

The Artful Mind

*Cognitive Science and the Riddle
of Human Creativity*

EDITED BY MARK TURNER

OXFORD
UNIVERSITY PRESS

2006

I

Art and Cognitive Evolution

Merlin Donald

This chapter offers an overview of the cognitive principles of art, the origins of art, and the cognitive function of art. Art is an activity that arises in the context of human cultural and cognitive evolution. Its sources include not only the most abstract integrative regions of the brain but also the communities of mind within which artists and audiences live. The interaction of these sources creates complex cultural-cognitive domains, which are reflected in art. Art and artists are active players in the co-evolution of culture and cognition.

In this chapter, I use the word *art* to refer to a wide class of expressive forms and media, including music, dance, theater, various multimedia categories (such as opera and cinema), painting, sculpture, aspects of the built environment, and architecture. The word can reasonably be extended to include most forms of written literature. I do not include any of the broader applications of the word *art*—as in, for instance, the art of mathematics, engineering, baseball, or carpentry. It may be said that there is an art to performing virtually any activity elegantly or well (including art: there is an art to good art, one might say), but that is another matter. Here I am concerned with the origins and functions of artistic forms and media themselves, rather than with issues of artistic creation, merit, beauty, or transcendence.

What cognitive principles govern art? And where should we begin a cognitive exploration of its origins? There is no consensus on this, but a few guidelines might help establish the territory to be explored.

(1) Art should be regarded as a specific kind of *cognitive engineering*. As a first principle, art is an activity *intended to influence the minds of an audience*. It involves the deliberate construction of representations that affect how people (including the artist) view the world. This reflects a very deep human tendency for the reciprocal control of attention, which carries with it a propensity to deliberately *engineer* the experiences of others (especially of our own progeny and peers). Joint and reciprocal control of attention is the foundation of human social communication; just as parents guide their children's attention to certain aspects of the world, most artists attempt to control their audience's attention, leading it by the hand, so to speak, into a carefully engineered experience. To achieve this, the artist must be an effective pedagogue, anticipating the audience's reactions (this principle applies even if the artist wants to elicit an apparently unpredictable result, in which case, of course, uncertainty itself is engineered into the outcome).

(2) Art is always created in the context of *distributed cognition*. Human cultures can be regarded as massive distributed cognitive networks, involving the linking of many minds, often with large institutional structures that guide the flow of ideas, memories, and knowledge. Artists are highly placed within these cultural-cognitive networks, often serving as the creative engine that drives much of the enterprise. They influence the cognitive activity of their particular tribe or generation (for artists, like everyone else, are situated in space and time), both by preserving and by modifying its symbols, images, and other expressive forms. In a sense, they are one with the network: they derive their most basic ideas and techniques, as well as their inspiration, from it, and must operate within the limitations it imposes.

(3) Art is *constructivist* in nature, aimed at the deliberate refinement and elaboration of mental models and worldviews. These are the natural products of cognition itself, the outcome of the brain's tendency to strive for the integration of perceptual and conceptual material over time. The term *large-scale neural integration* refers to the nervous system's cross-modal unification of many sources of experience into a single abstract model or percept. The canonical example of this kind of integration is event-perception, which can unify a blur of millions of individual sensations of sight, sound, touch, taste, smell, and emotions into unitary event-percepts. This ability is very limited in simple organisms, where the "stimulus" of behavior is often an uncomplicated one-dimensional property, such as a pheromone or a color, but it is common and very highly developed, in most social mammals and especially in human beings, where it has evolved into a very abstract capacity to integrate not only the raw materials of experience but also the constituents of memory itself.

Thus, a dog is able to understand complex social events, such as “begging” behavior or “submission,” which involve socially relativistic percepts that unfold over time. Humans, of course, navigate much more abstract versions of social behavior, which culminate in worldviews that frame their interpretation of events. The Stoic, scientific, Puritan, and Romantic worldviews share a basis in the need to achieve abstract integration of smaller events. Such worldviews are collective, or cultural, products of the inherent drive toward integration.

Large-scale integration might be regarded as the major adaptive advantage conveyed by the complex of special brain capacities often labeled *conscious processing* (Donald 2001). As the nervous system’s capacity for conscious processing evolved, selected species achieved increasingly more abstract kinds of cognitive integration, which gave an accordingly wider temporal and spatial range to their behavior. Hence these species’ ability to perceive distant, complex, and very abstract events that occur in the social environment, such as changing alliances, whose complexity exceeds the perceptual capacities of simpler creatures. In humans, this constructive integrative capacity evolved into a communally shared capacity: human culture is essentially a distributed cognitive system within which worldviews and mental models are constructed and shared by the members of a society. Artists are traditionally at the forefront of that process, and have a large influence on our worldviews and mental models.

(4) Most art is *metacognitive* in nature. *Metacognition* is, by definition, self-reflection. Art is self-reflective. The artistic object compels reflection on the very process that created it—that is, on the mind of the artist, and thus of the society from which the artist emerged. Ultimately, art derives from the innate human capacity for self-observation. That is why art has been so instrumental in defining cultural periods and in providing tribes, of whatever size and complexity, with their self-identifying symbols and allegories. Art is thus inherently metacognitive in its cognitive function on both the individual and social levels. Though the term *metacognition* customarily refers to individual self-reflection, I use it especially to denote art’s crucial role as a collective vehicle for self-reflection and as a shared source of cultural identity.

At various points in human cultural history, artists and writers have built comprehensive metacognitive systems that served to reflect on society and human nature; typical examples are the complex pictorial representations of knowledge so common in medieval European alchemy, and the multitude of very large Italian paintings that tried to sum up the conventions of Renaissance social order. These artistic objects reflected the predominant mental models and worldviews of those societies back to their members, and placed

artists in a position of considerable metacognitive influence, even though they derived their material from the society itself. The power of the artists arose because they often subtly (and sometimes not so subtly) altered the prevailing images and worldviews of their societies in a highly selective manner. The worldviews of communities have often been permanently changed through the efforts of a single artist (e.g., Verdi's revolutionary impact on nineteenth-century Italian politics). On such occasions, art sits high in the hierarchy of cultural-cognitive governance. Traditional religions have long recognized (and, consequently, relied on) the cognitive influence of art. Much the same can be said of modern secular states, such as Maoist China, and of modern corporations. The social-reflective role of art has always been controversial. But the ferocity of the arguments revolving around this topic testifies to the fundamental nature of art's contribution to the collective processes of thought, memory, and perception in society. This contribution is evident in the art of Christianity, Buddhism, and Islam, which conveys highly formal, integrated worldviews. It is also evident in the chaotic and fluid imagery of modern secular society, which conveys many different worldviews.

(5) Art is a *technology-driven* aspect of cognition. Although it may have begun as a natural expression of our collective need to represent reality, the media of artistic expression affect what can be represented, and these media differ tremendously between societies. The effect of technology on art is far-reaching. Technology affects the kinds of cognitive networks artists can construct, in part by setting limits on the kinds of ideas and images that can be represented and created. Major works of art constitute a crucial part of society's attempt to engineer, manipulate, and reflect on its own experience and occasionally to fabricate *de novo* its defining ideas and images. In historical context, technique and technology are central in defining what artists do and what choices they can make. Moreover, technology can actually alter the properties of the distributed cognitive systems of society and change the nature of the cognitive work that is done.

(6) The role of the artist, viewed as a component in a distributed cognitive system, is *not necessarily fixed*. As the system goes, so goes the role of art—and, indeed, the very definition of art. Elsewhere (Donald 1991, 2001), I have argued that symbolic technology (including the many technologies involved in making art) can deeply affect the architecture of cognition, both inside the head and outside, in the social network. In particular, such innovations as writing systems, new graphic media, and external memory systems can change the kind of art, and the range of worldviews, that are possible *because they influence memory itself, through both the media of storage and the pathways of retrieval*. Symbolic technologies ultimately enabled Brunelleschi to

build the dome of Santa Maria del Fiore in Florence. Similarly, they enabled Rodin to conceive of, and cast, his bronzes, while setting limits on what he could represent. Technology often determines the parameters of thought and creation (mathematical thought is a particularly clear example of this—mathematics is all about finding the right set of symbols to capture an idea).

This point has been largely missed in cognitive theories of art. When one is dealing with a distributed network of many individuals linked together, rather than an isolated individual, as a major source of creativity, the properties of the network, particularly those of network memory, become highly relevant. These are typically affected much more by technology than by the properties of biologically defined memory in the individual, which are largely fixed in the genome.

(7) Art is always aimed at a *cognitive outcome*. The conventional engineering of, say, a bridge or a drug compound is aimed at a specific physical outcome. In contrast, art is aimed at a specific cognitive outcome. It is designed to engineer a state of mind in an audience (even in cases of extreme narcissism where the only intended audience is the artist). The work is judged by its success in achieving this aim. Thus, in its ends, art is essentially different from other kinds of engineering, because its purpose is primarily cognitive. Cathedrals, and films, are specific kinds of cognitive machines. Their major social functions are cognitive: they influence memory, shape public behavior, set social norms, and modify the experience of life in their audiences. In these terms, the various techniques and media of art are a small but important part of the larger evolutionary trajectory of the human mind.

Art Viewed in an Evolutionary Context

Art is universal to all societies and unique to humans. Inevitably, when a phenomenon is both universal and species-unique, the question of its evolutionary origins arises. Within the reach of evolutionary theory, human evolution is special, and unusually complex, because it entails the co-evolution of biological and cultural forces. Art is central to that process, and one of the most interesting phenomena of human culture.

The cognitive domains of human cultural and cognitive evolution have emerged in three cascading stages, which I have labeled, successively, as Mimetic (~2 million years ago), Mythic (~150 thousand years ago), and Theoretic (last 2 thousand years, approximately) (Donald 1991, 1993, 1998a, 2001). These dates are only rough approximations; it is the sequence, rather than the specific dates, that is important. The progression is cumulative and

conservative, with each preceding stage remaining in place, and continuing to serve its specialized cognitive function in human society, as each new stage emerges. Even though art is a relatively recent development in the long history of the human species, it has an investment in all these cognitive domains, and its many forms reflect the very rich cognitive accumulations of human culture. Indeed, in many instances art has been a major factor in evolving these domains, and constitutes our primary evidence in determining the nature of prehistoric culture.

Because evolution is conservative, the modern mind retains all previous stages within its complex structure. The Mimetic domain (of which I shall have more to say later) comprises gesturing, pantomime, dance, visual analogy, and ritual, which evolved early and formed an archaic layer of culture based mostly on action-metaphor. Mimesis allowed for the spread of tool-making technology and fire-tending, through imitation and ritual. It also set the stage for the much later evolution of spoken language.

Mythic culture is based on spoken language, and especially on the natural social product of language, storytelling. Most societies have a specific subset of stories that acquire the status of myths, and these play a governing role in defining how to behave in a given culture. Myths also preserve notions of authority, gender, and morality. Mythic culture retains a subsidiary mimetic dimension, manifested in ritual, costume, and gesture, which are epitomized in various forms of art. We might even say that the mimetic dimension tends to fall under the governance of myth; thus the art and ritual of Christian civilization have been greatly concerned with the mythic content of that civilization. The same applies to Islamic, Jewish, Buddhist, and Hindu art. Traditional religion has often been the core institution for the regulation of what might be called "high" Mythic culture, and art has fallen under that kind of regulation in many societies.

Theoretic culture is a more recent historical development. It started very slowly, with the emergence of sophisticated writing technologies and scientific instruments, and then, after a long gestation period, became (somewhat) dominant in Western society after the Enlightenment. Over the past few hundred years, however, it has evolved very rapidly. Theoretic culture is symbol-based, logical, bureaucratic, and heavily dependent on external memory devices, such as writing, codices, mathematical notations, scientific instruments, books, records, and computers. It is the culture of government, science, and technology, and of many forms of art. In a global context, relative to the influence of the Mimetic and Mythic domains, Theoretic culture is still a minority culture. However, it is disproportionately influential because of its place in the distributed cognitive systems that determine such things as our

collective representation of the past and our tribal and class identities. Of necessity, even Theoretic institutions retain a Mimetic and Mythic element; human society cannot function without these more basic forms of representation, which carry out specific kinds of cognitive work. Whereas Theoretic modes of thought are dominant in planning, science, technology, and government, Mythic and Mimetic forms continue to dominate the vast majority of human transactions, including those that take place in the political and interpersonal domains.

Even though art is a relatively recent development in the long cognitive history of the human species, its forms reflect all these cognitive and cultural domains. The diversity of art, and its modern proliferation of forms, reflect the rich historical background of modern cognition and culture. Table 1.1 illustrates this point, by mapping various current artistic forms onto the proposed major domains of human cultural-cognitive emergence.

Note that this process is cumulative and scaffolded. By implication, the breakthrough adaptation, and the one from which all else that is distinctive about the human mind follows, is mimesis. The strong form of my hypothesis about art might be phrased as follows: the new is always and inevitably scaffolded on the old, and as a result, art is ultimately a reflection of the

TABLE 1.1

External Form	Cognitive Domain
Pantomime	Whole-Body Mimetic
Prosody, Chant	Vocal Mimetic
Most Rituals	Whole-Body and Vocal Mimetic
Acting, Body Language	Facial, Vocal, Whole-Body Mimetic
Costume, Dress, Makeup	Technologically Amplified Mimetic
Most Styles of Painting	Visual Mimetic
Sculpture, Crafted Objects	Visual, Tactile, Kinematic Mimetic
Popular Music	Auditory Mimetic
Oral Storytelling	Linguistic/Mythic
Epic Oral Poetry	Linguistic/Mythic
Lyric Poetry	Linguistic/Mythic
Novels, Other Extended Narratives	Linguistic/Mythic
Traditional Architecture	Mimetic/Mythic
Comic Books, Cartoons	Mixed Mimetic/Mythic
Formal Public Ritual and Spectacle	Mixed Mimetic/Mythic
Cinema, Opera, Theater	Mixed Mimetic/Mythic
Modern Architecture	Mixed Mimetic/Mythic/Theoretic
Modern Painting	Mixed Mimetic/Mythic/Theoretic
Modern Poetry and Music	Mixed Mimetic/Mythic/Theoretic

deepest and most ancient form of human expression, mimesis. This hypothesis is discussed further in a later section.

Art, Neuroscience, and Distributed Networks

Before embarking on this section, I should offer a caveat about the uses of neuroscience in this kind of very broad cognitive theorizing. All things cognitive—and art is no exception—are ultimately products of brain activity. It may seem to follow that, to understand art, we need only track its origins to some specific brain structure or function, such as the neural systems underlying human vision or human mimetic capacity. While there is undoubtedly some truth in this, the situation is not so simple.

When we speak of the mind, we usually invoke a theoretical entity called the *cognitive process*, which can be broken down into various component functions, such as perception, working memory, spatial attention, lexical search, episodic recall, and so on. Any complex mental task, including the production and viewing of art, is made up of chains of these cognitive components, arranged in functional architectures, or operational hierarchies that resemble the algorithms of computation. A major objective driving what might be termed the cognitive deconstruction of artistic experience is to analyze the functional architecture of its underlying component operations.

The act of looking at a painting, for example, might be deconstructed into a series of very brief components, each of which produces a “glimpse” of the object. These components include such things as moving the eyes, fixating and focusing them, processing the fixated image, storing that image in some form of temporary, or buffer, memory, and synthesizing the whole series of remembered images into a unified perception of the painting. This perception might then be subjected to further scrutiny in working memory. The sequence might be repeated and reflected upon many times before the viewer acquires any “expertise” or familiarity with the painting. This process provides the higher interpretative centers of the brain with multiple frames spread out over time, much like a cinematic sequence. This is true even if the object is a static thing, such as a sculpture, because such objects are always viewed in several glimpses taken over time, from various distinct fixations from different angles and distances.

It is evident that this type of cognitive sequence, which is typical of everyday cognition as well as of the experiencing of art, entails a complex and somewhat idiosyncratic series of brain operations. Some of the neural activity that drives these operations (to date, only the most elementary ones) can be

observed by electrical recording and brain imaging (see, for example, Zeki, this volume). Predictably, most works of art activate many brain regions and engage a variety of neural resources, depending on the modality of the artistic medium and the type of representation offered. Every creative or interpretative act, regardless of its input modality or conceptual demands, can be broken down, or deconstructed, in this way, into its neuro-cognitive “atoms.”

In every case, these will translate into a series of elementary brain operations that unfold in a complex sequence. The sequences will be quite different for various kinds of cognition, and for dissimilar artistic media, but the component operations will be basically alike. These complex sequences can become habitual and automatic. Thus, my reaction to one of my favorite paintings, Gustav Klimt's *Hope 1* (located in the National Gallery of Canada, Ottawa), always follows a familiar course: my gaze starts in one of a few possible places, and moves around the painting in a fairly predictable order, with emphasis on certain key features. These features lead me to a certain state of mind, and elicit memories which govern how I see the painting. This is a well-studied aspect of visual perception, and involves little or nothing by way of operations that are unique to the artistic experience.

The uniqueness of the artistic experience produced in my brain by that painting can undoubtedly be traced back, if not to the elementary components in the sequence, then to the high-level neural consequences of the sequence of meanings and associations uniquely triggered by the painting. Such sequences, which I have referred to elsewhere as *Condillac sequences* (Donald 2001), lead to, and sustain, the cognitive endpoint of the artistic experience: a unified state of awareness that such a work of art ideally sets up in my (or any viewer's) mind. Unfortunately, neither brain imaging technology nor neurobiology has solved the problem of how to measure, let alone model, these abstract chains of meanings or the specific states of awareness they induce. The technology to do this may come in the future, but it is not yet available.

However, the real limitation of this approach is not our lack of knowledge about the physical basis of Condillac sequences, or states of consciousness; presumably it will eventually be within our powers to advance our knowledge in these areas. A more serious long-term limitation of any strictly neuroscientific solution lies in the fact that the common component processes of experience in the nervous system are not the only drivers behind the experience of art. It may be argued that the most important drivers are largely cultural, or cognitive-cultural, and depend not only on what is experienced, but also on interpretative algorithms that may be peculiar to individuals or societies and have no fixed neural instantiation. These algorithms are embedded in the “distributed” cognitive processes of social networks.

Distributed networks constitute a higher level of cognitive control that exists on a social level, and exerts a tremendous pull on the minds of the individuals in the network. It combines the memory storage capacities of many brains with whatever memory technology a given social network has at its disposal, and weaves these into a cognitive system that extends far beyond the individual brain. Within such a system, the location of memory itself is problematic. Memory can reside anywhere in the network. Perceptions can emerge and undergo major transformation anywhere in the network. Representations become a shared resource, and the sources of creative change can be found in many different locations at once. By definition, the neural component of distributed cognition is almost impossible to track down. Moreover, it is not clear that tracking down the neural responses of participating brains would extend our understanding of network-level cognition itself, except perhaps by clarifying the nature of the interface between brain and network.

Large distributed cognitive networks, such as those commonly found in corporations, can achieve cognitive objectives that exceed the capacity of individual brains. This is especially true of memory retrieval and storage, but it also extends to thought and representation. Distributed cognition can exploit the specialized talents of individuals by combining them into a collective cognitive organ; in theory, such an organism has at its disposal all of the relevant capabilities of an entire population, plus whatever additional cognitive power technology can contribute to the system. In other words, the system prevails, and even the most brilliant individual's intellectual contribution will be judged by the standards of the system itself.

But even the distributed system is not the ultimate arbiter of artistic experience. There is an unpredictable, sometimes rather quirky individual contribution to the interpretation of any work of art, and despite careful crafting by the artist, a work of art itself can never be entirely in control of the neural end-state it produces in a given recipient. Individual memory is so complex as to become unpredictable in practice, and it is the way Condillac sequences are juxtaposed in the memory of the individual viewer that will ultimately lead the viewer to a specific end-state. It is unrealistic to expect that a common pattern of neural processing will ever suffice to "explain" or predict individual reactions to art.

Artists might insist that the main driver of artistic experience is the engineer of that experience—namely, the artist—and this holds partly true. Certainly, the way the artist manipulates events so as to set up an end-state in the minds of the audience starts the process running, and some technique (such as those of film) can be extremely compelling in controlling the audience's experience. But the brain might deconstruct the world presented by the

artist in many different ways, and through many different paths, while the goals and methods of the artist are largely set by larger social-cognitive networks that are distributed. The artist controls only a fraction of this process.

The major underlying challenge for cognitive science is not to discover all the possible cognitive processing paths by which artistic experience comes to be; that would be impossible, and pointless. Such an endeavor would not be unlike a particle physicist's trying to track every electron in, say, a roomful of people at a cocktail party. Why would one want to do this? It would explain nothing about cocktail parties or people. Nor would such an analysis explain a work of art. It is the very source of art-based cognition we should be chasing here, and that source will not be found in either the brain or social networks, taken by themselves.

Therefore the relevant research question is: What question should we be asking of art with empirical brain research? One answer might be: We should ask how art has historically enriched or modified the cognitive processes of human beings, both individually and collectively. To a cognitive scientist, art represents a singular, rather peculiar way of knowing the world. Art attacks the mind, not usually through its logical or analytic channels, but more commonly through its senses, passions, and anxieties. Under the distant guidance of the artist, the brains of the viewers gather the disparate pieces of evidence placed before them, while they draw on their own experiences to reconstruct the artist's intent. The challenge for the scientist is to interpret the cognitive source of the audience's perception of the worldview intended in the work. This can rarely be reduced to the solving of a simple static stimulus, or to any moment frozen in time. It almost always entails the integration of many complex perceptions over many viewings. Such interpretations are inherently dynamic in nature, and mostly, they engage large-scale neural integration over time.

This is done by an unknown integrative process, in what we euphemistically call the "higher regions" of the mind, where the work is ultimately interpreted. In terms of the laws of higher neural processing, we have no idea how this final step is achieved. We know much about the neural principles underlying such processing, and we know roughly which geographic regions are involved, but we still have no adequate theory of how large-scale parallel neural networks can create such an abstract and detailed conceptualization of the world.

We do know, however, that many species have roughly the same elements of sensory and perceptual intelligence as we do, despite having produced nothing like what human beings call art. The basic processes of the nervous system are very similar in monkeys, apes, and humans, and the overall design

of the brain is virtually identical. The human brain is much larger than those of apes and monkeys in certain areas, but as far as we have been able to determine, it has no qualitatively new regions or features. This might tempt us to think that the primate brain is a good starting point for a cognitive theory of art, and there is probably some gold to be mined by such studies. However, this is a self-limiting strategy and cannot explain much about the interpretation of art, since it avoids the central question: What makes human so different?

The answer seems to lie elsewhere—and not entirely in the brain by itself. In the case of human beings, there is an additional factor that must be taken into account in explaining art: the distributed cognitive processes of culture. Human culture is uniquely cognitive in its function. Human culture is a marketplace of ideas and images, feelings and impressions. Indeed, it is a vast cognitive network in its own right. The cultural network introduces an entirely new element to human life: immersion in a cognitive collectivity, or community of mind. This is perhaps the primary source of the enormous cognitive differences between human beings and our closest genetic relatives. Monkeys and apes solve the world alone; we do not. Human culture is based on the sharing of mental representations, and we are tethered to that network. It allows us to achieve things that are far beyond the capabilities of an ape or, for that matter, a socially isolated human brain.

Artists may sometimes have the illusion of separateness, of isolation from society. But in reality they have always been society's early warning device. The best of them are connected, and more deeply enculturated than most. It follows that the sources of their creativity, although partly personal, are also public, outside the nervous system, in the distributed system itself—that is, in culture, which encompasses, but supersedes, the individual nervous system.

The Evolutionary Origins of Art

The various expressive domains of art correspond roughly to major stages in the cognitive and cultural evolution of the human species. In previous publications (Donald 1991, 1993, 1998a, 2001) I have argued that art is an inevitable by-product of mimesis—a primordial, and truly human, cognitive adaptation that occurred very early in hominid prehistory and became the signature feature of the human mind. Mimesis had enormous cognitive consequences on the group level, resulting in a characteristically human form of communicative culture that later increased its influence with the emergence of language.

Mimesis is an analogue or holistic style of thought that is more basic to our uniquely human way of thinking than language or logic. Indeed, on present evidence language and logic evolved much later, from a mimetic platform. Mimesis is a foundation skill that arrived early in evolution, and defined the human style. The components of mimetic cognition are present to some degree in primates, but are vastly more developed in humans. This makes mimetic culture a logical, but radical, extension of the primate mind. It remains an important force in human affairs, and produces such typically human cognitive patterns as ritual, skill, gesture, tribal identification, personal style, and public spectacle. It explains our irresistible tendency to imitate one another and conform to patterns of group behavior, especially group emotional expression. It sets the tone of human social life, and it is the ultimate driving force behind art, which might be viewed as the ultimate refinement of the mimetic mode.

Mimesis is an innate capacity, and its universality allows human society to function smoothly. Then again, the mimetic tendency to copy others and conform is also a potentially fatal flaw that might someday destroy the human race; but that is quite another question. If humanity had somehow managed to evolve language and symbolic thought without first establishing an evolutionary platform for it in mimetic cognition, we would have very different minds. And very different cultures.

What is mimesis? The easiest answer to this question is simply to list some of the behaviors it encompasses. The term *mimesis* describes a cluster of capacities that were made possible by a single neuro-cognitive adaptation. They go together historically because they share certain key neural components. The four central mimetic abilities are mime, imitation, gesture, and the rehearsal of skill. Human beings are uniquely good at these. Apes have some small degree of competence in these areas, and this strengthens the case that these capacities might have been subjected to selection pressure early in hominid evolution, primarily to improve our ancestors' ability to obtain a high-quality diet in a changing environment.

Mimesis seems to have evolved as a cognitive elaboration of embodiment in patterns of action. Its origins lie in a redistribution of frontal-cortical influence during the early stages of the evolution of species *Homo*, when the prefrontal and parts of the premotor cortex expanded enormously in relative size and connectivity. The cognitive significance of this lies in the fact that, in virtually all social mammals, the frontal regions are concerned with the control of action and behavior, as opposed to the posterior areas, which are broadly concerned with the elaboration of perception. The disproportionate expansion of frontal influence gave hominids greatly improved motor control. More

important, the expansion of the prefrontal cortex was crucial in improving conscious self-regulation and metacognition. This created a new metacognitive field, a greatly expanded and differentiated working memory, in which hominids could observe themselves as actors, and rehearse and refine whatever they were doing. This also gave them some ability to reflect on the cognitive process itself, and the option of deliberately reflecting on, and shaping, their own actions.

The latter point is worth some elaboration. Only human beings reflect on their own actions, and modify them accordingly. Human children pass large amounts of time in skill-related play—that is, in rehearsing and altering their own actions. For instance, they might spend an entire afternoon improving their ability to bounce a ball, skip stones, make faces, assume odd postures, or create novel sounds. No other creature does anything like this. Many species engage in play, of course, and innate skills need to be exercised frequently by developing organisms. But most species play in a stereotyped manner, and cannot generate truly novel patterns or engage in role-playing or imaginary games. It is as if their attention were fixed on the external world, and unable to redirect itself toward the internal world of action. That is a great limitation because it precludes what humans know as culture. If attention is exclusively outward-directed, then motor activity, generated internally, remains fixed and stereotyped. And this rings true when examining what virtually all other mammals can do. They appear much less self-conscious than humans. Their awareness is other-directed, not self-directed.

Mimesis is therefore the direct result of consciously examining our own embodiment, of the brain using its body as a reduplicative device. The cognitive engine of this expressive skill is a much more powerful working memory space, an inner theater where imaginary actors play with actions and expressions, and where the embodied self performs various possible roles in the social world. It is also a place where self-initiated actions can be judged, altered, and exposed to internal critical scrutiny. The outcome of this remarkable process is a characteristically human capacity for reenacting even in a nonverbal, gestural, fuzzy, quasi-symbolic manner. A child's simple pantomime of a tea party or bedtime is a good example. It is an imaginary *playback* that tries to reduplicate an aspect of perceived reality, but alters reality in the process. *Reality* does not in fact look anything like its putative reenactment, and every successive mimetic act in such a sequence will become another variation on the initial reenactment. The metacognitive part of the mimetic mind can reflect on this scenario, which can be altered until the child judges it to be *right*. Unlike the stereotyped play of animals, the details of such a performance are never fixed. Mimetic expressions, even the simplest

them, are inherently creative and somewhat arbitrary. Mimesis can produce a virtual infinity of specific forms, even in the simplest reenactment, charade, or pantomime.

Moreover, mimetic expressions can potentially engage any part of the body. Unlike the songs of birds, they are not limited to one sense modality. Rather, mimesis is truly amodal, and can map virtually any kind of event-percept onto virtually any set of muscles, using many different specific read-outs. This leads to flexible analogue motor expressions, or action-metaphors. I might normally indicate anger with my face and low-level voice modulations, but at a distance I can substitute larger body gestures and very different sounds to achieve the same communicative effect. In a boardroom I might limit my expression of the same emotion to polite finger-tapping or searing glances. The point is that a mimetic production is never limited to one set of muscles or one fixed set of expressive forms. Mimetic creativity is *domain-general* or supra-modal, and fully accessible to consciousness. It meets all the criteria for what Fodor called a *nonmodular adaptation* (Fodor 1983) because it can range across all the perceptual and motor domains given to the actor's awareness. It creates a very abstract mimetic *mapping* of an act model onto a perceptual model, and this capacity allows the actor to use any part of the body to formulate and transmit intentions, ideas, and skills.

At the same time, mimesis is the supporting adaptation of many other human endeavors. It enables athletes, skilled craftsmen, and other performers to refine their skills by generating variations on their actions and selecting the most successful ones. Mimesis is always an attempt to reduplicate some aspect of reality in action, and in the case of skilled rehearsal, the rehearsal itself is a mimetic act: the performer is imitating his or her own previous actions, and creating variations of those actions. The result is a personal repertoire that can be altered toward achieving some ideal of action. This is the cognitive path to a multitude of human skills. People acquire an incredible number of skills in a lifetime—they play sports and music, drive, and talk, to mention a few—and all these skills have been learned and improved through mimetic action.

Mimesis is the original source of human culture—that is, communities of mind linked together in a public expressive domain. Taken together, the mimetic actions of a small group of primate actors will inevitably generate a social theater of some complexity, and a rudimentary version of human culture, limited in its range of expression. On a larger scale, the same abilities will establish the implicit customs and folkways of a truly human culture. Even in the absence of language, this process carries out its work, as happens in communities of nonsigning deaf people. Mimetic role-playing and fantasy constitute a basis for a limited worldview, but one that is at least partially

public, and subject to some degree of cultural change. When this capacity was amplified through an interaction with spoken language, the expressive potential of mimesis was fully realized, resulting in an expressive culture of great power.

Where did mimesis come from? Our closest relatives are the chimpanzees, with whom we shared a common ancestor five or six million years ago and whose genes are very close to ours. But while chimps and humans have virtually similar cognitive capacities, chimps are very different from humans. We have traveled an inordinate distance, and this needs an explanation. It is true that our brains have tripled in volume, doubling their number of neurons, and that certain brain areas have expanded disproportionately. But there do not seem to be any new neural modules or neurochemical transmitters in the human brain. The most radically novel factor in our evolution is culture itself, as a collective storehouse of knowledge, and our brains have evolved specifically for living in culture. We are the species that made cultures into distributed cognitive systems, and those systems have reshaped our brains. In fact, the human brain cannot realize one of its key design potentials, symbolic cognition, without extensive cultural programming.

If we concede that human infants get language and all the tools of symbolic thought from culture, then we should ask: Where did cultures come from? What generated them *de novo* in the wild? The answer is: mimetic action. Apes are notoriously poor at mimetic action. A species cannot generate a culture until it can escape the autochthonous solipsism of the central nervous system and generate a common cultural space that can accumulate knowledge. Apes never managed to do this, primarily because they are so poor at gesture and imitation, and virtually incapable of deliberately self-supervising the rehearsal of their own actions to refine them. However, they have some of the key elements of mimetic ability, and this provided natural selection with the opportunity, once conditions gave fitness value to improved mimetic skill, to nudge and shape archaic hominids in the direction they eventually took.

The importance of mimesis can be seen in the limitations of even the most brilliant enculturated apes, who can manage symbol use much more easily than the gestural or skill-related dimensions of human culture. It may seem odd that Kanzi (the star performer of enculturated chimpanzees, who can segment the speech stream, understand some of the rudiments of grammar, and employ a vocabulary of several hundred symbols) cannot manage even a simple iconic gesture or engage in the kind of role-playing common in two-year-old children. Nor can he play basketball, as his trainer observed. But this is not odd at all; it is entirely consistent with what I have said about the crucial importance of mimesis in human cognition.

The central role of mimesis is relevant to determining the cognitive role of art in human history and prehistory because all art is essentially mimetic in style. Even literature, which appears to depend more on language than on mimesis for its superficial forms, is ultimately shaped by mimetic tendencies emanating from the deepest part of the writer's mind. This idea was articulated very clearly by the eminent critic Erich Auerbach a generation ago (1953/2003). In a similar vein the French philosopher René Girard recognized the role of mimesis in forming the fundamental dramatic tensions driving human social life (Girard, 1979).

In short, art is the expressive culmination of the most ancient domain of the human mind, as manifested in the rituals, public actions, and gestures that characterize any human society. It is woven into the deepest layer of meaning that can be called uniquely human. The power of mimetic expression can be furthered by technology, but the roots of that very special expressive style go deep into the earliest evolutionary layer of human emergence.

Summary and Conclusion

Art is a distinctively human form of cognitive activity that is characterized by the following features.

1. Art is aimed at influencing the minds of an audience, and may therefore be called a form of cognitive engineering.
2. It always occurs in the context of distributed cognition.
3. It is constructivist in nature, aimed at the deliberate refinement and elaboration of worldviews.
4. Most art is metacognitive in its role—that is, it engages in self-reflection, both individually and socially.
5. The forms and media of art are technology-driven.
6. The role of the artist and the local social definition of art are not necessarily fixed and are products of the current social-cognitive network.
7. Nevertheless, art, unlike most conventional engineering, is always aimed at a cognitive outcome.

Viewed in an evolutionary context, art originated in the earliest stages of hominid evolution, the so-called Mimetic phase. Newer forms have been scaffolded onto the older ones, and as human beings have evolved complex languages and technologies, artists have developed new forms that contain within them all the elements of our evolutionary history. Every newly evolved

artistic domain has a unique combination of these elementary components. Surveyed as a whole, the domains of art ultimately reflect the entire evolved structure of the human cognitive-cultural system. The challenge to cognitive scientists and neuroscientists is to develop a methodology that will allow them to fathom the abstract amodal processes of large-scale neural integration that transform the complex representations imposed by artists on their audiences into meaningful experiences. The ultimate engine of art, and the common force that makes art so distinct in its cognitive style from science, is mimesis. Therefore the genesis of art will not be understood, even in principle, until the neural and cognitive principles and mechanisms of mimesis are better understood.

REFERENCES

- Auerbach, E. 1953/2003. *Mimesis: The representation of reality in western literature* (fiftieth anniversary edition). Princeton, N.J.: Princeton University Press.
- Donald, M. 1991. *Origins of the modern mind: Three stages in the evolution of culture and cognition*. Cambridge, Mass.: Harvard University Press.
- . 1993. Précis of *Origins of the Modern Mind* with multiple reviews and author's response. *Behavioral and Brain Sciences*, 16:737–91.
- . 1998a. Hominid enculturation and cognitive evolution. In *Cognition and material culture: The archaeology of symbolic storage*, ed. C. Renfrew and C. Scarre, 7–17. Cambridge, England: McDonald Institute of Archaeological Research.
- . 1998b. Mimesis and the executive suite: Missing links in language evolution. In *Approaches to the evolution of language: Social and cognitive bases*, ed. J. R. Hurford, M. Studdert-Kennedy, and C. Knight, 44–67. Cambridge, England: Cambridge University Press.
- . 2001. *A mind so rare: The evolution of human consciousness*. New York: W. W. Norton.
- Fodor, J. 1983. *Modularity of mind: An essay on faculty psychology*. Cambridge, Mass.: MIT Press.
- Girard, R. 1979. *Violence and the sacred*. Baltimore, Md.: Johns Hopkins University Press.