

here, and their replies should be on the way. In another 15 or 20 years we will receive their message and meet our future. Let us be neither surprised nor disappointed if its form is that of Artoo Detoo, the bright, personable canister packed with silicon chips.

THE AI REVOLUTION IS ON

Steven Levy

Steven Levy, a senior writer for *Wired*, has reported on digital technology for a quarter century. His articles have appeared in the *New York Times*, *Rolling Stone*, and *Macworld*, as well as *Wired*, where this selection appeared on December 27, 2010. Levy is the author of several books, including *In the Plex: How Google Thinks, Works and Shapes Our Lives* (2011), *Hackers* (2010), *Insanely Great: The Life and Times of the Macintosh, the Computer That Changed Everything* (1994), and *Artificial Life: The Quest for a New Creation* (1992). Here, he distinguishes between two visions of artificial intelligence: the kind of AI currently in place (using clever algorithms to access massive databases) and the predictions of futurists who foresee the emergence of thinking machines that will mark our "next step in the evolutionary process."

Diapers.com warehouses are a bit of a jumble. Boxes of pacifiers sit above crates of onesies, which rest next to cartons of baby food. In a seeming abdication of logic, similar items are placed across the room from one another. A person trying to figure out how the products were shelved could well conclude that no form of intelligence—except maybe a random number generator—had a hand in determining what went where.

But the warehouses aren't meant to be understood by humans; they were built for bots. Every day, hundreds of robots course nimbly through the aisles, instantly identifying items and delivering them to flesh-and-blood packers on the periphery. Instead of organizing the warehouse as a human might—by placing like products next to one another, for instance—Diapers.com's robots stick the items in various aisles throughout the facility. Then, to fill an order, the first available robot simply finds the closest requested item. The storeroom is an ever-shifting mass that adjusts to constantly changing data, like the size and popularity of merchandise, the geography of the warehouse, and the location of each robot. Set up by Kiva Systems, which has outfitted similar facilities for Gap, Staples, and Office Depot, the system can deliver items to packers at the rate of one every six seconds.

The Kiva bots may not seem very smart. They don't possess anything like human intelligence and certainly couldn't pass a Turing test.³ But they represent

³The "Turing Test" as conceived by mathematician Alan Turing (1912–1954): If a human interrogator communicates via computer with both another human being and a computer and cannot tell, from the printed or on-screen responses, which is which, the machine can be said to exhibit humanlike behavior and therefore passes the Turing Test.

a new forefront in the field of artificial intelligence. Today's AI doesn't try to re-create the brain. Instead, it uses machine learning, massive data sets, sophisticated sensors, and clever algorithms to master discrete tasks. Examples can be found everywhere: The Google global machine uses AI to interpret cryptic human queries. Credit card companies use it to track fraud. Netflix uses it to recommend movies to subscribers. And the financial system uses it to handle billions of trades (with only the occasional meltdown).

This explosion is the ironic payoff of the seemingly fruitless decades-long quest to emulate human intelligence. That goal proved so elusive that some scientists lost heart and many others lost funding. People talked of an AI winter—a barren season in which no vision or project could take root or grow. But even as the traditional dream of AI was freezing over, a new one was being born: machines built to accomplish specific tasks in ways that people never could. At first, there were just a few green shoots pushing up through the frosty ground. But now we're in full bloom. Welcome to AI summer.

5 Today's AI bears little resemblance to its initial conception. The field's trailblazers in the 1950s and '60s believed success lay in mimicking the logic-based reasoning that human brains were thought to use. In 1957, the AI crowd confidently predicted that machines would soon be able to replicate all kinds of human mental achievements. But that turned out to be wildly unachievable, in part because we still don't really understand how the brain works, much less how to re-create it.

So during the '80s, graduate students began to focus on the kinds of skills for which computers were well-suited and found they could build something like intelligence from groups of systems that operated according to their own kind of reasoning. "The big surprise is that intelligence isn't a unitary thing," says Danny Hillis, who cofounded Thinking Machines, a company that made massively parallel supercomputers. "What we've learned is that it's all kinds of different behaviors."

AI researchers began to devise a raft of new techniques that were decidedly not modeled on human intelligence. By using probability-based algorithms to derive meaning from huge amounts of data, researchers discovered that they didn't need to teach a computer how to accomplish a task; they could just show it what people did and let the machine figure out how to emulate that behavior under similar circumstances. They used genetic algorithms, which comb through randomly generated chunks of code, skim the highest-performing ones, and splice them together to spawn new code. As the process is repeated, the evolved programs become amazingly effective, often comparable to the output of the most experienced coders.

MIT's Rodney Brooks also took a biologically inspired approach to robotics. His lab programmed six-legged buglike creatures by breaking down insect behavior into a series of simple commands—for instance, "If you run into an obstacle, lift your legs higher." When the programmers got the rules right, the gizmos could figure out for themselves how to navigate even complicated terrain. (It's no coincidence that iRobot, the company Brooks cofounded with his MIT students, produced the Roomba autonomous vacuum cleaner, which doesn't initially know the location of all the objects in a room or the best way to traverse it but knows how to keep itself moving.)

The fruits of the AI revolution are now all around us. Once researchers were freed from the burden of building a whole mind, they could construct a rich bestiary of digital fauna, which few would dispute possess something approaching intelligence. "If you told somebody in 1978, 'You're going to have this machine, and you'll be able to type a few words and instantly get all of the world's knowledge on that topic,' they would probably consider that to be AI," Google cofounder Larry Page says. "That seems routine now, but it's a really big deal."

- 10 Even formerly mechanical processes like driving a car have become collaborations with AI systems. "At first it was the automatic braking system," Brooks says. "The person's foot was saying, I want to brake this much, and the intelligent system in the middle figured when to actually apply the brakes to make that work. Now you're starting to get automatic parking and lane-changing." Indeed, Google has been developing and testing cars that drive themselves with only minimal human involvement; by October, they had already covered 140,000 miles of pavement.

In short, we are engaged in a permanent dance with machines, locked in an increasingly dependent embrace. And yet, because the bots' behavior isn't based on human thought processes, we are often powerless to explain their actions. Wolfram Alpha, the website created by scientist Stephen Wolfram, can solve many mathematical problems. It also seems to display how those answers are derived. But the logical steps that humans see are completely different from the website's actual calculations. "It doesn't do any of that reasoning," Wolfram says. "Those steps are pure fake. We thought, how can we explain this to one of those humans out there?"

The lesson is that our computers sometimes have to humor us, or they will freak us out. Eric Horvitz—now a top Microsoft researcher and a former president of the Association for the Advancement of Artificial Intelligence—helped build an AI system in the 1980s to aid pathologists in their studies, analyzing each result and suggesting the next test to perform. There was just one problem—it provided the answers too quickly. "We found that people trusted it more if we added a delay loop with a flashing light, as though it were huffing and puffing to come up with an answer," Horvitz says.

But we must learn to adapt. AI is so crucial to some systems—like the financial infrastructure—that getting rid of it would be a lot harder than simply disconnecting HAL 9000's modules. "In some sense, you can argue that the science fiction scenario is already starting to happen," Thinking Machines' Hillis says. "The computers are in control, and we just live in their world." Wolfram says this conundrum will intensify as AI takes on new tasks, spinning further out of human comprehension. "Do you regulate an underlying algorithm?" he asks. "That's crazy, because you can't foresee in most cases what consequences that algorithm will have."

In its earlier days, artificial intelligence was weighted with controversy and grave doubt, as humanists feared the ramifications of thinking machines. Now the machines are embedded in our lives, and those fears seem irrelevant. "I used to have fights about it," Brooks says. "I've stopped having fights. I'm just trying to win."