

AI & Humans Working Together

What Happens Next - 07.15.23

Larry Bernstein:

Welcome to What Happens Next. My name is Larry Bernstein.

What Happens Next is a podcast which covers economics, education, and culture.

Today's Topic is AI & Humans Working Together – Not Independently.

Our speaker is Tom Malone who is a Professor of Management at MIT's Sloan School and the Director of the MIT Center for Collective Intelligence. Tom is the author of the book *Superminds: The Surprising Power of People and Computers Thinking Together*.

AI is the rage, and I want to learn from Tom about how we will interact with AI-based algorithms to become a Supermind by combining the best that humans can do with the awesome computing power of AI working symbiotically together.

Let's begin today's podcast with Tom's opening six-minute remarks.

Thomas Malone:

Let me summarize the two main messages in my book *Superminds*. The first is that we should be spending less time thinking about people or computers, and a lot more time thinking about people and computers. Less time, for instance, thinking about how computers are going to take away jobs and more time thinking about how people and computers together can do things that could never be done before. The second message is about how to see ghosts. Now, I don't mean real ghosts, but I do mean powerful entities that are all around us that are often invisible unless you know how to look.

These ghosts are what I call Superminds, which I define as groups of individuals acting together in ways that seem intelligent. Now, by this broad definition, Superminds really are all around us all the time. For instance, every hierarchical company is a kind of Supermind, a group of individuals acting together in ways that seem intelligent. Every democracy is a Supermind, whether it's in a club, a company or some other kind of organization. A very important kind of Supermind are the markets for goods and services, and communities, whether it's a neighborhood scientific community or some other kind of group. Now, here's something a lot of people don't realize. Almost everything we humans have ever accomplished was done by groups of people working together, often over time and space.

Computers have the potential to make these Superminds much smarter. Think of Wikipedia, where thousands of people and computers all over the world have created the largest encyclopedia. I think that's a good example of how computers can make Superminds smarter. It's also possible to use computers in ways that make Superminds stupid, like when fake news influences voters in a democracy.

It's fashionable these days to be pessimistic about how AI will affect society, but I have a more optimistic view. Computers have the potential to help us create human-computer super-lines that are smarter than anything we've ever seen before.

Computers can do things better than people like arithmetic and pattern recognition, and humans can do the rest. Perhaps even more importantly, we can also use computers, the internet and zoom to create what I call hyperconnectivity. That is connecting people to other people and often to computers, at much larger scale.

Now, I think we often overestimate the potential of AI in all this, perhaps because it's easy to imagine computers as smart as people because our science fiction is full of it. But unfortunately, it's much easier to imagine such computers than to create them. And I think it's likely to be at least many, many decades before we reach full human level artificial intelligence.

On the other hand, I think we often underestimate the potential of hyperconnectivity. Perhaps that's because in a certain sense, it's probably easier to create hyperconnectivity than to imagine it. For instance, we've already created the most massively hyperconnected group of people our planet has ever known with billions of people connected to the internet. But it's still hard for us to imagine how to use this hyperconnectivity. I think we need to move from thinking about humans in the loop to thinking about computers in the group. That's my message.

Larry Bernstein:

A Supermind are individuals working together to make an intelligent group. And in your book, you use Xerox repairmen that interacted using a message board to solve problems as an example of a successful Supermind group. Tell us about that ethnographic study.

Thomas Malone:

Well, you might think that fixing copying machines is a one-person job, that the person who goes out there should have the knowledge needed to diagnose and repair these machines. But a former colleague of mine at Xerox Park named Julian Orr did do an ethnographic study that you were referring to where he talked to these repair people and found that the work, they did was often much more cooperative, collaborative, collective than you might think.

They ended up creating an online knowledge base of different kinds of problems and ways of solving them and gave awards to the people who came up with the most useful examples in that knowledge base. That's an example of how it can make your organization, in this case, the group of repair people, much more intelligent than they would otherwise have been.

Larry Bernstein:

In your book, you describe five different Superminds: hierarchies, communities, democracies, markets, and ecosystems. Tell us about that.

Thomas Malone:

There are five different kinds of Superminds for making group decisions. The hierarchy where group decisions are essentially made by delegating them to individuals lower in the hierarchy. And those decisions made lower can always be overruled by people higher in the hierarchy. That's an effective way of organizing a lot of activities and making group decisions in a way that's pretty robust. It certainly worked well for us in the last century or so. The Xerox repair people were an example of what I call a community, a group of individuals that make decisions through informal consensus, often based on shared norms and reputations. That's also very robust, in fact, that's probably the oldest form of Supermind that we humans have used.

There are democracies where you have a formal voting process, and the individuals or the decisions that get the most votes are the ones that are the decisions of the whole group. Another important Supermind is the markets that exchange goods and services have an interesting way of making group decisions, which is the group decision is the combination of pairwise agreements between individual buyers and sellers. So, they don't all have to agree on any one thing.

If two people get together and agree, they want to exchange things, they do, and the total of all those is the group decision of the market. So those are the four kinds of human forms of cooperation that do require cooperation to work. But if you don't have any cooperation among the groups, group members, then you have the fifth kind of Supermind, which I call the ecosystems. So, in an ecosystem, the group decision is made by the law of the jungle. Whoever has the most power gets what they want and the survival of the fittest. So, I would say those five types of Superminds for making group decisions account for most things we see in the world around us. Almost everything we see what happens in a company, what happens in a country, those are all combinations of those different kinds of Superminds.

Larry Bernstein:

I spent my career trading fixed income securities. And the amazing thing is all you need to know is the price and a description of the security to invest. You don't need to know the seller, or why it is for sale, or any other aspect, it is all about price. Tell us about the role of prices and markets.

Thomas Malone:

Absolutely. You've just described the magic of the market or what Adam Smith called the invisible hand of the market. And it is hard to believe and almost magical if you think about it, that this very, very decentralized form of decision making can do a pretty good job of allocating resources, goods, people's time, et cetera, in an efficient way.

If those are commodities, then those descriptions are well known. But if you're trading things that aren't commodities. Each thing has some individual characteristics like the services of a particular human with expertise or a particular automobile with special options you need to know what it is that's being sold at that price. But I think that combination of prices and product descriptions does allow a very information efficient and decentralized way of making lots of very complicated decisions with pretty good outcomes.

Larry Bernstein:

Wikipedia is a fantastic example of a community Supermind. I heard that one day one of the random contributors downloaded all the US census data for every town in the US to the Wikipedia sites for each individual town. This is a community that applies knowledge and technology to make a superb final product that is constantly evolving.

Thomas Malone:

Well, I think that's a great example. One that's been a favorite of mine for a long time. Wikipedia is primarily a community. It operates with a kind of informal consensus. Essentially anybody on Wikipedia can change anything anytime they want to. And if somebody else thinks it should be changed, they can change it too. And so it keeps changing until everyone who cares enough to look is satisfied with what's there, so that's a literal consensus. Part of the secret of Wikipedia was that they used those well-known aspects of community decision making, like norms about what's good, what kind of articles we need, what things we shouldn't have in articles.

Wikipedia codifies many of those norms very explicitly on the site. People who edit Wikipedia have reputations within the Wikipedia community. Part of what motivates people is their desire to have a reputation among people they care about in the community of Wikipedia editors. What is unusual about Wikipedia is the degree to which they've used modern communication and computational technology to let a community operate at a scale and in a way that would've been unimaginable before.

The number of active Wikipedia contributors it's something like 50,000. So, imagine 50,000 people in a giant football stadium, and they all had paper and pencil, and they were trying to write an encyclopedia by consensus. It would've just been impossible for enough people to see all the things people wrote. They would've probably resorted to a hierarchical structure with

editors and subeditors and so forth. And it seems to me almost impossible that they could have done anything remotely as good and as fast as Wikipedia.

Larry Bernstein:

Years ago, I was watching the Chicago Bears in a playoff game, and the starting quarterback Jay Cutler got injured on a play and his backup entered the game. I didn't know anything about his replacement so I looked him up on Wikipedia, and it had all his stats from college and pro ball, but most incredibly it stated that he had entered that playoff game 15 seconds ago. This truly was a real-time encyclopedia. What are the implications of being up to date to the minute?

Thomas Malone:

Yes, absolutely. So, it makes it possible to do so many things so quickly. Another point here is that if you didn't have the communication and editing technology that Wikipedia is built upon, you couldn't have anything like that larger scale of community operating effectively. You would probably have to resort to some kind of hierarchical structure to reduce the amount of communication that's needed. One of the big lessons here is that the cheap communication and computation makes it possible to organize collective human activity in ways that are not only faster or wider or bigger, but often it makes it possible to organize them using a community instead of a hierarchy using a market instead of a democracy. We can do things faster and in different, more decentralized ways.

Larry Bernstein:

Some organizations like Wikipedia are successful because they have rules and norms that allow for growth and use decentralized expertise to solve problems. Occupy Wall Street failed maybe because it lacked a hierarchical structure and community norms to succeed. What norms were missing from Occupy Wall Street that undermined its efforts?

Thomas Malone:

I use the example of Occupy Wall Street in my book, and it's a cool example because I was there personally to witness part of what went on. I was on sabbatical at NYU at the time of Occupy Wall Street, and a couple of weekends I went down and listened to the protestors discussing things. The most interesting session I went to was one where several dozen people were trying to come up with a mission statement for Occupy Wall Street. Anybody could walk in off the street and be part of it. In fact, Michael Moore, the movie director, was one of the people sitting there in the room with me, and a bunch of other people. I came in and they were very near the beginning of the proposed mission statement.

So far they have agreed on saying, we believe in a fair and just society. Somebody said, I think it should really be, we believe in a truly fair and just society. So, they were following a set of rigid norms for consensus decision making that included things like having almost everyone must

agree. And if this is not a consensus decision, then you must keep going until you find something that everyone will agree to, or at least be willing to not leave the group. So given those norms of consensus decision making that the occupied members were following, they started discussing this question of whether we should add the word truly to the mission statement.

And there was one person there who said that if we don't do this, she would leave the group. And people went back and forth for at least two hours trying to decide whether to use the word truly in there. So, the person who was originally blocking the consensus eventually said, oh, well, never mind, it's okay with me. You can go ahead. But then there was a long debate about whether the rules allowed that? What do the rules say we can do in this situation or not do in this situation. Very little of the discussion was about is the word truly a good thing or not.

I left there feeling very pessimistic about the possibility that this group would ever come up with a coherent mission statement. Years later when I was writing my book, I scanned the web to see if I could find any mission statement for Occupy Wall Street and couldn't find one anywhere. So, it's likely that they didn't manage to do that which I think was a pity.

They were certainly very well intentioned. I thought later as I was writing the book, if this group, instead of trying to do everything in a face-to-face group had been able to use the technology of Wikipedia, they may well have come up with a very good mission statement. That's an example of how new technology can create new ways of organizing human groups.

Larry Bernstein:

Life is very complicated, and institutions work for reasons that are not obvious and they get the kinks out. Parliamentary procedures are a fast and easy way to manage a group meeting instead of coming up with something new on the fly that seems better or fairer.

Thomas Malone:

I completely agree with that. Norms can be very useful, very valuable, and we need them. The norms are the key components of communities as decision making mechanisms. Humans, ever since our days of hunting and gathering bands, have used norms as a very important way of organizing our group activity. Some norms are kind of arbitrary. It doesn't matter that much what they are as long as everyone agrees. Like, how many days should there be in the week? There's no single right answer to that. It just should be something that everyone agrees on.

If you had to choose between parliamentary procedure and consensus decision making, and if you were in a large group, especially a group of people who may not have been kind of aligned in terms of what they wanted, then it's almost certain that the parliamentary procedure methods would work more effectively than the group consensus methods. But if you're in a group where deep dedication to the group decision is critical, like if you're about to go into war and

everybody's life is going to be on the line and one way of getting a deep commitment of the group members if everyone essentially had to agree on something before the group did it. There are trade-offs. I'm not saying that either one is always better, depends on the situation, depends on how much dedication is needed from the people in the group, depends on how quickly you need to make the decision. In the case of the Occupy protestors, they were using a very formal and rigid form of consensus. Most communities have a more flexible, informal kind of consensus decision making, and that seems to work well in a lot of cases.

Larry Bernstein:

How will AI change the nature of work?

Tom Malone:

The main way that computer technology has changed things is by reducing the cost of communication. It made it possible for people all over the world to communicate much faster, cheaply, and easily. Now with a new generation of generative AI, it will be possible for computers to participate in ways that weren't ever possible before. Computers will be able to play a role like what used to be played by humans. For instance, in Wikipedia, computer technology will allow the use of little bots, little programs to find curse words or other objectionable content and remove it from Wikipedia articles automatically. They've also had simple tools that would correct grammar automatically. So, I think it will now become possible for more intelligent bots based on generative AI technology to understand English and respond in English in ways like what a human would do.

Larry Bernstein:

Moore's law is an observation that the number of transistors in an integrated circuit doubles every 18 months to two years. And this doubling also increases processing speeds, database sizes and reduces the cost of computing. How will Moore's law impact the power of AI in the near future?

Tom Malone:

So there's several things going on here. One is Moore's Law was a remarkably stable and accurate prediction of progress in computer technology for many decades, long after people thought it would last. People were always saying, "well now we reached the end of Moore's Law," but then some new technology would make it possible to keep going. It's not a law of nature or physics, it's a description of a very complicated human, economic, technological process of how these things get made and how they get priced and so forth.

It's been a remarkable phenomenon. I'm not a deep expert on this, but my understanding is that Moore's law hasn't been operating as reliably in the last few years as it did for many decades before. One way we've dealt with that, however, is by going to parallel processing in a much

bigger way. Moore's law is about how fast a single processor can process things, or how much memory a single processor can have, but it's always been possible to have multiple processors working in parallel. They can't do everything that a single processor can do as fast.

In other words, if you make a single processor five times faster than it can do anything five times faster, if you have five parallel processors for some tasks, they can do things five times as fast, as long as those tasks are all parallelizable. In other words, you can do things without worrying about the things going on in the other four processors. In general, you can't do everything that way. But there are some things like graphic processing, which can be done in a highly parallel way. And the new AI chips that Nvidia and other companies are selling originally were created or designed for graphical processing, but now they can be used for many of these AI programs, which are highly parallelizable. So that's a way of getting around Moore's Law. Even if you can't make the individual processors twice as fast every 18 months, you can just create many more of them and use software and algorithms that allow things to be done in parallel to a large degree. And that's a big part of how a lot of these recent generative AI programs have been able to do what they do.

Larry Bernstein:

How will these more powerful computing machines change AI's performance?

Tom Malone:

Nobody knows for sure. There are two possibilities. One is that to go from the AI we have today to human level AI, all we need is more processing power. So, the deep learning algorithms that these systems use depend on adjusting billions of parameters through learning from many more billions of examples of text and images and so forth. What they found is that when you use the same algorithms, but you just have more parameters you get better performance. So, again, I'm not a deep expert on this, I don't think this is likely, but it's possible that just getting enough more parameters will make these current algorithms so smart that you can't tell them apart from humans. I think some other much deeper algorithmic advances will be needed, but at least to some degree, just having more and more processors that can manage more and more parameters will almost certainly continue to increase the intelligence of these AI algorithms.

Larry Bernstein:

How should we think abstractly about how AI works?

Tom Malone:

Here is a kind of intuition for how these current generative AI systems work. A lot of people when they see a program that you type something in English into it and it comes back with a very articulate answer in English, they think, wow, that sounds just like a human. But I think a better model for how these things work is not that they're like humans underneath and they have

human emotions and human motivations and stuff, it's like what they really are is just much bigger versions of the auto complete that you have when you type in something in the Google search bar.

So what these algorithms are trying to do is give the words that are here so far, they're basically just trying to predict what's the most likely next word. It turns out that when they have these billions of parameters that are learned, that recognize various patterns, they can do an amazingly good job of predicting the next words.

They must have learned a lot of things about how the world works and how humans talk about things by capturing those very complex patterns in probabilistic form, but they're not thinking like humans do, or at least they're not thinking like humans do consciously.

An interesting possibility is that the way these current generative AI systems work may be more like how human unconsciousness works. Unconscious thinking may work more like those algorithms and our conscious thinking where we're doing logical reasoning that we can explain in words is more like what was called classical AI. That is if Julius Caesar was a Roman and all Romans drank wine, then Julius Caesar must have drunk wine.

So that would be an example of a logical syllogism that classical AI was all about. The recent progress in AI has been driven largely by these machine learning systems, very deep neural nets that have these billions of parameters that learn things, but in a way that's very hard, even for the human programmers who create these systems to understand the reason you said this word at that point is because these billion parameters had a greater value for that word than any other word. That's not a very satisfying explanation to humans. It's not the way we think we think, but much of what we humans do really occurs at a kind of unconscious level that we can't describe at all.

For instance, how do you recognize your mother's face? You can say something about it, she had dark hair and a short nose, but it's very difficult to describe another person's face in a way that someone who's never seen the person before would recognize them just from your description.

So, facial recognition occurs at a kind of unconscious level that we can't articulate explicitly in a way that other people can understand. And these new systems are maybe capturing human intelligence better than they're capturing the intelligence that we can think about and talk about consciously.

Some of my colleagues like Josh Tenenbaum at MIT have talked about what's called neuro-symbolic computing, a combination of the neural net approach, which is what today's most advanced systems are using, and the logical symbolic kind of reasoning that classical AI have used. Maybe what we need is some combination of those two kinds of computing, and that might

be more like what we humans do. We can't talk very much about the things we do in a purely neural way, but we must have some kind of combination of those two kinds of reasoning going on in our minds.

Larry Bernstein:

How can AI help us make better decisions and make us more productive?

Tom Malone:

The much more powerful changes will result not just from simple substitution of automated processing for human labor but from rethinking how we do things in the first place.

Take Wikipedia as an example. Before Wikipedia, the way encyclopedias got done was through a hierarchical process of editors and reviewers and expert contributors. Wikipedia made it possible to do far better, far bigger, far more accessible encyclopedia in a completely different way. So how can we invent new organizations that do things in very different ways, create very different kinds of products, that take advantage of the potential that these new technologies make possible. We need to invent new technologies that do new things far better, far cheaper, whatever. I think it may be just as important for us to think of innovative new ways of organizing human work, innovative new ways of producing the same old products and services better and producing new products and services that couldn't even be done before.

Another example of that would be Google Search, a service that would've been completely impossible 30 years ago. It's now one of the biggest companies because it figured out how to do something that was very valuable for people in a way that couldn't have been remotely done before, but it now is. So, I think it's figuring out how to do things like that that will be needed to get the big payoffs from these new technologies in terms of productivity.

Larry Bernstein:

As I get older, I find it more difficult to adapt to new technologies. What does that mean for the adoption of AI?

Tom Malone:

I think you put your finger on another one of the factors that's needed for us to take advantage of the real potential of these new technologies. It's not just reinventing organizations and products and services, it's also changing people's ability to work with and take advantage of the new potentials on an individual basis. One reason for optimism about how it may not take forever for some kinds of changes, but new generations may be needed before they become widely adopted.

Larry Bernstein:

I end each episode with a note of optimism. What are you optimistic about as it relates to AI and technology?

Tom Malone:

I'm optimistic that we will figure out how to do new and better things with AI, not just cheaper than what we used to do, but new ways that people and computers together can do things that were never possible before. I think 10 years from now, it will be hard for us to imagine how we will ever survive without AI, blah, blah, blah. I can't tell you exactly what AI, blah, blah, blah will be, but that's my optimistic expectation for the future.

Larry Bernstein:

Thanks to Tom for joining us today.

If you missed last week's show, check it out. The topic was Gerrymandering as well as Economic Liberty.

Our first speaker was retired Federal Judge Gary Feinerman who discussed the recent case in North Carolina related to gerrymandering and whether the state courts can be the final arbiter for state redistricting maps.

Our second speaker was Renee Flaherty who is the attorney who successfully argued a case with the Georgia Supreme Court. This is about the limits of state authority to regulate occupations. Georgia demanded that women who teach new mothers how to breastfeed meet minimum education requirements that would prevent otherwise qualified women from counseling lactation care.

Renee is an attorney with the Institute for Justice, a not-for-profit, that challenges government overreach on licensing and regulation as well as infringement on individual property rights.

You can find our previous episodes and transcripts on our website whathappensnextin6minutes.com. Please subscribe to our weekly emails and follow us on Apple Podcasts or Spotify.

Thank you for joining me, good-bye.