

**Between Chomsky and BRUTUS.
Can Machines be Creative?**
by
Dragoș Avădanei

Chomsky believed that lying beneath the astonishing linguistic abilities of humans is a universal grammar, represented by deep generative structures that nobody really knows how they got to be there, i.e. in their own modules within the brain and developing, largely autonomously, from human cognition. Then came Gerald Edelman, a neuroscientist, who believed that meaning does not reside in one site of the human brain, “but is typically a dynamic and variable pattern of connections over many elements”(Turner); our subjective experience of thought and sensation arise from the simultaneous activation of many different overlapping systems of neurons, called maps, which influence and reinforce one another. And then came Mark Turner, who uses the second author, Edelman, to tell Chomsky that he was simply wrong, and that it is not grammar which inhabits the deepest region of the mind’s linguistic capacities, but parable and the ability to tell stories.

Keywords: *cognitive science, Herbert Simon, Chomsky, creativity, computers.*

Nobel Prize Winner (in economics) Herbert Simon (1916-2001) was one of the most influential American scientists of the 20th century, whose interest and research efforts ranged from cognitive psychology to economics, to public administration and the philosophy of science; he is counted among the fathers of such diverse domains as artificial intelligence, information processing, problem-solving and decision-making, organization theory, complex systems. He is remembered for such quotes as: “There are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until—in a visible future—the range of problems they can handle will be coextensive with the range to which the human mind has been applied...” (encouraging or scary?); or—“Information consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention, and a need to allocate that attention efficiently among the overabundance of information sources that might consume it.”

But, of course, he is remembered for his many books and contributions, published—the former—beginning with 1947 and including *Administrative Behavior, Models of Man, The Sciences of the Artificial, Models of Discovery, Models of Thought, Models of Bounded Rationality, Models of My Life*.

In so far as we are concerned here, he is remembered for a special issue (vol. 4, 1995) of the *Stanford Humanities Review* dedicated to the topic “Where Cognitive Science Meets Literary Criticism” and including a position paper by

himself, titled “Literary Criticism: A Cognitive Approach,” thirty-three peer commentaries coming from well-known specialist in English, Foreign Languages, Philosophy, Computer Science, Computer Engineering, Cultural Studies, Humanities in general, Technology, Mathematics, Semiotics, and Neuropsychology (about 60,000 words in all), and a final reply signed again by Simon. A presentation of the problems in this issue looks like a good introduction to our tentative survey of cognitive literary criticism (whose very existence—we shall see—is questioned as yet). Our source has been the Internet, and hence the absence of page numbers for the quotations.

The main question is that of the relationship between literary criticism and cognitive science and how/if they can be useful to each other. The larger question obviously is the one raised half a century ago by C. P. Snow, i.e. that of the two cultures, of the humanities and the sciences. Developments in AI that came after Snow changed significantly the direction of the question in that many theories developed to describe a certain cognitive ability in cognitive science (solving certain types of mathematical problems or the competence to play sophisticated games like chess, for instance) have been transformed into computational models whose practical results, the machines, can reproduce that specific skill; so that the program’s results are the best assessment of the theory and its explanatory power; the performance of a chess-playing program becomes thus the best measure of the theory’s power to explain the phenomenon of chess-playing. Along these lines, our own question would be whether the production of literature can be transformed into a computational model so that the explanatory power of literary theory and criticism could be assessed by it. In other words, if the literary critic can show how the mind of the writer works—and he can certainly know it if the writer is a machine—then his cognitive job is finished before its very beginning. We shall return to this possibility.

Towards the end of his paper, Simon becomes so confident as to say that “criticism can be viewed (imperialistically) simply as a branch of cognitive science”; only his confidence is not shared by a number of respondents and the question is far from being *simple*. Generally speaking, Simon’s aim is, first, to provide a precise, science-based definition of meaning understood in operational terms, and, second, to show how his theoretical account can be applied to the explanation of literary texts; thirdly, an implicit aim is to define meaning in such a way as to advance his program of simulating human intelligence with computers. His basic message is strongly optimistic as regards the potential of a cognitive approach to literary theory and criticism so that his proposal rests on the assumption that there is a congruence between the structure of texts and the structure of minds: since “literary criticism concerns... the meanings of, in, and evoked by literary texts” and “cognitive science concerns thinking,” meaning and thinking are obviously the concern of both.

In a text the meaning may have three sources: the author's meaning, the meaning of the text, and the meaning that derives from a reading of the text. And he goes on to explore the intended meaning (oscillating between intension and intention) without any hint at all that he might be aware of Wimsatt and Beardsley's intentional fallacy. Next he appropriately approaches the problem of context, which includes "the memory of surrounding elements of the text." And thus, "the meaning of the text... will be a function of the memory contents that are accessed by recognition of words," and this recognition is given by the power of association. As we move into a text, like Stendhal's *Chartreuse de Parme*, the meaning of each sentence or unit is "expanded by knowledge of the meaning of the other/s." Thus: recognition, memory, association, and context.

He next turns to what he knows best, i.e. the symbolic processes a computer can execute by using his own "physical symbol system hypothesis" according to which symbols can be represented by patterns of electromagnetism in computers:

"The basic processes that a computer can perform with symbols are to *input* them into memory, *combine* and *reorganize* them into symbol *structures*, *store* such structures over time, *erase* them, *output* them through motor processes, *compare* pairs of symbols for equality and inequality, and '*branch*' (behave conditionally on the outcome of such texts). The physical symbol system hypothesis asserts that possessing these processes is the necessary and sufficient condition for a system to be capable of thinking". (our italics)

This is part of the larger discipline of artificial intelligence, and the implication is that if a computer can do all these things it can also write, and those who developed "the writer" are sure to know how it works. But this will be the subject of a separate section.

Without any allusion to I. A. Richards and C. K. Ogden's *The Meaning of Meaning: A Study of the Influence of Language upon Thought and the Science of Symbolism* (1923), Simon returns to the problem of meaning, distinguishes between potential and actual meaning and gives us further comments on contexts (in memory of in the universe, in larger texts and in culture). He then also distinguishes between contexts (depending on writer's and reader's prior knowledge) and schemas, or local contexts that grow out of the information found in the text itself.

The context or contexts of writing (historico-biographical, cultural, social, etc.) are paralleled by the reader's contexts, the former obviously determining the author's meaning (Simon's author is never dead), the latter—the reader's.

Once again, without showing any sign that he knows about William Empson's *Seven Types of Ambiguity* (1930), Simon focuses his attention on ambiguity (multiple meanings, enigmas, options...), simply concluding that it is

inexhaustible, “a permanent lode of treasure for scholars.” And so criticism becomes part of the work as a whole, in fact, part of the authorship: “Shakespeare must now share... authorship with all those who have commented on him borrowed from him plagiarized him, been compared with him, distanced themselves from him.” Shakespeare’s meaning is the sum total of these meanings, coming from as many critical contexts. If we, by any chance, accept the theory that critical thinking has always preceded creative thinking, then we will also know that *post-facto* scholarship and criticism only come to complete the cycle and take us back to the beginning: criticism in search of its own roots, before creativity developed in-between.

But Simon does not go this far and prefers to return to the story grammars of machines—“accounts of the structures of tales and the processes that understand the tales by discovering these structures...” only to hope that, *in time*, a bridge will be created between the two cultures:

“Professional competence is a domain of the humanities, like competence in a domain of science, requires the accumulation of a great deal of specialized knowledge. We cannot expect to master the content of more than a very few domains in any great depth. What we can hope to do is to work toward a common understanding of the mental processes that all of us use to extract meanings. /the meanings are there in advance—to be extracted!/ However distinct and dissimilar the domains, our minds, fashioned from the same raw stuff and employing the same basic symbolic processes, must have a great deal in common that we can share.”

Which says more about possible bridges than about cognitive literary criticism. When he comes back, at the end, to reply to commentaries, Herbert Simon expresses his conviction that “experiments and computer simulations are a principal contribution of cognitive science to literary criticism.” So the question remains as to what computer simulations can teach us about the nature of human thought in general, and literary thought in particular, and here the next question is whether computer simulations can go beyond cognition into the realms of emotion, motivation, and aesthetic judgment. One may want to look into it, before we return to Simon.

Between Chomsky and BRUTUS

Chomsky believed that lying beneath the astonishing linguistic abilities of humans is a universal grammar, represented by deep generative structures that nobody really knows how they got to be there, i.e. in their own modules within the brain and developing, largely autonomously, from human cognition.

Then came Gerald Edelman, a neuroscientist, who believed that meaning does not reside in one site of the human brain, “but is typically a dynamic and variable pattern of connections over many elements”(Turner); our subjective

experience of thought and sensation arise from the simultaneous activation of many different overlapping systems of neurons, called maps, which influence and reinforce one another (see also Antonio Damaso's model of "convergence": the brain integrates information across various sensory modalities).

And then came Mark Turner, who uses the second author, Edelman, to tell Chomsky that he was simply wrong, and that it is not grammar which inhabits the deepest region of the mind's linguistic capacities, but parable and the ability to tell stories, which means that mind is literary *before* it is linguistic; Edelman's overlapping systems of neurons are called blended spaces, and it is in such as these that disparate elements of parables come together to form meaning. The conclusion here, which then becomes an assumption for his seminal book *The Literary Mind: the Origins of Thought and Language* (OUP, 1996; see also Jerry Hobbs' *Literature and Cognition*, Stanford, 1990), is both challenging and convincing: story projection and parable precede grammar; language follows from these mental capacities as a consequence, and thus language is the child of the literary mind. Thus it is worth quoting in full from p.5 of his "Preface":

"The literary mind is the fundamental mind... Story is the basic principle of mind. Most of our experiences, our knowledge, and our thinking is organized as stories. The mental scope of story is magnified by *projection*—one story helps us make sense of another. The projection of one story into another is parable... We interpret /think, invent, plan, decide, reason, imagine, persuade/ every level of our experience by means of parable. Language is not the source of parables, but instead its complex product... Parable is the root of the human mind—of thinking, knowing, acting, creating and plausibly even speaking... Narrative imagining—story—is the fundamental instrument of thought. Rational capacities depend upon it. It is our chief means of looking into the future, of predicting, of planning and of explaining. It is a literary capacity indispensable to human cognition generally. This is the first way in which the mind is essentially literary."

And:

"The literary mind—the mind of stories and parables—is not peripheral but basic to thought. *Story* is the central principle of our experience and knowledge. *Parable*—the projection of story to give meaning to new encounters—is the indispensable tool of everyday reason. Literary thought makes everyday thought possible... ; the basic issue for cognitive science is the nature of literary thinking."

For cognitive science in general, not just for cognitive literary criticism. Therefore, the two ways in which the human mind is essentially literary consist in

the story being a “basic principle of mind,” a fundamental cognitive capacity since our thinking is organized as stories and, second, in the projection of one story onto another, i.e. parable, which

“is also, like the story, a fundamental instrument of mind... The essence of parable is its intricate combining of two of our most basic forms of knowledge: story and projection. This classic combination produces one of our keenest mental processes of construing meaning. The evolution of the genre of parable... follows inevitably from the nature of our conceptual systems.” (Ch.I)

Thus cognitive science in general depends upon it since “if we want to study the everyday mind, we can begin by turning to the literary mind exactly because the everyday mind is essentially literary.” (*ibid.*)

And Turner is not alone in his belief; thus, in R. S. Wyer’s *1995 Knowledge and Memory: The Real Story* (Hillsdale, NJ), Roger C. Schank and Robert Ableson share the view that narrative is central and ubiquitous in human cognition and that human knowledge, all of it, is based on stories; the same year Schank devoted a whole book to the topic (*Tell Me a Story: Narrative and Intelligence*, Evanston, IL) while Daniel C. Dennett had previously explained (in *Consciousness Explained*, Boston, 1991) that thinking consists basically in the telling of parallel stories.

Here we can introduce the following argument: if the human mind is the literary mind, and this mind functions in terms of story and parable, then any mind able to tell a story has to be a literary mind: in other words, if such a literary mind were to invent a machine that is able to tell stories, then that human mind must first have a story about this invention and, second, the invention must have a literary mind.

In one particular case, if two computer scientists decide to spend many years of their lives to devise a computer program that tells stories, they must have a story about how such a story-minded program can come into being. They (Selmer Bringsjord and David A. Ferrucci, *Artificial Intelligence and Literary Creativity: Inside the Mind of Brutus, as Storytelling Machine*, Mahwah, NJ: Lawrence Erlbaum Associates, 2000) know “that it seems plausible that narrative does stand at the heart of cognition in *any* domain, whether it’s air traffic control, medical diagnosis, pedagogy, or corporate decision making.” (Ch. 1.2) What they do not seem to know is the path to follow in building a storytelling program that must be conceived and implemented in terms of a narratively organized mind.

Stimulated by Mary Boden’s work on computers and creativity (*The Creative Mind: Myths and Mechanisms*, New York: Basic Books, 1991), Bringsjord and Ferrucci propose to answer her four questions about computers and what they could or might do, now or in the future: Can computational ideas

help us understand how human creativity is possible? (Boden answers “Yes,” B/F answer “No”). Could computers do things which at least *appear* to be creative? (Yes, Yes) Could a computer ever appear to *recognize* creativity? (Yes, Yes) Could computers themselves ever *really* be creative? (No, No)

For a definition of creativity they go back to E. P. Torrance’s test of 1966 (*The Torrance Test of Creative Thinking: Technical-Norms Manual*, Princeton, NJ: Personnel Press), p. 47:

“/Creative thinking/… is the process of sensing difficulties, problems, gaps in information, missing elements, something askew; making guesses and formulating hypotheses about these deficiencies; evaluating and testing these guesses and hypotheses; possibly revising and retesting them; and finally communicating the results.”

Combined with Turner and others, creative thinking is the process of being able to imagine a story, and this involves subjectivity—impossible to plant into a computer (John Searle, *The Rediscovery of Mind*, Cambridge: MIT Press, 1992)—and creating imagery that readers would respond to—and that again is impossible (Ellen J. Esrock, *The Reader’s Eye: Visual Imaging as Reader Response*, Baltimore: Johns Hopkins, 1994). The third impossibility is that a point of view cannot be formalized in computational terms: “It’s hard to see how one can engineer a machine with the capacity to occupy the point of view of a creature of fiction if one doesn’t know what a creature of fiction is...” (p.75)

However, this is not enough to discourage Bringsjord and Ferrucci from “realizing a (seemingly) literally creative machine” (p.81), and they decide about their first step, which is that of selecting one of the several immemorial themes: unrequited love, fanaticism, revenge, jealousy, self-deception, infatuation, hatred, alienation, despair, triumph... and betrayal. A second step would consist in mathematizing one or more of these themes, and since they chose “betrayal” they decided to achieve a BRUTUS (name of program) architecture that could contain a “thematic concept instantiation”; but the caveat is there all of the time, namely that “there is something in the human sphere that exceeds computation” (p.91), and that something is not a process in physics, a mathematical theorem, a chemical formula, a medical diagnosis, or an astronomical theorem, but the telling of a story, and a literary story at that, not any of the previously mentioned tasks, which are also stories—we have seen. And maybe that something that cannot be formalized is interestingness, for which they have to return to A. Church’s 1936 thesis, according to which what can be effectively computed is co-extensive with what can be algorithmically computed, and the question becomes that of finding an algorithm for interestingness; not really necessary, since there is an easier way out, that of finding another authority (E. Mendelson) who demonstrated, in 1990, that Church’s thesis is unprovable, so no search for an algorithm is at stake.

Well, then, how about beauty (see Simon, here and next); since this also seems difficult, it immediately becomes irrelevant, “because one can exhaustively analyze cognition (and replicate it on a machine) without bothering to grapple in earnest with this concept.” (p.120)

And how about emotions: Yes, they can also be ignored, since one (John Pollock, *Cognitive Carpentry: A Blueprint for How to Build a Person*, Cambridge: MIT Press, 1995) can build an artificial person that could, without feeling any fear, compute the need to quickly flee a lion. Which means that such a “person” could do without love or hatred, by having in it computations that simulate actions and attitudes humans perform when they love or hate. Interestingness, beauty, and emotion (plus Simon’s motivation) can be only experienced in the first person, while knowledge about all of them can be had in the third person; and thus we return to the question of point of view with knowledge coming from an objectified point of view. And we can take a great leap here and say that you cannot replace a story with a criticism of that story, and thus cognitive criticism is at bay.

Knowing all these, but reluctantly accepting them, Bringsjord and Ferrucci decide to follow a story about story-making, i.e. not building a program that can tell stories, but one that creates the illusion of doing so: “we carefully operate under the belief that human (literary) creativity is beyond computation—and yet strive to craft the *appearance* of creativity from suitably configured computation”. (p.149). They know that we all—writers and readers alike—are in a tight spot here, so the decision is to appeal to history: J. Meeham’s first story generator, TALE-SPIN (1981) which produces, among its best stories, one like “Hunger”:

“Once upon a time John Bear lived in a cave, John knew that John was in his cave. There was a beehive in a maple tree. Tom Bee knew that the beehive was in the maple tree. Tom was in his beehive. Tom knew that Tom was in his beehive. There was some honey in Tom’s beehive. Tom knew that the honey was in Tom’s beehive. Tom had the honey. Tom knew that Tom had the honey. There was a nest in a cherry tree. Arthur Bird knew that the nest was in the cherry tree. Arthur was in his nest. Arthur knew that John was in his cave...”

Then came Scott Turner’s MINSTREL (*The Creative Process: A Computer Model of Storytelling*, Hillsdale, NJ: Lawrence Erlbaum Associates, 1994) based upon the idea that creativity is a matter of solving problems—which, we have seen is mistaken from the start.

And thus history is of no real help, so our authors have to rely on their own research into computer programs and stories, attempting to meet several desiderata already encountered: to spark the readerly imaging; to process one of

the immortal themes; to do something that is uncomputable, i.e. interestingness; develop story grammars (Simon); avoid mechanical prose (like “Hunger”).

The most complex of all—or, probably, the handiest of all—is the story grammar, which they borrow from P. W. Thorndyke (*Cognitive Psychology*, New York: Academic Press, 1977): story (probably narrative and discourse), setting, theme, plot, characters, time, event, state, episode, actions/attempts, outcome/resolution + language generation. These and others are known to us all from the many theories and poetics of fiction (see next) in which stories were deconstructed and reconstructed (see, for instance, Seymour Chatman’s complex diagram). But there is another factor, namely that there are many dimensions over which a story can vary, so architectural differentiation for a story generation system has to be devised, i.e. “for each aspect of the story that can vary, there.../has to be...a corresponding distinct component of the technical architecture that can be parameterized to achieve different results.” (p.161)

Consequently, BRUTUS’s anatomy, its technical architecture is decomposed into two distinct levels: *the knowledge level* and *the process level* (we can now remember that a couple of dozen pages back, the statement was that you can know *about* something being interesting without being able to create something interesting); the knowledge is domain knowledge (people, places, events...), linguistic knowledge (sentences, words, etc.), and literary knowledge. With some domain knowledge and a dash of linguistic knowledge, “a story generation system can cough up a story,” but this would be a weak story, looking like the TALE-SPIN “Hunger” story above, or “more like a laundry list.” (p.168) But the darkest part of what they are groping for is literary knowledge, i.e. “the high art of storytelling.” Domain knowledge can offer a pool of story elements, and so can linguistic knowledge, while the second level, the process level can provide four lines of development: thematic concept instantiation, plot generation, story structure expansion, and language generation. These all come into data structures called frames, which are organized hierarchically: relations are established, production rules are processed by a reasoning engine, agents, events, beliefs, goals, plans, and actions are introduced; characters are given proactive and reactive behavior, words are selected from the dictionary pool, word formation, derivational and inflectional morphology come to help, and generative grammars ensure agreement, punctuation and the like.

And now the literary knowledge, i.e. ways of using words and phrases to achieve literary objectives that are, once again, generating imagery in the reader’s mind, suggesting a character’s landscape of consciousness, a producing a certain mood, positive or negative, secure or anxious, for the reader. These are achieved by the three types of literary association: iconic features, literary modifiers, and literary analogs (or metaphors).

In order that literary and linguistic knowledge could be linked, literary augmented grammars are used in BRUTUS, grammars that are based upon literary constraints, i.e. various parts of speech are categorized and associated with one another by a variety of classification and association rules. Then plot is developed through simulation, and, finally, story structure expansion is programmed, and since Bringsjord and Ferrucci's story becomes too complicated, we prefer to leave it at that (grammar hierarchies, choices, levels, taxonomies, terminals, paragraph grammars, scenarios, variability and variables, etc.)

The brief sample story generated is titled "Betrayal in Self-Deception," "Self-Betrayal," or simple "Betrayal":

"Dave loves the university of Rome. He loves its studious youth, ivy-covered clock towers and its sturdy brick. David wanted to graduate. Prof. Hart told Dave, 'I will sign your thesis at your defense.' Prof. Hart actually intends to thwart Dave's plans to graduate. After Dave completed his defense, and the chairman of Dave's committee asked Prof. Hart to sign Dave's thesis, Prof. Hart refused to sign. Dave was crushed."

Generated by whom? A confusing answer is given two pages later, where sample stories and variations of the previous one are also given: "As you read them now, try to call upon what you have read in this book so that you can 'demystify' the fact that they can be generated by a 'mere' machine."

The generative process thus remains ambiguous: knowledge and the formalizations of knowledge do not seem to be enough on one level, while the process level lacks the essential components: interestingness, subjectivity, point of view, esthetic judgment, and emotion; cognition does not seem to be enough, for the time being at least, in understanding the literary mind. We may now want to see what Simon's peer commentators have to say about it.

Comments on Simon's Position

Among the thirty-three peer commentaries to Simon's proposition, one finds reactions of all sorts, from negative attitudes and complete refusal to acceptance and suggestions to meet on middle ground; therefore, most of the responses are either negative or ambiguous, and we shall look at them in this order.

Stefano Velotti, for instance, condemns Simon's imperialistic attitude of cognitive science ("criticism can be viewed... simply as a branch of cognitive science...") and compares cognitive science with a self-deceiving emperor who thinks that it is sufficient to reduce the world to a map in order to conquer it. What Simon proposes is a view of literature between a Rorschach test (inkblots that would prompt the reader to project every kind of personal associations into it) and as a treasure hunt (for meanings, obviously). His definition of meaning is circular because cognitive science is, after all, only a set of theories:

“What makes literature literature is the fact that it exists or lets emerge—through determinate meanings—the human experience of general meaningfulness (sense, perception, awareness, feeling) that makes theories of meaning possible. All the particular meanings of a text, every image-meaning or emotion-meaning (to repeat Simon’s terminology) are at the same time vehicles or, better, exempla of that very condition that cannot be said *per se* in a particular meaning, but only felt, perceived, questioned. This way of looking at literature is not to be found in the ‘hundreds of flowers’ Simon would like to let bloom.”

Most of the other negative responses focus on this question of meaning, but also on the other elements in Simon’s menu, i.e. intentionality, context, ambiguity and evocation. As a matter of fact, one is imperiously tempted—as Simon himself seems to be in the end—to read all of these responses as commentaries on the meaning of meaning, and of the value of various contexts, and on the kinds of readings that can be applied to a text (including Simon’s and the others’). Thus, referring to the “current orthodoxy now known as cognitivism,” Brian Rotman thinks that Simon never heard of the “intentional fallacy,” did not understand that Chomsky’s generative grammar (closely associated with the cognitivism approach) contributed nothing of value to the reading of texts, and finally that he, Simon, “is either gesturing to an enterprise more complex than he conveys here or he seriously underestimates his audience.”

Mukesh J. Patel’s assumption is that meaning *evocation* is not particularly well understood in cognitive science, and thus the whole approach “seems to have omitted from consideration the notion that a large part of the debate and difference among literary critics has to do with the social, cultural, ethical, and political implications of the interpretation of the text; the debate is not merely confined to differences of opinion on the correct or acceptable reading of a text. The wider implication of evoked meanings matters, and on that cognitive science can only remain mute.”

In his turn, Paul Miers thinks that Simon contradicts himself and what he presents is a kind of “disembodied dogma cloaked in the voice of passive agency so characteristic of official science.” A voice coming now from an advocate of the dominant role for symbol processing within artificial intelligence, who, together with Alan Newell posited, in 1976, that being a physical symbol system is both a necessary and sufficient condition for being intelligent. And this, shows Kevin B. Korg is Simon’s major weakness, i.e. assuming the physical symbol system hypothesis and thus implying that symbolic representations suffice to capture all of the semantic content that is accessible to us. And hence, one like Don Byrd can confidently state that the scientist and the poet are, in fact, figures of an unresolvable dualism: “for one meaning has to do with symbolic exchange, the return of a symbol for a symbol, and for the other, meaning has to do with the

destruction of the symbol system altogether and its replacement with the experience of value. Literary art is only incidentally representational; its processes are only incidentally involved with information processing.” And further on: “a poem communicates no information; it does not reduce uncertainty /Blaga’s corolla of wonders/… Imagination substitutes a world where things are important or unimportant for a world where things are true or false.” This may come from the fact that the poem is not an object but an event—“it does not mean,” but happens, and thus the critic is not simply an interpreter, but a performer.

Still Don Byrd also proposes a transition to the middle-ground position: “If we are to develop a useful, interdisciplinary relationship between those working with literary forms and those working with computer simulations, it will be necessary to begin with the recognition that language does not broadly translate from one discipline to the other, It will be necessary to find a common ground outside of both disciplines.” And this is an anticipation of Robert Pogue Harrison’s complex metaphysical question:

“Is there some way in which that which literature says without saying so preserves in its text the impenetrability of the phenomenal world as well as the inscrutability of our presence in it—an inscrutability that cognitive science can neither account for nor acknowledge, given that our access to the world takes place ultimately beyond the bounds of conceptualization or at best takes place at the edges of intelligibility where conceptualization struggles, but fails, to maintain its grasp of the world?”

Quoting his own *Reading Minds: The Study of English in the Age of Cognitive Science* and assuming, once again, that the everyday mind may be essentially literary (so, cognitivism and the study of mind being one, cognitivism and study of literary mind are one), Mark Turner is confident that there is *surely* a wide expanse of ground common to literary criticism and cognitive science. And a caveat: “That Simon seems bold to us in imagining a connection between cognitive science and literary criticism is a reminder of how dismembered the humanities have become.”

Long the same lines, Helga Wild is ready to propose another kind of (imperialistic) relationship: “Is it not that literary theory would just as well underlie cognitive science and provide the principles of its functioning? After all, the knowledge and achievements of science come to us as descriptions, case studies, and histories, in article and book form, in short, as texts. And is not the function of the literary critic to make sure that this act does not disappear and be forgotten in the fictions that are thereby produced?” Precisely: if the literary mind is the everyday mind (including that of scientists), literary criticism and literary theory underlie the functioning of mind in general, i.e. of cognitive science.

The idea of process (see Don Byrd above) is stressed upon by Richard Vinograd, who proposes that we think of meaning as something dynamically produced: “meaning doesn’t reside in the text, or in the author’s mind, or in the reader’s mind, but is continuously produced in the process of interaction between reader and text.” And: “Reading and meaning are not exactly located: they occur in the text as much as in the mind. We might say that in reading, the mind is engaged in the process of the text. Or even: as much as the text is in the mind, the mind is in the text”. Even though she thinks that Herbert Simon’s mechanistic model, “blurring the differences between symbol systems in silicon chips and in flesh and blood is inherently repellent,” Janet H. Murray is also skeptical of cognitive science’s imperialism, does not expect “that our complex and richly textured emotional life will be captured by quantitative or mechanistic models,” but would still welcome the collaboration proposed by Simon between literary critics who could “learn the extent to which their concepts can be made ‘precise’ without reductiveness” and cognitive scientists who could “test the limits of their very powerful forms of representation.”

Most of the things are reserved for the future: “The development of a cognitive approach to literary criticism—the project of Aristotle and of I. A. Richards and of Herbert Simon—has much yet to accomplish. We needn’t wait for artificial minds to come into being... for the work to proceed.” (Paul Johnston) So, in spite of the fact that Bringsjord and Ferrucci let us with a sense of uncertainty as to the potential of computerized programs to replicate all the important components of a creative literary mind, cognitive meaning can still be taken as representing a large part of the study of literary criticism; in other words, we need not wait for the development of a good or great story generating program in order to see that there is a lot of communication going on between the critical mind, the literary mind, and the... cognitive mind.

Overview and Conclusions

Herbert Simon and his peer commentators provide a rather ambiguous road into the possibilities of a joint exploration of literary criticism and cognitive science. In their turn Bringsjord and Ferrucci do not succeed entirely in persuading about BRUTUS and its narrative talents. However, cognitive science and literary criticism do have a number of things in common, though, as we shall presently see, there still are a number of unknowns that require further attention.

Instead of Simon’s explanations, rather verbose in one view, one could prefer the formulation of Jon K. Barwise (mathematician and philosopher, 1942-2000) and John Perry’s equation in their 1983 *Situation and Attitudes* (Cambridge, MIT Press) where the notion of situation semantics (rules that determine the information context in a language) is introduced. The knowns in the equation are S (the symbol, or the expression, or the *text* as such), and the conventional restrictions or constraints R; the three unknowns for the reader of S

are R (conventional constraints), c (circumstances shared by author and reader), and P (propositional content the author wants to convey). The task of literary interpretation is to use the available information about the unknowns (biographical and historical data, culture, etc.) to circumscribe the range of their possible values: “One is no longer tempted to think that all the possible information one can extract from a statement is somehow part of its content. Information about each of the parameters in the equation gives information about all the others--...
CR(S, C) = P.”

On condition we remember that literary art’s processes are only incidentally involved with information processing. (see Don Byrd above) Anyway, we are here with an equation that patches nicely—if not really convincing—the paths between Snow’s and Simon’s two cultures.

Bibliography

Allman, J. M., (1999) *Evolving Brains* New York: Scientific American Library.

Barkow, J. H. et all., (eds.), (1992) *The Adapted Mind*, New York: OUP.

Baron-Bohen, S., (1995) *Mindblindness*, Cambridge: MIT Press.

Beardsley, Monroe C., (1970) *The Possibility of Criticism*, Detroit: Wayne State Univ. Press.

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

Booth, Wayne C., (1988) *The Company We Keep: An Ethics of Fiction*, Berkeley: California UP.

Bortulussi, Marisa and Peter Dixon, (2003) *Psychonarratology: Foundations for the Empirical Study of Literary Response*, New York: Cambridge UP.

Bringsjord, Selmer and David A. Perrucci, (2000) *Artificial Intelligence and Literary Creativity*, Mahwah, NJ: Lawrence Erlbaum Associates.

Chatman, Seymour, (1978) *Story and Discourse*, Ithaca: Cornell UP.

Donald, M., (1991) *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition*, Cambridge, MA: Harvard Univ. Press.

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

Eco, U., (1979) *The Role of the Reader: Explorations in the Semiotics of Texts*, Bloomington: Indiana UP.

Esrock, E. J., (1994) *The Reader’s Eye: Visual Imaging as Reader Response*, Baltimore: Johns Hopkins.

Genette, G., (1980) *Narrative Discourse*, Ithaca: Cornell UP.

Gerrig, Richard J., (1993) *Experiencing Narrative Worlds: On the Psychological Activities of Reading*, New Haven: Yale UP.

Hartman, G. H., (1975) *The Fate of Reading and Other Essays*, Chicago: Univ. of Chicago Press.

Hays, D. G., (1981) *Cognitive Structures*, New Haven: HRAF Press.

Hirsch, E. D., Jr., (1967) *Validity in Interpretation*, New Haven: Yale UP.

Hobbs, J. R., (1990) *Literature and Cognition*, Stanford: CSLI.

Hofstadter, D. R., (1995) *Fluid Concepts and Creative Analogies*, New York: OUP.

Holland, N. N. , (1968) *The dynamics of Literary Response*, New York: OUP.

Iser, Wolfgang, (1978) *The Act of Reading*, London: Routledge and Kegan Paul.

Juhl, P. D., (1980) *Interpretation: An Essay in the Philosophy of Literary Criticism*, Princeton: Princeton UP.

Moravec, H., (1999) *Robot: Mere Machine to Transcendental Mind*, Oxford: OUP.

Neisser, U., (1976) *Cognition and Reality*, San Francisco: W. H. Freeman.

Nelson, Cary, (ed.), (1986) *Theory in the Classroom*, Champaign: Univ. of Illinois Press.

Papert, S., (1980) *Mindstorms: Children, Computers, and Powerful Ideas*, New York: Basic Books.

Parker, David, (1994) *Ethics, Theory, and the Novel*, Cambridge: Cambridge UP.

Pinker, S., (1997) *How the Mind Works*, New York: W. W. Norton.

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

Pribram, K., (1971) *Languages of the Brain*, Englewood Cliffs, NJ: Prentice Hall.

Schank, R., (1995) *Tell Me a Story*, Evanston: Northwestern Univ. Press.

Sperber, Dan et. all., (eds), (1995) *Causal Cognition: A Multidisciplinary Debate*, New York: OUP.

Spivak, G. Ch., (1993) *Outside in the Teaching Machine*, New York: Routledge.

Torrance, E. P., (1966) *The Torrance Tests of Creative Thinking*, Princeton: Personnel Press.

Turkle, S., (1984) *The Second Self: Computers and the Human Spirit*, New York: Basic Books.

Turner, M., (1996) *The Literary Mind*, New York: OUP.

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

Whiten, Andrew, (ed.), (1991) *Natural Theories of Mind: Evolution, Development, and Simulation of Everyday Mindreading*, Cambridge: Basil Blackwell.

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

Zwann, Rolf A., (1993) *Aspects of Literary Comprehension: A Cognitive Approach*, Amsterdam: Benjamins.

(“Alexandru Ioan Cuza” University, Iași)