



ARTIFICIAL_INTELLIGENCE September 21, 2010, 12:02AM EST

AI's Time Has Arrived

After decades of failure, artificial intelligence is ready for the masses, owing to tectonic advances in computing power

By [Gary Morgenthaler](#)

After 50 years of frequent failure and narrow success, artificial intelligence (AI) is going mainstream. A confluence of trends—cloud computing, smart phones, expanded broadband capability, improved AI algorithms, plus the steady Moore's Law expansion of raw processing power—is producing a vast acceleration in AI capability. Not only are AI's individual subdisciplines—speech recognition, natural language understanding, machine learning, computer vision, etc.—improving, they are beginning to work in concert. We are, finally, starting to approach the subtlety of real human intelligence. In the process, we are moving from the merely good to the uncanny.

Within the next three to five years, you will be able to:

- **Use virtual personal assistants (VPAs)** to manage your business and social calendars. VPAs will effortlessly find the most convenient time for a four-person meeting next week or reserve you a 2 p.m. tee time at the club. They will know your social graph and habitual patterns and, increasingly, like any good executive assistant, make helpful suggestions. ("Do you want me to invite your accountant to this meeting?")
- **Text ahead to a descendant of today's Roomba** to "vacuum the downstairs" confident in the knowledge that it will be able recognize and avoid both the antique desk and the cat.
- **Indulge your taste for robatayaki**, barbeque typically served at family-owned Japanese restaurants where English proficiency is minimal. You'll be able to locate the nearest such restaurant to your Osaka hotel, determine its hours of operation, and use your smartphone translator to order in serviceable Japanese.
- **Ask search engines questions of fact and get definitive answers**—not millions of blue links. Instead of mere keywords—say, "Ricky Henderson," "league leader," "walks," and "stolen bases"—you'll ask, "In what years did Ricky Henderson lead the American League in both walks and stolen bases?" (Answer: 1982, 1983, 1989, and 1998.)

Within the same period, changes to the IT competitive landscape will prove just as profound:

- **Most new smartphone applications will have to be voice-enabled** rather than thumb-typed. Customers will insist on the capability to search, text, e-mail, schedule, collaborate, and purchase just by talking to their phones.
- **AI will emerge as the key differentiator in smartphone competition** and with it, overall IT leadership.

To grasp the full magnitude of the tectonic shift underway requires some understanding of AI's checkered past. Founded as a discipline in 1956, AI has engendered lofty hopes ("Within 20 years, computers will be able to do anything a man can do," according to Herb Simon, one of the discipline's four founders) and bitter

disillusionment. Despite decades of intriguing laboratory demonstrations, AI usually fell well short of real-world reliability.

Take AI's natural language understanding of a phrase such as "I want to drop off my car." Does the speaker want to turn over parking of his car to a valet, push the car off a cliff, or perhaps injure yourself by falling off the car? A meaning easily grasped by most Americans has proven exasperatingly hard for AI to comprehend consistently. Better natural-language-understanding algorithms help, but the ultimate solution almost always comes down to hundreds of millions of iterations of speech patterns with almost as many individual English speakers. In AI, above all, what matters for real-world applicability is brute force computing power.

Significant AI improvement has thus had to await the extra processing power predicted by Moore's Law. By the late 1990s, processing power had reached the point where, with the right algorithms, computers could reliably recognize the several hundred basic words used to reserve airline seats or make stock transactions. As a result, successful speech recognition companies emerged. Along with speech recognition, AI has demonstrated commercial success in other relatively narrow subdisciplines, such as logistics, data mining, and medical diagnostics.

ADVENT OF THE CLOUD

Until recently, that has been the status of AI: gradual, partial success coupled with research and development so segregated into silos that practitioners rarely communicate with one another. Enter cloud computing. Although valued primarily for its efficiency in data administration—and certainly not new—AI developers have begun to take advantage of it only in recent years. In so doing, they have found new ways to work with all that siloed AI information and, essentially, recreate a broader AI discipline.

First, by creating deep pools of easily accessed data, cloud computing facilitates data mining and crowd sourcing. These techniques intelligently sort the millions of facts about everyday life that humans know, i.e. "common sense." Companies such as [True Knowledge](#) are automatically scouring Wikipedia, Freebase, and other databases for "facts"—all painstakingly entered and curated by individual contributors. AI-based common-sense reasoning, although not yet here, appears within reach.

Second, those same deep pools of data facilitate rapid feedback. Machine-learning algorithms that automatically recognize complex patterns and make intelligent decisions based on data can now be applied against these vast data sets and use that experience to continuously enhance performance.

THE OLD, LABORIOUS WAY

That's a monumental improvement. Ten years ago, Nuance, a speech recognition company in which I invested and on whose board I served, had to collect results laboriously from discrete voice response platforms decentralized throughout the population. For example, the company would capture how speakers in Boston and rural Tennessee would each pronounce the word "fire." Then Nuance engineers would periodically reconsolidate that information into a new main data set that would go out every 18 months to customers as a new software release.

Today, Nuance and other speech recognition companies, such as [Vlingo](#), simply draw on data pools in the cloud to compare and analyze constantly hundreds of millions of utterances and feed results back into their systems in near real time. One result: commercially available dictation software that can capture the tens of thousands of words in the whole English language with near 100 percent accuracy. Other AI comparison machines, such as those for facial recognition, also learn and steadily improve their products in similar fashion.

Third, and most significantly, cloud computing both turbocharges Moore's Law processing speeds and enables those narrow AI subdisciplines to communicate and work in concert. From 1990 to 2020, Moore's Law progress alone will increase the processing power of a single server by a factor of 1 million. By 2020, in contrast, cloud computing could multiply AI-related processing power by a factor of 1-billion. The combination of cloud computing and open source programs such as the Yahoo! ([YHOO](#))-developed Hadoop, permits AI systems to run data and algorithms across multiple servers simultaneously. For AI, that means we can split off those cumbersome silo subdisciplines (machine learning, speech recognition, dialog management, etc.), process them in parallel, and neatly reassemble them into something that makes sense.

PREPARE FOR BIONIC SWARMS

In the future, swarms of computers, like colonies of ants or flocks of starlings, will be directed via cloud computing toward global problem solving. (If this sounds to you something like the "bionic swarms" from William Gibson's 1984 science fiction novel, *Neuromancer*, you're right.) AI researchers already have an official name for this phenomenon: Swarm Intelligence. Darpa recognized this opportunity and its potential for future military action early on. Five years ago, it began contracting more than \$200 million with [SRI International](#) and leading university research subcontractors to reassemble AI subdisciplines into an integrated whole. The goal: a mobile, automated knowledge assistant for battlefield command and control.

But it's not only in battle that mobile computing is key. Because of its size, the smartphone has limited ability for information capture and display. Users must gather data from the cloud and find a way to make decisions on the fly. Any AI applications that help users quickly turn complexity into simplicity thus become essential.

With financing from my firm, San Jose (Calif.) startup [Siri](#) licensed the Darpa-funded SRI technology and developed a virtual personal assistant that allowed users to access services verbally, such as restaurant reservations and movie tickets, from the Web. Apple ([AAPL](#)) bought Siri earlier this year. (Not disclosing terms was a condition for the deal.) Expect to see more AI-based mobile technology quite soon. Besides Apple, both Google ([GOOG](#)) and Microsoft ([MSFT](#)) have made intelligence-at-the-interface a key R&D focus. Google's recent announcement of its Voice Actions app for Android, seen by analysts as a response to Apple's Siri acquisition, should really be understood more broadly as the second of many salvos in the coming Smartphone AI Wars.

And what of the larger future of AI? Optimists, such as Raymond Kurzweil in *The Singularity is Near*, foresee an AI utopia in which human and machine intelligence combine. Pessimists, such as Samuel Butler in his prescient *Darwin Among the Machines* (1863), forecast a dystopian future akin to SkyNet in Terminator 3. For now, however, it's clear that the decade of mainstream AI has arrived. The question for businesses and investors alike is: "What's your AI strategy?"

Gary Morgenthaler is a partner at Morgenthaler Ventures in Menlo Park, Calif.