

## **Lifebox Immortality & How We Got There**

By [Rudy Rucker](#) and [Leon Marvell](#)

[This paper appeared in *Re:live Media Art Histories 2009*, Edited by Sean Cubitt and Paul Thomas, published by The University of Melbourne & Victorian College of the Arts and Music, available [online](#) as a Creative Commons file. This article is Copyright © Leon Marvell and Rudy Rucker, 2010.]

Now that the lifebox is so ubiquitous in these last, fading hours of the 21<sup>st</sup> century, it is opportune to examine how we got here and where the lifebox came from. This short history begins by re-presenting a historically important paper sourced from the recollections of *Rudy Rucker's lifebox* — perhaps the most famous and beloved of lifeboxes still in existence. Rucker's paper was originally presented at the *Re:live, the Third International Conference on the Histories of Media Art, Science and Technology* in the early years of the 21<sup>st</sup> century: November of 2009 in Melbourne, Australia.

### **Lifebox Immortality<sup>i</sup>**

Rudy Rucker

One of the most venerable dreams of science fiction is that people might become immortal by uploading their personalities into some kind of lasting storage. Once your personality is out of your body in a portable format, it could perhaps be copied onto a fresh tank-grown blank human body, onto a humanoid robot or, what the heck, onto a pelican with an amplified brain. Preserve your software, the rest is meat!

In practice, copying a brain would be very hard, for the brain isn't in digital form. The brain's information is stored in the geometry of its axons, dendrites and synapses, in the ongoing biochemical balances of its chemicals, and in the fleeting flow of its electrical currents. In my early cyberpunk novel *Software*, I wrote about some robots who specialized in extracting people's personality software — by eating their brains. When one of my characters hears about the repellent process, "[His] tongue twitched, trying to flick away the imagined taste of the brain tissue, tingly with firing neurons, tart with transmitter chemicals."<sup>ii</sup>

In this paper I'm going to talk about a much weaker form of copying a personality. Rather than trying to exactly replicate a brain's architecture, it might be interesting enough to simply copy all of a

person's memories, preserving the interconnections among them.

We can view a person's memory as a hyperlinked database of sensations and facts. The memory is structured something like a website, with words, sounds and images combined into a superblog with trillions of links.

I don't think it will be too many more years until we see a consumer product that makes it easy for a person to make a copy of their memory along these lines. This product is what I call a lifebox.<sup>iii</sup>

My idea is that your lifebox will prompt you to tell it stories, and it will have enough low-level language recognition software to be able to organize your anecdotes and to ask you follow-up questions. As the interviews progress, the lifebox's interviewer-agent harks back to things that you've mentioned, and creates fresh questions pairing topics together. Now and then the interviewer-agent might throw in a somewhat random or even dadaistic question to loosen you up.

As you continue working with your lifebox, it builds up a database of the facts you know and the tales you spin, along with links among them. Some of the links are explicitly made by you, others will be inferred by the lifebox software on the basis of your flow of conversation, and still other links are automatically generated by looking for matching words.

And then what?

Your lifebox will have a kind of browser software with a search engine capable of returning reasonable links into your database when prompted by spoken or written questions from other users. These might be friends, lovers or business partners checking you out, or perhaps grandchildren wanting to know what you were like.

Your lifebox will give other people a reasonably good impression of having a conversation with you. Their questions are combed for trigger words to access the lifebox information. A lifebox doesn't pretend to be an intelligent program; we don't expect it to reason about problems proposed to it. A lifebox is really just some compact digital memory with a little extra software. Creating these devices really shouldn't be too hard and is already, I'd say, within the realm of possibility — it's already common for pocket-sized devices to carry gigabytes of memory, and the terabytes won't be long in coming.

I discussed the lifebox at some length in my Y2K work of futurology, *Saucer Wisdom*<sup>iv</sup>, a book in the form of a novel, framed in terms of a character named Frank Shook who has a series of glimpses into the future — thanks to some friendly time-traveling aliens who take him on a tour in their tiny flying saucer. (And, no, I'm not a UFO true believer, I just happen to think they're cute and enjoyably

archetypal.)

You might visualize a lifebox as a little black plastic thing that fits in your pocket. It comes with a light-weight clip-on headset with a microphone and earphone. It's completely non-technical, anyone can use a lifebox to create their life story, to make something to leave for their children and grandchildren.

In my novel, my character Frank watches an old man using a lifebox. His name is Ned. White-haired Ned is pacing in his small back yard — a concrete slab with some beds of roses — he's talking and gesturing, wearing the headset and with the lifebox in his shirt pocket. The lifebox speaks to him in a woman's pleasant voice.

The marketing idea behind the lifebox is that old duffers always want to write down their life story, and with a lifebox they don't have to write, they can get by with just talking. The lifebox software is smart enough to organize the material into a shapely whole. Like an automatic ghost-writer.

The hard thing about creating your life story is that your recollections aren't linear; they're a tangled banyan tree of branches that split and merge. The lifebox uses hypertext links to hook together everything you tell it. Then your eventual audience can interact with your stories, interrupting and asking questions. The lifebox is almost like a simulation of you. And over time, a lifebox develops some rudimentary simulations of its individual audience members as well—the better to make them feel they're having conversations with an intelligent mind.

To continue his observations, my character Frank and his friends skip forward in time until past when Ned has died and watch two of Ned's grandchildren play with one of the lifebox copies he left behind.

Frank watches Ned's grandchildren: little Billy and big Sis. The kids call the lifebox "Grandpa," but they're mocking it too. They're not putting on the polite faces that kids usually show to grown-ups. Billy asks the Grandpa-lifebox about his first car, and the lifebox starts talking about an electric-powered Honda and then it mentions something about using the car for dates. Sis — little Billy calls her "pig Sis" instead of "big Sis" — asks the lifebox about the first girl Grandpa dated, and Grandpa goes off on that for awhile, and then Sis looks around to make sure Mom's not in earshot. The coast is clear so she asks some naughty questions about Grandpa's dates. Shrieks of laughter. "You're a little too young to hear

about that stuff," says the Grandpa-lifebox calmly. "Let me tell you some more about the car."

My character Frank skips a little further into the future, and he finds that lifeboxes have become a huge industry. People of all ages are using lifeboxes as a way of introducing themselves to each other. Sort of like home pages. They call the lifebox database a *context*, as in, "I'll send you a link to my *context*." Not that most people really want to spend the time it takes to explicitly access very much of another person's full context. But having the context handy makes conversation much easier. In particular, it's now finally possible for software agents to understand the content of human speech — provided that the software has access to the speakers' contexts.

Coming back to the idea of saving off your entire personality that I was initially discussing, there is a sense in which saving only your memories is perhaps enough, as long as enough links among your memories are included. The links are important because they constitute your *sensibility*, that is, your characteristic way of jumping from one thought to the next.

On their own, your memories and links aren't enough to generate an emulation of you. But if *another person* studies your memories and links, that other person can get into your customary frame of mind, at least for a short period of time. The reason another person can plausibly expect to emulate you is that, first of all, people are universal computers and, second of all, people are exquisitely tuned to absorbing inputs in the form of anecdotes and memories. Your memories and links can act as a special kind of software that needs to be run on a very specialized kind of hardware: another human being. Putting it a bit differently, your memories and links are an emulation code.

Certainly exchanging memories and links is more pleasant than having one's brain microtomed and chemically analyzed, as in my novel *Software*.

I sometimes study an author's writings or an artist's works so intensely that I begin to at least imagine that I can think like them. I even have a special word I made up for this kind of emulation; I call it *twinking*. To *twink* someone is to simulate them internally. Putting it in an older style of language, to *twink* someone is to let their spirit briefly inhabit you. A *twinker* is, if you will, like a spiritualistic medium channeling a personality.

Over the years I've twinked my favorite writers, scientists, musicians and artists: Robert Sheckley, Jack Kerouac, William Burroughs, Thomas Pynchon, Frank Zappa, Kurt Gödel, Georg Cantor, Jorge Luis Borges, Edgar Allan Poe, Joey Ramone, Phil Dick, Peter Bruegel, etc. The

immortality of the great ones results from faithful twinkling by their aficionados.

Even without the lifebox, if someone doesn't happen to be an author, they can make themselves twinkable simply by appearing in films. Thomas Pynchon captures this idea in a passage imagining the state of mind of the 1930s bank-robber John Dillinger right before he was gunned down by federal agents outside the Biograph movie theater in Chicago, having just seen *Manhattan Melodrama* starring Clark Gable.

John Dillinger, at the end, found a few seconds' strange mercy in the movie images that hadn't quite yet faded from his eyeballs — Clark Gable going off unregenerate to fry in the chair, voices gentle out of the deathrow steel *so long, Blackie ...* there was still for the doomed man some shift of personality in effect — the way you've felt for a little while afterward in the real muscles of your face and voice, that you were Gable, the ironic eyebrows, the proud, shining, snakelike head — to help Dillinger through the bushwhacking, and a little easier into death.

v

The effect of the lifebox would be to make such immortality accessible to a wider range of people. Most of us aren't going to appear in any movies, and even writing a book is quite hard. Again, a key difficulty in writing any kind of book is that you somehow have to flatten the great branching fractal of your thoughts into a long line of words. Writing means converting a hypertext structure into a sequential row — it can be hard even to know where to begin.

As I've been saying, my expectation is that in not too many years, great numbers of people will be able to preserve their software by means of the lifebox. In a rudimentary kind of way, the lifebox concept is already being implemented as blogs. People post journal notes and snapshots of themselves, and if you follow a blog closely enough you can indeed get a feeling of identification with the blogger. And many blogs already come with search engines that automatically provide some links. Recently the cell-phone company Nokia started marketing a system called *Lifeblog*, whereby a person can link and record their daily activities by using a camera-equipped cell phone. And I understand that the Hallmark corporation, known for greeting cards, is researching an on-line memory-keeping product.

Like any other form of creative endeavor, filling up one's lifebox will involve dedication and a fair amount of time, and not everyone will feel like doing it. And some people are tongue-tied or

inhibited enough to have trouble telling stories about themselves. Certainly a lifebox can include some therapist-like routines for encouraging its more recalcitrant users to talk. But lifeboxes won't work for everyone.

What about some science fictional instant personality scanner, a superscanner that you wave across your skull and thereby get a copy of your whole personality with no effort at all? Or, lacking that, how about a slicer-dicer that purees your brain right after you die and extracts your personality like the brain-eaters of *Software*? I'm not at all sure that this kind of technology will ever exist. In the end, the synaptic structures and biochemical reactions of a living brain may prove too delicate to capture from the outside.

I like the idea of a lifebox, and I have vague plans to try and make one for myself. *Rudy's Lifebox* — my personal pyramid of Cheops. I see it as a website or a cloud-based application that includes a large database with all my books, all my journals, some years of blog entries, and a connective guide/memoir — with the whole thing annotated and hyperlinked. And I might as well throw in my photographs, videos and sound-recordings — I've taken thousands of photos over the years.

It should be feasible to endow my lifebox with enough interactive abilities; people could ask it questions and have it answer with appropriate links and words. Off-the-shelf Google site-search box does a fairly good job at finding word matches. And it may be that the Wolfram|Alpha search engine—which purportedly has some measure of natural language comprehension—can soon do better.

For a fully effective user experience, I'd want my lifebox to remember the people who talked to it. This is standard technology—a user signs onto a site, and the site remembers the interactions that the user has. In effect, the lifebox creates mini-lifebox models of the people it talks to, remembering their interests, perhaps interviewing them a bit, and never accidentally telling the same story twice—unless prompted to.

If I'm dead by the time my lifebox begins receiving heavy usage, then in some sense I'm not all that worried about getting paid by my users. Like any web or cloud-based application, one could charge a subscription fee, or interrupt the information with ads.

If I use my lifebox while I'm still alive, some other options arise. I might start letting my lifebox carry out those interview or speaking gigs that I don't have the time or energy to fulfill. Given that many bits of this paper, "Lifebox Immortality," are in fact excerpted and reshuffled from my other writings, it's conceivable that my lifebox actually wrote this paper.

Moving on, my lifebox could be equipped to actively go out and post things on social

networking sites, raising my profile on the web and perhaps garnering more sales of my books and more in-person speaking invitations. This could of course go too far—what if my lifebox became so good at emulating me that people preferred its outputs to those of own creaky and aging self?

But I don't, however, see any near-term lifebox as being a living copy of its creator. At this point, my lifebox will just be another work of art, not so different from a bookshelf of collected works or, once again, like a searchable blog.

Looking further ahead, how *would* one go about creating a human-like intelligence? That is, how would we animate a lifebox so as to have an artificial person?

A short answer is that, given that our brains have acquired their inherent structures by the process of evolution, the likeliest method for creating intelligent software is via a simulated process of evolution within the virtual world of a computer. There is, however, a difficulty with simulated evolution — even with the best computers imaginable, it may take an exceedingly long time to bear fruit.

An alternate hope is that there may yet be some fairly simple model of the working of human consciousness which we can model and implement in the coming decades. The best idea for a model that I've seen is in Jeff Hawkins and Sandra Blakeslee, *On Intelligence*.<sup>vi</sup> Their model describes a directed evolution based upon a rich data base that develops by continually moving to higher-level symbol systems.

For now in any case, it would help the progress of AI to create a number of lifeboxes. It may well be that these constructs can in fact serve as hosts or culture mediums where we can develop fully conscious and intelligent minds.

But for now, even without an intelligent spark, a lifebox can be exceedingly lifelike.

At the very least—as Leon Marvell will point out—we've invented a great new medium.

## **How We Got There**

Leon Marvell

*The mind is a fractal hypertext and the self is a looping recursion within this hypertext.*

Can we even trace who originally said this? In a world composed almost entirely of data-noise, the sources of innovation have become inconsequential and our innovators appear as mere specters

haunting the edges of our collective dreams. The notion of the self as a looping recursion within a hypertext-society-of-mind is so pervasive in contemporary society that it is almost a cliché, yet it is my aim to trace the connections between this idea and the lifebox — to unearth the hypertextual pattern within the ideo-technological network that gave rise to the lifebox.

Only a few years before Rudy Rucker delivered the preceding paper, another writer of speculative fiction, Robert J. Sawyer, had written a novel called *Mindscan*<sup>vii</sup> in which, inspired by the speculations of Ray Kurzweil's *The Age of Spiritual Machines* (1999), he envisioned individuals being able to upload their psyches into artificially produced bodies: "...the locations, interconnections, and contents of all the somas, axons, dendrites, pre-synaptic vesicles, neurotransmitter concentrations, and other neural components and levels" would be exactly replicated such that the "entire organisation can then be re-created on a neural computer of sufficient capacity, including the contents of its memory."<sup>viii</sup>

In the process imagined by Sawyer a quantum fog is injected into the cerebrum of the person wishing to have their consciousness replicated. A 4-D "snap-shot" of the individual's psyche is captured within the quantum effluvium, and this psychical pattern is instantaneously transferred to the artificial body. As the artificial body's braincase is similarly permeated with quantum fog, instant transfer of consciousness occurs through the exploitation of the phenomenon of quantum entanglement. In Sawyer's time mathematician Roland Penrose and medical scientist Stuart Hameroff were proposing that consciousness was an effect of the macroscopic, coherent superposition of quantum states in microtubules within the cerebrum. In effect, the human brain operated as a massively parallel quantum computer and consciousness was a particular outcome of quantum wave-state collapse. It is highly probable that these speculations were the secondary inspiration behind Sawyer's vision of the transfer of consciousness being effected through the exploitation of the quantum entanglement.<sup>ix</sup>

While Penrose and Hameroff's theory was considered both highly theoretical and rather eccentric in the late 20<sup>th</sup> and very early 21<sup>st</sup> century, the furious progression of technological prostheses in the middle to late 21<sup>st</sup> century has at least confirmed the utility of these speculations. We now live in an era in which bio-computers utilising Penzweil architecture capable of zettabyte-and-beyond memory are networked across the globe and inner planets, and individuals enabled with enough G8 Kredits can upload, modify and maintain their lifeboxes in what amounts to a virtual perpetuity. The Afterlife was never easier.

Yet despite the success of Penzweil bio-computation, Rucker was alarmingly prescient in arguing that, "In the end, the synaptic structures and biochemical reactions of a living brain may prove



too delicate to capture from the outside.”<sup>x</sup> As we now know only too well, the No-Cloning theorem was proven correct in the middle of this century and the disastrous consequences of those wayward experiments are so well known, so notorious, that we shall refrain from further consideration of them here.

In the past century no artificial versions of selves have been produced such that they are indistinguishable from their antecedent, original version (as both Sawyer and Rucker had imagined), but it is certainly the case that lifeboxes are now possessed of powerful emulation software — more powerful than Rucker could have imagined in the early 21<sup>st</sup> century — operating within the wetware carapaces of these ubiquitous portable devices. This emulation software is powerful enough to enable networked lifeboxes to create their own communities of complex hive-minds and to undertake the direction of low-level societal computing and urban redesign, as is well known. We await the perhaps unanticipated consequences of these operations with both excitement and not a little trepidation.

Rather than dwelling on the obvious, in this paper I want to draw attention to the ideo-technological history that underpins the present functioning of the lifebox within contemporary society: for here Rucker’s notion that a lifebox without “some radically more powerful software... would just be another work of art...” has proven to be not quite so prescient. The radically more powerful software indeed came into being, with the result that the lifebox has become perhaps the ultimate artistic technology. We now live in an era in which everybody is indeed getting their 15 minutes worth, and Lifebox-Dandyism has been the fad for over a decade, totally supplanting early forms of social networking platforms and becoming the pre-eminent form of non-proximal mediated communication.

The ideo-technological history I will be tracing is inspired by two notions explored in Rucker’s contribution to this paper: *twinking* and the fractal branching structure of the human mind.

Rucker’s concluding remark that “even without an intelligent spark, a lifebox can be exceedingly lifelike” captures the spirit of the successive years of lifebox development and also gestures towards the less obvious history of lifelike technologies. Historical records show that in the mid-1960s a computer program called ELIZA was designed to run a natural language processing emulation called DOCTOR. This was one of the first instances of so-called expert systems, in this case a comparatively simple example of pattern matching software. Despite the primitiveness of the software the program was so successful in its simulation of a specific human interaction scenario, that of a doctor (the computer) and a patient (the software’s interlocutor), that it fooled many people into thinking that they were interacting in a virtual consulting room with an actual doctor. Simply put, many people were convinced that the

computer possessed the "spark of intelligence" that Rucker speaks of.

When individuals were told that the lifelike impression of ELIZA was an illusion produced by emulation software, many desperate patients refused to believe it, thinking that there must have been a cognizing human being hidden within the interaction, not a lifeless program run on a machine. Others smugly opined that the program was the perfect example of how a computer could never be programmed with human-like intelligence, as it merely parroted human speech patterns. Both of these responses derive from a belief that there is something *quintessentially* human that no silicon-based machine could ever successfully emulate or reproduce: that which the ancient philosopher Descartes called the *cogito*, the conscious self, itself a sign of a spiritual substance, the 'soul', that made humans unique.

The progressive refinements of computing technology behind the development of the lifebox — specifically the development of bio-circuitry and the consequent enabling of quantum computation at the micro-cellular level — have led us to reject the idea of this quintessence and consequently reify the notion that, contrary to the Cartesian viewpoint, the mind is a fractal hypertext and the self is a looping recursion within this hypertext. Two "outsider" hypertextual *contexts* of the late 20<sup>th</sup>, early 21<sup>st</sup> century will to be invoked here to further my enquiry into the ideo-technological history of the lifebox: Theodore (Ted) Nelson and Douglas Hofstadter.

In 1965 Ted Nelson in the *Proc. 20th Nat. Conf. Assoc. Computing Machinery* stated, "Let me introduce the word 'hypertext' to mean a body of written or pictorial material interconnected in such a complex way that it could not conveniently be presented or represented on paper." We know that two events inspired this idea: his reading of Vannevar Bush's article from the *Atlantic Monthly* first published in 1945, *As We May Think*, in which Bush conceived of an artificial device that would connect associative trails between texts for archival purposes, and his reading of Samuel Taylor Coleridge's poem *Xanadu* wherein he discovered the image of a huge storehouse of memories in the form of Kublai Khan's Pleasure Dome. These twin encounters can be seen as the initiatory factors behind his life's work.

Nelson's conception of the hypertextual was of a system of non-linear, non-synchronic inscription that allowed one to constellate meanings in localized, open clusters, following one's own associative trails rather than passively allowing those imposed by the source documents. Nelson invented a neologism for this process that would illuminate the interconnectedness of ideas, a word that revealed the relations between science, music, literature, visual arts and the moving image. For Ted Nelson everything was (and is) "deeply intertwined."

Following his initial insight, in the late 1960s Nelson spent time at Brown University in Providence, Rhode Island, helping to build a hypertext system. By the early years of the 21<sup>st</sup> century however he had come to regret that formative involvement:

That project dumbed down hypertext to one-way, embedded, non-overlapping links. Its broken and deficient model of hypertext became by turns the structure of the NoteCards and HyperCard programs, the World Wide Web, and XML.<sup>xi</sup>

Nelson realised that primitive systems such as the World Wide Web and XML coding were instances of striated data patterns that served only to support the Commissars' reassertion of their hegemony. Of XML coding in particular he noted,

It gratuitously imposes hierarchy and sequence wherever it can, and is very poor at representing overlap, parallel cross-connection, and other vital non-hierarchical media structures that some people do not wish to recognise. I believe humanity went down the wrong path because of that project at Brown.<sup>xii</sup>

His life-long project *Xanadu* was to be a global electronic dissemination system that would have created a vast, labyrinthine library available for all to access — a virtual Library of Alexandria, but without the Pharaoh's flunkies barring your entrance if you were not of the learned elite. In the early years of the 21<sup>st</sup> century this project eventually morphed into Xanadu<sup>®</sup> Space<sup>™</sup>, an attempt to sculpt data streams in a graphical space of three virtual dimensions. Nelson imagined that this would become the *sine qua non* of the virtual social networking spaces that had begun to emerge in and around the year 2003:

Envision social networking done this way: imagine your personal profile as a flying document in space, with thousands of connections streaming off in all directions, where you can spin various wheels or whatever to zoom in and have different parts of the network light up or disappear.<sup>xiii</sup>

In this new hypertextual space one's personal flying profile would take the form of a "live

document", yet it would be a document unlike any document that had preceded Nelson's conception: "My style of hypertext would allow you to create your own mesh of insightful structures in a live document, as you explore. A document is not a file and nor is it necessarily a sequence. It is a structure".<sup>xiv</sup>

Here one is reminded of the sage advice of another writer of speculative fiction from the period we are discussing, William Gibson. He once famously said that if one wanted to write science fiction, then one should write about the next 15 minutes. Metaphorically speaking, the distance between Nelson's dream of a living, hypertextual document that was a constantly morphing personal profile is only 15 minutes away from Rucker's conception of the lifebox. The "document" Nelson imagined became the lifebox of today: not a file, not a directory (the very word 'directory' would no doubt have made Nelson reach for his revolver<sup>xv</sup>) nor even a sequence in time, but rather a spatialised structure that we now popularly refer to as a "context" following the popular dissemination of much of Rudy Rucker's terminology in about the middle of this century.

If there is one single key to unlocking the ideo-technological history I am discussing, then it is this: Nelson's use of the word "structure". For expedience I will deliberately conflate this term with the word "pattern" while simultaneously recalling the words of the *pater familias* of cybernetic theory, Norbert Wiener,

We are not stuff that abides, but patterns that perpetuate themselves.<sup>xvi</sup>

We now know that the 'stuff' mentioned by Wiener — that is, the classical conception of matter — is in actuality a pattern with an event structure. Rocks and glaciers are recursive events, it is just that they are very slow events. And while it is now generally conceded that selves are evanescent looping recursions within a fractal hypertextual space, it may easily be forgotten that this was not always the case.

As far as one can ascertain, the proximal hypertext, the *context*, for these notions is Douglas Hofstadter. In the early 21<sup>st</sup> century Hofstadter conceived the idea that the self is a recursive loop, a strange loop, as he called it.<sup>xvii</sup> To provide a picture of what he intended by this, he often said that a self is rather like a smile. A smile isn't a thing, it's not composed of 'stuff', it's a pattern. He formulated his idea somewhat like a Zen koan:

So what is this thing called a smile? ...A smile persists for a while, and then vanishes. Where is your smile when it's not on your face? It's a potential. [It's] a pattern — like a whirlpool or a tornado.<sup>xviii</sup>

Furthermore a smile "can exist in different media, on different substrates if you prefer. I see it in the mirror, in photographs. And, again, a bit of it is on my children's faces if they happen to be smiling. So if someone asks: "Your smile yesterday and your smile today: which one is the 'real' smile?" I'd reply: neither, both are genuine, my smile comes in multiple instances."<sup>xix</sup>

The recursive instances that we call the self are reflexive, fractal patterns that include an image of themselves. It is this self-image that is the key to Hofstadter's phrase "strange loop". Strange loops occur when, traversing any hierarchical system, one eventually discovers that one is back at the beginning again. Selves are thus irreducibly self-reflexive, paradoxical, and ironical—in a funny kind of way, of course.<sup>xx</sup>

By insisting that a self, like a human smile, is a recursive event pattern, Hofstadter sought to "get across that "I" can exist in multiple spots in the world, that it can flicker in and out of existence the way a smile can."<sup>xxi</sup> Even so, "A person's smile changes over a lifetime, from childhood to old age. Yet people may say: "I still see the same smile I could see 50 years ago."<sup>xxii</sup>

If we consider the foregoing in light of what one might call the 'translation problem' raised by Rucker in his paper on the lifebox re-presented here, Hofstadter provides another angle that allows us to understand how the problem was eventually solved. He proposes that when a novel is translated from one language to another, nobody gets really upset and yells that the translation is a lie — despite the fact that not a single word of the original language remains. He notes that this is because a novel is not only a sequence of words, it is a complex pattern comprised of characters, events, places, cultures and literary style:

And one essential in preserving its identity across media or languages, in deciding whether a translation really is *Eugene Onegin*, for example, is the "grain size", the resolution. A summary isn't a novel, it's too coarse.<sup>xxiii</sup>

Thus we can see how the translation problem was early connected to the resolution problem. Now that we are familiar with the notion that selves are a form of emulation code — now that we are

living in an actual world that was at one time only a fictional world inhabited by Frank Shook and his friends — we recognize that we are all high-order twinklers who daily utilize a comparatively low-order twinker, the lifebox. Yet with each successive iteration of the wetware we have seen the *context* thus produced become more and more lifelike, so much so that many of the *cognoscenti* amongst the Lifebox-Dandy set often participate in events solely through the agency of their lifebox — or even relegate the task of producing artworks to their devices. The finer the grain-size, the more lifelike the lifebox has become.

Yet as our century as progressed it has become clear that there may well be a catastrophic bifurcation point beyond which the emulation of a self cannot proceed, a level of resolution, of grain, that is impossible to achieve. What is this mysterious point? Some have suggested that the quantum wave-state collapse represents this point. But perhaps the wave state-collapse signals a tipping-point, a threshold to another level of understanding. No doubt we might never know the answer to this question, and perhaps it is nonsensical to ask it in the first place. One might conclude however that at this late 21<sup>st</sup> century juncture the lifebox may have to doff its hat to the masters of old, to the painters and sculptors and multimedia artists of the centuries before and including the early years of our own.

When we stand before the dendritic paintings of Jackson Pollock for example, the chaotic surfaces emulating the microtomed sections of his brain, the chance neural firings of feet, heart, blood, hands and brush are there forever etched into the evanescent matter of the canvas, and no matter at which resolution we set our forensic micrometers, we will never capture this complexity and will apprehend the very being of Pollock always ahead of us, travelling ever onwards towards the infinite.

---

<sup>i</sup> This material is adapted and expanded from a section of Rudy Rucker, *The Lifebox, the Seashell and the Soul*, (Basic Books, New York, 1995).

<sup>ii</sup> Rudy Rucker, *Software*, (Ace Books, New York 1982), p. 36. Reprinted in *The Ware Tetralogy* (Prime Books, Rockville, Maryland 2010). In quantum information theory there's a quite different kind of discussion concerning whether it would be possible to precisely copy any physical system such as a brain. The so-called No-Cloning Theorem indicates that you can't precisely replicate a system's quantum state without destroying the system. If you had a quantum-state replicator, you'd need to destroy a brain in order to get a quantum-precise copy of it. This said, it's quite possible that you could create a behaviorally identical copy of a brain without having to actually copy *all* of the quantum states

involved.

<sup>iii</sup> I first used the word in a short story, "Soft Death" (*The Magazine of Fantasy and Science Fiction*, September, 1986).

<sup>iv</sup> Rudy Rucker, *Saucer Wisdom*, (Tor Books, New York 1999) pp. 57 - 59.

<sup>v</sup> Thomas Pynchon, *Gravity's Rainbow*, (Viking Press, New York 1973) p. 516.

<sup>vi</sup> Jeff Hawkins and Sandra Blakeslee, *On Intelligence*, (Times Books, New York 2004).

<sup>vii</sup> Robert J. Sawyer, *Mindscan* (Tor Books, New York 2005).

<sup>viii</sup> *ibid*, page 43

<sup>ix</sup> But see Rucker's note (ii) above.

<sup>x</sup> *Vide* the first half of this paper.

<sup>xi</sup> "Lost in hyperspace", *New Scientist*, issue 2561, 22 July 2006, page 26.

<sup>xii</sup> *ibid*.

<sup>xiii</sup> "Living online: The internet could be so much better", *New Scientist*, issue 2569, 16 September 2006, page 55.

<sup>xiv</sup> *New Scientist* magazine, issue 2569, 16 September 2006, page 55.

<sup>xv</sup> We are not really sure what the antique term 'revolver' actually means; yet we find the locution rather quaint and perhaps of philological/historical interest.

<sup>xvi</sup> Norbert Wiener, *The Human Use of Human Beings* (Avon Books, New York 1967), page 130.

<sup>xvii</sup> Douglas Hofstadter, *I am a Strange Loop*, (Basic Books, New York 2007).

<sup>xviii</sup> "In the end, we are all part of one another", *New Scientist*, issue 2594, 10 March 2007, page 46-48.

<sup>xix</sup> *ibid*.

<sup>xx</sup> According to Hofstadter's memoirs, at least.

<sup>xxi</sup> "In the end, we are all part of one another", *New Scientist*, issue 2594, 10 March 2007, page 46-48.

<sup>xxii</sup> *ibid*.

<sup>xxiii</sup> *ibid*