

WIRED

By Analogy

A talk with the most remarkable researcher in artificial intelligence today, Douglas Hofstadter, the author of *Gödel, Escher, Bach*.

By Kevin Kelly

In 1979, Douglas Hofstadter burst into public consciousness with a book so out of the ordinary it won a Pulitzer Prize for its young first-time author. Titled *Gödel, Escher, Bach: An Eternal Golden Braid*, it was a rambling 777-page magnum opus crowded with dense diagrams about feedback loops, witty parables about knowledge, and intellectual puns that had mathematicians in stitches. And it was all electronically typeset by its hackerish author. The book became a bestseller, and in the 1980s, it could be found on the bookshelves of anyone who had anything to do with the new fields of cybernetics, artificial intelligence, and computer science - although how many read it, or understood more than a third of it is unclear. But nearly every other author writing since in the same area has referred at least once to the legendary *Gödel, Escher, Bach*.

As abruptly as Hofstadter burst on the scene, he disappeared. He began the book while a graduate student, and after its wild and unexpected success he went back to the lab and spent decades attempting to make his ideas work in real systems. He occasionally published his results in very technical journals, but he wrote little for the dedicated lay reader. Now, after 15 years, Hofstadter has published another major book. It is a technical tome, yet it still has humor: *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*.

Those readers hoping for the witty high of *Gödel, Escher, Bach* will be disappointed: this book is a workhorse. It attempts to lay out the elemental code of creativity - how a creative leap, no matter how small, is made. Hofstadter's ideas are, as usual, so original and so off-kilter that few yet know what to make of them. In his quiet way, Hofstadter is the most remarkable researcher in artificial intelligence today. He was interviewed by *Wired's Kevin Kelly*.

***Wired:* Looking back, what would you say your first book, *Gödel, Escher, Bach*, was really about?**

Hofstadter:

I can see that many people didn't understand what the book was trying to do. Many people took it to be nothing but a title. They would look at the three words - Gödel, Escher, Bach - and if they knew who these people were, they'd say, Oh, this is a book about mathematics, art, and music.

What *Gödel, Escher, Bach* was really about - and I thought I said it over and over again - was the word I. Consciousness. It was about how thinking emerges from well-hidden mechanisms, way down, that we hardly understand. How not just thinking, but our sense of self and our awareness of consciousness, sets us apart from other complicated things. How understanding self-reference could help explain consciousness so that someday we might recognize it inside very complicated structures

such as computing machinery. I was trying to understand what makes for a self, and what makes for a soul. What makes consciousness come out of mere electrons coursing through wires.

And yet many people treated the book as just some sort of big interdisciplinary romp whose point was simply to have fun. In fact, the fun was merely icing on the cake. Originally, the book was purely about the way the proof of Gödel's theorem kept cropping up in the middle of a fortress - Principia Mathematica by Bertrand Russell and Alfred North Whitehead - that was designed to keep it out. I thought, Here's a structure that attempts to keep out self-knowledge, but when things get sufficiently complex and sufficiently tangled, all of a sudden - whammo! - it's got self-representation in it. That to me was the trick that underlies consciousness.

So, at first, there were no dialogs, no jokes, no wordplay, and no references to Escher or Bach. But as I typed the manuscript up in '74, I decided it was written in an immature style. I decided to insert the dialogs and the Escher so that the playfulness became a kind of a secondary - but extremely important - part of the book.

Many people focused on those things and treated the book as a big game-playing thing. I had been aiming to have the book reach philosophers, people who thought about the mind and consciousness, and a small number actually saw what I was getting at, but most people just saw the glitter. At the time, I felt I'd lost a great deal by writing a book like that so early in my career, because I was no longer taken seriously by anybody.

Wired: Really?

Hofstadter:

Of course, I'm exaggerating when I say "by anybody," because many of the people I admired most in cognitive science and philosophy did take me seriously. But the majority of people characterized the book as a work of popularization or as a long-winded but amusing presentation of Gödel's theorem that threw in a bunch of crazy things like molecular biology, Zen, paintings, and puns.

On the other hand, I suppose if the philosophers had all written, "This is the tome of the ages about consciousness," but said nothing about the wordplay, I still would have felt hurt and sad. In either scenario I'd feel a little slighted. [Laughs.]

What the book did do was excite a lot of young people. Hundreds of people have written to me saying it launched them on a path of studying computer science or cognitive science or philosophy. And that's always nice. But often it's treated as fluff.

Wired: Is that why your new book, Fluid Concepts and Creative Analogies, has a more workaday, serious, this-is-real-science feel?

Hofstadter:

I guess I came to the point at which, after many years of research in this field, I felt I had to have a book that addressed all these themes in a way that was as direct and forthright as possible. So, in a sense, you could say I was trying to counter what I had done earlier. But it's more than that. I didn't put in a ton of glitter; nevertheless, the book is still fairly playful.

Wired: True: if you were trying to be staid, you weren't completely successful! A lot of your wit still seeps through in the new book. I was happy to see that, because it seems to me that part of what you're trying to say is that wit and humor are essential to intelligence.

Hofstadter:

I surely can't imagine an intelligence that wasn't capable of finding humor. I once held a humor workshop with a group of people I'd known for a long time. I thought these people would understand

what I meant when I said "humor." But a lot of participants kept bringing up things like ethnic jokes, people falling on bananas, and violence; somebody gets water sprayed in the face, and people laugh. But that wasn't what I meant. What I meant by humor is what I now call "slippage humor." I know it sounds like somebody slipping on a banana peel, but it isn't. It means the slippage of concepts. It means intellectual humor - humor based on the play of ideas. Maybe it's what one could call wit - clever remarks. That's the kind of humor I'm most interested in.

Wired: In your books, you regularly mention shifts, fluidity, slippage, and jumps. There is this constant activity as things seem to leap out of one place and into another. The one thing that seems to connect your books is the shifts that move from one place to another without an immediate, evident transition.

Hofstadter:

That's deeply related to my view of what thinking is all about. It's the ability to look at complex situations and strip away things that don't count. Let me give an example: A few people in artificial intelligence have talked about how we can look at a very complex series of events and then sum it up by retrieving from our minds a proverb - something like, "The pot is calling the kettle black." The ability to do that - to filter out, get to the core - can seem like a large, almost arbitrary leap from one place to another. But stripping away the irrelevancies gives you a little conceptual nugget; then you just step to the next one, which is closely related. It turns out that the next nugget you've stepped to takes you to some other domain - after you put back all the irrelevancies.

So, to reproduce that leaping process on a machine would present us with something very powerful.

But the crux of the process is the act of stripping away the irrelevancies. To get a computer to imitate this, most of my colleagues simply spoon-feed the computer the gist of two different situations. The gists have already been extracted, though, so it's really kind of a pointless exercise at that point.

Do you equate thinking with creativity? Are they the same in your book?

My work has often been a search for what makes special minds, not just ordinary minds. I'm much more curious about what makes for a great mathematical discovery, a great discovery in physics, or a great pun. I'm much more curious about that than I am about ordinary dinner-table conversation.

Nonetheless, I feel they're closely related. It's almost as if by having language we jump to a metalevel of knowledge that's intimately mixed in with the perceptual knowledge of the world around us. Once you have language, concepts are sort of self-referential. That's what fascinates me.

Wired: In your new book, you mention that you reject metaphors that understand intelligence as a search through a space. It's an interesting idea, because in some sense, this notion of the search space is the only decent idea that AI has come up with so far. It's a wonderful metaphor for understanding.

Hofstadter:

I don't think the word search is dead wrong. It's more the connotations that have come to surround it. Tons of people have spent their lives talking about very formal and mathematical ways of searching a space that's already well defined. All you need to do, they say, is hunt around in it using very efficient, mathematically optimized techniques. To me that has seemed so completely off the track that I now hesitate to even use the word search. There's no such thing as creation, there's no such thing as invention, there's only discovery. That's my personal feeling.

Wired: So, where are you headed? Are you trying mostly to understand human creativity? Or are you trying to replicate it?

Hofstadter:

Well, you catch me in a tricky dilemma here. I want any computer program that my students and I work on to delight me with its cleverness. I want it to outclever its programmers. But at the same time, if after 10 or 20 years of work my program composed a great novel, made a series of great mathematical discoveries, or wrote lots of great one-liners, I would be terribly distressed. I feel the human spirit is infinitely more complex than anything that we're going to be able to create in the short run. And if we somehow did create it in the short run, it would mean that we aren't so complex after all, and that we've all been tricking ourselves.

Take music. Some of my favorite composers are Chopin, Bach, and Fauré. There is something so deep about them. Or take Billie Holiday, singing with some of her accompanists in the 1930s - playing and improvising. Now, if all that incredible poignancy turns out to be something that can be mass marketed on a chip, it will destroy my image of something very deep and sacred to the human spirit. I'll just have to eat my words and say, "Well, I guess all that complexity was just another kind of circuitry we can manufacture." You want to write some new Bach music? Just type in b-a-c-h, hit carriage return, and, five minutes later, you'll have a whole new Mass.

If that were the case, I would be devastated. If such minds of infinite subtlety and complexity and emotional depth could be trivialized by a small chip, it would destroy my sense of what humanity is about - what humans are about, what love is about, what caring about people is about, and what humor is.

Wired: But what if it took several centuries of hard work to make a device that could spit these things out? And what if it wasn't a chip, but something about the size of a refrigerator? Would that make you feel better?

Hofstadter:

Only if the refrigerator had emotions as complex as ours. Let me read you a bit from Gödel, Escher, Bach: "A 'program' which could produce brilliant music would have to wander around the world on its own, fighting its way through the maze of life and feeling every moment of it. It would have to understand the joy and loneliness of a chilly night wind, the longing for a cherished hand, the inaccessibility of a distant town, the heartbreak and regeneration after a human death.

It would have to have known resignation and world-weariness, grief and despair." The refrigerator couldn't be stationary.

It would have to acquire its abilities by struggling in the world. By failing, succeeding, and socializing.

Wired: Do you think your own research into this realm of creativity has helped you become more creative?

Hofstadter:

No. I think it's channeled my creativity in certain directions. Take Fats Waller - a great jazz musician. He just immersed himself in jazz. And he became more and more creative as he immersed himself more deeply. A deep immersion in anything makes you more creative, if you're creative to begin with. Constantly honing your faculties enhances you.

But I'm a bit wary of the idea that by thinking about creativity I can tweak the mechanisms that are responsible for my creativity. Some people say that by studying parallel processing, parallel hardware, parallel this and that, we develop a whole new paradigm for thought that will make us think differently. That's a really weird claim. It seems no more likely to me than that it will alter the mechanisms of our digestive tract. We can't modify those things! There's a fine line between what it means to have your thoughts shaped and having your thought processes shaped.

Wired: You often talk about the essences of things, as in your attempt to create a program that can extract the stylistic essence of the letter a from many kinds of alphabets. You talk about reducing

complex things to their cores. Are you a Platonist? Do you believe all things have essences and cores?

Hofstadter:

Things are context-dependent. But if intelligence arose on another star system and somehow became aware of a situation that I too was aware of - and thought I really deeply understood - my hunch is the alien intelligence would understand the same thing I did. That's why we believe that prime numbers are going to be found by other civilizations.

I believe in the idea that intelligent creatures, wherever they are, when faced with complex situations, will see similar things in them - if they're really intelligent. Life is such that we have a certain way of filtering situations and finding what's at their core. Good intelligences, the kind that really survive well, will probably all be very similar to one another in their ability to do this.

Wired: Your new book is all about analogies and trying to understand how our minds understand statements like "Amsterdam is the San Francisco of Europe." If we make an analogy and say Richard Feynman was the coyote trickster of physics, then who is Douglas Hofstadter? Who are you? What's your role?

Hofstadter:

Hmmm. I guess I'm the type for whom analogy is the driving force behind the way I think. And I've managed to convince myself that analogy is really at the core of thinking - not just for myself, but for other people, too. I'm trying to put forth a vision of thought that involves - if you don't want to say "analogy-making" you can say "stripping away irrelevancies to get at the gist of things." I feel I've discovered something essential about what thinking is, and I'm on a crusade to make it clear to everybody. n

Kevin Kelly is executive editor at Wired.

[Copyright](#) © 1993-2004 The Condé Nast Publications Inc. All rights reserved.

[Copyright](#) © 1994-2003 Wired Digital, Inc. All rights reserved.