Welcome to MIT's *ComputerScienceandArtificialIntelligenceLabsAlliances* podcast. I'm Kara Miller. On today's show, companies face a thorny addition problem. How do you combine humans and computers and get something more than the sum of its parts?

Often, the most important factor in business success is not just how productive are you, not how much you can increase your productivity, but how intelligent are you.

Thomas Malone, the founding director of the MIT Center for Collective Intelligence talks about the race to reshape corporations and the reason that the marriage of human and computer can be both incredibly powerful and incredibly challenging.

So sometimes, the hierarchy says, we want you to do X. And the community says, nah. We don't really want to do X. And that makes it really hard for the hierarchical managers to cause the changes they want to change and think should be easy. It turns out it's much harder if there are communities involved in the company that don't want to do that.

In an era of AI, how do you up your game? Try computing with a human touch. That's coming up next.

[MUSIC PLAYING]

Think for a minute about Starbucks at 7:45 on a Tuesday morning. Super busy, people lining up, orders pouring in from mobile phones, baristas looking at screens, talking to each other, handing out chai lattes and cappuccinos. And you've seen this before.

Most of us look at that scene and think steamy chaos. Tom Malone, on the other hand, sees what he calls a super mind.

Which, by the way, I use that word to mean groups of individuals-- they can be people. They can be computers. They can be combinations of both. I define a super mind as a group of individuals acting together in ways that seem intelligent.

The productivity, the interaction, the connectedness on display in Starbucks on a Tuesday morning is incredible when you look at it through Malone's eyes. There's a very delicate interplay between people and computers that results in a level of output out that no single person could accomplish on their own.

And I think that's a concept that we don't really think about that much. But it's a very important way of thinking many, many things that are important in our world, not just groups of humans and computers, but groups of humans all by themselves. I would say that super minds include hierarchical corporations, and democracies, and markets, and communities.

So if you think about it that way, almost everything we humans have ever accomplished was not done by individual humans working all alone. It was done by super minds. Malone is a professor of management at the MIT Sloan School of Management. He has helped found four software companies, holds 11 patents.

And he's the author, among other things, of the book *Superminds.* He believes that as we rush into an era of AI and machine learning, it is worth taking a step back to rethink the relationship between humans and computers. Humans have been teaming up with outside entities, mostly other humans, to create super minds for as long as we've been around.

We've teamed up to get food, to build bridges, to make tools, things we just could not have done as solo actors. We were smarter together. And computers have been one of the standout tools we've made.

They can conjure up spreadsheets for us or provide us with digital calendars. But they can also serve as assistants, suggesting ways we might fix our spelling or finish a sentence. And Malone says we're moving into a moment when computers will operate as more than tools and assistants.

Though, they will be that, too. They will be peers and managers. So imagine, he says, you're driving down the street alongside a self-driving car. All of a sudden, a computer is your peer.

You're a driver. It's a driver. And we may increasingly see computers break up large tasks and delegate them amongst different workers. That's already been tried.

It's likely to become increasingly common. Malone consults with businesses around the world, and believes we are far too concerned about being displaced by computers. But we spend too little time thinking about working with computers to create ever more powerful super minds.

It's not people or computers that we have to think about. It should be people and computers. So we shouldn't be thinking about how many jobs are computers going to take away from people so much. We should be thinking more about what people and computers can do together that could never be done before.

So if we stick for-- I want to talk about that. But just let's stick for a minute with the concern. How much do you hear from business leaders or it could be more from middle managers about that worry that this time is different.

And there's been particular concern about AI, I think, in the sense of like, well, maybe this is a technology that will really eat into white collar work in a way that-- for sure, we've seen factories automate. And we've seen blue collar jobs that have been lost. But maybe this is a different pool of people this time. I just wonder what you hear in terms of concern.

Well, I think that's not just hearing, but I think it's true that the new technologies, especially generative AI most recently, will change the work that many white collar workers do, not just blue collar workers. Though, it'll change some of that, too. So I think that part is true.

People certainly talking a lot about it. I think a lot of people are worried about it. But I think we worry too much and, perhaps, in the wrong way. Many people's jobs will change in ways that will probably make the jobs more interesting.

Some people's jobs may go away, altogether. But I think most of those people will find opportunities to do more interesting jobs. That's an optimistic perspective. Certainly, it won't be true for everyone. But I think it's much more true than many people worry about.

I mean, I always think about all these jobs that, let's say, existed around-- 1900, like people worked at making carriages, and people worked making hats, when women and men wore hats a lot more. And like all these things that almost nobody does, I mean, it's not like somebody doesn't make hats, but not that many people do that anymore.

But as you say, we don't have a 50% unemployment rate. So lots of things exist now that somebody in 1900 would have been surprised was a job, like graphics designer, or something.

Exactly, exactly, yeah. I like to talk about what my grandfather would have thought when-- in the 1930s, he moved to the farm where I grew up in New Mexico. And if you had told him that someday his grandchildren's friends would have job titles like website designer, and software developer, and social media influencer, and so forth, he would have said what does that even mean?

And he would certainly not have believed that those people's jobs were going to replace jobs of people he knew like shoemaker, or plumber, or, for that matter, farmer, actually. As you probably know, in 1800, about 90 something percent of the US population employment was in agriculture. Today, it's less than 2%.

Wow.

So part of why this is hard for people to think about is that it's easy to imagine the jobs that could be destroyed because they exist today. You may know people who have them. It's much harder to imagine the jobs that will be created because they will seem as strange to us now as the jobs that we have today would have seemed to my grandfather.

I'd love to talk, if we can, about a couple of specific situations in business where you feel like the super mind approach has been applied. It can be in different sectors or whatever. But so we can really bring it down to that concrete level.

So let me start with Google Search. This is not a new example. But I think it's one that illustrates the kinds of things that are possible.

You might think of Google Search as a purely technological thing. Let's start by observing that Google Search would be of no use whatsoever if there weren't a bunch of information on the internet created by humans that the Google Search technology could index, and then help you search.

So Google Search is not a purely technical device or invention. It's an invention that only has value because of the combination between humans creating information, technology for organizing and helping to organize and search that information, and humans using that information.

In a somewhat similar way, ChatGPT, which again, people think of as a technology taking stuff away from humans, ChatGPT couldn't exist, wouldn't begin to have the value it does if there hadn't been millions and millions of pages of text and other material on the internet that ChatGPT could learn the patterns in.

And those pages were created by humans. Again, something that we might be tempted to think of as pure technology isn't that, at all. It's a synergistic combination or the result of a synergistic combination of what people can do and what computers can do.

I know you work with a lot of companies. I wonder if you think that there's a difference between the way that kind of-- we might call them like digital native companies, like a Netflix, or an Amazon, that sort of was born in a pretty recent moment. And they were almost born to work on computers versus--

Obviously, there's much older companies. When you think about finance, you think about automotive companies, or oil and gas, or whatever. Do you feel like there's a difference in how those two types of companies think about human-computer collaboration?

There's certainly a lot of differences in those kinds of companies in many ways. It's kind of a well known phenomenon in the history of business that companies take on a character that is, in some sense, determined by the era in which they were founded. For instance, I just heard about a study very recently about companies that were founded in the last three to five years.

Where, it said that they had a much higher proportion of their workers who worked at home or even lived far away from whatever offices the company had much higher than in companies founded in an earlier era. So I think that's just one of many examples of the ways in which new companies very often take advantage of the new technologies that are then available at the time they're founded.

And they're able to do that in a way and without as much difficulty as companies founded before who have already built up their kind of habits, their cultures, their structures based on an earlier generation of technology. Oh, and, I guess, the key point I wanted to make is that I think that even though it's often said that we need to be--

And often believe-- that we need to be innovative in technologies, I think it's, perhaps, just as important or maybe even more important to be innovative in the ways we use those technologies to organize work. In other words, it's just as important to invent new kinds of organizations as it is to invent new technologies that make those organizations possible.

Just as one kind of simple example of that, think about Amazon. Amazon, in a sense, invented-- they weren't literally the first person or first company to do it, but they, in a sense, really did invent what we think of as modern online electronic commerce. That was a new way of doing things, a new way of organizing things.

It couldn't have existed without the internet, and the databases, and so forth. But those technologies alone wouldn't have had nearly the value they do now if there hadn't been inventions like the kind of online electronic commerce that Amazon pioneered. And now we have companies in more and more industries that are almost more a tech company than they are a company in that industry.

I mean, some people would say, for instance, that Tesla is not just an automobile company. It's really more of a technology company.

Well, I think it's a really important point because, in some ways, I wonder if you worry about gaps opening up because as-- you gave the example of Amazon. In some ways, it is a tech company. A lot of people think of it as a tech company.

But it's a retail company like who are its competitors? I mean, Walmart, in some ways. And Netflix and Tesla, as you say, in some ways, they're tech companies. But in some ways, they're not tech companies, at all. And so I wonder if you feel like you worry that because some companies are doing better with that collaboration, gaps are opening up?

Oh, yes, I think those are all great examples of this phenomenon, Netflix, Amazon, Facebook, in a sense, it's either a publishing company, or an entertainment, company, or something--

That's true. That's true.

But all of these companies are ones that, in some sense, revolutionized their industries by using technology and organizing themselves in ways that were appropriate for doing that. They revolutionized their industries by inventing a new way of organizing work that took advantage of computers in ways that would never have been possible before.

So when businesses come to you, when you consult, or when you talk to folks and they worry about, let's say, they are an older company, or they're not one of these few companies that was maybe founded recently at this sort of tech human interface, what are ways in which they could use your observations on the idea of the super mind and use it to drive business decisions?

Well, especially in a rapidly changing economy, where innovations may matter as much as things like economies of scale and so forth used to matter in the past, often, the most important factor in business success is not just how productive are you, not how much you can increase your productivity, but how intelligent are you.

How intelligent is your company? How capable is it of sensing and responding to what's going on in the environment in intelligent and rapid ways? So I think that's useful way of thinking about things no matter what you're trying to do. And it's likely to become more and more useful and more and more industries.

So in terms of how companies can use this perspective of super minds, specifically to think about how to use technology, I actually have a number of concepts in the *Supermind's* book about that. And since the book was published, we've developed a methodology we call supermind design methodology.

Which kind of codifies some of those concepts in a way that we think will make it easy for people to ask themselves provocative questions to help think about that. For instance, the super mind design methodology is organized around a set of moves or conceptual moves, which are essentially questions you can ask yourself.

Some of them are fairly simple, like what are the parts of this. We call that zoom in. What is this a part of? We call that zoom out. What are analogies to this? We call that analogies.

But some of them are more specifically focused on how to organize a super mind. One is called groupify, which suggests you take whatever your problem is and think about, well, how could this be done by a hierarchy? Probably, we've already thought about that.

So how could this problem be solved by a market? Or how could this problem be solved by a democracy? Or how could it be solved by a community? And just thinking about those things, often, triggers ideas people might never have had before that might be innovative ways of organizing their organizations.

And that might be a very good idea about something to do. So one way then is to use this methodology. But we've also encoded this methodology in a new system we call super mind ideator. It's based on the GPT technology.

So it's based on generative AI. But on top of that underlying software, we've added some prompting, and some fine tuning, and some user interfaces. So you can just type in your problem, whatever it is, and the system will then automatically generate for you a whole bunch of ideas about things you can think of or things you could do to solve that problem.

So in a sense, it automates some of the creative thinking that people could do using just the methodology alone. But our expectation is not that you're going to type in the problem. And this system is going to give you the answer.

It's that you're going to type in the problem and this system is going to inspire you to think about a lot of things you might not have thought about. And you might come up with even better ideas than it came up with or pick a few good ones and realize that many of the ideas the system came up with are stupid. But you're still way better off than if you hadn't had anything to trigger your innovative thinking like that.

Right, right, right even if 18 of the 20 ideas are bad, that might leave you with two good ones. And that's not so bad.

Exactly, exactly.

It also occurs to me that even putting computers aside, obviously, we've talked a lot about computers. But when you talk about, oh, is this a problem, a business problem that a higher-- like how would a hierarchy solve it? Well, it's probably been done that way because that's how businesses are organized.

Generally, there's a boss. And there's whatever people under that person. You know, how would a democracy like more people being more equal, solve it, or whatever. It also suggests that cultures within organizations, again, putting computers aside, but just how the people deal with each other is extremely important.

Because, if the boss has all the answers, that may limit the number of answers you're going to get, right?

Exactly, exactly. So I think you're absolutely right that different styles, different types of super minds, like the ones we were listing, have different cultures and different kind of assumptions about humans and what their roles are. I think it's also interesting to observe, by the way, that even which of those methods you can use feasibly is affected by the technologies you have.

For instance, one of the reasons that hierarchies are so common today is that they do a pretty good job of economizing on the amount of communication you need. Basically, people send problems up the hierarchy. And then they get directions down the hierarchy.

And you just need a few links for that to happen. But for a democracy to work, everybody, in some sense, has to be informed about what the problems are or what the questions are. And everybody has to express an opinion through their vote.

And somebody has to count all that. And so there's a lot more communicating and counting needed to make a democracy work effectively than to make a hierarchy work effectively.

And potentially disagreeing may factor into that, too. There's a lot more disagreeing when there's more equality than when one person decides.

Exactly, but the advantage of using something like a democracy is that you get a lot more knowledge brought to bear on the problem, a lot more perspectives brought to bear on the problem. And you do have a way of making a final decision, just by counting the votes. So you can do that without computers.

But it's a very time consuming, expensive process, which is why we only vote every few years for most things in Democratic governments. But computers make it completely feasible, cheap, and easy if we want to do all that whole process much faster, much more easily, and possibly in ways that lead to a much better decisions.

So, of course, that doesn't mean we should make all decisions by democracies. There are lots of other factors to consider. But I think one of the key points is that these new technologies, in addition to being able to do some of the tasks that used to be doable only by humans, these new technologies can also enable new ways of organizing work because of things like reducing communication costs.

Now I know you're very excited about the possibilities for AI. But I wonder if you also feel like companies are too eager to think like AI is going to change everything. Like we can just throw a lot of compute power at a problem, and it'll be solved.

Yeah, I'm optimistic about what can happen. That's for sure. But I think right now, as with 1999 and the internet, we are in danger of and probably in a hype cycle that may have taken something that is, in fact, very good with very high promise and exaggerated its benefits and especially the speed with which they'll happen.

So we're probably setting ourselves up for a little bit of disillusionment somewhere along the way. But also, sure enough, basically, almost everything we thought about as happening next year in 1999, did eventually happen. It just took another 10 or, in some cases, 20 years before it was really a big deal.

And I think that's probably similar to what's happening with generative AI right now. Yes, some amazing things are now possible. And it's easy to see those. And it's easy to extrapolate all the other things that could be done.

But I think we often underestimate the time it takes to adjust human organizations to do things in a new way, the time it takes for people to learn new skills, the time it takes for people to just be willing to try things. Not everybody wants to be a first mover.

So I think there's real value there. I think the companies and other organizations that don't take advantage of this are in danger of being left behind. But I think it's also likely that we'll have some overoptimism about this for a little while longer and then some disillusionment.

Do you think it's going to be hard for leaders to get their companies to buy into a really new way of thinking about and using technology? Like is that transition, both in terms of people adjusting to a new approach, but also maybe learning as part of that things that they don't know right now?

Is that going to also slow things down? Is that going to be a tougher sell than maybe leaders think?

Well, if you're trying to make a change in a company that requires people to do things in a new way, especially if it requires them to learn new stuff that's not easy to learn or takes them out of their comfort zone, et cetera, that's usually not easy in any company for any reason. One of the well known, I guess, you could say morals or something in the organization theory, is how hard organizational change is, especially.

And I think this is probably telling, especially when you're at the top of the organization and trying to get the people down below you to change. That's harder than you might think, in part, because hierarchies aren't perfect order followers. People have their own incentives, their own ideas about what should be done.

In fact, to go back to the super mind metaphor, every real company has not only a hierarchy. But it has a kind of labor market. And every company has a set of communities that are part of it. And sometimes, the super minds fight with each other.

So sometimes, the hierarchy says, we want you to do X. And the community says, nah, we don't really want to do X. And that makes it really hard for the hierarchical managers to cause the changes they want to change, and think should be easy.

It turns out it's much harder if there are communities involved in the company that don't want to do that. So I think those things are hard no matter what, whether it's because of technology, or different working hours, or almost anything.

We've been talking mostly about an idea that you have that you've brought to companies to leaders that they've read about. I wonder what you've learned, maybe just in the last year or two as people have really-- as you said, like the super mind, in some ways, has really found its moment. But what have you learned from people at companies about this concept that maybe you didn't even think of when you were writing the book?

That's an interesting question. I can tell you one thing we've learned-- maybe two things from our research that's been surprising to me. One is we thought that it would be easy to combine people and computers in ways that led to huge improvements in performance.

It turns out to be a lot harder than I thought to do that. We've been doing experiments and measuring that. And yes, you can find improvements. But often, you find-- in fact, we've looked at other people's studies. And we often find that you combine people and computers, and it does worse.

Interesting.

So I think, for me, it's certainly too early to say it's a bad idea to combine people and computers. That must be wrong because there are examples where those combinations are amazingly powerful. And we've just been talking about some of them today.

Lots of others we could talk about, as well. But I think it means it's harder than it seems. And I think it means, in part, we need new theories or new guidelines for when and how to do this.

So I don't have all the answers to what those are. We've got some ideas that we're playing with that I think are very promising. But I think that's one thing I've learned. Another thing I've learned.

And this is more from our research than from people on the ground, but it is that it's surprisingly hard to evaluate these things, especially when you're talking about not how many spreadsheets did you complete this week. But how good is this advertising copy? Or how good is this short story?

Or how good is this image, this picture that's been generated? Of course, you know that's not trivial, but it's turned out to be surprisingly hard to do good evaluations, I think.

Yeah, that makes sense.

In terms of what have I learned from talking to people on the ground, I guess, I've been surprised at the variety of comments I hear. Some people I talk to say, oh, we thought about doing this, and it's just not working. I'm not going to work.

And I hear other people say, wow, I use ChatGPT for everything now. And it's incredibly helpful.

That's so interesting.

I think we still have to find our way in terms of where the real opportunities are.

Right, I have a final question for you, which is a little tangential, but I think does fit into some of what we've been talking about. I wonder how much you think about distraction? The computer scientists Cal Newport at Georgetown has written a lot about this, kind of the notion that, though, there's a lot to be gained, obviously, by joining together humans and computers, there's also a lot that can be taken away from your ability to concentrate and think deep thoughts and actually make progress on things.

And I wonder if you think about that kind of trade off of like that we're more and more always plugged in, but does that subtract something from our abilities?

Yeah, I think it's a very good question. Many of our technologies today like mobile phones that have newspapers on them and they have little things that come up every time. There's a headline that you might be interested in that shows you on the lock screen.

And I find myself checking the newspapers several times a day, at least, when I shouldn't be doing that, really. I probably check my email much more often than I need to. So I think it does affect our productivity.

In some sense, of course, we could always just say, oh, no, I'm not going to do that. But it's hard to not do that. In this particular case about distractions from mobile devices, we are, unfortunately, fighting against the power of market in which very wealthy companies have lots of incentives to capture our attention, whether it's good for us or not.

I think that actually has another moral to the story, which is the reason this happens is because those wealthy companies are supported by advertising. So their real customers are not the users. They're the advertisers.

If we had a different business model, like if they were supported by subscriptions, rather than advertising, they'd be much more motivated to do what their actual users wanted them to do. And so that's an example of how the kind of super mind in play can have a big effect on the outcomes that arise.

Thomas Malone is director of the MIT Center for Collective Intelligence. He's also a Patrick J. McGovern professor of Management at the MIT Sloan School of Management. Tom, Thank you so much. This is great.

Thank you, my pleasure.

[MUSIC PLAYING]

And if you're interested in learning more about the CSAIL Alliance's program and the latest research at CSAIL, please visit our website at cap.csail.mit.edu and grab this podcast on Spotify, Apple Music, or wherever you get your podcasts. I'm Kara Miller.

Our show is produced by Matt Purdy and Nate Caldwell with help from Audrey Woods. Tune in next month for a brand new edition of the CSAIL Alliances podcast, and stay ahead of the curve.

[MUSIC PLAYING]