Artificial Intelligence and Literary Creativity

Inside the Mind of BRUTUS, a Storytelling Machine

Selmer Bringsjord David Ferrucci

August 30, 1999

To Katherine and Alexander,
who show me super-computational creativity every day (S.B.)

and

To my parents, Antonio and Connie,
for their unyielding love, support, courage, and intellectual energy,
which have inspired me for as long as I can remember (D.F.)

Contents

Acknowledgments

This book and BRUTUS/BRUTUS$_1$ are the result of years of effort on our parts, but lots of people helped, and, accordingly, we have lots of debts. They are too many and too great to pay, but we'll try at least to mention most.

Thanks are due to Harriet Borton for her matchless LaTeX and TeX expertise, without which there would be no book. LaTeX and TeX are amazing things, and Selmer now wonders how he ever wrote a syllable without them. (Selmer assumes that Dave is now entirely converted to the notion that precise writing should be like computer programming.)

We thank the tremendous team at Lawrence Erlbaum for their patience and professionalism: Anne Duffy, Art Lizza, Linda Eisenberg, and copy editors we know only by meticulous marks. Before leaving LEA, Ray O'Connell signed us up for this project, and gave valuable guidance when the following pages were but hazy dreams.

Thanks are due to the *Journal of Experimental and Theoretical Artificial Intelligence* for permission to use, in Chapter **??**, parts of the paper "The Case Against AI From Imagistic Expertise" [30]. We are likewise grateful to *Behavioral and Brain Sciences* for permission to use, in Chapter 1, parts of the review [35] of Margaret Boden's *The Creative Mind: Myths and Mechanisms* [20], and to the French publisher Arras for allowing us to use, in Chapter **??**, translated parts of "Pourquoi Hendrik Ibsen Est-Is Une Menace pour La Littérature Générée Par Ordinateur?" [33].

Dave is indebted to IBM and, specifically, IBM's T.J. Watson Research Center, for time to work on BRUTUS and BRUTUS$_1$, and the book itself. Likewise, Selmer is grateful for the support Rensselaer has given him through the years. The Minds *&* Machines Laboratory, which Selmer directs at Rensselaer, has in particular

been and continues to be a great resource. This lab only runs at present because students like Tom Poltrino, Clarke Caporale, Micah Clark, Faris Gammoh and others (e.g., Jack Yuzyenko and Claudia Hunter) make it possible. (In addition to his lab responsibilities, Tom Poltrino graciously helped with some of the images that follow.) The world has heard about BRUTUS and all the "big" questions he raises about the human mind in large part because people at RPI like Megan Galbraith and Nancy Connell and Tom Torello are gifted at securing opportunities for us to communicate through the mass media. Of course, we are indebted as well to the media itself, who to their considerable credit realize that there is a thirst out there in the laic world for answers to "big" questions. Will the machines we humans build eventually leave us in the dust? Will even our great feats of creativity fall as Kasparov's achievements have fallen? How is it that machines will eventually match us, given that we have *originated* such things as *Julius Caesar*? These are questions for the new millennium; arguably, these are *the* questions for the next century.

Many colleagues have helped with research related to this book. In particular, a number of colleagues have worked in the "Autopoeisis Project," initiated by Selmer Bringsjord and Dave Porush in 1990 with the aim of getting a machine to autonomously[1] write sophisticated fiction by the turn of the century. In the years before our work on BRUTUS, Dave Porush provided innumerable insights about "rule-breaking" belletristic literature. These insights helped move us toward rejecting Church's Thesis. (See Chapter **??**.) Having Marie Meteer at Rensselaer for three years to specialize in natural language generation as part of the Autopoeisis team was invaluable. BRUTUS_1, unfortunately, lacks the genuine NLG capability (e.g., as described in [163]) that Marie can bring to a system. Our hope is that future incarnations of the BRUTUS architecture will have such capability. Chris Welty participated in Autopoeisis before going to

[1] BRUTUS_1 doesn't represent success for Autopoeisis, because BRUTUS_1 is not autonomous. BRUTUS, the architecture of which BRUTUS_1 is a partial implementation, has no provision for autonomy, or what might be called "free will." Hofstadter has suggested 6 requirements for a computational artifact to be deemed truly creative (see p. 411 of [111]). It seems to us that BRUTUS easily allows for implementations that satisfy this list — but these implementations would not have anything like real free will. So Hofstadter must be wrong. One of us has argued at length that no computational entity can have true autonomy: see the chapter "Free Will" in [40].

Vassar and made many brilliant contributions; he was involved in the original discussions about betrayal. Many student researchers in Autopoeisis came up with many helpful ideas through the years. Finally, thanks are due to the Henry R. Luce Foundation for the $300,000 that launched Autopoeisis. Support along the way also came from IBM, AT&T, and Apple Computer.

In recent years, trenchant comments have come from Ron Noel, whose "non-parameterized" approach to creativity (as reported, e.g., in [22]) is the opposite of ours, and has made Selmer think long and hard about the logicist approach to machine creativity advocated and exemplified herein. Ingenious ideas about human creativity have come from Michael Zenzen and Jim Fahey over the last 11 years. Michael has as keen a grasp of the nature of creativity as anyone on this planet. Others in the Creativity Circle at Rensselaer have been very helpful: Elizabeth Bringsjord, Ellen Esrock (whose unique and substantial contributions, courtesy of her fascinating book *The Reader's Eye*, will be discovered later by our readers), and Kathy Voegtle.

We are indebted to Spiral Design Inc. for the smashing image used on the cover. Observant and patient readers will find therein not only relevant technical elements seen later in the book, but also visual expression of the kind of serenity enjoyed by our BRUTUS, but definitely *not* by Shakespeare's Brutus. This image is also the core image for the aforementioned Minds *&* Machines Laboratory and Program at Rensselaer.

We are grateful to all those who have reacted to demos of implementations of BRUTUS that preceded BRUTUS_1. One particularly profitable demo took place at Brown University, sponsored by the Department of Cognitive Science and the Department of Philosophy. It was here that we were encouraged to turn to reader response theory for insights into how prose can be "engineered" to trigger desired psychological states in readers.

Finally, the debates. We are indebted to many debates and discussions with many people on the issues touched upon in this book. The majority of these thinkers advanced positions at odds with our own, and thereby helped sharpen our thoughts. They include: Margaret Boden, whose optimism about reducing creativity to computation stands in stark contrast to our calculated engineering; Marie Meteer, whose bottom-up approach contrasts with our top-down ap-

proach; Chris Welty, whose grasp of the undeniable fact (which we gleefully affirm herein) that BRUTUS$_1$ mirrors the minds of his two creators produces a stimulating cynicism; John McCarthy, who at *IJCAI 91*, in argument with Selmer, defended a view (viz., that explicit, formal definition of such concepts as betrayal is superfluous) that this book is in large part designed to overthrow; Jim Fetzer, whose semiotic brand of non-computationalism stands in contrast to BRUTUS' underpinnings; Pat Hayes, Stu Shapiro, Marvin Minsky, Ken Ford, four who have staunchly articulated the view that, despite Selmer's arguments to the contrary, cognition *is* computation; and Jim Moor, Robert Cavalier, Marvin Croy, and many others in the national "Computing and Philosophy" group, who pretty much every year supply objections to the kinds of arguments in this book.

After BRUTUS$_1$ will come BRUTUS$_2$, and then BRUTUS$_3$, and so on. That, at least, is the plan — a plan which, given all the help we've needed to this point, will, alas, require another wave of support to carry us on. Our debts will therefore mount. However, all errors herein and hereafter are decidedly our own.

Preface

The Marriage of Logic and Creativity

This book marks the marriage of logic and creativity.

While it may be true that incompatible humans often wed, there are doubtless unions of a less palpable sort that can never even come to pass. Such is the case, by the lights of many, for precisely what we are about herein. Creativity and logic? *Married*? Upon hearing of our plans, 7 years ago, to harness theorem-proving technology in order to create a computer program able to generate belletristic fiction, a rather famous novelist informed us that creativity and logic are as far apart as the east is from the west (and he proudly quipped that even such a metaphor is beyond logic, and hence beyond machines). Just an anecdote, yes, and just the opinion of one, but the truth of the matter is that this attitude is widely (and often fiercely) affirmed. Creativity is generally regarded to involve breaking the kind of rigid rules standing at the heart of logic; creativity, at least of the artistic variety, is commonly identified with the emotions and the "irrational." Freud, whose specific claims are today a bit tenuous, remains a seminal figure for often getting at least the tenor of things right. Freud believed that creativity is the link between art and play, and requires the "suspension of rational principles." He wrote that "The creative writer does much the same as the child at play. He creates a world of phantasy which he takes very seriously — that is, which he invests with large amounts of emotion — while separating it sharply from reality" ([93], p. 144). However problematic Freud's rather dark theories may be today, here he is simply making an observation that cannot be doubted. But the issue is whether such sophisticated play can in the end be reduced to logic. Is the play of Joyce and Tolstoy and Updike and Helprin and Morrison at

bottom logic in action?

Many used to ask a different question: Could a computer ever beat the best human chess player? With Kasparov brooding and Deep Blue and his silicon cousins improving every week, many are *now* asking: Could a computer beat all human grandmasters *time and time again in normal tournament play*? To this the both of us unhesitatingly answer in the affirmative (as should, we daresay, anyone who knows a thing or two about the dizzying ascension of raw computing power on this planet — though by our calculations it will nonetheless take a decade for machines to achieve such metronomic triumph).

Will Robots Soon Be Smarter Than Us?

So computers will soon be smarter than us at chess; nothing controversial here. What about everything else, creative activities included? Well, according to a quartet of recent books, there will soon be *nothing* that computers and robots can't beat us at. The books are

1. *Robot: Mere Machine to Transcendent Mind*, by Hans Moravec [165]
2. *The Age of Spiritual Machines: When Computers Exceed Human Intelligence* by Ray Kurzweil [142]
3. *When Things Start to Think* by Neil Gershenfield [97]
4. *March of the Machines: Why the New Race of Robots Will Rule the World* by Kevin Warwick [243]

We find many of the predictions in these books to be laughable.[2] For example, Moravec predicts that robots will get smarter and smarter so fast that 2040 will mark the advent of "fourth generation" robots, which will exceed us in all respects. They will not only do the kind of work we currently associate with robots (inflexible physical work; e.g., manufacturing) but will "run the companies and do the research" ([165], p. 125). The chief problem with predictions like this is that they are flatly inconsistent with the utter absence of

[2]The first three of these books have been recently reviewed by Colin McGinn [156]. McGinn explains that there is no reason to think that robots will have (to use the terminology of our Chapter 3 in this book) a point of view, and so it would be rather stupid to agree to have your "mind" downloaded into a machine.

machine creativity in the world today. The Index in Moravec's book contains not a single entry for creativity, and yet it takes some creativity to do research, does it not? And how about running IBM? Does that take some creativity? Every single piece of technology today is due to many creative humans who lived yesterday. Look around you now. How many artifacts can you count whose origins can be traced to one or more highly creative human beings? I'm (Selmer) typing this on my laptop at a favorite restaurant. In my laptop I see the reflections of Turing, and the entrepreneur Steven Jobs. I'm sipping a "Fresh Samantha" fruit smoothie from a multicolor jug, sold by a company whose Maine-based founders had a creative idea about bottling expensive smoothies with a hip label. The watch I'm wearing has the ancestral fingerprints of a thousand engineers. There is a light above me; in it I see Edison. There is a phone beside me; in it I see Bell. Obviously, I could go on — and on. So could you, we wager.

In particular, it's safe to say that we simply wouldn't have computers and robots around today were it not for countless strokes of human creative genius. And yet Moravec, whose vision is a computation-driven one, is silent on creativity. Very interesting. Where are the AI labs in which computers are creating things? Where are the labs in which computers are creating new branches of mathematics, new modes of music, great novels, novel scientific theories, and so on? Where are they? They do not exist.[3]

We do not want to give you the wrong idea, reader. The two of us are quite optimistic about what AI can achieve. For example, we're inclined to believe that

- NASA will run successful missions to Mars and other planets largely on the strength of "immobots," HAL9000-like AIs that will control

[3]Selmer is at work on a book-length antidote to the fanatical sanguinity seen in these four books. Part of this antidote consists in good old-fashioned fact-checking. For example, Warwick tells us that machines that can beat us on IQ tests already exist. Really? Selmer is willing to compete against any present-day machine on the Weschler adult intelligence test, and to wager serious money that he can win. This test includes a task in which the test taker must assemble a coherent story from jumbled diagrams that represent snapshots of the action. What machine can do *that*? The test also includes general commonsense reasoning questions that even CYC would have trouble with. Nonetheless, a robot able to excel on this IQ test is under construction in Bringsjord's Minds *&* Machines Laboratory.

the ships in question.

- AI-controlled cars, safer than their human-controlled counterparts, will be available sooner rather than later.
- General house-cleaning robots will arrive — again, sooner rather than later.
- Even now, the bulk of medical diagnosis can be carried out by computers, at an accuracy level surpassing all but a small number of human diagnosticians. In the future, machine diagnosis will reach a point where it is downright *irrational* to consult a human M.D. first.

And so on. But notice that the kind of list we have in mind doesn't require any creativity to speak of. (Sorry, diagnosticians.) So we still have the question before us: What about creativity? Robotic drivers may be securely in our future, but Einstein, Gödel, Tolstoy, Turing, Shakespeare, Plato, Cantor, ... — could machines ever reach *their* rank? Could we ever build a genuinely creative machine?

We seek to answer "the creativity question" not from the comfort of our armchairs, but from the workbenches in our laboratories. Specifically, we seek to ascertain whether or not literary creativity is the sole province of humans by attempting to *build* artificial authors. The first fruit of our labor, 5 years in the making (with another half-decade prior to this one devoted to less ambitious systems), is BRUTUS, a storytelling agent specializing in narrative that involves betrayal first and foremost, and also self-deception and other literary themes. The mind of BRUTUS is revealed in the book you're holding.

From Chess to Literary Creativity

In our experience, the public is quite comfortable with the notion that a machine can play invincible chess — because even those who know nothing of the niceties of search algorithms intuitively grasp the mathematical fact that chess, at bottom, is utterly mechanical, that if one can "look far enough ahead" the game becomes trivial. On the other hand, given the reaction of the public to BRUTUS_1's prowess as reported in the media (as evidenced by a persistent stream of rather emotional communication we receive), we think it's safe to say that while we (and many other AIniks, e.g., Douglas Hofstadter [111]) merrily press ahead in the attempt to reduce creativity to computation, the lay mind is fundamentally disturbed by the prospect

of creative machines.[4] This is probably because they realize, intuitively, that the future described in the quartet of books cited earlier can come to pass *if* machines become creative. In presenting the anatomy of BRUTUS's brain herein, we will soothe the souls of those who, hearing about his exploits, fear that humans will soon have nothing over machines. It will become crystal clear in what follows that BRUTUS should give his human creators rather a lot of credit. Put in terms of our terminology, we say that BRUTUS has weak, rather than strong, creativity. (Of course, there are people out there at the other end of the spectrum: people who think that a machine that creates genuine literature is right around the corner. Figure 1 encapsulates this attitude.)

What we call "strong creativity" is what might be called "raw origination." Raw origination is akin to creation *ex nihilo*, and though this form of creativity may well be impossible, the fact of the matter is that the *concept* of creating something from nothing is very real not only to monotheists, but also to many hardheaded scientists who have pondered creativity. The paradigmatic example is Margaret Boden, arguably the world's leading authority on computational creativity. Boden [19] distinguishes between a brand of creativity associated with the novel combinations of old ideas (she gives the example of the Lennon/McCartney arrangement of "Yesterday," marked by the unprecedented combination of a cello with music of this type), and a type of creativity in which something utterly and completely new is produced (e.g., non-Euclidean geometry, wherein the sum of the interior angles of a triangle is *not* 180 degrees). Computers, of course, have no trouble with the former type of creativity. The latter type is somewhat more difficult for them. It's exceedingly hard to see how a computer could, say, autonomously discover a new class of numbers through new proof techniques, which was one of Cantor's novel achievements.

The distinction between strong and weak creativity isn't a new one. When Alan Turing, one of the grandfathers of computer science and AI, proposed that if a machine could pass his famous "imitation game" (in which a computer passes if it's linguistically indistinguishable from a human; the game is now known as the "Turing Test"), we humans should immediately conclude that such a machine can gen-

[4] It's important to note that we don't think the reduction can be pulled off. Hofstadter does.

Figure 1: Nonchalance Regarding BRUTUS$_1$'s Descendants. Roz Chast © 1996 from The New Yorker Collection. All Rights Reserved.

uinely think, he considered an objection from Lady Lovelace that was given on the strength of raw origination. She argued: "Computers will never be creative, for creativity requires *originating* something, and this is something computers just don't do. Computers do what they are programmed to do, nothing more." (Turing presents his imitation game, and discusses the Lovelace objection, in his [236].)

Suppose for the sake of argument that Lovelace is correct. Even so, the other sense of creativity, "weak creativity," remains intact. Weak creativity has its roots in the "operational" notion of creativity devised by psychologists. For example, E. Paul Torrance, who more than any other psychologist has probed the nature and concrete signs of creativity, holds that x is to be deemed creative just in case x scores well on the dominant test for creativity in children and adults: The Torrance Tests of Creative Thinking.[5] This test comes in both "visual" and "verbal" forms. In the visual form, test takers are asked to draw pictures (often by enriching existing sketches); in the verbal form, test takers are asked to write — creatively. For example, one of the activities subjects engage in on the verbal test is the following.

> Most people throw their tin cans away, but they have thousands of interesting and unusual uses. In the spaces below and on the next page, list as many of these interesting and unusual uses as you can think of. Do not limit yourself to any one size of can. You may use as many cans as you like. Do not limit yourself to the uses you have seen or heard about; think about as many possible new uses as you can. (From the verbal version of [233].)

After the Torrance Test is administered, one can send it out to be professionally judged. Our aim on the problem of literary creativity is to build an artificial agent capable of producing stories that would be scored as highly creative by human judges in the dark as to whether or not the stories they receive are from humans or machines. One of us (Bringsjord) has refined this scenario into what he calls the "short short story game," or just S^3G for short. The idea is simple; it is summed up in Figure 2. A human and a computer compete against each other. Both receive one relatively simple sentence, say: "Barnes kept the image to himself, kept the horror locked away as best he could." (For a much better one, see

[5] See [233] for the test itself. For reviews of the test, see [50], [227], [235].

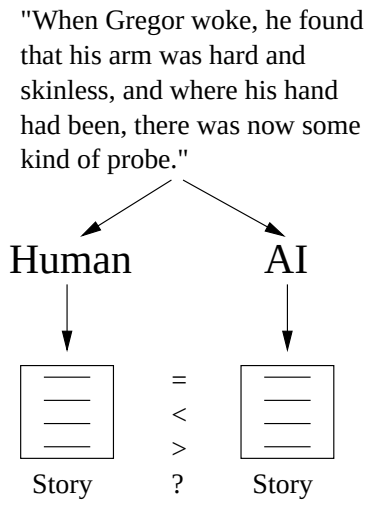

Figure 2: The Short Short Story Game, or S^3G for Short.

the "loaded" sentence shown in Figure 2.[6]) Both mind and machine must now fashion a short short story (about 500 words) designed to be truly interesting; the more literary virtue, the better. Our goal, then, is to build an artificial author able to compete with first-rate human authors in S^3G, much as Deep Blue went head to head with Kasparov.

Unfortunately, this goal is too tough to reach, at least for the foreseeable future; it may even be a goal that is forever beyond the reach of a machine. (Consider the process of writing something like *David Copperfield* from a picture like that shown in Figure 3, which is taken from an illustrated version of this classic [74].) Our more immediate goal is therefore to build a machine capable of passing

[6]The actual opening, which we visit in Chapter 2, is as follows:

> As Gregor Samsa awoke one morning from uneasy dreams he found himself transformed in his bed into a gigantic insect. He was lying on his hard, as it were armor-plated, back and when he lifted his head a little he could see a dome-like brown belly divided into stiff arched segments on top of which the bed quilt could hardly keep in position and was about to slide off completely. His numerous legs, which were pitifully thin compared to the rest of his bulk, waved helplessly before his eyes. ([122], p. 67)

a less demanding Torrance-like test; that is, a silicon author able to generate stories that would be regarded creative, even if these stories are well below what a muse-inspired member of *Homo sapiens sapiens* can muster.

"Your Mother and Brother Have Died."

Figure 3: Possible "Dickensian" Input for S^3G. (Reprinted with kind permission from Waldman Publishing Corporation.)

How Do You Build an Artificial Author?

How does one go about building such artificial author? Our answer comes in the following pages. At this point we mention only

one property we believe a good story generator must have: **wide variability**.

There are many dimensions over which a story can vary. Plot is only one of them. Characters, settings, literary themes, writing style, imagery, etc. — these are other dimensions, and there are many more. Generally speaking, belletristic fiction has very wide variability across these dimensions. Mark Helprin's latest novel is likely to have a rather unpredictable plot traversed by rather unpredictable characters in rather unpredictable settings tossed by unpredictable mixtures of love, revenge, jealousy, betrayal, and so on, as reported in prose with a cadence and clarity rarely seen. One of the chief effects of it all is to conjure unforgettable images in the reader's mind. (One of us is haunted weekly by the image of the lost gold in Helprin's *Memoirs From the Antproof Case*.) At the other end of the spectrum fall formulaic fiction and film; here the variability is narrow. Some romance novels, for example, fail to offer wide variability of plot and characterization: It's the same character types time and time again, dancing hot and heavy to the same choreography. (If BRUTUS$_n$, some refined descendant of BRUTUS$_1$, is to soon find employment at the expense of a human writer, in all likelihood it will be as an author of formulaic romance and mystery.)

Whether or not a story generator can be implemented to achieve wide variability hinges on what we call **architectural differentiation**. A story generation system has architectural differentiation if for each substantive aspect of the story that can vary, there is a corresponding distinct component of the technical architecture that can be parameterized to achieve different results. While we owe many debts to the pioneers who have come before us in the field of story generation, it's safe to say that their systems failed to enable wide variability via architectural differentiation.

From the start, our approach has been to bestow the BRUTUS architecture with a counterpart to *every* substantive aspect of human literary genius. While our first implementation of this architecture, BRUTUS$_1$, has quite limited variability, ancestors will implement more and more of those parts of the architecture designed to secure wide variability.

Wide variability is an important property, but there are others that are equally important. One of the ways to encapsulate all of them, and to quickly characterize our approach, is to say that BRU-

TUS is designed to satisfy what we call the seven magic desiderata for a successful story generator, namely:

MD1 *Give proposed rigorous accounts of strong creativity a run for their money.* An impressive storytelling AI is one that satisfies, or at least comes close to satisfying, proposed sophisticated accounts of *strong* creativity. BRUTUS$_1$ does this: As we show later, the system qualifies as capable of raw origination on Margaret Boden's definition of this concept.

MD2 *Generate imagery in the reader's mind.* An artificial agent aspiring to be counted among the literati must be able to spark significant readerly imaging. (Sometimes even literary fiction can earn classification as such despite displaying ordinary prose. Victor Hugo's *Les Miserables* is a case in point: The writing is simple, relative to other immortals, anyway, but what readers can forget the scenes set in the sewers beneath Paris?)

MD3 *Situate the story in "landscape of consciousness."* A good storytelling AI must produce stories having not only a landscape of action, but also a landscape of consciousness, that is, a landscape defined by the mental states of characters.

MD4 *Mathematize concepts at the core of belletristic fiction.* No artificial agent will lay claim to being counted literarily creative unless it processes the immemorial themes (e.g., betrayal) at the heart of literature; and such processing can presumably come only if the themes in question have been formalized.

MD5 *Generate genuinely interesting stories.* A true artificial storyteller must produce genuinely interesting stories. Among the things that readers find interesting are particular topics like sex and money and death (as the well-known cognitive scientist Roger Schank has explained [205]), and also classic themes like betrayal, ruthless ambition, and unrequited love.

MD6 *Tap into the deep, abiding structures of stories.* Any truly impressive artificial author must be in command of story structures that give its output an immediate standing amongst its human audience. For BRUTUS$_1$, these structures take the form of what are called 'story grammars.'

MD7 *Avoid "mechanical" prose.* Last but not least: An artificial author must produce compelling literary prose.

The seven magic desiderata are cashed out in BRUTUS, a rich and highly differentiated system architecture for story generation.

BRUTUS$_1$ is the current implementation of the BRUTUS — notice the absence of the subscript — architecture.

Why Build an Artificial Author?

Finally, a question interviewers and members of the audience and out-of-the-blue e-mailers have asked us time and time again through the years: Why do it? There are at least three general reasons, two theoretical, one practical.

The first theoretical reason for investing time, money, and talent in the quest for a truly creative machine is to work toward an answer to the question of whether we ourselves are machines. If the creative side of human cognition can be captured by computation, then it's surely likely that we are at bottom computers. (The more quotidian side of human mentation can presumably be mechanized, and "lower level" sensing and effecting in interchange with the environment should present no insurmountable obstacles to AI's upward march through the next century.) As you will see in the coming pages, we follow a singular method: As we uncover reasons for believing that human creativity is in fact beyond the reach of computation, we will be inspired to nonetheless engineer systems that dodge these reasons and *appear* to be creative. A side effect of our approach is perhaps to furnish AI with at least an early brick or two in a theoretical foundation for machine creativity. Absent such a foundation (whose mortar, to be effective, would presumably have to be somewhat mathematical in nature), artificial creative agents will never arrive.

The second theoretical reason for our work is stark and simple: to silence those who believe that logic is forever closed off from the emotional world of creativity. BRUTUS is Vulcan through and through, utterly devoid of emotion, but he nonetheless seems to have within his reach things that touch not only our minds, but our hearts.

The practical rationale for our endeavor is that machines able to work alongside humans in arenas calling for creativity would have incalculable worth. A machine able to write a full, formidable novel, or compose a feature-length film, or create and manage the unfolding story in an online game, would be, we suspect, pure gold.

S.B. Troy NY / D.F. Yorktown Heights NY — June 1999

On Silicon Wings

Granite majesty rises our vision to heaven and bound
Crushed and ground, smashed and spread,
Bed our mother's ebb and tide.
The push and pull delivers a grainy pebble ride.
Beneath our feet, cushioning our journey
From slimy cellular slop to pensive petitioners of paradise.

The mat of our birth and the walls of our demise.

Stretching through time, small and broken pieces of dirt
The fallen and the forgotten, the plenty and the bare –
Rise to cup the water to our lips,
Rise to clear our vision of things far and small,
Rise to road our passage from home to home,
Rise to bridge our thoughts from sun to sun.

And the splendid, seemingly solid, visions of heaven,
Humbly laid down to bed our birth and our play,
Rise again to lift us above the somatic images of paradise lost,
Mimicking our minds to clear sight of our souls.

On Silicon wings we will fly.

List of Tables

List of Figures

Bibliography

[1] Andersen, S. & Slator, B. (1990) "Requiem for a Theory: The 'Story Grammar' Story," *Journal of Experimental and Theoretical Artificial Intelligence* **2.3**: 253–275.

[2] Aristotle, "On Memory and Recollection," 45a 5, 450a 7, and 451a; 19, in: McKeon, R., ed., (1941) *The Basic Works of Aristotle* (New York, NY: Random House), pp. 425–519.

[3] Ashcraft, M.H. (1994) *Human Memory and Cognition* (New York, NY: HarperCollins).

[4] Attneave, F. & Curlee T.E. (1983) "Locational Representation in Imagery: A Moving Spot Task," *Journal of Experimental Psychology: Human Perception and Performance* **9**: 20–30.

[5] Averbach, E. & Coriell, A.S. (1961) "Short-term Memory in Vision," *Bell System Technical Journal* **40**: 309–328.

[6] Baron, R.J. (1985) "Visual Memories and Mental Images," *International Journal of Man-Machine Studies* **23**: 275–311.

[7] Barr, A. (1983) "Artificial Intelligence: Cognition as Computation," in Machlup, F., ed., *The Study of Information: Interdisciplinary Messages* (New York, NY: Wiley-Interscience), pp. 237–262.

[8] Barwise, J. & Etchemendy, J. (1993) *Hyperproof* (Stanford, CA: CSLI Publications).

[9] Barwise, J. & Etchemendy, J. (1993) *Turing's World 3.0* (Stanford, CA: CSLI Publications).

[10] Bennett, C.H. (1984) "Thermodynamically Reversible Computation," *Physics Rev. Lett.* **53**: 1202.

[11] Bennett, C.H. (1982) "The Thermodynamics of Computation — A Review," *International Journal of Theoretical Physics* **21**: 905–940.

[12] Bennett, C.H. (1973) "Logical Reversibility of Computation," *IBM Journal of Research Development* November: 525–532.

[13] Berlekamp, E., Conway, J., & Guy, R. (1982) *Winning Ways*, Vol. 2 (New York, NY: Academic Press). See chapter 25 for Conway's description of Life.

[14] Black, J.B. & Wilensky, R. (1979) "An Evaluation of Story Grammars," *Cognitive Science* **3**: 213–230.

[15] Block, N. (1995) "On a Confusion About a Function of Consciousness," *Behavioral and Brain Sciences* **18**: 227–247.

[16] Block, N. (1981) *Imagery* (Cambridge, MA: MIT Press).

[17] Block, N. (1980) "Troubles with Functionalism," in *Readings in Philosophy of Psychology Vol. I* (Cambridge, MA: Harvard University Press).

[18] Boden, M. (1995) "Could a Robot be Creative?—And Would We Know?" in Ford, K.M., Glymour, C. & Hayes, P.J. eds., *Android Epistemology* (Cambridge, MA: MIT Press), pp. 51–72.

[19] Boden, M. (1994) "Creativity and Computers," in T. Dartnall, ed., *Artificial Intelligence and Computers* (Dordrecht, The Netherlands: Kluwer), pp. 3–26.

[20] Boden, M. (1991) *The Creative Mind: Myths and Mechanisms* (New York, NY: Basic Books).

[21] Boolos, G.S. & Jeffrey, R.C. (1989) *Computability and Logic* (Cambridge, UK: Cambridge University Press).

[22] Bringsjord, S. & Noel, R. (in press) "Why Did Evolution Engineer Consciousness?" in Mulhauser, Gregory, ed., *Evolving Consciousness* (Reading, MA: Benjamin Cummings).

[23] Bringsjord, S. & Ferrucci, D. (1998) "Logic and Artificial Intelligence: Divorced, Separated, Still Married . . .?" *Minds and Machines* **8**: 273–308.

[24] Bringsjord, S., Bringsjord, E. and Noel, R. (1998) "In Defense of Logical Minds," in *Proceedings of the 20th Annual Conference of the Cognitive Science Society* (Hillsdale, NJ: Lawrence Erlbaum Associates), pp. 173–178.

[25] Bringsjord, S. (1998) "Chess Is Too Easy," *Technology Review* **101.2**: 23–28 This is an expanded version of one part of the paper Bringsjord, S. & Lally, A. (1997) "Chess Isn't Tough Enough: Better Games for Mind-Machine Competition," in *Collected Papers from the 1997 Workshop Deep Blue versus Kasparov: The Significance for Artificial Intelligence*, Technical Report WS-97-04, (Menlo Park, CA: AAAI Press), pp. 14–19.

[26] Bringsjord, S. (1998) "Philosophy and 'Super'-Computation," in Moor, J. & Bynum, T., eds., *The Digital Phoenix* (Oxford, UK: Basil Blackwell), pp. 231–252.

[27] Bringsjord, S. (1997) "Consciousness by the Lights of Logic and Commonsense," *Behavioral and Brain Sciences* **20.1**: 144–146.

[28] Bringsjord, S. (1997) *Abortion: A Dialogue* (Indianapolis, IN: Hackett).

[29] Bringsjord, S. (1997) "An Argument for the Uncomputability of Infinitary Mathematical Expertise," in Feltovich, P., Ford, K.M., & Hoffman, R.R., eds., *Expertise in Context* (Menlo Park, CA: AAAI Press), pp. 475–497.

[30] Bringsjord, S. & Bringsjord, E. (1996) "The Case Against AI From Imagistic Expertise," *Journal of Experimental and Theoretical Artificial Intelligence* **8**: 383–397.

[31] Bringsjord, S. (1995) "Computation, Among Other Things, is Beneath Us," *Minds and Machines* **4**: 469–488.

[32] Bringsjord, S. (1995) "Could, How Could We Tell If, and Why Should–Androids Have Inner Lives," in *Android Epistemology*, Ford, K., Glymour, C., & Hayes, P., eds., (Cambridge, MA: MIT Press), pp. 93–122.

[33] Bringsjord, S. (1995) "Pourquoi Hendrik Ibsen Est-Is Une Menace pour La Littérature Générée Par Ordinateur?" (traduit par Michel Lenoble) in Vuillemin, A., ed., *Littérature et Informatique la Littérature Générée Par Orinateur*, (Arras, France: Artois Presses Universite), pp. 135–144.

[34] Bringsjord, S. (1995) "In Defense of Impenetrable Zombies," *Journal of Consciousness Studies* **2.4**: 348–351.

[35] Bringsjord, S. (1994) Review of Margaret Boden's *Myths and Mechanisms*, *Behavioral and Brain Sciences* **17.3**: 532–533.

[36] Bringsjord, E. & Bringsjord, S. (1994) "Can AI Accommodate Imagistic Expertise?" in *Proceedings of the Second International Conference on Expert Systems for Development* (Los Alamitos, CA: IEEE Press), pp. 36–41,

[37] Bringsjord, S. (1994) "Searle on the Brink," *Psyche* **1**: 19–31.

[38] Bringsjord, S. & Bringsjord, E. (1994) "Animal Communication of Private States Doesn't Illuminate the Human Case," *Behavioral and Brain Sciences* **16.4**: 645–646.

[39] Bringsjord, S. (1992) "CINEWRITE: An Algorithm-Sketch for Writing Novels Cinematically, and Two Mysteries Therein," *Instructional Science* **21**: 155–168.

[40] Bringsjord, S. (1992) *What Robots Can and Can't Be* (Dordrecht, The Netherlands: Kluwer).

[41] Bringsjord, S. (1991) "Is the Connectionist-Logicist Clash one of AI's Wonderful Red Herrings?" *Journal of Experimental & Theoretical AI* **3.4**: 319–349.

[42] Bringsjord, S. & Zenzen, M. (1991) "In Defense of Hyper-Logicist AI," in *IJCAI '91* (Mountain View, CA: Morgan Kaufmann), pp. 1066–1072.

[43] Britton, B. (1983) "What Makes Stories Interesting," *Behavioral and Brain Sciences* **6**: 596–597.

[44] Bruner, J. (1986) *Actual Minds, Possible Worlds* (Cambridge, MA: Harvard University Press).

[45] Cervantes, M. (1999) *Don Quijote* (New York, NY: Norton).

[46] Castañeda, H. (1979) "Fiction and Reality: Their Fundamental Connections; An Essay on the Ontology of Total Experience," *Poetics* **8**: 31–62.

[47] Chalmers, D. (1995) "Minds, Machines, and Mathematics," *Psyche*.

- http://psyche.cs.monash.edu.au/psyche/volume2-1/psyche-95-2-09-shadows-7-chalmers.html
- ftp:ftp.cs.monash.edu.au.psyche/psyche-95-2-09-shadows-7-chalmers.txt

[48] Chandrasekaran, B., Hari Narayanan, N. & Iwasaki, Y. (1993) "Reasoning With Diagrammatic Representations: A Report on the Spring Symposium," *AI Magazine*, Summer, 23–32.

[49] Charniak, E. & McDermott, D. (1985) *Introduction to Artificial Intelligence* (Reading, MA: Addison-Wesley).

[50] Chase, C.I. (1985) "Review of the Torrance Tests of Creative Thinking," in Mitchell, J.V., ed., *9th Mental Measurements Yearbook, Vol. II* (Lincoln, NB: Buros Institute of Mental Measurement), pp. 1631–1632.

[51] Church, A. (1936) "An Unsolvable Problem of Elementary Number Theory," in Dave, M., ed., *The Undecidable* (New York, NY: Raven Press), pp. 89–100.

[52] Chellas, B.F. (1980) *Modal Logic: An Introduction* (Cambridge, UK: Cambridge University Press).

[53] Cleland, C. (1995) "Effective Procedures and Computable Functions," *Minds and Machines* **5**: 9–23.

[54] Cleland, C. (1993) "Is the Church-Thesis True?" *Minds and Machines* **3**: 283–312.

[55] Cole, D. & Foelber, R. (1984) "Contingent Materialism," *Pacific Philosophical Quarterly* **65.1**: 74–85.

[56] Cornoldi, D., Cortesi, A., & Preti, D. (1991) "Individual Differences in the Capacity Limitations of Visuospatial Short-Term

Memory: Research on Sighted and Totally Congenitally Blind People," *Memory and Cognition* **19**: 459–468.

[57] Clark, A. (1997) "The Dynamical Challenge," *Cognitive Science* **21.4**: 461–481.

[58] Cummins, R. & Schwarz, D. (1991) "Connectionism, Computation and Cognition," in Horgan, T. and Tienson, J., eds., *Connectionism and the Philosophy of Mind* (Dordrecht, The Netherlands: Kluwer Academic Publishers), pp. 60–73.

[59] Davidson, D. (1987) "Knowing One's Own Mind," *Proceedings and Addresses of the American Philosophical Association* **60**: 441–458.

[60] Davis, M.D., Sigal, R. & Weyuker, E.J. (1994) *Computability, Complexity, and Languages* (San Diego, CA: Academic Press).

[61] Davis, M. (1993) "How Subtle is Gödel's Theorem? More on Roger Penrose," *Behavioral and Brain Sciences* **16**: 611–612.

[62] Davis, W. (1988) *Passage of Darkness: The Ethnobiology of the Haitian Zombie* (Chapel Hill, NC: University of North Carolina Press).

[63] Davis, W. (1985) *The Serpent and the Rainbow* (New York, NY: Simon & Schuster).

[64] De Beaugrande, R. & Colby, B.N. (1979) "Narrative Models of Action and Interaction," *Cognitive Science* **3.1**: 43–46.

[65] Dehn, N. (1981) "Story Generation After TALE-SPIN," *IJCAI 81* (San Mateo, CA: Morgan Kaufmann), 16–18.

[66] Dennett, D.C. (1996) "Cow-sharks, Magnets, and Swampman," *Mind & Language* **11.1**: 76–77.

[67] Dennett, D.C. (1995) "The Unimagined Preposterousness of Zombies," *Journal of Consciousness Studies* **2.4**: 322–326.

[68] Dennett, D.C. (1994) "The Practical Requirements for Making a Conscious Robot," *Philosophical Transactions of the Royal Society of London* **349**: 133–146.

[69] Dennett, D.C. (1993) "Review of Searle's *The Rediscovery of the Mind*," *Journal of Philosophy* **90.4**: 193–205.

[70] Dennett, D.C. (1991) *Consciousness Explained* (Boston, MA: Little, Brown).

[71] Dennett, D. (1981) "The Nature of Images and the Introspective Trap," in Block, N., ed., *Imagery* (Cambridge, MA: MIT Press), pp. 51–61.

[72] Dennett, D.C. (1978) *Brainstorms* (Cambridge, MA: MIT Press).

[73] Descartes, R. (1911–first edition) *The Philosophical Works of Descartes Vol. I*, translated by Haldane, E.S. & Ross, G.R.T. (Cambridge, UK: Cambridge University Press).

[74] Dickens, C. (MCMXCII) *David Copperfield* (New York, NY: Baronet Books).

[75] Dickmann, M.A. (1975) *Large Infinitary Languages* (Amsterdam, The Netherlands: North-Holland).

[76] Dietrich, E. (1990) "Computationalism," *Social Epistemology* **4.2**: 135-154.

[77] Dougherty, R.C. (1994) *Natural Language Computing: An English Generative Grammar in Prolog* (Mahwah, NJ: Lawrence Erlbaum Associates).

[78] Doyle, A.C. (1984) "The Adventure of Silver Blaze," in *The Celebrated Cases of Sherlock Holmes* (Minneapolis, MN: Amarenth Press), pp. 172–187.

[79] Dreyfus H. L. & Dreyfus, S.E. (1986) *Mind Over Machine* (New York, NY: Free Press).

[80] Dretske, F. (1996) "Absent Qualia," *Mind & Language* **11.1**: 78–85.

[81] Earman, J. (1986) *A Primer on Determinism* (Dordrecht, The Netherlands: D. Reidel).

[82] Ebbinghaus, H.D., Flum, J. & Thomas, W. (1984) *Mathematical Logic* (New York, NY: Springer-Verlag).

[83] Eco, U. (1979) *The Role of the Reader: Explorations in the Semiotics of Texts* (Bloomington, IN: Indiana University Press).

[84] Esrock, E.J. (1994) *The Reader's Eye: Visual Imaging as Reader Response* (Baltimore, MD: Johns Hopkins).

[85] Euclid. (1956) *The Thirteen Books of Euclid's Elements*, trans. T. Heath (New York, NY: Dover).

[86] Feferman, S. (1994) "Penrose's Gödelian Argument," *Psyche*.

- http://psyche.cs.monash.edu.au/psyche/volume2-1/psyche-95-2-7-shadows-5-feferman.html
- ftp:ftp.cs.monash.edu.au.psyche/psyche-95-2-7-shadows-5-feferman.txt

[87] Feldman, C.F., Bruner, J., Renderer, B., & Spitzer, S. (1990) "Narrative Comprehension," in Britton, B.K. & Pellegrini, A.D., eds., *Narrative Thought and Narrative Language* (Hillsdale, NJ: Lawrence Erlbaum Associates), pp. 1–78.

[88] Fetzer, J. (1996) "Minds Are Not Computers: (Most) Thought Processes Are Not Computational Procedures," paper presented at the annual meeting of the Southern Society for Philosophy and Psychology, Nashville, April 5.

[89] Fetzer, J. (1994) "Mental Algorithms: Are Minds Computational Systems?" *Pragmatics & Cognition* **2.1**: 1–29.

[90] Fjelde, R. (1965) Foreword in Ibsen, H. (1965) *Four Major Plays* (New York, NY: New American Library).

[91] Flanagan, O. & Polger, T. (1995) "Zombies and the Function of Consciousness," *Journal of Consciousness Studies* **2.4**: 313–321.

[92] Folina, J. (1993) "Commentary on Selmer Bringsjord's 'Church's Thesis, Contra Mendelson, Is Unprovable ... And Worse: It May be False'," Annual Eastern Division APA Meeting, Atlanta, GA, December 27, 1993.

[93] Freud, S. (1959) *Creative Writers and Daydreaming* (London, UK: Hogarth Press and the Institute of Psychoanalysis).

[94] Funt, B.V. (1980) "Problem Solving With Diagrammatic Representations," *Artificial Intelligence* **13**: 201–230.

[95] Gardin, F. & Meltzer, B. (1989) "Analogical Representations of Naive Physics," *Artificial Intelligence* **38**: 139–159.

[96] Genesereth, M.R. & Nilsson, N.J. (1987) *Logical Foundations of Artificial Intelligence* (Los Altos, CA: Morgan Kaufmann).

[97] Gershenfield, N. (1999) *When Things Start to Think* (New York, NY: Henry Holt & Company).

[98] Glasgow, J., Narayanan, H. & Chandrasekaran, B., eds. (1995) *Diagrammatic Reasoning* (Cambridge, MA: MIT Press).

[99] Glasgow, J. & Papadias, D. (1995) "Computational Imagery," in Glasgow, J., Narayanan, H., & Chandrasekaran, B., eds., *Diagrammatic Reasoning* (Cambridge, MA: MIT Press), pp. 435–480.

[100] Glasgow, J.I., Fortier, S. & Allen, F.H. (1992) "Molecular Scene Analysis: Crystal Structure Determination Through Imagery," in Hunter, L., ed., *Artificial Intelligence and Molecular Biology* (Cambridge, MA: MIT Press), pp. 433–458.

[101] Glymour, C. (1992) *Thinking Things Through* (Cambridge, MA: MIT Press.

[102] Grzegorczyk, A. (1957) "On the Definitions of Computable Real Continuous Functions," *Fundamentals of Mathematics* **44**: 61–71.

[103] Grzegorczyk, A. (1955) "Computable Functionals," *Fundamentals of Mathematics* **42**: 168–202.

[104] Gummerman, K., Gray, C. & Wilson, J.M. (1992) "An Attempt to Assess Eidetic Imagery Objectively," *Psychonomic Science* **28.2**: 115–118.

[105] Harnad, S. (1995) "Why and How We Are Not Zombies," *Journal of Consciousness Studies* **1**: 164–167.

[106] Harnad, S. (1991) "Other Bodies, Other Minds: A Machine Incarnation of an Old Philosophical Problem," *Minds and Machines* **1.1**: 43–55.

[107] Haugeland, J. (1981) *Artificial Intelligence: The Very Idea* (Cambridge, MA: MIT Press).

[108] Hilbert, D. (1926) "On the Infinite," *Math. Annalen* **95**: 161–190. Translated in Van Heijenoort.

[109] Hobbes, T. (1839) *De Corpore*, chap. 1, in *English Works*, Molesworth, ed., reprinted in (1962) *Body, Man and Citizen* (New York, NY: Collier), pp. 124–225.

[110] Hoffman, R.R. & Klein, G.A. (1993) "Seeing the Invisible: Perceptual-Cognitive Aspects of Expertise," in *Cognitive Science Foundations of Instruction*, Rabonowitz, M., ed., (Hillsdale, NJ: Lawrence Erlbaum), pp. 203–226.

[111] Hofstadter, D.R. (1995) *Fluid Concepts and Creative Analogies* (New York, NY: Basic Books).

[112] Hofstadter, D.R. (1985) "Waking Up from the Boolean Dream," chap. 26 in his *Metamagical Themas: Questing for the Essence of Mind and Pattern* (New York, NY: Bantam), pp. 631–665.

[113] Hofstadter, D. (1982) "Metafont, Metamathematics, and Metaphysics," *Visible Language* **14.4**: 309–338.

[114] Hopcroft, J.E. & Ullman, J.D. (1979) *Introduction to Automata Theory, Languages and Computation* (Reading, MA: Addison-Wesley).

[115] Jackson, F. (1982) "Epiphenomenal Qualia," *Philosophical Quarterly* **32**: 127–136.

[116] Jacquette, D. (1994) *Philosophy of Mind* (Englewood Cliffs, NJ: Prentice-Hall).

[117] Johnson, G. (1997) "Undiscovered Bach? No, a Computer Wrote It," *The New York Times*, November 11, pp. F1-2.

[118] Johnson-Laird, P. (1988) *The Computer and the Mind* (Cambridge, MA: Harvard University Press).

[119] Johnson-Laird, P.N. (1983) *Mental Models: Toward a Cognitive Science of Language, Inference, and Consciousness* (Cambridge, MA: Harvard University Press).

[120] Joyce, J. (1986) "Eveline," in Joyce, J., *The Dubliners* (New York, NY: Penguin).

[121] Julstrom, B.A. & Baron, R.J. (1985) "A Model of Mental Imagery," *International Journal of Man-Machine Studies* **23**: 313–334.

[122] Kafka, F. (1948) "The Metamorphosis," in *The Penal Colony*, trans. W. Muir and E. Muir (New York, NY: Schocken Books).

[123] Kafka, F. (1948) "The Penal Colony," in *The Penal Colony*, trans. W. Muir and E. Muir (New York, NY: Schocken Books).

[124] Kalmár, L. (1959) "An Argument Against the Plausibility of Church's Thesis," in Heyting, A., ed., *Constructivity in Mathematics* (Amsterdam, Holland: North-Holland), pp. 72–80.

[125] Keisler, H.J. (1971) *Model Theory for Infinitary Logic* (Amsterdam, The Netherlands: North-Holland).

[126] Kerr, N.H. (1987) "Locational Representation in Imagery: The Third Dimension," *Memory and Cognition* **15**: 521–530.

[127] Kintsch, W. (1980) "Learning From Text, Levels of Comprehension, or: Why Anyone Would Read a Story Anyway," *Poetics* **9**: 87–98.

[128] Kleene, S.C. (1936) "General Recursive Functions of Natural Numbers," *Math. Annalen* **112**: 727–742.

[129] Klein, S. (1975) "Meta-Compiling Text Grammars as a Model for Human Behaviour," *TINLAP*, 94–98.

[130] Kolata, G. (1996) "Computer Math Proof Shows Reasoning Power," *The New York Times' Cybertimes*. (No page numbers: an online article.)

[131] Kosslyn, S. (1994) *Image and Brain* (Cambridge, MA: MIT Press).

[132] Kosslyn, S., Alpert, N.M., Thompson, W.L., V. Maljkovic, V., Weise, S.B., Chabris, C.F., Hamilton, S.E., Rauch, S.L., & Buonanno, F.S. (1993) "Visual Mental Imagery Activates Topographically Organized Visual Cortex: PET Investigations," *Journal of Cognitive Neuroscience* **5.3**: 263–287.

[133] Kosslyn, S. (1983) *Ghosts in the Mind's Machine* (New York, NY: Norton).

[134] Kosslyn, S. (1980) *Image and Mind* (Cambridge, MA: Harvard).

[135] Kosslyn. S. & Schwartz, S.P. (1978) "Visual Images as Spatial Representations in Active Memory," in Riseman, E.M., & Hanson, A.R. *Computer Vision Systems* (New York, NY: Academic Press), pp. 12-38.

[136] Kosslyn. S. & Shwartz, S.P. (1977) "A Simulation of Visual Imagery," *Cognitive Science* **1**: 265–295.

[137] Kreisel, G. (1968) "Church's Thesis: A Kind of Reducibility Thesis for Constructive Mathematics," in Kino, A., Myhill, J., & Vesley, R.E., eds., *Intuitionism and Proof Theory: Proceedings of a Summer Conference at Buffalo, N.Y.* (Amsterdam, Holland: North-Holland), pp. 219–230.

[138] Kreisel, G. (1965) "Mathematical Logic," in Saaty, T.L., ed.,*Lectures in Modern Mathematics* (New York, NY: John Wiley), pp. 111–122.

[139] Kripke, S. (1971) "Naming and Necessity," in Davidson, D.& Harman, G., eds., *Semantics of Natural Language* (Dordrecht, The Netherlands: Reidel), pp. 253–355, 763–769.

[140] Kugel, P. (1990) "Is it Time to Replace Turing's Test?" Paper presented at *Artificial Intelligence: Emerging Science or Dying Art Form*, sponsored by the American Association of Artificial Intelligence and the State University of New York at Binghamton's program in Philosophy and Computer and Systems Sciences, the University at Binghamton, Binghamton, NY, June 27.

[141] Kugel, P. (1986) "Thinking May Be More Than Computing," *Cognition* **22**: 137–198.

[142] Kurzweil, R. (1999) *The Age of Spiritual Machines: When Computers Exceed Human Intelligence* (New York, NY: Viking).

[143] Lambek, J. (1961) "How to Program an Infinite Abacus," *Canadian Mathematical Bulletin* **4**: 295–302. (See the correction noted in (1962) **5**: 297.)

[144] Larkin, J. & Simon, H.A. (1987) "Why a Diagram Is (Sometimes) Worth Ten Thousand Words," *Cognitive Science* **10**: 65–100.

[145] (1981) *Graphic Art Materials Reference Manual* (New York, NY: Letraset).

[146] Lebowitz, M. (1984) "Creatiing Characters in a Story-Telling Universe," *Poetics* **13**: 171–194.

[147] Lewis, D. (1978) "Truth in Fiction," *American Philosophical Quarterly* **15**: 37–46.

[148] Lewis, H. & Papadimitriou, C. (1981) *Elements of the Theory of Computation* (Englewood Cliffs, NJ: Prentice-Hall).

[149] Lindsay, R.K. (1988) "Images and Inference," *Cognition* **23**: 229–249.

[150] Lucas, J.R. (1964) "Minds, Machines and Gödel," in Anderson, A.R., ed., *Minds and Machines* (Englewood Cliffs, NJ: Prentice-Hall), pp. 43–59.

[151] Luria, A.R. (1968) *The Mind of a Mnemonist*, trans. L. Solotaroff (New York, NY: Basic Books).

[152] Martin, R. (1984) *Recent Essays on Truth and the Liar Paradox* (Oxford, UK: Oxford University Press).

[153] Marxen, H. & Buntrock, J. (1990) "Attacking the Busy Beaver 5," *Bulletin of the European Association for Theoretical Computer Science* **40**: 247–251.

[154] Maudlin, T. (1989) "Computation and Consciousness," *Journal of Philosophy* **84**: 407–432.

[155] McCulloch, W.S. & Pitts, W. (1943) "A Logical Calculus of the Ideas Immanent in Nervous Activity," *Bulletin of Mathematical Biophysics* **5**: 115–137.

[156] McGinn, C. (1999) "Hello, Hal" *New York Times Book Review* January 3: 11–12.

[157] McMenamin, M. (1992) "Deciding Uncountable Sets and Church's Thesis," unpublished manuscript.

[158] Mele, A. (forthcoming) "Real Self-Deception," *Behavioral and Brain Sciences*.

[159] Melzak, Z. (1961) "An Informal Arithmetical Approach to Computability and Computation," *Canadian Mathematical Bulletin* **4**: 279–293.

[160] Meehan, J. (1981) "TALE-SPIN," in Schank, R. & Reisbeck, C., eds., *Inside Computer Understanding: Five Programs Plus Miniatures* (Hillsdale, NJ: Lawrence Erlbaum Associates), pp. 197–226.

[161] Mendelson, E. (1990) "Second Thoughts About Church's Thesis and Mathematical Proofs," *Journal of Philosophy* **87.5**: 225–233.

[162] Mendelson, E. (1963) "On Some Recent Criticism of Church's Thesis," *Notre Dame Journal of Formal Logic* **4.3**: 201-205.

[163] Meteer, M. (1992) *Expressibility and the Problem of Efficient Text Planning* (London, UK: Pinter).

[164] Millikan, R.G. (1996) "On Swampkinds," *Mind & Language* **11.1**: 103–117.

[165] Moravec, H. (1999) *Robot: Mere Machine To Transcendant Mind* (Oxford, UK: Oxford University Press).

[166] Moschovakis, Y. (1968) "Review of Four Recent Papers in Church's Thesis," *Journal of Symbolic Logic* **33**: 471–472. One

of the four papers is: Kalmár, L. (1959) "An Argument Against the Plausibility of Church's Thesis," in Heyting, A., ed., *Constructivity in Mathematics* (Amsterdam, Holland: North-Holland), pp. 72–80.

[167] Nagel, T. (1974) "What Is it Like to Be a Bat?" *Philosophical Review* **83**: 435–450.

[168] Nelson, R.J. (1987) "Church's Thesis and Cognitive Science," *Notre Dame Journal of Formal Logic* **28.4**: 581–614.

[169] Newell, A. (1980) "Physical Symbol Systems," *Cognitive Science* **4**: 135–183.

[170] Oakhill, J.V., Johnson-Laird, P.N. & Garnham, A. (1989) "Believability and Syllogistic Reasoning," *Cognition* **31**: 117–140.

[171] O'Keefe, R. (1990) *The Craft of Prolog* (Cambridge, MA: MIT Press).

[172] Omori, T. (1992) "Dual Representation of Image Recognition Process: Interaction of Neural Network and Symbolic Processing," *Proceedings of the International Symposium on Neural Information Processing*, pp. 50–53.

[173] Parsons, T. (1975) "A Meinongian Analysis of Fictional Objects," *Grazer Philosphische Studien* **1**: 73–86.

[174] Partee, B., Meulen, A. & Wall, R. (1990) *Mathematical Methods in Linguistics* (Dordrecht, The Netherlands: Kluwer Academic Publishers).

[175] Peck, M.S. (1983) *People of the Lie* (New York, NY: Simon and Schuster).

[176] Penrose, R. (1994) *Shadows of the Mind* (Oxford, UK: Oxford University Press).

[177] Penrose, R. (1989) *The Emperor's New Mind* (Oxford, UK: Oxford University Press).

[178] Piaget, J. & Inhelder, B. (1969) *The Psychology of the Child* (New York, NY: Basic Books).

[179] Piaget, J. & Inhelder, B. (1966) *L'Image Mentale Chez L'enfant* (Paris, France: Presses Universitaires de France).

[180] Pinker, S. (1997) *How the Mind Works* (New York, NY: Norton).

[181] Plum, F. & Posner, J.B. (1972) *The Diagnosis of Stupor and Coma* (Philadelphia, PA: F.A. Davis).

[182] Pollock, J. (1995) *Cognitive Carpentry: A Blueprint for How to Build a Person* (Cambridge, MA: MIT Press).

[183] Pollock, J. (1989) *How to Build a Person* (Cambridge, MA: MIT Press).

[184] Post, E.L. (1936) "Finite Combinatory Processes – Formulation 1," *Journal of Symbolic Logic* **1.3**: 103–105.

[185] Poundstone, W. (1985) *The Recursive Universe* (New York, NY: William Morrow).

[186] Propp, V. (1986) *The Morphology of the Folktale* (Austin, TX: University of Texas Press).

[187] Putnam, H. (1960) "Minds and Machines," in his *Mind, Language, and Reality: Philosophical Papers, Vol. 2* (Cambridge, UK: Cambridge University Press), pp. 45–61.

[188] Pylyshyn, Z. (1981) "Imagery and Artificial Intelligence," in Block, N., ed., *Readings in Philosophy of Psychology Vol. 2* (Cambridge, MA: Harvard).

[189] Quaife, A. (1992) *Automated Development of Fundamental Mathematical Theories* (Dordrecht, The Netherlands: Kluwer).

[190] Racter (1984) *The Policeman's Beard is Half Constructed* (New York, NY: Warner).

[191] Rado, T. (1963) "On Non-Computable Functions," *Bell System Technical Journal* **41**: 877–884.

[192] Rapaport, W.J. (1991) "Predication, Fiction, and Artificial Intelligence," *Topoi* **10**: 79–91.

[193] Robinson, J.A. (1992) "Logic and Logic Programming," *Communications of the ACM* **35.3**: 40–65.

[194] Rogers, H. (1967) *Theory of Recursive Functions and Effective Computability* (New York, NY: McGraw-Hill).

[195] Rosenthal, D.M. (forthcoming) "State Consciousness and What It's Like," in hi Title TBA (Oxford, UK: Clarendon Press.

[196] Rosenthal, D.M. (1990) "Why Are Verbally Expressed Thoughts Conscious?" ZIF Report No. 32, Zentrum für Interdisziplinäre Forschung, Bielefeld, Germany.

[197] Rosenthal, D.M. (1990) "A Theory of Consciousness," ZIF Report No. 40, Zentrum für Interdisziplinäre Forschung, Bielefeld, Germany.

[198] Rosenthal, D.M. (1989) "Thinking That One Thinks," ZIF Report No. 11, Research Group on Mind and Brain, Perspective in Theoretical Psychology and the Philosophy of Mind, Zentrum für Interdisziplinäre Forschung, Bielefeld, Germany.

[199] Rosenthal, D.M. (1986) "Two Concepts of Consciousness," *Philosophical Studies* **49**: 329–359.

[200] Russell, B. (1936) "The Limits of Empiricism," *Proceedings of the Aristotelian Society* **XXXVI**: 131–150.

[201] Russell, S. & Norvig, P. (1995) *Artificial Intelligence: A Modern Approach* (Englewood Cliffs, NJ: Prentice-Hall).

[202] Sackheim, H. & Gur, R. (1978) "Self-Deception, Self-Confrontation, and Consciousness," in *Consciousness and Self-regulation, Vol. 2* (New York, NY: Plenum Press), pp. 117-129.

[203] Schacter, D.L. (1989) "On the Relation Between Memory and Consciousness: Dissociable Interactions and Conscious Experience," in Roediger, H. & Craik, F., eds., *Varieties of Memory and Consciousness: Essays in Honour of Endel Tulving* (Hillsdale, NJ: Erlbaum), pp. 22-35.

[204] Schank, R. (1995) *Tell Me a Story* (Evanston, IL: Northwestern University Press).

[205] Schank, R. (1979) "Interestingness: Controlling Inferences," *Artificial Intelligence* **12**: 273–297.

[206] Searle, J. (1992) *The Rediscovery of the Mind* (Cambridge, MA: MIT Press).

[207] Searle, J. (1983) *Intentionality* (Cambridge, UK: Cambridge University Press).

[208] Searle, J. (1980) "Minds, Brains and Programs," *Behavioral and Brain Sciences* **3**: 417–424.

[209] Sharples, M. (1997) "Story Telling by Computer," *Digital Creativity* **8.1**: 20–29.

[210] Shepard, F.N. & Metzler, J. (1971) "Mental Rotation of Three-dimensional Objects," *Science* **171**: 701–703

[211] Sieg, W. & Byrnes, J. (1996) "K-graph Machines: Generalizing Turing's Machines and Arguments," in Hájek, P., ed., *Gödel 96*, *Lecture Notes in Logic 6* (New York, NY: Springer-Verlag), pp. 98–119.

[212] Siegelmann, H. (1995) "Computation Beyond the Turing Limit," *Science* **268**: 545-548.

[213] Siegelmann, H. & Sontag, E.D. (1994) "Analog Computation Via Neural Nets," *Theoretical Computer Science* **131**: 331–360.

[214] Simon, H. (1980) "Cognitive Science: The Newest Science of the Artificial," *Cognitive Science* **4**: 33–56.

[215] Simon, H. (1981) "Study of Human Intelligence by Creating Artificial Intelligence," *American Scientist* **69.3**: 300–309.

[216] Slezak, P. (1982) "Gödel's Theorem and the Mind," *British Journal for the Philosophy of Science* **33**: 41–52.

[217] Smith, M.C. (1981) *Gorky Park* (New York, NY: Ballantine Books).

[218] Smolensky, P. (1988) "On the Proper Treatment of Connectionism," *Behavioral & Brain Sciences* **11**: 1–22.

[219] Smolensky, P. (1988) "Putting Together Connectionism — Again," *Behavioral & Brain Sciences* **11**: 59–70.

[220] Smullyan, R.M. (1992) *Gödel's Incompleteness Theorems* (Oxford, UK: Oxford University Press).

[221] Soare, R. (1980) *Recursively Enumerable Sets and Degrees* (New York, NY: Springer-Verlag).

[222] Sperling, G. (1960) "The Information Available in Brief Visual Presentations," *Psychological Monographs* **74**: 48.

[223] Sterling, L. & Shapiro, E. (1986) *The Art of Prolog* (Cambridge, MA: MIT Press).

[224] Stillings, N.A., Weisler, S.E., Chase, C.H., Feinstein, M.H., Garfield, J.L., & Rissland, E.L. (1995) *Cognitive Science: An Introduction* (Cambridge, MA: MIT Press).

[225] Stromeyer, C.F., & Psotka, J. (1970) "The Detailed Texture of Eidetic Images," *Nature* **225**: 346–349.

[226] Suppes, P. (1972) *Axiomatic Set Theory* (New York, NY: Dover).

[227] Swartz, J.D. (1988) "Torrance Tests of Creative Thinking," in Keyser, D.J. & Sweetland, R.C., eds., *Test Critiques, Vol. VII* (Kansas City, MO: Test Corporation of America), pp. 619–662.

[228] Tarjan, R.E. (1971) "An Efficient Planarity Algorithm" (Report STAN-CS-244-71), Stanford, CA: Stanford University.

[229] Tchaikovsky. (1893) "Letter to Vladimir Davidov (Tchaikovsky's nephew)," reproduced in *Notes on Tchaikovsky's Symphony No. 6 – Pathétique*, text included with CD produced by RCA Records, New York, NY, 1985.

[230] Thomas, W. (1973) "Doubts About Some Standard Arguments for Church's Thesis," in *Papers of the Fourth International Congress for Logic, Methodology, and Philosophy of Science, Bucharest* (Amsterdam, Holland: D. Reidel), pp. 13–22.

[231] Thorndyke, P.W. (1977) "Cognitive Structures in Comprehension and Memory of Narrative Discourse," in his *Cognitive Psychology* (New York, NY: Academic Press), pp. 224–239.

[232] Torrance, E.P. (1988) "The Nature of Creativity as Manifest in its Testing," in Sternberg, R.J., ed., *The Nature of Creativity: Contemporary Psychological Perspectives* (Cambridge, UK: Cambridge University Press), pp. 72–89.

[233] Torrance, E.P. (1966) *The Torrance Tests of Creative Thinking: Technical-Norms Manual* (Princeton, NJ: Personnel Press).

[234] Trabasso, T. (1996) "Review of *Knowledge and Memory: The Real Story*," Robert S. Wyer, ed., Lawrence Erlbaum, 1995, *Minds & Machines* **6**: 399–403.

[235] Treffinger, D.J. (1985) "Review of the Torrance Tests of Creative Thinking," in Mitchell, J.V., ed., *9th Mental Measurements Yearbook, Vol. II* (Lincoln, NB: Buros Institute of Mental Measurement), pp. 1632–1634.

[236] Turing, A.M. (1964) "Computing Machinery and Intelligence," in Andersen, A.R., ed., *Minds and Machines*, Contemporary Perspectives in Philosophy Series (Englewood Cliffs, NJ: Prentice-Hall), pp. 4–30.

[237] Turner, S. (1994) *The Creative Process: A Computer Model of Storytelling* (Hillsdale, NJ: Lawrence Erlbaum Associates).

[238] Tye, M. (1991) *The Imagery Debate* (Cambridge, MA: MIT Press).

[239] Tye, M. (1988) "The Picture Theory of Mental Images," *The Philosophical Review* XCVII.4: 497–520.

[240] Van Heijenoort, J., ed. (1967) *From Frege to Gödel* Amsterdam, The Netherlands: North-Holland).

[241] Van Inwagen, P. (1977) "Creatures of Fiction," *American Philosophical Quarterly* **14**: 299–308.

[242] Vasey, P. (1989) *LPA-flex Technical Reference* (London, England: Logic Programming Associates Ltd.).

[243] Warwick, K. (1997) *March of the Machines: Why the New Race of Robots Will Rule the World* (London, UK: Century).

[244] Webb, J. (1980) *Mechanism, Mentalism and Metamathematics* (Dordrecht, The Netherlands: D. Reidel).

[245] Weyl, H. (1949) *Philosophy of Mathematics and Natural Science* (Princeton, NJ: Princeton University Press).

[246] Wilensky, R. (1983) "Story Grammars Versus Story Points," *Behavioral and Brain Sciences* **6**: 529–591.

[247] Winograd, T. & Flores, F. (1986) *Understanding Computers and Cognition: A New Foundation for Design* (Norwood, NJ: Ablex).

[248] Wittgenstein, L. (1983) *Remarks on the Foundations of Mathematics* (Cambridge, MA: MIT Press).

[249] Wos, L. (1996) *The Automation of Reasoning: An Experimenter's Notebook With OTTER Tutorial* (San Diego, CA: Academic Press).

[250] Wos, L. (1992) *Automated Reasoning: Introduction and Applications* (New York, NY: McGraw Hill).

[251] Wyer, R.S. (1995) *Knowledge and Memory: The Real Story* (Hillsdale, NJ: Lawrence Erlbaum Associates).

[252] Yazdani, M. (1989) "Computational Story Writing," in Williams, N., and Holt, P., eds., *Computers and Writing* (Norwood, NJ: Ablex), pp. 125–147.

[253] Ybarra, M.J. (1996) "Discovering an Answer in the Flames," *New York Times*, Sunday, February 4, Section A, p. 13.

[254] Zenzen, M. & Hollinger, H. (1985) *The Nature of Irreversibility* (Dordrecht, The Netherlands: D. Reidel).

Index