At Issue:

Will researchers develop a computer that duplicates human intelligence in the foreseeable future?

Hans Moravec
Professor of computer science, Carnegie Mellon University

The field of artificial intelligence began with efforts to capture the surface layers of conscious human thought. It produced programs that play chess, prove theorems or solve other individual, narrow intellectual problems, but this "top-down" route has yet to produce anything resembling overall human intelligence. Meanwhile, the related field of robotics produced machines with behavior something like that of simple animals. The "bottom-up" route seems to be following a path roughly paralleling the evolution of natural intelligence, but about 10 million times as fast. Today, the smartest robots have control systems comparable to the nervous systems of the earliest vertebrates, like tiny fish. Early computers seemed powerful only because they were applied to tasks, like arithmetic, that humans do enormously inefficiently. On the other hand, our evolving ancestors lived or died by their physical and social interactions, and we inherit powerfully efficient equipment for those functions. Computers are far from those strengths, but are gaining so fast they should match us, even there, in less than 50 years.

Comparing modest assemblies of neurons, like the vertebrate retina, or control ganglia, in insects and other invertebrates with efficient computer programs that provide approximately the same functions for robots suggests that the work of a thousand neurons can be matched by a well-written program running at 1 million instructions per second (1 MIPS). High-end personal computers today afford a merely insectlike 500 MIPS, but the power is doubling each year. At that pace, home computers will be powerful enough to host humanlike intelligence in just a few decades.

Robot perception and navigation have advanced this decade to where research robots are now cruising hallways and roads autonomously for hours at a time. I think we are at the threshold of an evolution of mass-marketed robots, that, in coming decades, may go something as follows: Specialized autonomous utility robots will soon be among us, cleaning, transporting and eventually performing other manual tasks. They will be followed by a first generation of universal robots, with mental power comparable to a small lizard, able to host programs for many different physical applications. A second generation, with a mind like a small mouse, will adapt to specific situations by conditioned learning. A third, somewhat like a monkey, will model the physical and social world, so that it can mentally rehearse its tasks, invent new variations, observe and imitate others and explain its actions in physical and social terms.

A fourth, humanlike generation will add abstract reasoning and generalization to the repertoire. Then the real fun begins.

Stuart Shieber
Professor of computer science, Harvard University

In 1950, the British mathematician Alan Turing devised what has come to be known as the Turing Test for machine intelligence, based on a judge's ability or inability to distinguish between a computer and a person in dialog with the two. The Turing Test is popularly thought of as the Holy Grail of research in artificial intelligence. A computer able to pass the Turing Test is in essence what many think of when they think of "duplicating human intelligence with a computer."

I have good news and bad news. The bad news is that like the Holy Grail of Arthurian legend, the Turing Test will, I expect, remain beyond our grasp for quite a while, undoubtedly not achievable during our lifetimes. Philosophers like John Searle and Hubert Dreyfus [of the University of California at Berkeley] think that a machine passing the test is impossible, although I am agnostic on this point. Various AI luminaries and media pundits have been predicting the imminent arrival of intelligent machines, from "just around the corner" to "just a few decades" for many years now, at least since the 1950s. Like prognostications of the end of the world, as the predicted dates pass, the predictions become necessarily more inaccurate. The problem of how intelligence works and how it can be duplicated artificially is tremendously difficult. Denying this fact is simple hubris. The issue is not one of insufficient computer power; even if we had computers faster by orders of magnitude (which we undoubtedly will), we would need to know how to make use of the resources to duplicate human intelligence or allow computers to learn it on their own.

The good news is that like the Holy Grail, the goal itself is less important than the quest. Indeed, many if not most AI researchers view the Turing Test as an exceedingly poor goal for current research in the field. The study of AI, by engaging some of the brightest minds in computer science on arguably the hardest problems in the field, can claim credit for time-sharing computers, windowed interfaces, computer dictation, medical diagnostic systems, financial-industry mechanization and Deep Blue. The technologies on which these systems were based were not developed by researchers trying directly to build artificially intelligent Turing-test passers, but through myriad attacks on varied problems in understanding particular types of knowledge, reasoning, learning and intelligent behavior.

It is important not to gauge progress in AI on progress in passing the Turing Test. One can, and we do, have tremendous progress in the former, both in theory and in practice, without approaching the Grail itself.