



Superminds: How humans and machines can work together

by Jim Guszczka and Jeff Schwartz

PHOTOGRAPHY BY BEN GEBO

Deloitte.
Insights

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee, and its network of member firms, each of which is a legally separate and independent entity. Please see <http://www.deloitte.com/about> for a detailed description of the legal structure of Deloitte Touche Tohmatsu Limited and its member firms. Please see <http://www.deloitte.com/us/about> for a detailed description of the legal structure of the US member firms of Deloitte Touche Tohmatsu Limited and their respective subsidiaries. Certain services may not be available to attest clients under the rules and regulations of public accounting. For information on the Deloitte US Firms' privacy practices, see the US Privacy Notice on Deloitte.com.

Copyright © 2019. Deloitte Development LLC. All rights reserved.

We naturally think of “intelligence” as a trait belonging to individuals. We’re all—students, employees, soldiers, artists, athletes—regularly evaluated in terms of personal accomplishment, with “lone hero” narratives prevailing in accounts of scientific discovery, politics, and business. Similarly, artificial intelligence is typically defined as a quest to build individual machines that possess different forms of intelligence, even the kind of general intelligence measured in humans for more than a century.

Yet focusing on individual intelligence, whether human or machine, can distract us from the true nature of accomplishment. As **Thomas Malone**, professor at MIT’s Sloan School of Management and director of its Center for Collective Intelligence notes: “Almost everything we humans have ever done has been done not by lone individuals, but by groups of people working together, often across time and space.”

Malone, the author of 2004’s *The Future of Work* and a pioneering researcher in the field of collective intelligence, is in a singular position to understand the potential of AI technologies to transform workers, workplaces, and societies. In this conversation with Deloitte’s **Jim Guszczka** and **Jeff Schwartz**, he discusses a vision outlined in his recent book *Superminds*—a framework for achieving new forms of human-machine collective intelligence and its implications for the future of work.

Superminds and collective intelligence

JIM GUSZCZA, US CHIEF DATA SCIENTIST, DELOITTE CONSULTING LLP: Let's start by defining our terms. Can you tell us what a "supermind" is, and how you define collective intelligence?

THOMAS MALONE, DIRECTOR, MIT CENTER FOR COLLECTIVE INTELLIGENCE: A "supermind" is a group of individuals acting collectively in ways that seem intelligent, and collective intelligence essentially has the same definition. For many years, I defined collective intelligence as groups of individuals acting collectively in ways that seem intelligent. But I think it's probably more useful to think of collective intelligence as the *property* that a supermind has.

GUSZCZA: So collective intelligence is a kind of emergent property of a group of individuals?

MALONE: Yes, and it doesn't always have to be a group of people. Collective intelligence is something that can emerge from a group that includes people and computers. Or it could be a group of only computers, or of bees or ants or even bacteria. Collective intelligence is a very general property, and superminds can arise in many kinds of systems, although the systems I've mostly talked about are those that involve people and computers.

JEFF SCHWARTZ, PRINCIPAL AND US LEADER, FUTURE OF WORK, DELOITTE CONSULTING LLP: Even before talking about collective intelligence, you make an important distinction between two kinds of intelligence, right?

MALONE: Yes. In a very broad sense, you could say intelligence is the ability to achieve goals. There are other ways of defining intelligence, but that one is useful for our purposes. And this suggests two more specific kinds of intelligence. The first is *specialized intelligence*: the ability to achieve specific goals in specific situations. The other is *general intelligence*:

the ability to achieve a wide range of goals in a wide range of situations.

GUSZCZA: And if I understand correctly, that distinction is important to understanding the capabilities of today's AI systems.

MALONE: Yes. Something many people don't realize is that even the most advanced AI programs today have only specialized intelligence. For instance, the IBM Watson program that beat the best human players on *Jeopardy!* couldn't even play tic-tac-toe, much less chess. It was very specialized for the task of playing that specific game. And, similarly, a self-driving car that may be great at staying on the road in the middle of traffic can't begin to take objects off a shelf in a warehouse and put them in a box. Each of these programs has only specialized intelligence. In contrast, even a five-year-old child has more general intelligence than the most advanced computer programs today. A child can carry on a much more sensible conversation about a much wider range of topics than any computer program today, and operate more effectively in an unpredictable physical environment.

GUSZCZA: A fundamental insight from artificial intelligence research in the past 60 years is while computers are often good at things that are hard for humans, many things that come naturally even to young children are very difficult for computers.

MALONE: To a first approximation, that's right. We have often fallen into an assumption that intelligence is one-dimensional: You can have more or less intelligence, but there's only one dimension to it. We will increasingly come to understand there are many dimensions of intelligence, and many different kinds of intelligence possible through different combinations of those dimensions. So it's much more complicated than just "more or less." It's a whole space. There may, for example, be as many different kinds of intelligence as there are species of living things on our planet.

Now, it's certainly the case that for some kinds of intelligence—doing arithmetic, for instance—computers are way better than people. And over the past decade, computers have become much better than people at certain kinds of pattern recognition made possible by machine learning. But that doesn't mean computers are smarter than people at everything, by any means. It just means that, for this particular kind of thinking, if you want to call it that, computers are way better than people. But there are plenty of other things that people are better at than computers.

Measuring group intelligence

SCHWARTZ: Early in your book *Superminds*, you discuss the characteristics of intelligent groups. Can you say a bit about this?

MALONE: We were essentially trying to develop an IQ test for groups. IQ tests measure the general—not specialized—intelligence of *individuals*, and they've been around for about a century. It turns out to be an empirical fact that people who perform well at a certain task, such as reading, also on average perform well at other things, such as math or three-dimensional figure rotations. In other words, someone's ability to do one mental task is correlated with their ability to do very many others. This is the broad general intelligence of *individuals* that traditional intelligence tests measure.

But, as far as we could tell, nobody had tried to create a test of the general intelligence of *groups*. We wanted to see whether there was a similar kind of general intelligence for groups, and we found that yes, in fact, there is. It appears there is for groups—just as for individuals—a single statistical factor that predicts how well a group will do on a wide range of very different tasks. We call this factor *collective intelligence*—it's a way of measuring what you might call *general* collective intelligence. We thought it was pretty interesting to show that such a factor exists and that it's possible to measure it.

What many people found even more interesting was what we found to be correlated with group intelligence. At first, we worried that the intelligence of individual group members would be pretty much the only thing that determined how smart the group was. But we found the correlation between the group's collective intelligence and the individual intelligence of the group *members* was only moderate. In other words, just having a bunch of smart people isn't enough to make a smart group. Instead, we found three other characteristics that were significantly correlated with the group's collective intelligence.

The first was the degree to which the people in the group had what you might call social intelligence or social perceptiveness. We measured this by showing people pictures of other people's eyes, and asking them to guess what emotion the person in the picture was feeling. It turns out that when a group has a bunch of people who are good at this, the group is, on average, more collectively intelligent than when it doesn't. The second factor was how evenly people participated in the group's conversations. If you have one or two people in a group who dominate the conversation, then, on average, the group is less collectively intelligent than when people participate more evenly. And, finally, we found the group's collective intelligence was correlated with the proportion of women in the group. Having more women was correlated with more intelligent groups. It's important to understand, though, that the factor about female membership was mostly explained statistically by the factor about social intelligence. So one possible interpretation is that what you need for a group to be collectively intelligent is to have a number of people in the group who are high on that measure of social intelligence.

We don't think this is the final word; we believe there are many other factors affecting what makes a group smart. But this is at least an intriguing set of suggestions about the kinds of things that can help make groups smart.

GUSZCZA: It often seems organizations reward individual performance, but hope for good teamwork. Is there enough of a movement toward actually trying to cultivate practices and standards around forming smart teams in large organizations?

MALONE: There is a great deal of work that could be done here. As you say, most evaluations in organizations still rest on individuals, yet the organizations' results depend almost entirely on teams. We could certainly do much more evaluation of teams and, perhaps even more importantly, we could do much more systematic analysis of what helps make teams work better. One of the things we are now increasingly in a position to do is to capture vastly more data about who's on a team, who does what work, and how well the team's work turned out. So there's a lot of really interesting work that can be done to build more evidence-based results about—and guidelines for—how to create effective teams.

Humans in the loop, computers in the group

GUSZCZA: The discussion of collective intelligence leads us back to a major theme of your book: that it's more useful to think of AI in terms of humans and computers complementing one another within the context of smart groups, rather than viewing it as a zero-sum game in which humans just fill in for what computers can't yet do.

MALONE: We have spent way too much time thinking about people *versus* computers, and not nearly enough time thinking about people *and* computers. Way too much time thinking about what jobs computers are going to take away from people, and not nearly enough time thinking about what people and computers can do together that could never be done before. As we develop a deeper

understanding of the space of possible kinds of intelligences, we'll be able to talk more precisely about that.

SCHWARTZ: How should we be thinking about and exploring the different ways that people and machines will work together in the future?


“We have spent way too much time thinking about people *versus* computers, and not nearly enough time thinking about people *and* computers.”

MALONE: One thing people often talk about when discussing people and computers is to say that we need to have “humans in the loop.” That usually means that computers are going to be doing almost everything, but we'd better have some people around in case something goes wrong. But I think it's much more useful to start with the ways humans have accomplished almost everything we've ever accomplished in our history: in groups. These groups of humans are examples of what I call superminds. They can be companies, or armies, or families, or many other kinds of things. Almost everything we humans have ever done has been done not by lone individuals, but by groups of people working together, often across time and space. This includes everything from inventing language to making the turkey sandwiches I usually have for lunch.

So rather than start with the “human-in-the-loop” concept of “one person, one computer,” let's start with the human groups we've used to accomplish almost everything and add computers into those groups. When we do that, computers can use their specialized intelligence to do the things they do better than people, and people can use their general

intelligence to do the things computers can't do very well yet.

Even more importantly, we can also use computers to create what I call *hyperconnectivity*: connecting people at a scale and in rich new ways that were never possible before. If you think about it, almost everything we use computers for today is really some form of this. Most people use computers primarily for email or looking at the Web or word processing or social media or various things like that, none of



which really involve much artificial intelligence or even much computation in the sense of arithmetic or logical reasoning. These uses of computers today are really almost entirely about connecting people to other people. And I don't think that's going to change anytime soon.

In fact, I think we often overestimate the potential of artificial intelligence, perhaps because it's so easy for us to imagine computers as intelligent as people. But unfortunately, it's much harder to create such machines than to imagine them. On the other hand, I think we often underestimate the potential of hyperconnectivity, perhaps because in a certain sense it's easier to create hyperconnected systems than to imagine them. We've already created the most massively hyperconnected groups the world has ever known, with billions of people connected to the internet. But it's hard for us to imagine what they can do already, much less what they will be able to do in the future.

A phrase I like to use to summarize all this is that we need to move from thinking about "humans in the loop" to "computers in the group."

Four types of human-computer collaboration

SCHWARTZ: You discuss multiple ways that computers can be "in the group," as tools, assistants, peers, managers, and so on. Could you give a few examples of that?

MALONE: We already know a lot about the different roles people can have relative to each other in groups. So that gives us at least some language for thinking about the roles computers can have as well. The most obvious one, and the one people talk about the most, is computers playing the role of *tools*. For example, when you're using a computer as a word processor or a spreadsheet, the computer is doing exactly what you tell it to do, and is more or less subject to your constant attention. As with

other kinds of tools, the computer doesn't do much unless you're there telling it exactly what to do.

The next level up is what you might call an *assistant*. We certainly use people as assistants for other people. And computers are increasingly taking on that role. Unlike a tool, the assistant often has more autonomy, takes more initiative in helping achieve your goals, and may know things you don't know to help you achieve your goals more effectively.

GUSZCZA: Here at Deloitte, many of us have been doing data science and predictive analytics for about 20 years. One of our applications has been building predictive algorithms to help insurance underwriters better select and price risks, or claims adjusters better handle insurance claims. For the simplest cases the computer just completes the task. For intermediate cases, the human might need to disambiguate some inputs. The human then spends more time on the complex cases that require context and common sense and judgment. Would this be an example of an assistant?

MALONE: That's a great example of an assistant. The computer can actually do some of the tasks cheaper, faster, and often better than the person, just as an electric saw can cut things faster than a person can. But, unlike the electric saw, the underwriting assistant can also take more initiative when handling straightforward cases. You could even say that things like the autocorrect function in text messaging is an example of an assistant that can take a little more initiative—often with amusingly off-the-mark results!

The next level up is what you might call a *peer*. We'll increasingly see examples of computers acting as peers for people in many kinds of situations. One of my favorite examples is from a research project I did several years ago with Yiftach Nagar. We trained machine learning predictive algorithms to predict the next plays in American football games, and then let the computers participate in prediction markets along with humans.

SCHWARTZ: And what about machines as managers?

MALONE: That's the last kind of possibility in this spectrum. People can get freaked out about this, but if you think about it, we already have machines as managers in many situations that seem very normal. In the old days, police officers directed traffic at busy intersections. Today, stoplights do this, and we think nothing of it. It seems completely natural and normal, as I think it should. It's quite likely we'll see more and more examples of machines doing things like using algorithms to figure out the sequence of tasks that need to be done, predicting which person is best suited to do each task, and automatically routing the task to that person.

Another thing managers often do is evaluate the work of the people who report to them. In some cases, computers can easily evaluate people's work. An example from the realm of science is a system called Foldit. Foldit helps scientists discover new ways of folding protein molecules in three dimensions to have certain medicinal or other properties. It turns out people are better than computers at figuring out new three-dimensional ways of folding molecules, but computers are much better than people at *evaluating* the potential energy that's relevant here. The Foldit system has helped make significant progress in developing ways of treating AIDS, for instance, by using this combination of people to generate possibilities and computers to evaluate those possibilities. This is another example of a computer acting as a kind of manager, in this case evaluating the work of the people. Once again, nobody thinks there is anything particularly strange about this, and I believe we'll see lots of such examples.

Humanizing work

SCHWARTZ: We've discussed what machines can do well, and also what groups and superminds can do well. As you think about the types of capabilities

“I don’t think there is a fixed definition of what it means to be human or to humanize. That’s a malleable thing that changes as the animals, machines, and other things around us in the world change.”

and skills we humans need to develop and double down on in the coming decades, what will the *human* dimension of work look like? This is a big discussion on what it means to humanize work, and what skills and capabilities are required.

MALONE: Our concept of what it means to be human is affected by what else is in the world around us. A few hundred years ago, only humans could do arithmetic. So, it was in a certain sense “humanizing” to do arithmetic. But now that machines can do arithmetic way better than we can, we don’t think of arithmetic calculation as a human-like activity anymore. And, more generally, as computers do more of the things that used to be doable only by people, we’ll come to think of those things as not part of what it means to be human. The point is I don’t think there is a fixed definition of what it means to be human or to humanize. That’s a malleable thing that changes as the animals, machines, and other things around us in the world change.

With respect to computers in particular, I get a little frustrated with people who say things like, “Well, computers will never be really creative” or “They’ll never be able to have deep interpersonal skills.” Yes, they will be able to do some of those things more and more over time. It’s very difficult to draw a hard-and-fast line around things computers will *never* be able to do.

But, as a practical guide, there are some things people are likely to be able to do better than computers for the foreseeable future. One is using general intelligence, which we’ve already discussed. A second is interpersonal skills, which we spoke of as being especially important for the collective intelligence of human groups. Even though computers can do some kinds of interpersonal things already and will do more of them over time, it’s going to be quite a while before computers have the kind of broad interpersonal skills that people do. One thing that we’ll end up paying people more for is their interpersonal abilities.

In medicine, for instance, increasingly there will be algorithms able to process all kinds of lab test results and millions of case examples in their knowledge base and do a pretty good job of diagnosing human illnesses. They will probably be able to do this better than most human physicians could even when the human is sitting in the room with the patient. But there’s still going to be a need for humans in the room with the patient. People will be needed to gather the information for online diagnoses and to help provide some of the needed treatments. Perhaps most importantly, people will be needed to provide the kind of human contact and sympathy that’s an important part of the healing process. In this and many other kinds of work, people’s interpersonal skills will probably become even more important than we expect today.

A third dimension where it will probably be some time before computers come close to humans is certain kinds of physical skills, such as operating effectively in complicated and unpredictable physical environments. There are already robots that work just fine on assembly lines where everything is very cut and dried and very routine. But think about the physical skills needed, for example, to be a plumber. You've got to figure out how to open a particular kind of a cabinet under a sink and know how to move the different shapes of bottles and cans and whatever other stuff is under there, and you've got to figure out how to maneuver around weirdly shaped pipes and maybe cut part of a wall open to get at something in a weird old-fashioned building. All kinds of complicated and unpredictable physical skills are needed that machines aren't likely to have anytime soon.

Those are three examples of where there will be continuing needs for humans for the foreseeable future: general intelligence, social intelligence, and physical intelligence. I believe people's jobs will increasingly be humanized in the sense that they'll include more of those things that humans do better than machines.

But suppose there comes a day when computers and physical robots can do *everything* that people can do, better and cheaper. That's probably at least many decades away. But even if that day comes, I think there will still be some things that we'll want people to do, such as keeping us company. Even today, why do we go to see live actors perform a play when we can actually see a higher-quality performance on our TV anytime we want? Why do we go to a football game and watch humans try to move a ball down a field with a combination of other humans? I'm pretty sure it would be easy to make a machine that could do that better than people can, but I don't think it would be as entertaining to watch machines play football against other machines as it is to watch humans play football against other humans. I think there will always be a desire for humans to do some things simply because they're humans.

Implications for organizations

SCHWARTZ: What are some implications for public institutions and business leaders as they try to operationalize this?

MALONE: Superminds are the entities that accomplish almost everything in our world. Every company in the world is a supermind. Every democratic government is a supermind. Every army, neighborhood, scientific community, club. Every market where you buy and sell things is a supermind. Superminds have been around at least as long as people have, and when you learn to recognize them, you realize they run our world. Almost everything we've done has been done by superminds.

In the book, I discuss five different types of decision-making superminds: *hierarchies*, *markets*, *democracies*, *communities*, and what I call *ecosystems*. Thinking about which kinds of superminds are relevant for different kinds of situations is potentially a very powerful way of thinking about many of our societal problems.

We can understand many of the things that happen in society as resulting from the interplay of different types of superminds: Laws are enforced by hierarchical governments, which are chosen by democratic election processes that in some sense reflect the values of broader communities. Understanding this interplay between communities, democracies, and hierarchical governments provides a way of thinking more systematically about what should be done by which kinds of superminds.

For instance, if we want to deal with the problem of fake news, we could try to let *markets* do it on their own (which hasn't worked very well so far), we could let government *hierarchies* regulate it, or we could rely on *community*-based reputations by, for example, letting broadly respected organizations use online systems to rate the credibility of different news sources.



“Superminds have been around at least as long as people have, and when you learn to recognize them, you realize they run our world.”

SCHWARTZ: And can you ask a similar set of questions at the level of companies?

MALONE: Yes. For example, most decisions in companies today are made by the corporate hierarchy, but which decisions might be better made by some kind of democracy? Are there decisions currently being made by managers that could be better made by combining the votes of people who really know the situation? Are there decisions, such as how much of which products to make, that could be made better by some kind of internal market rather than by a managerial hierarchy? And, whether we realize it or not, many decisions in a company are made by communities—a kind of informal consensus involving community norms. People often call that the “culture” of an organization, but I think “community” is another good word for that. Once again, this way of looking at the world gives us a systematic framework for thinking about how best to design and combine these different kinds of superminds.

Perhaps the simplest and most broadly applicable implication is just the very idea that each company is a supermind. Realizing this gets us thinking about a) how we’re kind of “in this together”, and b) how could we make our superminds smarter. Traditionally, we’ve spent a lot of time thinking about how to make companies more productive. But the measures of productivity were mostly developed as a way of capturing the things important in a manufacturing economy. As we increasingly move into what you might call a knowledge-based economy, productivity in some sense still measures things that are important. But many other things that are becoming more important are probably more usefully thought of as *intelligence*, not *productivity*. So how could we create more intelligent companies, more intelligent organizations? The idea of superminds is a pretty natural way of understanding our companies, the markets that interrelate them, and so on. It’s useful for managers

to think in terms of us being part of a supermind, and how we can make our superminds smarter.

SCHWARTZ: To zoom out still further, what are the implications of superminds for the future of work?

MALONE: For individuals, there are at least two kinds of implications. The first is, if you want to accomplish almost anything in the world and if you're realistic about it, you need to be thinking about how to work with superminds to achieve whatever you want. In some sense we already know that, but this gives us a more systematic framework for thinking about it.

The other, perhaps even more personal, view is that we as individuals are all part of many powerful superminds. And all of these superminds are part of one giant global supermind. So not only our fate as individuals, but our fate as humanity really depends on the choices our global supermind makes. We should be hoping we can influence our global supermind to make choices that are not just *smart* but also *wise*. To do that, we should be thinking about what values are most important to us—what values we think are most wise—and how can we help support and shape our superminds to achieve those values. ●

JIM GUSZCZA is Deloitte Consulting LLP's US chief data scientist. He is based in Santa Monica, California.

JEFF SCHWARTZ, a principal with Deloitte Consulting LLP, is Deloitte's global leader for Human Capital Marketing, Eminence, and Brand and the US Consulting leader for the Future of Work. He is based in New York.

Read more on www.deloitte.com/insights

Why humans and computers think better together

Some have voiced fears that artificial intelligence could replace humans altogether. But that isn't likely. A more valuable approach may be to view machine and human intelligence as complementary, with each bringing its own strengths to the table.

Visit www.deloitte.com/insights/cognitive-collaboration



Deloitte.

Insights

 Follow @DeloitteInsight #DeloitteReview

 www.linkedin.com/company/deloitte-insights

Subscribe to receive email updates at deloitte.com/insights