AI Magazine Volume 13 Number 3 (1992) (© AAAI)

Letters

■ Editor:

As a communication scholar, I am well aware that many traditionalists view the respective disciplines of communication and computer science as unrelated. My recently completed master's thesis argues against this view. Many concepts from the field of communication have been used by artificial intelligence researchers and scholars in the development of AI. The central argument of my perspective is that artificial intelligence is the communication link between humans and multimedia.

I view intelligence as a necessary element for communication to occur, and artificial intelligence has the potential to provide the current multimedia with that intelligence. AI is concerned with, among other issues, designing systems that possess the humanly perceived tools associated with intelligent human behavior. These tools are the various methods that the human mind uses to process information. We could divide these tools into four very general categories: text and symbols, full-motion video, animation and graphics, and language and sound. The first category includes the printed word and any symbols (icons) that the human mind recognizes as meaningful information. The second category includes the short clips of real-life visual memories that the human mind stores over time. These are the short movies that the human mind records from the outside world and then stores for future use. The third category includes all the cognitive maps, images, graphics, pictures, etc. that the human mind creates as an aid to understanding concepts, or just for entertainment. They are always "in-house" creations and not full-motion video recordings from the outside world. The fourth category includes audio recordings from the outside world, as well as internal applications for processing language and speech information.

The multimedia revolution is being led by communication scholars, non-AI computer scientists, art majors, and even business consultants. This latest computer revolution has taken shape only within the past five years. The AI revolution has been going on for nearly four decades. These two revolutions have been operating independently with limited success, instead of together with potentially phenomenal success. The multimedia revolution has successfully broken into the marketplace on all levels, but lacks the key component (symbolic reasoning) needed for the next evolutionary stage. Meanwhile, the AI revolution has the potential for producing that key component, but lacks the marketplace success of the multimedia revolution. By transcending traditional academic boundaries, all parties involved will gain fresh perspectives, as well as valuable knowledge. In addition, each party may conclude that the other guys aren't so strange after all.

Matthew R. Cronin, M.S.
Communication Theory and Research
Florida State University

■ Editor:

I was pleased to see that you described the "imitation game" correctly (*AI Magazine* 13(2): 92), i.e., the game is to determine which of the responses is from the woman and which from the computer imitating the woman. I used the imitation game when I taught AI to non-techies at San Jose State University, San Jose, California.

A male and female student are sent out to the corridor with the instruction that the male is to fake being the female. The class composes questions that are sent out by courier and the responses are identified only as A and B. The class continues to ask questions until the students feel that they know which of A and B is the woman. We then play the game again with new participants, this time with the woman faking being the man.

Many interesting AI and sociological aspects soon reveal themselves.

First of all, the class realizes that only the women students are capable of phrasing questions that detect when a man is imitating a woman. Similarly, only men can tell when a woman is faking being a man. Secondly, the class soon realizes that sexual differences are much reduced from Turing's time—women know about sports and men often can answer cooking questions. Women may not be able to answer questions about sewing or cosmetics and many men no longer know much about car engines.

The unisex revolution makes the game particularly difficult to play, and the class can find it quite frustrating to phrase the right questions, but occasionally a bright student hits upon a devastating technique: Ask a false question. For example, the question "What is a Lipetz head screw driver?" asked when a woman is trying to imitate a man will often result in the answers:

A: It is a screw driver used for Lipetz-head screws.

B: I never heard of it. Do you mean a Phillips head screw driver?

Once this false question approach is discovered, few students can successfully fool the class. I expect that imitation programs would do even worse.

Oscar Firschein
Software and Intelligent Systems Office, DARPA

■ Editor:

Robert Epstein's article on the Turing Competition (*AI Magazine* 13(2): 81:95) provided a one-sided discussion of the Turing Test, treating it as an important issue of current AI thought. However, this is clearly contrary to the general research thrust of AI (check any AAAI conference proceedings) and rightly so. The Turing Test is an inappropriate criterion for artificial intelligence mainly because it is purely a behavioral test. Since the first description of the test was published during the 1950's, the heyday of psychological behaviorism, this is not surprising. But the Turing Test faces the same problems that have since been noted about psychological behaviorism. Behaviorism is simply insufficient to express all that is going on with intelligence. For example, Putnam describes a super-spartan who experiences great pain but does not display any manifestation of pain. Perhaps the more critical problem with behaviorism is that it is unable to handle complex behavior since it does not utilize mental states of the agent. Even seemingly simple examples like eating a hamburger cannot be adequately described using stimulus-response pairs since there are too many possible causes for eating a hamburger. Perhaps when I'm hungry I eat hamburgers, but sometimes I eat tacos; and I might also eat a hamburger when I'm not hungry but feeling blue; etc.

In the sense that the Turing Test is a purely behavioral test, it is not sufficient. A robot that passes the Turing Test may be quite unsatisfactory; for example, it could have randomly selected the correct responses. The unsatisfactory nature of the Turing Test is currently being seen in today's computer chess programs and even in the "whimsical" winner of this year's Loebner prize. The programs with the best behavior are not the ones that necessarily are the most intelligent but those with the best hardware and VLSI implementations or most memory.

The crucial point is that in addition to behavioral requirements, normative requirements must be considered. We cannot be solely interested in results, but how the results were achieved. How should the mind work? Simon made the same point when he distinguished between two types of rationality in economic systems: *procedural* (where the agent undertakes some kind of deliberation to select an action) and *substantive* (where the agent somehow selects the rational action without deliberation). Procedural rationality depends on the process that generates the behavior while substantive rationality does not; it is only dependent on the characteristics of the environment in which it takes place. Simon argues that procedural rationality is preferable because it provides a theory of rational choice that has predictive and explanatory power. Possessing a theory, the agent can apply it to new situations and learn by refining its theory. Providing this normative theory is the new test and challenge of AI.

Gary Ogasawara
Computer Science Division
UC Berkeley

Gary Ogasawara's letter raises an interesting point with respect to the validity of the Turing Test. But it overlooks a couple of important reasons why the Turing Test may still be the best test we can devise to answer the question, "Can machines think?" The first is that it would be very dangerous to rely on a test that makes specific assumptions about the mechanisms that must be used to achieve intelligence. Ten years ago it might not have felt risky to do that since the received position in AI was that intelligence is the result of symbol manipulation. But recent advances in subsymbolic (e.g., connectionist) systems, in reactive systems, and in hybrid approaches make it clear that we must be open minded as we search for computational paradigms that can provide some or all of the basis of intelligent behavior. The second issue we need to consider is the sheer size and complexity of the job of passing the Turing Test. It is true that simpler tasks can be achieved with simplistic mechanisms, such as the ones that were used by the winning program in this year's competition. But that program did not actually pass the test. And there is no evidence that suggests that a program with that architecture ever can. In fact, every computational device that anyone has ever come up with has so far fallen short of passing the test. If passing the test just required fast hardware, that would not be true. While there may be more subtle tests that will eventually be able to be used to measure degrees of machine intelligence, the Turing Test still serves to define a point that will surely be a key landmark in the development of systems that can help us think.
—Elaine Rich, Coeditor.

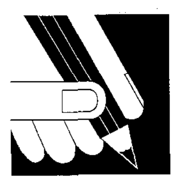

AI Magazine welcomes letters to the editor. If you want to respond to something you read in this issue, please email your letter to aimagazine@aaai.org, or mail it to the Managing Editor at the AAAI office. Please send your letter quickly so that it might be included in the issue now being prepared for publication.

Correction

The telephone number for *Presence: Teleoperators and Virtual Environments* (page 108, Summer magazine) was incorrect. The correct telephone number is (617) 253-2889.