YOU ARE NOT A GADGET — Jaron Lanier

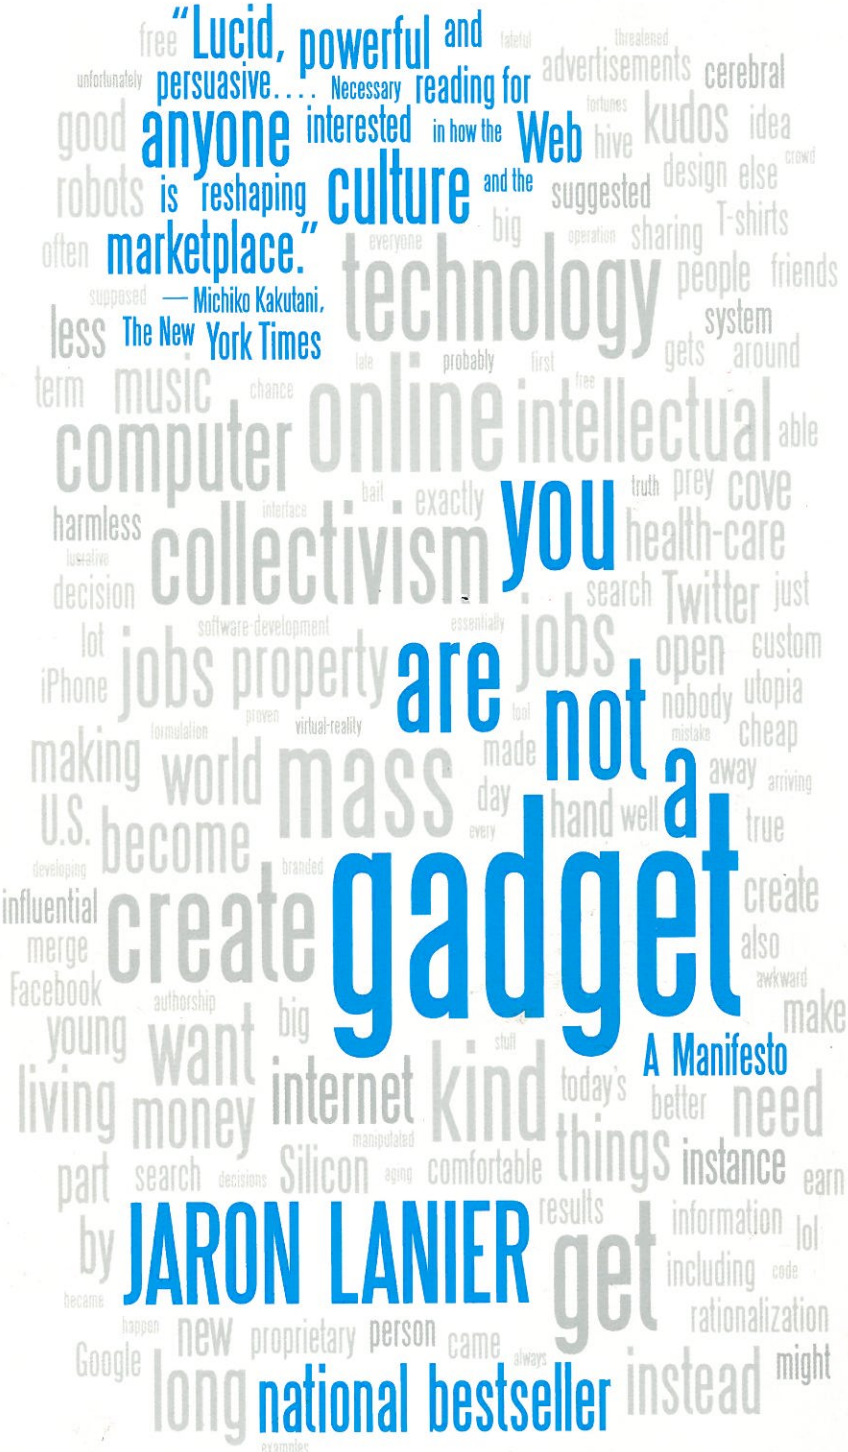

From my point of view, this type of design feature is nonsense, since you end up having to work more than you would otherwise in order to manipulate the software's expectations of you. The real function of the feature isn't to make life easier for people. Instead, it promotes a new philosophy: that the computer is evolving into a life-form that can understand people better than people can understand themselves.

Another example is what I call the "race to be most meta." If a design like Facebook or Twitter depersonalizes people a little bit, then another service like Friendfeed—which may not even exist by the time this book is published—might soon come along to aggregate the previous layers of aggregation, making individual people even more abstract, and the illusion of high-level metaness more celebrated.

Information Doesn't Deserve to Be Free

"Information wants to be free." So goes the saying. Stewart Brand, the founder of the *Whole Earth Catalog*, seems to have said it first.

I say that information doesn't deserve to be free.

Cybernetic totalists love to think of the stuff as if it were alive and had its own ideas and ambitions. But what if information is inanimate? What if it's even less than inanimate, a mere artifact of human thought? What if only humans are real, and information is not?

Of course, there is a technical use of the term "information" that refers to something entirely real. This is the kind of information that's related to entropy. But that fundamental kind of information, which exists independently of the culture of an observer, is not the same as the kind we can put in computers, the kind that supposedly wants to be free.

Information is alienated experience.

You can think of culturally decodable information as a potential form of experience, very much as you can think of a brick resting on a ledge as storing potential energy. When the brick is prodded to fall, the energy is revealed. That is only possible because it was lifted into place at some point in the past.

In the same way, stored information might cause experience to be revealed if it is prodded in the right way. A file on a hard disk does indeed contain information of the kind that objectively exists. The fact that the

bits are discernible instead of being scrambled into mush—the way heat scrambles things—is what makes them bits.

But if the bits can potentially mean something to someone, they can only do so if they are experienced. When that happens, a commonality of culture is enacted between the storer and the retriever of the bits. Experience is the only process that can de-alienate information.

Information of the kind that purportedly wants to be free is nothing but a shadow of our own minds, and wants nothing on its own. It will not suffer if it doesn't get what it wants.

But if you want to make the transition from the old religion, where you hope God will give you an afterlife, to the new religion, where you hope to become immortal by getting uploaded into a computer, then you have to believe information is real and alive. So for you, it will be important to redesign human institutions like art, the economy, and the law to reinforce the perception that information is alive. You demand that the rest of us live in your new conception of a state religion. You need us to deify information to reinforce your faith.

The Apple Falls Again

It's a mistake with a remarkable origin. Alan Turing articulated it, just before his suicide.

Turing's suicide is a touchy subject in computer science circles. There's an aversion to talking about it much, because we don't want our founding father to seem like a tabloid celebrity, and we don't want his memory trivialized by the sensational aspects of his death.

The legacy of Turing the mathematician rises above any possible sensationalism. His contributions were supremely elegant and foundational. He gifted us with wild leaps of invention, including much of the mathematical underpinnings of digital computation. The highest award in computer science, our Nobel Prize, is named in his honor.

Turing the cultural figure must be acknowledged, however. The first thing to understand is that he was one of the great heroes of World War II. He was the first "cracker," a person who uses computers to defeat an enemy's security measures. He applied one of the first computers to break a Nazi secret code, called Enigma, which Nazi mathematicians

had believed was unbreakable. Enigma was decoded by the Nazis in the field using a mechanical device about the size of a cigar box. Turing reconceived it as a pattern of bits that could be analyzed in a computer, and cracked it wide open. Who knows what world we would be living in today if Turing had not succeeded?

The second thing to know about Turing is that he was gay at a time when it was illegal to be gay. British authorities, thinking they were doing the most compassionate thing, coerced him into a quack medical treatment that was supposed to correct his homosexuality. It consisted, bizarrely, of massive infusions of female hormones.

In order to understand how someone could have come up with that plan, you have to remember that before computers came along, the steam engine was a preferred metaphor for understanding human nature. All that sexual pressure was building up and causing the machine to malfunction, so the opposite essence, the female kind, ought to balance it out and reduce the pressure. This story should serve as a cautionary tale. The common use of computers, as we understand them today, as sources for models and metaphors of ourselves is probably about as reliable as the use of the steam engine was back then.

Turing developed breasts and other female characteristics and became terribly depressed. He committed suicide by lacing an apple with cyanide in his lab and eating it. Shortly before his death, he presented the world with a spiritual idea, which must be evaluated separately from his technical achievements. This is the famous Turing test. It is extremely rare for a genuinely new spiritual idea to appear, and it is yet another example of Turing's genius that he came up with one.

Turing presented his new offering in the form of a thought experiment, based on a popular Victorian parlor game. A man and a woman hide, and a judge is asked to determine which is which by relying only on the texts of notes passed back and forth.

Turing replaced the woman with a computer. Can the judge tell which is the man? If not, is the computer conscious? Intelligent? Does it deserve equal rights?

It's impossible for us to know what role the torture Turing was enduring at the time played in his formulation of the test. But it is undeniable that one of the key figures in the defeat of fascism was destroyed, by our

side, after the war, because he was gay. No wonder his imagination pondered the rights of strange creatures.

When Turing died, software was still in such an early state that no one knew what a mess it would inevitably become as it grew. Turing imagined a pristine, crystalline form of existence in the digital realm, and I can imagine it might have been a comfort to imagine a form of life apart from the torments of the body and the politics of sexuality. It's notable that it is the woman who is replaced by the computer, and that Turing's suicide echoes Eve's fall.

The Turing Test Cuts Both Ways

Whatever the motivation, Turing authored the first trope to support the idea that bits can be alive on their own, independent of human observers. This idea has since appeared in a thousand guises, from artificial intelligence to the hive mind, not to mention many overhyped Silicon Valley start-ups.

It seems to me, however, that the Turing test has been poorly interpreted by generations of technologists. It is usually presented to support the idea that machines can attain whatever quality it is that gives people consciousness. After all, if a machine fooled you into believing it was conscious, it would be bigoted for you to still claim it was not.

What the test really tells us, however, even if it's not necessarily what Turing hoped it would say, is that machine intelligence can only be known in a relative sense, in the eyes of a human beholder.*

The AI way of thinking is central to the ideas I'm criticizing in this

*One extension of the tragedy of Turing's death is that he didn't live long enough to articulate all that he probably would have about his own point of view on the Turing test.

Historian George Dyson suggests that Turing might have sided *against* the cybernetic totalists. For instance, here is an excerpt from a paper Turing wrote in 1939, titled "Systems of Logic Based on Ordinals": "We have been trying to see how far it is possible to eliminate intuition, and leave only ingenuity. We do not mind how much ingenuity is required, and therefore assume it to be available in unlimited supply." The implication seems to be that we are wrong to imagine that ingenuity can be infinite, even with computing clouds, so therefore intuition will never be made obsolete.

Turing's 1950 paper on the test includes this extraordinary passage: "In attempting to construct such machines we should not be irreverently usurping His power of creating souls, any more than we are in the procreation of children: rather we are, in either case, instruments of His will providing mansions for the souls that He creates."

book. If a machine can be conscious, then the computing cloud is going to be a better and far more capacious consciousness than is found in an individual person. If you believe this, then working for the benefit of the cloud over individual people puts you on the side of the angels.

But the Turing test cuts both ways. You can't tell if a machine has gotten smarter or if you've just lowered your own standards of intelligence to such a degree that the machine seems smart. If you can have a conversation with a simulated person presented by an AI program, can you tell how far you've let your sense of personhood degrade in order to make the illusion work for you?

People degrade themselves in order to make machines seem smart all the time. Before the crash, bankers believed in supposedly intelligent algorithms that could calculate credit risks before making bad loans. We ask teachers to teach to standardized tests so a student will look good to an algorithm. We have repeatedly demonstrated our species' bottomless ability to lower our standards to make information technology look good. Every instance of intelligence in a machine is ambiguous.

The same ambiguity that motivated dubious academic AI projects in the past has been repackaged as mass culture today. Did that search engine really know what you want, or are you playing along, lowering your standards to make it seem clever? While it's to be expected that the human perspective will be changed by encounters with profound new technologies, the exercise of treating machine intelligence as real requires people to reduce their mooring to reality.

A significant number of AI enthusiasts, after a protracted period of failed experiments in tasks like understanding natural language, eventually found consolation in the adoration for the hive mind, which yields better results because there are real people behind the curtain.

Wikipedia, for instance, works on what I call the Oracle illusion, in which knowledge of the human authorship of a text is suppressed in order to give the text superhuman validity. Traditional holy books work in precisely the same way and present many of the same problems.

This is another of the reasons I sometimes think of cybernetic totalist culture as a new religion. The designation is much more than an approximate metaphor, since it includes a new kind of quest for an afterlife. It's so weird to me that Ray Kurzweil wants the global computing cloud to

scoop up the contents of our brains so we can live forever in virtual reality. When my friends and I built the first virtual reality machines, the whole point was to make this world more creative, expressive, empathic, and interesting. It was not to escape it.

A parade of supposedly distinct "big ideas" that amount to the worship of the illusions of bits has enthralled Silicon Valley, Wall Street, and other centers of power. It might be Wikipedia or simulated people on the other end of the phone line. But really we are just hearing Turing's mistake repeated over and over.

Or Consider Chess

Will trendy cloud-based economics, science, or cultural processes outpace old-fashioned approaches that demand human understanding? No, because it is only encounters with human understanding that allow the contents of the cloud to exist.

Fragment liberation culture breathlessly awaits future triumphs of technology that will bring about the Singularity or other imaginary events. But there are already a few examples of how the Turing test has been approximately passed, and has reduced personhood. Chess is one.

The game of chess possesses a rare combination of qualities: it is easy to understand the rules, but it is hard to play well; and, most important, the urge to master it seems timeless. Human players achieve ever higher levels of skill, yet no one will claim that the quest is over.

Computers and chess share a common ancestry. Both originated as tools of war. Chess began as a battle simulation, a mental martial art. The design of chess reverberates even further into the past than that—all the way back to our sad animal ancestry of pecking orders and competing clans.

Likewise, modern computers were developed to guide missiles and break secret military codes. Chess and computers are both direct descendants of the violence that drives evolution in the natural world, however sanitized and abstracted they may be in the context of civilization. The drive to compete is palpable in both computer science and chess, and when they are brought together, adrenaline flows.

What makes chess fascinating to computer scientists is precisely that