



The new division of labor

On our evolving relationship with technology

ANALYTICS AND INFORMATION MANAGEMENT

The world is changing, fast. Customers are demanding a more personalized experience. Competitors are appearing out of nowhere. Governments and regulators are demanding faster and more accurate reports. The most important answers to your complex business challenges may be hiding in your data. Those who can make sense of the vast quantities of information at their disposal are coming out on top. But in a world where the amount of information is doubling every two years, how can we keep up? For more information, visit us on <https://www2.deloitte.com/au/en/pages/technology/solutions/enterprise-information-management.html>

HUMAN CAPITAL

Our services leverage design-thinking, digital innovation, and data-driven industry insights to help you navigate today's business challenges. A new wave of human resource, talent, and organization priorities require HR teams to innovative talent, leadership, and change programs. For more information, visit us on https://www2.deloitte.com/au/en/pages/human-capital/topics/human-capital.html?icid=top_human-capital

Contents

A matter of trust | 2

Integrating human and machine | 4

The limitations of AI | 8

The value of labor? | 10

Conclusion | 13

Endnotes | 14

A matter of trust

WHEN THE COMPUTERIZED word processor supplanted the typewriter, responsibility for producing day-to-day business correspondence shifted from typist to manager. This made document production much more efficient as it eliminated the round trip from Dictaphone¹ through typing pool and back to the author for approval. This increase in efficiency didn't occur overnight though. Word processors may have promised to make managers more productive, but many initially balked at using them because it didn't suit their personal work habits, nor did they see it as a task fit for their station.

Crafting their own documents meant that managers needed to acquire composition and proofreading skills, as well as the more technical skills (such as typing and functional digital literacy) required to operate a word processor. But perhaps the most fundamental change they had to make was in their mindsets. To effectively use a word processor, they had to build a *relationship* with this new technology, trusting that colorful squiggly underlines actually indicated errors, that saved documents could be retrieved at will, and that the machine wouldn't crash (well, not too often anyway). Building this trust—integrating the new tool into their personal work habits—took time, as people were naturally loath to abandon old attitudes and behaviors.²

Most managers did come to embrace word processors, with the result that these types of machine not only became widespread, but fundamentally changed the way business correspondence is produced. Along with this change came a shift in social norms. Typing is now con-

sidered an important skill to be taught to children in school, while executives no longer see typing as beneath their role.

Without these changes in social norms, and without the shift in attitude from anxiety to acceptance, word processors may have remained on the fringes of society, or at least not have been used to their full potential. We might have simply replaced the typing pool's typewriters with word processors, preserving the old, inefficient Dictaphone-to-typist-to-author process for producing correspondence. Instead, we changed our entire approach to correspondence production—and were able to take better advantage of the new technology as a result.

Putting technology to effective use isn't only about recognizing the superiority of a new tool, or even just about learning how to use it. It's also a matter of emotional acceptance and social validation—factors that are at least as important as the intellectual understanding that the new technology is “better.” This is true for much more than word processors. From shipping containers³ to smartphones,⁴ both work habits and social norms had to change before the core technology could have a transformative impact.

Putting technology to effective use isn't only about recognizing the superiority of a new tool, or even just about learning how to use it. It's a matter of emotional acceptance and social validation.

We are now facing the same situation with another potentially transformative technology: artificial intelligence (AI). While AI-based technologies are rapidly gaining ground in both the workplace and society at large, it can be argued that they are overwhelmingly being used in a suboptimal fashion—to automate tasks that have traditionally been performed by humans, who then are considered redundant and eliminated. In a previous essay, we argued that this is unnecessarily limiting AI’s potential utility.⁵ There’s a growing body of evidence that solutions created collaboratively by humans and machines are different from, and superior to, solutions created by either humans or machines individually.⁶ If this is true, then it follows that humans, so far from being replaceable, are essential partners in realizing AI’s optimal value.⁷

Could our current social norms and attitudes be responsible for our failure to experiment with using AI-based tools collaboratively to define and solve problems *together*? It’s at least possible. Certainly, there’s no shortage of anxiety about technology in general and AI in particular. Even apart from the existential fear that robots will take our jobs, many people are concerned that the digital technology pervading society has (or will soon) become autonomous, emancipated, and independent from us, its parents. Some believe that technology is the prime factor shaping society and our lives, and that it is leading us in undesirable directions. Nor is it clear how we might trust the technology that’s seen as a contributor to “fake news,” among other problems. In some people’s opinion, the development of AI is the endgame for humans, the final nail in the coffin of a species that has unwittingly ceded control of its fate to its mechanical creations.

The root of these fears may lie in the instrumental relationship that we typically have with technology. We conceive of ourselves as tool “users,” the active agents, while our tools are relegated to the status of the “used.” We’re so accustomed to framing of our relationship with tools according to this instrumental “user/used” dichotomy that it’s difficult to

imagine it being any other way. Because of this, the only roles we can see for ourselves with respect to AI are that of either the “users,” the masters, or the “used,” the slaves—and with the technology seemingly about to gain a mind of its own, it’s anything but certain that we’ll wind up on top.

But AI is different from previous generations of tools, and not because it may be poised to develop sentience and take over the world. The difference is that our relationship with AI-enabled tools has the potential to evolve beyond being an instrument for particular tasks—and *that*, ironically, is because of AI’s very ability to mimic having “a mind of its own.”⁸ Because AI embeds reasoning, often quite complex reasoning, and the ability to act, in a tool, the tool is no longer just a passive instrument serving our needs. Rather than simply using such a tool, we’re now *interacting* with it, allowing it to make and act on decisions on our behalf, considering its responses, and changing our intentions as a result.

It’s possible that if we persist in approaching AI-based tools as mere instruments rather than as actors with (albeit limited) autonomy and agency in their own right, we will fail to realize AI’s full value—the equivalent of replacing typewriters with word processors while continuing to send correspondence through the typing pool. On the other hand, if we pursue a new, different relationship with AI—one that capitalizes on the technology’s decision-making capabilities—we can envision a new division of labor between humans and intelligent machines that takes better advantage of the strengths of each. It is through this process of constructing our new relationship that we will come to develop the ethics, the social norms, that constrain the agency of both humans and machines, determining what machines can do, what humans can direct the machines to do, and what is beyond the pale. What this new relationship with AI could look like, and what it might imply for society’s view of how work is valued and rewarded, is the subject of the rest of this essay.

Integrating human and machine

PRIOR TO THE industrial revolution, we had a simple, mechanical relationship with technology. Work implied craft and the craftsperson. Indeed, the final task of a journeyman (a craftsperson on the journey between apprentice and master) before graduating was to construct a *masterwork* demonstrating their proficiency in all aspects of their craft. Technology was little more than technique, with the craftsperson typically responsible for making and maintaining their own tools.

This relationship bifurcated in the industrial era, separating technology from technique as the technology and the tools that represented it became more complicated. Some workers specialized in technology—creating tools—while others focused on using them—technique. Each of these relationships require their own particular knowledge and skills. A weaver might operate a power loom, to pick one example, but typically would not know how to construct one. Cloth made by a mechanist, conversely, would be of poor quality no matter how detailed their knowledge of a loom’s inner workings. This divergence in roles has accelerated as technology has become increasingly complicated, resulting in tools such as computers or even the modern motor car that most of us use but few understand.⁹

In both the preindustrial and industrial eras, however, the tools in question, whether a craftsperson’s chisel or a machinist’s lathe, remained *objects*. All decision-making authority remained with the human, whether tool maker or tool user; the person was always the active agent. AI has the potential to change that, and in doing so, to drive another step

change in the way we relate to our now-intelligent tools.

AI enables us to codify decisions via algorithms: Which chess piece should be moved, what products are best to populate this investment portfolio? These decisions are made in response to a changing environment—our chess opponent’s move, a change in an individual’s circumstances. In fact, it’s the environmental change that prompts the action. This responsiveness to external stimuli is why it is often more natural to think of AI as mechanizing behaviors rather than tasks.¹⁰ “Task” implies a piece of work to be done regardless of the surrounding context; a behavior is a response to the world changing around us, something we do in response to appropriate stimulus, with the same stimulus in different contexts triggering different behaviors.

In both the preindustrial and industrial eras, decision-making authority remained with the human. AI has the potential to change that, and in doing so, to drive another step change in the way we relate to our now-intelligent tools.

Solutions that contain a set of automated behaviors have a degree of autonomy and agency, as they will react to some changes in the environment around them (constrained by what their set of

behaviors enables) without an operator’s intervention. This autonomy and agency might be relatively benign, such as a face-recognition behavior automatically tagging new holiday snapshots with the names of family members that it identifies, and getting a few wrong. Or it might be much more consequential, like a military drone that also has the authority to apply lethal force.

In any case, just as we sense and respond to the environment changing around us, so will the intelligent digital systems we have created. We will weave the digital behaviors of these systems with our own human behaviors, reacting to changes they make in the world at the same time as they react to changes that we make. It will be as if these solutions exist along with us in the organization chart, their autonomy and agency affecting our own. Rather than “digital solutions,” they can more properly be thought of as “digital agents.”

As we interact with digital agents to get work done—discovering different ways to effectively combine their capabilities with ours—our relationship with technology will be primed to evolve again to account for the sharing of authority between worker and tool. This, we posit, will add a new, tripartite layer on top of the “maker/user” dichotomy that arose in the industrial era:

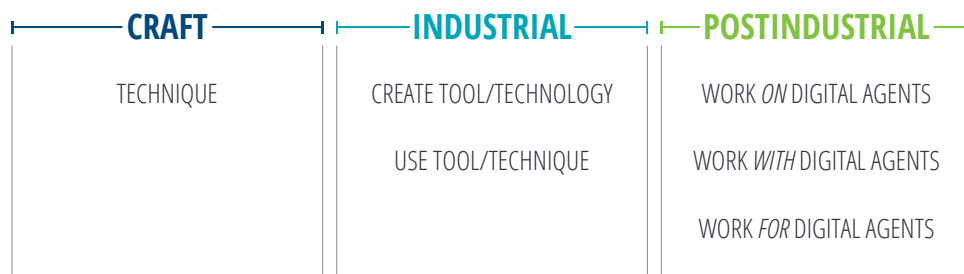
First, we can expect humans to work *on* machines, similar to how a manager might teach a worker how to perform a task. This might be (as is often assumed) via “coding,” such as when a developer takes engineering knowledge and programs it into a computer-aided design tool.¹¹ Knowledge can also be captured by example. A “truck driver” can teach an autonomous truck how to park in a particular loading bay by manually guiding it in the first time. A robot chef learns to cook a meal by observing, and then copying, the actions of a human.¹² Or our digital agent might be taught by preference. This is an iterative process where the agent responds to a range of stimuli, with a human trainer selecting the responses that best match what the agent needs to achieve.¹³ Over time the agent will infer what the most appropriate response is.

Just as we sense and respond to the environment changing around us, so will the intelligent digital systems we have created.

Second, humans will find themselves working *with* machines, when digital behaviors are used to

FIGURE 1

Humans can work *on*, *with*, or *for* digital agents



Source: Deloitte analysis.

complement our own human behaviors. An obvious example is a tumor-identification behavior used to augment an oncologist's ability to diagnose skin cancer by helping them locate potential tumors.¹⁴ In this category are automated behaviors that help us observe the environment around us. Machine vision might detect driver fatigue or distraction, either via image recognition or by observing brain wave activity. We can bolster our decisions by using digital behaviors that enumerate the available options, helping us weigh them objectively, and suggest the most likely course of action: anything from which movie to watch next to how best to treat a medical condition. Or digital behaviors can send the results of our decisions back out into the world. This might be as simple as setting a price, or as complex as steering an autonomous car or crafting a news story.¹⁵

Third, we will work *for* machines, when our activities are directed by a machine and the quality of our work gauged by them. Machines may delegate actions that they are unable to perform to a human. A human can drive a car (at least, until the cars can drive themselves) to deliver a package or person under the direction of a machine, as with many ridesharing services.¹⁶ Rather than piece work, a machine might specify a sequence of tasks, creating a schedule for a human worker that optimizes their time, and then tracks progress. In any case, the work is specified, and quality measured by the machine, without negotiation with the human.

Further, if we're to understand our possible future relationship with intelligent technology, then

we also need to consider that the relationship won't be one-sided. That is, we won't just work *on*, *with*, or *for* digital agents; digital agents will also work *on*, *with*, or *for* us—or the relationship may be roughly equal and collaborative. (We might say that a semi-autonomous car is assisting the human driver, for example, but we can just as well consider the human to be assisting the semiautonomous car.)¹⁷

Machines may delegate actions that they are unable to perform to a human.

Consequently, we can take our model for post-industrial work and make it two-dimensional (figure 2). The horizontal axis has columns for each of the working *for*, *with*, and *on* roles that we've already discussed, and represents the different ways authority is divided between human and machine. The new vertical axis represents the division of agency, and is also broken into three.

The top row of figure 2 captures how *our* relationship with technology is changing, where the machines have the agency while ours is limited to supporting them. The bottom row of figure 2 is the reverse, capturing how technology's relationship with us is changing, where we have the agency and the machines' agency is limited to supporting us. In the middle we have human-machine collaboration, where the two work as equals.

FIGURE 2

Our postindustrial relationship with technology



Source: Deloitte analysis.

The limitations of AI

ALL NINE MODES of the relationship we've described in the previous section represent a mix of human and machine behaviors, whether the human or the machine takes the lead in working *on*, *with*, or *for* each other. That doesn't mean that the two parties are interchangeable; AI has its own inherent limitations, meaning we expect a difference in the *kind* of behaviors that humans and machines contribute. Humans "think" whereas machines crunch numbers. That difference means that humans, with their ability to adapt to an almost infinite range of contexts, can behave in ways that create *new* knowledge—forming new goals, defining new problems, and framing new contexts in which to act. AI-enabled machines, on the other hand, apply knowledge and depend on the context or contexts provided them if they are to behave in the expected fashion. Their contribution is to embody *existing* knowledge, evincing it in behaviors that help to achieve a defined set of goals in a specified range of contexts—but they are unable to form new goals or redefine the context in which they are pursued as they are constrained to viewing the world in the way that their creators have framed it.

To understand why, consider that, before a task can be automated, it must be mechanized, the actions required to prosecute the task realized in gears, pulleys, levers, and pistons. To be mechanized, a task must be precisely defined and the context within which the mechanism will operate carefully controlled. Mechanized tasks might be precise but they are also inflexible, requiring everything to be *just so*.

Artificial intelligence has a similar dynamic.

Just as tasks can be mechanized and automated, algorithms that give rise to behaviors can similarly

be mechanized and then automated. For Alpha Go, Google's champion Go computer, this was done with a Markov Decision Process (a mathematical framework for managing decisions when outcomes are partly under the control of the decision-maker).¹⁸ Once mechanized, we can apply the magic of digital computing to give the algorithm a life of its own—to turn it into a digital agent. However, if our digital agent is to be reliable, then we need to carefully control the context (the data inputs, any outputs, and the time and resources allowed) it operates in. Even small disturbances to this context can result in dramatically different, and potentially incorrect, answers.

AI has its own inherent limitations, meaning we expect a difference in the kind of behaviors that humans and machines contribute.

For instance, a complex solution such as an autonomous car contains many mechanized behaviors: lane following, obstacle avoidance, separation maintenance, and so on. Despite their apparent sophistication, however, these solutions contain relatively few behaviors compared to a human, and each behavior is comparatively unsophisticated and narrowly defined. An autonomous car cannot use eye contact with other drivers, for example, to determine who will take the right of way at a confusing intersection. Consequently, the current generation of (semi) autonomous cars can only operate in tightly constrained environments, and they require humans to augment them when their limited set of behaviors fails.

This is why mechanizing and then automating algorithmically determined behaviors will not result in an “AI,” an artificial being, or “AGI,” an artificial general intelligence, that is capable of creating new knowledge.¹⁹ The set of behaviors is too limited, individual behaviors are too narrowly defined, and the digital agent too inept at acquiring new behaviors without human intervention for the machine to be considered intelligent. The autonomy exhibited by an AI algorithm is limited, as the algorithm will only consider the data that it is configured to consider, and

will only reason in the way it’s designed. Its agency is similarly limited as it can only see the world in terms of how its creators framed it, can only choose among the behaviors configured into it, and can only affect the world (have the agency) in ways that we allow it to. And the solution can only be assumed to be operating correctly if the context it operates in is exactly the same as the one its behaviors were designed for (such as on a freeway where vehicles don’t stop unexpectedly). Otherwise it will fail.

The value of labor?

IF WORK IS divided along the lines of knowledge creation versus knowledge application—with humans responsible for the former and digital agents for the latter—the role that (human) workers play in production will change, as well as the relationships between a firm and its customers and between the firm and its workers.

Let's first consider the change in workers' role in production. In the past, production was divided into a series of specialized tasks, with humans responsible for complex tasks, and simple tasks delegated to machines.²⁰ Human and machine worked in sequence. In the future, if we choose to relate to AI in the manner described above, work will be allocated to humans and machines, not along the lines of complex vs. simple, but according to whether the work requires knowledge to be created vs. knowledge to be applied. Humans will make sense of the world and machines will plan and then deliver suitable solutions, with production built around discovering, framing, and solving problems.²¹ Human and machine work together.

In the future, work will be allocated to humans and machines according to whether the work requires knowledge to be created vs. knowledge to be applied.

An example can be seen in the challenge of selling retirement financial products.²² No one wants to buy financial products as their ultimate goal; rather, they want to fund a happy retirement. The problem is that it's hard to describe just what a

happy retirement is for a particular individual. At best they'll have a vague idea—playing golf, traveling, or investing time in a neglected hobby, are all common scenarios. First, the individual needs to determine that what they think will make them happy will *actually* make them happy. Next, they need to establish reasonable expectations, based on their anticipated savings and future earnings. Then they need to consider how they might change their attitudes and behaviors today—taking lunch to work rather than buying it—to change their retirement. Finally, once they know what will actually make them happy, have reasonable expectations, and adopt suitable attitudes and behaviors, they will know their income streams, timelines, and appetite for risk. At this point it's easiest to just press the 'robo-invest' button.

In this example we can see a new division of labor. The human works with the client to discover what the client's "happy retirement" might be, to frame and define the problem, as it were. They're responsible for managing the interaction between what's desired (the client's dreams) and what's possible (financial constraints and opportunities) to discover what's practical (the "happy retirement"), engaging in something like a Socratic dialogue. The AI-enabled robo-adviser is responsible for developing and updating an investment plan based on the client's requirements as they are discovered and evolve. Finally, an operations platform executes the trades required to populate the portfolio. The human discovers, the digital agent plans, and more conventional automation delivers.

We can see similar shifts across the postindustrial landscape. Bus drivers (an example from our

last essay)²³ won't spend their time driving buses (after all, the buses can drive themselves), but will spend time managing disruptions, supporting riders, and optimizing the operation of the bus network. Factory workers might have roles within the factory, but the value they provide is in discovering ways to improve the factory's operation, something a highly automated factory cannot do on its own.²⁴ Sales staff will work with clients to understand how the store can help them with whatever problem brought the client across the store's threshold (whether or not that threshold is digital or physical) and build their relationship with the store, rather than operate the till and check prices. And so on ...

In this envisioned division of labor between human and intelligent machine, the bulk of the value is created in the discovery and framing of the problem, as planning and execution are automated and commoditized. But just because society *could* divide work between humans and machines in this way doesn't mean that such a state of affairs is inevitable. On the contrary, achieving this new division of labor will require profound change from most businesses, because in today's society and economy, the work that businesses define as valuable—and pay people to perform—is *not* necessarily that of discovering and framing the problem.

In the retirement adviser example, the client receives the most value from the Socratic dialogue in which they discover what their happy retirement might be, and which helps them develop the attitudes and behaviors that will get them there. From the client's point of view, the mechanical process of populating an investment portfolio has relatively little value, as it represents commoditized knowledge and skills. Investment firms, however, frame their relationship with the client in terms of a process that takes client details (investment goal, appetite for risk, income streams) and creates and maintains an investment portfolio to suit. The

value for the firm is in the tasks that match client funds to investment products—activities that the client places very little value in. The firm's desire to measure the number of clients served by an adviser, and the funds invested in each instance, is also disconnected from when the client derives the most value from the relationship.

In this envisioned division of labor between human and intelligent machine, the bulk of the value is created in the discovery and framing of the problem, as planning and execution are automated and commoditized.

Ideally, from the client's point of view, the firm would approach the client at the start of their career when they are forced to first consider saving for their retirement. While a young person's understanding of what a "happy retirement" might look like may be vague, it's important for them to develop the attitudes and behaviors that will set them on the right path as the magic of compound interest and share market returns needs time to work. However, an investment firm can expect to derive little revenue from a client at this time in their lives. The metrics it uses to measure financial advisers—investments made, funds under management—instead drive the advisers to approach individuals near the peak of their earning capacity (in their late 40s and early 50s, once the kids are out of home), as these individuals have the most money to invest. Unfortunately, this is often too late, as their investment plans have already been made, or they've established a lifestyle which makes it difficult to free up funds to invest. Nor can they invest for the long term.

For the firm to derive as much value as possible from the relationship, it needs to establish

the relationship early, at the start of the career, investing the time and the effort to instill the attitudes and behaviors that will carry the client into retirement. The relationship also needs to be maintained, as the client's understanding of what their happy retirement might be evolves, or their life circumstances change. This is where value for the client is being created—and where early investment in the relationship will eventually pay off for the firm—yet investment firms typically do not reward advisers well for this kind of activity. This disconnect between where and how value is created for clients, and where value is realized by the firm, might be a contributing factor to the breakdown in trust between financial institutions and their clients.

It's our view that humans and intelligent machines work together optimally when humans create knowledge, machines apply knowledge, and then humans explain the decisions that have been

made. This approach can also create more value for customers, changing the relationship between firm and customer. However, if we don't value and reward humans for doing the kind of work that comes with an optimal human-AI relationship—if our social and economic systems persist in framing work in terms of tasks completed, and to value labor in terms of its ability to prosecute these tasks—then we can expect AI solutions to continue to be used as they often are today: as cost-cutting enablers, substitutes for humans instead of partners with humans. If we don't ensure that our digital agents' behaviors exhibit our own core values and combine effectively with those of our employees, or if we unnecessarily constrain their behaviors so that they cannot adapt to the evolving context the firm finds itself in, then we won't be able to realize their potential or that of our employees. And as long as this remains the case, we can expect the potential value in this new generation of solutions to go unrealized.

Conclusion

AT THE DAWN of the industrial revolution, Benjamin Franklin observed that “man is a tool-making animal.”²⁵ Homo faber, the tool-making human, seemed to be supplanting homo sapiens, the thinking human. Some 200 years later we are increasingly asking ourselves if we have successfully managed our tools, or if they are managing us. Max Frisch, writing in the aftermath of the Second World War, felt that the later was true, and we were being unmade by our own technological ability.²⁶ This is over simplifying our relationship with technology though, a relationship that dates back at least to the invention of cooking.²⁷ We have evolved with our technology and our technology has evolved with us, as has the relationship between us and our tools. Each shapes the other.

Before our current instrumental relationship, with us playing the role of tool maker and tool user, we had a more craft-based relationship that didn’t make such a strong distinction between the making and using of tools. Today, the development of digital systems—systems that embody automated reasoning, with a (limited) degree of agency—is forcing our relationship with technology to evolve again. It’s often written that humans will be replaced by ever more capable machines, or that machines will augment humans. But both these points of view are limited, as they’re rooted in the idea that we are either tool user or tool maker. In the future we can expect to have a more active relationship with technology. Our tools are no longer simply *our tools*; they are taking on a life of their own. A digital agent could easily be your boss, coworker, and subordinate all at the same time, or it might take the lead with you supporting it; you might lead while the

digital system supports you, or the work might be a collaborative effort.

Today, the development of digital systems—systems that embody automated reasoning, with a (limited) degree of agency—is forcing our relationship with technology to evolve again.

We will have a much richer and more nuanced relationship with digital systems, our so-called “digital agents,” than we did with the tools of the industrial era. Digital agents, like all technology, are more precise and relentless than humans making them in many circumstances, more reliable. They are limited, however, as they cannot make sense of the world like a human can, and are unable to step outside the algorithmic box they live in. Digital agents can learn to be more precise and efficient over time, but they can’t learn how to do something differently. Humans, on the other hand, notice the new or unusual and engage in the collaborative processes that results in new insights and knowledge. Our new relationship with technology will be founded on this difference.

However, successfully adopting the next generation of digital tools, autonomous tools to which we delegate decisions and that have a limited form of agency, requires us to acknowledge this new relationship. At the individual level, forming a productive relationship with these new digital tools requires us to adopt new habits, attitudes,

and behaviors that enable us to make the most of these tools. At the enterprise level, the firm must also acknowledge this shift, and adopt new definitions of value that allow it to reward workers for contributing to the uniquely human ability to create

new knowledge. Only if firms recognize this shift in how value is created, if they are willing to value employees for their ability to make sense of the world, will AI adoption deliver the value they promise.

Endnotes

1. Dictaphone was a brand of dictation machine that can trace its heritage back to the Volta Laboratory established by Alexander Graham Bell in Washington, D.C. in 1881. The name was trademarked in 1907, and spun out as a separate company in 1923. Since then “Dictaphone” has become a generic term for referring to any dictation machine.
2. Sometimes this resistance to adopting new work habits can leave one far back on the technology adoption curve, such as George R.R. Martin who uses a 1970s word processor (WordStar 4.0) to write *Game of Thrones*. See: Neha Prakash, “George R.R. Martin writes ‘Game of Thrones’ on a ‘70s word processor,” Mashable, May 15, 2014.
3. When we think of shipping containers we typically picture the 20- or 40-foot long ISO 6346 intermodal container that became a common sight with the development of the world-wide intermodal container network in the 1950s and 1960s. The intermodal shipping container, as a technology, however, dates back at least to the early 1800s, and was first standardized in the 1933. It wasn’t until Malcom McLean, a trucking magnate, developed the current containerization system, requiring the development of new docks and loading systems, and changes in labor relations, that the intermodal container became a common sight. See: Marc Levinson, *The Box: How the Shipping Container Made the World Smaller and the World Economy Bigger* (Princeton University Press, 2016).
4. The iconic smartphone, the exemplar of all modern smartphones, is the iPhone 3G, released in 2008. The smartphone category, though, dates back to the early 1990s, with smartphones prior to the iPhone 3G only seeing limited adoption. With widespread adoption came the need to develop social norms governing smartphone use. Initially, for example, it was common for owners to use their smartphones whenever they wanted, though over time social norms emerged limiting their use, such as the norm of setting one’s smartphone aside during a dinner conversation.

5. Peter Evans-Greenwood, Alan Marshall, and Matthew Ambrose, *Reconstructing jobs: Creating good jobs in the age of artificial intelligence*, Deloitte Insights, July 8, 2018.
6. Garry Kasparov, "The chess master and the computer," *New York Review of Books*, February 11, 2010.
7. Peter Evans-Greenwood, Harvey Lewis, and Jim Guszczka, "Reconstructing work: Automation, artificial intelligence, and the essential role of humans," *Deloitte Review* 21, July 31, 2017.
8. Machines that seem to have a mind of their own are quite an old phenomenon. Alan Turing developed in 1950 what came to be known as the Turing Test, a test to see whether or not a machine could exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human. The test was defeated in the 1960s by ELIZA which used a collection of rules triggered off by keywords to generate seemingly intelligent responses.
9. It's not that our old craft-based relationship with technology was supplanted by these new industrial relationships, of course. Both exist simultaneously—but the industrial relationship has come to dominate.
10. One of the central points of our previous essay, "Reconstructing work," was that work enabled by AI should be constructed around problems and behaviors, rather than products and tasks. See: Evans-Greenwood, Lewis, and Guszczka, "Reconstructing work."
11. Alpha Zero is even an example of this, though it's often claimed that the solution taught itself to play chess. In reality, the rules of chess were encoded into the features that feed the neural networks and Monte Carlo tree search at the heart of Alpha Zero, and then two instances were played off against each other to explore these rules and find optimal moves.
12. Moley Robotics has developed a robotic kitchen featuring a fully functional robot integrated into a professional kitchen, which can be taught by example.
13. Paul Christiano et al., "Deep reinforcement learning from human preferences," Cornell University, January 12, 2017.
14. We might call this the *Where's Wally* problem. While humans might be better at classifying potential tumors (as machines have a higher false-positive rate than humans), we are not as good at finding potential tumors in an image as machines. We believe that we see the entire image when, really, we only see a small area where our attention is focused as it meanders over the image, with our mind filling in the rest of the details to create the illusion of seeing the entire image. This is why it's difficult to find Wally in an image filled with similar likenesses, as we only see Wally if we happen to look directly at him. AI, on the other hand, methodically scans the entire image. It might not be as accurate at identifying something as complex as a tumor, but it has a much better chance of finding it.
15. Tim Simonite, "Robot Writing Moves from Journalism to Wall Street," *MIT Technology Review*, accessed February 19, 2018.
16. The main distinction between a taxi dispatch service and a ridesharing service is that a dispatch service will announce jobs while a ridesharing service assigns jobs, and if you refuse too many jobs you're banned from the service. With a dispatch service, your assignments are your own concern, whereas with ridesharing the computer effectively controls them. Ridesharing services also use ratings to measure performance, where the computer collects the ratings and computes the result by using a somewhat opaque algorithm.
17. Alex Roy, "The half-life of danger: The truth behind the Tesla Model X crash," *The Drive*, April 16, 2018.
18. It's worth noting that the Markov Decision Process (MDP) was invented in the late 1950s, and was initially used in process engineering, and is otherwise unrelated to AI. The AI community first became interested in the MDP in the early 2000s.
19. An artificial general intelligence (AGI) is a machine with sufficient intelligence to successfully perform any intellectual task that a human being can.

20. We note that the difference between complex and simple tasks is a somewhat arbitrary definition, as what was complex in the past might well be manageable today and simple in the near future.
21. Evans-Greenwood, Lewis, and Guszczka, "Reconstructing work."
22. Evans-Greenwood, Marshall, and Ambrose, *Reconstructing jobs*.
23. Ibid.
24. Craig Trudell, Yuki Hagiwara, and Ma Jie, "'Gods' edging out robots at Toyota facility," *Japan Times*, April 7, 2014.
25. As quoted in James Boswell, *Life of Samuel Johnson* (1791) (entry for April 7, 1778).
26. Max Frisch, *Homo Faber: A report* (San Diego: Harcourt Brace Jovanovich, 1994).
27. Claude Lévi-Strauss argued that our definition of ourselves as not being animals is consistently related to the techniques and technologies of food preparation. See Claude Lévi-Strauss, *Le cru et le cuit* (Paris: Plon, 1964).

About the Deloitte Center for the Edge

The Deloitte Center for the Edge, which is part of our global research network, is designed to help businesses profit from emerging technology opportunities.

Technology developments strike at the very heart of a business and finding strategies for corporate growth is essential. Our mission is to identify and explore opportunities that aren't yet on senior management's agenda but should be. While we're focused on long-term trends and opportunities, we also look at the implications for near-term action.

As a client, you can benefit not only from the Centre's international research, but also from the Australian chapter's research on issues affecting the local operating environment.

About the authors

PETER EVANS-GREENWOOD is currently a fellow at The Deloitte Australia Center for the Edge—helping organizations embrace the digital revolution through understanding and applying what is happening on the edge of business and society. Evans-Greenwood has spent 20 years working at the intersection between business and technology. These days he works as a consultant and strategic adviser on both business and technology sides of the fence. He is based in Melbourne.

ROBERT HILLARD is Deloitte's chief strategy & innovation officer and is responsible for positioning the firm to tackle the disruption of technology, new competitors, challenging economic conditions, and changing regulatory priorities. Throughout his career, he has retained his focus on technology issues, and is passionate about shaping the firm's response to the huge changes to Australia's society and economy as a result of technology, digital disruption, and the shift to an information economy. He is particularly known for his work in business transformation, digital disruption, and information management through governance or board roles with professional and industry associations, his books, media, and client work. Hillard is based in Melbourne.

ALAN MARSHALL is the national lead partner in Deloitte's Analytics and Cognitive practice in Australia, working with clients to establish cognitive and analytic platforms that augment human decision-making and drive automation. Marshall is a thought leader in the application of automation and artificial intelligence in an organizational context. He is also well known as a specialist in developing strategies and solutions for clients that utilize data assets to improve yield on decision-making or productivity in process execution. His experience is wide ranging and includes driving outcomes for clients across the oil & gas, mining, utilities, and higher education sectors. He is based in Perth.

Contacts

Alan Marshall

National lead, Analytics and Cognitive
Partner
Deloitte Consulting Pty Ltd
+61 8 9365 8139
almarshall@deloitte.com.au

David Brown

National lead, Human Capital
Partner
Deloitte Consulting Pty Ltd
+61 2 9322 7481
davidbrown@deloitte.com.au

Deloitte.

Insights

Sign up for Deloitte Insights updates at www.deloitte.com/insights.



Follow @DeloitteInsight

Deloitte Insights contributors

Editorial: Junko Kaji, Ramani Moses, Preetha Devan, Rupesh Bhat, and Nairita Gangopadhyay

Creative: Molly Woodworth

Promotion: Ankana Chakraborty

Cover artwork: Emily Moreano

About Deloitte Insights

Deloitte Insights publishes original articles, reports and periodicals that provide insights for businesses, the public sector and NGOs. Our goal is to draw upon research and experience from throughout our professional services organization, and that of coauthors in academia and business, to advance the conversation on a broad spectrum of topics of interest to executives and government leaders.

Deloitte Insights is an imprint of Deloitte Development LLC.

About this publication

This publication contains general information only, and none of Deloitte Touche Tohmatsu Limited, its member firms, or its and their affiliates are, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your finances or your business. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser.

None of Deloitte Touche Tohmatsu Limited, its member firms, or its and their respective affiliates shall be responsible for any loss whatsoever sustained by any person who relies on this publication.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as "Deloitte Global") does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the "Deloitte" name in the United States and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.

About Deloitte Australia

In Australia, the member firm is the Australian partnership of Deloitte Touche Tohmatsu. As one of Australia's leading professional services firms, Deloitte Touche Tohmatsu and its affiliates provide audit, tax, consulting, and financial advisory services through approximately 6,000 people across the country. Focused on the creation of value and growth, and known as an employer of choice for innovative human resources programs, we are dedicated to helping our clients and our people excel. For more information, please visit Deloitte's web site at www.deloitte.com.au.

Liability limited by a scheme approved under Professional Standards Legislation.

Member of Deloitte Touche Tohmatsu Limited© 2019 Deloitte Touche Tohmatsu.