

COLLECTIVE INTELLIGENCE

A Conversation with Thomas W. Malone [11.21.12]

As all the people and computers on our planet get more and more closely connected, it's becoming increasingly useful to think of all the people and computers on the planet as a kind of global brain.

THOMAS W. MALONE is the Patrick J. McGovern Professor of Management at the MIT Sloan School of Management and the founding director of the MIT Center for Collective Intelligence. He was also the founding director of the MIT Center for Coordination Science and one of the two founding co-directors of the MIT Initiative on "Inventing the Organizations of the 21st Century".

COLLECTIVE INTELLIGENCE

Pretty much everything I'm doing now falls under the broad umbrella that I'd call collective intelligence. What does collective intelligence mean? It's important to realize that intelligence is not just something that happens inside individual brains. It also arises with groups of individuals. In fact, I'd define collective intelligence as groups of individuals acting collectively in ways that seem intelligent. By that definition, of course, collective intelligence has been around for a very long time. Families, companies, countries, and armies: those are all examples of groups of people working together in ways that at least sometimes seem intelligent.

It's also possible for groups of people to work together in ways that seem pretty stupid, and I think collective stupidity is just as possible as collective intelligence. Part of what I want to understand and part of what the people I'm working with want to understand is what are the conditions that lead to collective intelligence rather than collective stupidity. But in whatever form, either intelligence or stupidity, this collective behavior has existed for a long time.

What's new, though, is a new kind of collective intelligence enabled by the Internet. Think of Google, for instance, where millions of people all over the world create web pages, and link those web pages to each other. Then all that knowledge is harvested by the Google technology so that when you type a question in the Google search bar the answers you get often seem amazingly intelligent, at least by some definition of the word "intelligence."

Or think of Wikipedia, where thousands of people all over the world have collectively created a very large and amazingly high quality intellectual product with almost no centralized control. And by the way, without even being paid. I think these examples of things like Google and Wikipedia are not the end of the story. I think they're just barely the beginning of the story. We're likely to see lots more examples of Internet-enabled collective intelligence—and other kinds of collective intelligence as well—over the coming decades.

If we want to predict what's going to happen, especially if we want to be able to take advantage of what's going to happen, we need to understand those possibilities at a much deeper level than we do so far. That's really our goal in the MIT Center for Collective Intelligence, which I direct. In fact, one way we frame our core research question there is: How can people and computers be connected so that—collectively—they act more intelligently than any person, group or computer has ever done before? If you take that question seriously, the answers you get are often very different from the kinds of organizations and groups we know today.

We do take the question seriously, and we are doing a bunch of things related to that question. The first is just trying to map the different kinds of collective intelligence, the new kinds of collective intelligence that are happening all around us in the world today. One of the projects we've done we call "mapping

the genomes of collective intelligence". We've collected over 200 examples of interesting cases of collective intelligence ... things like Google, Wikipedia, InnoCentive, the community that developed the Linux open source operating system, et cetera.

Then we looked for the design patterns that come up over and over in those different examples. Using the biological analogy, we call these design patterns "genes," but if you don't like the analogy or the metaphor, you can just use the word "design patterns." We've identified so far about 19 of these design patterns—or genes—that occur over and over in these different examples.

For instance, the community of people that developed the Linux open source operating system embodies what we call the "crowd" gene, because anyone who wants to can contribute new modules for the Linux operating system. But that community also embodies what we call the "hierarchy" gene, because Linus Torvalds and a few of his friends and lieutenants decide—essentially hierarchically—which of the modules that people send in will actually be included in the new versions of the system. So that's the genomes of collective intelligence project.

Another thing we are doing is creating new examples of collective intelligence. I'm personally spending the most time on a project to address the problem of global climate change. We've created an online platform and now a community of almost 4,000 people using that platform to come up with proposals for what to do about global climate change. We call this the Climate CoLab and we've had several years of people developing online proposals, and working with other people, in many cases, whom they find on the site to develop these proposals. For the last couple of years, the winners of our annual contests have presented their ideas in briefings at the United Nations in New York and on Capitol Hill in Washington, D.C.

Another project we're doing is one that tries to measure collective intelligence. If we think it's important, as we do, it would sure be nice to be able to measure it. The approach we're taking in this project is one of using the same statistical techniques that are used to measure individual intelligence, but applying those techniques to measure the intelligence of groups. We found that, just as with individuals, there is a single statistical factor that predicts how well a given group will do on a very wide range of different tasks.

Interestingly, when we did this work, we thought that there might be such a factor, but that it would really just be essentially the intelligence of the individual people in the group. What we found was that the average and the maximum intelligence of the individual group members was correlated, but only moderately correlated, with the collective intelligence of the group as a whole.

If it's not just putting a bunch of smart people in a group that makes the group smart, what is it? We looked at bunch of factors you might have thought would affect it: things like the psychological safety of the group, the personality of the group members, et cetera. Most of the things we thought might have affected it turned out not to have any significant effect. But we did find three factors that were significantly correlated with the collective intelligence of the group.

The first was the average social perceptiveness of the group members. We measured social perceptiveness in this case using a test developed essentially to measure autism. It's called the "Reading the Mind and the Eyes Test". It works by letting people look at pictures of other people's eyes and try to guess what emotions those people are feeling. People who are good at that work well in groups. When you have a group with a bunch of people like that, the group as a whole is more intelligent.

The second factor we found was the evenness of conversational turn taking. In other words, groups where one person dominated the conversation were, on average, less intelligent than groups where the speaking was more evenly distributed among the different group members.

Finally, and most surprisingly to us, we found that the collective intelligence of the group was significantly correlated with the percentage of women in the group. More women were correlated with a more intelligent group. Interestingly, this last result is not just a diversity result. It's not just saying that you need groups with some men and some women. It looks like that it's a more or less linear trend. That is, more women are better all the way up to all women. It is also important to realize that this gender effect is largely statistically mediated by the social perceptiveness effect. In other words, it was known before we did our work that women on average scored higher on this measure of social perceptiveness than men.

This is the interpretation I personally prefer: it may be that what's needed to have an intelligent group is just to have a bunch of people in the group who are high on this social perceptiveness measure, whether those people are men or women. In any case, we think it's an interesting finding, one that we hope to understand better and one that already has some very intriguing implications for how we create groups in many cases in the real world.

The way we define intelligence for the purpose of the studies I was just describing to you is essentially the same way psychologists define intelligence at the individual level. It turns out that when you give a bunch of people a bunch of different mental tasks, whether it's solving crossword puzzles or doing mental rotation of figures or doing arithmetic or all kinds of things like that, when you give a bunch of people a bunch of those tasks, it turns out that some people do better on most of them and others do worse on most of them.

In a technical sense, what that says is that there's a single principal component in a factor analysis that explains a significant portion of the variance. It didn't have to be that way. It could have been, for instance, that some people are good at verbal things and some people are good at mathematical things. If you're good at verbal things, you might be worse at math on average, and vice versa. That might have been the case, but it turns out to not be true empirically.

In fact, one of the most well documented and in some ways surprising results in all of psychology is the fact that there is this single factor of intelligence that correlates with an individual's performance across a very wide range of tasks. In fact, even though there are other meanings of the word intelligence in English, a very important element or nuance of the word "intelligence" in English is just that it's somebody who's good at a lot of mental things, somebody who is good at learning quickly, at adapting to new situations, at doing a bunch of things.

The most intelligent person is not the one who's best at doing any specific task, but it's the one who's best at picking up new things quickly. That's essentially the definition we used for defining intelligence at the level of groups as well. We said that a group is intelligent if it's able to perform well on a wide range of different tasks. It was actually performance that we were looking at.

We had a bunch of groups in our case, about 200 groups from two to five people who came into our laboratory and did about five or ten different tasks together as a group. It depended on the different conditions how many exact tasks were done, but the tasks included things like brainstorming how many uses for a brick could you come up with. Or one of the kinds of tasks was doing some problems from an IQ test, but doing those problems as a group rather than as individuals.

Another one was planning a shopping trip with various kinds of constraints, and another one was building buildings out of Lego's according to a fairly complicated set of rules about how you could

combine things. After we let all these different groups do all these different kinds of tasks, then we'd do all the correlations of the different tasks done by different groups, and we determine a way of weighting the scores on those different tasks that maximally predicts how well the groups will do on all of the tasks.

The way we operationally define the group's collective intelligence is as a weighted average of their scores on all these different tasks. Thus, it's a way of predicting how well they'll do on many different kinds of tasks. The project I just described on measuring collective intelligence was done with a set of collaborators at different institutions. The two co-principal investigators were Anita Woolley at Carnegie Mellon and Chris Chabris at Union College, and then we had several other collaborators as well.

In general, all of this is part of a unit at MIT called the MIT Center for Collective Intelligence, of which I am the director. The Center for Collective Intelligence is housed in the Sloan School of Management, but it includes faculty members from many different parts of MIT such as the Media Lab, the Brain and Cognitive Science department and the Computer Science and AI lab as well as the Sloan School of Management. We have faculty members, we have graduate students, some post-docs, some research scientists, a number of people from many different disciplines and many different perspectives all grappling with these questions of what is collective intelligence? How can we measure it? How can we understand it? How can we create it?

I have a pretty unusual background that's led me into doing all of these strange things. My undergraduate degree was called Mathematical Sciences, basically applied math and computer science. In graduate school, I did a master's degree in something called engineering-economic systems, a program at Stanford that no longer exists by that name. My PhD was in cognitive and social psychology, also from Stanford.

If you want me to be a little bit grandiose about this, I'll tell you what I wrote on my college application when I was a senior in high school. I said I wanted to help solve the problems created by technology changing faster than society could adapt. Pretty grandiose thing for a 17 year old to say, and I forgot about it after I had written it. But then years later, I looked back at the degrees and the kinds of things I was doing, and I realized I was trying to do what I had said on my college applications. The technology I wanted to focus on was information technology. How could it not just change things that we have to adapt to, but how could it enable new possibilities that we could take advantage of?

The specific kind of societal problem I focused on primarily was how to organize work, how to organize businesses, other kinds of organizations, even society's efforts to do things. In a sense, I have tried to do what I set out to do as a high school senior. That also is related to one of the most recent projects that I'm focusing on, which is the Climate CoLab, where we're not just talking about how work can be organized, which was the main focus of my book called "The Future of Work," but now focusing on another specific problem created by technology changing faster than society could adapt. Very literally, all the kinds of technologies that cause carbon emissions are changing our planet in ways that nearly all scientists who have looked at the question agree are causing some changes that we need to adapt to, that we haven't done a very good job of adapting to. One of the reasons why I think the Climate CoLab project has particular resonance for me personally is because it's another way of trying to solve the kind of problem I said as a high school senior I wanted to.

Why are we doing all this work? There are at least three answers. The first is, as scientists, we want to understand how the world works, and in particular, how the world of groups of people and computers work together. How human societies and human networks work. Second, we want to help businesses, governments and other kinds of organizations know how to work better themselves. How can we create more intelligent organizations, more intelligent businesses, more intelligent governments, more intelligent societies?

Third, in a way, we are trying to understand how our whole world and society is evolving in a way that I think is making us more collectively intelligent. You could say that the Internet is one way of greatly accelerating the connections among different people and computers on our planet. As all the people and computers on our planet get more and more closely connected, it's becoming increasingly useful to think of all the people and computers on the planet as a kind of global brain.

Our future as a species may depend on our ability to use our global collective intelligence to make choices that are not just smart, but also wise.

I think this may be the way we are evolving, and I think in some kind of deep philosophical sense, one of the ways our work may be helpful is in understanding and I hope accelerating this move toward a more collectively intelligent society, a more collectively intelligent planet.

What's the science here? In a sense, we're trying to understand scientifically how groups of humans work together now using the means we have and have had for connecting humans to each other, face to face communication, telephone, Internet, et cetera. More importantly, perhaps, we're also trying to understand the science behind the deeper phenomena of humans working together or humans and computers working together in ways that will help us understand how to create new kinds of human or human and computer cooperatives or collective intelligences. So in that sense, the boundary between science and engineering begins to blur.

Science is about understanding what is, engineering is about how to create what you want to be. But they're clearly related to each other. Understanding better how the world works helps you shape the world in ways you want it to be, and often trying to shape the world in ways you want helps you understand fundamental scientific questions about how the world is in ways that you might never have thought of asking before. Another way of thinking about the question of what's the science here is to relate what we're doing in collective intelligence.

We just had the first academic conference on collective intelligence in April of 2012 held at MIT. I was one of the two co-organizers. We had a number of very interesting speakers and people there, and many people said it was one of the best conferences they'd ever been to. There's a sense that there is a field catalyzing here, a field congealing here.

One way of asking the question of what's the science here is to relate this to some of the other fields and names that are being talked about today. For instance, there are a lot of people talking about computational social science. One way of distinguishing between computational social science and collective intelligence is to say that computational social science is essentially about methodologies: new ways of answering social science questions that have been with us for some time. We have lots of new kinds of data, for instance, and new ways of analyzing data enabled by computers that help us attack long-term social science questions like how do networks form and evolve and so forth.

Collective intelligence as a field, instead of focusing on a methodology, focuses on a set of questions, a set of phenomenon about those questions. Collective intelligence, as the name implies, is about the phenomenon of intelligence as it arises in groups of individuals—whether those individuals are individual people or whether they are organizations, companies, or markets.

The angle of intelligence provides a useful forcing function for thinking about the scientific questions. You can define intelligence in various ways. We did it in one way with our experiments on measuring group intelligence. You can define it in other ways as having to do with problem solving or perception or memory or learning. But those are all examples of phenomena we associate with the concept of intelligence. Analyzing these phenomena when they happen in groups rather than just inside individual

brains is a very evocative, a very provocative and a very productive way of framing some quite interesting scientific questions.

We had about 20 invited speakers at our conference on collective intelligence. The entire list is at the web site CI2012.org. Some of the speakers include: Yochai Benkler, Jonathan Zittrain from Harvard, Rob Miller from MIT, Bob Kraut from Carnegie Mellon, Anita Woolley and Chris Chabris whom I've mentioned earlier from CMU and Union College, Lada Adamic from University of Michigan, Scott Page from University of Michigan. I'm sure I'm forgetting some, but a number of people who have very interesting things to say about the phenomenon collective intelligence from very different disciplines.

Some of them social scientists, sociologists, social psychologists, some economists, like Colin Camerer and interestingly, biologists like Deborah Gordon from Stanford and Iain Couzin from Princeton.

I started college as a freshman in 1970, and I remember in my sophomore year, I read an article in Scientific American about artificial intelligence by Marvin Minsky from MIT. I was really excited about his vision for artificial intelligence and that was in fact one of the things that sparked my interest in computer science and computers as a technology that was changing the world.

I was worried from reading Marvin's article that if things were going as fast as it seemed from the way he wrote about it, that we might have solved all the problems of artificial intelligence by the time I graduated from college, and there wouldn't be anything left for me to do. I talked to a friend of mine who was an upper classmen, and he reassured me that no, it wouldn't all be solved by 1974, and he was certainly right about that. Of course, many of the problems of artificial intelligence still haven't been solved now, 40 years later.

In fact, I think that the approach we're taking in the Center for Collective Intelligence or generally in the field of collective intelligence is one that has real promise in helping to realize the ultimate ambitions of the field of artificial intelligence without requiring all of the intelligence to be provided by the computers. In other words, there are lots of intelligent things that, of course, humans can do, that we know humans can do. There are some intelligent things that we have now learned to program computers to do, but still many intelligent things that computers can't do yet.

But when you put human intelligence and computer intelligence together and say that it's not cheating to go have people solve part of the problem, that that's really the goal, then I think whole new vistas of ways of creating intelligent systems begin to appear. That's another way of saying what I think is one of the key possibilities of the field of collective intelligence.

Perhaps a closing question for us to ponder is how do these new kinds of intelligence help us understand what it means to be human in the first place, and what our role as humans on the planet is. I think most people would agree that certainly not the only thing, but one of the most important things about humans is their intelligence. We think of ourselves as intelligent beings. We compare ourselves to animals and other inhabitants of this planet and think of ourselves as probably the most intelligent beings on the planet.

Even just to understand that aspect of our humanity, it's useful to think about what intelligence is, what the concept means and how it can occur. This perspective of collective intelligence gives us some deep and powerful ways of doing that. It raises questions, for instance, like how could we recognize intelligence if we saw it? Some biologists have tried to study intelligence in other animals.

How can you tell whether a dog or a cat or an ant is intelligent? How intelligent compared to each other or to humans? In the case of ants, for instance, maybe the right unit of analysis is not the individual ant, but the colony. Some of the work Deborah Gordon of Stanford is doing, for instance, is very related to

measuring the intelligence of colonies of ants. Once we start thinking along those lines, and also observing other artificial entities on our planet, like computers, which exhibit more and more of this same—or at least a different form of the phenomenon that we might want to call intelligence—once we start seeing these different kinds of intelligence, it becomes clear that intelligence does arise in groups of individuals.

The groups of ants can very usefully be viewed as intelligent probably more than the individual ants themselves. It's clearly possible to view groups of humans and their artifacts, their computational and other artifacts, as intelligent collectively as well. That perspective raises not only deep and interesting scientific questions, but also raises what you might think of as even philosophical questions about what we humans are as groups, not just as individuals.

You might well argue that human intelligence has all along been primarily a collective phenomenon rather than an individual one. Most of the things we think of as human intelligence really arise in the context of our interactions with other human beings. We learn languages. We learn to communicate. Most of our intellectual achievements as humans really result not just from a single person working all alone by themselves, but from interactions of an individual with a culture, with a body of knowledge, with a whole community and network of other humans.

I think and I hope that this approach to thinking about collective intelligence can help us to understand not only what it means to be individual humans, but what it means for us as humans to be part of some broader collectively intelligent entity.