

6. Imaging the Future	91
MICHAEL I. POSNER AND DANIEL J. LEVITIN	
7. What to Do with Your Own Personal Brain Scanner	111
ALAN GEVINS	
8. The Deep and Surface Structure of Memory and Conscious Learning: Toward a 21st-Century Model	127
KARL H. PRIBRAM	
9. What Are Brains For?	157
MICHAEL S. GAZZANIGA	
III PSYCHOLOGY (MEMORY, THEORY, AND COGNITION) IN THE 21ST CENTURY	173
10. The Future of Cognitive Psychology?	175
HENRY L. ROEDIGER III	
11. The Memory Trainers	199
GAY SNODGRASS	
12. On Future Psychological Categories	235
JEROME KAGAN	
13. The Goal of Theory in Experimental Psychology	253
GEORGE SPERLING	
14. How to Prepare for Our Future of Totally Unexpected Opportunities	265
NEAL E. MILLER	
15. The Future of Psychology	271
HANS J. EYSENCK	

IV	THE SCIENCE OF THE MIND	303
	16. Mind Sciences and the 21st Century	305
	ROBERT L. SOLSO	
	About the Authors	325
	Name Index	341
	Subject Index	351

Foreword

A Personal View of 20th-Century Psychology: With an Eye to the 21st Century

ERNEST R. HILGARD

Almost a decade ago I finished writing *Psychology in America: A Historical Survey* in which I chose to concentrate on 20th-century psychology in Americas where the growth had been particularly remarkable and where I knew many of the participants and developments personally. Little did I know then that I might be asked to write a foreword to a book with the title *Mind and Brain Sciences in the 21st Century*. You will find in this book a collection of stimulating papers written by many of the people who have had a central role in the shaping of psychology as we now know it. Time, of course, will decide the validity of the prophesies.

As a historian of psychology I have been far more accustomed to writing about the past and what it means than about the future, but welcome the chance to participate in this worthy venture as, after all, “the past is (truly) prologue.” We can, of course, learn much from the past as we plot the future of the science of the mind and I would like to offer a brief, personalistic commentary of where we have been, with an eye to where we might be going.

About 100 years ago, psychologists attempted to establish the field as a science among the more recognized sciences that had flourished during that remarkable century. These psychologists, perhaps more so than other scientists, sought to establish a “new psychology” as a systematic body of knowledge, independent and unified according to basic scientific principles and practices. Their ties with philosophy encouraged this desire for a systematic stance. During the early 20th century, the major competing theories in American psychology took the form of schools, each of which wanted to dominate. They had reached their height by the early 1930s, when the summarizing books appeared and contrasts between schools were sharpened. Behaviorism had gradually become ascendant by that time, primarily based on the proposed

objectivity of its methods. There had, however, been little systematic development of behaviorism through detailed experimentation designed to justify the theory, rather than to serve other purposes. John B. Watson, who continued to popularize behaviorism, tied it loosely to the conditioned reflex, to extreme environmentalism, and offered proposals for child rearing that were based on a minimum of experimentation and were without careful deduction from formulated theory.

Most psychologists engaged in investigatory work during this period, even though theories of small scope had rather little to do with their day-to-day conception of what it meant to be a psychological scientist. To the typical investigator the general characteristics of science seemed well-known and familiar: careful observations, often aided by instruments; quantitative measurements; appropriate controls so that the variables under study were surely the effective ones; precisely reported procedures so that others could verify the results by repeating what was done; and so on. What was important was the new knowledge generated. Such knowledge was not a mere description of isolated facts; the search was always for some “lawfulness” in the relationships studied. If the lawfulness applied over a larger domain, then it fell in the range of “theoretical explanation.” For many investigators, even today, all this seems plain and straightforward. Investigators were then, as now, problem solvers, and problems in a smaller and more accessible domain were more likely to be solved than the very large ones. Now, as we are about to enter the third millennium it would seem that even though the problems will be vastly more complicated, the matter of psychologists breaking down complex problems into manageable questions and solving them might still be an appropriate strategy in advancing knowledge.

Shortly after the midpoint of the 20th century the climate was changing in psychology as well as in the logic of science. Best known was Thomas Kuhn’s viewpoint that immediately became popular in which he suggested that science does not advance primarily by accumulation but instead by revolutionary changes in the form of new ways of looking at things, described by him as “paradigm shifts.” His proposal struck a very responsive chord. A social-psychological byproduct of this interpretation was that older scientists were so deeply wedded to their paradigms that the newer ones often failed to get acceptance until the older scientists had died and a younger generation took over. There is, in other words, a conservative tendency with respect to those theories that had provided a framework in which science had already advanced.

The advances of science within such an established framework were called by Kuhn “normal science.” That is, some of the work of scientists in improving their instruments and methods, or in exploring new lines of investigation, require no shifts in the basic paradigms. For example, the relativity of space and time, which makes Newtonian space-time obsolete, does not make obsolete ordinary mapping of the earth’s surface according to Euclidean geometry or ordinary timing of events by clocks or stopwatches, because the space-time continuum has appreciable effects only for events in space near the speed of light. Before a science has reached the stage of elegant theories of wide scope and while it still depends mostly on normal science with somewhat limited theory, it is said to be in a *preparadigmatic* stage. Most of those who tried to apply Kuhn’s theory to psychology have agreed that psychology is in that stage.

My own experience in psychology came in a single course with Professor Madison Bentley at the University of Illinois in the first quarter of the 20th century, while I was majoring in chemical engineering. Naturally, I learned from him the psychology most identified with Wilhelm Wundt, by way of E. B. Titchener, who Bentley had worked with at Cornell and later succeeded. This introspective psychology, with its emphasis upon sensation, was commonly known as “structuralism” and was a very influential school of thought in the early years. It remained influential as I shifted my interest to psychology in the course of my Ph.D. studies at Yale, leading to the completion of my degree in 1930.

At Yale, the theoretical position in psychology had shifted to “functionalism,” a viewpoint derived more from William James and especially from James R. Angell and John Dewey at the University of Chicago. This position was represented most strongly among my teachers by Edward S. Robinson, who had come to Yale from a faculty position at Chicago. It was a congenial environment because Angell, a former Chicago colleague of Dewey, was the President of Yale.

In the meantime, the new orientation, known as “behaviorism,” had developed under Watson at The Johns Hopkins University and other universities, such as at The Ohio State University under A. P. Weiss. To some extent American psychology retained its behaviorist coloring for some 50 years between 1913 and 1963. Still, there were many who were not converted, and held a less rigid position, neither Titchnerian nor Watsonian, and described a psychology of the “middle road,” as characterized by the very influential

Robert S. Woodworth. There were, of course, other systems, listed by Woodworth in his book *Contemporary Schools of Psychology* (1931) and by Edna Heider in *Seven Psychologists* (1933). Much later the influential development of cognitive psychology and cognitive neuroscience occurred.

Models as theories became prominent while the psychology of learning was still the focus of American psychological theory, and new mathematical tools began to be used in dealing particularly with the changes taking place during repetitive learning, and with related problems, employing statistical concepts in the special form of stochastic models after 1950. The concept of information theory began to enter in the early 1950s, but the new cognitive psychology did not become prominent until cybernetics emerged a few years later. Then information *processing* began to be substituted for information *theory* and to supersede the prevalent SR (stimulus-response) approach. Thinking in terms of models, such as those of short-term and long-term memory, and feedback relations within them, made it possible for the new cognitive psychology to be recognized as something different from a rediscovery of old ways in psychology's history of studies of perception, memory, imagination, language, and thinking, but now there was a new orientation toward these topics, and new instrumentation, such as the high-speed computer, to give a different view to these old topics. The revitalization at the core of psychology was remarkable, and refreshed as well the many specialized areas of developmental psychology, the psychology of personality, social psychology, and clinical psychology.

In retrospect, the roots of the change can be found much earlier, but the suddenness and widespread adoption of the new standpoint shows how unpredictable such events are in the history of science. Cognitive psychology flourished around 1960 and by 1985 was 25 years old—a very long time for a popular orientation in psychology to retain its ascendancy over novel successors.

The picture of psychology at present, if it is accepted that psychology is still in a preparadigmatic stage as mentioned earlier, is that there will be a kind of pluralism for a time while theories are being tried out over somewhat limited domains of data. The procedures of refinement of the accuracy of operations of the presenting of results so that they can be replicated by others, following the generally acceptable logic of science, will continue to make some gains through accumulation. The collection of new facts will not be simply descriptive, because the smaller theoretical models will produce a selection of facts and perhaps result in significant discoveries in the special fields under investigation.

Today psychology can be conceived of as a large family of many psychologies, unified by social practice and the disciplinary structure of universities. The more typical development has been of subdisciplines interstitial between psychology and other scientific subspecialties, resulting in sensory science; cognitive science, with such subspecialties as linguistics and artificial intelligence; behavioral biology; neuroscience; and life-span development. All of the applied branches are by their very natures interdisciplinary. Why, then, should not their theories be expected to reflect the common problems with which they are involved with their associates from other disciplines?

If the growing edges of science lie in these closely related subspecialties, new disciplines may emerge, as indeed they have in such areas as biochemistry and astrophysics. For example, social psychology belongs both to psychology and to sociology—not a different social psychology for each—and neurocognition, which would be a part of neurology and cognition and brain sciences in general, would seem to be a most natural arrangement. The social organization of sciences may be maintained by tradition, loyalty, and commitment. Science as an intellectual enterprise has criteria of self-renewal and change within its very structure, and the free-market interplay of theories may prove to be the path that will be taken.

In the 100 years of modern psychology, despite many differences, psychologists have continued to have enough in common that we can all identify ourselves as part of a common tradition, distinguishable from related ones such as anthropology and sociology, even though we are interrelated in many ways.

What binds us together are agreement upon a preference for experimental approaches, the use of appropriate statistics in determining the reliability and validity of findings, and a preference for theories that integrate such findings. We have attained status as a legitimate social science and also a biological science, depending upon the subfields under consideration. While we may expect changes, our role as a legitimate member of the scientific disciplines appears to be assured.

Sections of this foreword are drawn from Hilgard, E. R. (1987). *Psychology in America: A historical survey*. New York: Harcourt, Brace, Jovanovich.

Preface

There is a strong need among people to ritualize anniversaries. In 1992, for example, we Americans commemorated Columbus's fortuitous discovery of the "New World," 500 years after he and his mates set foot on dry land in the Caribbean. Also, in 1992, American psychologists celebrated the 100th anniversary of the founding of the American Psychological Association. Traditionally, as the world is about to turn the corner on a new century, there is a palpable increase in people's curiosity about where we have been and where we might be going. When the planet is about to enter a new millennium some people seem to get overly nostalgic; a few scholars may wax philosophically; and other folks get a little giddy as they look forward to a new era. But, to put these matters into perspective, it is only a birthday based on an artificial time continuum. Nevertheless, it is a good excuse for all of us—especially those interested in the science of the mind—to reflect on our past and speculate about our future. That simple premise is the basis for the three-volume series of the future of mind sciences. The first volume, *The Science of the Mind: 2001 and Beyond*¹ appeared in 1995 (in collaboration with Dominic Massaro) and, like the present publication, contained brief essays by distinguished scholars on their visions of the future.

The 20th century has been extraordinary by all measures. Those who lived through this period, such as Ernest Hilgard (born in 1903) who wrote the foreword, and those who have been interested in its intellectual history know, certainly, that the face of science and technology (as well as other features of our lives) has changed radically over the past 100 years and, if the past is an indication of the future, changes during the next century might be even more extreme. (A contrary view is that we have changed the world in such immoderate ways that humans may feel like aliens and, as such, rebel against the artificial and return to more earthy pleasures, which include the appreciation of the natural environment, the arts, music, literature and, in general, a

humanistic Renaissance.) One reason the game of predicting the future is so fascinating is the outcome is never certain. However, given an understanding of the resources available and the forces that shape our destiny, an approximation—or an “educated guess,” in today’s parlance—of what the world might be, is possible. Furthermore, the more serious side of guessing what the future might hold for us is that we might avoid some potentially disastrous pitfalls and maximize the chances for survival and enlightenment. Futures are created, not preordained.

In the initial letter of invitation to prospective contributors to this book I suggested that the papers should be tightly composed, well reasoned, and readable by the informed layperson, as well as stimulating to colleagues and students. When the essays began to arrive it was clear the contributors, in addition to being outstanding scientists, were also talented writers—you will find some beautifully crafted papers in this book. In a few cases the writer pleaded ignorance about the future on the basis of past histories of predictions, but then, happily for us, plunged ahead and made several predictions anyway. Others wrote fantastic papers, bordering on the whimsical or sci-fi to express their thoughts about the future, while others selected a single theme and told of their expectations for its fate. In every case, the contributors to this book worked enthusiastically and with good humor on the task. Even though some of the papers seem to picture the other side of Oz, they all have very serious implications for the type of psychology we might find on the other side of 2000. Each author has won my profound thanks, but more important, has won the appreciation of generations of people who will learn how scientists did psychology throughout the 20th century and what those scientists predicted the next century might bring.

I had an idea as to what types of essays individual authors might produce based on their previous writings. And, for the most part, those expectations were fulfilled. In a few instances however, unexpected essays were submitted, which, in every case, added richness to the book. Each chapter was edited for content, aptness, and style and returned for revision. Thus, the eventual organization of the book was fashioned around the chapters submitted rather than forcing chapters to conform to a preconceived taxonomy. Three major categories emerged: Consciousness and the 21st Century has chapters by Bernard Baars, Carl Sagan and Ann Druyan, Richard Thompson, and Endel Tulving. The second section, Brain and Mind in the 21st Century, contains chapters by Edward Smith, Michael Posner and Dan Levitin, Alan Gevins, Karl Pribram, and Michael Gazzaniga. The third major category is called Psychology

(Memory, Theory, and Cognition) in the 21st Century and has chapters by Henry Roediger, Gay Snodgrass, Jerome Kagan, George Sperling, Neal Miller, and Hans Eysenck. The final section, called *The Science of the Mind*, has an overview chapter by Robert Solso.

Many people contributed to this book who are gratefully acknowledged here. Each of the authors is recognized for writing an interesting and prophetic chapter. In addition, Jerry Weinstein of The MIT Press gave support, encouragement, and moved the manuscript through what to an outsider seems like a labyrinthine editorial process. Betty Stanton, senior editor of The MIT Press/Bradford Books, a most sagacious editor indeed, has offered her wise editorial advice and counsel. I am grateful to my graduate students at the University of Nevada, Reno, who read and discussed most of the manuscripts and gave interesting perspectives on the contents. Finally, I wish to acknowledge Kim Beal, who has served as a resident editor who made cogent remarks on most of the chapters, including my own, and to Alan Rees, who not only read and commented on the chapters but also drafted a copy of the biographical sketches.

The assembling of talented psychologists and brain scientists for this volume was not done (entirely) for intellectual amusement but to give insight into the types of worlds which might emerge. These eminent scholars have given you their visions of a world that might be, not what is determined to be. Fortunately, with knowledge of the existing resources, the forces of society and politics, and comprehension of nature, it is possible to plan future societies in which the community of humans may achieve greater understanding of who we are and where we are going. A major goal of *The Science of the Mind* project is to encourage contemporary scientists and scholars to consider alternative worlds, to make choices about the allocation of resources (both human and physical), and to develop plans to maximize the actualization of favorable conditions and avoid many of the problems we endured during the 20th century.

R.L.S

May 1996

Lake Tahoe, Nevada

Note

1. Solso, R. L., & Massaro, D. W. (Eds.) (1995). *The science of the mind: 2001 and beyond*. New York: Oxford University Press.

I

CONSCIOUSNESS AND THE 21ST CENTURY

1

Psychology in a World of Sentient, Self-Knowing Beings: A Modest Utopian Fantasy

BERNARD J. BAARS

There is no more important quest in the whole of science probably than the attempt to understand those very particular events in evolution by which brains worked out that special trick that enabled them to add to the cosmic scheme of things: colour, sound, pain, pleasure, and all the other facets of mental experience.

—Roger Sperry

Suppose that in the 21st century, psychology—or more likely biopsychology—actually began to work? And suppose that it began to deal with the reality of our personal experience? What might that world be like?

Western thought began with the Socratic injunction “Know thyself” and Asian philosophies have pursued a similar aim throughout history. Science has allowed us to know the world with stunning depth and accuracy, changing the conditions of life for humans more than any other development since the invention of agriculture, 10 millenia ago. But it has been extraordinarily difficult to turn the powerful lens of science inward, toward ourselves, so much so that for most of this century psychology and the brain sciences have tended to evade the very existence of our own experience. That denial is now on the way out, and we are beginning to gain a firmer understanding of consciousness, voluntary control, and even self.

Many scientific students of consciousness believe that the necessary evidence and theory have begun to build in a sustained way in the last 10 years or so, and that this process is now accelerating. In a decade or so we may have an early biopsychology of consciousness. Not an ultimate understanding (if such a thing exists), and certainly not one that is immune to further shifts, but some genuine scientific knowledge nevertheless. If that prediction is anywhere near

accurate, humans will be confronted with something entirely new, because the only sciences we know are sciences of *otherness*, of things outside of ourselves. What would this unprecedented event mean for everyday life?

Is It a Clearly Defined Phenomenon?

For many decades it was customary in psychology to dismiss personal experience, because there seemed to be no way to obtain solid evidence about it, and theory was simply inconceivable. Many of us now believe that that was a self-fulfilling prophecy, a sort of avoidant confirmation bias. If you think a subject is impossible to study, you won't study it, thereby "proving" that it is impossible to study. Avoidant thinking is not unusual in the history of science, but like any unreasonable fear, we must eventually confront the things we have been avoiding, or give up the whole enterprise.

Table 1.1 shows 25 polarities that have been found useful in the last decade or two in studies of human cognition. They range from "attended vs. unattended information" and "explicit vs. implicit memory," all the way to "waking EEG compared to deep sleep and coma." Each polar pair of terms is fundamental for a considerable scientific literature. To a considerable extent, all 25 polarities can be captured under two major headings: conscious and unconscious phenomena. The realization that we have been "speaking consciousness" all of our lives has been coming home to many of us, and the consensus seems to be that we might as well call the thing by its proper name. That is not to say that we have a well-established grasp of the problem today, but we seem to be making incremental progress, just as we have been able to make progress on other formidable topics such as attention, perception, and language.

Evidence

In its major features consciousness is not a subtle thing. When humans are not conscious, our bodies wilt, our eyes roll up in their orbits, our brain waves become large, slow, and regular, and we cannot read a sentence like this one.

While the outer signs of consciousness are pretty clear, it is our inner life that counts for most of us. At this instant you and I are surely conscious of some aspects of the act of reading—the shape of *these letters* against the white texture of *this page*, and the inner sound of *these words*. But we are probably not

Table 1.1
Polar pairs of terms used in contemporary psychology

Associated with Consciousness	Associated with Unconsciousness
1. Explicit cognition	1. Implicit cognition
2. Immediate memory	2. Long-term memory
3. Controlled processes	3. Automatic processes
4. Novel, informative, and significant information	4. Old, predictable, or insignificant information
5. Attended stimulation	5. Nonattended stimulation
6. Declarative memory (beliefs)	6. Procedural memory (skills)
7. Autobiographical memory	7. Semantic memory
8. Supraliminal stimulus processing	8. Subliminal stimulus processing
9. Recalled memories	9. Stored memories
10. Explicit goals, decisions, problem-solving	10. Implicit goals, decisions, incubation
11. Available memories	11. Unavailable memories
12. Stimuli in implicit learning	12. Learned pattern in implicit learning
14. Rehearsed item in working memory	14. Unrehearsed items in working memory
15. Current images	15. Images in memory, or which are automatic
17. Wakefulness and dreams (rapid EEG)	17. Deep sleep and coma (slow EEG)
18. Wide access to mental functions	18. Local access to mental routines
19. Voluntary control	19. Automatisms used by voluntary system
20. Explicit reasoning	20. Unconscious inferences
21. Focal contents	21. Fringe experiences (feelings of familiarity, etc.)
23. Auto-noetic (E. Tulving)	23. Noetic memory
24. Intact reticular formation and intralaminar nuclei	24. Lesioned reticular formation and intralaminar nuclei
25. Explicit ideas, beliefs, etc.	25. Presupposed or implicit ideas

aware of the touch of the chair, of a certain background taste, the subtle balancing of our body against gravity, a flow of conversation in the background, or the syntax of *this phrase*; nor are we aware of the fleeting present of only a few seconds ago, of our affection for a friend, or the multiple meanings of words, as in *this case*. There is nevertheless good evidence that such unconscious events are actively processed in our brains, every waking moment of the day.

The contents of consciousness include the immediate perceptual world; inner speech and visual imagery; the fleeting present and its fast-fading traces in immediate memory; bodily feelings like pleasure, pain and excitement; surges of emotional feelings; autobiographical memories; clear and immediate intentions, expectations, and actions; explicit beliefs about oneself and the world; and concepts that are abstract but focal. In spite of decades of behavioristic avoidance, few would quarrel with this list today.

These examples illustrate the meaning of the word “consciousness” we are trying to understand: that is, focal consciousness of easily described events, like “I see a printed page,” or “I am imagining my mother’s face.” A great body of objective evidence shows that conscious contents like these can be reported *as conscious* with great accuracy under the right conditions. These conditions include immediate report, freedom from distraction, and some means for corroborating the report. These are standard laboratory conditions in thousands of experiments in perception, memory, language, attention, and imagery. In all these conditions subjects will tell us they are conscious of certain events, and it is always a good idea as a psychologist to listen to your subjects.

The Rebirth of Consciousness in Science

Outside of psychology the greatest impetus to consciousness research has come from prominent researchers in biology and brain science, notably Francis Crick, Gerald Edelman, Rodolfo Llinás, Michael Gazzaniga, and many others. Crick, Edelman, and Gazzaniga in particular have been influential in encouraging brain researchers to work on the empirically accessible parts of the problem, of which there are now quite a few.

In psychology the renewal of consciousness studies has its roots in the cognitive revolution of the 1970s, which set the framework for much of contemporary psychology. After a long series of very careful experiments on what

we now call “working memory” and “selective attention,” psychologists found themselves making careful *inferences* to explain the observations. Making inferences about explanatory entities is a standard gambit in science, of course. “Atoms” were entirely inferential to Dalton, though today we can observe atomic lattices with electron micrographs. The existence of the planet Pluto was first inferred from perturbations in the orbit of the more visible outer planets. The vast depth of geological time was inferred from the fossil record, the layering of rock on an exposed mountain side, and the decay of carbon 14. The list of inferential constructs simply goes on and on. Science simply cannot grow without making inferences, “going beyond the information given,” as Jerome Bruner has phrased it. There is no reason for psychology to be different.

George Miller’s famous paper on the “magical number seven, plus or minus two” provides a nice illustration. It appeared in 1956, at the start of the “cognitive revolution” in psychology. Alfred Binet, the French pioneer psychologist, already knew that people cannot keep in mind more than five to nine unrelated numbers. It seemed to be a more-or-less random fact, a little like the observation that many people, when asked to choose a number between 1 and 10, will choose 7. Miller was smart enough to notice that we are actually “persecuted by an integer”—that the magic number reappears over and over again in our evidence, whenever we study the way human beings deal with unrelated sets of words, numbers, colors, short phrases, tones, and rating categories. Of course, we can always “chunk” sets of unrelated numbers by learning predictable number sequences, chunks like 1900, 1914, and 1776. But then the immediate memory limit turns into seven plus or minus two *chunks*. The magic number simply pops out in a different form.

A vast outpouring of research followed Miller’s 1956 paper, by and large confirming a powerful and unexpected pattern of results that stayed solid across many different kinds of items and conditions. Eventually that solid behavioral pattern came to be understood to reflect the operation of a “buffer memory,” an interface between the timing of events in the outer world and in memory. *A sentence like this one cannot be understood unless somehow we can store the underlined words for several seconds, while we wait for the rest of the sentence to arrive, with the information needed to complete a coherent thought.* The words *a coherent thought* in the last sentence came perhaps 5 to 10 seconds after the subject, *a sentence like this one*. Yet the later words must be integrated with the earlier ones in a way that makes sense. We can even switch the beginning and ending of the sentence by writing a paraphrase: *The information needed to*

complete **a coherent thought** may arrive several seconds after the subject of **a sentence like this one has** been introduced. Our brains are constructed so that the order of subject and object does not matter very much. But that seems to imply the existence of a short-term memory for sentences, one that allows words arriving at different times to be organized in a meaningful structure, regardless of which arrived first.

Now we can do many other things with immediate memory. We can easily interfere with our ability to understand the italicized sentence above, by reading it again while keeping in mind a few arbitrary numbers, let's say 31, 15, and 11. Try it—you'll see that it now becomes quite difficult to understand. But that suggests that short-term memory for language and for numbers is the same thing, or at least that they make use of overlapping mental resources. Dozens of other manipulations are possible, and they create a robust and broadly understandable pattern of results.

Constructs like "short-term memory" were soon followed by many others: "semantic networks," "imagery," "implicit knowledge," "mental grammars," and the like. Today we routinely use our data to index underlying explanatory entities. It must be done carefully, to avoid circularity. But once we have found a vast array of evidence indicating there is such a thing as immediate memory, or elemental oxygen, we can use the construct to explain new observations, to tie all of the evidence together into a cohesive story.

Making inferences about psychological constructs is a crucial step for understanding consciousness. Once psychologists developed the habit of postulating inferred constructs to explain and simplify a solid pattern of observations, there was no more *principled* objection to thinking of consciousness in just the same way, as a theoretical construct based on reliable, public evidence.

The cognitive approach to consciousness can be understood strictly in terms of public evidence, the things we can all agree upon. But consciousness is special, of course. Each of us has some useful access to our own experience that is not shared by others. Some things, especially perceptual events, can be accessed consciously with extraordinary precision—witness 150 years of solid scientific work in sensory psychophysics. Mental images can be reported quite accurately, as shown by Stephen Kosslyn and others. Inner speech has been found to be quite reliable in studies of mental problem solving and spontaneous thought.

Other aspects of mental functioning are very hard to operationalize with private reports. Long-term intentions are notoriously difficult to report accu-

rately; witness the self-deception we all engage in annually with New Year's resolutions. We can feel utterly convinced that we intend to drop those 15 pounds, but apparently it is difficult to distinguish between long-term goals that will be carried out and those that are destined to fade in a few weeks.

From a scientific point of view we cannot share our personal experience directly with one another, so that we deal publicly only with *descriptions* of experiences. But for many well-studied phenomena the subjective and objective evidence converges so well that the distinction has little practical meaning. We can all understand perceptual events, like the reader's experience of *this page*, from a subjective or an objective point of view. Indeed, perception researchers tacitly recognize the close mapping between subjective and objective descriptions of experimental stimuli, when they routinely "run themselves" in any new experiment. We could pretend that perceptual events do not apply to our own experience, and that we are only exploring the objective behavioral and brain processes of an utterly unknown species infesting the surface of the fourth planet of Sol. But that pretense is not necessary, because in most experimental situations the inside and outside perspectives dovetail so well. This pattern of convergence persuades me at least that in practice, the famous gap between mind and body is a bit of a myth.

Contrastive Analysis; Treating Consciousness as a Variable

Earth gravity is so constant in our experience that its very existence goes unnoticed. Historically, it took a great effort of imagination to understand that gravitational attraction could be different elsewhere in the universe than on earth. Newton's ability to take that imaginative leap made it possible to solve the ancient mystery of planetary motion. In the same way, 19th-century naturalists had to *learn* to think of species as changeable—animals and plants appear to be immutable, after all. Prior to Darwin few naturalists believed in the gradual evolution of species, but seeing species as varying over time made it possible to understand the living world.

William James, along with most educated people in the 19th century, could not imagine consciousness as a variable, because he believed it was the *only* proper topic for psychology. Nothing was in any way commensurate with it. Yet he knew about the principle of comparison. In discussing the functions of consciousness he wrote:

The study . . . of *the distribution* of consciousness shows it to be exactly such as we might expect in an organ added for the sake of steering a nervous system grown too complex to regulate itself. (Emphasis added)

That is the essence of the experimental method, but it was not conceptually possible to apply it to conscious experience until surprisingly recently, perhaps as recent as the last 20 years. James could not apply the standard scientific strategy of comparison to consciousness as such, because he believed that it was the sole instance of mentality, while *unconscious* events were “only physical.” Most people in the 19th century could not imagine that conscious and unconscious events could be compared. “Unconscious intelligence” was a bizarre oxymoron to our greatgrandparents. Consciousness was the crown of human reason; unconsciousness was a different metaphysical substance.

We know something about the reader’s experience of *this word*, and we have some evidence about its unconscious representation when the same word is provided subliminally, or in an unattended channel, or when it has become automatic through repetition. Indeed, we now have robust evidence that for many conscious events, we can find unconscious analogues. We know that unattended speech is processed at least up to the meaning of individual words, that subliminal words can stimulate mental processes related to the meaning of those words, and that the reader’s mind even now is computing the syntax of *this sentence*, quite unconsciously. That is, we can find comparison conditions in which the content of input is held constant, while consciousness can be varied. Thus we can treat consciousness as a variable across a great variety of cases.

When we run these comparisons in detail, we obtain the pattern of results shown in table 1.2. The left column shows that conscious processes tend to be computationally inefficient, which is to say that a conscious symbolic computation, like mental arithmetic, seems to be slow and vulnerable to error and interference. Of course the range of conscious contents is vast: imagine all the sensory contents, all the mental images, the memories, ideas, and so on. Conscious events are always internally consistent, they occur serially, one after the next, and have limited capacity. In the right-hand column are unconscious comparison conditions. If we compare automatic computation, like analyzing the syntax of a sentence, with conscious computation, it is clear that automatic processes tend to be highly efficient, fast, and robust. There is much evidence

Table 1.2
Capabilities of conscious and unconscious processes

Conscious processes	Unconscious processes
1. Computationally inefficient: e.g., mental arithmetic.	1. Very efficient in routine tasks: e.g., syntax.
Many errors, relatively low speed, and mutual interference between conscious processes.	Few errors, high speed, and little mutual interference.
2. Great range of contents.	2. Each routine process has a limited range of contents.
3. High internal consistency at any single moment, seriality over time, and limited processing capacity.	3. The set of routine, unconscious processes is diverse, can sometimes operate in parallel, and together has great processing capacity.

that each unconscious patch of neural tissue is specialized in just a single function, but that all the specialized patches operate in parallel, all at the same time. Taking the unconscious specialized tissues of the brain together, they have vast processing capacity.

This pattern of evidence has been interpreted in a framework called Global Workspace (GW) theory (Baars, 1988). GW theory presents a “theater model” in which consciousness requires a central workspace, much like the stage of a theater. The theory is based on the belief that, like the cells of the human body, the detailed workings of the brain are widely distributed. There is no centralized command that tells neurons what to do. Just as each cell in the body is controlled by its own molecular code, the adaptive networks of the brain are controlled by their own aims and contexts. To organize this vast distributed domain, there is a network of neural patches that work together to display conscious events. Today the best candidates for these loci of conscious experience may be the sensory projection areas of the cortex, where the great neural radiations coming from the eyes, the ears, and the body first reach the surface of the brain. A few small structures of the core brain stem and midbrain are essential to consciousness, but great quantities of tissue elsewhere in the brain can be lost without causing a loss of conscious experience. Conscious contents appear to be disseminated globally to a great multitude of networks throughout the brain that are unconscious but that have observable conscious consequences downstream.

As it happens, all unified theories of cognition today are theater models. GW theory derives from the integrative modeling tradition of Allan Newell,

Herbert A. Simon, John Anderson, and others in cognitive science. It is consistent with models of working memory by Alan Baddeley, of the mind's eye by Stephen Kosslyn, of explicit knowledge after brain damage by Daniel Schacter and Morris Moscovitch, the thalamocortical searchlight elaborated by Francis Crick, and society models outlined by Michael Gazzaniga and Marvin Minsky. The brain implications of GW theory have been explored by James Newman and myself. British mathematician John G. Taylor and others are working to apply modern "neural net" models to the problem. The convergence of ideas today is simply astonishing.

It seems now that the failure of 19th-century psychology was not primarily a matter of *empirical* difficulties in dealing with questions like "imageless thought." Rather, it seems now that it was a *conceptual* inability to understand the very idea of an intelligent unconscious. I believe that James's often desperate inner struggles with mind-body issues stem largely from this single conceptual block: his inability, and that of his entire generation, to deal with consciousness as a variable with a natural comparison condition. Such conceptual blocks are of course utterly routine historically. One can see the history of science as a process of grappling with, yielding to, and finally overcoming conceptual blocks much like this one.

Behaviorists after James's death in 1910 rejected the whole topic because it seemed rife with endless, useless perplexities. They made little progress on it because they avoided *both* conscious *and* unconscious processes. They, too, were unable to see consciousness as a variable. Neither James nor Skinner could apply the experimental method to the most humanly important topic of all.

If we live at an historic time for the study of human experience, it is not just that we have more facts but because we can treat consciousness as something that exists to a greater or lesser degree, something with a comparison condition. When we repeat a single word to the point where its meaning fades from consciousness, consciousness is a *dependent* variable. When we observe the effects of an anesthetic on learning, consciousness is an *independent* variable. Most theoretical constructs in science can be explored either as dependent or independent variables, and consciousness is no exception. In a slightly different vocabulary, conscious experience is both a cause and an effect. It participates in a network of inferred and observable relationships with the rest of reality.

Today, spectacular new imaging techniques are giving us remarkable insight into the living brain. Experimental methods have been honed to great