Distributed Processes, Distributed Cognizers and Collaborative Cognition

Stevan Harnad
Chaire de recherche du Canada
Centre de neuroscience de la cognition (CNC)
Université du Québec à Montréal
Montréal, Québec, Canada H3C 3P8
http://www.crsc.uqam.ca/en/index2_en.html
harnad@uqam.ca

and

Department of Electronics and Computer Science
University of Southampton
Southampton SO17 1BJ
United Kingdom
http://www.ecs.soton.ac.uk/~harnad/
harnad@ecs.soton.ac.uk

ABSTRACT: Cognition is thinking; it feels like something to think, and only those who can feel can think. There are also things that thinkers can do. We know neither how thinkers can think nor how they are able do what they can do. We are waiting for cognitive science to discover how. Cognitive science does this by testing hypotheses about what processes can generate what doing ("know-how") This is called the Turing Test. It cannot test whether a process can generate feeling, hence thinking -- only whether it can generate doing. The processes that generate thinking and know-how are "distributed" within the heads of thinkers, but not across thinkers' heads. Hence there is no such thing as distributed cognition, only collaborative cognition. Email and the Web have spawned a new form of collaborative cognition that draws upon individual brains' real-time interactive potential in ways that were not possible in oral, written or print interactions.

In our age of "virtual reality," it is useful to remind ourselves now and again that a corporation cannot literally have a head-ache, though its (figurative) "head" (CEO) might. And that even if all N members of the Board of Directors have a head-ache, that's N head-aches, not one distributed head-ache. And that although a head-ache itself may not be localized in one point of my brain, but distributed across many points, the limits of that distributed state are the boundaries of my head, or perhaps my body: the head-ache stops there, and so does cognition. If a mother's head-ache is her three children, then her children get the distributed credit for causing the state, but they are not part of the state. And if the domestic economic situation is a "head-ache," that distributed state is not a cognitive state, though it may be the occasion of a cognitive state within the single head of a single head of state (or several cognitive states in the individual heads of several individual heads of state).

The problem, of course, is with the vague and trendy word "cognition"(and its many cognates: *cognize*, *cogitate*, Descartes "cogito," and of course the Hellenic forebears: *gnosis*, *agnosia* and *agnostic*). William of Occam urged us not to be profligate with

"entities": *Entia non sunt multiplicanda praeter necessitatem*. But entities are just "things"; and we can presumably name as many things as we can think of. Cognition, at bottom, is, after all, *thinking*.

So there are naturally occurring things (e.g., people and animals) that can *think* (whatever "thinking" turns out to be). Let's call that "Natural Cognition (NC)" And there are artificial things (e.g., certain machines), that can *do* the kinds of things that naturally thinking things can do; so *maybe* they can think too: "Artificial Cognition (AC)" And then there are *collections* of naturally thinking things – or of naturally thinking things plus artificially thinking things – that can likewise do, collectively, the kinds of things that thinking things can do individually, so maybe they can think, collectively, too: "Distributed Cognition" (DC).

But what about the head-aches, and Descartes, and Occam? Descartes was concerned about what we can be absolutely *sure* about, beyond the shadow of a doubt. He picked out the entity we want to baptize as "thinking": It's *whatever is going on in our heads when we are thinking* that we want to call "thinking:" We all know *what* that is, when it's happening; no way to doubt that. But we can't be sure *how* we do it (cognitive science will have to tell us that). Nor can we be sure that others can do it (but if they are sufficiently like us, it's a safe bet that others can think too). And we can't even be sure the thinking – however it works – is actually going on within our heads rather than elsewhere (but there too, it's a safe bet that it's all happening inside our heads, or no riskier than the bet that we have heads at all).

So it seems that head-aches and cognition cover the same territory: heads ache and heads think, and just as it is self-contradictory to deny that my head's aching (when it's aching), it's self-contradictory to deny that my head's thinking (when it's thinking): How do I *know* (for sure) in both cases? Exactly the same way: it *feels like* something to have a head-ache and it *feels like* something to think: No feeling: no aching. And, by the same token: No feeling: no thinking.

No feeling: no thinking? What about Freud, and the thoughts being thought by my unconscious mind? Well, we have Occam's authority to forget about extra entities like "unconscious" minds unless we turn out to need them to explain what there is to explain. Let's say the jury's still out on that one, and one mind seems enough so far. But can't one and the same mind have both conscious and unconscious (i.e., felt and unfelt) thoughts?

Let's try that out first on head-aches: Can I have both felt and unfelt head-aches? I'd be inclined to say that – inasmuch as a head-ache is associated with something going wrong in my head: a constricted blood vessel, for example – I might have a constricted blood vessel without feeling a head-ache. But do I want to call that an unfelt head-ache? By the some token, if I feel a head-ache *without* a constricted blood vessel, do I then want to say I don't have a head-ache after all? Surely the answer is No in both cases: There are no unfelt head-aches, and when I feel a head-ache, that's a head-ache no matter what else is or is not going on in my head. That's what I *mean* by a head-ache.

So now what about thinking? Can I be thinking when I don't feel I'm thinking? We've frankly confessed that we don't yet know *how* we think; we're waiting for

cognitive science to discover and tell us how. But are we ready to accept -- when we are (say) thinking about nothing in particular -- that we are in fact thinking that "the cat is on the mat," because (say) a brain scan indicates that the activity normally associated with thinking that particular thought is going on in our brains at this very moment? Surely we would say – as with the constricted blood vessel when I do not have a head-ache : "Maybe that activity is going on in my head right now, but that is not what I am thinking! So call it something else: a brain process, maybe. But it's not what I'm thinking unless I actually feel I'm thinking it (at the time, or once my attention is drawn to it).

Perhaps too many years of Freudian profligacy have made you feel that quibbling about whether or not that brain process is a "thought" is unwarranted. Then try this one: Suppose you *are* thinking that "the cat is on the mat," but one of the accompanying brain processes going on at that moment is the activity normally associated with thinking that "the cat is *not* on the mat"? Are you still prepared to own to be the thinker of that unfelt thought, the opposite of the one you feel you are thinking?

Perhaps the years of Freudian faith in the existence of alter egos co-habiting your head have made you ready to accept even that contradiction (but then I would like to test your faith by drawing your attention to a brain process that says you want to sign over all your earthly property to me, irrespective of what might be your conscious feelings about the matter!). At the very least, our subjective credulity about unconscious alter egos -- cohabiting the same head but of a different mind about matters than our own -- has its objective limits.

A more reasonable stance (and Occam would approve) is to agree that we know *when* we are thinking, and *what* we are thinking (that thought), just as we know when we are feeling a head-ache and what we are feeling (that head-ache), but we do now know *how* we are thinking, any more than we know what processes underlie and generate a head-ache. And among those unknown processes might be components that *predict* that we may be having *other* thoughts at some later time -- just as a vasoconstriction without a head-ache may be predictive of an eventual head-ache (or stroke), and a process usually associated with thinking that "the cat is *not* on the mat" may be predictive of eventually thinking that the cat is not on the mat, even though I am at the moment still thinking that the cat is on the mat.

So we are waiting for cognitive science to provide the functional explanation of thinking just as we are waiting for neurovascular science to provide the functional explanation of head-aches. But we have no doubt about when we are and are not thinking -- and almost as little doubt about what we are and are not thinking -- as we have about when we are and are not having a head-ache. There is plenty of scope for unfelt accompanying or underlying *processes* here; just not for unfelt thoughts (or head-aches).

This also brings us back to the three candidate forms of thinking: natural/individual (NC), artificial (AC), and collective/distributed (DC). It is important to stress the "collective" aspect of distributed cognition, because of course there is already a form of "distributed" cognition in the natural/individual case (NC). This is thinking that is taking place within one's own head, but with the associated processes being

distributed across the brain in various ways. It is not very useful to speak of this as "distributed cognition" (DC) at all (though we will consider a hypothetical variant of it later): It is clearly distributed *processes* that somehow underlie and generate the thinking; some of those processes might have felt correlates, many of them may not. For there is another correlate of thinking, apart from the processes that accompany and generate the thinking, and that is the *doing* (or rather the doing *capacity* and *tendency*) that likewise accompanies thinking.

Here it is useful (and Occam would not disapprove) to admit a near-synonym of thinking – namely, *knowing* -- that is even more often used as the anglo-saxon cognate term for the Latinate "cognition" than "thinking" is. But "knowing" has a liability: it has a gratuitous bit of certainty about it that over-reaches itself, going beyond Descartes' careful delineation of what is certain and necessarily true from what is just highly probably true: When I feel a head-ache, I *know* I have a head-ache, but I merely *think* I have vasoconstriction. That's fine. Both count as cognizing something. But when I think "the cat is on the mat," I certainly don't *know* the cat is on the mat, yet I'm still cognizing. [Strictly speaking, even when I *see* the cat is on the mat, I don't know it for sure, in the way that I know for sure that I have a head-ache or that 2+2=4, or that it looks/feels-as-if the cat is on the mat -- but there's no need here to get into the irrelevant philosophical puzzles ("the Gettier problems") about the differences between knowing something and merely thinking something that happens to be true.]

The reason knowing is sometimes a useful stand-in for thinking is that it ties cognition closer to action: There is "knowing-that" – which is very much like "thinking-that," as in thinking/knowing that "the cat is on the mat." And there is "knowing-how," as in knowing how to play chess or tennis. (Know-how has no counterpart when we speak only about thinking rather than knowing.) Skill or know-how is something I *have*, and its "proof" is in *doing* it (not just in thinking I can do it). Now an argument can be made for the fact that know-how is not cognition at all. Know-how may (or may not) be acquired consciously and explicitly; but once one has a bit of know-how, one simply has it; conscious thinking is not necessarily involved in its exercise (though one usually has to be awake and conscious to exercise it).

But if know-how were excluded from cognition because it did not necessarily involve conscious thinking, then we would have to exclude all the other unconscious processes underlying the "how" of thinking itself! So whereas thinking itself is the necessary and sufficient condition for being a cognitive system, thinking in turn has necessary conditions of its own too, and most of those are unconscious processes.

The same is true of know-how: it is generated by unconscious processes, just as thinking is. Know-how may or may not be acquired, and if acquired, it may or may not be acquired via thinking (though one almost certainly must be awake and thinking while acquiring it); and know-how may or may not be exercised via thinking (though one almost certainly must be awake and thinking while exercising it). Moreover, just about all thinking (including knowing-that) also has a know-how dimension associated with it. If I think that "the cat is on the mat" is true then I know it follows that the "the cat is not on the mat" is false. Thoughts are not punctate. They have implications. And the implications are part of the know-how implicit in the thought itself. I know how to reply to (many) questions about the whereabouts of the cat if I

think that "the cat is on the mat." And the know-how goes beyond the bounds of thinking and even talking about what I think: it includes doing things in the world. If I think that "the cat is on the mat," I also know how to go and find the cat!

All this belaboring of the obvious is intended to bring out the close link between thinking capacity and doing capacity (via know-how). Which brings us to the second case of cognition: "artificial cognition" (AC). If there are things other than living creatures (e.g., certain machines) that can *do* the kinds of things that living/thinking things can do, then *maybe* they can think too. Note the "maybe." It is quite natural to turn to machines in order to explain the "how" of cognition. Unlike the know-how of the heart or the lungs, the brain's know-how is unlikely to be discoverable merely from observing what the brain can do and what's going on inside it while it is doing so. That might have been sufficient if all the brain could do was to move (in the gross sense of navigating in space and manipulating objects). But the brain can do a lot of subtler things than just walking around and fiddling with objects: it can perceive, categorize, speak, understand and think. It is not obvious how the know-how underlying all those capacities can be read off of brain structure and function. At the very least, trying to design machines that can also do what brains (of humans and animals) can do is a way of testing theories about how such things can be done *at all*, any which way. In addition, it puts the power of both computation and neural simulation in the hands of the theorist.

So whether or not machines will ever be able to think -- and please remember that by "think" we mean being able to have the feeling, rather like a head-ache, that we have agreed with Descartes to call "thinking"-- we in any case need machines in order to study and explain the *how* of thinking – the know-how that normally underlies and accompanies the feeling.

Can machines think (Turing 1950)? The answer depends in part on the rather more arbitrary notion we have of what a machine is, compared to our clear, Cartesian notion of what thinking is: Is a system that is designed and assembled out of biomolecules by people a machine? What if some of its components are synthetic? All of its components? What if toasters grew on trees – would they then not be machines? Maybe there is no natural kind corresponding to "machine." Maybe all autonomous, moving/functioning systems, whether man-made or nature-made, are machines, and the only substantive question is: which kinds of machines can do which kinds of things?"

So we are not guaranteed to be right about machine thinking: There will be no Cartesian certainty there, as with out own. Maybe machines will never be able to *feel*: maybe they will only be able to *do*. So research on artificial cognition is, strictly speaking, research on what sorts of machine processes can successfully generate the kinds of know-how that we have -- the ability to *do* the kinds of things that we can do, including, of course, speaking, and replying (verbally, as well as responding with other coherent actions) to what is said. That is the methodological idea behind the Turing Test: Thinking is as thinking does. Once a machine can do anything a person can do, and do it in a way that is indistinguishable to any person from the way any other person does it, do we have any better grounds for doubting whether the machine thinks than we have for doubting whether any person other than myself thinks? There is no Cartesian certainty in any case but my own.

So as soon as we turn from the "whether" question to the "how" question about thinking, we must have recourse to the know-how of machines in order to test theories about the know-how of our brains. That puts artificial cognition on a methodological and epistemic par with natural cognition. We have already agreed, moreover, that natural cognition is "distributed" over various parts of the brain, but that it is more accurate to refer to this as the distributed processes underlying or generating cognition (like the distributed processes underlying or generating a head-ache) rather than as "distributed cognition." No doubt machine processes that can do as natural cognition does will be distributed too, across the insides of the machine: Or across several machines? We have admitted that the notion of "machine" is fuzzy. By the same token, is the notion of "one machine" not equally fuzzy (Harnad 2003b)?

We have agreed that the machine must be autonomous; it must be able to do whatever it can do *on its own*, without any help from us, otherwise it is partly *our* capacities that are being exhibited and tested in our joint performance capacity. But whereas a machine's know-how must be independent of ours, does it need to be independent of the know-how of other machines? We are in the same situation here as in the case of the distributed brain processes inside our heads: In the case of our heads, we are talking about the distributed processes that generate two things: our thinking (1) and our know-how (2).

The thinking (1) is a unitary, felt state, like a head-ache, but the physical processes generating it are distributed – distributed, however, only within my head. It is a logical possibility that the physical processes generating my thought that "Venus is a far-away planet" consist of a widely distributed state that includes my head and Venus and perhaps a lot of other components outside my head, just as it is a logical possibility that the physical processes generating my seaside head-ache consist of a widely distributed state that includes my head and the sun and perhaps a lot of other components outside my head. But the probability of such distributed out-of-body feeling-states is of about the same order as the probability of telekinesis, clairvoyance or reincarnation -- or of the possibility that no one other than myself feels. So let us agree to ignore such far-fetched logical possibilities: The limits of my thinking-states are the limits of the processes going on in my brain.

(A bit later, we will consider hypothetical future out-of-body brain prostheses – analogues of artificial kidney machines that perform some brain functions extracorporeally, via physical links or telemetry. But first we need to return to machine know-how.)

Do we have any reason, with machines, to assume that artificial cognition cannot be distributed across machines?

First, remember that there is an element of Cartesian certainty about natural cognition that has been replaced by mere Turing probability in the case of artificial cognition. The element that has been replaced is actually the essence of thinking, which is that thinking is a form of feeling (1), but a form of feeling that is also closely associated with a capacity and propensity for doing, i.e., with know-how (2). The Turing Test is based purely on (2), with (1) being taken on faith: faith in the same telepathic "mind-reading" powers that each of us uses every day to detect whether and what other

people are thinking (Baron-Cohen et al. 2000, Nichols & Stich 2004; Premack & Woodruff 1978).

So if (i) thinking is as thinking does -- and hence (ii) if one machine can do anything and everything a thinking person can do, then it can think -- then what if one machine cannot do it all, but two, three or ten jointly can? We allowed that the brain process generating natural cognition could be distributed across the brain: why can't the machine processes generating artificial cognition be distributed across multiple machines?

We are partly up against the arbitrariness of what we mean by (one) "machine" again. It is easy to individuate and count Cartesian thinkers: Each one has a mind (and brain) of his own. Ask them and they will tell you so. You can take a roll call; and if you are one of those thinkers, you know you are not all or part of any of the other thinkers or vice versa. And "mind-reading" aside, the only thoughts (and the only head-aches) each thinker is privy to are his own; and these all occur within the confines of his head. But how do we individuate machines at all (even setting aside the question of whether they can think)?

A toaster seems well individuated: It's a device that bronzes bread. But that seems to pick out an entity only because we are interested in bronzing bread. If the bread-bronzer were just a component of a more complicated Rube-Goldberg device -- that drops bread into the bronzing component, which then pops up and hits a bell that triggers a lead ball to roll down a tube onto a roll of toothpaste, whose contents this squeezes onto a rotating electric brush, triggering it into motion – do we have a toaster here or a toothbrush? And how many "machines" are involved?

Individuating machines is not quite as hopeless as this suggests, however, for there are still two non-arbitrary criteria we can use, one sufficient to individuate machines and the other sufficient to individuate "thinking" machines (AC): First, the machine must be *autonomous*: Given its Input (I) it must generate its Output (O) without any outside help (otherwise its boundaries would be indeterminate and it would not have been individuated). Second, its I/O task should be the same as that of a natural living kind: The Turing Test.

So whatever autonomous system can pass the Turing Test counts as one thinking machine: it has the know-how of the corresponding natural thinker (us). (It should be obvious that if one of the components of this machine were itself a natural thinker, that would be cheating, because the whole purpose of the Turing Test is to explain *how* natural thinkers can think, by designing an artificial thinker using components for which we already know how they work. Using a natural thinker as one of the components would just compound the mystery, and leave the "explanation" ungrounded.)

But apart from not including any unexplained components, the Turing-Test-passing machine is free to be any autonomous system that can successfully pass the Turing Test. This entails a lot of constraints already, for our I/O capacity consists entirely of things that we do in space and time with our bodies. So the candidate would have to be a robot; and since it must be able to do anything and everything we can do, in real time, and indistinguishably from the way we do it (it can't navigate a room by sending

out a parallel proxy in all directions, nor can it take a lifetime to make a chess move), it has its work cut out for it. Still, there is no reason that all of its hardware would need to be located inside the robot. The autonomous system that passes the Turing Test could be a distributed one, with some of it functions inside the robot, others in a remote control station.

Even in the case of human cognition, the future possibility of remote prostheses is not out of the question. What is not negotiable, however, is the autonomy of the system and the unity of feeling, hence thinking: A real brain with synthetic remote prostheses could in principle have a distributed head-ache, with the feeling state literally taking place both in and outside the head. (Remove or inactivate either component and the head-ache vanishes.) So, by the same token, both natural and artificial cognition could be distributed in this sense: the generating processes – already "distributed" spatially within the brain – could have their spatial distribution widened beyond the confines of the person's or robot's head. Nothing really radical about that. But the natural thinker would still be thinking its own individual thoughts, and feeling its own individual head-aches.

About the robot there is no way to know for sure (without actually *being* the robot) whether it is indeed thinking (rather than merely doing: i.e., exhibiting the know-how that is normally generated along with the thinking in the case of thinkers, but without the thinking, because the robot feels nothing at all, neither thoughts nor head-aches). But if the Turing-Test passing robot is indeed thinking, hence feeling, then it too will be thinking its own individual thoughts and feeling its own individual head-aches, whether its hardware is distributed remotely or all contained locally.

Let us call these two relatively uncontroversial forms of distributedness the "distributed processing" that generates cognition, rather than "distributed cognition" (DC), which we have reserved for the third putative kind of cognition. And let us summarize what is certain, what is probable, and what is possible:

It is certain that I think. It is highly probable that other people think too. It is highly probable that my thinking occurs only within my own head, but that the processes generating my thinking are distributed within my brain. (It is not even clear how "local" as opposed to distributed the brain processes corresponding to a thought would have to be in order to be "nondistributed": surely even if a thought were generated by the presence of a single molecule, the molecule itself is distributed in space!) It is highly improbable that anything other than humans and other animals can think today. It is highly probable (for the same reason that it is highly probable that other people think too) that anything that can do anything and everything a person can do (indistinguishably from a person, for a lifetime) can think too. So it is highly probable that a robot that could pass the Turing Test would be able to think. It is possible for the processes generating the thinking in both humans and robots to be distributed more widely than just within their respective heads.

Now what about the possibility of true distributed collective cognition, where there is thought generated by distributed processes, some or all of whose constituents are themselves natural or artificial thinkers?

Let us set aside the trivial, unproblematic cases first:

If two people are talking, that's not DC, that's a conversation; same if they're emailing; same if it's N people. Let's call that Collaborative Cognition: CC.

If a person uses a computer or a database, that's not DC, that's human/machine interaction, computation and consultation; with N people jointly using N computers or databases, it's again CC, not DC.

If N people use N computers to gather or process data, that's not DC, that's human/machine interaction and human/human collaboration, i.e., CC; same if N people jointly write and revise a text, or N texts.

If a robot controlled by N people, or N people plus N computers, passes the Turing Test, that too is human/machine interaction and human/human collaboration. It is neither DC nor AC; just NC plus CC (and the Turing Test has not been passed)

If a robot controlled by N computers passes the Turing Test, that is not DC but AC. If the autonomous system consisting of the robot plus the N computers not only has know-how, but it also thinks, then that is distributed processing generating thought.

If N robots that can pass the Turing Test email to one another, use computers, gather and process data and jointly write and revise texts, that too is CC, not DC.

So, thus far, *nothing* is DC. What would it take to generate genuine DC, rather than merely distributed processes generating NC or AC, or distributed NC and AC cognizers collaborating to produce CC? The head-ache test is the decisive one: If the autonomous system consisting of the NC and AC cognizers (plus any other constituents you may wish to add) somehow becomes the kind of system capable of feeling a head-ache, then it is the kind of system capable of thinking a thought, and its constituents have collectively managed to generate DC.

Things nearly as wondrous have happened: If the (distributed) nonliving components and processes inside living single-celled organisms are analogous to the distributed processes that generate NC (and perhaps eventually AC), then the single multi-cellular organism that is generated by the distributed single cells and other components of which it is composed would be analogous to DC: A living thing constituted out of living things is like a thinking thing constituted out of thinking things. But the latter is highly improbable.

Not to close on an improbable note: Even if there is no DC, but only CC, wondrous things can still arise from it. We could say that all human civilization and knowledge to date already arises from CC: cumulative, collective, collaborative know-how. But with the age of the computer and the Internet, the power and possibilities of CC take a quantum leap. Consider the milestones that have accorrred in cognitive evolution (Harnad 1991; Cangelosi & Harnad 2002):

The first cognitive milestone was the evolution of language, millions of years ago, through organic adaptive change in our brains that allowed human cognizers to communicate and collaborate digitally and symbolically, instead of just instrumentally and through sensorimotor imitation, as other species do. This was the greatest

cognitive milestone of all, for with it came not only the full power of language to express, describe and explain just about anything, but implicit in it also (although only to be exploited much, much later in human history) was the power of computation to simulate and model just about anything. Language co-evolved with the power of thinking itself (the "language of thought"), and indeed the speed of conversation and the speed of thought are of roughly the same magnitude, allowing cognizers to interdigitate their thoughts, collaborating synchronously, in real time (local CC). Language also allowed human knowledge to be formulated explicitly and to be passed on by word of mouth (the "oral tradition"). This was a form of serial collaboration and cumulation, with each successive narrator elaborating the cumulative record in his own way (distal CC).

Being oral, language provided a lot of scope for real-time colloquy and collaboration, but being dependent on serial hearsay for transmission and preservation, its cumulative record was labile and unreliable. So the next cognitive milestone was the invention of the lapidary medium: writing .This allowed the fruits of human collaboration and thinking to be faithfully recorded, preserved and transmitted speaker-independently – "off-line,"so to speak ("verba volant; scripta manent"). The offline, asynchronous, written medium thereby became far more powerful than the online, synchronous oral medium for the dissemination, reliability and longevity of human knowledge; but it lacked much of that real-time interactivity for which language and the speed of thought had co-evolved. Hence writing fell out of phase with the potential speed and power of interactive online thought -- although it did at the same time foster the skills of solo offline thought: the written tradition.

Print was the third cognitive milestone, radically extending the reach of the written tradition, now scribbler-independently, but still out of phase with the full speed, power and interactivity of real-time, interdigitating thought. Cognitive collaboration was still either oral and synchronous (leaving no record, until the advent of real-time audio recording) or written and asynchronous, hence far slower and less interactive. Nor did the type-writer or even the word-processor bridge the temporal gap between parallel and serial cognitive collaboration.

So many potential cycles of productive interaction were lost: until the temporal gap between the conversational speed of interdigitating thought for which our brains are adapted and the much slower tempo of dissemination of written text was at last bridged again by email and the Internet, the fourth cognitive milestone: "scholarly skywriting" (Harnad 2003a). It is now possible for a text to be written, transmitted and responded to in real time, at almost conversational speed (i.e., the speed of thought), as if it were all being written in the sky, for all to see and respond to -- in real time if they wish. Perhaps just as important, it is possible to *quote/comment* text (by living and active or even long-dead authors) and to branch that collaborative interaction instantaneously to many other potential interlocutors, and potentially the whole planet, through email, hypermail, blogs, and web archives.

Now it was never the strength of the oral tradition to have several people speaking at once. Conversation is optimal when it is serial and one-on-one, or with several interlocutors turn-taking – again serially, but in real time. Moreover, not everyone has (or should have) something to say about everything. So there are no doubt constraints and optima that will emerge with skywriting as the practice develops. But right now,

the problem is not an excess or embarrassment of skywritten riches, producing an un-navigable din, but a dearth of online scholarly content and CC: Most of cyberspace is still devoted to trivial pursuit, not to CC.

This will soon change: Skywriting itself is one of its own sure rewards: It was the presence of an audience that inspired the eloquence of the bard, the oracle and the sage in the days of the oral tradition. Writing in the skies, instantly visible to one's peers, is one incentive for scholarly CC. So is the prospect (and provocation) of "creative disagreement" (Harnad 1979, 1990). The likelihood of their texts being seen, scrutinized , criticized , used, applied and built-upon by their peers inspires scholars both to skywrite and to be careful and rigorous; having their skywritings criticized or elaborated in turn inspires further iterations of skywriting. Soon shared research-data and joint data-analyses too will become part of the skywriting. This is all CC.

The impact of scholarly writing was already being measured and rewarded in Gutenberg days (by counting journal citations); skywriting offers many new ways of monitoring, measuring, maximizing, evaluating and rewarding the impact of CC through the analysis of (distributed!) patterns in downloads, citations, co-citations, co-authorships, and even co-text (Brody & Harnad 2005).

All of this is CC. It is the fruit of the collective , interactive know-how of many individual thinkers. If it goes wrong, it will inspire many individual head-aches, not one distributed one. And if it inspires pride, that will be felt by many individual cognizers, not one distributed one.

REFERENCES

Baron-Cohen, S., Tager-Flusberg, H. & Cohen, D.J. (2000) (eds.) Understanding other minds. N.Y.: Oxford University Press

Brody, T. and Harnad, S. (2005) Earlier Web Usage Statistics as Predictors of Later Citation Impact. JASIST (in press) http://eprints.ecs.soton.ac.uk/10713/

Cangelosi, A., Greco, A. & Harnad, S. (2002) Symbol Grounding and the Symbolic Theft Hypothesis. In: Cangelosi, A. & Parisi, D. (Eds.) Simulating the Evolution of Language. London, Springer.
http://cogprints.org/2132

Harnad, S. (1979) Creative disagreement. The Sciences 19:18 - 20.
http://eprints.ecs.soton.ac.uk/10852/

Harnad, Stevan (1990) Creativity: Method or Magic?
http://cogprints.org/1627/

Harnad, Stevan (1991) Post-Gutenberg Galaxy: The Fourth Revolution in the Means of Production of Knowledge. Public-Access Computer Systems Review 2(1):pp. 39-53. http://cogprints.org/1580/

Harnad, Stevan (1995) Interactive Cognition: Exploring the Potential of Electronic Quote/Commenting, in Gorayska, B. and Mey, J.L., Eds. Cognitive Technology: In Search of a Humane Interface, pages pp. 397-414. Elsevier. http://cogprints.org/1599/

Harnad, Stevan (2003a) Back to the Oral Tradition Through Skywriting at the Speed of Thought. Interdisciplines.
http://cogprints.org/3021/

Harnad, S. (2003b) Can a Machine Be Conscious? How? Journal of Consciousness Studies. http://cogprints.org/2460/

Harnad, S. (2005) Searle's Chinese Room Argument. To appear in Encyclopedia of Philosophy (Macmillan).
http://eprints.ecs.soton.ac.uk/10424/01/chineseroom.html

Nichols, S. & Stich (S) 2004. Mindreading. Oxford: Oxford University Press.

Premack, D. & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? Behavioral & Brain Sciences, 4, 515-526.

Turing, A. M. (1950) Computing Machinery and Intelligence. Mind 49:433-460.
http://cogprints.org/499/

KEYWORDS: Cognition, computation, artificial intelligence, Turing Test, neural networks, collaboration, robotics, consciousness, feeling, thinking, Descartes, mind-reading, open access, interoperability

STEVAN HARNAD, born in Hungary, did his undergraduate work at McGill and his doctorate at Princeton and is currently Canada Research Chair in Cognitive Science at University of Quebec/Montreal and adjunct Professor at Southampton University, UK. His research is on categorisation, communication and cognition. Founder and Editor of Behavioral and Brain Sciences, Psycoloquy and CogPrints Archive. He is Past President of the Society for Philosophy and Psychology, Corresponding Member of the Hungarian Academy of Science, and author and contributor to over 150 publications.