

The Cogsci Moment

Concepts of Information i218
Geoff Nunberg

March 31, 2011



Agenda

The emergence of cognitive science

On "formal symbol manipulation"

Minds as information processing devices: the Turing Test

Searle's response

Chomsky's review of Verbal Behavior



Defining Cognitive Science: Gardner's 5 criteria

Fundamental:

Concern with mental representations, "a level of analysis wholly separate from the biological or neurological... and the sociological or cultural" (accept supervenience but not reductionism)

Computer is central to the understanding of mind

Methodological

De-emphasis of affect, social context, culture [and the M-word – GN]

Interdisciplinarity (e.g., Sloan initiative)

Philosophical roots of agenda: "it is virtually unthinkable that cognitive science would exist... had there not been a philosophical tradition dating back to the time of the Greeks.

[Add: formalization of theories (GN)]



The historical background

Behaviorism and the rejection of "introspectionism" and mentalist notions (mind, desires, intentions, etc.) for observables.

Problems accounting for plans, complex behaviors, speaking... (playing an arpeggio, hammering a nail, speaking) in terms of reflexes, stimulus-response, etc.



The conceptual background

Turing: programs as models of cognitive processes, "thinking machines."

Logical models of neuronal processes

Information theory: info. as 1st class citizen; formal equivalence of signaling mechanisms



Minds as Information Processing Devices

Intelligent systems exhibit their intelligence by achieving goals in the face of different and changing environments. Intelligent behavior is adaptive, hence must take on strikingly different forms when the environments are correspondingly different. [hence "artificial"]

Intelligence is mind implemented by any patternable kind of matter.

Task = finding invariants in adaptive systems; the "inner environment"

Herbert A. Simon, "Cognitive Science: The Newest Science of the Artificial," 1980



Minds as Information Processing Devices



Thought as FSM. Device manipulates symbols purely on basis of shape; computers as “numeral crunchers.” (Watson doesn't know anything about Toronto, just "toronto.")

Thinking is carried out by the brain using the same basic symbol-manipulating processes that are used by computers: reading symbols writing them, storing them in memory, copying them, comparing them for equality or inequality.... The availability of these processes provides the necessary and sufficient conditions for a system to exhibit intelligence. Simon





AI: Some Early Achievements

1956: Logical Theorist (Newell & Simon) proves theorems from Principia Mathematica using list processing, procedures intended to simulate human thought (working backwards, analogy, etc.)

1964: Bobrow's STUDENT solves word problems.

The gas consumption of my car is 15 miles per gallon. The distance between Boston and New York is 250 miles...

1970: Winograd's SHRDLU "understands" relations among elements in a block world.



AI: Some Early Achievements

1970: Winograd's SHRDLU "understands" relations among elements in a block world.

Person: Find a block which is taller than the one you are holding and put it into the box.

Computer: By "it", I assume you mean the block which is taller than the one I am holding.

Computer: OK.

Person: What does the box contain?

Computer: The blue pyramid and the blue block.

Person: What is the pyramid supported by?


Computer: The box.

Person: How many blocks are not in the box?

Computer: Four of them.

Person: Is at least one of them narrower than the one **Person** told you to pick up?

Computer: Yes, the red cube.





Minds as Information Processing Devices

It might have been necessary a decade ago to argue for the commonality of the information processes that are employed by such disparate systems as computers and human nervous systems. The evidence for that commonality is now overwhelming. Wherever the boundary is drawn, there exists today a science of intelligent systems that extends beyond the limits of any single species. Simon, 1980

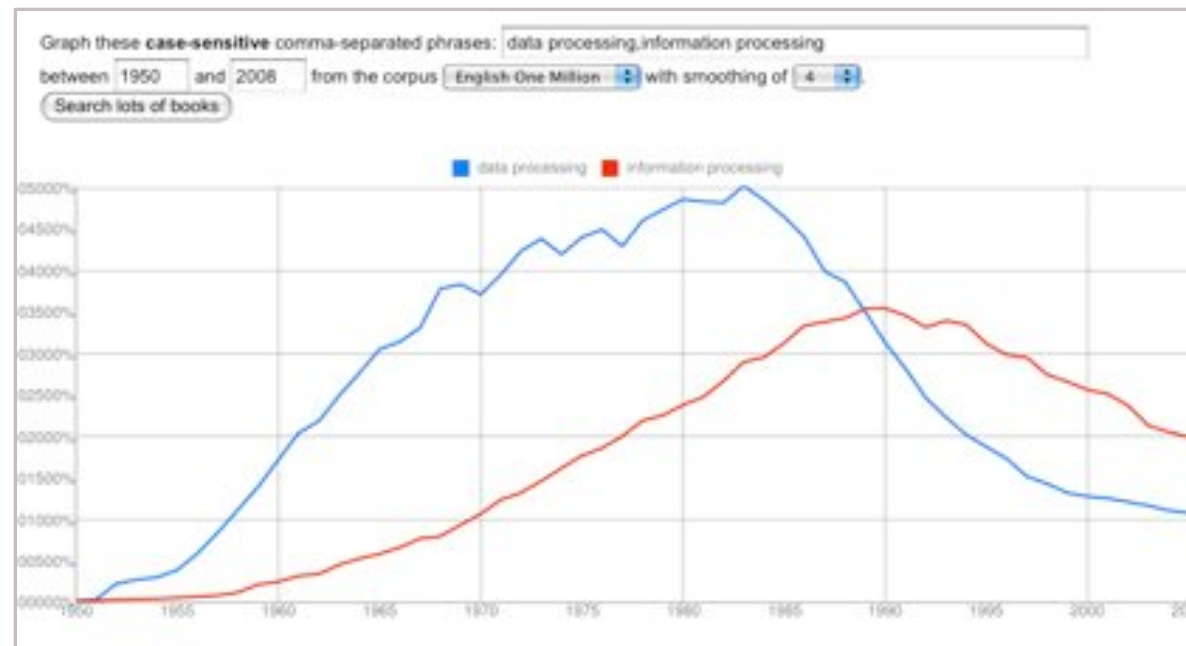


Aside: The ambiguity of "processing"



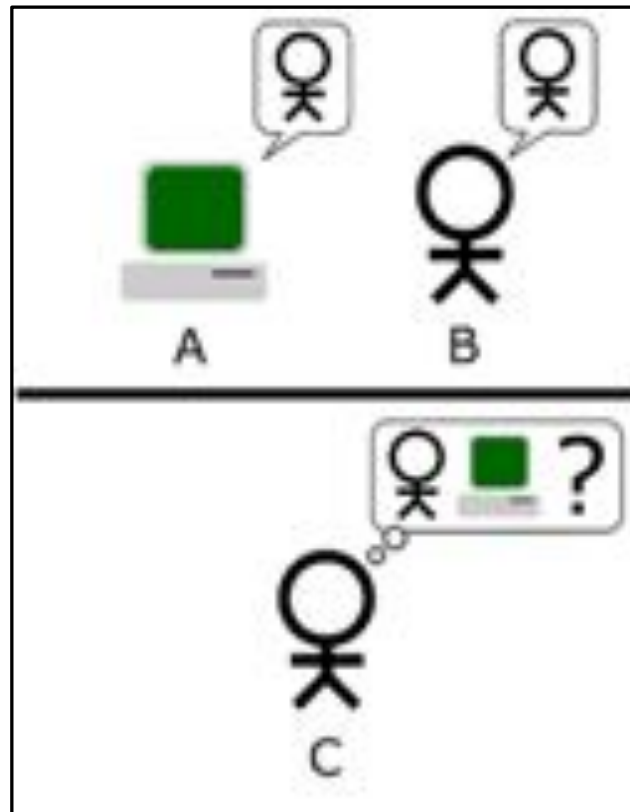
Is the object of "process" the input or the output?

OED: To register or interpret (information, data, etc.);
Computing to operate on (data) by means of a program
Contrast *processed food* vs *food processor*





The Turing Test





The Turing Test



I propose to consider the question, "Can machines think?" This should begin with definitions of the meaning of the terms "machine" and "think." The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous... Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words. Alan Turing, "Computing Machinery and Intelligence," 1950



The Turing Test



I believe that in about fifty years' time it will be possible to programme computers, with a storage capacity of about 10^9 , [=10 gigabytes] to make them play the imitation game so well that an average interrogator will not have more than 70 percent chance of making the right identification after five minutes of questioning.at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted.



Responses to Turing



Too easy: cf Joseph Weizenbaum's Eliza:

Men are all alike

IN WHAT WAY?

They're always bugging us about something or other.

CAN YOU THINK OF A SPECIFIC EXAMPLE?

Well, my boyfriend made me come here.

YOUR BOYFRIEND MADE YOU COME HERE

He says I'm depressed much of the time

I AM SORRY TO HEAR YOU ARE DEPRESSED

It's true. I am unhappy

DO YOU THINK COMING HERE WILL HELP YOU NOT TO BE UNHAPPY

I need some help, that much seems certain.

WHAT WOULD IT MEAN TO YOU IF YOU GOT SOME HELP

Perhaps I could learn to get along with my mother.

TELL ME MORE ABOUT YOUR FAMILY

My mother takes care of me.



Responses to Turing

Robert French: computers can't duplicate subcognition

flugblog as name of a breakfast cereal, a computer company, etc.

"Rate dry leaves as a hiding place"/"Rate these jokes for funniness..."

?= Turing's "argument from informality of behaviour"?



Responses to Turing

What Computers can't do...

Emotions/consciousness:


G. Jefferson, Lister Oration.: "No machine could feel (and not merely artificially signal, an easy contrivance) pleasure at its successes, grief when its valves fuse, be warmed by flattery, be made miserable by its mistakes.... be angry or depressed when it cannot get what it wants."

But if the machine can seem to have emotions...: "an easy contrivance"? How does one tell they're not real?

Turing: "This argument appears to be a denial of the validity of our test. According to the most extreme form of this view the only way by which one could be sure that machine thinks is to be the machine and to feel oneself thinking."



What would it take to persuade you?



Interrogator: In the first line of your sonnet which reads 'Shall I compare thee to a summer's day', would not 'a spring day' do as well or better?

Computer: It wouldn't scan.

Interrogator: How about 'a winter's day'? That would scan all right.

Computer: Yes, but nobody wants to be compared to a winter's day.

Interrogator: Would you say Mr. Pickwick reminded you of Christmas?

Computer: In a way.

Interrogator: Yet Christmas is a winter's day, and I do not think Mr Pickwick would mind the comparison

Computer: I don't think you're serious. By a winter's day one means a typical winter's day, rather than a special one like Christmas.



Responses to Turing



Intentionality: Searle's Chinese Room

Suppose that I'm locked in a room and given a large batch of Chinese writing. Suppose that ...to me, Chinese writing is just so many meaningless squiggles. Now suppose I am given a second batch of Chinese script together with a set of rules for correlating the second batch with the first batch.





Responses to Turing

Intentionality: Searle's Chinese Room

...Unknown to me, the people who are giving me all of these symbols call the first batch "a script," they call the second batch a "story. ' and they call the third batch "questions." Furthermore, they call the symbols I give them back in response to the third batch "answers to the questions." and the set of rules in English that they gave me, they call "the program."



On "formal symbol manipulation"

Suppose also that after a while I get so good at following the instructions and the programmers get so good at writing the programs that from the external point of view... my answers to the questions are absolutely indistinguishable from those of native Chinese speakers. Nobody just looking at my answers can tell that I don't speak a word of Chinese.





On "formal symbol manipulation"

I have inputs and outputs that are indistinguishable from those of the native Chinese speaker, and I can have any formal program you like, but I still understand nothing...

...formal symbol manipulations by themselves don't have any intentionality; they are quite meaningless; they aren't even *symbol* manipulations, since the symbols don't symbolize anything. In the linguistic jargon, they have only a syntax but no semantics





Could there be a Chinese room?

Cf difficulties of machine translation

Syntax + vocabulary

Mettons marie dans la chaise/le fauteuil.

s'asseoir sur le fauteuil/s'installer dans le fauteuil

Cloche/sonnette/timbre

Fenetre/vitrine

Etiquette/sticker/label/tag

The man looked at the girl with the telescope
(blue vs. penetrating eyes).

Pronouns

The police banned the demonstration of the feminists. They were afraid of violence.

Since the dictionary is constructed on the basis of the text that is being processed, *it/he* need refer to only a small amount of context to resolve ambiguities.



Mind as information processing: "The Magical Number 7"



"The Magical Number 7"

"I have been persecuted by an integer. For seven years this number has followed me around, has intruded in my most private data, and has assaulted me from the pages of our most public journals."

"... I shall begin my case history by telling you about some experiments that tested how accurately people can assign numbers to the magnitudes of various aspects of a stimulus. In the traditional language of psychology these would be called experiments in absolute judgment. Historical accident, however, has decreed that they should have another name. We now call them experiments on the capacity of people to transmit information."




"The Magical Number 7"

If it is a good communication system, however, there must be some systematic relation between what goes in and what comes out. ...If we measure this correlation, then we can say how much of the output variance is attributable to the input and how much is due to random fluctuations or "noise" introduced by the system during transmission. So we see that the measure of transmitted information is simply a measure of input-output correlation.



"The Magical Number 7"


The situation can be described graphically by two partially overlapping circles. Then the left circle can be taken to represent the variance of the input, the right circle the variance of the output, and the overlap the covariance of input and output. I shall speak of the left circle as the amount of input information, the right circle as the amount of output information, and the overlap as the amount of transmitted information.





"The Magical Number 7"

If the human observer is a reasonable kind of communication system, then when we increase the amount of input information the transmitted information will increase at first and will eventually level off at some asymptotic value. This asymptotic value we take to be the *channel capacity of the observer*:





"The Magical Number 7"

~2.6 (1.9-3.9) bits seems to be channel capacity for absolute judgments of pitch, loudness, taste, size of square, hue, brightness...

Number of items increases as dimensions of stimulus increase (faces, words, etc.)

Short term memory shows analogous limits, but wrt chunks, not bits: recoding features into phonemes into words; dots & dashes into letters, words; binary digits into base-4 digits...

"the concepts and measures provided by the theory of information provide a quantitative way of getting at some of these questions."



The fortunes of MTC in psychology

2 features of Miller's approach:

mind as info processing device; focus on capacities not context

mathematical theory of info quantity

Although Shannon's information theory is alive and well in a number of fields, after an initial fad in psychology during the 1950s and 1960s it no longer is much of a factor, beyond the word *bit*, in psychological theory...



The elements of choice in information theory are absolutely neutral ... a pure, unstructured set whose elements are functionally interchangeable.... the stimuli of psychological experiments are to some degree structured, and so, in a fundamental way, they are not in any sense interchangeable. [complexities of multidimensionality, recoding, etc.]

R. Duncan Luce, "Whatever Happened to Information Theory in Psychology?" *Rev. Gen. Psych.*, 2003



Are Mental Representations Propositional?

Cf work by Roger Shepard et al. on mental rotation







The role of mental representations: Chomsky on Skinner



Chomsky's Review of *Verbal Behavior*



Verbal Behavior: 1957 book by B. F. Skinner outlining view of linguistic behavior in terms of "operant conditioning" ("radical behaviorism").

Much of the time, a man acts only indirectly upon the environment from which the ultimate consequences of his behavior emerge. Instead of going to a drinking fountain, a thirsty man may simply "ask for a glass of water"—that is, may engage in behavior which produces a certain pattern of sounds which in turn induces someone to bring him a glass of water....The glass of water reaches the speaker only as the result of a complex series of events including the behavior of a listener.

Linguistic acts are "operants" of various types ("mands," "tacts," "echoic," etc.) whose strength and frequency (and probability?) are determined by stimuli and reinforcements.



Chomsky's Review of *Verbal Behavior*



Chomsky's review of *Verbal Behavior* published in *Language* in 1959

What is so surprising is the particular limitations [Skinner] has imposed on the way in which the observables of behavior are to be studied.... One would naturally expect that prediction of the behavior of a complex organism (or machine) would require, in addition to information about external stimulation, *knowledge of the internal structure of the organism, the ways in which it processes input information and organizes its own behavior.* These characteristics of the organism are in general a complicated product of inborn structure, the genetically determined course of maturation, and past experience..



Chomsky's Review of *Verbal Behavior*

A typical example of stimulus control for Skinner would be the response to a piece of music with the utterance Mozart or to a painting with the response Dutch. These responses are asserted to be "under the control of extremely subtle properties" of the physical object or event. Suppose instead of saying Dutch we had said *Clashes with the wallpaper, I thought you liked abstract work, Never saw it before, Tilted, Hanging too low, Beautiful, Hideous, Remember our camping trip last summer?...* Skinner could only say that each of these responses is under the control of some other stimulus property of the physical object. ...This device is as simple as it is empty.... the word *stimulus* has lost all objectivity in this usage



Chomsky's Review of *Verbal Behavior*

The phrase "X is reinforced by Y ...is being used as a cover term for "X wants Y," "X likes Y," "X wishes that Y were the case," etc. Invoking the term reinforcement has no explanatory force, and any idea that this paraphrase introduces any new clarity or objectivity into the description of wishing, liking, etc., is a serious delusion.



Chomsky's Review of *Verbal Behavior*

Skinner defines the process of confirming an assertion in science as one of “generating additional variables to increase its probability,” and more generally, its strength. If we take this suggestion quite literally, the degree of confirmation of a scientific assertion can be measured as a simple function of the loudness, pitch, and frequency with which it is proclaimed, and a general procedure for increasing its degree of confirmation would be, for instance, to train machine guns on large crowds of people who have been instructed to shout it.



Chomsky's Triumph(?)



“Chomsky's review has come to be regarded as one of the foundational documents of the discipline of cognitive psychology, and even after the passage of twenty-five years it is considered the most important refutation of behaviorism.”
Frederick Newmeyer, 1986

“...Noam at his best, mercilessly out for the kill, daring, brilliant, on the side of the angels ... in the same category as St. George slaying the dragon.” Jerome Bruner, 1983



Chomsky's Triumph(?)



"From the perspective of most behavior analysts, Chomsky's review of Skinner's *Verbal Behavior* was a kind of ill-conceived dam in the progress of science, a rhetorically effective but conceptually flawed document that would eventually be overborne." David C Palmer, 2006 "A Half Century of Misunderstanding"

... and in the end, mentalism plays a limited role in the triumph of "Chomskianism" in linguistics.