

## Chapter 1

### **The Coming Evolution**

Ben Goertzel & Stephan Vladimir Bugaj

#### **1. STEPS TOWARD THE SINGULARITY**

Nearly everyone who has seriously thought about the evolution of technology over the next few hundred years has come to the same conclusion: We live at a crucial point in history – an incredibly exciting and frightening point; a point that is stimulating to the point of excess, intellectually, philosophically, physically and emotionally. A number of really big technologies are brewing. Virtual reality, which lets us create synthetic worlds equal in richness to the physical world, thus making the Buddhist maxim “reality is illusion” a palpable technical fact. Biotechnology, allowing us to modify our bodies in various ways, customizing our genes and jacking our brains, organs and sense organs into computers and other devices. Nanotechnology, allowing us to manipulate molecules directly, creating biological, computational, micromechanical, and other kinds of systems that can barely be imagined today. Artificial intelligence, enabling mind, intelligence and reason to emerge out of computer systems – thinking machines built by humans. And advances in unified field theory in physics will in all likelihood join the party, clarifying the physical foundation of life and mind, and giving the nanotechnologists new tricks no one has even speculated about yet.

Even immortality is not completely out of the question. As Eric Drexler argued in *Engines of Creation* (1987), nano-scale robots could swarm through your body repairing your aging cells. Or more simply, appropriate

medication could put an end to cell death – after all, amoebas don't age, why should we? Or, as Hans Moravec depicted in his classic book *Mind Children* (1990), brain scan technology combined with AI could have us uploading our minds into computers once our bodies wear down. It sounds like science fiction, but it's entirely scientifically plausible: these would be immense engineering achievements, but wouldn't violate any known scientific laws. A lot of seemingly impossible things will soon be possible.

This is what some recent pundits have labelled "The Singularity." Ray Kurzweil is writing a book called *The Singularity is Near*, but he's only the latest in a series of thinkers to take up this theme. The term Singularity as used in this context was introduced by sci-fi writer and futurist Vernor Vinge in the 1970's.

Bill Joy, Chief Scientist of Sun Microsystems, one of the leading companies pushing forward the current phase of the tech revolution, recently wrote an article in *Wired* painting this same kind of future, but with a markedly dystopian bent

(<http://www.wired.com/wired/archive/8.04/joy.html>). He believes all this amazing technological development will happen, and he finds it intensely scary. It's difficult to blame him, actually. The potential for abuse of such technologies is obvious. We have to hope that an ethical evolution comparable to the technological evolution will occur at the same time, beautifully synchronized. This is essentially what Ray Kurzweil foresees in *The Age of Spiritual Machines* (1999). Kurzweil think its all going to happen in the next 20 years. By comparison, I consider myself conservative: I think it may take 50 or so, although in 20 years we'll certainly have moved a long way from where we are now.

The Internet is part of this heady mix. It's a low-tech virtual reality itself, sucking more and more of our time and attention away from the physical world. It's both a brain and a perceptual world for artificial intelligence systems. And it's a tool for accelerating technical progress in every possible direction, enabling unprecedentedly efficient communication and information retrieval.

This evolving network of long-term possibilities is wonderful and amazing to think about, but, on the other hand, it's almost too big for any one mind or small group of minds to grapple with. Imagine a bunch of pre-linguistic proto-humans trying to comprehend what the advent of language was going to do to them. That's basically the kind of situation we're in! Nevertheless, in spite of the difficulty intrinsic in foresight and the impossibility of planning any kind of revolution in advance, least of all the technological and psychocultural kind, there's something we can do beside sit back and watch as history leads us on. We can focus on particular aspects of the revolution we're wreaking, understanding these aspects as part of the

whole, and also focusing hard on the details, remembering that, historically, some of the subtlest and most profound general humanistic insights have come from dealing with very particular issues.

What I've been working on, for the past few years, the Webmind AI system, is what I believe will be the first component of the emerging tech revolution to fall fully into place: Internet-based artificial intelligence. Of course, I realize this is a gutsy statement. The field of artificial intelligence isn't all that fashionable these days; and this is understandable enough. AI's biggest claim to fame isn't any of its particular achievements – beating Kasparov at chess, diagnosing diseases better than doctors, mastering integral calculus – but rather its incredibly persistent habit of over-promising and under-delivering. This is always fatal in the business world, and not very favorable among academic circles either. But I believe that AI is finally ready to outgrow its history of big brags and false starts. AI's time has finally come.

Within the next few years, I believe, there will emerge the first real AI systems. The Webmind system that we're building at Webmind Inc. will be one of them, but there may well be others. And within the next few decades this technological advance will induce fundamental changes in the human condition, transforming the way we view ourselves as thinking beings, the way we interact with each other through electronic communication networks – the way we work, the way we feel, the way we live.

Of course, AI is only part of the huge transformation we humans are bringing down upon ourselves. Eventually AI will cease to exist as a discipline in itself, becoming diffused in the general matrix of technological innovation. Every new technology will be intelligent in one way or another, and AI will inform nanotech, biotech, virtual reality, unified physics, chemistry, refrigerators, toasters – you name it. But everything has to start somewhere. Life originally started with a few thousand molecules huddling together, perhaps inside a water droplet, developing primitive reproductive and metabolic abilities. Life has gone far beyond that now, growing to encompass things as complex and fantastic as you and me. But the simplest incarnations of life still have something to teach us, because they share so many properties with what has grown from them. Like proto-cells in the primordial soup, the AI programs of the next decade will be the first, primitive, early versions of a whole new order of being. Viewing them with enough imagination, one will be able to get at least a murky, muddled glimpse of what's awakening within and around us.

## 1.1 The Coming AI Revolution

The ongoing acceleration of AI development is obvious from technical advances in the Internet industry over the last few years. Take the case of search engines for example – or in the lingo of computer science, “information retrieval” tools. At the tail end of the technological race for intelligent information retrieval, one has standard search engines like Alta Vista and Yahoo. These have basically no intelligence whatsoever. They rely on pure bulk of information. Then you have moderately intelligent search tools made by firms like LexiQuest and Autonomy, that apply various specialized algorithms to grasp something of the meaning of texts. This is where things have stood for a long time. But then, over the last year, something new has arisen: A dozen or so start-up firms have come to prominence with technology that goes beyond this, and tries to understand text in a more thorough and flexible way, building subtle “semantic maps” describing of the meanings of documents. There are shortcomings in all this work: Talavara, for example, has one of the best AI search systems around, but it focuses too much on the syntactic analysis of documents, and uses a semantic map that isn’t nearly as flexible as the corresponding structures in the human mind. Similarly, WorldFree’s Know-All product does a good job of answering a variety of questions, but it uses an overly rigid “ontology” for representing knowledge – a fixed set of categories, not a flexible one like exists in the human mind.

None of these firms have really come to grips with the problem of reasoning on the semantic maps that their AI systems glean from reading text. But this is the next step. Over the next two years, you’ll likely see firms coming out with sophisticated reasoning engines hooked up to their syntax understanding and semantic mapping engines. Then, a couple years after that, people will realize that reason isn’t enough – that you need intuition as well – and the problem of synthesizing reason and intuition in a single flexible adaptive intelligent system will become paramount. And so on. The Webmind system provides a way of leapfrogging much of this incremental development, because it’s been designed up-front with a synergetic model of all the mind’s functions, rather than incrementally adding functions based on competitive business needs. However, the main point where the evolution of AI is concerned isn’t the coolness of Webmind or any other particular AI system – it’s that we now have practical commercial problems, such as Web search, that are being solved by AI technologies; and that because of this, AI technologies are improving at a fever pitch. Even the major search engine companies – the technology dinosaurs of the Internet business – have started to jump on board, with

AltaVista releasing an AI-supercharged smart search site last month. It isn't all pie in the sky anymore. The PR departments of these various firms have a way of avoiding the word "AI", but that's exactly what it is. Of course, it isn't real AI yet – not truly adaptive, flexible, self-aware intelligence – but the path from current AI products to a real thinking computer program isn't all that difficult to map out, if you have a feel for the terrain.

What has made this onslaught of AI improvements possible? AI researchers haven't gotten any more brilliant ... and the technical teams of these various AI Internet firms don't tend to have any profound new insights into digital mind. Rather, they're mostly implementing ideas that have been around in academia for quite a while. What's sparked the current burst of development is, quite simply, *hardware* – the tremendous accelerations in computer hardware that have occurred over the past two decades. There have been vague theories on how to make AI work for a long time, but without hardware adequate for implementing the theories, there was little incentive to make them precise and work out all the details. Now the hardware is there, and what isn't there will be there in a few years. Computers with gigabytes of RAM cost only a few thousand dollars, and networking technology lets anyone build their own distributed supercomputer. We take for granted that we need to buy a new computer every two years because the old ones so rapidly get useless – but think about how amazing that is! What if the same were true of cars or refrigerators or musical instruments?

This hardware acceleration won't bring us biotech or nanotech or unified physics; it won't even bring us virtual reality, which requires much better human sensory interfaces than we have today. But it will bring us artificial intelligence. And AI, once we have it, will help launch the other component technologies of the emerging tech revolution. AI will help accelerate biotech, by helping us to understand how the DNA sequences mapped out in the Human Genome Project actually combine to give rise to self-organizing organisms like you and me. Biotech will allow us to jack computers into our bodies in new ways, enabling truly visceral virtual reality experiences. Eventually maybe AI's will be pressed into service solving the numerous hard engineering problems required to make nanotech actually work. Everything will fall into place during the next 100 years or so; but I suspect that Internet intelligence will be the thing that leads the way, coming first and making it easier for the other things to happen.

Creating Internet Intelligence

Wild Computing, Distributed Digital Consciousness, and  
the Emerging Global Brain

Goertzel, B.

2002, XVI, 330 p., Hardcover

ISBN: 978-0-306-46735-6