Chapter 13

Culture, mind and education

Jerome Bruner

Jerome Bruner, born 1915, was for many decades before his death in 2016 'the grand old man' of American learning and cognitive research and theory, and for a long time he could claim to be the only researcher in the field who had personally known both Vygotsky and Piaget. In the 1940s and 1950s, he made detailed studies on perception, thinking and cognition, and after the so-called 'Sputnik-shock' in 1957, Bruner was appointed chairman of the commission that was set up to fundamentally reconstruct the American school system. Later he laid the groundwork for the concept of the science-centred curriculum. His last important book in the educational area, The Culture of Education, *from 1996, can certainly still be viewed as a relevant and contemporary contribution, and the following chapter is made up of the two first programmatic sections of that book, which probably will stand as the most durable work of his vast production.*

Computationalism and culturalism

The essays in [The *Culture of Education*] are all products of the 1990s, expressions of the fundamental changes that have been altering conceptions about the nature of the human mind in the decades since the cognitive revolution. These changes, it now seems clear in retrospect, grew out of two strikingly divergent conceptions about how mind works. The first of these was the hypothesis that mind could be conceived as a computational device. This was not a new idea, but it had been powerfully reconceived in the newly advanced computational sciences. The other was the proposal that mind is both constituted by and realized in the use of human culture. The two views led to very different conceptions of the nature of mind itself and of how mind should be cultivated. Each led its adherents to follow distinctively different strategies of inquiry about how mind functions and about how it might be improved through "education." The first or *computational* view is concerned with *information processing:* how finite, coded, unambiguous information about the world is inscribed, sorted, stored, collated, retrieved, and generally managed by a computational device. It takes information as its given, as something already settled in relation to some preexisting, rule-bound code that maps onto states of the world. This so-called "well-formedness" is both its strength and its shortcoming, as we shall see. For the process of knowing is often messier and more fraught with ambiguity than such a view allows.

Computational science makes interesting general claims about the conduct of education (Segal et al. 1985, Bruer 1993, Chi et al. 1988), though it is still unclear what specific lessons it has to teach the educator. There is a widespread and not unreasonable belief that we *should* be able to discover something about how to teach human beings more effectively from knowing how to program computers effectively. One can scarcely doubt, for example, that computers provide a learner with powerful aids in mastering bodies of knowledge, particularly if the knowledge in question is well defined. A well-programed computer is especially useful for taking over tasks that, at last, can be declared "unfit for human production." For computers are faster, more orderly, less fitful in remembering, and do not get bored. And of course, it is revealing of our own minds and our human situation to ask what things we do better or worse than our servant computer.

It is considerably more uncertain whether, in any deep sense, the tasks of a teacher can be "handed over" to a computer, even the most "responsive" one that can be theoretically envisioned. Which is not to say that a suitably programmed computer cannot lighten a teacher's load by taking over some of the routines that clutter the process of instruction. But that is not the issue. After all, books came to serve such a function after Gutenberg's discovery made them widely available (Ong 1991, Olson 1994).

The issue, rather, is whether the computational view of mind itself offers an adequate enough view about how mind works to guide our efforts in trying to "educate" it. It is a subtle question. For in certain respects, "how the mind works" is itself dependent on the tools at its disposal. "How the *hand* works," for example, cannot be fully appreciated unless one also takes into account whether it is equipped with a screwdriver, a pair of scissors, or a laser-beam gun. And by the same token, the systematic historian's "mind" works differently from the mind of the classic "teller of tales" with his stock of combinable myth-like modules. So, in a sense, the mere existence of computational devices (and a theory of computation about their mode of operating) can (and doubtless will) change our minds about how "mind" works, just as the book did (Olson 1994).

This brings us directly to the second approach to the nature of mind – call it *culturalism*. It takes its inspiration from the evolutionary fact that mind could not exist save for culture. For the evolution of the hominid mind is linked to the development of a way of life where "reality" is represented by a symbolism shared by members of a cultural community in which a technical-social way of life is both organized and construed in terms of that symbolism. This symbolic mode is not only shared by a community, but conserved, elaborated, and passed on to succeeding generations who, by virtue of this transmission, continue to maintain the culture's identity and way of life.

Culture in this sense is *superorganic* (Kroeber 1917). But it shapes the minds of individuals as well. Its individual expression inheres in *meaning making*, assigning meanings to things in different settings on particular occasions. Meaning making involves situating encounters with the world in their appropriate cultural contexts in order to know "what they are about." Although meanings are "in the mind," they have their origins and their significance in the culture in which they are created. It is this cultural situatedness of meanings that assures their negotiability and, ultimately, their communicability. Whether "private meanings" exist is not the point; what is important is that meanings provide a basis for cultural exchange. On this view, knowing and communicating are in their nature highly interdependent, indeed virtually inseparable: however much the individual may seem to operate on his or her own in carrying out the quest for meanings, nobody can do it unaided by the culture's symbolic systems. It is culture that provides the tools for organizing and understanding our worlds in communicable ways. The distinctive feature of human evolution is that mind evolved in a fashion that enables human beings to utilize the tools of culture. Without those tools, whether symbolic or material, man is not a "naked ape" but an empty abstraction.

Culture, then, though itself man-made, both forms and makes possible the workings of a distinctively human mind. On this view, learning and thinking are always *situated* in a cultural setting and always dependent upon the utilization of cultural resources (see e.g. Bruner 1990). Even individual variation in the nature and use of mind can be attributed to the varied opportunities that different cultural settings provide, though these are not the only source of variation in mental functioning.

Like its computational cousin, culturalism seeks to bring together insights from psychology, anthropology, linguistics, and the human sciences generally, in order to reformulate a model of mind. But the two do so for radically different purposes. Computationalism, to its great credit, is interested in any and all ways in which information is organized and used – information in the well-formed and finite sense mentioned earlier, regardless of the guise in which information processing is realized. In this broad sense, it recognizes no disciplinary boundaries, not even the boundary between human and non-human functioning. Culturalism, on the other hand, concentrates exclusively on how human beings in cultural communities create and transform meanings.

I want to set forth in this chapter some principal motifs of the cultural approach and explore how these relate to education. But before turning to that formidable task, I need first to dispel the shibboleth of a necessary contradiction between culturalism and computationalism. For I think the apparent contradiction is based on a misunderstanding, one that leads to gross and needless overdramatization. Obviously the approaches are very different, and their ideological overspill may indeed overwhelm us if we do not take care to distinguish them clearly. For it surely matters ideologically what kind of "model" of the human mind one embraces (Brinton 1965). Indeed, the model of mind to which one adheres even shapes the "folk pedagogy" of schoolroom practice. Mind as equated to the power of association and habit formation privileges "drill" as the true pedagogy, while mind taken as the capacity for reflection and discourse on the nature of necessary truths favors the Socratic dialogue. And each of these is linked to our conception of the ideal society and the ideal citizen.

Yet in fact, neither computationalism nor culturalism is so linked to particular models of mind as to be shackled in particular pedagogies. Their difference is of quite a different kind. Let me try to sketch it.

The objective of computationalism is to devise a formal redescription of *any* and *all* functioning systems that manage the flow of well-formed information. It seeks to do so in a way that produces foreseeable, systematic outcomes. One such system is the human mind. But thoughtful computationalism does *not* propose that mind is like some particular "computer" that needs to be "programmed" in a particular way in order to operate systematically or "efficiently." What it argues, rather, is that any and all systems that process information must be governed by specifiable "rules" or procedures that govern what to do with inputs. It matters not whether it is a nervous system, or the genetic apparatus that takes instruction from DNA and then reproduces later generations, or whatever. This is the ideal of artificial intelligence (AI), so-called. "Real minds" are describable in terms of the same AI generalization – systems governed by specifiable rules for managing the flow of coded information.

But, as already noted, the rules common to all information systems do not cover the messy, ambiguous, and context-sensitive processes of meaning making, a form of activity in which the construction of highly "fuzzy" and metaphoric category systems is just as notable as the use of specifiable categories for sorting inputs in a way to yield comprehensible outputs. Some computationalists, convinced a priori that even meaning making can be reduced to AI specifications, are perpetually at work trying to prove that the messiness of meaning making is not beyond their reach (McClelland 1990, Schank 1990). The complex "universal models" they propose are sometimes half-jokingly referred to by them as "TOEs," an acronym for "theories of everything" (Mitchell 1995). But though they have not even come near to succeeding and, as many believe, will probably never in principle succeed, their efforts nonetheless are interesting for the light they shed on the divide between meaning making and information processing.

The difficulty these computationalists encounter inheres in the kinds of "rules" or operations that are possible in computation. All of them, as we know, must be specifiable in advance, must be free of ambiguity, and so on. They must, in their ensemble, also be computationally consistent, which means that while operations may alter with feedback from prior results, the alterations must also adhere to a consistent, prearranged systematicity. Computational rules may be contingent, but they cannot encompass unforeseeable contingencies. Thus Hamlet cannot (in AI) tease Polonius with ambiguous banter about "yonder cloud shaped like a camel, nay 'tis backed like a weasel," in the hope that his banter might evoke guilt and some telltale knowledge about the death of Hamlet's father.

It is precisely this clarity, this prefixedness of categories, that imposes the most severe limit on computationalism as a medium in which to frame a model of mind. But once this limitation is recognized, the alleged death struggle between culturalism and computationalism evaporates. For the meaning making of the culturalist, unlike the information processing of the computationalist, is in principle interpretive, fraught with ambiguity, sensitive to the occasion, and often after the fact. Its "ill-formed procedures" are like "maxims" rather than like fully specifiable rules (Sperber and Wilson 1986, Grice 1989). But they are hardly unprincipled. Rather, they are the stuff of *hermeneutics*, an intellectual pursuit no less disciplined for its failure to produce the click-clear outputs of a computational exercise. Its model case is text interpretation. In interpreting a text, the meaning of a part depends upon a hypothesis about the meanings of the whole, whose meaning in turn is based upon one's judgment of meanings of the parts that compose it. But a wide swath of the human cultural enterprise depends upon it. Nor is it clear that the infamous "hermeneutic circle" deserves the knocks it gets from those in search of clarity and certainty. After all, it lies at the heart of meaning making.

Hermeneutic meaning making and well-formed information processing are incommensurate. Their incommensurability can be made evident even in a simple example. Any input to a computational system must, of course, be encoded in a specifiable way that leaves no room for ambiguity. What happens, then, if (as in human meaning making) an input needs to be encoded according to the context in which it is encountered? Let me give a homely example involving language, since so much of meaning making involves language. Say the input into the system is the word *cloud*. Shall it be taken in its "meteorological" sense, its "mental condition" sense, or in some other way? Now, it is easy (indeed necessary) to provide a computational device with a "look-up" lexicon that provides alternative senses of *cloud*. Any dictionary can do it. But to determine *which* sense is appropriate for a particular context, the computational device would also need a way of encoding and interpreting all contexts in which the word *cloud* might appear. That would then require the computer to have a look-up list for all possible contexts, a "contexticon." But while there are a finite number of words, there are an infinite number of contexts in which particular words might appear. Encoding the context of Hamlet's little riddle about "yonder cloud" would almost certainly escape the powers of the best "contexticon" one could imagine!

There is no decision procedure known that could resolve the question whether the incommensurability between culturalism's meaning making and computationalism's information processing could ever be overcome. Yet, for all that, the two have a kinship that is difficult to ignore. For once meanings are established, it is their formalization into a well-formed category system that *can* be managed by computational rules. Obviously one loses the subtlety of context dependency and metaphor in doing so: *clouds* would have to pass tests of truth functionality to get into the play. But then again, "formalization" in science consists of just such maneuvers: treating an array of formalized and operationalized meanings as if they were fit for computation. Eventually we come to believe that scientific terms actually were born and grew that way: decontextualized, disambiguated, totally "look-uppable."

There is equally puzzling commerce in the other direction. For we are often forced to interpret the output of a computation in order to "make some sense" of it – that is, to figure out what it "means." This "search for the meaning" of final outputs has always been customary in statistical procedures such as factor analysis where the association between different "variables," discovered by statistical manipulation, needed to be interpreted hermeneutically in order to "make sense." The same problem is encountered when investigators use the computational option of parallel processing to discover the association between a set of coded inputs. The final output of such parallel processing similarly needs interpretation to be rendered meaningful. So there is plainly some complementary relationship between what the computationalist is trying to explain and what the culturalist is trying to interpret, a relationship that has long puzzled students of epistemology (von Wright 1971, Bruner 1985).

In an undertaking as inherently reflexive and complicated as characterizing "how our minds work" or how they might be made to work better, there is surely room for two perspectives on the nature of knowing (von Wright 1971). Nor is there any demonstrable reason to suppose that without a single and legitimately "true" way of knowing the world, we could only slide helplessly down the slippery slope that leads to relativism. It is surely as "true" to say that Euclid's theorems are computable as to say, with the poet, that "Euclid alone has looked on beauty bare."

A theory of mind

To begin with, if a theory of mind is to be interesting educationally, it should contain some specifications for (or at least implications bearing on) how its functioning can be improved or altered in some significant way. All-or-none and once-for-all theories of mind are not educationally interesting. More specifically, educationally interesting theories of mind contain specifications of some kind about the "resources" required for a mind to operate effectively. These include not only instrumental resources (like mental "tools"), but also settings or conditions required for effective operations – anything from feedback within certain time limits to, say, freedom from stress or from excessive uniformity. Without specification of resources and settings required, a theory of mind is all "inside-out" and of limited applicability to education. It becomes interesting only when it becomes more "outside-in," indicating the kind of world needed to make it possible to use mind (or heart!) effectively – what kinds of symbol systems, what kinds of accounts of the past, what arts and sciences, and so on. The approach of computationalism to education tends to be inside-out – though it smuggles the world into the mind by inscribing bits of it in memory, as with our earlier dictionary example, and then relies on "lookup" routines. Culturalism is much more outside-in, and although it may contain specifications about mental operations *eo ipso*, as it were, they are not as binding as, say, the formal requirement of computability. For the approach of the computationalist to education is indeed bound by the constraint of computability – that is, whatever aids are offered to mind must be operable by a computational device.

When one actually examines how computationalism has approached educational issues, there seem to be three different styles. The first of these consists in "restating" classical theories of teaching or learning in a computable form. But while some clarity is gained in so doing (for example, in locating ambiguities), not much is gained by way of power. Old wine does not improve much for being poured into differently shaped bottles, even if the glass is clearer. The classic reply, of course, is that a computable reformulation yields "surplus insight." Yet "association theory," for example, has gone through successive translations from Aristotle to Locke to Pavlov to Clark Hull without much surplus yield. So one is justifiably impatient with new claims for veiled versions of the same as with many so-called parallel distributed processing (PDP) "learning models" (Rumelhart and McClelland 1986).

But in fact, computationalism can and does do better than that. Its second approach begins with a rich description or protocol of what actually transpires when somebody sets out to solve a particular problem or master a particular body of knowledge. It then seeks to redescribe what has been observed in strict computational terms. In what order, for example, does a subject ask for information, what confuses him, what kinds of hypotheses does he entertain? This approach then asks what might be going on computationally in devices that operate that way, for instance, like the subject's "mind." From this it seeks to reformulate a plan about how a learner of this kind might be helped – again within the limits of computability. John Bruer's interesting book *Schools for Thought* (1993) is a nice example of what can be gained from this fresh approach.

But there is an even more interesting third route that computationalists sometimes follow. The work of Annette Karmiloff-Smith (1979, 1992) provides an example if taken in conjunction with some abstract computational ideas. All complex "adaptive" computational programs involve redescribing the output of prior operations in order both to reduce their complexity and to improve their "fit" to an adaptation criterion. That is what "adaptive" means: reducing prior complexities to achieve greater "fitness" to a criterion (Mitchell 1995, Crutchfield and Mitchell 1994). An example will help. Karmiloff-Smith notes that when we go about solving particular problems, say language acquisition, we characteristically "turn around" on the results of a procedure that has worked locally and try to redescribe it in more general, simplified terms. We say, for example, "I've put an s at the end of that noun to pluralize it; how about doing the same for *all* nouns?" When the new rule fails to pluralize *woman*, the learner may generate some additional ones. Eventually, he ends up with a more or less adequate rule for pluralizing, with only a few odd "exceptions" left over to be handled by rote. Note that in each step of this process that Karmiloff-Smith calls "redescription," the learner "goes meta," considering how he is thinking as well as what he is thinking about. This is the hallmark of "metacognition," a topic of passionate interest among psychologists – but also among computational scientists.

That is to say, the rule of redescription is a feature of *all* complex "adaptive" computation, but in the present instance, it is also a genuinely interesting *psychological* phenomenon. This is the rare music of an overlap between different fields of inquiry – if the overlap turns out to be fertile. So, REDESCRIBE, a TOE-like rule for adaptive computational systems that also happens to be a good rule in human problem solving, may turn out to be a "new frontier." And the new frontier may turn out to be next-door to educational practice.

So the computationalist's approach to education seems to take three forms as noted. The first reformulates old theories of learning (or teaching, or whatever) in computable form in the hope that the reformulation will yield surplus power. The second analyzes rich protocols and applies the apparatus of computational theory to them to better discern what might be going on computationally. Then it tries to figure out how the process can be helped. This, in effect, is what Newell, Shaw, and Simon did in their work on the General Problem Solver, and what is currently being done in studies of how "novices" become "experts" (Chipman and Meyrowitz 1993). Finally there is the happy fortuity where a central computational idea, like "redescription," seems to map directly onto a central idea in cognitive theory, like "metacognition."

The culturalist approaches education in a very different way. Culturalism takes as its first premise that education is not an island, but part of the continent of culture. It asks first what function "education" serves in the culture and what role it plays in the lives of those who operate within it. Its next question might be why education is situated in the culture as it is and how this placement reflects the distribution of power, status, and other benefits. Inevitably, and virtually from the start, culturalism also asks about the enabling resources made available to people to cope and what portion of those resources is made available through "education," institutionally conceived. And it will constantly be concerned with constraints imposed on the process of education – external ones like the organization of schools and classrooms or the recruitment of teachers and internal ones like the natural or imposed distribution of native endowment, for native endowment may be as much affected by the accessibility of symbolic systems as by the distribution of genes.

Culturalism's task is a double one. On the "macro" side, it looks at the culture as a system of values, rights, exchanges, obligations, opportunities, and power. On the "micro" side, it examines how the demands of a cultural system affect those who must operate within it. In that latter spirit, it concentrates on how individual human beings construct "realities" and meanings that adapt them to the system, at what personal cost, with what expected outcomes. While culturalism implies no particular view concerning inherent psycho-biological constraints that affect human functioning, particularly meaning making, it usually takes such constraints for granted and considers how they are managed by the culture and its instituted educational system.

Although culturalism is far from computationalism and its constraints, it has no difficulty incorporating its insights – with one exception. It obviously cannot rule out processes relating to human meaning making, however much they do not meet the test of computability. As a corollary, it cannot and does not rule out subjectivity and its role in culture. Indeed, as we shall see, it is much concerned with intersubjectivity – how humans come to know "each other's minds." In both these senses, culturalism is to be counted among the "sciences of the subjective." And, in consequence, I shall often refer to it as the "cultural psychological" approach, or simply as "cultural psychology." For all that it embraces the subjective in its purview and refers often to the "construction of reality," cultural psychology surely does not rule out "reality" in any ontological sense. It argues (on epistemological grounds) that "external" or "objective" reality can only be known by the properties of mind and the symbol systems on which mind relies (Goodman 1978).

A final point relates to the place of emotion and feeling. It is often said that all "cognitive psychology," even its cultural version, neglects or even ignores the place of these in the life of mind. But it is neither necessary that this be so nor, at least in my view, is it so. Why should an interest in cognition preclude feeling and emotion (see e.g. Oatley 1992)? Surely emotions and feelings are represented in the processes of meaning making and in our constructions of reality. Whether one adopts the Zajonc view that emotion is a direct and unmediated response to the world with subsequent cognitive consequences or the Lazarus view that emotion requires prior cognitive inference, it is still "there," still to be reckoned with (Zajonc 1980, 1984, Lazarus 1981, 1982, 1984). And as we shall see, particularly in dealing with the role of schools in "self-construction," it is very much a part of education.

References

Brinton, C. (1965): *The Anatomy of Revolution*. New York: Vintage Books.

Bruer, J.T. (1993): *Schools for Thought: A Science of Learning in the Classroom*. Cambridge, MA: MIT Press.

Bruner, J. (1985): "Narrative and Pragmatic Modes of Thought." In E. Eisner (ed.): *Learning and Teaching the Ways of Knowing: Eighty-Fourth Yearbook of the National Society for the Study of Education*. Chicago: Chicago University Press, pp. 97–115.

Bruner, J. (1990): *Acts of Meaning*. Cambridge, MA: Harvard University Press.

Chi, M.T.H., Glaser, R., and Farr, M.J. (eds.) (1988): *The Nature of Expertise*. Hillsdale, NJ: Lawrence Erlbaum.

Chipman, S., and Meyrowitz, A.L. (1993): *Foundations of Knowledge Acquisition: Cognitive Models of Complex Learning*, vols. 1 and 2. Boston: Kluwer Academic Publishers.

Crutchfield, J. P, and Mitchell, M. (1994): *The Evolution of Emergent Computation*. Santa Fe Institute Technical Report 94–03–012, Santa Fe, NM: Santa Fe Institute.

Goodman, N. (1978): *Ways of Worldmaking*. Indianapolis, IN: Hackett.

Grice, H.P (1989): *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Karmiloff-Smith, A. (1979): *A Functional Approach to Child Language: A Study of Determiners and Reference*. Cambridge, UK: Cambridge University Press.

Karmiloff-Smith, A. (1992): *Beyond Modularity: A Developmental Perspective on Cognitive Science*. Cambridge, MA: MIT Press.

Kroeber, A.L. (1917): "The Superorganic." *American Anthropologist*, 19 (2), pp. 163–213.

Lazarus, R.S. (1981): "A Cognitivist's Reply to Zajonc on Emotion and Cognition." *American Psychologist*, 36, pp. 222–223.

Lazarus, R.S. (1982): "Thoughts on the Relations between Emotion and Cognition." *American Psychologist*, 37 (9), pp. 1019–1024.
Lazarus, R.S. (1984): "On the Primacy of Cognition." *American Psychologist*, 39 (2), pp. 124–129.
McClelland, J.L. (1990): "The Programmable Blackboard Model of Reading." In D.E. Rumelhart and J.L. McClelland: *Parallel Distributed Processing: Psychological and Biological Models*, vol. 2. Cambridge, MA: MIT Press, pp. 122–169.
Mitchell, M. (1995): "What Can Complex System Approaches Offer the Cognitive Sciences?" Paper presented at the Annual Meeting of the Society for Philosophy and Psychology, State University of New York at Stony Brook, June 10.
Oatley, K. (1992): *Best Laid Schemes: The Psychology of Emotions*. Cambridge, UK: Cambridge University Press.
Olson, D.R. (1994): *The World on Paper: The Conceptual and Cognitive Implications of Writing and Reading*. Cambridge, UK: Cambridge University Press.
Ong, W.J. (1991): *Orality and Literacy: The Technologizing of the Word*. London: Routledge.
Rumelhart, D.E., and McClelland, J.L. (eds.) (1986): *Parallel Distributed Processing: Explanations in the Microstructure of Cognition*, vols. 1 and 2. Cambridge, MA: MIT Press.
Schank, R.C. (1990): *Tell Me a Story: A New Look at Real and Artificial Memory*. New York: Scribner.
Segal, J.W, Chipman, S.F., and Glaser, R. (1985): *Thinking and Learning Skills: Volume 2, Relating Instruction to Research*. Hillsdale, NJ: Lawrence Erlbaum.
Sperber, D., and Wilson, D. (1986): *Relevance: Communication and Cognition*. Cambridge, MA: Harvard University Press.
von Wright, G.H. (1971): *Explanation and Understanding*. Ithaca, NY: Cornell University Press.
Zajonc, R.B. (1980): "Feeling and Thinking: Preferences Need No Inferences." *American Psychologist*, 35 (2), pp. 151–175.
Zajonc, R.B. (1984): "On the Primacy of Affect." *American Psychologist*, 39 (2), pp. 117–123.