

Deloitte Review



Reconstructing Work

Automation, artificial Intelligence,
and the essential role of humans

PESSIMIST OR OPTIMIST?

WILL pessimistic predictions of the rise of the robots come true? Will humans be made redundant by artificial intelligence (AI) and robots, unable to find work and left to face a future defined by an absence of jobs? Or will the optimists be right? Will historical norms reassert themselves and technology create more jobs than it destroys, resulting in new occupations that require new skills and knowledge and new ways of working?

The debate will undoubtedly continue for some time. But both views have been founded on a traditional conception of work as a collection of specialised tasks and activities performed mostly by humans. As AI becomes more capable and automates an ever-increasing proportion of these tasks, is it now time to consider a third path? Might AI enable *work itself* to be reconstructed?

It is possible that the most effective use of AI is *not* simply as a means to automate more tasks, but as an enabler to achieve higher-level goals, to create more value. The advent of AI makes it possible indeed, desirable – to reconceptualise work, not as a set of discrete tasks laid end to end in a predefined process, but as a collab-

orative problem-solving effort where humans define the problems, machines help find the solutions, and humans verify the acceptability of those solutions.

CONSTRUCTING WORK

PRE-INDUSTRIAL work was constructed around the product, with skilled artisans taking responsibility for each aspect of its creation. Early factories (commonly called *manufactories* at the time) were essentially collections of artisans, all making the same product to realise sourcing and distribution benefits. In contrast, our current approach to work is based on Adam Smith's division of labor,¹ in the form of the task. Indeed, if we were to pick one idea as the foundation of the Industrial Revolution it

would be this division of labour: Make the coil spring rather than the entire watch.

Specialisation in a particular task made it worthwhile for workers to develop superior skills and techniques to improve their productivity. It also provided the environment for the task to be mechanised, capturing the worker's physical actions in a machine to improve precision and reduce costs. Mechanisation then begat automation when we replaced human

This impressive history of industrial automation has resulted not only from the march of technology, but from the conception of work as a set of specialised tasks.

power with water, then steam, and finally electric power, all of which increased capacity. Handlooms were replaced with power looms, and the artisanal occupation shifted from weaving to managing a collection of machines. Human computers responsible for calculating gunnery and astronomical tables were similarly replaced with analogue and then digital computers and the teams of engineers required to develop the computer's hardware and software. Word processors shifted responsibility for document production from the typing pool to the author, resulting in the growth of departmental IT. More recently, doctors responsible for interpreting medical images are being replaced by AI and its attendant team of technical specialists.²

This impressive history of industrial automation has resulted not only from the march of technology, but from the conception of work as a set of specialised tasks. Without specialisation, problems wouldn't have been formalised as processes, processes wouldn't have been broken into well-defined tasks, and tasks wouldn't have been mechanised and then automated. Because of this atomisation of work into tasks (conceptually and culturally), jobs have come to be viewed largely as compartmentalised collections of tasks. (Typical corporate job descriptions and skills matrices take the form of lists of tasks.) Job candidates are selected based on their knowledge and skills, their ability to prosecute the tasks in the job

description. A contemporary manifestation of this is the rise of task-based crowdsourcing sites – such as TaskRabbit³ and Kaggle,⁴ to name only two – that enable tasks to be commoditised and treated as piecework.

DOES AUTOMATION DESTROY OR CREATE JOBS?

A I demonstrates the potential to replicate even highly complex, specialised tasks that only humans were once thought able to perform (while finding seemingly easy but more general tasks, such as walking or common sense reasoning, incredibly challenging). Unsurprisingly, some pundits worry that the age of automation is approaching its logical conclusion, with virtually all work residing in the ever-expanding domain of machines. These pessimists think that robotic process automation⁵ (RPA) and such AI solutions as autonomous vehicles will destroy jobs, relegating people to filling the few gaps left in the economy that AI cannot occupy. There may well be more jobs created in the short term to build, maintain, and enhance the technology, but not everyone will be able to gain the necessary knowledge, skills, and experience.⁶ For example, it seems unlikely that the majority of truck, bus, or taxi drivers supplanted by robots will be able to learn the software development skills required to build or maintain the algorithms replacing them.

Further, these pessimists continue, we must consider a near future where many (if not all)

low-level jobs, such as the administrative and process-oriented tasks that graduates typically perform as the first step in their career, are automated. If the lower levels of the career ladder are removed, they will likely struggle to enter professions, leaving a diminishing pool of human workers to compete for a growing number of jobs. Recent advances in AI prompt many to wonder just how long it will be before AI catches up with the majority of us. How far are we from a future where the only humans involved with a firm are its owners?

Of course, there is an alternative view. History teaches that automation, far from destroying jobs, can and usually does create new jobs, and not just those for building the technology or training others in its use. This is because increased productivity and efficiency, and the consequent lowering of prices, has historically led to greater demand for goods and services. For example, as the 19th century unfolded, new technology (such as power looms) enabled more goods (cloth, for instance) to be produced with less effort;⁷ as a consequence, prices dropped considerably, thus increasing demand from consumers. Rising consumer demand not only drove further productivity improvements through progressive technological refinements, but also significantly increased demand for workers with the right skills.⁸ The optimistic view holds that AI, like other automation technologies before it, will operate in much the same way. By automating more and more complex tasks, AI could potentially

reduce costs, lower prices, and generate more demand – and, in doing so, create more jobs.

THE PRODUCTIVITY PROBLEM AND THE END OF A PARADIGM

OFTEN overlooked in this debate is the assumption made by both camps that automation is about using machines to perform tasks traditionally performed by humans. And indeed, the technologies introduced during the Industrial Revolution progressively (though not entirely) did displace human workers from particular tasks.⁹ Measured in productivity terms, by the end of the Industrial Revolution, technology had enabled a weaver to increase by a factor of 50 the amount of cloth produced per day;¹⁰ yet a modern power loom, however more efficient, executes the work in essentially the same way a human weaver does. This is a pattern that continues today: For example, we have continually introduced more sophisticated technology into the finance function (spreadsheets, word processing, and business intelligence tools are some common examples), but even the bots of modern-day robotic process automation complete tasks in the conventional way, filling in forms and sending emails as if a person were at the keyboard, while “exceptions” are still handled by human workers.

We are so used to viewing work as a series of tasks, automation as the progressive mechanisation of those tasks, and jobs as collections of tasks requiring corresponding skills, that it is difficult to conceive of them otherwise. But

The lesson here is that human and machine intelligence are different in complementary, rather than conflicting, ways. While they might solve the same problems, they approach these problems from different directions.

there are signs that this conceptualisation of work may be nearing the end of its useful life. One such major indication is the documented fact that technology, despite continuing advances, no longer seems to be achieving the productivity gains that characterised the years after the Industrial Revolution. Short-run productivity growth, in fact, has dropped from 2.82 (1920–1970) to 1.62 percent (1970–2014).¹¹ Many explanations for this have been proposed, including measurement problems, our inability to keep up with the rapid pace of technological change, and the idea that the tasks being automated today are inherently “low productivity.”¹² In *The Rise and Fall of American Growth*,¹³ Robert Gordon argues that today’s low-productivity growth environment is due to a material difference in the technologies invented between 1850 and 1980 and those invented more recently. Gordon notes that prior to the Industrial Revolution mean growth was 1.79 percent (1870–1920),¹⁴ and proposes that what we’re seeing today is a reversion to this mean.

None of these explanations is entirely satisfying. Measurement questions have been de-

bated to little avail. And there is little evidence that technology is developing more rapidly today than in the past.¹⁵ Nor is there a clear reason for why, say, a finance professional managing a team of bots should not realise a similar productivity boost as a weaver managing a collection of power looms. Even Robert Gordon’s idea of one-time technologies, while attractive, must be taken with a grain of salt: It is always risky to underestimate human ingenuity.

One explanation that hasn’t been considered, however, is that the industrial paradigm itself - where jobs are constructed from well-defined tasks - has simply run its course. We forget that jobs are a social construct, and our view of what a job is, is the result of a dialogue between capital and labour early in the Industrial Revolution. But what if we’re heading toward a future where work is different, rather than an evolution of what we have today?

SUITABLE FOR NEITHER HUMAN NOR MACHINE

CONSTRUCTING work around a pre-defined set of tasks suits neither human nor machine. On one hand, we

have workers complaining of monotonous work,¹⁶ unreasonable schedules, and unstable jobs.¹⁷ Cost pressure and a belief that humans are simply one way to prosecute a task leads many firms to slice the salami ever more finely, turning to contingent labour and using smaller (and therefore more flexible) units of time to schedule their staff. The reaction to this has been a growing desire to recut jobs and make them more human, designing new jobs that make the most of our human advantages (and thereby make us humans more productive). On the other hand, we have automation being deployed in a manner similar to human labour, which may also not be optimal.

The conundrum of low productivity growth might well be due to both under-utilised staff and under-utilised technology. Treating humans as task-performers, and a cost to be minimised, might be conventional wisdom, but Zeynep Ton found (and documented in her book *The Good Jobs Strategy*) that a number of firms across a range of industries – including well-known organisations such as Southwest Airlines, Toyota, Zappos, Wegmans, Costco, QuikTrip, and Trader Joe’s – were all able to realize above-average service, profit, and growth by crafting jobs that made the most of their employees’ inherent nature to be social animals and creative problem-solvers.¹⁸ Similarly, our inability to realise the potential of many AI technologies might not be due to the limitations of the technologies themselves, but, instead, our insistence on treating them as

independent mechanised task performers. To be sure, AI can be used to automate tasks. But its full potential may lie in putting it to a more substantial use.

There are historical examples of new technologies being used in a suboptimal fashion for years, sometimes decades, before their more effective use was realized.¹⁹ For example, using electricity in place of steam in the factory initially resulted only in a cleaner and quieter work environment. It drove a productivity increase only 30 years later, when engineers realised that electrical power was easier to distribute (via wires) than mechanical power (via shafts, belts, and pulleys). The single, centralised engine (and mechanical power distribution), which was a legacy of the steam age, was swapped for small engines directly attached to each machine (and electrical power distribution). This enabled the shop floor to be optimised for workflow rather than power distribution, delivering a sudden productivity boost.

A NEW LINE BETWEEN HUMAN AND MACHINE

THE question then arises: If AI’s full potential doesn’t lie in automating tasks designed for humans, what is its most appropriate use? Here, our best guidance comes from evidence that suggests human and machine intelligence are best viewed as complements rather than substitutes²⁰ – and that humans and AI, working together, can

achieve better outcomes than either alone.²¹ The classic example is freestyle chess. When IBM's Deep Blue defeated chess grandmaster Garry Kasparov in 1997, it was declared to be "the brain's last stand." Eight years later, it became clear that the story is considerably more interesting than "machine vanquishes man." A competition called "freestyle chess" was held, allowing any combination of human and computer chess players to compete. The competition resulted in an upset victory that Kasparov later reflected upon:

The surprise came at the conclusion of the event. The winner was revealed to be not a grandmaster with a state-of-the-art PC but a pair of amateur American chess players using three computers at the same time. Their skill at manipulating and "coaching" their computers to look very deeply into positions effectively counteracted the superior chess understanding of their grandmaster opponents and the greater computational power of other participants. Weak human + machine + better process was superior to a strong computer alone and, more remarkably, superior to a strong human + machine + inferior process... Human strategic guidance combined with the tactical acuity of a computer was overwhelming.²²

There are historical examples of new technologies being used in a suboptimal fashion for years, sometimes decades, before their more effective use was realised.

The lesson here is that human and machine intelligence are different in complementary, rather than conflicting, ways. While they might solve the same problems, they approach these problems from different directions. Machines find highly complex tasks easy, but stumble over seemingly simple tasks that any human can do. While the two might use the same knowledge, how they use it is different. To real-

ise the most from pairing human and machine, we need to focus on how the two interact, rather than on their individual capabilities.

TASKS VERSUS KNOWLEDGE

RATHER than focusing on the task, should we conceptualise work to focus on the knowledge, the raw

material common to human and machine? To answer this question, we must first recognise that knowledge is predominantly a social construct,²³ one that is treated in different ways by humans and machines.

Consider the group of things labeled "kitten." Both human and robot learn to recognise "kitten" the same way:²⁴ by considering a labeled set of exemplars (images).²⁵ However, although kittens are clearly *things in the world*, the concept of "kitten" – the knowledge, the

identification of the category, its boundaries, and label – is the result of a dialogue within a community.²⁶

Much of what we consider to be common sense is defined socially. Polite behavior, for example, is simply common convention among one's culture, and different people and cultures can have quite different views on what is correct behavior (and what is inexcusable). How we segment customers; the metric system along with other standards and measures; how we decompose problems into business processes and the tasks they contain, measure business performance, define the rules of the road, and drive cars; regulation and legislation in general; and the cliché of Eskimos having dozens, if not hundreds, of words for snow,²⁷ all exemplify knowledge that is socially constructed. Even walking – and the act of making a robot walk – is a social construct²⁸ as it was the community that identified “walking” as a phenomenon and gave it a name, ultimately motivating engineers to create a walking robot, and it's something we and robots learn by observation and encouragement. There are many possible ways of representing the world and dividing up reality, to understand the nature and relation

of things, and to interact with the world around us, and the representation we use is simply the one that we agreed on.²⁹ Choosing one word or meaning above the others has as much to do with societal convention as ontological necessity.

Socially constructed knowledge can be described as *encultured* knowledge, as it is our culture that determines what is (and what isn't) a kitten, just as it is culture that determines what is and isn't a good job. (We might even say that knowledge is created between people, rather than within them.) Encultured knowledge extends all the way up to formal logic, math, and hard science. Identifying and defining a phenomenon for investigation is thus a social process, something researchers must do before practical work can begin. Similarly, the rules, structures, and norms that are used in math and logic are conventions that have been agreed upon over time.³⁰ A fish is a fish insofar as we all call it a fish. Our concept of “fish” was developed in dialogue within the community. Consequently, our concept of fish drifts over time: In the past “fish” included squid (and some other, but not all, cephalopods), but not in current usage. The concepts that we use to

Humans experience the world in all its gloriously messy and poorly defined nature, where concepts are ill-defined and evolving and relationships fluid.

think, theorise, decide, and command are defined socially, by our community, by the group, and evolve with the group.

KNOWLEDGE AND UNDERSTANDING

HOW is this discussion of knowledge related to AI? Consider again the challenge of recognising images containing kittens. Before either human or machine can recognise kittens, we need to agree on what a “kitten” is. Only then can we collect the set of labeled images required for learning.

The distinction between human and machine intelligence, then, is that the human community is constantly constructing new knowledge (labeled exemplars in the case of kittens) and tearing down the old, as part of an ongoing dialogue within the community. When a new phenomenon is identified that breaks the mold, new features and relationships are isolated and discussed, old ones reviewed, concepts shuffled, unlearning happens, and our knowledge evolves. The European discovery of the platypus in 1798 is a case in point.³¹ When Captain John Hunter sent a platypus pelt to Great Britain,³² many scientists’ initial hunch was that it was a hoax. One pundit even proposed that it might have been a novelty created by an Asian taxidermist (and invested time in trying to find the stitches).³³ The European community didn’t know how to describe or classify the new thing. A discussion ensued, new evidence was sought, and features identified, with the community eventually deciding that the platypus

wasn’t a fake, and our understanding of animal classification evolved in response.

Humans experience the world in all its gloriously messy and poorly defined nature, where concepts are ill-defined and evolving and relationships fluid. Humans are quite capable of operating in this confusing and noisy world; of reading between the lines; tapping into weak signals; observing the unusual and unnamed; and using their curiosity, understanding, and intuition to balance conflicting priorities and determine what someone actually meant or what is the most important thing to do. Indeed, as Zeynep Ton documented in *The Good Jobs Strategy*,³⁴ empowering employees to use their judgment, to draw on their own experience and observations, to look outside the box, and to consider the context of the problem they are trying to understand (and solve), as well as the formal metrics, policies, and rules of the firm, enabled them to make wiser decisions and consequentially deliver higher performance. Unfortunately, AI doesn’t factor in the unstated implications and repercussions, the context and nuance, of a decision or action in the way humans do.

It is this ability to refer to the context around an idea or problem – to craft more appropriate solutions, or to discover new knowledge to create (and learn) – that is uniquely human. Technology cannot operate in such an environment: It needs its terms specified and objectives clearly articulated, a well-defined and

fully contextualised environment within which it can reliably operate. The problem must be identified and formalised, the inputs and outputs articulated, before technology can be leveraged. Before an AI can recognise kittens, for instance, we must define what a kitten is (by exemplar or via a formal description) and find a way to represent potential kittens that the AI can work with. Similarly, the recent boom in autonomous vehicles is due more to the development of improved sensors and hyper-accurate maps, which provide the AI with the dials and knobs it needs to operate, than the development of vastly superior algorithms.

It is through the social process of knowledge construction that we work together to identify a problem, define its boundaries and dependences, and discover and eliminate the unknowns until we reach the point where a problem has been defined sufficiently for knowledge and skills to be brought to bear.

A BRIDGE BETWEEN HUMAN AND MACHINE

If we're to draw a line between human and machine, then it is the distinction between *creating* and *using* knowledge. On one side is the world of the unknowns (both known and unknown), of fuzzy concepts that cannot be fully articulated, the land of the humans, where we work together to make sense of the world. The other side is where terms and definitions have been established, where the problem is known and all variables are quantified, and automa-

tion can be applied. The bridge between the two is the social process of knowledge creation.

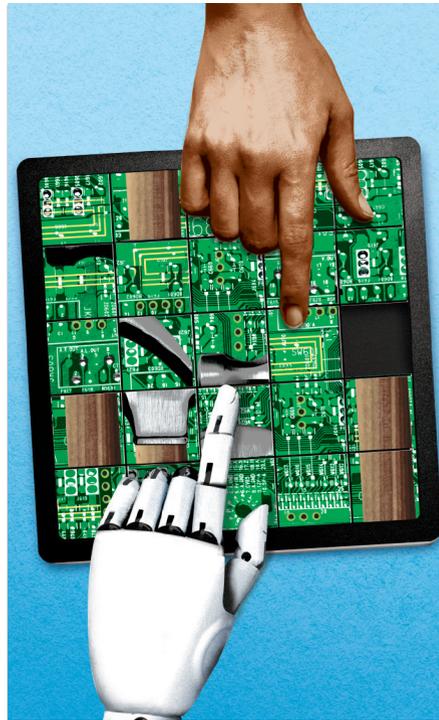
Consider the question of what a “happy retirement” is: We all want one, but we typically can't articulate what it is. It's a vague and subjective concept with a circular definition: A happy retirement is one in which *you're* happy. Before we can use an AI-powered robo-advisor to create our investment portfolio, we need to take our concept of a “happy retirement” through grounding the concept (“what will *actually* make me happy, as opposed to what I *think* will make me happy”), establishing reasonable expectations (“what can I expect to fund”), to attitudes and behaviors (“how much can I change my habits, how and where I spend my money, to free up cash to invest”), before we reach the quantifiable data against which a robo-advisor can operate (investment goals, income streams, and appetite for risk). Above quantifiable investment goals and income streams is the social world, where we need to work with other people to discover what our happy retirement might be, to define the problem and create the knowledge. Below is where automation – with its greater precision and capacity for consuming data – can craft our ultimate investment strategy. Ideally there is interaction between the two layers – as with freestyle chess – with automation enabling the humans to play what-if games and explore how the solution space changes depending on how they shape the problem definition.

RECONSTRUCTING WORK

THE foundation of work in the pre-industrial, craft era was the product. In the industrial era it is the task, specialised knowledge, and skills required to execute a step in a production process. Logically, the foundation of post-industrial work will be the problem – the *goal to be achieved*³⁵ – one step up from the solution provided by a process.

If we're to organise work around problems and successfully integrate humans and AI into the same organisation, then it is management of the problem definition – rather than the task as part of a process to deliver a solution - that becomes our main concern.³⁶ Humans take responsibility for shaping the problem – the data to consider, what *good* looks like, the choices to act – which they do in collaboration with those around them and their skill in doing this will determine how much additional value the solution creates. Automation (including AI) will support the humans by augmenting them with a set of digital *behaviors*³⁷ (where a *behavior*

is the way in which one acts in response to a *particular* situation or stimulus) that replicate *specific* human behaviors, but with the ability to leverage more data and provide more precise answers while not falling prey to the various cognitive biases to which we humans are prone. Finally, humans will evaluate the appropriateness and completeness of the solution provided and will act accordingly.



Indeed, if automation in the industrial era was the replication of *tasks* previously isolated and defined for humans, then in the post-industrial era, automation might be the replication of isolated and well-defined *behaviors* that were previously unique to humans.

INTEGRATING HUMANS AND AI

CONSIDER the challenge of eldercare. A recent initiative in the United

Kingdom is attempting to break down the silos in which specialised health care professionals currently work.³⁸ Each week, the specialists involved with a single patient – health care assistant, physiotherapist, occupational therapist, and so on – gather to discuss the patient.

Each specialist brings his or her own point of view and domain knowledge to the table, but as a group they can build a more comprehensive picture of how best to help the patient by integrating observations from their various specialties as well as discussing more tacit observations that they might have made when interacting with the patient. By moving the focus from the tasks to be performed to the problem to be defined – *how to improve the patient's quality of life* – the first phase of the project saw significant improvements in patient outcomes over the first nine months.

Integrating AI (and other digital) tools into this environment to augment the humans might benefit the patient even more by providing better and more timely decisions and avoiding cognitive biases, resulting in an even higher quality of care. To do this, we could create a common digital workspace where the team can capture its discussions; a whiteboard (or blackboard) provides a suitable metaphor, as it's easy to picture the team standing in front of the board discussing the patient while using the board to capture important points or share images, charts, and other data. A collection of AI (and non-AI) *digital behaviors* would also be integrated directly into this environment. While the human team stands in front of the whiteboard, the digital behaviors stand behind it, listening to the team's discussion and watching as notes and data are captured, and reacting appropriately, or even responding to direct requests.

Data from tests and medical monitors could be fed directly to the board, with predictive behaviors keeping a watchful eye on data streams to determine if something unfortunate is about to happen (similar to how electrical failures can be predicted by looking for characteristic fluctuations in power consumption, or how AI can be used to provide early warning of struggling students by observing patterns in communication, attendance, and assignment submission), flagging possible problems to enable the team to step in before an event and prevent it, rather than after. A speech-to-text behavior creates a transcription of the ensuing discussion so that what was discussed is easily searchable and referenceable. A medical image – an MRI perhaps – is ordered to explore a potential problem further, with the resulting image delivered directly to the board, where it is picked up by a cancer-detection behavior to highlight possible problems for the team's specialist to review. With a diagnosis in hand, the team works with a genetic drug-compatibility³⁹ behavior to find the best possible response for this patient and a drug-conflict⁴⁰ behavior that studies the patient's history, prescriptions, and the suggested interventions to determine how they will fit in the current care regime, and explore the effectiveness of different possible treatment strategies. Once a treatment strategy has been agreed on, a planning behavior⁴¹ converts the strategy into a detailed plan – taking into account the urgency, sequencing, and preferred providers for each intervention – listing

the interventions to take place and when and where each should take place, along with the data to be collected, updating the plan should circumstances change, such as a medical imaging resource becoming available early due to a cancellation.

Ideally, we want to populate this problem-solving environment with a comprehensive collection of behaviors. These behaviors might be predictive, flagging possible events before they happen. They might enable humans to explore the problem space, as the chess computer is used in freestyle chess, or the drug-compatibility and drug-conflict AIs in the example above. They might be analytical, helping us avoid our cognitive biases. They might be used to solve the problem, such as when the AI planning engine takes the requirements from the treatment strategy and the availability constraints from the resources the strategy requires, and creates a detailed plan for execution. Or they might be a combination of all of these. These behaviors could also include non-AI technologies, such as calculators, enterprise applications such as customer relationship management (CRM) (to determine insurance options for the patient), or even physical automations and non-technological solutions such as checklists.⁴²

UNIQUELY HUMAN

IT'S important to note that scenarios similar to the eldercare example just mentioned exist across a wide range of both blue- and

white-collar jobs. The Toyota Production System is a particularly good blue-collar example, where work on the production line is oriented around the problem of improving the process used to manufacture cars, rather than the tasks required to assemble a car.

One might assume that the creation of knowledge is the responsibility of academy-anointed experts. In practice, as Toyota found, it is the people at the coalface, finding and chipping away at problems, who create the bulk of new knowledge.⁴³ It is our inquisitive nature that leads us to try and explain the world around us, creating new knowledge and improving the world in the process. Selling investment products, as we've discussed, can be reframed to focus on determining what a happy retirement might look like for this *particular* client, and guiding the client to his or her goal. Electric power distribution might be better thought of as the challenge of improving a household's ability to manage its power consumption. The general shift from buying products to consuming services⁴⁴ provides a wealth of similar opportunities to help individuals improve how they consume these services, be they anything from toilet paper subscriptions⁴⁵ through cars⁴⁶ and eldercare (or other medical and health services) to jet engines,⁴⁷ while internally these same firms will have teams focused on improving how these services are created.

Advances (and productivity improvements) are typically made by skilled and curious practitio-

ners solving problems, whether it was weavers in a mill finding and sharing a faster (but more complex) method of joining a broken thread in a power loom or diagnosticians in the clinic noticing that white patches sometimes appear on the skin when melanomas regress spontaneously.⁴⁸ The chain of discovery starts at the coalface with our human ability to notice the unusual or problematic – to swim through the stream of the unknowns and of fuzzy concepts that cannot be fully articulated. This is where we collaborate to make sense of the world and create knowledge, whether it be the intimate knowledge of what a happy retirement means for an individual, or grander concepts that help shape the world around us. It is this ability to collectively make sense of the world that makes us uniquely human and separates us from the robots – and it cuts across all levels of society.

If we persist in considering a job to be little more than a collection of related tasks, where value is determined by the knowledge and skill required to prosecute them, then we should expect that automation will eventually consume *all* available work, as we must assume that any

well-defined task, no matter how complex, will eventually be automated. This comes at a high cost, as while machines can learn, they don't in themselves, create new knowledge. An AI tool might discover patterns in data, but it is the humans who noticed that the data set was interesting and then inferred meaning into the

patterns discovered by the machine. As we relegate more and more tasks to machines, we are also eroding the connection between the problems to be discovered and the humans who can find and define them. Our machines might be able to learn, getting better at doing what they do, but they won't be able to reconceive what *ought* to be done, and think outside their algorithmic box.

There is a third option, though: one where we move from building jobs around processes and tasks, a solution that is optimal for neither human nor machine, to building jobs around problems.

CONCLUSION

AT the beginning of this article, we asked if the pessimists or optimists would be right. Will the future of work be defined by a lack of suitable jobs for much of the population? Or will historical norms reassert themselves, with automation creating more work than it destroys? Both of these options are quite possible since, as we often for-

get, work is a social construct, and it is up to us to decide how it should be constructed.

There is a third option, though: one where we move from building jobs around processes and tasks, a solution that is optimal for neither human nor machine, to building jobs around problems. The difficulty is in defining production as a problem to be solved, rather than a process to be streamlined. To do this, we must first establish the context for the problem (or contexts, should we decompose a large production into a set of smaller interrelated problems). Within each context, we need to identify what is known and what is unknown and needs to be discovered. Only then can we determine for each problem whether human or machine, or human and machine, is best placed to move the problem forward.

Reframing work, changing the foundation of how we organise work from *task to be done* to *problem to be solved* (and the consequent reframing of automation from the replication of *tasks* to the replication of *behaviors*) might provide us with the opportunity to jump from the industrial productivity improvement

S-curve⁴⁹ to a post-industrial one. What drove us up the industrial S-curve was the incremental development of automation for more and more complex tasks. The path up the post-industrial S-curve might be the incremental development of automation for more and more complex behaviors.

The challenge, though, is to create not just jobs, but *good* jobs that make the most of our human nature as creative problem identifiers. It was not clear what a good job was at the start of the Industrial Revolution. Henry Ford's early plants were experiencing nearly 380 percent turnover and 10 percent daily absenteeism from work,⁵⁰ and it took a negotiation between capital and labor to determine what a good job should look like, and then a significant amount of effort to create the infrastructure, policies, and social institutions to support these good jobs. If we're to change the path we're on, if we're to choose the third option and construct work around problems whereby we can make the most of our own human abilities and those of the robots, then we need a conscious decision to engage in a similar dialogue. ●

Peter Evans-Greenwood is a fellow at the Deloitte Centre for the Edge, Australia, supported by Deloitte Consulting Pty. Ltd.

Harvey Lewis is a director with Deloitte MCS Limited and lead for cognitive computing in the Technology Consulting practice in the United Kingdom.

James Guszcza is the US chief data scientist for Deloitte Consulting LLP.

Endnotes

1. The concept of the division of labour – the deconstruction of the problem into a set of sequential tasks, with participants specialising in particular tasks – has a long history, one reaching all the way back to Plato, though it seems to be Adam Smith that most people associate with the idea. It was his 1776 book, *An Inquiry into the Nature and Causes of the Wealth of Nations* (more commonly known as *The Wealth of Nations*), in which Adam Smith posited that enabling workers to concentrate and specialise on their particular tasks leads to greater productivity and skills. It's worth noting that Smith foresaw many of today's problems when he observed that dividing labor too finely can lead to "the almost entire corruption and degeneracy of the great body of the people . . . unless government takes some pains to prevent it." Alexis de Tocqueville made the same point more bluntly when he stated (in his 1841 book, *Democracy in America: Volume I*) that "Nothing tends to materialise man, and to deprive his work of the faintest trace of mind, more than extreme division of labor."
2. For a thoughtful discussion of the application of such AI methods to radiology and the potential impact on practitioners, see Siddhartha Mukherjee, "A.I. versus M.D.," *New Yorker*, April 3, 2017, <http://www.newyorker.com/magazine/2017/04/03/ai-versus-md>.
3. TaskRabbit (www.taskrabbit.com) provides an online and mobile marketplace for everyday tasks – such as cleaning, handyman work, and moving – that matches consumers with freelance labor.
4. Kaggle (www.kaggle.com) is an online platform for analytics and predictive modelling that enables companies and researchers to post their data and run competitions with freelance data scientists to provide the best data models.
5. Robotic process automation (RPA) is an approach to automating common clerical tasks by creating *software robots* that replicate the actions of human clerical workers interacting with the user interface of a computer system, operating on the user interface in the same way that a human would. Common tasks for these software robots are data entry or transfer, such as an auditor extracting financial transactions from a client's bookkeeping system and entering them into the audit system.
6. The increasing difficulty individuals find in maintaining the knowledge and skills required is often attributed to a combination of a decreasing half-life of knowledge and the red queen effect. The half-life of knowledge is a concept attributed to Fritz Machlup, and was intended to capture the feeling that knowledge ages much more rapidly today than it did in the past (the analogy made between nuclear decay and the erosion of knowledge is awkward at best). More precisely, it is defined as the time that has to elapse before half the knowledge or facts in a particular domain are superseded or shown to be false. In 2008, Roy Tang determined that the half-life of knowledge was 13 years for physics, 9 for math, and 7.1 years for psychology and history. The term is inherently imprecise due the challenges in cleanly defining a domain and identifying (and discriminating between) the knowledge and facts it contains. The red queen effect refers to an evolutionary hypothesis that proposes that organisms must constantly change and adapt, or be overtaken by other organisms that change and adapt faster in a constantly changing environment. The effect is named after the Red Queen in Lewis Carroll's *Through the Looking-Glass*.
7. It's interesting to note that early punch-card looms - where the pattern to be woven was encoded in a series of punch cards – were a precursor of the modern digital computer.
8. Refer to J. Bessen, *Learning by Doing: The Real Connection between Innovation, Wages, and Wealth* (Yale University Press, 2015), for a thorough discussion of the relationship between the initial invention of a new automation technology and the subsequent incremental improvement of the technology by workers identifying better work practices and improvements, and how the productivity improvements reduced cost which, in turn, resulted in higher demand. The first power looms, for example, improved productivity by a factor of 2.5, while the subsequent incremental improvements lifted the factor up to 50 by the end of the Industrial Revolution.
9. It's commonly claimed that the only example of a job that has been entirely eliminated by technology is that of the elevator attendant, though it's interesting to note that this job was also created by technology.
10. Bessen, *Learning by Doing*.
11. Taken from Robert J. Gordon, figure 1-1 ("Annualized growth rate of output per person, output per hour, and hours per person, 1870–2014"), *The Rise and Fall of American Growth: The U.S. Standard of Living Since the Civil War* (Princeton University Press, 2016).

12. This is the “automation paradox”: When computers start doing the work of people, the need for people often increases. Rather than replace the human, these solutions still require human oversight. If automation is being used for tasks where human workload or cognitive load is low, then it can complicate situations when human workload is high. A good example is aircraft autopilots, where routine tasks were handed off to automation, leaving the pilot to deal with the tricky scenarios, such as landing or negotiating with air traffic control. The relationship between pilot and plane has changed, and pilots find it unsettling when the automation is not operating flawlessly. Something as simple as a sensor icing up might cause the autopilot to disengage, surprising the crew and nudging them onto a path that leads to a fatal mistake. Joe Pappalardo, contributing editor at *Popular Mechanics* magazine, points out that “catastrophic failures don’t happen as often but they are more catastrophic when they do.” Pilot error is the notional cause for roughly 50 percent of fatal accidents, but the source of this error might be the interface between human and automation. An entirely manual system was more robust as it lacked this human-computer hand-off. As Pappalardo concludes, “If something went wrong in the 1970s, there was a chance you could land it.” See Finlo Rohrer and Tom de Castella, “Mechanical v human: Why do planes crash?,” *BBC News Magazine*, March 14, 2014, <http://www.bbc.com/news/magazine-26563806>.
13. Gordon, *The Rise and Fall of American Growth*.
14. Ibid, figure 1–1 (“Annualized growth rate of output per person, output per hour, and hours per person, 1870–2014”).
15. We mistake what is unfamiliar as something that is new in and of itself. Many of the AI technologies considered part of cognitive computing are not new. The statistical approach to machine translation originated in the late 1980s. The groundwork for artificial neural networks was established by Donald Hebb in the ‘40s, refined in the ‘90s when key innovations such as back propagation were developed, and became practical mid-2000s when hardware and data sets caught up. Many of the technologies considered part of cognitive computing have similarly long histories. Compare this to the development of motion pictures. As a child, Charlie Chaplin performed in three large music halls an evening. By 1915, 10 years later, he could be seen in thousands of halls across the world. It took radio only 10 years from the launch of the first commercial radio station in 1920, to reach 80 percent of homes. Just 8 percent of urban American households had electricity in 1907. By 1929, 85 percent had electricity. After a long gestation as various inventors attempted to use coal gas to fuel a self-propelled engine, Karl Benz successfully trailed a two-stroke gasoline engine on New Year’s Eve of 1879 (just 10 weeks after Edison had perfected the electric light bulb). Just over 20 years later in 1906, Wilhelm Maybach developed a six-cylinder engine that powered a car with equivalent power and function to a modern compact. With that the car took off, taking only another 20 years to rocket from effectively zero percent ownership to 60 percent, after which it took a more leisurely pace as it asymptotes toward today’s figure of roughly 80 percent. Today’s technology environment, however, is highly entailed—new technologies depend on earlier ones, and as time passes and society accretes new technologies, the technologies themselves become more complex as they depend on a greater number of prior developments and resources. Google Translate appeared in 2006 as this is when Google’s engineers had finally obtained a data set that could exercise the statistical algorithms the service was based on, algorithms proposed in the ‘80s. Autonomous cars quickly flipped from pie in the sky to *you’ll be able to buy one real soon* once better sensors were developed and comprehensive electronic road maps were compiled, accurate down to the centimeter. And so on.
16. Anthropology professor David Graeber explored the phenomenon of what he termed *bullshit jobs* in his 2013 essay *On the phenomenon of bullshit jobs*. He noted that many clerical jobs are unfulfilling, with the workers responsible for them feeling that their labor is unproductive and pointless, their work unnecessary. See David Graeber, “On the phenomenon of bullshit jobs,” *Strike*, 2013, <http://strikemag.org/bullshit-jobs/>.
17. Similar to how Ford’s early factories were experiencing 380 percent turnover and 10 percent daily absenteeism from work in their first years of operation.
18. Z. Ton, *The Good Jobs Strategy: How the Smartest Companies Invest in Employees to Lower Costs and Boost Profits* (New Harvest, 2014).
19. Bessen provides many fascinating examples that show how the development of *know-how*, the knowledge of how to make best use of technology, provides the majority of the productivity improvement attributed to a new technology, with the invention of the technology itself providing a much more modest productivity boost. Bessen, *Learning by Doing*.
20. This is the theme of the authors’ previous work; see Jim Guszczka, Harvey Lewis, and Peter Evans-Greenwood, “Cognitive collaboration: Why humans and computers think better together,” *Deloitte Review*

- 20, January 23, 2017, <https://dupress.deloitte.com/dup-us-en/deloitte-review/issue-20/augmented-intelligence-human-computer-collaboration.html>.
21. Ibid.
 22. Garry Kasparov, "The chess master and the computer," *New York Review of Books* 57, no. 2 (2010): pp. 1–6, www.nybooks.com/articles/2010/02/11/the-chess-master-and-the-computer.
 23. While the preponderance of our knowledge might be socially constructed not all knowledge is. We experience our own heartbeat, for example, without intervention, though identifying, delineating, and naming this phenomenon "heartbeat" was the result of social construction. The book *Introduction to New Realism* by Maurizio Ferraris is recommended to the more philosophically minded readers, as a sound definition of the position taken by this report: Maurizio Ferraris, *Introduction to New Realism* (Bloomsbury Academic, 2015).
 24. Indeed, it was the development of AI tools that enabled us to do things such as recognising images of kittens plucked from the Internet, which has caused so many conceptions as prior to that tacit knowledge was considered the exclusive domain of humans.
 25. It's interesting to note that children need to see much fewer images of kittens than AI to learn the category. Humans, in general, require less data to learn than machines.
 26. Care must be taken not to confuse the *thing* (ontology) with our *knowledge* of the thing (epistemology). The thing – *kitten*, perhaps – is clearly an immalleable object in the world, but our knowledge of the thing is socially constructed. It is the *knowledge* that we work with, that we capture in mechanisms and automate. A machine learning tool doesn't *kitten* (operate on the object directly), it *recognises* kittens (operates on its knowledge of what a kitten is). Similarly, an autonomous car doesn't interact with stoplights directly, it relies on its knowledge of stoplights and the signals from its various sensors to interpret the environment around it. It is the imperfect nature of this interpretation process that causes autonomous cars to make mistakes (just as humans do). It's for this reason that Nietzsche repeatedly wrote, "There are no facts, only interpretations" in the margins of his notebooks.
 27. Franz Boas, in his book *Handbook of American Indian Languages*, discusses how languages don't necessarily draw lines between the lexemes in semantic fields in the same places as other languages. Canadian Inuit separates falling snowflakes (for which the qana-root is used) from snow lying on the ground (for which the api- root is used), just as English separates water running along (as in river) from water standing still (as in lake), and so on. He was stressing that this arbitrariness of lexical denotation boundaries was something the two languages had in common, not that Inuit was quantitatively unusual and made quantitative claims on the number of different words the American Eskimos have for snow. See Franz Boas, *Handbook of American Indian Languages*, 1911, pp. 179–222.
 28. We should note that "walking" is also an example of *embodied* knowledge. Embodied knowledge depends on the configuration of one's body (robot or human), and one's ability will depend on the synergies between knowledge and body. Usain Bolt's training partner, Yohan Blake, has a strikingly similar technique and cadence to Bolt, but is a few centimeters shorter and consequently doesn't travel quite as far with each stride. This is also why teaching a robot how to walk is a challenging task. It's not that we don't understand how walking works, it is the difficulty in building a suitable body and dealing with the complex computations required. This is where techniques such as reinforcement learning are powerful, as they enable us to teach the robot by example, rather than having to explicitly define all the processes and calculations required. It is more difficult to transmit embodied knowledge than formal knowledge (math or logic), as the knowledge is only useful to the recipient if they have the same hardware.
 29. The cognitive scientist Richard Nisbett's book, *The Geography of Thought*, provides examples of how concepts, categories, and judgments vary across cultures. See Richard Nisbett, *The Geography of Thought: How Asians and Westerners Think Differently . . . and Why* (Free Press, 2003). A brief introduction to these ideas can be seen in Lera Boroditsky, "How the languages we speak shape the ways we think," video, <https://www.youtube.com/watch?v=VHulvUwgFWo>.
 30. This "social accretion over time" is the reason for many of the complexities and quirks of mathematical notation.
 31. To pick one example of many novel Australian creatures, such as the kangaroo, emu, or drop bear.
 32. Captain John Hunter was the second governor of New South Wales.
 33. As the platypus specimens arrived in England via the Indian Ocean, naturalists suspected that Chinese sailors, known for their skill in stitching together hybrid creatures, were playing a joke on them.
 34. Ton, *The Good Jobs Strategy*.

35. As opposed to *work to be done*, which represents an inherently task-based view of work.
36. We should note here that shifting our focus from *process to problem* enables us to make processes malleable, rather than being static. AI technologies already exist – and are, in fact, quite old – that enable us to assemble a process incrementally, in real time, enabling us to more effectively and efficiently adapt to circumstances as they change. This effectively hands responsibility for defining and creating processes over to the robots: Yet another complex skill is consumed by automation.
37. We note that behaviors are not necessarily implemented with AI technologies. Any digital (or, indeed, non-digital) technology can be used.
38. Matthew Price, The health workers that help patients stay at home, BBC News, February 8, 2017, <http://www.bbc.com/news/health-38897257>.
39. Personalized genetic medicine promises to avoid dangerous drug reactions by matching the drug to be used to the patient's genetic code. See Dina Maron, "A very personal problem," *Scientific American*, 2016, <https://www.scientificamerican.com/article/a-very-personal-problem/>.
40. Rule and constraint satisfaction engines are a well-established area of AI, dating back to the 1970s.
41. The first planning engine, STRIPS (Stanford Research Institute Problem Solver), was developed in 1971 by Richard Fikes and Nils Nilsson at SRI International.
42. Checklists have long been used as powerful tools to ensure quality. For more details, see Atul Gawande, *The Checklist Manifesto: How to Get Things Right* (Metropolitan Books, 2009).
43. We assume that knowledge and innovation flow downhill, from basic research or the lone inventor to praxis, though this is not true. While basic research and invention do result in new innovations, it is more common for knowledge to emerge bottom-up, the result of people solving problems and building on what had come before. For a good overview of a complex topic, see Daniel Sarewitz, "Saving science," *New Atlantis*, no. 49 (spring/summer 2016): pp. 4–40, <http://www.thenewatlantis.com/publications/saving-science>.
44. A trend known as servitization, the conversion of products into value-added services. The classic example is Rolls Royce's TotalCare program, where airlines pay for engine operating hours rather than buy (or lease) the engines themselves. Customers pay a fixed rate for each hour the engine is available for operation, while Rolls Royce monitors the engines remotely and takes responsibility for improving, repairing, or replacing broken engines. TotalCare was first formalized in the 1980s. Since then servitisation has moved into the consumer sphere.
45. Who Gives a Crap (<https://au.whogivesacrap.org>) provides what are effectively toilet paper subscriptions.
46. A range of services has emerged – such as GoGet (<https://www.goget.com.au>) and Flexicar (<http://flexicar.com.au>) – that enables individuals to rent cars by the hour, with the car housed in a parking space nearby.
47. Rolls Royce TotalCare, mentioned in endnote No. 44, enables airlines to buy operating hours ("hot air out the back of the engine") rather than purchase or lease the engines. TotalCare, and similar services, are considered one of the key enablers of the low-cost airline industry.
48. See Mukherjee, "A.I. versus M.D.," for an insightful discussion on the relationship between machine learning and diagnosticians.
49. An S-curve, also known as a sigmoid, is a line with the rough shape of an "S" leaning to the right. Starting horizontal, the line gradually curves up to a linear middle section, before curving back down to become horizontal again. S-curves are commonly used to represent technology development or adoption, as they mirror the slow-fast-slow nature of these processes.
50. Ford Motor Company, "100 years of the moving assembly line," <http://corporate.ford.com/innovation/100-years-moving-assembly-line.html>, accessed April 14, 2017; Michael Perelman, *Railroad-ing Economics: The Creation of the Free Market Mythology* (Monthly Review Press, 2006), pp. 135–136.

Local Future of Work Experts



Pam Maharaj

Human Capital Leader Africa
pammaharaj@deloitte.co.za
+27 82 458 2518



Janine Nel

Organisation Transformation & Talent Leader
jnel@deloitte.co.za
+27 82 789 8103



Andre Vermeulen

Leader of Human Resources Technology
andvermeulen@deloitte.co.za
+27 74 347 0138



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. Please see www.deloitte.com/about for a more detailed description of DTTL and its member firms

Deloitte provides audit, consulting, financial advisory, risk management, tax and related services to public and private clients spanning multiple industries. With a globally connected network of member firms in more than 150 countries and territories, Deloitte brings world-class capabilities and high-quality service to clients, delivering the insights they need to address their most complex business challenges. Deloitte's more than 225 000 professionals are committed to making an impact that matters

This communication contains general information only, and none of Deloitte Touche Tohmatsu Limited, its member firms or their related entities (collectively, the "Deloitte Network") is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser. No entity in the Deloitte network shall be responsible for any loss whatsoever sustained by any person who relies on this communication