

EdgeThe Third Culture

[Home](#)[About Edge](#)[Features](#)[Edge Editions](#)[Press](#)[Edge Search](#)[Subscribe](#)

The hive mind is for the most part stupid and boring. Why pay attention to it?

The problem is in the way the Wikipedia has come to be regarded and used; how it's been elevated to such importance so quickly. And that is part of the larger pattern of the appeal of a new online collectivism that is nothing less than a resurgence of the idea that the collective is all-wise, that it is desirable to have influence concentrated in a bottleneck that can channel the collective with the most verity and force. This is different from representative democracy, or meritocracy. This idea has had dreadful consequences when thrust upon us from the extreme Right or the extreme Left in various historical periods. The fact that it's now being re-introduced today by prominent technologists and futurists, people who in many cases I know and like, doesn't make it any less dangerous.

DIGITAL MAOISM:

The Hazards of the New Online Collectivism [5.30.06]

By Jaron Lanier

An Edge Original Essay



Introduction

In "Digital Maosim", an original essay written for *Edge*, computer scientist and digital visionary Jaron Lanier finds fault with what he terms the new online collectivism. He cites as an example the Wikipedia, noting that "reading a Wikipedia entry is like reading the bible closely. There are faint traces of the voices of various anonymous authors and editors, though it is impossible to be sure".

His problem is not with the unfolding experiment of the Wikipedia itself, but "the way the Wikipedia has come to be regarded and used; how it's been elevated to such importance so quickly. And that is part of the larger pattern of the appeal of a new online collectivism that is nothing less than a

resurgence of the idea that the collective is all-wise, that it is desirable to have influence concentrated in a bottleneck that can channel the collective with the most verity and force. This is different from representative democracy, or meritocracy. This idea has had dreadful consequences when thrust upon us from the extreme Right or the extreme Left in various historical periods. The fact that it's now being re-introduced today by prominent technologists and futurists, people who in many cases I know and like, doesn't make it any less dangerous".

And he notes that "the Wikipedia is far from being the only online fetish site for foolish collectivism. There's a frantic race taking place online to become the most "Meta" site, to be the highest level aggregator, subsuming the identity of all other sites".

Where is this leading? Lanier calls attention to the "so-called 'Artificial Intelligence' and the race to erase personality and be most Meta. In each case, there's a presumption that something like a distinct kin to individual human intelligence is either about to appear any minute, or has already appeared. The problem with that presumption is that people are all too willing to lower standards in order to make the purported newcomer appear smart. Just as people are willing to bend over backwards and make themselves stupid in order to make an AI interface appear smart (as happens when someone can interact with the notorious Microsoft paper clip,) so are they willing to become uncritical and dim in order to make Meta-aggregator sites appear to be coherent."

Read on as Jaron Lanier throws a lit Molotov cocktail down towards Palo Alto from up in the Berkeley Hills...

—JB

THE REALITY CLUB

THE REALITY CLUB: Responses to Lanier's essay from Douglas Rushkoff, Quentin Hardy, Yochai Benkler, Clay Shirky, Cory Doctorow, Kevin Kelly, Esther Dyson, Larry Sanger, Fernanda Viegas & Martin Wattenberg, Jimmy Wales, George Dyson, Dan Gillmor, Howard Rheingold

PRESS: [Slashdot](#), [Sueddeutsche Zeitung](#), [perlentaucher.de](#), [Die Welt](#), [Inside Higher Ed](#), [Sueddeutsche Zeitung](#), [Arte & Letters Daily](#), [BoingBoing](#), [New York Times](#), [Open Democracy](#), [The Guardian](#) [[click here](#)]

DIGITAL MAOISM

(JARON LANIER:) My Wikipedia entry identifies me (at least this week) as a film director. It is true I made one experimental short film about a decade and a half ago. The concept was awful: I tried to imagine what Maya Deren would have done with morphing. It was shown once at a film festival and was never distributed and I would be most comfortable if no one ever sees it again.

In the real world it is easy to not direct films. I have attempted to retire from directing films in the alternative universe that is the Wikipedia a number of times, but somebody always overrules me. Every time my Wikipedia entry is corrected, within a day I'm turned into a film director again. I can think of no more suitable punishment than making these determined Wikipedia goblins actually watch my one small old movie.

Twice in the past several weeks, reporters have asked me about my filmmaking career. The fantasies of the goblins have entered that portion of the world that is attempting to remain real. I know I've gotten off easy. The errors in my Wikipedia bio have been (at least prior to the publication of this article) charming and even flattering.

Reading a Wikipedia entry is like reading the bible closely. There are faint traces of the voices of various anonymous authors and editors, though it is impossible to be sure. In my particular case, it appears that the goblins are probably members or descendants of the rather sweet old *Mondo 2000* culture linking psychedelic experimentation with computers. They seem to place great importance on relating my ideas to those of the psychedelic luminaries of old (and in ways that I happen to find sloppy and incorrect.) Edits deviating from this set of odd ideas that are important to this one particular small subculture are immediately removed. This makes sense. Who else would volunteer to pay that much attention and do all that work?

The problem I am concerned with here is not the Wikipedia in itself. It's been criticized quite a lot, especially in the last year, but the Wikipedia is just one experiment that still has room to change and grow. At the very least it's a success at revealing what the online people with the most determination and time on their hands are thinking, and that's actually interesting information.

No, the problem is in the way the Wikipedia has come to be regarded and used; how it's been elevated to such importance so quickly. And that is part of the larger pattern of the appeal of a new online collectivism that is nothing less than a resurgence of the idea that the collective is all-wise, that it is desirable to have influence concentrated in a bottleneck that can channel the collective with the most verity and force. This is different from representative democracy, or meritocracy. This idea has had dreadful consequences when thrust upon us from the extreme Right or the extreme Left in various historical periods. The fact that it's now being re-introduced today by prominent technologists and futurists, people who in many cases I know and like, doesn't make it any less dangerous.

There was a well-publicized study in *Nature* last year comparing the accuracy of the Wikipedia to *Encyclopedia Britannica*. The results were a toss up, while there is a lingering debate about the validity of the study. The items selected for the comparison were just the sort that Wikipedia would do well on: Science topics that the collective at large doesn't care much about. "Kinetic isotope effect" or "Vesalius, Andreas" are examples of topics that make the *Britannica* hard to maintain, because it takes work to find the right authors to research and review a multitude of diverse topics. But they are perfect for the Wikipedia. There is little controversy around these items, plus the Net

provides ready access to a reasonably small number of competent specialist graduate student types possessing the manic motivation of youth.

A core belief of the wiki world is that whatever problems exist in the wiki will be incrementally corrected as the process unfolds. This is analogous to the claims of Hyper-Libertarians who put infinite faith in a free market, or the Hyper-Lefties who are somehow able to sit through consensus decision-making processes. In all these cases, it seems to me that empirical evidence has yielded mixed results. Sometimes loosely structured collective activities yield continuous improvements and sometimes they don't. Often we don't live long enough to find out. Later in this essay I'll point out what constraints make a collective smart. But first, it's important to not lose sight of values just because the question of whether a collective can be smart is so fascinating. Accuracy in a text is not enough. A desirable text is more than a collection of accurate references. It is also an expression of personality.

For instance, most of the technical or scientific information that is in the Wikipedia was already on the Web before the Wikipedia was started. You could always use Google or other search services to find information about items that are now wikified. In some cases I have noticed specific texts get cloned from original sites at universities or labs onto wiki pages. And when that happens, each text loses part of its value. Since search engines are now more likely to point you to the wikified versions, the Web has lost some of its flavor in casual use.

When you see the context in which something was written and you know who the author was beyond just a name, you learn so much more than when you find the same text placed in the anonymous, faux-authoritative, anti-contextual brew of the Wikipedia. The question isn't just one of authentication and accountability, though those are important, but something more subtle. A voice should be sensed as a whole. You have to have a chance to sense personality in order for language to have its full meaning. Personal Web pages do that, as do journals and books. Even *Britannica* has an editorial voice, which some people have criticized as being vaguely too "Dead White Men."

If an ironic Web site devoted to destroying cinema claimed that I was a filmmaker, it would suddenly make sense. That would be an authentic piece of text. But placed out of context in the Wikipedia, it becomes drivel.

Myspace is another recent experiment that has become even more influential than the Wikipedia. Like the Wikipedia, it adds just a little to the powers already present on the Web in order to inspire a dramatic shift in use. Myspace is all about authorship, but it doesn't pretend to be all-wise. You can always tell at least a little about the character of the person who made a Myspace page. But it is very rare indeed that a Myspace page inspires even the slightest confidence that the author is a trustworthy authority. Hurray for Myspace on that count!

Myspace is a richer, multi-layered, source of information than the Wikipedia, although the topics the two services cover barely overlap. If you want to research a TV show in terms of what people think of it, Myspace will reveal more to you than the analogous and enormous entries in the Wikipedia.

The Wikipedia is far from being the only online fetish site for foolish collectivism. There's a frantic race taking place online to become the most "Meta" site, to be the highest level aggregator, subsuming the identity of all other sites.

The race began innocently enough with the notion of creating directories of online destinations, such as the early incarnations of Yahoo. Then came AltaVista, where one could search using an inverted database of the content of the whole Web. Then came Google, which added page rank algorithms. Then came the blogs, which varied greatly in terms of quality and importance. This led to Meta-blogs such as Boing Boing, run by identified humans, which served to aggregate blogs. In all of these formulations, real people were still in charge. An individual or individuals were presenting a personality and taking responsibility.

These Web-based designs assumed that value would flow from people. It was still clear, in all such designs, that the Web was made of people, and that ultimately value always came from connecting with real humans.

Even Google by itself (as it stands today) isn't Meta enough to be a problem. One layer of page ranking is hardly a threat to authorship, but an accumulation of many layers can create a meaningless murk, and that is another matter.

In the last year or two the trend has been to remove the scent of people, so as to come as close as possible to simulating the appearance of content emerging out of the Web as if it were speaking to us as a supernatural oracle. This is where the use of the Internet crosses the line into delusion.

Kevin Kelly, the former editor of *Whole Earth Review* and the founding Executive Editor of *Wired*, is a friend and someone who has been thinking about what he and others call the "Hive Mind." He runs a Website called *Cool Tools* that's a cross between a blog and the old *Whole Earth Catalog*. On *Cool Tools*, the contributors, including me, are not a hive because we are identified.

In March, Kelly reviewed a variety of "Consensus Web filters" such as "Digg" and "Reddit" that assemble material every day from all the myriad of other aggregating sites. Such sites intend to be more Meta than the sites they aggregate. There is no person taking responsibility for what appears on them, only an algorithm. The hope seems to be that the most Meta site will become the mother of all bottlenecks and receive infinite funding.

That new magnitude of Meta-ness lasted only a month. In April, Kelly reviewed a site called "popurls" that aggregates consensus Web filtering sites...and there was a new "most Meta". We now are reading what a collectivity algorithm derives from what other collectivity algorithms derived from what collectives chose from what a population of mostly amateur writers wrote anonymously.

Is "popurls" any good? I am writing this on May 27, 2006. In the last few days an experimental approach to diabetes management has been

announced that might prevent nerve damage. That's huge news for tens of millions of Americans. It is not mentioned on popurls. Popurls does clue us in to this news: "Student sets simultaneous world ice cream-eating record, worst ever ice cream headache." Mainstream news sources all lead today with a serious earthquake in Java. Popurls includes a few mentions of the event, but they are buried within the aggregation of aggregate news sites like Google News. The reason the quake appears on popurls at all can be discovered only if you dig through all the aggregating layers to find the original sources, which are those rare entries actually created by professional writers and editors who sign their names. But at the layer of popurls, the ice cream story and the Javanese earthquake are at best equals, without context or authorship.

Kevin Kelly says of the "popurls" site, "There's no better way to watch the hive mind." But the hive mind is for the most part stupid and boring. Why pay attention to it?

Readers of my previous rants will notice a parallel between my discomfort with so-called "Artificial Intelligence" and the race to erase personality and be most Meta. In each case, there's a presumption that something like a distinct kin to individual human intelligence is either about to appear any minute, or has already appeared. The problem with that presumption is that people are all too willing to lower standards in order to make the purported newcomer appear smart. Just as people are willing to bend over backwards and make themselves stupid in order to make an AI interface appear smart (as happens when someone can interact with the notorious Microsoft paper clip,) so are they willing to become uncritical and dim in order to make Meta-aggregator sites appear to be coherent.

There is a pedagogical connection between the culture of Artificial Intelligence and the strange allure of anonymous collectivism online. Google's vast servers and the Wikipedia are both mentioned frequently as being the startup memory for Artificial Intelligences to come. Larry Page is quoted via a link presented to me by popurls this morning (who knows if it's accurate) as speculating that an AI might appear within Google within a few years. George Dyson has wondered if such an entity already exists on the Net, perhaps perched within Google. My point here is not to argue about the existence of Metaphysical entities, but just to emphasize how premature and dangerous it is to lower the expectations we hold for individual human intellects.

The beauty of the Internet is that it connects people. The value is in the other people. If we start to believe that the Internet itself is an entity that has something to say, we're devaluing those people and making ourselves into idiots.

Compounding the problem is that new business models for people who think and write have not appeared as quickly as we all hoped. Newspapers, for instance, are on the whole facing a grim decline as the Internet takes over the feeding of curious eyes that hover over morning coffee and even worse, classified ads. In the new environment, Google News is for the moment better funded and enjoys a more secure future than most of the rather small

number of fine reporters around the world who ultimately create most of its content. The aggregator is richer than the aggregated.

The question of new business models for content creators on the Internet is a profound and difficult topic in itself, but it must at least be pointed out that writing professionally and well takes time and that most authors need to be paid to take that time. In this regard, blogging is not writing. For example, it's easy to be loved as a blogger. All you have to do is play to the crowd. Or you can flame the crowd to get attention. Nothing is wrong with either of those activities. What I think of as real writing, however, writing meant to last, is something else. It involves articulating a perspective that is not just reactive to yesterday's moves in a conversation.

The artificial elevation of all things Meta is not confined to online culture. It is having a profound influence on how decisions are made in America.

What we are witnessing today is the alarming rise of the fallacy of the infallible collective. Numerous elite organizations have been swept off their feet by the idea. They are inspired by the rise of the Wikipedia, by the wealth of Google, and by the rush of entrepreneurs to be the most Meta. Government agencies, top corporate planning departments, and major universities have all gotten the bug.

As a consultant, I used to be asked to test an idea or propose a new one to solve a problem. In the last couple of years I've often been asked to work quite differently. You might find me and the other consultants filling out survey forms or tweaking edits to a collective essay. I'm saying and doing much less than I used to, even though I'm still being paid the same amount. Maybe I shouldn't complain, but the actions of big institutions do matter, and it's time to speak out against the collectivity fad that is upon us.

It's not hard to see why the fallacy of collectivism has become so popular in big organizations: If the principle is correct, then individuals should not be required to take on risks or responsibilities. We live in times of tremendous uncertainties coupled with infinite liability phobia, and we must function within institutions that are loyal to no executive, much less to any lower level member. Every individual who is afraid to say the wrong thing within his or her organization is safer when hiding behind a wiki or some other Meta aggregation ritual.

I've participated in a number of elite, well-paid wikis and Meta-surveys lately and have had a chance to observe the results. I have even been part of a wiki about wikis. What I've seen is a loss of insight and subtlety, a disregard for the nuances of considered opinions, and an increased tendency to enshrine the official or normative beliefs of an organization. Why isn't everyone screaming about the recent epidemic of inappropriate uses of the collective? It seems to me the reason is that bad old ideas look confusingly fresh when they are packaged as technology.

The collective rises around us in multifarious ways. What afflicts big institutions also afflicts pop culture. For instance, it has become notoriously difficult to introduce a new pop star in the music business. Even the most

successful entrants have hardly ever made it past the first album in the last decade or so. The exception is American Idol. As with the Wikipedia, there's nothing wrong with it. The problem is its centrality.

More people appear to vote in this pop competition than in presidential elections, and one reason for this is the instant convenience of information technology. The collective can vote by phone or by texting, and some vote more than once. The collective is flattered and it responds. The winners are likable, almost by definition.

But John Lennon wouldn't have won. He wouldn't have made it to the finals. Or if he had, he would have ended up a different sort of person and artist. The same could be said about Jimi Hendrix, Elvis, Joni Mitchell, Duke Ellington, David Byrne, Grandmaster Flash, Bob Dylan (please!), and almost anyone else who has been vastly influential in creating pop music.

As below, so above. *The New York Times*, of all places, has recently published op-ed pieces supporting the pseudo-idea of intelligent design. This is astonishing. *The Times* has become the paper of averaging opinions. Something is lost when American Idol becomes a leader instead of a follower of pop music. But when intelligent design shares the stage with real science in the paper of record, everything is lost.

How could the *Times* have fallen so far? I don't know, but I would imagine the process was similar to what I've seen in the consulting world of late. It's safer to be the aggregator of the collective. You get to include all sorts of material without committing to anything. You can be superficially interesting without having to worry about the possibility of being wrong.

Except when intelligent thought really matters. In that case the average idea can be quite wrong, and only the best ideas have lasting value. Science is like that.

The collective isn't always stupid. In some special cases the collective can be brilliant. For instance, there's a demonstrative ritual often presented to incoming students at business schools. In one version of the ritual, a large jar of jellybeans is placed in the front of a classroom. Each student guesses how many beans there are. While the guesses vary widely, the average is usually accurate to an uncanny degree.

This is an example of the special kind of intelligence offered by a collective. It is that peculiar trait that has been celebrated as the "Wisdom of Crowds," though I think the word "wisdom" is misleading. It is part of what makes Adam Smith's Invisible Hand clever, and is connected to the reasons Google's page rank algorithms work. It was long ago adapted to futurism, where it was known as the Delphi technique. The phenomenon is real, and immensely useful.

But it is not infinitely useful. The collective can be stupid, too. Witness tulip crazes and stock bubbles. Hysteria over fictitious satanic cult child abductions. Y2K mania.

The reason the collective can be valuable is precisely that its peaks of intelligence and stupidity are not the same as the ones usually displayed by individuals. Both kinds of intelligence are essential.

What makes a market work, for instance, is the marriage of collective and individual intelligence. A marketplace can't exist only on the basis of having prices determined by competition. It also needs entrepreneurs to come up with the products that are competing in the first place.

In other words, clever individuals, the heroes of the marketplace, ask the questions which are answered by collective behavior. They put the jellybeans in the jar.

There are certain types of answers that ought not be provided by an individual. When a government bureaucrat sets a price, for instance, the result is often inferior to the answer that would come from a reasonably informed collective that is reasonably free of manipulation or runaway internal resonances. But when a collective designs a product, you get design by committee, which is a derogatory expression for a reason.

Here I must take a moment to comment on Linux and similar efforts. The various formulations of "open" or "free" software are different from the Wikipedia and the race to be most Meta in important ways. Linux programmers are not anonymous and in fact personal glory is part of the motivational engine that keeps such enterprises in motion. But there are similarities, and the lack of a coherent voice or design sensibility in an esthetic sense is one negative quality of both open source software and the Wikipedia.

These movements are at their most efficient while building hidden information plumbing layers, such as Web servers. They are hopeless when it comes to producing fine user interfaces or user experiences. If the code that ran the Wikipedia user interface were as open as the contents of the entries, it would churn itself into impenetrable muck almost immediately. The collective is good at solving problems which demand results that can be evaluated by uncontroversial performance parameters, but it is bad when taste and judgment matter.

Collectives can be just as stupid as any individual, and in important cases, stupider. The interesting question is whether it's possible to map out where the one is smarter than the many.

There is a lot of history to this topic, and varied disciplines have lots to say. Here is a quick pass at where I think the boundary between effective collective thought and nonsense lies: The collective is more likely to be smart when it isn't defining its own questions, when the goodness of an answer can be evaluated by a simple result (such as a single numeric value,) and when the information system which informs the collective is filtered by a quality control mechanism that relies on individuals to a high degree. Under those circumstances, a collective can be smarter than a person. Break any one of those conditions and the collective becomes unreliable or worse.

Meanwhile, an individual best achieves optimal stupidity on those rare occasions when one is both given substantial powers and insulated from the results of his or her actions.

If the above criteria have any merit, then there is an unfortunate convergence. The setup for the most stupid collective is also the setup for the most stupid individuals.

Every authentic example of collective intelligence that I am aware of also shows how that collective was guided or inspired by well-meaning individuals. These people focused the collective and in some cases also corrected for some of the common hive mind failure modes. The balancing of influence between people and collectives is the heart of the design of democracies, scientific communities, and many other long-standing projects. There's a lot of experience out there to work with. A few of these old ideas provide interesting new ways to approach the question of how to best use the hive mind.

The pre-Internet world provides some great examples of how personality-based quality control can improve collective intelligence. For instance, an independent press provides tasty news about politicians by reporters with strong voices and reputations, like the Watergate reporting of Woodward and Bernstein. Other writers provide product reviews, such as Walt Mossberg in *The Wall Street Journal* and David Pogue in *The New York Times*. Such journalists inform the collective's determination of election results and pricing. Without an independent press, composed of heroic voices, the collective becomes stupid and unreliable, as has been demonstrated in many historical instances. (Recent events in America have reflected the weakening of the press, in my opinion.)

Scientific communities likewise achieve quality through a cooperative process that includes checks and balances, and ultimately rests on a foundation of goodwill and "blind" elitism — blind in the sense that ideally anyone can gain entry, but only on the basis of a meritocracy. The tenure system and many other aspects of the academy are designed to support the idea that individual scholars matter, not just the process or the collective.

Another example: Entrepreneurs aren't the only "heroes" of a marketplace. The role of a central bank in an economy is not the same as that of a communist party official in a centrally planned economy. Even though setting an interest rate sounds like the answering of a question, it is really more like the asking of a question. The Fed asks the market to answer the question of how to best optimize for lowering inflation, for instance. While that might not be the question everyone would want to have asked, it is at least coherent.

Yes, there have been plenty of scandals in government, the academy and in the press. No mechanism is perfect, but still here we are, having benefited from all of these institutions. There certainly have been plenty of bad reporters, self-deluded academic scientists, incompetent bureaucrats, and so on. Can the hive mind help keep them in check? The answer provided by experiments in the pre-Internet world is "yes," but only provided some signal

processing is placed in the loop.

Some of the regulating mechanisms for collectives that have been most successful in the pre-Internet world can be understood in part as modulating the time domain. For instance, what if a collective moves too readily and quickly, jittering instead of settling down to provide a single answer? This happens on the most active Wikipedia entries, for example, and has also been seen in some speculation frenzies in open markets.

One service performed by representative democracy is low-pass filtering. Imagine the jittery shifts that would take place if a wiki were put in charge of writing laws. It's a terrifying thing to consider. Super-energized people would be struggling to shift the wording of the tax-code on a frantic, never-ending basis. The Internet would be swamped.

Such chaos can be avoided in the same way it already is, albeit imperfectly, by the slower processes of elections and court proceedings. The calming effect of orderly democracy achieves more than just the smoothing out of peripatetic struggles for consensus. It also reduces the potential for the collective to suddenly jump into an over-excited state when too many rapid changes to answers coincide in such a way that they don't cancel each other out. (Technical readers will recognize familiar principles in signal processing.)

The Wikipedia has recently slapped a crude low pass filter on the jitteriest entries, such as "President George W. Bush." There's now a limit to how often a particular person can remove someone else's text fragments. I suspect that this will eventually have to evolve into an approximate mirror of democracy as it was before the Internet arrived.

The reverse problem can also appear. The hive mind can be on the right track, but moving too slowly. Sometimes collectives would yield brilliant results given enough time but there isn't enough time. A problem like global warming would automatically be addressed eventually if the market had enough time to respond to it, for instance. Insurance rates would climb, and so on. Alas, in this case there isn't enough time, because the market conversation is slowed down by the legacy effect of existing investments. Therefore some other process has to intervene, such as politics invoked by individuals.

Another example of the slow hive problem: There was a lot of technology developed slowly in the millennia before there was a clear idea of how to be empirical, how to have a peer reviewed technical literature and an education based on it, and before there was an efficient market to determine the value of inventions. What is crucial to notice about modernity is that structure and constraints were part of what sped up the process of technological development, not just pure openness and concessions to the collective.

Let's suppose that the Wikipedia will indeed become better in some ways, as is claimed by the faithful, over a period of time. We might still need something better sooner.

Some wikipotians explicitly hope to see education subsumed by wikis. It is at

least possible that in the fairly near future enough communication and education will take place through anonymous Internet aggregation that we could become vulnerable to a sudden dangerous empowering of the hive mind. History has shown us again and again that a hive mind is a cruel idiot when it runs on autopilot. Nasty hive mind outbursts have been flavored Maoist, Fascist, and religious, and these are only a small sampling. I don't see why there couldn't be future social disasters that appear suddenly under the cover of technological utopianism. If wikis are to gain any more influence they ought to be improved by mechanisms like the ones that have worked tolerably well in the pre-Internet world.

The hive mind should be thought of as a tool. Empowering the collective does not empower individuals — just the reverse is true. There can be useful feedback loops set up between individuals and the hive mind, but the hive mind is too chaotic to be fed back into itself.

These are just a few ideas about how to train a potentially dangerous collective and not let it get out of the yard. When there's a problem, you want it to bark but not bite you.

The illusion that what we already have is close to good enough, or that it is alive and will fix itself, is the most dangerous illusion of all. By avoiding that nonsense, it ought to be possible to find a humanistic and practical way to maximize value of the collective on the Web without turning ourselves into idiots. The best guiding principle is to always cherish individuals first.

Jaron Lanier is a film director. He writes a monthly column for Discover Magazine.

[Jaron Lanier's Edge Bio Page](#)

THE REALITY CLUB

On "Digital Maoism: The Hazards of the New Online Collectivism" By Jaron Lanier

Responses to Lanier's essay from Douglas Rushkoff, Quentin Hardy, Yochai Benkler, Clay Shirky, Cory Doctorow, Kevin Kelly, Esther Dyson, Larry Sanger, Fernanda Viegas & Martin Wattenberg, Jimmy Wales, George Dyson, Dan Gillmor, Howard Rheingold

[[Proceed to the discussion...](#)]

John Brockman, Editor and Publisher
Russell Weinberger, Associate Publisher
[contact: editor@edge.org](mailto:editor@edge.org)
Copyright © 2006 By [Edge Foundation, Inc](#)
All Rights Reserved.