art. We associate art with that stuff made by people that invokes a strong aesthetic reaction, and we know firsthand how powerful this reaction can be. The aesthetic impact of art, through awe, the sublime, or specific emotions,

tends to make us endow it with sacredness, deeper meaning, and purpose. This demands sensitivity from fields such as computational creativity. Engineers promising to "solve"' art, or tame it, like the legacy of the Muzak company, easily become the focus of ridicule or contempt for overreach and simplification. Music frequently invokes such passion. For many, Confucius is right that "music produces a kind of pleasure which human nature cannot do without," it is the "universal language of mankind." This is usually mentioned with the implication that it can stimulate understanding between cultures, and music and art are commonly described as things that "make us human," notwithstanding the circularity of argument therein. To quote Robin Williams in the 2007 film *August Rush*, "You know what music is? God's little reminder that there's something else besides us in this universe, a harmonic connection between all living beings, everywhere, even the stars."

I would add one more element to Gaut's list. In many situations art is associated with an explicit rejection of other purpose, most notably manifest in the bohemian creed "l'art pour l'art," or "art for art's sake." For Oscar Wilde, "A work of art is useless as a flower is useless. A flower blossoms for its own joy. We gain a moment of joy by looking at it. That is all that is to be said about our relations to flowers."[8] Under certain circumstances things that serve some other purpose, such as a television commercial, are treated as lesser examples of art. This may be more than just a unique perspective from a particular romantic viewpoint; it crops up elsewhere and may have a more universal quality grounded in evolutionary theories of display.

The philosopher of art Denis Dutton,[9] for example, looks at common properties of what he calls high art, forming a "conspicuous consumption" view in which art: involves the use of expensive materials; is time consuming to create, either directly or via the accumulation of skill; is remote from any possible use; is intellectually challenging; and is possibly fleeting.

Following the structure of chapter 3, I will work through a number of considerations of human artistic behavior beginning with a narrow view of the individual and expanding outward to incorporate aspects of social behavior. In these sections, I interleave work from psychology, social science, creative practice, and evolutionary theory. As this review navigates topics that are at times disputed, conflicting, speculative, or weakly formulated, it does not work toward any sort of unified theory. The result, instead, is a set of sketches of human artistic behavior at different scales

and in different contexts, grounded in the framework developed so far. The purpose of this is to provide a set of potentialities that any analysis of computational creativity practice can benefit from taking into account.

Evolutionary Explanations

In discussing art and psychology, it is useful to specifically consider the question of evolutionary psychology: how and why, if at all, we might have evolved specific psychological traits. An evolutionary perspective bridges the gap between the mechanical world of physical causality and the intentions and goals of lived human experience. It allows for functions and teleofunctions to be well grounded and nonmystical. Work in evolutionary psychology is often speculative and hard to support by experimental evidence: brains do not fossilize in useful ways, and archaeological artifacts give only very indirect clues as to how we thought and behaved, let alone how those thoughts and behaviors were manifest in genes. Nevertheless, a look at the spectrum of possible evolutionary psychology factors bearing on our understanding of art provides a useful grounding for the most fundamental questions of what art is and why we engage with it.

There are five broad classes of evolutionary narrative relating to human artistic behavior. The first four are concerned with human cognitive traits, while the last takes a cultural or biocultural evolutionary perspective.

Consequentialist theories posit that artistic behavior itself does not have any evolutionary value or functional origin but instead arises from a cocktail of other adaptations that cause us to gain pleasure from certain configurations. Steven Pinker[10] has most famously made this case for music, which he contrasts with language—a cognitive capacity clearly evolved with an adaptive function. His now notorious cheesecake analogy makes the point very well. Does cheesecake have an adaptive function? he asks. Quite evidently not, because cheesecake itself did not exist in the era of our most recent adaptive development. A taste for sugar and fat, however, quite clearly do. Cheesecake should rightly be understood as a human innovation that exploits these traits and others besides to create a supercharged stimulus. Music likewise. This argument has Occam's razor on its side.

Cognitivist theories appeal to the intellectual and aesthetic stimulation associated with artistic behavior, and posit that we engage in forms of artistic behavior in order to serve some form of either cognitive development, reasoning, or exploration and discovery. For example, one evolutionary

explanation for our love of storytelling is as a tool that enables us to think through social scenarios.[11] Other possible explanatory scenarios include ways to structure thought, ways to scaffold the learning of language, and forms of pretend play that enable an exploration of physics, patterns, social relations, and so on.

Cohesionist theories look at ways that creative domains may function to bind individuals together into cohesive social groups. Particularly in music, evidence from anthropology and social psychology supports some sort of group fitness effect.[12] This has been posited as a reason we are expert entrainers (we have an ability to hold a beat, sharing temporal experience with each other). Our ability to act collectively in time and in tune has been compared to the chorusing of wolves[13] and a psychological effect of boundary loss has been posited,[14] whereby we literally lose a sense of self in the process of singing together.

Cohesionist theories are considered problematic because they are potentially *group-selectionist*. On its face, Darwinian evolution explains the selection of fit *individuals* well, but not of fit groups; adaptations that are beneficial to groups of individuals must be robust in the face of freeloading individuals. This is the wider problem of the evolution of altruism.[15] If singing music in a group stimulates more pro-group behavior in individuals, then the fittest individuals would be those who could fake the cohesive behavior without wasting effort in contributing—in other words, they are evolutionary tricksters or freeloaders. But this selfish-gene objection is not unsurpassable. Theoretical solutions for group selection do exist, and models of gene-culture interaction such as those of Boyd and Richerson[16] and the niche construction theorists[17] also suggest possible conditions in which pro-cohesion selection pressures can emerge.

A common approach to solving the problem of altruistic behaviors is by appeal to inclusive fitness,[18] or kin selection: the principle that any behavior that benefits one's closest kin is also likely benefiting one's genes and can therefore be selected for. Models that take into account population structure show that altruism caused by inclusive fitness can expand to apply to large populations.[19] This fits with another strand of cohesionist thinking around music, that posits primitive musical behaviors—basic gestures and the ability to rhythmically entrain—as essential to the development of intersubjectivity, and hence language and sociality, in caregiver-infant interactions. Dissanayake,[20] for example, develops a multifaceted

theory that brings together the development of intersubjectivity in earliest childhood with the role of art, broadly conceived as a cultural act of *making special*, in constructing markers and loci for the development of group identity.

Competitivist theories posit scenarios in which artistic behavior might form the basis for interindividual competition. The most commonly explored instance of this is in sexual selection, typically with the male members of any species (who tend to invest less in their offspring due to the uncertainty and reduced physical cost of fatherhood) competing to attract females to mate with. Sexual selection may take the form of an unconstrained runaway train of coevolution toward bizarre and arbitrary aesthetic features, but many theorists believe that it will always be bound by the logic of *honest signaling* or the *handicap principle*,[21] according to which displays have evolved to be indicators of "good genes." Various evolutionary psychologists, most outspoken among them Geoffrey Miller,[22] have argued that baroque and apparently functionless social behavior in humans—music, art, storytelling, anything with the potential to show off one's skill in the construction of complex and elegant forms—is likely to be sexually selected behavior. Whereas for Pinker the apparent uselessness of art and music is because it actually serves no function, for Miller this apparent uselessness, coupled with the great time invested in it, is the signature of sexual selection and honest signaling theory: showing off through exuberance.

However, in humans there is good reason to believe that competition in creative (and other) domains is not only manifest in the attraction of mates. Late night dance floor exertions might have a more blatant connection to sexual display, but some of Miller's anecdotal themes (serenading at balconies, the promiscuity of male stars like Jimi Hendrix, and so on) are more tenuous. Looking at rich and famous musicians is confounding exactly because they are rich and famous. As others such as Paul Bloom[23] have noted, in the rich cultural context of humans a more general *social* selection theory is just as powerful at explaining the emergence of human artistic behavior. It is just as effective to attract a business partner with a display of creativity or aesthetic sophistication (and hence gain financial success with which to attract a mate) than it is to directly attract a mate.

Another body of competitivist theory includes social brain and Machiavellian intelligence hypotheses,[24] which posit that the evolution of human intelligence has largely been driven by selection pressures for social intelligence

such as the ability to deceive, manipulate, and form stable alliances with others. This can invoke a runaway evolutionary process, with humans living in larger groups and consequently having greater incentives, and needing more sophisticated methods, for social intelligence.

Emergentist theories take a different tack from the other four. Rather than focus on how the human cognitive machinery underlying art has come about, they look at domains of art practice as complex systems that may be self-generating. Thus, rather than there being an identifiable external cause underlying these domains (musical skill indicates good genes, aesthetic pleasure arises from solving perception problems, and so on), the system as a whole is emergent. Luhmann,[25] for example, looked at cultural systems of art production and consumption as *autopoietic* systems, sustaining themselves through a metabolic process. A human being might be enculturated into an aesthetic cultural system, compete for success in that system, and then act to reinforce that system, much as in the social dynamics of a cult.

It can also be seen how learned or culturally emergent behaviors can drive further biological evolution. This effect, known as the Baldwin effect after its initial identification by James Baldwin, has been demonstrated in various species, most famously in humans in the case of the evolution of lactose tolerance in subpopulations of humans. This is closely related to the theory of niche construction discussed in chapter 3.

These five bodies of theory are diverse and piecemeal. Anyone looking at the biological foundations of human artistic behavior should reasonably expect that it is not a simple whole with a simple explanation, but a complex assemblage constructed from multiple evolutionary pathways. Claims to identify a simple origin story for art are likely to be suspect.

The Innate and the Universal

No matter how seductive the approach has been to computer scientists, it has long been untenable to uphold a strict theory of universal beauty, in which some set of measurable properties of an aesthetic artifact tells us how any human observer might form an aesthetic response to it. A growing body of work in psychology points to the extreme plasticity of the human brain, dismissing the wishful thinking that aesthetic objects can be understood simply in terms of universal features such as symmetry, complexity, and harmonic ratios. At the same time, it is equally problematic to think of brains as blank slates shaped only by lived experience. The plasticity

of brains comes from powerful, genetically evolved learning capacities which have specific properties. Even if our brains were perfect blank slates (whatever that could possibly mean), the lived experiences of humans have many universal features: horizons, sun and stars, gravity, visual depth, the harmonics of resonant bodies, the blues and greens of nature, the feeling of water, animal behaviors, faces and bodies, sex and sexes, laughter and tears. Whatever the extent to which we are shaped by our genes or our environments, humans across the world and across millennia have much that is common in their behavior, and nontrivial perceptual universals do abound in human behavior: our recognition of facial expressions, the prosody of speech and the sounds we make to babies, the low-level mechanics of visual and sonic object perception, the tendency to perceive metrical structure in rhythmic patterns, and so on.

There are also various widely agreed upon aspects of the brain's overall structure. One is that the brain can be divided into two distinct thinking systems.[26] The first system is instinctive, subconscious, and designed for rapid response in critical situations. This system is known as the fast system or *System 1*. The second system is conscious, more general purpose, and analytical, dealing with problems that have not been encountered before and need to be thought through. This is known as the slow system or *System 2*. Several theories of aesthetic response, as discussed below, depend on this architecture, often explaining aesthetic phenomena in terms of interactions between Systems 1 and 2.[27] Another widely agreed upon architectural aspect is that myriad low-level neural processing components, such as for edge detection in visual processing, coexist and interact. Each of these individual elements and the overall architecture of such elements may be more or less fixed across cultures and individuals. In visual information processing, this model, sometimes known as the pandemonium model, has been used to explain how we manage to perceive objects regardless of their orientation, by abstracting the component elements of objects from the visual field and reconstituting them as higher-level concepts.

But what can we say about universal aspects of aesthetic perception? Are there things that all humans find beautiful, regardless of their lived experience? In classical theories of universal beauty, beauty is understood via the mathematical properties of an aesthetic object. Candidate properties include symmetry, various measures relating to complexity including fractal dimension, frequency distributions, and information content, and

the existence of various significant ratios such as the golden ratio and the square root of five. Historically, researchers have come at these principles less from a grounding in psychology than from a preexisting belief in a mathematical basis for universal beauty. This is confounded by a potential cultural circularity: the fact that classical examples that are revered as great works have to greater or lesser extents been consciously created according to these mathematical principles.

George Birkhoff's influential 1933 work presented an early effort to develop an empirical study of aesthetics. Birkhoff explains:

> The typical aesthetic experience may be regarded as compounded of three successive phases: (1) a preliminary effort of attention, which is necessary for the act of perception, and which increases in proportion to what we shall call the complexity (*C*) of the object; (2) the feeling of value or aesthetic measure (*M*) which rewards this effort; and finally (3) a realization that the object is characterized by a certain harmony, a symmetry, or order (*O*), more or less concealed, which seems necessary to the aesthetic effect.[28]

For Birkhoff, the relation between *M*, *C*, and *O* is the surprisingly simple equation $M=O/C$. He explains this by analogy to a return on investment in business. "In each business there is involved a certain investment *i* and a certain annual profit *p*. The ratio *p/i* … is regarded as the economic measure of success." For Birkhoff, the challenge for an artist is to make order out of complexity, as the businessman strives to make profits out of investment.

More recently, studies of complexity have been the focus of psychological research. Alex Forsythe and colleagues demonstrate a relationship between measures of human judgments of beauty and the visual complexity of images. The latter is measured by how compressible the images are, with GIF compression providing the most effective correlation with human judgments.[29] Their results support the theory of psychologist Daniel Berlyne[30] that there is an aesthetic preference for a sweet spot of complexity; we like things that are not too simple, not too complex.

Similar results have been obtained in music. Richard Voss and John Clarke[31] examined the structure of audio signals in music as well as language. They found that these signals exhibited a consistent pattern of complexity, described by the distribution of frequencies within the signal, which exhibit a 1/f law, which states that the presence of any frequency is inversely proportional to the frequency itself. This distribution is also known as pink noise, as compared to a flat distribution, where all frequencies are equally

present (white noise), and a $1/f^2$ distribution (brown noise). Intuitively, a sequence exhibiting white noise will be too unpredictable to be enjoyable, whereas a brown noise sequence will be too predictable to be enjoyable. Voss and Clarke's results do not apply directly to the sound spectra themselves but to derived aspects of the signal: the power and the approximated overall frequency. It is the variation in these signals that exhibits the 1/f property. Voss and Clarke then generated musical sequences with these various distributions and showed that the 1/f sequences were considered the most pleasing. Bill Manaris and colleagues have similarly examined this law applied to note sequences in symbolic representations of music.[32]

Studies of complexity appeal to our perceptual ability to make sense of the world. I will continue this discussion below in the context of a more adaptive variant of the types of aesthetic theory just outlined. Meanwhile, other candidate fundamentals of beauty have been looked at in terms of more specific potential evolutionary origins. One example is the importance of facial and bodily symmetry in indicating fitness. Symmetry in body and face morphology can be an indicator of health or good genes and is believed to be a good predictor of perceived attractiveness for this reason.[33]

Reflectional symmetry of static images is only one form of symmetry. Other researchers have shown symmetry of dance moves to be attractive cross-culturally.[34] In both cases, it should be added that human faces and human bodies likely form a specific, cognitively, innately encoded category. We are far more sensitive to human bodies and faces than to other objects, with dedicated capacities for recognizing facial expressions, remembering hundreds or thousands of distinct individuals, and modeling and empathizing with bodily actions. So we might expect that a preference for symmetry falls into a class of specific body-related cognitive functions, rather than generalizing to other object forms.[35]

Others have attempted to identify biologically evolved universals at higher composite levels. The landscape preference hypothesis posits that we find certain kinds of landscape scenes aesthetically appealing because they represent suitable landscapes for habitation.[36] Preferable features, it posits, include open spaces with clustered trees, the presence of water, a vantage on the horizon, animals, and some degree of diversity or moderate complexity.[37] The supporting evidence suggests that the landscape preference is more strongly felt in children but is slowly overridden by learned preferences.[38] Such thinking stems from the late twentieth century turn

in evolutionary psychology, pushing back against a view of the mind as a blank slate that gains its structure and knowledge through learning, to one in which the mind has multiple dedicated evolved subunits, sometimes described as analogous to a penknife. According to such a view, something as specific as a landscape preference might plausibly have evolved in some corner of the brain. Appeal to other examples of innate responses to visually encoded forms support the basic premise: "The system that identifies snake-ness in the visual array, coupled to fear-releasing circuits, is an adaptation that lowered deaths due to venomous snake bites among our ancestors."[39]

John Tooby and Leda Cosmides, chief proponents of this penknife view of the evolved brain, in which many such specific adaptive behaviors can be identified, have other such adaptive examples that apply to high-level phenomena. They offer our pleasure response to fiction as another candidate.[40] Stories provide a platform for exploring imaginary social scenarios, they argue, and so we have evolved a pleasure for hearing and telling stories. Essentially, narrative engagement helps us experiment with potentially real future scenarios just as the pretend play of the young of many species helps prepare them for fight and flight tactics. It is a form of social practice. As with the landscape theory, some compelling evidence is provided, but so far we have little idea how such preferences might be encoded cognitively.

More generally, Tooby and Cosmides appeal to nonconsequentalist approaches to explaining artistic behavior: "Almost all of the phenomena that are central to the humanities are puzzling anomalies from an evolutionary perspective. Chief among these are the human attraction to fictional experience (in all media and genres) and other products of the imagination … involvement in the imaginative arts appears to be an intrinsically rewarding activity, without apparent utilitarian payoff."[41]

The strong adaptationism of evolutionary psychologists like Tooby and Codmides is commonly criticized as being *overly* adaptationist. Anti-adaptationist arguments, most notably those of Stephen Gould and Richard Lewontin, point out that many behaviors are unlikely to have direct biologically adapted origins and really do not warrant them.[42] They may instead be what Gould and Lewontin call *spandrels*,[43] stemming from an adaptive origin that does not directly relate to the behavior in question. Partly, this is an appeal to Occam's razor, but it is also simply a critique on the plausibility of certain attempts at evolutionary origin seeking. The phrase "just-so

stories," named after Rudyard Kipling's famous short stories, is the pejorative term used to describe such speculative evolutionary storytelling. For reasons of programmatic simplicity, the lower-level the adaptive trait, the more plausible it is in evolutionary terms. An innate preference for simple harmonic ratios is more plausible than an explicitly encoded preference for complex symphonies or composite landscape scenes.

Consider, in the case of a landscape preference, what a good job conscious thought does of habitat selection, at least in modern humans.[44] "There's water, that's good! And there are trees, also good! But I'm not sure about that bear." Could such thought be subjugated to some subconscious urge that drives you blindly toward a preferred landscape? From an informational point of view it is hard to grasp the task of encoding such preference in a way that bypasses conscious categorical thought and goes straight into a pleasure response. Nevertheless, we do find ourselves rapt by certain landscapes and natural scenes in a way that does indeed appear to bypass conceptual or conscious thought, a phenomenon that warrants some explanation.

This brings us back to the Pinkerian perspective, where we identify low-level components contributing to aesthetic responses, whatever their evolutionary origins, but don't expect to find coherent evolutionary adaptations at higher levels. For example, we might like horizons to be horizontal because we like to be upright but not because they form part of a safe environment.

The Pinkerian view of music as a set of cultural technologies built to exploit existing pleasure responses leads us to look at how the basic work of making sense of our world provokes these pleasure responses. In auditory perception, it is widely recognized that we have a series of innate perceptual capacities, such as for clustering sounds according to common origins (streaming and source separation), tracking beats, and recognizing pitch intervals. Some of these capacities have fairly uncontroversial evolutionary adaptive origins based in making sense of the world. Source separation for example, is essential to our ability to analyze auditory scenes. We see this most clearly in the "cocktail party effect," the ability to train our attention on one voice in a room, amid a babble of people talking. Music often plays with such effects, mixing timbres and their movement through time in such a way as to construct complex auditory scenes where sonic objects merge and separate. Here Pinker's analogy is at its strongest: chefs and composers alike weave together known responses in original ways.

We also know that naturally occurring sounds have common structural properties.[45] When you hit a hollow tree or pluck a guitar string, the sound has a percussive initial phase, also known as the transient phase, which is characterized by a high noise content for a short duration—you've just excited a physical system and thrown it into a chaotic frenzy. After this transient phase the remainder of the sound is defined by resonance: a regular oscillation (hence non-noisy, with a clear tone) that emerges out of the energetic movement of the initial impact, which gradually loses energy and dissipates.

Another way to make sound is to scrape something or blow air across it. In both cases you get a noisy, grainy sound caused by myriad tiny interactions (imagine dragging something through gravel or wind blowing through trees causing leaves to move). In some cases (blowing across a bottle or scraping a string with a bow) the system also feeds back on itself so that those myriad tiny interactions coordinate to stimulate the natural resonances within the system they inhabit, and hence, again, create more regular, less noisy, sounds—tones such as those of a flute.

Furthermore, when sound-making objects do resonate they tend to do so with signature spectral properties. Stably oscillating objects produce sets of harmonically related frequencies. A harmonic frequency relation is simply one where the ratio between frequencies is simple, such as 1:2 (octave) or 2:3 (perfect fifth). Harmonically related frequencies reinforce each other, so it is natural that this should be the product of a resonating process. Some metal objects like bells, gamelans, and steel drums have more complex spectral signatures, as the material allows multiple resonances to co-occur, but strings, tubes, and vocal cords generally resonate with clear harmonic structures that define the *timbre* of the sound.

It is therefore a fairly uncontroversial suggestion that our auditory system may be specifically adapted to decoding these properties of naturally occurring sound, recognizing individual auditory events as percussive-resonant sequences, and grouping frequencies based on their harmonic relatedness. Note that *because* these types of sound reliably occur in nature, it is equally plausible that we exploit the environment to learn them, as that we are born with readymade perceptual systems that are already prepared for such an environment. A clear distinction between what *exactly* is encoded in genes and what is learned from the environment is not necessarily obvious. It has also been suggested that language and music mimic

these properties because they are so well established in our perceptual system. Indeed, in speech we identify fricative (scraping), plosive (percussive), and sonorant (resonant) elements, but given that language and song also use natural physical actions to make sound, it is somewhat unremarkable that our vocal systems also display such characteristics.

A slightly more mysterious aspect of our auditory perception is the ability to track beats, known sometimes as entrainment. Specifically, if I start clapping at regular intervals, you will be able to clap along, even if I gently speed up or slow down. If I stop, you will be able to continue at the tempo I have established. I can join back in, you can lead the tempo changes, and so on.[46] Other animals cannot do this, at least nowhere near as well as we can. Many can synchronize cyclical behaviors but not in the adaptive, tempo-invariant, predictive way that we do. This ability is constrained to a specific speed range; try to follow a beat at intervals smaller than 100ms or larger than 2s and you will not manage to successfully predict the time of the next beat, but keeping within this range you have an innate capacity to keep time.[47]

Compared to our auditory event detection, the environment doesn't offer such obvious reasons why we might have developed this entrainment ability. Many ideas have been suggested, drawing on pretty much anything that creates a regular beat: heartbeats and walking are example candidates that have been appealed to in evolutionary theories. One theory has it that we needed to time our vocal communication with our cadence, for better audibility. Social factors are also compelling sources of explanation. Even if it has hitchhiked on the perception of natural sounds, vocal communication in various species seems to have occupied, or it is perhaps better to say forged, a novel acoustic niche in which to operate. A body of theory revolves around the basic mechanisms of developing the type of intersubjective engagement that is the foundation for language learning.[48] Studies of mother-infant interaction have claimed to identify rhythmic engagement between the participants.[49] Another idea is that rhythmic entrainment establishes a sense of boundary loss among individuals in a group, enforcing group cohesion: we unite by sounding and moving in time.[50] However, a Pinkerian explanation is equally possible here: perhaps we simply evolved the ability to manage timed sequences very well, for example in learning to make tools, or perhaps we simply have the general intelligence necessary to learn this skill and find it compelling.

As we move to higher levels of structure, appeal to environmental factors becomes harder. Theories such as the group theoretic work of Gerald Balzano in music propose an aesthetic preference for kinds of minimally complex structures.[51] Balzano's particular focus was on the structure of diatonic scales in Western music, noting the way in which these scales form near-symmetric subsets of a larger set (the major scale is a subset of the chromatic scale and is asymmetrically distributed—think of the arrangement of black and white keys on a piano keyboard). Balzano argues for a general principle of form that can be manifest in different ways in different cultures. For example, while diatonic groups are prevalent in Western music, similar patterns exist in rhythmic manifestations in certain non-Western cultures.

Balzano's work resembles a form of structuralism, the late twentieth-century movement that viewed our brains as inclined toward certain structuring relations, which might be manifest in different ways in different cultural contexts. The leading structuralist anthropologist Claude Lévi-Strauss mapped the categorical relations expressed in South American myths as they varied from place to place, identifying a common structure around which a series of well-formed variations exists.[52] According to this view, we have preexisting mental structures that are filled by some culturally or environmentally specific content. This is just one way of combining notions of fixed universal and fluid learned aspects of perception.

Increasingly, however, a *Hebbian* paradigm of neural learning is proving dominant, in which we view the interaction between evolutionary and developmental components of neural circuitry as one in which the process of learning shapes the neural structure.[53] The theories described in the following sections continue this theme, supporting a richer sense of cultural and environmental adaptivity.

Adaptive Perspectives and Behaviors

A more complex and more plausible view of aesthetics focuses less on universals and more on the adaptive capacity of the mind. Accordingly, the universal aspects of our cognition are the learning mechanisms sitting under the surface rather than the specific perceptual and evaluative systems. Human cultures are, after all, highly adaptive to their environmental or historically contingent circumstances, and yet cultural and individual variations are constrained in some way by these fixed structures of the mind. This is a theme continued in modern psychology. With respect to

the aesthetic foundations of visual art, Vilayanur Ramachandran and William Hirstein propose ten psychological principles.[54] Each relates to evolutionarily and psychologically plausible properties of our visual system and appeals to the pleasure we obtain in the process of learning about and making sense of the world, much in the way that Pinker describes music as auditory cheesecake. The principles are: peak shift, perceptual grouping and binding, contrast, isolation, perceptual problem solving, symmetry, abhorrence of coincidence/generic viewpoint, repetition, rhythm and orderliness, balance, and metaphor.

I will not discuss each in detail, but consider some general properties. The elements in this list are more or less context dependent, combining a universal psychological principle with a cultural environmental context which dictates the specifics of the aesthetic response. For example, they explain the peak shift principle as follows:

> If a rat is taught to discriminate a square from a rectangle (of say, 3:2 aspect ratio) and rewarded for the rectangle, it will soon learn to respond more frequently to the rectangle. Paradoxically, however, the rat's response to a rectangle that is even longer and skinnier (say, of aspect ratio 4:1) is even greater than it was to the original prototype on which it was trained.[55]

This is how they apply the concept to art:

> Consider the way in which a skilled cartoonist produces a caricature of a famous face, say Nixon's. What he does (unconsciously) is to take the average of all faces, subtract the average from Nixon's face (to get the difference between Nixon's face and all others) and then amplify the differences to produce a caricature.[56]

The peak shift principle offers an explanation for why we don't just learn to like what we know, but are driven toward certain types of new and surprising forms. It makes for a good explanation for the extremes that emerging genres strive to achieve. Apply the same process as the caricaturist to a subgenre of, say, heavy metal, and you have a candidate explanation for why that subgenre pushes out at the extremes so readily, as in thrash or doom metal. It suggests a form of categorical learning: the rat hasn't just learned a shape but a concept, which itself defines a space of unexplored possibilities. The peak shift principle could be thought of as an example of where our ability to learn patterns is imperfect (the rat is not trained to generalize), but nevertheless drives aesthetic choice and a desire for certain forms of novelty.

Theories such as Balzano's and Ramachandran and Hirstein's form a bridge between preferences for simple mathematical qualities such as

symmetry or fractal dimension, and a more adaptive concept of aesthetic preference based on the ability to model one's environment adaptively. The latter types of theory are often constructed in the language of information and complexity, although we find parallels in the work of Csikszentmihalyi and specifically his theory of *flow*.[57] Csikszentmihalyi's starting point for the concept of flow is that we find pleasure in things that support our learning about the world, and those things that best support our learning strike a balance between two extremes: they are neither too easy or common, nor too hard or unrecognizable. Pleasure in performing tasks comes when the tasks are neither boring nor overly challenging. Boredom is the expression of the desire to experience new things and thus learn. When we achieve what he calls a flow state, caused by an optimum balance of task difficulty, we become intensely immersed in what we are doing to the point of being lost in the task and losing sense of time. We can see how complexity-based theories of aesthetics such as Balzano's model of tonal structure, serving to provide a platform for engaging tonal puzzles, might satisfy Csikszentmihalyi's principle: aesthetic structures that are not too radically complex but not tediously simple either.

Peak shift and flow relate to a more general response profile that is familiar in psychology and relevant to creativity, where novelty and pleasure are related by an inverted-U curve, named after Wilhelm Wundt: the Wundt curve.[58] The Wundt curve underlies a model of curiosity and scaffolded learning about the world. The basic idea is that we are stimulated by new things (perceptual experiences), as part of an adaptive tendency for curiosity; it is good to learn new things about the world. Specifically, we are interested by things that are neither too familiar nor too unusual.

It is clear why familiar things don't tell us anything new about the world, but less clear what the problem is with things that are too unfamiliar. The proposed answer is that they are too hard for us to deal with. Learning about the world has to be done in small steps: some things just don't make sense until we have the framework to make sense of them.

Such phenomena illustrate ways in which cognitive strategies for learning about the world and sense-making might lead to culturally or individually determined aesthetic preferences. I discussed in the previous chapter how copying behaviors enable cultural transmission and hence cultural evolution. These learning phenomena show how cultural artifacts might be driven to change in a runaway evolutionary process. Given a motivation

to learn about the world, then, we set the scene for cultural artifacts to undergo constant modification in response to a need for novel stimuli.

Daniel Berlyne was one of the first to formalize notions of boredom and curiosity. He defines human "epistemic curiosity" as "a drive reducible by knowledge-rehearsal."[59] In Berlyne's experimental research, he shows that absolute complexity of images combine with the degree of novelty to determine the "hedonic" response of individuals. More complex images become *more* pleasurable with greater exposure to them, due to familiarization, while less complex images become *less* pleasurable, due to boredom.[60] In more recent work the information theoretic dimension of such a line of research has come to the fore. For machine learning scientist Jürgen Schmidhuber,[61] the same idea of finding beauty in optimally complex learning material is presented as the need to get a good rate of *information compression*. To understand this we first note that learning about the environment necessarily involves data compression. Learning the underlying principles of form, rather than, say, keeping a complete log of every scene ever to pass in front of us, is essential if we are to efficiently remember things, let alone form a generalized understanding of the world. For example, when a child begins to learn language they first learn the set of phonemes of their native language (which differ between languages, dialects, and accents and hence must be learned), and this forms a discrete symbol set upon which the language is built. This is a form of data compression that mediates between the sound spectra we hear and the words and sentences we understand. If you later learn a distant language, like an English person learning Chinese, you need to add new phonemes to your palette, thus rearranging this compressed mental representation, and this may be a taxing process. Until you do so, you cannot correctly distinguish certain sounds in that language; they are inaudible to you. It is important to note that such learning can be done to some extent in an unsupervised manner, meaning that you just need to experience the world in order to learn this structure; you don't need to be told what the right classification is. Through repeated exposure we begin to notice more detail.

Thus according to this view we are constantly attempting to compress the data arriving via our perception of the world. A new musical melody would be understood in terms of the abstract representations derived from previous experience of melodies, and everything is jiggled around to fit this new melody into the model. For both Berlyne and Schmidhuber, then, in

different ways, the idea is that our success at information compression is a source of pleasure; we seek experiences that help us make better sense of the world. Just as with flow, something too familiar or simple is boring and unstimulating, but so is something too unusual or complex. More generally, these theories offer some form of explanation of curiosity.[62]

Relatedly, in David Huron's research into the evolutionary psychology of music, he takes a "neural Darwinism" approach, according to which the mind is constantly trying out candidate perceptual strategies to see which ones perform best, thus cultural factors can strongly influence the way we learn to perceive the world. According to a neural Darwinist view of music perception:

1. There are competing mental representations of sound.
2. Representations are shaped by exposure to the environment.
3. Representations are differentially favored depending on their predictive success.[63]

As an example, Huron considers the development of absolute pitch versus relative pitch in individuals. Absolute pitch perception is a skill that enables someone to recognize a musical pitch without previously being primed with a reference pitch. Most people perceive pitch relatively (also known as intervallic perception); we are able to perceive intervals between pitches, even if we don't know what those pitches are. A trained musician who does not possess absolute pitch will typically prime themselves with a reference pitch, after which they can deduce further pitches based on their relative distance to the reference pitch, for a short period of time. According to a neural Darwinist view, our brains are equipped to adopt either of these ways of perceiving pitch, based on which works out best in our own lived experience of the world, but at some point we adapt to using one or the other and this aspect of our perception becomes more or less locked in.

Huron argues that this selection is made based on what works best in our given environment. Absolute pitch can be a useful skill for professional musicians, whereas it can also be a burden and is largely unnecessary. When singing happy birthday at a party, it is fine just to lock in on whatever key emerges among the singers and to use intervals from there on. Huron's evidence is that absolute pitch emerges under specific social circumstances, in particular in societies where there is an established standard tuning. It is learned very early on, at a time when the brain might be wiring in these kinds of perceptual strategies. Supporting this, tests show

that absolute pitch performance is correlated to the exposure to specific pitches; someone with absolute pitch will identify more common pitches more quickly than less common pitches.

In these models, a set of cultural conditions might shape how our perception of the world forms. This is broadly consistent with the idea of forming successful "compressed" models of the world. Huron's central thesis is that these cognitive functions are there because we are forever trying to limit the unexpected, to model the world in such a way as never to be surprised. However, he argues, surprises themselves can be pleasurable exactly because they feed this learning function, through the interplay between our fast (System 1) and slow (System 2) cognition systems. System 2 appraises and feeds back into the System 1 response, and where contradictions occur, we are stimulated via a mechanism dubbed *contrastive valence*.

Putting all of this together, Huron proposes the imagination-tension-prediction-reaction-appraisal (ITPRA) model, designed to explain how aspects of musical experience are stimulating. The ITPRA theory takes into account five stages in the response to temporal sequences of auditory stimuli:

- I—*imagining* different outcomes, forming anticipation
- T—as anticipated event approaches, *tension* builds
- P—*prediction* response—feeling evoked in response to success of prediction
- R—At the same time as P, a fast *reaction* response (System 1)
- A—a slower *appraisal* response (System 2)

Such a model of perception again offers a framework in which new culturally determined experiences might dictate new neural perceptual frameworks, completely novel conceptual spaces, constructed on top of a fixed, high-level architecture, accounted for by sound evolutionary theory. The basic mechanism for musical pleasure is hardwired, but the ways in which different learned aspects of music are used to stimulate this system may vary considerably.

Geraint Wiggins and Jamie Forth[64] offer a computational model broadly related to Huron's ideas and drawing on a model of central cognitive information processing called the global workspace theory. The global workspace theory defines a top-level cognitive architecture in which a central cognitive system associated with conscious awareness interacts with a large number of simple, domain-specific subconscious systems, competing for

attention. Specifically, in their Information Dynamics of Thinking (IDyOT) model, the focus is again on temporal sequences relevant to music and also language, and the submodules considered are Markovian processes that attempt to predict what will happen next in a given situation. When certain conditions are met, any one of these submodules might trigger the attention of the global workspace.

Such multiple hypothesis models are appealing for a number of reasons, not least that they offer a neural correlate to the incubation-illumination model of creative cognition. More generally, they offer a plausible neural model of the relation between conscious and subconscious cognitive processes (whether or not this strictly corresponds to fast System 1 and slow System 2 thinking). They can also plausibly support a neural Darwinism approach, helping explain *how* the brain manages to adapt to cultural specifics. Lastly, they correspond well with studies of music perception. For example, it is understood that our ability to track beats is associated with the way that we perceive metrical structure in music: not as a single temporal representation but as an array of harmonically related sensitivities to different tempi. One theory of metrical perception is that an array of coupled resonant oscillators interact to form a structured temporal representation.[65]

Wiggins and Forth's model is also generative as well as perceptual. In the absence of stimuli, the submodules "freewheel." This is a natural thing for a sequential model to do if it is allowed to feed back on itself, making predictions based on its own previous sequence of predictions. So the same system that is used to anticipate events in the world can also operate as a natural generator of imagined events. Just as these models present themselves to the global workspace in the course of perceiving the world, they can also do so in the course of idle thought, leading to what Wiggins and Forth describe as "spontaneous creativity."

In the IDyOT model, the condition for submodules to attract attention from the global workspace is that of high information content. When a submodule produces high information content, this alerts the central system. Such high information content shares the same properties of things we find as aesthetically pleasing according to models such as the Wundt curve: not already known, but not so unusual as to be treated as noise.

In summary, an important set of theories explaining aesthetic perception focus on the sophisticated mechanisms we employ for modeling and making predictions about the world, that variously hypothesize:

- Mechanisms for compressing and organizing information into efficient representations that help us make predictions;
- Mechanisms that *stimulate* us to seek experiences that drive efficient representations;
- A number of low-level, domain-specific neural processing structures ranging from more hardwired systems to more data-driven neurally plastic systems, possibly competing to be adopted based on performance;
- A central system associated with conscious awareness, that interacts with myriad subsystems, for which the high-level architecture is fixed;
- Certain subsystems that as well as being involved in perception can also be set to freewheel, causing imagination.

This model is not mutually exclusive to the idea of additional dedicated modules evolved to perform specific functions. Thus a very broad view of how genes dictate the shape of the brain is that certain global architectural features are hardwired—the very lowest-level features (how neurons grow, fire, and adapt) at least are definitely hardwired—and in addition certain domain-specific submodules may be hardwired, or at least set up so as to be very likely to form in certain ways given normal stimuli.

Acting in Webs of Significance

We now have a psychological sketch, albeit painted in a broad brush, that complements the themes developed in chapter 3 to fill out an image of humans acting as art-making or art-experiencing agents. Such models of perception provide a solid grounding for understanding how it can be possible for different individuals to develop different views of the world, to be motivated to explore new experiences, and also to spontaneously imagine and generate new things. Accordingly, individuals learn by building models of their world, and are intrinsically motivated to do so, through a distributed process that is compatible with the essential properties of individual creative thought. Given that much of what stimulates us, that is, much of what we spend our time trying to model, is cultural—produced by other humans—rather than from the biological world, then the possibility of runaway cumulative culture developed in chapter 3 follows relatively easily. Inhabiting a world of tonal music, for example, means developing perceptual models of tonal music to which we are attuned.

Even if the underlying psychological mechanisms are grounded in the simplest evolutionary context—that of learning about one's physical environment—this is sufficient for explaining open-ended cultural evolution. New concepts or patterns that arise in the realm of culture, absent in our physical environments, have the potential to form and evolve through cultural processes. This fits into consequentialist and cognitivist evolutionary frameworks. We have seen some cohesionist and competitivist evolutionary explanations, involving more complex selective pressures grounded in social behavior, that provide even richer models of runaway cultural or biocultural evolution. Although considering the impact of such theories on how we perceive human culture is tantalizing, they remain relatively speculative.

Either way, the anthropologist Clifford Geertz's celebrated maxim holds true that "man is an animal suspended in webs of significance he himself has spun."[66] Cultural systems can be autopoietic: generating themselves and maintaining their structure through the feedback of individual human action. Peter Berger and Thomas Luckmann's *The Social Construction of Reality*,[67] first published in 1966, successfully characterized this capacity for social systems to create their own contexts for future interactions. We might update this to a more snappy computational analogy proffered by some theorists, that the conceptual materials that occupy our cultural lives define a sort of *virtual reality*—much more immersive than the contemporary technology of the same name—that runs on the hardware of highly plastic human brains. This virtual world is the most spectacular of constructed niches, in fact a tapestry of interlocking niches. Because individual success is predominantly determined by social interactions, we accept these niches as the contexts in which we must act. The strategies of modern humans getting by are sociocultural strategies: do well at school, make desirable goods, bet on the right horse, learn a language, dress smart, get noticed.

Art may be far from an innocent bystander in this tapestry of social concerns, playing many different roles in culturally constructed niches and their social dynamics. That our taste in art, music, and jokes is sculpted by a process of learning about the world establishes the foundations for the diversity or artistic groups and domains but is only part of the story.

Individual and group identity is central to the development of this theme. More than simply being something that we learn to like, forms of

art are things that we use to define ourselves, whether actively or subconsciously. The sociologist Pierre Bourdieu played an important role in introducing several useful terms and ideas to this area. He refers to *cultural goods* as these objects of exchange and identity formation, which as with my use of "art" can refer to practices and artifacts ranging from high art and music to tattoos and hairstyles. Bourdieu's empirical focus in his study of taste was on class identity, and one of his most important contributions is the idea that taste and its associated expertise in cultural fields become indicators of our class membership. We may not be expert composers, artists, or poets, but we become expert consumers, connoisseurs, of these fields, knowing and appreciating the finer detail:

> Taste classifies, and it classifies the classifier. Social subjects, classified by their classifications, distinguish themselves by the distinctions they make, between the beautiful and the ugly, the distinguished and the vulgar, in which their position in the objective classifications is expressed or betrayed.[68]

Two related cultural phenomena form the basis for Bourdieu's analysis: cultural competencies and cultural capital. *Cultural competencies* are those fine-tuned abilities that display depth of immersion in a field of practice. Competencies, Bourdieu argues, are not simply consciously and explicitly established through education but are established in "total, early, imperceptible learning, performed within the family from the earliest days of life." This acquisition must occur early on, and over a long duration, adding a cost that cannot be faked. This imperceptibly acquired cultural identity that is not easily faked or shaken off, Bourdieu terms *habitus*: "What are grasped through indicators such as educational level or social origin or, more precisely, in the structure of the relationship between them, are also different modes of production of the cultivated habitus, which engender differences not only in the competences acquired but also in the manner of applying them."[69]

For Bourdieu, *cultural capital* is that which is possessed by virtue of a lifelong enculturation giving access to this immersive learned experience (capital is "accumulated labor"). Bourdieu's analysis goes further, distinguishing between the types of cultural capital that can be acquired through education (educational capital), showing that education can provide *some* access to the more mainstream domains such as fine art or classical music by individuals who do not have access to such domains though family background. "Nothing more clearly affirms one's 'class,' nothing more infallibly classifies, than tastes in music."[70]

For Bourdieu, taste as a set of deeply felt preferences and interests is intimately tied up with competence in a field. Simply liking something is both an aspect of the underlying mechanism and a surface quality of the more functionally pivotal phenomenon of indicating competence in that area. Being able to reel off names of relevant artists and their various interconnections and histories, being able to recognize works of a given artist despite only subtle differences between them and others in their field, using the correct descriptive language and so on, these can be seen both a consequence of and a driver for love of a certain domain. Bourdieu elaborates on how such learned qualities can perform a critical function in a system of cultural economics (Bourdieu believed in a "general science of the economy of practices"—that is, economics applying elsewhere besides where money is found).

> The initial accumulation of cultural capital, the precondition for the fast, easy accumulation of every kind of useful cultural capital, starts at the outset, without delay, without wasted time, only for the offspring of families endowed with strong cultural capital; in this case, the accumulation period covers the whole period of socialization. It follows that the transmission of cultural capital is no doubt the best hidden form of hereditary transmission of capital, and it therefore receives proportionately greater weight in the system of reproduction strategies, as the direct, visible forms of transmission tend to be more strongly censored and controlled.[71]

As is evidenced from these quotes, Bourdieu's view of the cultural world, as Georgina Born outlines in a detailed analysis, centers around "competitive position-taking [by] … actors engaged in cultural production."[72] Born emphasizes Bourdieu's key contribution is this notion of the "field of cultural production" as a spatial domain in which human competition is played out: "The field … is a structured space of possible positions and trajectories, a social topology constituted through the competitive yet complementary position-taking of rival actors."[73] Born also notes the potential issues with cultural production being understood only as a struggle between classes. Bourdieu was perhaps too caught up with specific issues of French class dominance and maintenance, that drove and was in turn maintained by the "unequal distribution of cultural capital." A more generalized interpretation of Bourdieu's theory based on diverse forms of dynamic and competition may be more palatable.[74]

In a more recent critique of Bourdieu, Antoine Hennion questions the idea that "taste is conceived only as a passive social game, largely ignorant

about itself." He instead elevates taste to "a form of presence in the world." Hennion examines the activities of amateurs who, with limited resources of cultural capital, still exhibit the capability to construct themselves through various relations with cultural material. Consider how this amateur record collector behaves toward his records:

> Before putting them away he used to leave many new records in the bottom right-hand corner of his bookcase—until the day he had the idea of transforming this disorder into the basis of a system for arranging his records. From then on he let them move upwards to the left, depending on the last time he had listened to them. This is a typical invention of an amateur: his record library gradually changed into a reflection of his tastes. The amateur triumphed over the musicologist: his taste, not the history of music, governs his system of classification.[75]

Hennion also emphasizes how interaction between individuals around taste exhibit aspects of cohesion and competition. Taste influences social relations and vice versa:

> The amateur's constitution of his taste and the practical methods used to develop it are based upon the recurrent presence of a mediator, an initiator. The example set by a recognized amateur is crucial here, like when an older opera fan corrects the prejudices of a younger one who still despises Bellini or Auber: "Wait a little; you'll see what you'll think of them later"; or when a techno amateur shrugs at seeing his young mate's records and lets him hear "the right thing," far from all this "commercial" stuff. There is no taste as long as one is alone, facing objects; no amateur knows from the outset how to appreciate good things, or simply what he likes. Taste starts with the comparison with others' tastes.[76]

Cultural competency also has a more overt manifestation, obviously, in the skills of those actively involved in, or seeking to be involved in, creative production. There is no doubt that skill, whether by nature or nurture, is critical to artistic success, above and beyond *mere* taste. But again, there may be more non-obvious connections between skill, labor, display, and the form that creative production takes. Philip Auslander, for example, suggests that individual musicians aim to overtly *display* their skill in various ways, and thus musical styles are not neutral but may be designed to draw attention to this skill accordingly. Such display might need to be exaggerated in order to communicate skill to audiences who, unlike other musicians, may not be expertly able to identify skill otherwise: "What counts, ultimately, is audience perception, not actual degree of difficulty."[77]

Although the tantalizing potential for Bourdieu's and others' thinking to be examined as a hypothesis of evolutionary psychology, this avenue seems

not to have been explored in any depth. In the language of evolutionary theory, such cultural competency, like one's accent, would be recognized as an honest signal of one's background, hard or impossible to fake, because the finer details would be detected by anyone who was truly from that background. In my earlier classification, such a Bourdieuian evolutionary theory would be seen as competitivist,[78] bearing similarities to sexual selection theory, but importantly, operating at the level of a culturally constructed niche that has little (but not nothing) to do directly with the attraction of a mate. We have seen previously how evolutionary explanations of social functions are complicated by the selfish gene principle. For an evolutionary explanation to make sense it has to apply to the inclusive fitness of the genes. Costly signaling theory states that in order for musical or any other art preference, capacity, or knowledge to work as an identity marker it has to be hard to fake. It has to be an honest indicator of your inclusion of that group, something that can only be acquired under certain circumstances: over a long period of time, from an early age, or with privileged access to the right people. The logic goes that you spend your time in one group, and this time is marked by the identity markers that you acquire—certain tastes, styles, ways of doing things, and knowledge. It is not easy to enter another group and fake one's identity.

This can apply to highly salient things like accents, dialects, gestures, and walks, which are acquired in a contagious, subconscious manner. It also seems to fit the profile of fashion items like clothing and music (Hennion's critique aside): they are complex and fast changing, requiring one to keep up with what's going on. Also, insofar as class is concerned (though the same may be said of other types of social group), any competitive process (for Bourdieu it is the struggle between classes) also comes with a complementary cohesionist component, requiring some degree of cooperation between members of the group. In the wider picture, as selfish gene theory would predict, the system also involves power struggles between individuals or families, *within* classes, which constantly dictate what cultural traits are considered desirable indicators. Thus in Bourdieu's words "cultural (or linguistic) competence … is acquired in relation to a particular field functioning both as a source of inculcation and as a market."[79]

Developing a more general social psychology of music, David Hargreaves and Adrian North list cognitive, emotional, and social functions for music.[80] Social functions are split among self-identity, interpersonal relationships,

and mood. Their research supports the central claim that we use music (and the same might go for fashion and other forms of art) to mark identity and to detect other people's identity. Their work demonstrates how, in doing so, it serves both a group cohesion function and its converse, the marking of individuals as others. "Musical preference acts as a 'badge of identity' during adolescence."[81] Respondents to surveys indicated significant perceived correlations between perceived musical taste and the perception of traits in others, such as the importance of physical attractiveness, an interest in deeper meaning, feminist views, vanity, and hedonism. Relative musical taste also elicited consistent responses to questions like "Would you make fun of that person?" and "Will they be successful later in life?" leading North and Hargreaves to note that "the effects of musical preference on the person's social standing seem to be associated with the extent to which the style in question was prestigious in the eyes of the participants." North and Hargreaves frame these effects as part of a wider pattern of in-group/out-group behaviors extending beyond music but interacting with musical taste. Subcultural distinctions can be subtle. In one study, North and Hargreaves compare attitudes to indie pop and chart pop, which are sonically very similar and closely related. Indie pop "is intended to be more difficult to comprehend, often requiring the listener to be more active in determining its underlying meaning … and previous empirical research has established that these two styles elicit quantitatively different aesthetic responses from adolescents."[82] In a separate study, music was also identified as being *explicitly* used by people to create certain impressions with others.[83]

Other studies have implied that aspects of music taste, in particular the breadth of music taste, corresponds to personality traits.[84] Those with wider music tastes were seen to be more confident and having a stronger sense of self, but the causality underlying this correlation is not identified. The frenzied engagement we have with music in our early teens and late twenties plays into Hargreaves and North's model of musical identity formation.

More recently, in collaboration with Emery Schubert, North and Hargreaves have developed a theory that connects these social findings with neuroscientific models of associative memory. In this theory, associations are formed between musical style and perceptions of groups. Musical preference, in turn, is both influenced by and influences perceptions of identity. The "influenced by" part is particularly important. At some level we may strategically (albeit subconsciously) manipulate what we should like, but also

our group membership will simply dictate what music we are exposed to. Schubert, Hargreaves, and North propose that these associations are formed in lived experience: "When playing music at a campfire with friends for the first time, the network of associations with the environment (the campfire and atmosphere), the friends and the music will form new networks which represent the cooccurrence of the music and social context, and thus future experiences involving any or all of these components can lead to a large amount of activation spreading through the network at a subsequent activation involving any or all of those components (the campfire, the friends and/or the music)."[85] They also appeal to a widely studied potential for the contagion of mood ("emotional contagion") that would further support the ability for music to become an associative marker for positive and negative social experiences.

Schubert, Hargreaves, and North's theory is identified with prototype theory: people will like the most prototypical music that they hear, "that is, music that sounds most similar to their existing mental representations of musical styles and pieces."[86]

Theories of in-group and out-group associations have also been applied to the *creation* of artistic material, with an even greater overhead required to become an expert producer, as well as connoisseur. Edward Hagen and Gregory Bryant propose an evolutionary theory of music and dance as a "coalition signaling system," by which a group would use performance to signal their coalition strength and hence coordinated fighting ability to other groups. This, they propose, would be done through the development of complex dances that required skill and coordination and could not be learned easily. Thus a complex coordinated group dance would be an honest indicator of how long that group had spent together and how much time they had devoted to practice, at the same time as indicating individual traits like fitness and coordination. Again this would have both an internal cohesive component (the investment of time literally representing an investment in that group over another), an external competitive component (the use of cohesion to signal group strength), and would also be subject to internal competition (those individuals most capable of enacting the complex sequences would achieve greater success within the group).

Rhythm, in this context, could be seen as a structuring primitive that facilitated learning and coordination, upon which layers of rhythmic and melodic complexity could be built. It defines a space in which an exploration

of complexity can better take place, just as the discrete units that make up language do. According to such theories, music's complexity and diversity actually play a role in its social function. Hagen and Bryant's theory would require that dances are complex and original so that they are harder to learn and couldn't be transposed from one group to another. A similar expectation might exist in the context of Bourdieu's theory.

Kathryn Coe's ancestress hypothesis[87] takes a different but related stance, looking at the use of identity markers over long-term multigeneration lineages, applied to early evolutionary scenarios. She proposes that decorative designs and other cultural goods come to become associated with family groups, passed through lineages that lead back to specific ancestresses. Thus being born into a group immediately implies being born into an existing cultural identity, and the symbolic and decorative material that identifies the group is laden with the logic of group membership. "Mothers used art forms to anchor themselves and their kin to the father and his kin, and to promote the survival and reproductive success of kin and descendants. Individuals who abided by this strategy, accompanied by its strict codes of cooperation, left more descendants than did individuals who did not."[88]

In this case, again, we must consider not only one's attention to, and preference for, certain things over others as being driven through exposure effects but also being driven by identity and association—taste driven not only by immersion in a culture but by specific allegiances. If cultural competencies are so critical to getting by in culturally constructed niches, then rather than aesthetic preferences being influenced only by exposure to stimuli they should also be influenced by social factors that are external to the content itself. Whether you like something or not might depend on who made it, what has been said about it, or what other associations are made with it. This can include strategic or functional development of preferences based on social influences. Several experimental studies make the point.

Firstly, a bias has been observed toward preferring things that others prefer in general. In a famous study, Matthew Salganik, Peter Dodds, and Duncan Watts,[89] for example, looked at people's evaluations of pieces of music via an online evaluation task. They compared one scenario in which people rated pieces of music based purely on the content, with another scenario in which people were shown the current average preference ratings for each piece by other people in the study. They found that people were

strongly influenced by these ratings in the latter scenario, tending to align their judgment with the consensus. This results in a process of amplification in the second scenario. They observe that a piece of music still needs to have certain qualities to get liked in the first place, but once it becomes identified as a popular track, the positive feedback begins, and new evaluations are more likely to fall into line with the consensus. This means that once again luck plays a significant role in the success outcomes of creative work, this time amplified. It is a familiar idea that a talented artist just didn't quite make it because the buzz around their work just wasn't there, or similarly, that a widely praised artist was not essentially more talented than their peers but managed to capitalize on their luck and get the best out of their chances. An interesting aside is that algorithms have this same problem or can confound it by reinforcing this winner-takes-all effect. In automated recommendation systems, it is common to make recommendations based on popularity, and so new creators struggle to get a foothold, and unrecognized genius remains unrecognized. This is known as the "cold start" problem in recommender systems.

This idea is bluntly illustrated by a notorious experiment in which the acclaimed violinist Joshua Bell busked on the New York subway, dressed casually, performing repertoire that he performs frequently to sell-out audiences. According to a video of the episode, only one passerby paid any particular attention to his performance, and it turned out that person had seen him play the night before and recognized him. As the evolutionary psychologist Paul Bloom notes, "This experiment provides a dramatic illustration of how context matters when people appreciate a performance. … It is a clever demonstration, but perhaps not surprising. Everyone knows that the value of a painting shoots up if it is discovered to be by a famous artist, and plummets if it is discovered to be a fake. … Origins matter."[90] This phenomenon is a manifestation of Boyd and Richerson's frequency bias discussed in chapter 3, where we look to the surrounding population to make decisions about what we should be paying attention to.

Building on this is the idea that the prestige of an artist can in some cases be the more pivotal factor in the evaluation of a work than the content of the work itself. Again, we can appreciate from anecdotal experiences that this may be the case. A famous artist might produce something seemingly monstrous and flawed, but our attention is already drawn because they are a person of interest, the work cannot simply be ignored or cast off, and there

is much space for us to grow to like it. Questions such as "What is the significance or deeper meaning of the work?" or "Why has that artist chosen this subject?" might obscure the cruder and forever problematic question, "Is it good?" which would be thrown more easily at the wannabe beginner.

Bloom and colleagues[91] have studied a range of scenarios that attempt to isolate exactly where people derive value in the judgment of work, with a particular focus on those confounding factors that lie beyond the material properties of the work. One manifestation of this research is the question of why we think authenticity or originality is important in art. For Bloom, while our pleasure response is partly grounded in crude biological functional systems—the love of sugar, sex, or building a good model of the environment—it is always also heavily influenced by higher-level conceptual thought that is grounded in sociocultural concerns. His pithy contribution to the debate over what makes humans unique, then, is that *we like Tabasco sauce*. That is, we have the capability to like (or learn to like) things that we shouldn't like, things such as spices that on first experience cause us pain. According to Bloom's framework, the love of authenticity is a simple manifestation of the desire to convene with the *essence* of the artist, just as we seek autographs or pay to see our heroes in the flesh. An authentic painting has been physically touched by the artist, and although there may be no DNA trace worth hunting for, this is still meaningful to us.

It is a bizarre suggestion that we would care so deeply about such forms of essentialism, and it is unclear why we should. But Bloom's thesis of essentialism is applied convincingly to myriad cases: "Even a pleasure such as the satisfaction of hunger is affected by concerns about essence and history, moral purity and moral defilement. There is always a depth to pleasure." In each case that Bloom considers, we seem capable of reconfiguring our pleasure response, either on the fly (as with Salganik and colleagues' experiment, where additional information influences the evaluation there and then) or as our taste develops over time.

Bloom's essentialism hypothesis, as applied across different areas of pleasure, has a great depth of application. In one sense it fits well with those theories above such as Huron's and Schmidhuber's, where the pleasure response is associated with the task of building a good model of our environment. Essences in this sense are simply a manifestation of our categorical and formal understanding of the world. They inhabit the linguistic conceptual and symbolically grounded end of the spectrum, away from

lower-level qualities like harmonic content or structural complexity, all of which come together to form a holistic and multidimensional relation to cultural goods. When we think about things or people, we structure them according to various essential concepts that we believe them to be attached to. This can be as simple as tracking a single object while distinguishing it from other objects or seeing it undergoing transformations, something so intuitive that we might not notice that this requires an underlying mechanism. Bloom explains:[92]

> We can think about cherries as individual things. You can easily imagine a pair of cherries in a box, each soft, moist, red and tart, but you know there are two of them, not one. And this is not because we are merely sensitive to the magnitude of the properties—anyone can tell the difference between two small cherries and one big one. You can easily track an individual even if its properties are unstable, as when a caterpillar turns into a butterfly, or a frog into a prince … And if one takes a cherry, paints it green, injects it with salt, and freezes it solid, it now has none of the standard properties … but it doesn't *disappear*; the individual lives on even though its properties have changed.[93]

Bloom's essentialism can also be understood as serving a role as part of a system of social function. These essential entities that guide our model of the world and our subsequent behavior are generally about social relations. According to Bloom's thesis, basic pleasure functions such as a love of sugar collide, clashing and combining, with these more socially functional factors guiding pleasure. The contrast between these drives resemble the extremes of Maslow's hierarchy of needs, with physiological and safety needs at the bottom of the hierarchy and esteem and self-actualization needs at the top of the hierarchy.

Bloom's essentialism thesis is particularly pertinent when directly applied to people, roles, and relations. People have roles, status, and other values attached to them that influence how we locate them in the virtual constructions of culture. Shamans, priests, poets, and lawyers all have invisible properties attributed to them that affect how we interact with them, according to more or less formal rules.

It is timely to reiterate the plea at the beginning of this chapter to avoid thinking of art as something that fits a neat definition, is neatly bounded from the rest of life, or deserves a neat explanation. In all of the above scenarios there is an inescapable layering of contributing factors that is not easily reduced to a single simple narrative. For Pinker, musical pleasure is

a crafted hodgepodge that satisfies different evolutionary pleasure circuits. For Bloom, conflicts emerge between different forms of value, between levels in Maslow's hierarchy of needs, in the formation of taste. For Huron, fast and slow thinking systems feed back on each other to trigger different sorts of pleasure response to the satisfaction or breaking of expectations. Similarly, musicologist and evolutionary thinker Ian Cross identifies different levels upon which we interpret music: "The very low-pitched semitonal ostinato overlaid by a non-tonally related horn call at the outset of the film *Jaws*, overlaid on an otherwise fairly innocuous underwater scene, signifies to the listener/viewer that something big and unseen is out there in the water (only big things can produce low-frequency sounds) and that it may well be hunting (horn calls, in western culture, are conventionally interpreted in terms of hunting topics)—hence fear and perhaps terror may be wholly appropriate, and fairly universal, responses."[94] Here the literal size of the bass and the reference of the hunting horn invite a universal association and a culturally learned one, respectively. But for Cross, music is also successful in its social functions because of its great ambiguity of meaning.

Artistic Behavior as a Network Phenomenon

In chapter 3, I followed a path from thinking about social creativity as the sum of individual creative acts combining stochastically, to thinking about structured social processes supporting creativity, such as the constructive mythologizing described by Schaffer, or the design of urban environments to support serendipity and combinatorial creativity proposed by Landry.

Similarly, this chapter set off thinking about the effects of social learning on the formation of individual taste, looking at several social functions and strategies associated with taste, implicated in the structure of groups, to move on to looking at more substantial collective phenomena under two broad headings: generational effects, and network effects. When we think about styles, taste, genres or any other term to describe the qualities of creative artifacts, although the material in question may be intangible, loosely bounded, and ephemeral, it is safe to say that cultural production appears in clusters. Such groupings and their changes correspond to groupings and changes in the distributed creative systems of people and things, and vice versa.

These clusters of cultural practice can exhibit both adaptive and generative creativity depending on how concentrated each temporary grouping is

on a common goal, or how conflicting individual objectives are. It is important to emphasize the generative dimension here: aesthetic movements needn't have purpose in and of themselves, but may be both the emergent results of large numbers of individuals who do have purpose, and the evolving context in which they must form goals, tastes, and actions.

One cultural dynamic scenario touched on already is a generational cycle. Given the hypothesized identity dynamics of individual life cycles, as discussed above with reference to Hargreaves and North's work on youth music cultures, it would be natural to witness cyclical effects. Hargreaves and North argue that individuals acquire musical taste at a key receptive age and carry that taste through their lives. Correspondingly, the coming and going of genre movements are not the result of entire populations shifting taste but of a rolling turnover of individuals arriving, acquiring taste, and carrying it through their individual lives. A genre might form amid the creative energy of a new youth movement, take shape, and evolve through stages as its key proponents mature.[95] Like oscillatory physical processes—water waves, drops forming on a tap, the circular movement of a lava lamp—feedback effects can transform a gradual evolution into something more bumpy or explosive, with visible oscillations emerging from the substrate of human action.

Martindale[96] offers one such cyclical model in which genres of creative artifacts follow a common life cycle, taking an evolutionary perspective. Central to Martindale's theory is the idea of habituation, which we have encountered in various forms above. The crux of habituation is that we do not sustain interest in things that we are used to. Martindale asks what effect habituation has on the long-term progression of a creative domain. In Martindale's terms, the arousal potential of stimuli is reduced over time, and we should expect that changes are made that sustain this arousal potential, else through habituation the desired impact of an artform would be rendered ineffective. An initial proposal is to simply increase the stimulus intensity, but Martindale rejects this on the basis that cannot be increased indefinitely. Instead, creative production takes the form of a sequence of inspiration and elaboration. Inspiration is associated with primary-process thought, which is "free-associative, concrete, 'irrational,' and autistic," whereas elaboration is associated with secondary-process thought, which is "abstract, logical, and reality-oriented."[97] Martindale proposes "regression" as a core concept of the evolution of artforms, which involve the development toward more primary-process thinking:

> Novel ideas could emerge in two ways from the inspiration-elaboration process: holding the amount of elaboration constant, deeper regression (movement toward primary process thought) should lead to more free-associative thought and therefore increase the probability of original or remote combinations of mental elements. In other words, to produce a more novel idea one could regress to a more primary process level. Holding degree of regression constant, decreasing the amount of elaboration should lead to statements that are original by virtue of being nonsensical or nonsyntactic in varying degrees.

He goes on to hypothesize that regression is the core basis for transformation, but that it arrives eventually at an unsustainable extreme, at which point a radical stylistic change takes place:

> Across the time that a given style is in effect, we should expect works of art to have content that becomes increasingly more and more dreamlike, unrealistic, and bizarre. … Eventually, a turning point to this movement toward primary process thought during inspiration will be reached. At that time, increases in novelty would be much more profitably attained by decreasing level of elaboration—by loosening the stylistic rules governing the production of art works—than by attempts at deeper regression.[98]

Specifically, Martindale predicts that properties such as novelty, complexity, and variability should increase monotonically over time, while "measures of primary process content should exhibit cycles of increasing and decreasing density of [elements] indicative of regressive thought."[99] Other results of studies conducted by Martindale[100] suggest that the tendency to habituate varies among people, with more creative people habituating more slowly than less creative people, leading him to conclude that the desire for novelty is stronger than boredom as a creative driving force.

This thesis sets up a situation where the timing of genre revolutions matters—the drip on a tap needs to build before it will drop. Related to this is the question of who leads the revolution. Pioneering incumbents or unencumbered novices? If the latter, then individuals bidding to enter a field may find themselves well aligned with a rare window of opportunity to be part of a revolution; timing is everything.

More generally, those entering the field at different times might find a spectrum of strategies on offer. Are you one of the complexifiers who wants to push an old genre into new domains of richness, or a minimalist who wants to carve out a new space? The specific time and place at which an individual enters this cultural battleground matters, and this enriches our view of the formation of taste and individual objectives.

Likewise, the various frequency-dependent effects we have encountered can amplify and thus help define genres. Individuals entering a field don't branch out at random into new territory but engage with a well-established genre and reinforce its existence. But what happens at the point of revolution, where there is no bandwagon to jump onto? Is the emerging genre predetermined, the logical next step that is on the tip of everyone's tongue, or is there a conceptual vacuum that is filled opportunistically? There is some evidence for either, and the story is probably one of a combination of more or less deterministic factors, where certain overall pressures do preference some movements over others. For example, Nia et al.[101] argue that the historical evolution of violin designs toward louder and louder violins—achieved by optimizing among other things the shape of the sound holes cut into the violin body—exhibited a BVSR pattern, with different European violin workshops being differentially selected for success based on loudness. Loud violins have simple evolutionary advantages, and even if some tastes prefer other qualities, the power to bring larger numbers of people together in larger auditoria, to sound impressive, or literally to outcompete other instruments would all be factors that might drive this evolution *despite* people's preferences.

The same is true of recent history's notorious radio wars. Mix and mastering engineers can employ a range of techniques to get their tracks to sound louder than others, and they may do so for the success of the track *despite* their taste. Indeed, a number of these techniques degrade the quality of the music, not least by limiting its dynamic range. These examples may seem peripheral to the *real* creation of music but are just as relevant forms of creation and greatly influence the context in which music is made.

All artistic revolutions or gradualist transformations might be given such a narrative, but as with evolutionary theory, just-so stories can also be a seductive trap, and historical contingency, chance, and the potential arbitrariness of cultural goods may also play significant roles. Against a backdrop of relative technological and social determinism, the idea that the arts enable a more turbulent happenstance path through history, one that is even dictated by the free will of individual human agents, is a picture that resonates well for many people.

Using big data analytics, more recent studies have attempted to formalize the relational interactions between members of a field by studying network properties. We can trawl online databases to discover not only the

properties and success of cultural artifacts but the networks of relations that connect the producers of these artifacts. Gino Cattani and Simone Ferriani[102] use social network analysis to be able to determine the "position" of individual creators in a field. For any given social network we can define a centrality property and its converse: peripherality. Social movers and shakers are central nodes in social networks. Peripheral individuals have less influence.

In network analysis, there are various ways to define centrality, but a common one is to measure, for each node in the network, what the average distance is to all other nodes (the distance is measured as the shortest number of connections one has to follow to get to that node). The node with the shortest average distance to all other nodes is considered the most central, while the node with the longest average distance to all other nodes is the most peripheral. If a network looked like a spiderweb, then the node at the physical center of the web would also be the most central node, and those around the outside would be considered the most peripheral. Social networks rarely look like spiderwebs, more like fur balls. If a network were fully connected, with every node connected directly to every other node, it would reveal no difference between nodes, and there would be no structure to speak of.

Naturally occurring social networks lie somewhere between spiderwebs and fully connected networks. They are typified by a "scale-free" topology, where a small number of nodes have a large number of connections, and most nodes have a small number of connections (such networks also have a fractal nature). For example, a town mayor or innkeeper has daily interactions with a lot of people, is known about, and exerts influence in various ways, whereas the town's factory workers interact with a small number of people and have little influence.

Cattani and Ferriani's[103] analysis of the IMDB movie database attempts to discover how the creativity of individuals and teams within the film industry relates to their position in the social network (defined as the set of connections between individuals based on their collaborative history). Specifically, in exactly the same way that the Wundt curve dictates a peak of interest between things that are too familiar and things that are too unusual, the hypothesis that is considered is that a peak of creativity is found among individuals or teams who are placed neither too centrally (the old guard, the incumbents) nor too peripherally (the young Turks, the

insurgents). Creativity is measured here in terms of markers of prestige such as awards, also automatically gleaned from the database, rather than total sales or popularity. Cattani and Ferriani find some evidence for this hypothesized relationship by demonstrating that the relationship between centrality and creativity, so measured, follows an inverted-U relationship, which resembles the Wundt curve.

In a following study, they show that in fact there are significant differences between different communities involved in the evaluation of creativity with respect to network position. Whereas peers rate the creativity of *central* figures most highly, critics are much more willing to rate members of the *periphery* more highly. In this work, Cattani and Ferriani closely associate their thinking with Bourdieu, indicative of a general resurgence of interest in this approach to creative domains, in their words "a rich and vibrant tradition building on Bourdieu's pioneering insights treats cultural producers as engaged in an ongoing struggle to secure notoriety, prestige, and esteem from colleagues."[104] In particular the focus on an individual's *position* within a social network enables a truly situated study of creative individuals' goals, which can connect these goals with global effects.

> Because … judgments produce prestige hierarchies and affect field evolution, cultural fields are in a constant state of struggle between established and emerging actors who compete for symbolic distinction based on subjective rules of merit, and the vested interests and social objectives these rules embody. While incumbents work to defend and reproduce their views and impose consensus, challengers try to "break the silence of the doxa and call into question the unproblematic, taken-for-granted world of the dominant groups."[105]

More generally, this ties their work to oppositional models that pit an establishment against a wider population of outsiders. An interesting additional force influencing this dynamic is the establishment acting to lock in certain aesthetic criteria, defining a genre. For example, this could be seen as the main guiding process in Coe's model of ancestral groups competing for identity markers. Other big-data studies have contributed to this fleshing out of a model of creative competition. Noah Askin and Michael Mauskapf[106] study the success of pop songs and show that successful songs must differentiate themselves from the field; overly typical songs are less likely to succeed. They propose that "the pressures toward conformity and differentiation act in concert."[107]

This analysis does not say anything of the generational relations between individuals, but it is reasonable to assume that in general the more central players will have been around longer, establishing their centrality, while the peripheral players will be newcomers, trying to develop their connections. We can therefore imagine ways in which cyclic processes such as that identified by Martindale track the life histories of individual people.

Summary: Shifting Dimensions and Emergent Niches

The material in this chapter does not aim to offer a complete and coherent theory of artistic behavior. I have set out to gather together the most relevant theoretical strands on offer, which are diverse and may not always be entirely compatible, and present them in a way that helps frame how the algorithmic automation of creative processes can fit into human creative domains. Good models of what humans are doing when they are engaged in artistic creation or experience are important if we are to understand what algorithms can do well, where they fit into human networks of creativity, and how they need to be designed to interact with people.

These strands are disparate and interspersed with speculation, and no commitment has been made here to argue for any one theory or another, or to particularly defend any of these various theses. But there is a common theme throughout: human artistic behavior lives within a complex web of social motivations and interactions, and any model of human artistic behavior is complete only insofar as it is capable of modeling any actor's embedded relationship with this social context.

Nevertheless, the following features of these disparate theories can now be condensed into a loose framework for thinking about artistic behavior and its hybrid, networked agency.

1. The cognitive mechanisms we use to build an adaptive model of the world and form predictions about the future—evolved largely in the service of nonsocial learning—form a critical part of the understanding of aesthetic behavior. A picture is emerging of a structured high-level cognitive architecture and a series of low-level components that adapt in response to experiences, as well as a reward system that stimulates us to engage in certain experiences over others. There are numerous,

competing evolutionary explanations for artistic competencies, but the majority of mechanisms are best described as consequentialist, having nothing to do with art as functional behavior in itself, but rather positing artistic behaviors as emerging simply because preexisting cognitive structures motivate them.

2. Meanwhile, a speculative *functional* role for aspects of artistic behavior, grounded in the dynamics of group identity and cohesion, is supported by evidence from various sources. The passively attained habitus of individuals, immersed in a given cultural background, is coherent with an honest-signaling model of competitive evolution. Artistic behavior can be seen to serve the manifestation of groups in various ways, setting up the conditions for competition between groups, cohesion within groups, and also additional competition *within* groups, possibly taking a nested hierarchical and diffuse structure (groups within groups, and overlapping groups). This view potentially *subsumes* the go-to sexual selection theory of artistic behavior with a more comprehensive social selection model. It has roughly the same basic properties and implies at least one similar outcome: that individuals are in part attracted to one another based on tastes and competencies, as markers of identity.
3. Certain factors underlying aesthetic preference and artistic competencies *may* have other more specific evolved functions, such as the dedicated evolution of preferences for landscapes or stories or the ability to track beats and metrical patterns. If these theories happen to be validated in the future, they are still very likely to be isolated traits that do not explain artistic behavior in any depth. A possible exception, however, is that musical competencies, in particular rhythmic perception, may have evolved in a protolinguistic stage associated with the social function discussed in the previous point.
4. Artistic evaluation is necessarily multidimensional and dynamic. Interaction with any aesthetic object cannot be simplified to a single preference measure or given set of features. New information, including hidden (possibly in Bloom's *essentialist* sense) attributes of an object such as its provenance or popularity, can affect how we perceive it. As part of our learning about the world, we encounter radically new ways to think about things, from learning the rules of counterpoint to understanding the social relations that underlie the production of dance music remixes.

5. Bringing together the high plasticity of social learners, the functionality of artistic behavior with respect to group identity, and the related in-group and between-group dynamic factors causes certain social network and generational effects, such as cyclical patterns and roles assumed by individuals according to their location in social networks.

Bringing this all together with the concepts developed in chapter 3, we can characterize the relationship between creative activity and the creative domains themselves as a form of ongoing cultural niche construction: fluid and ecosystemic. A child is raised into a certain cultural environment, their plastic brain learning associations between styles and cultural artifacts and group identities. They may become a competitive agent in this creative landscape, seeking to successfully innovate and make their mark, if not as a cultural producer then as an active consumer who communicates and interacts with others through cultural consumption. Depending on their context—their specific social background, abilities, and timing—they may go mainstream or be radical, with various risks and payoffs associated with different strategies. They may seek alliances, perhaps in the form of artistic movements. All the while, their brains and the brains of those around them are adapting through ongoing exposure to aesthetic, cultural material. New associations are being formed, and their worldview might evolve rapidly through a cycle of feedback, interacting in a tight clique. The occupants of a subcultural niche will have trained themselves in a different pool of aesthetic data to the mainstream, drawn further into this constructed niche of domain specificity by pattern-hungry minds. Not only will this influence their judgment, but it will actually enable them to perceive things that others may not perceive, nuances and markers of value that only the trained can properly identify.

Specifically, something that was never before a factor of any great importance might suddenly become the new axis along which difference is perceived: the amount of swing, say, in a guitar lick, or the linearity of the unfolding of a film's narrative. This something may not be obvious to people, but that makes it all the more effective as a marker of identity, in the way that Bourdieu implied. For this reason, aesthetic evolution can be open-ended and constantly creative and radical.

Howard Becker's[108] sociological analysis of various art worlds outlines many examples of how such dynamics play out. Becker's emphasis is on the practical and the contextual minutiae that sustain forms of activity in

an art world. He notes that certain Western art music movements are often highly obscure and certainly have little mainstream appeal. Instead, they rely on the specific subcommunities of music and art students for whom obscure experimentation is essential creative nutrition. Even more pragmatically, when describing how art must be made for art galleries and art gallery culture, it must, at the very least, fit through the gallery door.

This impressionistic sketch of the social dynamics of creative domains may not be particularly detailed or concrete, but it should suffice as a sketch that is easy enough to hold in the back of one's mind as we now move on to consider what is happening in the world of computationally creative algorithms.

The open access edition of this book was made possible by generous funding from University of New South Wales.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Bown, Ollie, author.
Title: Beyond the creative species : making machines that make art and music / Oliver Bown.
Description: Cambridge, Massachusetts : The MIT Press, [2021] | Includes bibliographical references and index.
Identifiers: LCCN 2020017931 | ISBN 9780262045018 (hardcover)
Subjects: LCSH: Technology and the arts. | Arts—Technological innovations.
Classification: LCC NX180.T4 B68 2021 | DDC 700.1/05—dc23
LC record available at https://lccn.loc.gov/2020017931

10 9 8 7 6 5 4 3 2 1