

The Artful Mind

*Cognitive Science and the Riddle
of Human Creativity*

EDITED BY MARK TURNER

OXFORD
UNIVERSITY PRESS

2006

— -1
— 0
— +1

OXFORD

UNIVERSITY PRESS

Oxford University Press, Inc., publishes works that further
Oxford University's objective of excellence
in research, scholarship, and education.

Oxford New York
Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in
Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Copyright © 2006 by the Center for Advanced Study in
the Behavioral Sciences

Published by Oxford University Press, Inc.
198 Madison Avenue, New York, New York 10016

www.oup.com

Oxford is a registered trademark of Oxford University Press

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
electronic, mechanical, photocopying, recording, or otherwise,
without the prior permission of Oxford University Press.

Library of Congress Cataloging-in-Publication Data
The artful mind : cognitive science and the riddle of human creativity / edited
by Mark Turner.

p. cm.

Includes bibliographical references and index.

ISBN-13 978-0-19-530636-1

ISBN 0-19-530636-8

1. Art—Psychology. 2. Cognition. 3. Creative ability. I. Turner, Mark, 1954–
N71.A762 2006

701'.15—dc22 2005031824

9 8 7 6 5 4 3 2 1

Printed in the United States of America
on acid-free paper

—1 _____
0 _____
+1 _____

Acknowledgments

We are grateful to the Center for Advanced Study in the Behavioral Sciences, our host during the 2001–2002 academic year. As a token of our gratitude, we have assigned the copyright and the royalties for this book to the Center. The Center hosts the permanent supporting website for this book at <http://theartfulmind.stanford.edu>.

Generous support for our research was provided by the J. Paul Getty Grant Program. The Institute of Neuroesthetics and the Minerva Foundation arranged a conference at the University of California, Berkeley, that involved many of our members in January 2002. We are exceptionally indebted to Cynthia Read and Arthur Evenchik for their editorial work.

Mark Turner is grateful to Wildlife Education, Limited, for permission to reprint in his chapter the image from *Dinosaurs*, an issue of Zoobooks.

“The Neurology of Ambiguity,” by Semir Zeki, was first published in *Consciousness and Cognition*, Vol. 13, Issue 1 (2004), pages 173–196. It is reprinted with permission from Elsevier. The author is grateful for permission to reprint it here with minor changes.

Full color images that accompany this book are presented on the web at <http://theartfulmind.stanford.edu>. The site additionally presents supplementary documents, a history of the research activity surrounding this volume, and notes on further activities and research on the artful mind.

Contents

Contributors, ix

Prologue, xv

PART I Art and Evolution

- 1. Art and Cognitive Evolution, 3
Merlin Donald
- 2. The Aesthetic Faculty, 21
Terrence Deacon

PART II Art and Emotion

- 3. A Cognitive Account of Aesthetics, 57
Francis Steen
- 4. Composition and Emotion, 73
David Freedberg

PART III Art and the Way We Think

- 5. The Art of Compression, 93
Mark Turner

viii CONTENTS

- 6. The Cognitive Tango, 115
Lawrence M. Zbikowski
- 7. Dynamics of Completion, 133
Shirley Brice Heath

PART IV Art, Meaning, and Form

- 8. The Neuroscience of Form in Art, 153
George Lakoff
- 9. Form and Meaning in Art, 171
Per Aage Brandt
- 10. Slippages of Meaning and Form, 189
Stephen Murray

PART V Art and Sacred Belief

- 11. Making Relics Work, 211
Robert A. Scott
- 12. Architectural Space as Metaphor in the Greek Sanctuary, 225
Gloria Ferrari

PART VI Art and Ambiguity

- 13. The Neurology of Ambiguity, 243
Semir Zeki
- 14. Mastering Ambiguity, 271
Marc De Mey

Epilogue, 305

Index, 307

Prologue

The great riddle—of archaeology, cognitive science, neuroscience, anthropology, sociology, political science, linguistics, religious studies, and the humanities from literature and music to dance and art—is how we became human, how we acquired modern minds.

Human beings with mental architecture like ours came into existence only yesterday, evolutionarily speaking—perhaps fifty thousand years ago. At least, the archaeological record as we have it shows no robust evidence of cognitively modern behavior before that epoch. The staggering behavioral singularities that come with cognitively modern minds—advanced tool use, decorative dress, language, culture, religion, science, mathematics, art—present us with the greatest scientific embarrassment, for they appear to indicate a mysterious and unexplained discontinuity between us and the entire rest of Life.

To have a cognitively modern human mind is to be robustly artful, and conversely. This equivalence provides the inevitable starting point for a research program aimed at answering obvious yet hard questions: What is the evolutionary path from our remoter ancestors, who somehow lacked artful minds, to the existence of cognitively modern human beings, who cannot fail to be artful? How did the artful mind emerge? In a leap, or through slow development? What are the basic mental operations that make art possible for us now, and how do they operate? What neurobiology subtends these abilities? What is the interplay, in the phenomena of artfulness, between biological dispositions, individual experience, and cultural history?

—1
—0
—+1

xvi PROLOGUE

The individual human being, in form and movement, in thought and action, is a seamless intersection of powerful histories—phylogenetic history, individual development, and social and cultural history—all profoundly influential. A human being is a unified agency of biology, psychology, and social, environmental, and cultural patterns. And yet, the academic study of human beings is fragmented into scattered disciplines. How can science overcome this academic incoherence to launch a tradition of research in which neuroscientists, cognitive and developmental psychologists, archaeologists, vision scientists, evolutionary theorists, artists, art historians, semioticians, sociologists, and cultural historians join to explain the artful mind and its expression in cultures? How, in short, can inquisitive twenty-six-year-olds inspired to explain the artful mind discover a unified intellectual framework and institutional setting in which to begin thinking about it? Can their path be prepared to any degree by their elders, who lack their enviable plasticity and their exciting prospects, but who presumably command some of the knowledge, methods, and intuitions they might find useful?

This book has been designed with these goals and questions in mind. Its contributors collaborated over the course of a year, 2001–2002, at the Center for Advanced Study in the Behavioral Sciences. This year of collaborative residential research was made possible by a generous grant from the J. Paul Getty Grant Program. Some of the members of the group were in residence for the entire year, others for a month or two, a few for only short intervals here and there. We read each other's work; offered ideas, hints, and data; convened for more than thirty seminars; participated in three conferences; and conducted innumerable open-ended conversations over lunch and dinner or before and after screenings and shows.

This book is addressed to the next generation of scientists and scholars who seek to explain the wonders and mysteries of the artful mind. We hope it provides directions for a new field of research that, embryonic at present, can play an informative and eloquent role in answering the great riddle.

5

The Art of Compression

Mark Turner

Art is universal to our species. All human cultures show impressive, sustained, irrepressible impulses for artistic activity and understanding. Human art is possible because human beings differ mentally from other closely related species in having an advanced cognitive capacity for “double-scope integration.” This chapter focuses on the ways in which double-scope integration achieves conceptual compression, a hallmark of art.

Cognitively modern human beings have art, language, science, religion, refined tool use, advanced music and dance, fashions of dress, and mathematics. Blue jays, border collies, dolphins, and bonobos do not. Only human beings have what we have. This conspicuous Grand Difference constitutes a puzzling discontinuity in the evolution of life. How could these human singularities have emerged?

In *The Way We Think* (2002), and in earlier publications beginning in 1993, Gilles Fauconnier and I put forward the hypothesis that the Grand Difference arose in the following way. The basic mental operation of *conceptual integration*, also known as *blending*, has been present and evolving in various species for a long time, probably since early mammals, and there is no reason to doubt that many mammalian species aside from human beings have the ability to execute rudimentary forms of conceptual integration. Human beings evolved not an entirely different kind of mind, but instead the capacity for the strongest form of conceptual integration, known as *double-scope*

— -1
— 0
— +1

blending. Human beings are thus on a gradient with other species, but what a difference an extra step makes! Double-scope blending is the crucial incremental cognitive capacity that makes it possible for human beings to create and share art. (Technical introductions to the nature and mechanisms of blending can be found in Fauconnier and Turner 2002, 1998; Fauconnier 1997; Turner 1996, 2001, 2006.) In this chapter, I will focus on ways in which blending makes possible advanced forms of *conceptual compression* that are important for art.

As Merlin Donald observes, the brain strives “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. (Donald, this volume)

When we perceive a leaf twisting in the wind, we see it as one integrated leaf, one movement, one “wind.” When we look away and back, we think we see the “same” leaf before and after. This is a miraculous compression of perceptual diversity into unity. In all such cases, whether we are at rest or in action, we face a chaos of perceptual data. Bombarded by this diversity, we perform the highly impressive mental trick of compressing great ranges of it into manageable units. We parse an ocean of diversity quickly and reliably into a few elements coherently arrayed.

Typically, we are unaware that we face this perceptual diversity. When we look at the serene marble statue, it appears to us to be a single unit, without fragmentation, instability, or diversity, despite the fact that the perceptual data we are compressing to achieve this comforting and useful recognition of an abiding, unvarying statue are themselves shifty and uncoordinated.

At moments when we actually do manage to recognize that we confront shiftiness, we nonetheless feel—provided we are not in that instant afflicted with a cerebral hemorrhage, a drug-induced breakdown, or a chronic neural pathology—that the unities of the world shine through, fundamentals of perception, essentially impervious to accidents. We ascribe the tiny fraction of shiftiness that we do detect consciously to changes of viewpoint on our part or to motion or transformation on the part of the unity we perceive—events that, in our conception, leave the perceptual coherence of the world intact. The cloud moves in the wind, perhaps our view of it is blocked entirely while we walk past

the tree, and probably we are looking at the road for the most part anyway, but no matter, the cloud's unity is clear to us. This neurobiological creation of stability is profound and evident in everything we human beings do, despite our obliviousness to it. It is only under sedulous discipline during an ingenious experiment, for example, that we can begin to detect hints of the literal blind spots in our vision, caused by gaps where axons dive through the retina.

The neurobiological challenge of mastering spectacular perceptual diversity to achieve regularity and constancy is faced by very many species. Human beings, as I will discuss here, compound the difficulty of this challenge in a particular way. Over the last fifty thousand years, give or take (the dating is still being worked out in the archaeological record), human beings have demonstrated a remarkable ability to create new conceptual diversity. Human neurobiology makes it possible for our species, and only our species, to create great arrays of conceptual variety and yet to compress such arrays into manageable regularities. Explaining these operations of creation and compression is a key scientific problem in the study of the mind and the study of art. I focus on it in this chapter.

I begin with misleading but memorable examples, the first a detail from Pablo Picasso's 1907 *Les Femmes d'Alger (O. J.)* (see <http://theartfulmind.stanford.edu>). There are five women represented in *Les Femmes d'Alger (O. J.)*. The one in the lower right corner has a face presented in full frontal view, with two eyes, but a nose in profile. This "wild Squatter" consists of elements that come from "alternate visibilities," as Leo Steinberg describes:

The compression of space is greatest along the right margin, precisely where the rearward extension is deepest. And it is here, in the lower right corner, that the wild Squatter becomes the focus of intensified realization. In the development of this figure through its progressive stages, one discovers how Picasso gradually worked a straightforward backview, pigtail and all, towards ambiguity. While the end result is flat as a paper cutout, Picasso convokes alternate visibilities, relying in part on the punning scheme of the *Standing Nude*. In the final picture, an arm akimbo and one rising thigh fuse flat against a convenient curtain. Had these limbs been omitted, the rest of the figure might have been readable as a three-quarter backview with head jerked over the shoulder. But with these limbs retained, all three-quarter logic is thwarted. The figure becomes a full-splayed backview, in flat contradiction to the abrupt frontality of its head. And the violent wrench of her simultaneities more than makes

— -1
— 0
— +1

up for abstraction and flattening. It gives her pink flesh an aggressive immediacy, brought nearer still by the shameless impudence of the pose and the proximity of an implicated observer who knows every side of her. (Steinberg 1972, 172)

This image of a woman does not point to an abstraction, if by abstraction we mean what Semir Zeki means in *Inner Vision: An Exploration of Art and the Brain* (1999) and in his article “Artistic Creativity and the Brain” (2001, 51–52). Zeki observes in those works that the brain and art obey two supreme laws: the law of constancy—“registering the constant and essential characteristics of objects”—and the law of abstraction—“the process in which the particular is subordinated to the general, so that what is represented is applicable to many particulars.”

An example of abstraction is the operation of orientation tuning columns in the visual cortex. An orientation tuning column consists of neurons that respond preferentially to something linear at a specific angle. One of these orientation tuning columns, for example, prefers verticality. It will accordingly respond to anything vertical presented in the visual field—a pen, a finger, a lamppost. This is abstraction: the activation for verticality applies equally and completely to all the specific instances.

But the wild Squatter is not an abstract representation of a woman, in that sense of “abstraction.” For example, even her face—a small and relatively tame detail—seems to blend multiple views. She has a nose seen in profile, which comes from one view of the woman, but two eyes, which come from quite a different view of the woman. This representation cannot apply fully and equally to every view of the woman, in the way that the activation of the orientation tuning column for verticality can apply fully and equally to the pen, the finger, and the lamppost. On the contrary, the representation of the wild Squatter cannot apply fully to *any* of those views. So it is not an abstraction in that exact sense.

To be sure, there is abstraction involved in comprehending the wild Squatter. But the “alternate visibilities” do not point to an *abstraction* in Zeki’s sense. They point to a *compression*. The wild Squatter represents a compressed blend that includes elements from distinct input images—for example, from the frontal view and from the profile view in the case of the head, and from the backview as well for parts of her body, as Steinberg notes. As is typical in blending, the face of the wild Squatter leaves out many elements that could be found in the inputs to the blend. Projection from the inputs to the blend is selective. Additionally, the blend develops emergent properties that are not possessed by any of the input views. For example, in the blended face, but in

none of the inputs, we have a simultaneous view of elements from the profile view, the frontal view, and the backview.

These diverse elements from different inputs are not thrown into the blend arbitrarily. Each of the inputs is organized by a shared conceptual frame. The conceptual frame in this case is the same for all the inputs: it is the anatomical form of a human body—in fact, a female human body. The face, for example, makes use of the side and front views, both organized by the anatomical form of a female human body. If we emphasize this organizing conceptual frame as something generic over inputs, we produce a scheme something like the one illustrated in figure 5.1.

Both of these inputs are organized by the frame of the anatomical form of a female human body, and this frame is in the generic space. If we were thinking in terms of brain activation, we would say that each of the inputs activates the generic space. The blend shares this same organizing frame (see figure 5.2). It compresses elements from the separate inputs into that unified anatomical outline.

All four of these spaces are organized by the same general female anatomical form. The figure in the blend has a head and torso and all the usual appendages, and their part-whole relations of adjacency are for the most part preserved in the blend: the head is above the shoulders, one leg radiates from each hip, thigh and calf are joined by a knee, and so on. Steinberg naturally is drawn to the most violent wrench: “In the final picture, an arm akimbo and one rising thigh fuse flat against a convenient curtain.”




FIGURE 5.1. Cross-space mapping and generic space for the wild Squatter.

— -1
— 0
— +1

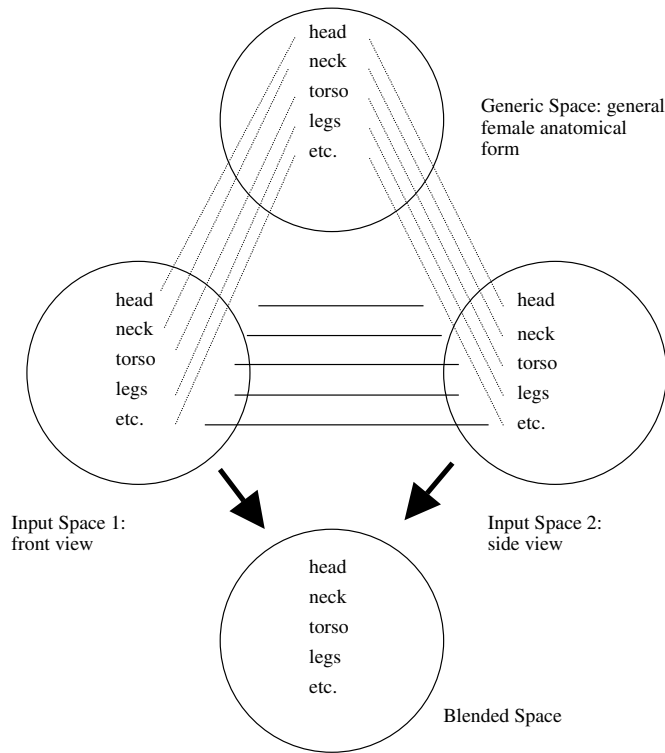


FIGURE 5.2. Part of the blending network for the wild Squatter.

The disparate views of the heads in the inputs are all projected to the one head in the blend. This is not a painting in which we have eyes on hands or a mane of hair on a doorknob. That would be quite a different blend. This painting has instead a clear morphological unity even as it compresses over diversity.

The blend is a compression of the input spaces, while the generic space is an abstraction over them. Again, if we are thinking in terms of activation, the blend activates partially both of the input spaces, and therefore activates the generic space to an exceptional degree. Indeed, the painting is remarkable for activating the entire mental network. We see just the same general method of compressing over diversity in Picasso's 1937 *Marie-Thérèse Walter* (see <http://theartfulmind.stanford.edu>).

Compression of viewpoint over time rather than space is exemplified in Marcel Duchamp's 1912 *Nu descendant un escalier* (see <http://theartfulmind.stanford.edu>). In this case, the compressed blend has elements that come from different temporal moments of watching the nude as it descends the staircase. In the blend, but in none of the inputs, we have an extremely

familiar conceptual unit, *the descent*, which remains connected to the different temporal moments. This unity-out-of-diversity can be expressed visually in Duchamp's fashion or linguistically by means of a definite noun phrase: "the descent." Duchamp's blend has emergent properties not possessed by any of the inputs. For example, in the blend, but in none of the inputs, we have a static form for the line of descent of the head.

Such compressions have occurred throughout human art since the Upper Paleolithic. A painting from the Hall of the Bulls in the Lascaux Cave (see <http://theartfulmind.stanford.edu>) shows the bull's head in exact profile, with one eye and one nostril, but also shows the horns from something like a three-quarters view. Similarly, we see the bull from the side but the bull's cloven hooves from nearer the front. This painting points not to an abstraction but to a compression over many quite diverse views.

In the cases of *Les Demoiselles d'Avignon*, *Marie-Thérèse Walter*, and *Nu descendant un escalier*, we recognize immediately that something is cockeyed, jumbled, or lumped. But in the case of the Lascaux bull, most people do not seem to be *consciously* aware that there is any compression at all until it is pointed out. This lack of recognition cannot be ascribed to ignorance of actual horns and hooves. I can provide evidence: although my office until recently sat at the summit of a hill in a coastal mountain range populated by wild animals, where I saw domesticated and wild hooved and horned animals in the flesh routinely, and although my children show me scores of photographic images in books and instructional films of bulls, bison, gazelles, reindeer, moose, and ibex, and although I would be quite astonished, and perhaps worried for the animal, were I to see an actual bull with horns and hooves disposed in this way, nonetheless I was not consciously aware of the oddly compressed nature of the Lascaux bull for a long time. The details of the integration network we use to understand it, with its selective projections from inputs to a compressed blend, are, as is typical, invisible to consciousness even though the product—the useful, single, human-scale, compressed bull—may be quite arresting.

Blending, even when it is remarkably creative in providing emergent structure in the blend, is always deeply conservative, anchored in what we know. The creativity is greater than we usually see, but also more profoundly anchored than we usually see. For example, although the Picasso and Duchamp and Lascaux blends are highly creative, they are rooted in what we know well—in what we, in important senses, already have. When we see a particular view of a person or an animal, we know that there are other views: we know when we see one side of the face that there is another, and that when we see the front of a person there is also a back that is not presently visible. When the person turns

— -1
— 0
— +1

to face us, we are not surprised at all to see these other sides of the person. Similarly, we know when we see someone in mid-descent down a staircase that the person is at other spots at other times; indeed, if we catch only one glimpse of the person descending, we conceptually complete the action, stretching before and after. The input spaces therefore are available to us already. The creativity comes in blending them in a way that is not otherwise available to us, providing in one static view elements from distinct, but known, views.

These phenomena of compression are well known within art history. Artists also often use compression as an explicitly avowed principle of theory and practice. But compression, far from being reserved for special and remarkable effects in art, is ubiquitous and indispensable throughout human understanding. Art, in this case, exploits a basic human mental capacity.

Consider an example not from art but from everyday journalism. On July 8, 1999, the *New York Times* reported that Hicham el-Guerrouj had broken the world record for the mile, with a time of 3:43.13. To convey at a glance the historical significance of the performance, the *New York Times* provided an illustration of a one-quarter-mile racetrack with six figures running on it. The six figures represent el-Guerrouj in a race against the five other runners—namely, the fastest milers from each decade since Roger Bannister broke the 4-minute barrier in 1954. El-Guerrouj is crossing the finish line as Bannister, trailing everyone else, is still 120 yards back. This illustration prompts us to construct a conceptual packet that blends structure from six separate input mental spaces, each with a one-mile race in which the world record is broken by a runner. The blend places all six runners on a single racetrack, with a single starting time, in a kind of mythic race.

The blend is a compression, giving us an immediately intelligible human-scale unity—that is, a single footrace, with a winner. Literally, the blend is “false,” obviously so in conjoining Hicham el-Guerrouj and Roger Bannister as competitors on the same track. It is also tacitly inexact in the way it locates the runners other than Hicham el-Guerrouj on the track. But it gives us a way of understanding the truth about the many different input spaces and their relationships. The integration network provides a compressed blend that lets us understand and remember at human scale a range of complicated knowledge that does not otherwise fit human-scale recognition. There is emergent dynamic structure in this compressed blend—namely, structure that cannot be found in any of the inputs: the blend is a simulation of a mythic race between legendary competitors, most of whom never in fact raced against each other. In this mythic race, Hicham el-Guerrouj “defeats” Roger Bannister by 120 yards. The compressed blend gives us the direct pleasure of immediately recognizing the competition and its conclusion.

All of the artworks and drawings I have adduced so far call up *mirror integration networks*. A mirror network is one in which all the inputs share a single organizing frame, which is projected also to the blend to organize it. The shared frame gives the inputs an analogical relation; each has the same set of roles. For example, we can frame George Washington and Abraham Lincoln using the conceptual frame *president of the United States of America*. The two resulting conceptual spaces are analogous since the frame-roles in one have identity links to the frame-roles in the other. In a mirror network, separate conceptual spaces share an organizing frame, which is also projected to organize a blended space.

It is extremely common for us to compress analogy links to identity links and identity links to unity. For example, there are striking and profound differences between me now and the male infant born to my mother many decades ago, but we construct analogy links between those two conceptual spaces. Those analogy links are almost always compressed to identity links, making not only the roles in those two spaces but also the values of the roles identical. That network with identity links is further compressed to create a blend in which there is a single unity, the person. The differences between the conceptual spaces—that is, between the baby and me now—are compressed to *change* in the blend. The result is a unified blended space in which there is a unity—the person—who undergoes change. We can thereby think of a complicated human life in the way we think of a leaf blowing in the wind: blending makes it possible in both cases for us to conceive of a single unit that undergoes change over time.

The two Picasso paintings also evoke mirror networks. Each of the input spaces is framed by the gross anatomical form of a woman, and there are identity links not only between the roles in each frame, such as *adult woman* in a *location*, but also between the values of those roles: it is the same woman, in the same place. But the details differ space to space, because each has a different view of the woman in that location. The blend keeps the organizing frame but projects conflicting details into it. The result is a mirror network. The analysis is identical for the Lascaux bull.

The painting of the nude descending the staircase evokes a mirror network in which each of the input spaces has the organizing frame of a person on a staircase, facing down, and engaged in movement. This frame is projected to the blend, but extended to incorporate several locations for the person.

The graphic illustration of the world record in the mile prompts for an integration network with six input spaces, each one framed by a mile footrace, on a standard oval racetrack, in a location and a time, and with a winner and

— -1
— 0
— +1

losers. That frame is projected to the blend, giving a mirror network that, as always, compresses various disparate conceptual spaces into one unified blended space.

In any mirror network, there is always the question of which of the conflicting details from the input spaces are absorbed in the blend, and how they are incorporated there. In some cases, the conflict between alternative details is resolved by projecting only one of them. For example, in the Lascaux bull network, the input spaces have several different views of the horns, but only one of them is projected to the blend. The bull is not equipped in the blend with several different sets of horns, each from a different view. That technical possibility is not deployed in this case.

The painting of the nude descending the staircase illustrates a different technical possibility. The frame for seeing a person heading down the stairs has only one view of the person available in any instant. But the blend receives from the disparate inputs many different and conflicting views, with the result that they become simultaneously visible in the blend. We are not confused. We know that each of the input spaces has only one view of the person. But the frame is extended in the blend to include several views.

The mythic footrace deploys yet a different technical possibility. Each of the input spaces has a winner and losers. None of the specific losers is projected to the blend, although the roles for losers are projected to the blend as part of the frame of a footrace. The specific winner (Hicham el-Guerrouj) in one of the input spaces is projected to inhabit the role *winner* in the blend, but the specific winners from the other input spaces (such as Roger Bannister) are projected to inhabit slots for *losers* in the blend. In this case, the organizing frame already has slots for *losers* into which these analogous winners can be placed. Every winner in an input space is projected to the blend, but only one of them inhabits the role *winner* there.

In all of these examples, even Duchamp's painting of the nude descending a staircase, the actual visual image is static. But the principles work the same for dynamic images. The solution of the famous riddle of the Buddhist monk presents a dynamic visual blend.¹ Here is the riddle:

A Buddhist monk begins at dawn one day walking up a mountain, reaches the top at sunset, and meditates at the top for several days until one dawn when he begins to walk back to the foot of the mountain, which he reaches at sunset. Making no assumptions about his starting or stopping or about his pace during the trips, prove that there is a place on the path that he occupies at the same hour of the day on the two separate journeys.

-1 ____
0 ____
+1 ____

An intuitive solution to the riddle comes from superimposing the two days—that is, imagining the Buddhist monk walking both up the path and down the path on the same day. Then there must be a place where he meets himself, and that place is clearly the one he would occupy at the same time of day on the two separate journeys. This is again a mirror network. Each of the two input spaces has the Buddhist monk walking on the mountain path from dawn to sunset. The blend has the same organization, but is now extended to have two people on the path, walking in opposite directions. The ascent from the first input and the descent from the second input are both projected to the blend, with the consequence that, in the blend, the monk is visibly in two places at the same instant, just as in Duchamp's painting of the nude descending the staircase, the body is visibly in several places at the same instant.

We are of course not confused. Just as we know that in the inputs for Duchamp's blend, the nude is in only one place at a time, so we know that in the inputs for the riddle of the Buddhist monk, the monk is in only one place at a time. But in the blend, there are two identical monks and the monk "meets" himself. The "meeting" is a compressed human-scale event inside the blend. It corresponds to an uncompressed set of links between the inputs. That is, we recognize that the "meeting" must decompress to a trio of identity links connecting the two input spaces: there is a spot on the path in the ascent that is connected by identity to a spot on the path in the descent; the time of day when the monk is located at that spot in the ascent is connected by identity to the time of day when the monk is located at that spot in the descent; and of course the monk in the first space is connected by identity to the monk in the second space. The monk's meeting himself in the blend is a human-scale compression of the input spaces and the relationships between them, and this compression allows us to have human-scale insight into something that is otherwise diverse, distributed, even impenetrable.

Now in this case, the image is not static. We run the blend as a dynamic simulation in which two people converge and then meet on a mountain path. I can only suggest this dynamic image with a static diagram (see figure 5.3). In this diagram, the generic space represents the conceptual structure that is taken as applying to both of the inputs. All four spaces have the monk traveling the mountain path. The blend has a further extension, with both monks, and the meeting.

Now consider a case where thousands of input spaces are potentially accessible by means of unpacking a blend. This blend lies in a conceptual network that provides an idea of a putative evolution from dinosaurs to birds. The artwork that prompts us to construct this conceptual integration network

— -1
— 0
— +1




FIGURE 5.3. Part of the blending network for the Buddhist monk.

is taken from an issue of Zoobooks, a publication for grade-school children (Wexo 1992).

This picture depicts a dinosaur chasing a dragonfly (see figure 5.4). We see the dinosaur at various stages of running along a path. Note that the dragonfly is always the same, the path is a single path, and the shadows show that the light is always coming from the same direction. We are given a single scene in which a single dinosaur is changing into a bird, which is at last able to catch and eat the dragonfly. It is also easy to understand the dinosaur as wanting to change into the bird, exactly so it can accomplish its goal.

We realize that at one level in the integration network behind this artwork, there are literally thousands of dinosaurs and birds, with no identity

-1 ____
0 ____
+1 ____




FIGURE 5.4. The evolution of birds from dinosaurs. From Zoobooks' *Dinosaurs*; reprinted with the permission of Zoobooks and Wildlife Education, Ltd.

connections between them. These individual organisms are compressed into representatives of generations. There are analogy and disanalogy connections between any two generations, given that there are differences generation to generation in the dinosaur phenotype. There are also cause-and-effect connections at that level, since one generation produces the next. At this level, we have different stages of the “identical” species, somewhat different in each space.

The compression to a “species” that “changes” is already quite an arresting compression. We know that none of the actual individual dinosaurs changed genetically, although of course each of them went through ontogenetic development. On the contrary, what happened was this: the very many different individual dinosaurs had different degrees of differential reproduction, but they all died, and the dinosaurs downstream were somewhat different, because of inheritance, variation, and selection. The typical representative of the species at one time differed from the typical representative generations later, but not because there was any change in any particular organism. If dinosaurs several generations later were a little more feathery, it is certainly not because any one dinosaur changed in that direction. On the contrary, the individual dinosaurs stayed just as they were, and all died off, and their replacements looked a little different.

— -1
— 0
— +1

But compression now gives us in the final blend not a species that changes but a unique dinosaur that changes. The analogies across all the decompressed spaces become *uniqueness* in the blend, and the disanalogies across them become *change* in the blend. In the blend, we have a unique element that undergoes change. This is just the kind of compression we use to see a leaf blowing in the wind or to understand that a person changes over a lifetime, but in the case of the dinosaurs, there are no identity links across all the individual dinosaurs compressed through this cascade of integrations. Human beings are not set up to understand mechanisms of evolution operating over evolutionary time scales. In the case of the dinosaurs and birds, we use the standard compression template that has at one level analogy and disanalogy across many spaces but also has a blend in which all the analogy links are compressed to uniqueness and all the disanalogy links are compressed to change for that unique element. This template of compression allows us to understand a diffuse range of meanings that is spatially and temporally very far from the kind of thing human cognition is set up to recognize. Compression brings this diffuse, complex, and otherwise impenetrable domain of evolution into a form that fits human understanding.

Of course, we are not fooled by the blend, any more than we believe that a Buddhist monk can in fact be in two places at once or that Hicham el-Guerrouj crossed the finish line 120 yards ahead of Roger Bannister or that when we look at a woman we can see something that resembles the painting of Marie-Thérèse Walter. But the blend is extremely useful as part of an integration network. It is important to see that the kind of compression used in the integration network for the evolution of dinosaurs is not restricted to children's textbooks. The best and most careful evolutionary biologists routinely conduct serious discussions of the way in which "a species" "changes" "over time" "to become" this or that; of how a species "acquired" this or that feature or "developed" or "grew" in this or that way; of how "it" "lost" this or that capacity; of how one organ "became" something else. All of this is quite false for the most decompressed understanding of the life and death over time of individual organisms and the differences between them. But the compressions are very useful. Evolutionary biologists are usually wary of ascribing intentionality to evolutionary "development," but their caution is usually a matter of reminding the audience that the intentional language can be cashed out to a level of understanding that does not involve intentionality. That is exactly right: the blend is useful because the compression it carries can help us access, activate, remember, and use the entire network in which it resides.

Mirror networks, or near-mirror networks such as the network for the evolution of dinosaurs into birds, are impressive displays of human creativity.

Made possible by our creative capacities for integration and compression, they in turn support the further work of those capacities, and this further work increases their conceptual reach.

Even more impressive are double-scope integration networks. In a double-scope network, the two inputs have different (and often clashing) organizing frames, and the blend has an organizing frame that receives projections from each of those organizing frames. The blend also has emergent structure of its own that cannot be found in any of the inputs. Sharp differences between the organizing frames of the inputs offer the possibility of rich clashes. Far from blocking the construction of the network, such clashes offer challenges to the imagination. The resulting blends can turn out to be highly creative.

The ability for double-scope blending seems to be available to children very early. For example, in Crockett Johnson's (1983) *Harold and the Purple Crayon*, written for three-year-olds, Harold uses his purple crayon to draw, and whatever he draws is real, although it is also clearly a child's drawing.

His world is a blend, of spatial reality and its representation. In the blend, the representation is fused with what it represents. When Harold wants light to go for a walk, he draws the moon, and so he has moonlight. The moon stays with him as he moves. This blend has two inputs. One input has elements of the real spatial world as we experience it and perceive it. One of those elements is the moon. The other input to the blend has conventional knowledge about drawing. In the input with the real moon, the moon cannot be created by drawing and it does not come into existence at someone's will. In the input with drawing, a drawn moon cannot emit moonlight or float along in the sky as the artist's companion. But in the blend, there is a special blended moon with special emergent properties: you can create it by drawing, and it gives light and "moves" with you.

The mechanisms of blending that give us this special blended moon are in operation throughout *Harold and the Purple Crayon* (see figure 5.5). When he needs to walk, Harold draws a path, and then sets off on his walk, taking his big purple crayon with him. When he wants to return home, he draws a window around the moon, positioning the moon where it would appear in his window if he were in his bedroom, and so he is automatically in fact in his bedroom and can go to bed.

Child Harold's blended world has new kinds of causality and event shape that are unavailable from either the domain of drawing or the domain of spatial living. The projection to this blend, and the completion and elaboration of the blend, are not algorithmic, not predictable from the inputs, but instead have considerable room for alternatives. For example, when one draws, one

— -1
— 0
— +1




FIGURE 5.5. Blending network for *Harold and the Purple Crayon*.

often makes practice sketches, erasures, and mistakes that do not count as part of the finished drawing. Which kinds of marks made with the purple crayon shall count as reality in the blend? The answer chosen by the author of the book is “all of them.” When Harold’s hand, holding the purple crayon, shakes as he backs away from the terribly frightening dragon, the resulting mark is a purple line of wavy scallops: “Suddenly he realized what was happening. But by then Harold was over his head in an ocean.”

The principle for connecting the purple sketches to elements of reality is, not surprisingly, image-schematic matching: if the sketch matches the iconic form of something, it is that thing. But it appears that this matching is constrained: a given purple sketch can be matched to exactly one reality. For example, once the wavy line is an ocean, Harold cannot transform the ocean into a cake by perceiving the wavy line as the icing on a cake. Yet in a differently conceived blend, in a different book, the character who does the drawing might possess the power to recast reality by perceiving the sketch first one way and then another.

In Harold’s blend, all of physical space is a piece of paper on which to draw. What are the possibilities in the blend of blank paper/empty space? Can

–1 ____
0 ____
+1 ____

Harold move as he wishes through it? The answer chosen by the author is that once something is drawn that gives Harold relative location, he is constrained by some of the physics of the real world. For example, once he draws the hull of a boat and part of the mast, he must climb the mast to draw the parts of the boat he could not reach from the ground. When he wants to find his house, he begins to draw a mountain so he can climb it to have a better view. He climbs the part he has drawn so he can draw more mountain to climb. But as he looks down over the other side of the mountain, he slips, and since he has been positioned with respect to the mountain, the blank space is now thin air, so he must be falling. He is obliged to draw a balloon to save himself from crashing.

Blends of the sort that we have looked at are found widely throughout art. In them, things are blended that do not in reality go together. From the point of view of evolution, to confuse things that should be kept distinct is like plucking forbidden fruit: we should not do it, on pain of death, quite literally, but also on pain of insanity. Yet, amazingly, with many mirror networks and all double-scope networks, we pluck that forbidden fruit. We put together what should be kept distinct.

Obviously, we should not confuse two different views of a woman or a bull. We should not confuse different moments in an event. We should not believe that world record holders in the mile who lived decades apart in fact all ran together on the same track at the same time. We should not believe that a Buddhist monk can in fact be in two places at once. We should not believe that a dinosaur in fact grew feathers and turned into a bird while chasing a dragonfly. We should not confuse the moon with a sketch of the moon, or paper with the world, or drawing with locomotion. But plucking this kind of forbidden conceptual fruit makes us extremely creative.

Art and literature are particularly specialized to take advantage of this kind of “forbidden fruit” blending. We recognize that the input spaces should be kept separate but nonetheless the blend is conceptually useful. It is true that sometimes this blending is aggressive and presents blends that are unharmonious or disconcerting or in some way out of joint. A mild example, studied with great insight by Per Aage Brandt (this volume), comes from the work of René Magritte. In *La Tentative de l’Impossible* (1928), Magritte produces a blend that follows the same initial lines as the blends we saw in *Harold and the Purple Crayon*. In this painting, we see a painter who is painting a woman. I do not mean that he is applying paint to the skin of an already existing woman, nor do I mean that he is in the conventional and straightforward sense merely applying paint to a canvas to make an image. On the contrary, the body of the woman comes into existence as he paints it but there is no obvious canvas. In

— -1
— 0
— +1

the blend of *La Tentative de l'Impossible*, as in the blend for Harold, reality is something one can draw on, and what one draws is real.

Forbidden-fruit integration gives us a way to keep things separate and yet to combine them. In the input spaces, they are quite separate. We understand the vital relations between them but do not compress them. Yet in the blend, we do exactly that. A forbidden-fruit integration network is a truly marvelous way to have your forbidden fruitcake and eat it, too. This is of course exactly what happens in *Harold and the Purple Crayon*. We understand very well the difference between the moon and a drawing of a moon, and in the input spaces, we understand their relationships of difference. But in the blend, they are integrated.

Every kind of art shows the power of forbidden-fruit integration. For example, if we survey high canon representations of the Annunciation, we “see” the Virgin holding, anachronistically, a lectionary, often opened to the narrative of the Annunciation. In Rogier van der Weyden’s Annunciation (see <http://theartfulmind.stanford.edu>), the medallion on the bed represents the Resurrection. The painting in this way gives us a compression of eternity, or at least, from not yet being born to being raised from the dead, all in one momentary scene.

We have no trouble interpreting this representation as evoking a blend of a young girl and the Mother of God, which is already a blend. The Virgin’s bedroom may additionally have features of a church—the lectionary stand and veil that are part of the furniture of an altar, trinitarian tracery windows in Broederlam’s version, a full Gothic church interior as in one of Jan van Eyck’s versions. Annunciations may have a representation of God in the upper left, although we do not interpret this to mean that God was just up and to the left of the bedroom. In the Mérode Altarpiece (see <http://theartfulmind.stanford.edu>), a homunculus already tolerating his own miniature cross is flying toward the womb of Mary. The representation evokes a blend of girl with Mother of God, bedroom with church, breath with life, and so on.

These Annunciations take concepts incredibly diffuse, foreign, and difficult to understand—eternity, divinity, theology, the Church, and the relationship of the immortal to the mortal—and compress them to an extremely familiar scene: a room, a woman reading in the room, and someone addressing her. What lies beyond human understanding is compressed to human scale. In this case, plucking the forbidden fruit gives us an understanding of God.

V. S. Ramachandran has encouraged us always to ask three particular questions when we discuss Art and the Mind: What? Why? and How? The What? question has by now been answered to an extent: forbidden-fruit blending happens widely throughout human art, science, religion, mathe-

matics, culture, and indeed throughout anything done by cognitively modern human beings, and it follows a set of constitutive principles and a set of governing principles. It is what makes us cognitively modern.

The Why? question is harder to answer because we do not have a time machine to take us back fifty or one hundred thousand years or more to obtain the evidence. I am nevertheless not uncomfortable proposing that forbidden-fruit integration is an extension of integration abilities possessed prior to its evolution, and that the power of the increased creativity resulting from the ability to do forbidden-fruit blending was extraordinarily adaptive. Even small increments in that mental ability would have conferred advantage, and so the natural selection story taking us from more rudimentary forms of conceptual integration to full forbidden-fruit integration is easy to imagine, if perhaps impossible to prove.

The How? question is daunting, since our ignorance about the neuroscience of higher-order thought is profound. Yet I have more confidence in the search for an answer here than I have in the attempts to address the Why? question, exactly because we can make actual observations, however indirect, and pursue normal scientific practices in search of an explanation of how our present brains accomplish forbidden-fruit integration. (By contrast, the answers to the Why? question are entirely speculative under present methods, and nearly all the evidence related to the Why? question has vanished from the earth.) I see two, or maybe two and a half, hypothetical sources for answers to the How? question. The first, obviously enough, has to do with neural binding of the sort we do during everyday perception, location-time collocation, and other mammalian integrations. The one-half is synaesthesia, which rates only one-half since it looks as if it is probably neural binding in a different mode. The last one is special-purpose forbidden-fruit integration, tightly restricted to certain narrow domains and behaviors. Chase play, for example, is common throughout the mammalian world and evidenced even in interactions between different species, such as a human child and a dog, or a polar bear and a wolf. In this behavior, allied organisms, usually a parent and an offspring, simulate predatory behavior. During that behavior, they are simultaneously activating motor patterns, attention patterns, and motivational structures that belong to two very disparate domains. It seems not an unreasonable hypothesis that the neural circuitry subtending binding, synaesthesia, or special-purpose forbidden-fruit blending might have gotten the ball rolling in the run-up to full cognitive modernity.

I cannot supply anything more than merely suggestive evidence for the following view of the cognitively modern human brain, but I am comfortable wondering whether the human brain is not a kind of vast bubble chamber,

____-1
____0
____+1

constantly trying to blend different things. Perhaps very many of these attempts are going on in our brains all the time. Perhaps almost any two things that are activated simultaneously become candidates for an attempt at blending. I imagine that most of these attempts fail almost immediately because the constitutive principles of blending are not fulfilled or the governing principles of blending are contravened or the integration networks do not attach themselves to any purpose we have. Of the relatively very few conceptual integration networks that are successful, only fewer still ever percolate into consciousness. But this constant attempt at blending provides a robust way of introducing a strong engine of variation into our conceptual systems. Almost all of those products of variation are selected against by governing principles or by pressures and affordances of our environments or by the absence of utility of any kind. But some of them, although they begin by blending structures that one might think have no business being blended, nonetheless provide quite powerful new conceptions.

The distributed cognition question is essential here. The great virtue of any one human being's coming up with any forbidden-fruit integration is that all the other human beings stand ready to understand it, incorporate it, and propagate it. In this way, culture is an incomparably larger bubble chamber than is an individual brain. With the entire species running forbidden-fruit experiments in this vast bubble chamber, there are at last wonderful possibilities for sustained, effective, and accretive creativity.

But these hypotheses are at the outer limit of what cognitive neuroscience, evolutionary biology, and art theory can at present investigate. An aggressive and sustained program of research is needed to explain how the brain accomplishes forbidden-fruit integration in art. Inevitably—and this is the central lesson from our studies—that program of research will require intensive and sustained collaboration across researchers in many different disciplines, to design the optimum scientific system of attack and to conduct it to its goals.

NOTE

1. A version of this riddle appears in Koestler 1964, 183–89. Koestler attributes the riddle to the psychologist Carl Dunker.

REFERENCES

Fauconnier, G. 1997. *Mappings in thought and language*. Cambridge, England: Cambridge University Press.
Fauconnier, G., and M. Turner. 1998. Conceptual integration networks. *Cognitive Science*, 22, no. 2 (April–June): 133–87.

–1 ____
0 ____
+1 ____

- . 2002. *The way we think: Conceptual blending and the mind's hidden complexities*. New York: Basic Books.
- Johnson, C. 1983 [1955]. *Harold and the purple crayon*. New York: Harper & Row.
- Koestler, A. 1964. *The act of creation*. New York: Macmillan.
- Steinberg, L. 1972. The Algerian women and Picasso at large. *Other criteria: Confrontations with twentieth-century art*. New York: Oxford University Press.
- Turner, M. 1996. *The literary mind: The origins of thought and language*. New York: Oxford University Press.
- . 2001. *Cognitive dimensions of social science: The way we think about politics, economics, law, and society*. New York: Oxford University Press.
- . 2006. The Blending Web site: <http://blending.stanford.edu>.
- Wexo, J. B. 1992. *Dinosaurs*. A volume of Zoobooks. San Diego: Wildlife Education, Ltd.
- Zeki, S. 1999. *Inner vision: An exploration of art and the brain*. New York: Oxford University Press.
- . 2001. Artistic creativity and the brain. *Science*, 293, no. 5527 (July 6): 51–52.