Cohabitation: Computation at 70, Cognition at 20

Stevan Harnad

Canada Research Chair in Cognitive Sciences
Université du Québec à Montréal
http://www.crsc.uqam.ca/
and
Department of Electronics and Computer Science
University of Southampton
http://www.ecs.soton.ac.uk/~harnad/

ABSTRACT: Zenon Pylyshyn cast cognition's lot with computation, stretching the Church/Turing Thesis to its limit: We had no idea how the mind did anything, whereas we knew computation could do just about everything. Doing it with images would be like doing it with mirrors, and little men in mirrors. So why not do it all with symbols and rules instead? Everything worthy of the name "cognition," anyway; not what was too thick for cognition to penetrate. It might even solve the mind/body problem if the soul, like software, were independent of its physical incarnation. It looked like we had the architecture of cognition virtually licked. Even neural nets could be either simulated or subsumed. But then came Searle, with his sino-spoiler thought experiment, showing that cognition cannot be all computation (though not, as Searle thought, that it cannot be computation at all). So if cognition has to be hybrid sensorimotor/symbolic, it turns out we've all just been haggling over the price, instead of delivering the goods, as Turing had originally proposed 5 decades earlier.

On Knowing How We Generate Our Know-How. In the 1960's, my teacher, D.O. Hebb, had a standard example for his undergraduate intro psych course to show that (what we would now call) "cognition" cannot all amount to just inputs and outputs plus the reward/punishment histories that shaped them, as behaviorism (then already in decline) had taught. He simply presented two single-digit numbers, the first bigger than the second – say, 7 and 2. Next he would pause; and then he would remind us how different our "response" to those two "stimuli" would be if earlier – even much earlier – we had been told "subtract" (or "add" or "multiply").

Now Hebb's example was not decisive – as no refutation of push/pull behaviorism can be decisive. The input sequence "subtract" "7" "2", generating the output "5" can be interpreted as a rote I/O sequence no matter how long the delay between the input "subtract" and the inputs "7", "2". But Hebb's point was a double one, to the effect that, first, surely there is something intervening between the command, the digits, and

the response; and, second, that surely whatever that intervening internal process is, that is the true object of study of (what would today be called) "cognitive science," and not just the inputs, outputs, and their reward histories.

Behaviorism Begged the Question. Or, to put it another way, the task of cognitive science is to explain what equipment and processes we need in our heads in order to be capable of being shaped by our reward histories into doing what we do. Skinner – whom Hebb had described (in part out of modesty, but surely also partly tongue-in-cheek) as the greatest contributor to psychology (sic) at the time – had always dismissed theorizing about how we are able to learn as being either unnecessary or the province of another discipline (physiology), hence irrelevant (to psychology; Harnad 1985; Catania & Harnad 1988). Cognitive science has since re-jigged the disciplinary boundaries somewhat, admitting neurosciences into the ecumenical circle, but it should be noted that Hebb's point was about internal processes in the head that underlie our capacity to do what we can do. There is still a Skinnerian uneasiness about counting the biomolecular details of brain implementation as part of cognitive science. We shall return to this when we discuss the hardware-independence of software.

In essence, Hebb's point was about question-begging: behaviorism was begging the question of "how?" How do we have the behavioral capacity that we have? What makes us able to do what we can do? The answer to this question has to be cognitive; it has to look into the black box and explain how it works. But not necessarily in the physiological sense. Skinner was right about that. Only in the functional, cause/effect sense. And regarding the functional explanation of our behavioral capacities, behaviorism had again been half-right, in its brief against its predecessor, introspectionism. Behaviorists had rightly pointed out that sitting in an armchair and reflecting on how our mind works will not yield an explanation of how it works (except of course in the sense that explanation in all disciplines originates from human observation and reflection).

For this, Hebb had a companion example to his 7/2 task. He would ask the intro class: "What was the name of your 3rd grade school-teacher?" When we triumphantly produced our respective answers, he would ask "How did you do it? How did you find the name?" He might have asked the same of addition and subtraction: "How is it that your head, having been told 'subtract', manages to come up with '5' in response to '7'/'2'? "

Beware of the easy answers: rote memorization and association. The fact that our brains keep unfailingly delivering our answers to us on a platter tends to make us blind (neurologists would call it "anosognosic") to the fact that there is something fundamental there that still needs to be accounted for. Memorizing single-digit sums, products and differences by rote, case by case, covers the trivial cases, but it does not generalize to the nontrivial ones. Surely we have not pre-memorized every possible sum, product and difference?

Is Computation the Answer? Computation already rears its head, but here too, beware of the easy answers: I may do long-division in my head the same way I do long-division on paper, by repeatedly applying a memorized set of symbol-manipulation rules – and that is already a big step past behaviorism – but what about

the things I can do for which I do not know the computational rule? Don't know it consciously, that is. For introspection can only reveal how I do things when I know, explicitly, <u>how</u> I do them, as in mental long-division. But can introspection tell me how I recognize a bird or a chair as a bird or a chair? How I play chess (not what the rules of chess are, but how, knowing them, I am able to play, and win, as I do)? How I learn from experience? How I reason? How I use and understand words and sentences?

Skinner got another famous piece of come-uppance from Chomsky (1959), who pointed out how question-begging was the very idea that our linguistic capacity (in particular our syntactic capacity), which Chomsky called our <u>competence</u>, can be explained as having been "shaped" by our history of hearing, speaking, reward and punishment. Grammar – at least the Chomskian "Universal Grammar" portion of it -- is a special case which I don't want to discuss here, because it seems to be complicated by a special condition called the "poverty of the stimulus" (Chomsky 1980) according to which the core grammatical rules are not learnable by trial and error and corrective feedback (i.e., reinforcement) based on the sounds the child hears and produces during the few years in which it learns language. That database is simply not rich enough for <u>any</u> inductive mechanism to learn the requisite rules on the basis of the data and the time allotted to the child; hence the child must already have the rules built-in in advance.

But Chomsky's lesson to Skinner applies beyond syntax. Vocabulary learning – learning to call things by their names – already exceeds the scope of behaviorism, because naming is not mere rote association: things are not stimuli, they are <u>categories</u>. Naming things is naming <u>kinds</u> (such as birds and chairs), not just associating responses to unique, identically recurring individual stimuli, as in paired associate learning. To learn to name kinds you first need to learn to identify them, to categorize them (Harnad 1996; 2005). And kinds cannot be identified by just rote-associating names to stimuli. The stimuli need to be processed; the invariant features of the kind must be somehow extracted from the irrelevant variation, and they must be learned, so that future stimuli originating from things of the same kind can be recognized and identified as such, and not confused with stimuli originating from things of a different kind. (Even "individuals" are not "stimuli," but likewise kinds, detected through their sensorimotor invariants; there are sensorimotor "constancies" to be detected even for a sphere, which almost never casts the identical shadow onto our sensory surfaces twice.)

So Chomsky already pointed out that it is not enough to say that learning words is just "verbal behavior," shaped by reward history. It is that too, but that much explains nothing. The question is: "How are we able to learn words, as shaped by our reward history? What is the underlying functional capacity?" (Chomsky called this the "competence" underlying our "performance"). The answer in the case of syntax had been that we don't really "learn" it at all; we are born with the rules of Universal Grammar already in our heads. In contrast, the answer in the case of vocabulary and categories is that we <u>do</u> learn the rules, but the problem is still to explain <u>how</u> we learn them: What has to be going on inside our heads that enables us to successfully learn, based on the experience or training we get, to identify categories, to which we can then attach a name?

Introspection Won't Tell Us. A misapplication of Wittgenstein (1953) (or perhaps a mistake of Wittgenstein's) is to conclude that if we cannot introspect the rules for categorizing things (today we would say "if their representation is not 'explicit'") then those rules do not exist. A more valid inference is that cognitive science cannot be done by introspection. If we are to explain our cognitive capacities, we must somehow come up with explicit hunches about how we are able to do what we can do, and then we have to test whether those hunches actually work: whether they can really delivery the behavioral goods. Our minds will have to come up with those hypotheses, as in every other scientific field, but it is unlikely that cognition will wear them on its sleeve, so that we can just sit in our armchairs, do the cognizing in question, and simply introspect <u>how</u> it is that we are doing it. In this respect, cognition is impenetrable to introspection (in a sense that is related to, but not quite the same as the sense that Zenon Pylyshyn (1980; 1999) had in mind with his "cognitive impenetrability" criterion, but I am getting ahead of myself).

One of the first candidate armchair theories of cognition was mental imagery theory: When we introspect, most of us are aware of images going on in our heads. (There are words too, but we will return to those later.) The imagery theorists stressed that, for example, the way I remember who my third-grade school-teacher was is that I first picture her in my head, and then I name her, just as I would if I had seen her. Today, after 3 decades of having been enlightened on this score by Zenon Pylyshyn's celebrated "Mind's Eye" critique of mental imagery in 1973, it is hard even to imagine that anyone could ever have failed to see this answer -- that the way I remember her name is by picturing her, and then identifying the picture -- as anything but empty question-begging. How do I come up with her picture? How do I identify her picture? Those are the real functional questions we are missing; and it is no doubt because of the anosognosia – the "picture completion" effect that comes with all conscious cognition -- that we don't <u>notice</u> what we are missing: We are unaware of our cognitive blind spots – and we are mostly cognitively blind.

It is now history how Zenon opened our eyes and minds to these cognitive blind spots and to how they help non-explanations masquerade as explanations. First, he pointed out that the trouble with "picture in the mind" "just-so" stories is that they simply defer our explanatory debt: How did our brains find the right picture? And how did they identify whom it was a picture of? By reporting our introspection of what we are seeing and feeling while we are coming up with the right answer we may (or may not) be correctly reporting the decorative accompaniments or correlates of our cognitive functions -- but we are not explaining the functions themselves. Who found the picture? Who looked at it? Who recognized it? And <u>how</u>? I first asked how *I* do it, what is going on in my head; and the reply was just that a little man in my head (the homunculus) does it for me. But then what is going on in the little man's head?

Discharging the Homunculus. Imagery theory leaves a lot of explanatory debts to discharge, perhaps an infinite regress of them. Zenon suggested that the first thing we need to do is to discharge the homunculus. Stop answering the functional questions in terms of their decorative correlates, but explain the functions themselves. Originally, Zenon suggested that the genuine explanation has to be "propositional" (Pylyshyn 1973) but this soon evolved into "computational" (Pylyshyn 1984). If I ask you who your 3rd grade school-teacher was, your brain has to do a computation, a computation that is invisible and impenetrable to introspection. The computation is done by our

heads implicitly, but successful cognitive theory must make it explicit, so it can be tested (computationally) to see whether it works. The decorative phenomenology that accompanies the real work that is being done implicitly is simply misleading us, lulling us in our anosognosic delusion into thinking that we know what we are doing and how. In reality, we will only know how when the cognitive theorists figure it out and tell us.

There were other aspects to Zenon's insight too, not all of them quite correct, in my opinion. One was the implication that words and propositions were somehow more explanatory and free of homuncularity than images. But of course one could ask the same question about the origin and understanding of words in the head as of the origin and understanding of pictures in the head. Let's say that Zenon could have replied that words and propositions are nevertheless closer to computability than images, hence closer to an explicit, testable, computable functional explanation (Harnad 1982). Perhaps; but images can be "computed" too. Here Zenon would agree, but pointing out that a computation is a computation either way. He had famously argued that Shepard's mental rotation task (Shepard & Cooper 1982) could in principle be performed computationally using something like discrete Cartesian coordinates and formulas rather than anything like continuous analog rotation.

But at that point the debate became one about optimality (which of the two ways was the most general and economic way to do it?) and about actuality (which of the two ways does the brain really do it?) and not about possibility in principle, or necessity. It had to be admitted that the processes going on in the head that got the job done did not have to be computational after all; they could be dynamical too. They simply had to do the job.

Zenon, in rightly resisting the functional question-begging of imagery theorists in favor of goods-delivering computational explanation, went a bit too far, first denying that noncomputational structures and processes could occur and explain at all, and then, when that proved untenable, denying that, if they did, they were "cognitive." Rightly impressed by the power of computation and of the Church/Turing Thesis (Teuscher 2004) – that just about anything was computable, and hence computationally simulable to as close an approximation as one liked – Zenon relegated everything that was noncomputational to the "noncognitive." It occurred "below the level of the architecture of the virtual machine" that does the cognizing, implemented instead in "informationally encapsulated" sensorimotor modules that were "cognitively impenetrable" – that is, not modifiable by what we know and can state explicitly in propositions and computations (Pylyshyn 1984). The criterion for what was to count as cognitive was what could be modified by what we knew explicitly; what could not be modified in that way was "subcognitive," and the domain of another discipline.

(The similarity here to Skinner's dismissal of "how" questions as pertaining to physiology rather than psychology is ironic, but I don't think it is damning: Zenon, after all, was trying to make us face the problem of functional explanation, whereas Skinner was avoiding it. Moreover, both Pylyshyn and Skinner were right in insisting that the details of the physical (hardware) implementation of a function were independent of the functional level of explanation itself – except that Skinner had no functional explanation, whereas Pylyshyn had an all-powerful one: computation.)

Computational Hegemony? But Zenon was not the first to get a little too carried away with the power of computation. I think his attempt to formulate an impenetrable boundary between the cognitive and the noncognitive – in the form of his cognitive impenetrability criterion and the functional autonomy of the architecture of the virtual machine on which the computations were being performed – was not as successful as his rejection of imagery as non-explanatory, his insistence on functional explanation itself, and his promotion of computation's pride of place in the explanatory armamentarium. Imagery was indeed non-explanatory in and of itself. But an internal dynamical system that could actually generate some of our behavioral capacity (e.g., visual rotation judgments) certainly could not be denied out of hand; and if, as both brain imaging data (Kosslyn 1994) and considerations of functional optimality subsequently suggested, dynamical analog rotational processes in the brain really do occur, then there are certainly no grounds for denying them the status of being "cognitive."

A very similar point can be made about Zenon's celebrated paper with Jerry Fodor, which pointed out that neural nets were (1) uninteresting if they were just a hardware for implementing a symbol (computional) system, (2) irrelevant (like other dynamical systems) if they could be simulated computationally, and (3) subcognitive if they could be "trained" into becoming a symbol system (which then goes on to do the real work of cognition) (Fodor & Pylyshyn 1988).

So far, this sorting and labeling of functional "modules" – and I use the word "modules" here loosely, without any assumptions about the degree to which they are truly independent of one another functionally (Fodor 1988) – is more an exercise in semantics or taxonomy than anything substantive, functionally speaking. But there is one substantive issue on which I think Zenon has quietly conceded without ever quite announcing it, and that is symbol grounding – the ultimate question about the relation between the computational and the dynamical components of cognitive function (Harnad 1990).

Computation and Consciousness. But first, let us quickly get rid of another false start: Many, including Zenon, thought that the hardware/software distinction spelled hope not only for explaining cognition but for solving the mind/body problem: If the mind turns out to be computational, then not only do we explain how the minds works (once we figure out what computations it is doing and how) but we also explain that persistent problem we have always had (for which Descartes is not to blame) with understanding how the mental states can be physical states: It turns out they are not physical states! They are computational states. And, as everyone knows, the computational "level" – the software that a machine is running – is independent of the dynamical physical level – the hardware of the machine on which the software is running. Not independent in the sense that the software does not need some hardware or other to run on, but in the sense that the physical details of the hardware are irrelevant. The same software can be run on countless, radically different kinds of hardware, yet the computational states are the same, when the same program is running. (I never thought there was much more to Zenon's "virtual machine" level than the hardware/software distinction, with a bit more interpretation.)

Well this does not solve the mind/body problem, for many reasons, but here I will only point out that it does not solve the problem of the relation between computational and dynamical processes in cognition either: Computations need to be dynamically implemented in order to run and to do whatever they do, but that's not the only computational/dynamical relationship; and it's not the one we were looking for when we were asking about, for example, mental rotation.

Computation is rule-based symbol manipulation; the symbols are arbitrary in their shape (e.g., 0's and 1's) and the manipulation rules are syntactic, being based on the symbols' shapes, not their meanings. Yet a computation is only useful if it is semantically interpretable; indeed, as Fodor and Pylyshyn (1988) have been at pains to point out, systematic semantic interpretability, indeed compositional semantics, in which most of the symbols themselves are individually interpretable and can be combined and recombined coherently and interpretably, like the words in a natural language -- is the hallmark of a symbol system. But if symbols have meanings, yet their meanings are not in the symbol system itself, what is the connection between the symbols and what they mean?

Grounding the Language of Thought. Here it is useful to think of propositions again, Pylyshyn's original candidate, as the prototypes of Fodor's (1975) "language of thought" – computation in both instances. The words in propositions are symbols. What connects those symbols to their referents? What gives them meaning? In the case of a sentence in a book, such as "the cat is on the mat," there is no problem, because it is the mind of the writer or reader of the sentence that makes the connection between the word "cat" and the things in the world we happen to call "cats," and between the proposition "the cat is on the mat" and the circumstance in the world we happen to call "cats being on mats." Let us call that mediated symbol-grounding: The link between the symbol and its referent is made by the brain of the user. That's fine for logic, mathematics and computer science, which merely use symbol systems. But it won't do for cognitive science, which must also explain what is going on in the head of the user; it doesn't work for the same reason that homuncular explanations do not work in cognitive explanation, leading instead to an endless homuncular regress. The buck must stop somewhere, and the homunculus must be discharged, replaced by a mindless, fully autonomous process.

Well, in Pylyshyn's computationalism, the only candidate autonomous internal function for discharging the homunculus is computation, and now we are asking whether that function is enough. Can cognition be just computation? The philosopher John Searle (1980) asked this question in his celebrated thought experiment. Let us agree (with Turing 1950) that "cognition is as cognition does" – or better, so we have a Chomskian competence criterion rather than a mere behaviorist performance criterion – that "cognition is as cognition can do." The gist of the Turing Test is that on the day we will have been able to put together a system that can do everything a human being can do, indistinguishably from the way a human being does it, we will have come up with at least one viable explanation of cognition.

Turing Sets the Agenda. Searle took Turing – as well as Pylyshyn – at their word. He said: Suppose we do come up with a computer program that can pass the Turing Test (TT). Will we really have explained cognition? In particular, will the system that passes the TT really cognize, will it really have a mind? The classical TT is conducted

by email (so you are not prejudiced by the way the candidate looks, which is irrelevant); it is basically a test – life-long, if need be – of whether the system has the full performance capacity of a real pen-pal, so much so that we would not be able to tell it apart from a real human pen pal. If it passes the test, then it really cognizes; in particular, it really understands all the emails you have been sending it across the years, and the ones it has been sending you in reply (Harnad, in press).

And of course, being implementation-independent, the winning software can be run on any hardware. If the TT-passing computational states are really the cognitive states, they will be the right cognitive states in every implementation. So Searle simply proposes to conduct the TT in Chinese (which he doesn't understand) and he proposes that he himself should become the implementing hardware himself, by memorizing all the symbol manipulation rules and executing them himself, on all the email inputs, generating all the email outputs. Searle's very simple point is that he could do this all without understanding a single word of Chinese. And since Searle himself is the entire computational system, there is no place else the understanding could be. So it's not there. The meanings are all just in the heads of the users again – the real Chinese pen-pals doing the T-Testing. So the TT-passing program is no more cognitive than any other symbol system in logic, mathematics or computer science. It is just a bunch of symbols that are systematically interpretable by us -- by users with minds. It has again begged the question of how the mind actually does what it does – or rather, it has failed to answer it.

Newton Still Available. So what is still missing, then, if computation alone can always be shown to be noncognitive and hence insufficient, by arguments analogous to Searle's? Searle thought the culprit was not only the insufficiency of computation, but the insufficiency of the Turing Test itself; he thought the only way out was to abandon both and turn instead to studying the dynamics of the brain. I think Searle went too far (Harnad 2001). There is still scope for a full functional explanation of cognition, just not a purely computational one. As we have seen, there are other candidate autonomous, non-homuncular functions in addition to computation, namely, dynamical functions such as internal analogs of spatial or other sensorimotor dynamics: not propositions describing them nor computations simulating them, but the dynamic processes themselves, as in internal analog rotation; perhaps also real parallel distributed neural nets rather than just symbolic simulations of them.

The root of the problem is the symbol-grounding problem: How can the symbols in a symbol system be connected to the things in the world that they are ever-so-systematically interpretable as being about: connected directly and autonomously, without begging the question by having the connection mediated by that very human mind whose capacities and functioning we are trying to explain! For ungrounded symbol systems are just as open to homuncularity, infinite regress and question-begging as subjective mental imagery is!

The only way to do this, in my view, is if cognitive science hunkers down and sets its mind and methods on scaling up to the Turing Test, for all of our behavioral capacities. Not just the email version of the TT, based on computation alone, which has been shown to be insufficient by Searle, but the full robotic version of the TT, in which the symbolic capacities are grounded in sensorimotor capacities and the robot itself (Pylyshyn 1987) can mediate the connection, directly and autonomously,

between its internal symbols and the external things its symbols are interpretable as being about, without the need for mediation by the minds of external interpreters.

We cannot prejudge what proportion of the TT-passing robot's internal structures and processes will be computational and what proportion dynamic. We can just be sure that they cannot <u>all</u> be computational, all the way down. As to which components of its internal structures and process we will choose to call "cognitive": Does it really matter? And can't we wait till we get there to decide?[1]

REFERENCES

Catania, A.C. & Harnad, S. (eds) (1988) The Selection of Behavior. The Operant Behaviorism of BF Skinner: Comments and Consequences. New York: Cambridge University Press.

Chomsky, Noam (1959) A Review of B. F. Skinner's Verbal Behavior. *Language* 35: 26-58 http://cogprints.org/1148/

Chomsky, N. (1980) Rules and representations. *Behavioral and Brain Sciences* 3: 1-61.

Fodor, J. A. (1975) The language of thought New York: Thomas Y. Crowell

Fodor, J. A. (1985) Precis of "The Modularity of Mind." *Behavioral and Brain Sciences* 8: 1 - 42.

Fodor, J. A. & Pylyshyn, Z. W. (1988) Connectionism and cognitive architecture: A critical appraisal. Cognition 28: 3 - 71.

Harnad, S. (1982) Neoconstructivism: A Unifying Constraint for the Cognitive Sciences, In: Language, mind and brain (T. Simon & R. Scholes, eds., Hillsdale NJ: Erlbaum), 1 - 11. http://cogprints.org/0662

Harnad, S. (1985) Hebb, D.O. - Father of Cognitive Psychobiology 1904-1985 *Behavioral and Brain Sciences* 8: 765. (Obituary) http://cogprints.org/1652/

Harnad, S. (1990) The Symbol Grounding Problem Physica D 42: 335-346. http://cogprints.org/0615/

Harnad, Stevan (1996) Experimental Analysis of Naming Behavior Cannot Explain Naming Capacity. *Journal of the Experimental Analysis of Behavior* 65: 262-264. http://cogprints.org/1605/

---

[1] One could ask whether grounded cognition ("sticky" cognition, in which symbols are connected to their referents, possibly along the lines of Pylyshyn's (1994) indexes or FINSTs) would still be computation at all: After all, the hallmark of classical computation (and of language itself) is that symbols are arbitrary, and that computation is purely syntactic, with the symbols being related on the basis of their own arbitrary shapes, not the shapes of their referents.

Harnad, S. (2001) What's Wrong and Right About Searle's Chinese Room Argument? In: M. Bishop & J. Preston (eds.) Essays on Searle's Chinese Room Argument. Oxford University Press. http://cogprints.org/1622/

Harnad, S. (2005) To Cognize is to Categorize: Cognition is categorization, in Lefebvre, C. and Cohen, H., Eds. *Handbook of Categorization*. Elsevier. http://eprints.ecs.soton.ac.uk/11725/

Harnad, S. (in press) The Annotation Game: On Turing (1950) on Computing, Machinery and Intelligence. In: Epstein, Robert & Peters, Grace (Eds.) *The Turing Test Sourcebook: Philosophical and Methodological Issues in the Quest for the Thinking Computer*. Kluwer http://cogprints.org/3322/

Kosslyn, S. M., Pinker, S., Smith, G. & Shwartz, S. P. (1979) On the demystification of mental imagery. Behavioral and Brain Sciences 2: 535 - 548.

Kosslyn, S.M. (1996) Image and Brain: The Resolution of the Imagery Debate MIT Press

Pylyshyn, Z. W. (1973) What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin* 80: 1-24.

Pylyshyn, Z. W. (1980) Computation and cognition: Issues in the foundations of cognitive science. *Behavioral and Brain Sciences* 3: 111-169.

Pylyshyn, Z. W. (1984) *Computation and cognition*. Cambridge MA: MIT/Bradford

Pylyshyn, Z. W. (Ed.) (1987) *The robot's dilemma: The frame problem in artificial intelligence*. Norwood NJ: Ablex

Pylyshyn, Z.W. (1994). Some primitive mechanisms of spatial attention. Cognition, 50, 363-384.

Pylyshyn, Z. W. (1999) Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences* 22: 341-364. http://www.bbsonline.org/documents/a/00/00/04/94/index.html

Pylyshyn, Z.W. (2002) Mental imagery: In search of a theory. *Behavioral and Brain Sciences 25*: 157-182 http://www.bbsonline.org/documents/a/00/00/19/46/index.html

Searle, John. R. (1980) Minds, brains, and programs. Behavioral and Brain Sciences 3: 417-457 http://www.cogsci.soton.ac.uk/bbs/Archive/bbs.searle2.html

Teuscher, Christof (2004) *Alan Turing: Life and Legacy of a Great Thinker*. Springer

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

Wittgenstein, L. (1953) *Philosophical investigations*. New York: Macmillan