

## Preface

The Internet as it exists today is just a little baby. But it's on the verge of a fundamental transition. Today it's a distributed network of content and software, serving diverse people diverse functions. Soon enough it will be a self-organizing intelligent system, with its own high-level coherent patterns, serving not only as a mind but as a world inhabited by a diversity of digital life forms.

My aim in this book is to get across some general concepts that I feel must become widely understood if we are to, collectively, bring the Internet past its current state and into the next phase.

Only a few of these ideas are my own inventions; this is a body of knowledge that has emerged from a large community of computer scientists, engineers, complexity scientists and associated thinkers over the last few decades. But I believe I've filled in a few crucial gaps, and highlighted interconnections between others' ideas in various novel ways.

This book straddles traditional categories in several respects. It's neither a philosophy book nor a science book nor an engineering book – it's a little bit of all of these. It's neither a narrowly-focused academic treatise nor a trade book aimed at the general audience – rather, it's aimed at readers with a strong intellectual background but not necessarily specialty training in computer science, complexity science, or other disciplines involved here.

By fleshing out and implementing the ideas given here and other related ideas, over the next decade, we will be able to transform the Internet into something far beyond what most people have ever imagined. This will include Internet-based AI programs that surpass human beings in general intelligence. It will include Internet Alife agents that buzz around from machine to machine like intelligent digital animals, fulfilling various tasks

and generally enjoying themselves. It will include coherent, self-organizing behaviors on the whole-Net scale, with specific implications for Internet functionality. For example, one day in 2012, the Net as a whole may have a big insight into human nature, which will cause all search engine queries regarding literature and psychology to come back with more intelligent answers.

Some of our expectations will of course remain unfulfilled – for example, the Turing test, which requires a computer to simulate a human in conversation, is unlikely to be passed by any digital system until we can engineer human-like bodies for our AI programs. Cranial jacks and true simulated reality may be further off than some cyber-visionaries would like to think, because research into this kind of biotech cannot proceed with the speed of purely computational work, for ethical as well as technological reasons. But in other areas, I predict, our computer programs will exceed our expectations dramatically. Our understanding of the collective dynamics of the human race will be greatly enhanced by interaction with AI systems linked in to the collective data resources of the Internet – a point that futurists have rarely focused on.

Yes, I admit it, I'm an all-out hi-tech optimist. Of course, I am well aware that technological optimism has been proved wrong before. AI pioneers of the 1950's and 60's predicted that true machine intelligence was right around the corner. On the other hand, the explosion of power in computer hardware has exceeded even the most optimistic serious predictions. So technological pessimism has been proved wrong before too. Now we have PC's with 4 gigabytes of RAM, mainframes with 100; and we have high-bandwidth network cable that can link an arbitrary number of computers together into a single computational process. The Internet, today, contains all the computational firepower needed to create a global distributed intelligent system, surpassing human intelligence in many ways, serving as both a mind and a world for digital organisms. The ideas in this book provide the key to making this potential real.

I've been nursing these grand ideas for a while, and others have been nursing them for longer, but this seems a timely time to present them in a systematic way. These days, more and more people are recognizing that the Internet is more than just a cool way to send naked pictures of Pamela Anderson to your buddies in Chechnya. But there doesn't yet seem to be a common conceptual framework for thinking about the Net and what it can become. My goal here is to provide such a framework, mainly on a conceptual level, digging into the details whenever it seems necessary.

What kind of conceptual framework am I talking about? One phrase I use for it is *wild computing*. So far computer programs have been tame, trained beasts, doing exactly what they're told exactly the way they're told to

– and when they don't, it's a bug that should be eliminated. Over the next few years this will change substantially. The Internet will become a mind in itself, and a world for digital organisms. Software living on the Net will have to behave flexibly and unpredictably in order to be useful, in order to survive. In fact, *Wild Computing* was the original title for this book, and a much briefer early draft was posted online several years ago, under that title.

The notion of wild computing implies a new kind of computer science and computer technology, one that focuses on self-organizing networks and emergent dynamics rather than algorithms and data structures (though algorithms and data structures will still be there, to be sure, just as the bit-level constructs of assembly-language were not vanquished by the advent of structured programming). Artificial intelligence is essential here, because humans don't have the time or the ability to deal with the glut of rapidly shifting information that modern networks bring. And the network computing revolution is essential, both because the Internet is a network (a network of networks), and because minds and brains and ecosystems are *necessarily* networks (networks of networks as well).

Java, the premier language of the new paradigm, is maturing into a central role in server-side network apps. Exciting new Java-derivative technologies like Objectspace Voyager and Jini and Javaspace are beginning to take off too. Electronic commerce has long since become real; and online AI is finally picking up speed – embedded in websites like Excite and amazon.com are sophisticated systems for learning and guessing user profiles. We are just a few years away from the situation where the various intelligent systems on the Net, like the ones inside these websites, are learning from each other rather than existing as islands of intelligence in a sea of inert text and data. The Internet, and the computing environment in general, is poised for the advent of real artificial intelligence. What is needed to make it happen is understanding on the part of the technical community – understanding of how the network of mind emerges from the underlying network of computers.

Next, a few words on how I came to the ideas in this book may be appropriate. Until 1997, I thought and wrote about these types of issues from an academic perspective. My PhD is in mathematics, but I held university appointments in math, computer science and psychology. I was seeking the broadest possible intellectual perspective from which I could understand all the deepest issues related the mind, the universe and the emerging sphere of high-technology. Then in 1997, I left academia and founded the start-up company Webmind Inc. (formerly called Intelligenesis; we just changed the name). Since that time I've been involved principally with the design, implementation and marketing of Webmind, an Internet-

based AI system that exemplifies many of the principles I'll discuss here. The last chapter of the book describes Webmind, though not in great detail – both because it's a proprietary system, and because that would be a long book in itself. (Actually, this book exists in two volumes – a nontechnical tome called *Digital Intuition* and a more mathematical treatment provisionally entitled *Webmind: The Architecture and Dynamics of a True AI* – but these are currently only being distributed internally within Webmind Inc.) Webmind Inc. dissolved in March 2001, but the AI work that was started within that firm continues, pushed on by myself and a group of “AI diehards,” mainly located in Brazil where Webmind Inc.'s largest office was located.

The Webmind project has been unbelievably challenging and exciting, but has left me fairly little time to reflect on the general lessons and principles underlying the work I'm doing, and even less time to write down these lessons and principles systematically. This book synthesizes some of the ideas along these lines that I have found time to write about. Most of the chapters originated as informal articles, written to clarify ideas to myself and distributed informally to various co-workers and acquaintances. Some of them embody ideas I had years and years ago, others are insights that only came to me as a consequence of my recent practical work. Overall, the writing of this book began in 1995, when I first started thinking about Internet intelligence, and continued until early 2000.

The final revision of the manuscript was done in early 2001, at which point the material seemed quite “old” to me, since I had long since digested it and moved on to the next phase of ideas. But that's often the case: by the time a book is published, the author has often moved on to ideas that will be presented in the next book. Not that, as I'm typing these words not in July 2001, I'm no longer excited by the ideas in these pages. But as I write these words I'm taking a break from my ongoing research work, which consists largely of taking the ideas in the book to the next level, beyond what's described here. I hope that some readers will be inspired by the material given here to develop it further in all manner of different directions.

I find that the kind of ideas discussed in this book possess an odd combination of timeliness and time-bound-ness. Aspects of the conceptual framework given here are timeless in the sense that all deep philosophy is timeless. Aspects are timeless in the weaker sense that powerful science is timeless. And other aspects are extremely closely tied to the technology and society of today. Like the human mind itself, a book like this operates on many different time scales.

Now, without further ado, let me explain the structure of the book.

Following the first chapter, which reviews the general themes of the book, I've divided the book into three parts, reflecting the most natural division and progression of the subject matter.

Part I presents a general theory of mind and intelligence. If we're going to talk about Internet intelligence, we'd better know what intelligence is first. There is some philosophical background here, intended to establish a working conceptual vocabulary for addressing the Internet, the mind, and their interrelation. There is some clarification of basic concepts regarding evolution, mind and complexity. Then, there is a review of my previously published ideas about the structure and dynamics of mind. These ideas come from my four previous academic books, but are summarized here in what I hope is a clearly crystallized yet nontechnical way. The focus is on the mind as a network, because in later chapters we'll be exploring how the mind-network synergizes with computer networks in various ways.

Part II deals with the Net. The Net as a mind. The Net as a new kind of economy. The Net as a world, as a complex, self-organizing system. Of course, depending on your point of view, you can see the Internet as just about anything – it's remarkably Protean (and proteanly remarkable!). The background of Part I is necessary to understand why the perspective on the Net taken here has particular importance and depth. The future of the Net, it's argued, has a phase transition in it: a transition from a distributed network of computers into a coherent, self-organizing world/mind, the likes of which we've never before encountered.

Finally, Part III discusses concrete steps that can be taken to bring this phase transition about. The next-phase Net will evolve, yes – but we are the agents of this evolution. It will evolve, initially as a consequence of human beings building software that causes it to evolve. I'll describe what kinds of software need to be built in order to bring about the evolution of the net to the next level. Basically there are three aspects: a universal communication protocol for Internet agents; simple Internet Alife forms that migrate around and do useful things; and Internet intelligences like Webmind. The web of all these, taken together, will create the intelligent Net of the future.

Even after all the thought I've put into dividing these ideas into parts and chapters, the contents of the book still jump around a fair bit. But I'm not going to apologize for this too much. Basically, this is a diverse body of thinking, and I have not attempted to force a false unity upon it. I am too well aware of the real underlying coherence here to want to obscure it under a glib, sugar-coated systematization. (As Nietzsche said, "The will to a system is a lack of integrity.") Instead, the book mirrors its subject matter, in that the topics presented are not locked together into a rigid linear series, but represent rather nodes in a network, interrelating with each other in