McKinsey&Company

What is the future of work?

December 2017

A new podcast series from the McKinsey Global Institute explores how technologies like automation, robotics, and artificial intelligence are shaping how we work, where we work, and the skills we need to work.

The future of work is one of the hottest topics in 2017, with conflicting information from various experts leaving plenty of room for debate around what impact automation technology like artificial intelligence (AI) and robotics will have on jobs, skills, and wages. In the first episode of the *New World of Work* podcast from the McKinsey Global Institute—which is being featured in the *McKinsey Podcast* series—MGI chairman and director James Manyika speaks with senior editor Peter Gumbel about what these technologies are, how they will change work, and what new research says we can expect.

Podcast transcript

Peter Gumbel: Hello, and welcome to the MGI podcast. This is our new series on work, the world of work, and the changing world of work. Today, for our first podcast on this issue, I'm with James Manyika, who is the chairman and director of the McKinsey Global Institute; he's also a senior partner at McKinsey and is based in the San Francisco office.

James, this issue of work and the future of work is one that you have been looking at for some time, with work on automation and with the latest report on jobs, *Jobs lost, jobs gained*. Perhaps, you can start off by telling us about the broader issues, and which ones you're focusing on.

James Manyika: Well, I think we're having an interesting time in our history and our economy around the future of work. It comes up in almost every conversation with students, workers, CEOs, and policymakers. It's the topic of the day.

And typically, when this topic comes up, there are three or four issues embedded within it. First, there's the question and discussion around the impact of artificial intelligence, automation on work and jobs, and whether we'll have enough work and jobs left after that.

A second part of the conversation is around the changing models for work and work structure. This involves questions around independent work, the gig economy, and what people sometimes refer to as fissured work—whether people work as outsourced services or not.

And whether any of those kinds of evolved work models are going to become the future, and whether people can work effectively and sustainably and earn living wages with enough support—in that kind of world of more varied types of work.

The third topic that comes up is the income question. We know that most advanced economies, over the last decade, have seen this huge stagnation of incomes, at least wage-driven incomes for workers and households.

And so that ties into the inequality debate, and whether people work and earn enough to be able to make a living or not. And the question then is, Will technology make that even worse as we look forward? And then, finally, people are often asking the question, Just how does a workplace actually change?

These are questions about how work will be organized and how it will look in terms of people working alongside machines. All of these questions are embedded in this big topic called "the future of work." The work we're about to release has been a continuation of what we've been doing recently on automation, specifically, and on the impact of artificial-intelligence and autonomous systems, on this question around jobs, and will there be enough work, or are we going to create enough work to make up for what we're going to lose.

Peter Gumbel: This issue of automation, and the idea that it's taking away work, is one that is not at all new. If you look back in history, we've seen centuries of concern about machines coming and taking over work. Why is this discussion coming up in such a big way today?

James Manyika: The conversation is accelerated and has become heated probably for a couple of big reasons. One is, in the last few years, we've seen spectacular demonstrations and progress with artificial intelligence, autonomous systems, and robotics in a way that has been quite extraordinary.

Some people would argue we've made more progress in those systems in the last 5 or 6 or 8 years than we've seen in the last 50 years. This apparent rate of progress is what's probably changed the conversation. The other reason the conversation has changed is that, at least on the surface of it, compared with automation of the past, there's a sense that maybe this time it feels a little different.

> "There's a question whether we're really doing anything different or not; but it *feels* different—when you have machines that are able to do pattern matching better than human beings."

Now there's a big debate whether it's really different or not, but the perception is that it's different. The perception is that, in the past, we were basically adding muscle or mechanization to what people did: if you were digging a ditch, you got assisted by a machine;

if you were doing mechanical work, you were getting automation help. And we've done that for a very long time. Then there was also the sense that, well, even when we automated other tasks that were not mechanical, we were automating fairly routine work. This is work that you could write a script for, or a set of algorithms for, and then you got a machine to follow those rules, and it did the work.

What feels different is that we seem to be building machines that aren't just about adding muscle or automating routine tasks, but they seem to be doing wholly new different things. Getting at things that look like tacit knowledge, tasks that look like they're cognitive tasks. Tasks that you can't write an algorithm for a machine to do; so that's with techniques like machine learning, where it looks as if the machines are actually *learning* to do something—they're not being scripted to do something—they're discovering patterns, they're discovering things themselves.

So this idea that we seem to be doing something technically different is one of the reasons, again, it's come to the fore. Now, there's a question whether, in effect, we're really doing anything different or not; but it *feels* different—when you have machines that are able to do pattern matching better than human beings, and machines that are able to discover novel solutions to things—then it feels different to people. And so, people start to worry about what's left for human beings to do.

Peter Gumbel: Can you just say a couple of words about the technology that underlies machine learning and why it is now suddenly making such rapid progress?

James Manyika: Machine learning is essentially a set of techniques that take advantage of neural networks. We feed neural networks a lot of data, and they build up, through what are called "training algorithms," patterns of what the data mean; they build structure and sense out of that.

Within machine learning, you have particular areas like deep learning, and then you have areas like supervised learning. Supervised learning is when you use labeled data. The machine keeps trying, and there's a reference model that says, "Yeah, you got it right, didn't get it wrong: that's a cat, that's not a dog; that's a door, that's not a table." It's labeled by human beings. So you have these techniques, supervised learning, unsupervised learning, where the machine self-corrects. The techniques themselves have made progress.

The neural networks largely come in two flavors: convolutional neural networks and recursive neural networks. And they're good at slightly different things. Because of those techniques, we've now been able to do things that make classification much easier.

Classification has typically been applied to image recognition, facial recognition, and things where you're able to classify and organize patterns. We've applied similar techniques to natural-language processing, where you can process blocks of data and interpret and learn meaning out of them.

That kind of progress in patterning and in machine learning has been applied to machine vision. Also, because of techniques like machine learning that have been applied to natural-language processing, we've made enormous progress. When you couple that with what I'll describe as "systems-level progress"—where you put together these different systems: sensor systems, image-recognition systems, and navigation algorithms—you start to get driverless cars or autonomous systems, where you're putting together a collection of capabilities to do something in the physical or real world.

If you look at why there is so much progress this time around—versus the last 20 or 30 years, because none of these are new techniques—the progress comes down to about three things. One is, yes, the algorithmic techniques have made some progress. But progress comes from two other areas especially.

One is the amount of compute power we now throw at these problems, whether it's at the silicon level—now we apply not just typical CPUs [central processing units], but we now add GPUs [graphics processing units], and also, we now add compute clusters in the cloud, and so forth—you now have enormous compute power you can throw at the problem.

Then you have the third factor, which is the availability of data. Which is, we all now routinely add billions of pictures to the cloud every day, we add all kinds of voice technology and voice data streams all day long. We have huge amounts of data available to make these training algorithms work. Put all that together, and it's no surprise we've had breakthrough progress in the last five years.

"The question is: What does this mean for work? That's where a lot of concerns and anxieties come up."

Peter Gumbel: In the last report that came out in January on automation, you talk about the different activities that machines are able to take on in the workplace. But I think, at the same time, it's quite striking to see how you talk about the benefits and the upside, and the productivity gains from automation. That's an element that you seem to want to stress. Are you saying, then, that people have automation wrong? That the fears are wrong? That we should look at the benefits more closely?

James Manyika: We should look at both; I think it's worth spelling out both sides. On the benefits side, if you look at this from the point of view of businesses, for example, that are going to use the automation and AI machine-learning techniques, the benefits are relatively clear.

There are all kinds of performance improvements from reducing error rates, being able to do predictions better, and being able to discover novel, new solutions or insights. The benefits in a use-case sense to businesses are hard to disregard and ignore.

That's going to drive and encourage businesses to adopt these techniques—and they are. The benefits to the economy are also clear, because we know that—associated with most

automation technologies in the past, and even today, and in the future—automation of these systems improves productivity.

This is one of the mechanisms and ways in which we improve economy-wide productivity. And, at a time when we need more productivity growth in the economy, it's hard to ignore the contributions that these technologies bring to productivity, and hence, the economy, and hence, economic growth.

It's also hard to dispute some of the potential benefits and utility to people as users. We have now grown comfortable using technology, whether that's in voice-recognition assistance or in other techniques that are useful, to us, as users. Those benefits are clear: to the users, to the economy, and to business.

The question is: What does this mean for work? I think that's where a lot of concerns and anxieties come up, as to the impacts on work. What we do know is that if you look at most of the available technologies that have demonstrated the biggest impact—and we've looked at over 2,000 activities—these are activities that workers in the economy do. If you organize those activities into roughly eight categories, you have three categories of activities that are very easy to automate with the available technologies.

That's activities that involve data collection of one sort of another, activities that involve data processing of one sort or another, and activities that involve doing physical work in highly structured and predictable environments. Those three kinds of categories of activities, out of the total of about eight categories of activities, those three activities make up something like 51 percent of economic activity—and what people and workers do in an advanced economy like the United States.

That's a big part of what people do. Now to be perfectly clear, saying that 51 percent of activities are relatively easy to automate does not say that 51 percent of jobs are going to go away. The job question is a very different one, because we know that any one job consists of 20 or 30 different kinds of activities, aggregated into that job.

When you then ask the question of how many jobs, occupations, have a fair share, majority, or 90 percent or 100 percent of their activities that are easy to automate, you get a much smaller number—5 percent of occupations.

But then, what you also see is a host of other occupations—by our count 60 percent of occupations—that have about a third of their constituent activities that are easy to automate. That tells you that we're probably going to have more jobs *change* than disappear.

Because that 60 percent is a big chunk of what people actually do. The question of the impact on jobs and work is a much, much more nuanced and complicated one. And keep in mind, by the way, that while we talk about the impact of these technologies on work, it's only *one* thing to look at—the questions of technical feasibility.

In other words, What's now technically possible to automate? That's an interesting question, but that's just the first of four or five questions. The other questions include, What's it going to cost to develop and deploy those technologies? How does that play into labor-market dynamics in terms of the relative cost of having people do that? What is the availability of people who can do that task instead of a machine? What is the quality needed? What are the skills associated with the labor force?

These labor-market dynamics are another important consideration, as well as other ultimate questions about regulation and social acceptability, and so forth. The question of what the rate of adoption will be, and the extent of adoption, depends on many more factors beyond just technical feasibility.

Peter Gumbel: How long will it be, based on your research, before we start seeing a critical mass of automation being adopted in workplaces?

James Manyika: It's going to go occupation by occupation, technology by technology, and activity by activity. We've done research in a sample of about 46 countries, which are a mixture of developed economies and developing economies, and across that sample, it looks as if by 2030 you could imagine a range that has a midpoint, something like 16 percent of occupations would have been automated—and there would be impact and dislocation as a result of these technologies.

Now that number has a very wide range: at the low end, it could be very little, and at the high end, it could go all the way up to about 30 percent. The reason for that range is that it depends on the rates of adoption, the nature of the country, the wage dynamics in that country, and the wage dynamics in the sectors in that country.

While I said the midpoint across 46 countries is 16 percent, the midpoint for advanced economies would be much higher than that—in the 20 percent range; whereas, the midpoint for developing economies will be much, much lower, simply because their wage rates are much, much lower.

If you're in Japan, you could expect a higher percentage of your workforce to become automated, whereas if you're in India, you could expect a much lower percentage of your workforce to become automated—largely due to the different labor-market dynamics.

Peter Gumbel: You did say that there are going to be some difficult transitions. How difficult will they be? What sort of transitions we talking about?

James Manyika: We're going to see a few different kinds of transitions. The first one is that the mixture of occupations is going to shift. We know that when you take into account the activities that are easy to automate, relatively, and the ones that are relatively harder to automate, it will result in some occupations growing more than others.

What do I mean by that? For example, occupations that involve a lot of data gathering, data processing, or physical work are going to decline. The relatively harder occupations and

activities to automate, like care work and work that requires empathy, judgment, and so forth, those occupations are going to rise.

The mix of occupations is going to shift substantially. That means that people are probably going to have to move and be transitioned from certain occupations into new occupations, ones that are going to be growing. So that's one kind of transition.

Another kind of transition is going to be the skill requirements. We know that the skill requirements are going to shift for a couple reasons. One, because people are moving to new occupations that are going to require higher skills, often, in order to do those occupations; we know the skill requirements are going to go up, if only because people are going to be working alongside highly capable and increasingly capable machines.

In order for people to keep up, adapt, and work alongside effectively with highly capable machines, they will require a very different set of skills. So the skill transitions are going to be quite substantial. That's why we're having a conversation now, and we're starting to have a conversation about retraining and reskilling, especially for mid-career workers, who may have grown up in one environment with a certain set of skills and are now having to move into new occupations. Or, even if they're in the same occupation, that occupation now requires a higher level of skills in order to be valued and continue to be effective. The skill transitions are quite substantial. A third transition that I think we're going to have to think about is the potential impact of all of this on incomes and wages.

We know that the occupations that are going to be growing are the ones that historically haven't had the highest wage structures associated with them. We know that work in manufacturing always had slightly higher wages compared with, for example, work in activities like care work, where there are teachers or elder-care workers, and so forth. And so, we know that the mix of occupations that are growing—unless we change our minds on how we think about the value of that work—have not historically had higher wage structures associated with them.

We're going to have to deal with that. We're also going to have to deal with the fact that as workers transition from one occupation to another, they may require all kinds of support—like dislocation support—as they move from where they are to where they're going to be. We're going to have to rethink that.

That's particularly important at a time when, historically, most economies in the OECD [Organisation for Economic Co-operation and Development] have not always supported worker transitions as robustly as they could. In fact, worker-dislocation support has declined over the last 30 years.

At a time when we're probably going to need it even more, we're going to have to change our minds about how we think about worker-dislocation support. So all of these are some of the transitions that I think we're going to have to grapple with. And this is a matter not just for governments and policymakers but also for businesses and private-sector leaders, who are going to have to think about they will retrain their workforce.

How do they help redeploy their workforce as occupations and work change? How do they redesign work structures inside companies, to support different and new kinds of ways of work, so that there's enough work for everybody to do, as we manage our way through these transitions?

I'm relatively less worried about the question of whether there will be enough work for everybody. Of course, one can imagine scenarios where that would be the case. But I'm more worried, and focused on, these transition questions for workers, around skills, occupations, and the income and wage effects. I think that's where the real hard work's going to be.

Peter Gumbel: The last question is, this sounds like a wake-up call you're delivering to say, "You need to prepare as a government or as a business leader." Are you hearing those conversations taking place already in policy and business circles? Or is it still relatively early days for that?

James Manyika: The good news is that many forward-looking business leaders and policymakers are, in fact, thinking about these questions. You see quite a few examples of CEOs who are leading reskilling efforts throughout their workforce.

You see CEOS who are starting to think about these deployment and redeployment questions in their workforce. You see policymakers who are also starting to think about the right ways to approach these. But it hasn't quite become the widespread conversation that it needs to be. But I think this is what we need to be talking about now in terms of how to prepare for these choices.

I think one of the things we've learned from the research we've done is, as we started in this conversation, that there's two sides to this: on the one hand, I think we want companies, governments, and countries to embrace these technologies because of all the benefits that they bring to business and to the economy. We have to have an embracing conversation.

At the same time, we also have to face up to the transitions and challenges, and we have to help workers manage their way through this transition. The answer, in my mind, is not to slow everything down—because the problem with that is, if you slow down these technologies, you're also putting dampers, if you like, on business dynamism and on economic growth. And that's not a good thing, either. You want to embrace and manage the transition. I think that's the simultaneous challenge that we have ahead of us.

Peter Gumbel: Great, thank you very much, James. ◻

James Manyika is the chairman and director of the McKinsey Global Institute and a senior partner based in McKinsey's San Francisco office. **Peter Gumbel** is a senior editor for MGI and is based in the London office.

Copyright © 2017 McKinsey & Company. All rights reserved.