

Reconstructing jobs

Creating good jobs in the age of artificial intelligence

Deloitte's Cognitive and Analytics practice is a global consultancy that drives data-, analytic-, and cognitive-enabled transformations. With our premier data-powered and platform-enabled services, we help clients achieve a cognitive advantage.

Contents

Creating good jobs | 2

Problems rather than processes | 4

Good for workers, firms, and society | 7

How will we shape the jobs of the future? | 10


Endnotes | 11

Creating good jobs

When it comes to work, workers, and jobs, much of the angst of the modern era boils down to the fear that we're witnessing the automation endgame, and that there will be nowhere for humans to retreat as machines take over the last few tasks.

THE MOST RECENT wave of commentary on this front stems from the use of artificial intelligence (AI) to capture and automate tacit knowledge and tasks, which were previously thought to be too subtle and complex to be automated. Is there no area of human experience that can't be quantified and mechanized? And if not, what is left for humans to do except the menial tasks involved in taking care of the machines?

At the core of this concern is our desire for *good jobs*—jobs that, without undue intensity or stress, make the most of workers' natural attributes and abilities; where the work provides the worker with motivation, novelty, diversity, autonomy, and work/life balance; and where workers are duly compensated and consider the employment contract fair. Crucially, good jobs support workers in *learning by doing*—and, in so doing, deliver benefits on three levels: to the worker, who gains in personal development and job satisfaction; to the organization, which innovates as staff find new problems to solve and opportunities to pursue; and to the community as a whole, which reaps the economic benefits of hosting thriving organizations and workers. This is what makes good jobs productive and sustainable for the organization, as well as engaging and fulfilling for the worker. It is also what aligns good jobs with the larger community's values and norms,



since a community can hardly argue with having happier citizens and a higher standard of living.¹

Does the relentless advance of AI threaten to automate away all the learning, creativity, and meaning that make a job a good job? Certainly, some have blamed technology for just such an outcome. Headlines today often express concern over technological

innovation resulting in bad jobs for humans, or even the complete elimination of certain professions. Some fear that further technology advancement in the workplace will result in jobs that are little more than collections of loosely related tasks where employers respond to cost pressures by dividing work schedules into ever smaller slithers of time, and where employees are being asked to work for longer periods over more days. As the monotonic progress of technology has automated more and more of a firm's function, managers have fallen into the habit of considering work as little more than a series of tasks, strung end-to-end into processes, to be accomplished as efficiently as possible, with human labor as a cost to be minimized. The result has been the creation of narrowly defined, monotonous, and unstable jobs, spanning knowledge work and procedural jobs in bureaucracies and service work in the emerging "gig economy."²

The problem here isn't the technology; rather, it's the way the technology is used—and, more than that, the way people *think* about using it. True, AI can execute certain tasks that human beings have historically performed, and it can thereby replace the humans who were once responsible for those tasks. However, just because we can use AI in this manner doesn't mean that we should. As we have

previously argued, there is tantalizing evidence that using AI on a task-by-task basis may not be the most effective way to apply it.³ Conceptualizing work in terms of tasks and processes, and using technology to automate those tasks and processes, may have served us well in the industrial era, but just as AI differs from previous generations of technologies in its ability to mimic (some) human behaviors, so too should our view of work evolve so as to allow us to best put that ability to use.

In this essay, we argue that the thoughtful use of AI-based automation, far from making humans obsolete or relegating them to busywork, can open up vast possibilities for creating meaningful work that not only allows for, but requires, the uniquely human strengths of sense-making and contextual decisions. In fact, creating good jobs that play to our strengths as social creatures might be *necessary* if we're to realize AI's latent potential and break us out of the persistent period of low productivity growth that we're experiencing today. But for AI to deliver on its promise, we must take a fundamentally different view of work and how work is organized—one that takes AI's uniquely flexible capabilities into account, and that treats humans and intelligent machines as partners in search of solutions to a shared problem.

Problems rather than processes

CONSIDER A CHATBOT—A computer program that a user can converse or *chat* with—typically used for product support or as a shopping assistant. The computer in the *Enterprise* from *Star Trek* is a chatbot, as is Microsoft’s Zo, and the virtual assistants that come with many smartphones. The use of AI allows a chatbot to deliver a range of responses to a range of stimuli, rather than limiting it to a single stereotyped response to a specific input. This flexibility in recognizing inputs and generating appropriate responses is the hallmark of AI-based automation, distinguishing it from automation using prior generations of technology. Because of this flexibility, AI-enabled systems can be said to display *digital behaviors*, actions that are driven by the recognition of what is required in a *particular* situation as a response to a *particular* stimulus.

We can consider a chatbot to embody a set of digital behaviors, how the bot responds to different utterances from the user. On the one hand, the chatbot’s ability to deliver different responses to different inputs gives it more utility and adaptability than a nonintelligent automated system. On the other hand, the behaviors that chatbots evince are fairly simple, constrained to canned responses in a conversation plan or limited by access to training data.⁴ More than that, chatbots are also constrained by their inability to leverage the social and cultural context they find themselves in. This is what makes chatbots—and AI-enabled systems generally—fundamentally different from humans, and an important reason that AI cannot “take over” all human jobs.

Humans rely on context to make sense of the world. The meaning of “let’s table the motion,” for

example, depends on the context it’s uttered in. Our ability to refer to the context of a conversation is a significant contributor to our rich behaviors (as opposed to a chatbot’s simple ones). We can tune our response to verbal and nonverbal cues, past experience, knowledge of past or current events, anticipation of future events, knowledge of our counterparty, our empathy for the situation of others, or even cultural preferences (whether or not we’re consciously aware of them). The context of a conversation also evolves over time; we can infer new facts and come to new realizations. Indeed, the act of reaching a conclusion or realizing that there’s a better question to ask might even provide the stimulus required to trigger a different behavior.

Chatbots are limited in their ability to draw on context. They can only refer to external information that has been explicitly integrated into the solution. They don’t have general knowledge or a rich understanding of culture. Even the ability to refer back to earlier in a conversation is problematic, making it hard for earlier behaviors to influence later ones. Consequentially, a chatbot’s behaviors tend to be of the simpler, functional kind, such as providing information in response to an explicit request. Nor do these behaviors interact with each other, preventing more complex behaviors from emerging.

The way chatbots are typically used exemplifies what we would argue is a “wrong” way to use AI-based automation—to execute tasks typically performed by a human, who is then considered redundant and replaceable. By only automating the simple behaviors within the reach of technology, and then treating the chatbot as a replacement for humans, we’re eliminating richer, more complex social and cultural behaviors that make interac-

tions valuable. A chatbot cannot recognize humor or sarcasm, interpret elliptical allusions, or engage in small talk—yet we have put them in situations where, being accustomed to human interaction, people expect all these elements and more. It’s not surprising that users find chatbots frustrating and chatbot adoption is failing.⁵

A more productive approach is to combine digital and human behaviors. Consider the challenge of helping people who, due to a series of unfortunate events, find themselves about to become homeless. Often these people are not in a position to use a task-based interface—a website or interactive voice response (IVR) system—to resolve their situation. They need the rich interaction of a behavior-based interface, one where interaction with another human will enable them to work through the issue, quantify the problem, explore possible options, and (hopefully) find a solution.

We would like to use technology to improve the performance of the contact center such a person might call in this emergency. Reducing the effort required to serve each client would enable the contact center to serve more clients. At the same time, we don’t want to reduce the quality of the service. Indeed, ideally, we would like to take some of the time saved and use it to improve the service’s value by empowering social workers to delve deeper into problems and find more suitable (ideally, longer-term) solutions. This might also enable the center to move away from break-fix operation, where a portion of demand is due to the center’s inability to resolve problems at the last time of contact. Clearly, if we can use technology appropriately then it might be possible to improve efficiency (more clients serviced), make the center more effective (more long-term solutions and less break-fix), and also increase the value of the outcome for the client (a better match between the underlying need and services provided).

If we’re not *replacing* the human, then perhaps we can *augment* the human by using a machine to automate some of the repetitive tasks. Consider oncology, a common example used to illustrate

this human-augmentation strategy. Computers can already recognize cancer in a medical image more reliably than a human. We could simply pass responsibility for image analysis to machines, with the humans moving to more “complex” un-automated tasks, as we typically integrate human and machine by defining handoffs between tasks. However, the computer does not identify what is unusual with this *particular* tumor, or what it has in common with other unusual tumors, and launch into the process of discovering and developing new knowledge. We see a similar problem with our chatbot example, where removing the humans from the front line prevents social workers from understanding how the factors driving homelessness are changing, resulting in a system that can only service old demand, not new. If we break this link between doing and understanding, then our systems will become more precise over time (as machine operation improves) but they will not evolve outside their algorithmic box.

Our goal must be to construct work in such a way that digital behaviors are blended with human behaviors, increasing accuracy and effectiveness, while creating space for the humans to identify the unusual and build new knowledge, resulting in solutions that are superior to those that digital or human behaviors would create in isolation. Hence, if we’re to blend AI and human to achieve higher performance, then we need to find a way for human and digital behaviors to work together, rather than in sequence. To do this, we need to move away from thinking of work as a string of tasks comprising a *process*, to envisioning work as a set of complementary behaviors concentrated on addressing a *problem*. Behavior-based work can be conceptualized as a team standing around a shared whiteboard, each holding a marker, responding to new stimuli (text and other marks) appearing on the board, carrying out their action, and drawing their result on the same board. Contrast this with task-based work, which is more like a bucket brigade where the workers stand in a line and the “work” is passed from worker to worker on its way

to a predetermined destination, with each worker carrying out his or her action as the work passes by. Task-based work enables us to create optimal solutions to specific problems in a static and unchanging environment. Behavior-based work, on the other hand, provides effective solutions to ill-defined problems in a complex and changing world.

To facilitate behavior-based work, we need to create a shared context that captures what is known about the problem to be solved, and against which both human and digital behaviors can operate. The starting point in our contact center example might be a transcript of the conversation so far, transcribed via a speech-to-text behavior. A collection of “recognize-client behaviors” monitor the conversation to determine if the caller is a returning client. This might be via voice-print or speech-pattern recognition. The client could state their name clearly enough for the AI to understand. They may have even provided a case number or be calling from a known phone number. Or the social worker might step in if they recognize the caller before the AI does. Regardless, the client’s details are fetched from case management to populate our shared context, the shared digital whiteboard, with minimal intervention.

As the conversation unfolds, digital behaviors use natural language to identify key facts in the dialogue. A client mentions a dependent child, for example. These facts are highlighted for both the human and other digital behaviors to see, creating a summary of the conversation updated in real time. The social worker can choose to accept the highlighted facts, or cancel or modify them. Regardless, the human’s focus is on the conversation, and they only need to step in when captured facts need correcting, rather than being distracted by the need to navigate a case management system.

Digital behaviors can encode business rules or policies. If, for example, there is sufficient data to determine that the client qualifies for emergency

housing, then a business-rule behavior could recognize this and assert it in the shared context. The assertion might trigger a set of “find emergency housing behaviors” that contact suitable services to determine availability, offering the social worker a set of potential solutions. Larger services might be contacted via B2B links or robotic process automation (if no B2B integration exists). Many emergency housing services are small operations, so the contact might be via a message (email or text) to the duty manager, rather than via a computer-to-computer connection. We might even automate empathy by using AI to determine the level of stress in the client’s voice, providing a simple graphical measure of stress to the social worker to help them determine

If we’re to blend AI and human to achieve higher performance, then we need to find a way for human and digital behaviors to work together, rather than in sequence.

if the client needs additional help, such as talking to an external service on the client’s behalf.

As this example illustrates, the superior value provided by structuring work around problems, rather than tasks, relies on our human ability to make sense of the world, to spot the unusual and the new, to discover what’s unique in this *particular* situation and create new knowledge. The line between human and machine cannot be delineated in terms of knowledge and skills unique to one or the other. The difference is that humans can participate in the social process of creating knowledge, while machines can only apply what has already been discovered.⁶

Good for workers, firms, and society

AI ENABLES US to think differently about how we construct work. Rather than construct work from products and specialized tasks, we can choose to construct work from problems and behaviors. Individuals consulting financial advisors, for example, typically don't want to purchase investment products as the end goal; what they *really* want is to secure a happy retirement. The problem can be defined as follows: What does a "happy retirement" look like; how much income is needed to support that lifestyle, how to balance spending and saving today to find the cash to invest and navigate and (financial) challenges that life puts in the road, and what investments give the client the best shot at getting from here to there? The financial advisor, client, and robo-advisor could collaborate around a common case file, a digital representation of their shared problem, incrementally defining what a "happy retirement" is and, consequently, the needed investment goals, income streams, and so on. This contrasts with treating the work as a process of "request investment parameters" (which the client doesn't know) and then "recommend insurance" and "provide investment recommendations" (which the client doesn't want, or only wants as a means to an end). The financial advisor's job is to provide the rich human behaviors—educator to the investor's student—to elucidate and establish the retirement goals (and, by extension, investment goals), while the robo-advisor provides simple algorithmic ones, responding to changes in the case file by updating it with an optimal investment strategy. Together, the human and robo-advisor can explore more options

(thanks to the power and scope of digital behaviors) and develop a deeper understanding of the client's

Rather than construct work from products and specialized tasks, we can choose to construct work from problems and behaviors.

needs (thanks to the human advisor's questioning and contextual knowledge) than either could alone, creating more value as a result.

If organizing work around problems and combining AI and human behaviors to help solve them can deliver greater value to customers, it similarly holds the potential to deliver greater value for businesses, as productivity is partly determined by how we construct jobs. The majority of the productivity benefits associated with a new technology don't come from the initial invention and introduction of new production technology. They come from learning-by-doing:⁷ workers at the coalface identifying, sharing, and solving problems and improving techniques. Power looms are a particularly good example, with their introduction into production improving productivity by a factor of 2.5, but with a further factor of 20 provided by subsequent learning-by-doing.⁸

It's important to maintain the connection between the humans—the creative problem identifiers—and the problems to be discovered. This is something that Toyota did when it realized that highly mechanized factories were efficient, but

they didn't improve. Humans were reintroduced and given roles in the production process to enable them to understand what the machines were doing, develop expertise, and consequently improve the production processes. The insights from these workers reduced waste in crankshaft production by 10 percent and helped shorten the production line. Others improved axel production and cut costs for chassis parts.⁹

This improvement was no coincidence. Jobs that are good for individuals—because they make the most of human sense-making nature—generally are also good for firms, because they improve productivity through learning by doing. As we will see below, they can also be good for society as a whole.

Consider bus drivers. With the development of autonomous vehicles in the foreseeable future, pundits are worried about what to do with all the soon to be unemployed bus drivers. However, rather than fearing that autonomous buses will make bus drivers redundant, we should acknowledge that they will find themselves in situations that only a human, and human behaviors, can deal with. Challenging weather (heavy rain or extreme glare) might require a driver to step in and take control. Unexpected events—accidents, road work, or an emergency—could require a human's judgment to determine which road rule to break. (Is it permissible to edge into a red light while making space for an emergency vehicle?) Routes need to be adjusted due to anything from a temporarily moved stop to modifying routes due to roadwork. A human presence might be legally required to, for example, monitor underage children or represent the vehicle at an accident.

As with chatbots, automating the simple behaviors and then eliminating the human will result in an undesirable outcome. A more productive approach is to discover the problems that bus drivers deal with, and then structure work and jobs around these problems and the kinds of behaviors needed to solve them. AI can be used to automate the simple behaviors, enabling the drivers to focus on more important ones, making the human-bus combination

more productive as a result. The question is: Which problems and decision centers should we choose?

Let us assume that the simple behaviors required to drive a bus are automated. Our autonomous bus can steer, avoiding obstacles and holding its lane, maintain speed and separation with other vehicles, and obey the rules of the road. We can also assume that the bus will follow a route and schedule. If the service is frequent enough, then the collection of buses on a route might behave as a flock, adjusting speed to maintain separation and ensure that a bus arrives at each stop every five minutes or so, rather than attempting to arrive at a specific time.

As with the power loom, automating these simple behaviors means that drivers are not required to be constantly present for the bus (or loom) to operate. Rather than drive a single bus, they can now “drive” a flock of buses. The drivers monitor where each bus is, how it's tracking to schedule, with the system suggesting interventions to overcome problems, such as a breakdown, congestion, or changed road conditions. The drivers can step in to pilot a particular bus should the conditions be too challenging (roadworks, perhaps, where markings and signaling are problematic), or to deal with an event that requires that human touch.

These buses could all be on the same route. A mobile driver might be responsible for four-to-five sequential buses on a route, zipping between them as needed to manage accidents or dealing with customer complaints (or disagreements between customers). Or the driver might be responsible for buses in a geographic area, on multiple routes. It's even possible to split the work, creating a desk-bound “driver” responsible for drone operation of a larger number of buses, while mobile and stationary drivers restrict themselves to incidents requiring a physical presence. School or community buses, for example, might have remote video monitoring while in transit, complemented by a human presence at stops.

Breaking the requirement that each bus have its own driver will provide us with an immediate productivity gain. If 10 drivers can manage 25

autonomous buses, then we will see productivity increase by a factor of 2.5, as we did with power looms: good jobs for the firm, as workers are more productive. Doing this requires an astute division of labor between mobile, stationary, and remote drivers, creating three different “bus driver” jobs that meet different work preferences: good jobs for the worker *and* the firm. Ensuring that these jobs involve workers as stakeholders in improving the system enables us to tap into learning-by-doing, allowing workers to continue to work on their craft, and the subsequent productivity improvements that learning-by-doing provides, which is good for workers *and* the firm.

These jobs don’t require training in software development or AI. They do require many of the same skills as existing bus drivers: understanding traffic, managing customers, dealing with accidents, and other day-to-day challenges. Some new skills will also be required, such as training a bus where to park at a new bus stop (by doing it manually the first time), or managing a flock of buses remotely (by nudging routes and separations in response to incidents), though these skills are not a stretch. Drivers will require a higher level of numeracy and literacy than in the past though, as it is a document-driven world that we’re describing. Regardless, shifting from manual to autonomous buses does not imply making existing bus drivers redundant en masse. Many will make the transition on their own, others will require some help, and a few will require support to find new work.

The question then, is: What to do with the productivity dividend? We could simply cut the cost of a bus ticket, passing the benefit onto existing patrons. Some of the saving might also be returned to the

community, as public transport services are often subsidized. Another choice is to transform public transport, creating a more inclusive and equitable public transport system.

Buses are seen as an unreliable form of transport—schedules are sparse with some buses only running hourly for part of the day, and not running at all otherwise; and route coverage is inadequate leaving many (less fortunate) members of society in public transport deserts (locations more than 800 m from high-frequency public transport). We could rework the bus network to provide a more frequent service, as well as extending service into under-serviced areas, eliminating public transport deserts.

Automation can result in jobs that are good for the worker, firm, *and* society as a whole.

The result could be a fairer and more equitable service at a similar cost to the old, with the same number of jobs. This has the potential to transform lives. Reliable bus services might result in higher patronage, resulting in more bus routes being created, more frequent services on existing bus routes, and more bus “drivers” being hired. Indeed, this is the pattern we saw with power looms during the Industrial Revolution. Improved productivity resulted in lower prices for cloth, enabling a broader section of the community to buy higher quality clothing, which increased demand and created more jobs for weavers. Automation can result in jobs that are good for the worker, firm, *and* society as a whole.

How will we shape the jobs of the future?

THERE IS NO inevitability about the nature of work in the future. Clearly, the work will be different than it is today, though how it is different is an open question. Predictions of a jobless future, or a nirvana where we live a life of leisure, are most likely wrong. It's true that the development of new technology has a significant effect on the shape society takes, though this is not a one-way street, as society's preferences shape which technologies are pursued and which of their potential uses are socially acceptable. Melvin Kranzberg, a historian specializing in the history of technology, captured this in his fourth law: "Although technology might be a prime element in many public issues, nontechnical factors take precedence in technology-policy decisions."¹⁰

The jobs first created by the development of the moving assembly line were clearly unacceptable by social standards of the time. The solution was for society to establish social norms for the employee-employer relationship—with the legislation of the eight-hour an example of this—and the development of the social institutions to support this new relationship. New "sharing economy" jobs and AI encroaching into the workplace suggest that we might be reaching a similar point, with many firms feeling that they have no option but to create *bad jobs* if they want to survive. These bad jobs can carry an economic cost, as they drag profitability down. In this essay, as well as our previous,¹¹ we have argued that these bad jobs are also preventing us from capitalizing on the opportunity created by AI.

Our relationship with technology has changed, and how we conceive work needs to change as a consequence. Prior to the Industrial Revolution, work was predominantly craft-based; we had an in-

strumental relationship with technology; and social norms and institutions were designed to support craft-based work. After the Industrial Revolution, with the development of the moving production line as the tipping point, work was based on task-specialization, and a new set of social norms and institutions were developed to support work built around products, tasks, and the skills required to prosecute them. With the advent of AI, our relationship with technology is changing again, and this automation is better thought of as capturing behaviors rather than tasks. As we stated previously, if automation in the industrial era was the replication of tasks previously isolated and defined for humans, then in this post-industrial era automation might be the replication of isolated and well-defined behaviors that were previously unique to humans.¹²

There are many ways to package human and digital behaviors—of constructing the jobs of the future. We, as a community, get to determine what these jobs look like. This future will still require bus drivers, mining engineers and machinery operators, financial advisors, as well as social workers and those employed in the caring professions, as it is our human proclivity for noticing the new and unusual, of making sense of the world, that creates value. Few people want financial products for their retirement fund; what they really want is a happy retirement. In a world of robo-advisors, all the value is created in the human conversation between financial advisors and clients, where they work together to discover what the clients' happy retirement is (and consequently, investment goals, incomes stream, etc.), not in the mechanical creation and implementation of an investment strategy based on predefined parameters. If we're to make the most of AI, realize

the productivity (and, consequently, quality of life) improvements it promises, and deliver the opportunities for operational efficiency, then we need to choose to create good jobs:

- Jobs that make the most of our human nature as social problem identifiers and solvers
- Jobs that are productive and sustainable for organizations
- Jobs with an employee-employer relationship aligned with social norms
- Jobs that support learning by doing, providing for the worker's personal development, for the improvement of the organization, and for the wealth of the community as a whole.

The question, then, is: What do we want these jobs of the future to look like?

Endnotes

1. Our definition of "a good job" is derived from measures of job quality developed by the OECD in its annual Employment Outlook report. This measure was first developed in 2014, and is based on the eight dimensions of well-being from the report by the Stiglitz-Sen-Fitoussi Commission. (See Joseph E Stiglitz, Amartya Sen, and Jean-Paul Fitoussi, *Report by the Commission on the Measurement of Economic Performance and Social Progress*, 2010.) The OECD defines job quality as a measure of earnings quality, labor market security, and the quality of the working environment. We extend this by requiring that a good job must be "good" for each of the stakeholders involved, the firm and society as a whole, as well as the worker.
2. Research by Zeynep Ton has shown that these jobs can also carry an economic cost, as did the bad jobs of the 1910s, as they drag profitability down. Prof. Ton's subsequent book, *The Good Jobs Strategy*, argued that creating good jobs can be an economic, as well as a moral, imperative.
3. 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.
4. Chatbots come in two types: retrieval and generative. Retrieval-based chatbots use natural language techniques to break down a user's message and match it with a predefined response. The chatbot will typically follow a conversation plan, much like how workers in a call center follow a script. Generative chatbots use techniques from machine translation to generate new responses based on previous messages; rather than translating from one language to another, it translates from request to a response.
5. Victor Luckerson, "The chatbot revolution is on hold," *Ringer*, April 18, 2017; Justin Lee, "Chatbots were the next big thing: What happened?," *Medium*, June 5, 2018.
6. This was a major theme in one of the authors' prior essay. See Evans-Greenwood, Lewis, and Guszczka, "Reconstructing work."
7. James Bessen, *Learning by Doing: The Real Connection between Innovation, Wages, and Wealth* (New Haven: Yale University Press, 2015).
8. *Ibid.*
9. Craig Trudell, Yuki Hagiwara, and Ma Jie, "'Gods' edging out robots at Toyota facility," *Japan Times*, April 7, 2014.
10. Melvin Kranzberg, "Technology and history: 'Kranzberg's laws,'" *Technology and Culture* 27, no. 3 (1986): pp. 544-60, DOI: 10.2307/3105385.
11. Evans-Greenwood, Lewis, and Guszczka, "Reconstructing work."
12. *Ibid.*

About the authors

PETER EVANS-GREENWOOD is a fellow at the Deloitte 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 advisor on both the business and technology sides of the fence. He is based in Melbourne.

ALAN MARSHALL is the National Analytics leader for Deloitte Australia, responsible for working with clients to establish cognitive and analytic platforms that augment human decision-making and drive automation. He has delivered outcomes for clients in the oil and gas, mining, and utilities industries with a unique perspective on how these industries will be transformed by robotics, IoT analytics, and automation. He is based in Perth.

MATTHEW AMBROSE is a manager who works with clients to help them make data-driven decisions. With a focus on public sector and education, he uses analytics to drive strategy development, customer retention initiatives, revenue assurance, cost reduction, asset replacement strategy, and business case development. He is based in Perth.

Contacts

Peter Evans-Greenwood

Fellow, Center for the Edge Australia
Deloitte Consulting Pty. Ltd.
+61 439 327 793
pevansgreenwood@deloitte.com.au


Alan Marshall

National Analytics leader, Australia
Partner
Deloitte Consulting Pty. Ltd.
+61 8 9365 8139
almarshall@deloitte.com.au

Deloitte.

Insights

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

 Follow @DeloitteInsight

Deloitte Insights contributors

Editorial: Ramani Moses, Junko Kaji, Abrar Khan

Creative: Kevin Weier, Anoop K R

Deployment: Alex Kawecki

Production: Preetha Devan

Cover artwork: Peter Horvath

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 their related entities (collectively the "Deloitte Network") is, by means of this publication, 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 publication.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee, and its network of member firms, each of which is a legally separate and independent entity. Please see www.deloitte.com/au/about for a detailed description of the legal structure of Deloitte Touche Tohmatsu Limited and its member firms.

Deloitte provides audit, tax, consulting, and financial advisory services to public and private clients spanning multiple industries. With a globally connected network of member firms in more than 150 countries, Deloitte brings world-class capabilities and high-quality service to clients, delivering the insights they need to address their most complex business challenges. Deloitte has in the region of 200,000 professionals, all committed to becoming the standard of excellence.

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
© 2018 Deloitte Touche Tohmatsu.