CUTTER
Business Technology Journal
Management, Innovation, Transformation
Vol. 33, No. 6, 2020

Beyond Automation: *AI, ML & RPA*

San Murugesan
Guest Editor

CUTTER Business Technology Journal

As business models for creating value continue to shift, new business strategies are constantly emerging and digital innovation has become an ongoing imperative. *Cutter Business Technology Journal* delivers a comprehensive treatment of these strategies to help your organization address and capitalize on the opportunities of this digital age.

Cutter Business Technology Journal is unlike academic journals. Each monthly issue, led by an expert Guest Editor, includes five to seven substantial articles, case studies, research findings, and/or experience-based opinion pieces that provide innovative ideas and solutions to the challenges business technology professionals face right now — and prepares them for those they might face tomorrow. *Cutter Business Technology Journal* doesn't water down or delay its content with lengthy peer reviews. Written by internationally known thought leaders, academics, and practitioners — you can be certain you're getting the uncensored perspectives of global experts.

You'll benefit from strategic insight on how the latest movements in digital innovation and transformation, artificial intelligence/machine learning, Internet of Things, blockchain, analytics, and cloud, to name a few, are changing the business landscape for both new and established organizations and how cutting-edge approaches in technology leadership, enterprise agility, software engineering, and business architecture can help your organization optimize its performance and transition to these new business models.

As a subscriber, you'll also receive the *Cutter Business Technology Advisor* — a weekly bulletin featuring industry updates delivered straight to your inbox. Armed with expert insight, data, and advice, you'll be able to leverage the latest business management thinking to achieve your organization's goals.

No other journal brings together so many thought leaders or lets them speak so bluntly — bringing you frank, honest accounts of what works, what doesn't, and why. Subscribers have even referred to *Cutter Business Technology Journal* as a consultancy in print and likened each month's issue to the impassioned discussions they participate in at the end of a day at a conference!

Get the best in thought leadership and keep pace with the technologies and business models that will give you a competitive edge — subscribe to *Cutter Business Technology Journal* today!

Founding Editor: Ed Yourdon
Publisher: Karen Fine Coburn
Group Publisher: Christine Generali
Managing Editor: Cindy Swain
Copy Editors: Jennifer Flaxman, Tara K. Meads
Production Editor: Linda Dias
Client Services: service@cutter.com

Cutter Business Technology Journal® *is published monthly by Cutter Consortium, an Arthur D. Little company, 37 Broadway, Suite 1, Arlington, MA 02474-5552, USA • Tel: +1 781 648 8700 • Fax: +1 781 648 8707 • Email: cbtjeditorial@cutter.com • Website: www.cutter.com • Twitter: @cuttertweets • Facebook: Cutter Consortium. ISSN: 2475-3718 (print); 2475-3742 (online).*

Subscription rates are US $485 a year in North America, US $585 elsewhere, payable to Cutter Consortium. Reprints, bulk purchases, past issues, and multiple subscription and site license rates are available on request.

☐ Start my print subscription to *Cutter Business Technology Journal* ($485/year; US $585 outside North America).

Name | Title

Company Address

City | State/Province | ZIP/Postal Code

Email (Be sure to include for weekly *Cutter Business Technology Advisor*)

Fax to +1 781 648 8707, call +1 781 648 8700, or send email to service@cutter.com.
Mail to Cutter Consortium, 37 Broadway, Suite 1, Arlington, MA 02474-5552, USA.

Request Online License Subscription Rates

For subscription rates for online licenses, email or call: **sales@cutter.com** or **+1 781 648 8700.**

CUTTER CONSORTIUM
Access to the Experts

Opening Statement

by San Murugesan, Guest Editor

Automation is no longer being applied only to industrial manufacturing processes; it now extends to several other areas across different domains. Automation is also becoming smarter, with added intelligence, and more sophisticated, with extended capabilities. It is rapidly advancing in a few new directions and being widely adopted in more ways than ever before. It helps enterprises to become agile and flexible and to collaborate across business units. In addition to enterprises, consumers are also embracing advanced automation in their everyday lives through the growing use of voice assistants and smart home ecosystems.

Traditionally, automation employed "feedback control" to perform a task or to keep a parameter of interest within a specified limit. It transformed industries, replacing much manual work. Then, robotic process automation (RPA) helped to automate and speed up key business workflows. Now, driven by artificial intelligence (AI), machine learning (ML), autonomous operational capabilities, smart materials, GPS data, and other technologies, traditional automation is being enhanced with several advanced features. These features include intelligence, the competency to deal with unknown or uncertain environments, and the ability to perform satisfactorily even with partial information and under a sudden increase in workload.

This enhanced automation, also known as *smart automation*, *hyper automation, intelligent automation*, and *intelligent RPA* (IRPA), is generating significant interest among researchers and developers in both business and industry. It is transforming all aspects of business, whether supply chain management, financial services, customer service, transportation and logistics, or marketing. While both promising and encouraging, it also raises a few technical, organizational, managerial, social, ethical, and regulatory issues and challenges that need to be satisfactorily addressed. In this issue of *Cutter Business Technology Journal* (*CBTJ*), we examine the emergence of the new face of automation and explore novel ways to address the various issues and challenges we encounter.

The Need for Intelligent Automation

Traditional simple automation faces limitations in modern enterprises: it is incapable of meeting sophisticated requirements or handling the intricacies of the environment in which it operates. The increasing complexity of enterprise operations, interdependence among processes, a growing need to comprehend what to do in advance, demands for better performance and real-time operation, and a requirement for awareness of the context in which it operates have stymied simple automation.

Intelligent automation, which uses AI, ML, and other technologies, can address these limitations and manage unforeseen challenges, such as those issues enterprises have faced recently due to the COVID-19 crisis. The three use cases outlined below illustrate how IRPA and bots help to address the unusual demands enterprises have faced:

1. An airline company received 120,000 ticket cancellation requests from passengers in the early weeks of the COVID-19 pandemic due to travel restrictions, border closures, and flight cancellations — a 4,000% increase from its typical 3,000 cancellations per month. The company used IRPA tools to build a bot to process claims. With the deployment of these bots, it was able to work through the vastly increased number of cancellation requests and also free employees to work on complex cases that required human decision making.[1]

2. The pandemic posed a major unforeseen challenge to a biopharmaceutical company, which now had to manage enough supplies to ensure continuity of a large number of global clinical trials. The company quickly created a bot to monitor the growth of the crisis and manage its inventory and supply chain readiness. The bot automatically generates reports, pulling data from the World Health Organization (WHO), and allows the company's leadership to make decisions in real time.[2]

3. Takeda, a pharmaceutical company, sped up the clinical trial process for a promising COVID-19 treatment by adopting RPA and using software bots. The company reduced the processing time, which involved collecting prospective patients' information and determining their suitability for the trial, to days instead of weeks.[3]

Intelligent Automation Drivers

Intelligent automation is gaining considerable interest due to both business drivers and technological drivers. Business drivers include operational cost reduction, improvements in process agility and flexibility, resilience to demand and supply variations, remote or autonomous operations, and desire for online fault detection and correction. Technical drivers are advances in AI, ML, and deep learning; continuous intelligence, which enables organizations to use real-time and historical data for critical decisions and actions that need to be taken in near-real time, ranging from milliseconds to minutes; availability of large data sets needed for training and validation of learning algorithms; conversational bots; Internet of Things; computer vision; blockchain; and cloud, fog, and edge computing.

Automation Strategy

For successful implementation and realization of the intended benefits of automation, enterprises need to develop and implement a holistic automation strategy. An enterprise's automation strategy should be part of — and aligned with — its IT and AI strategies.[4]

Business values and the desired benefits of automation, not fanciful desires, should drive enterprise automation strategy. Not all that can be automated has to be automated. An enterprise's automation strategy should also consider other salient factors such as viability, supporting systems and infrastructure, technological and business risks, technology maturity, other related business processes, organizational readiness, and the potential impact on employees. Employee re-skilling or upskilling and relevant training is an important aspect that should be part of an automation strategy. The strategy should also consider ethical aspects and regulatory requirements, where relevant.

While formulating their automation strategy, enterprises should be mindful that, as Bill Gates has been attributed as saying, "Automation applied to an efficient process will magnify the efficiency.… [A]utomation applied to an inefficient operation will magnify the inefficiency."

Upcoming Topics

How Technology and Business Leaders Help in a Crisis
Steve Andriole

Proactive Risk Management
Tom Teixeira

In This Issue

We present in this issue of *CBTJ* a set of five articles that provide actionable insights on topics of current interest to professionals and executives. Our first article begins by demonstrating an intelligent enterprise. Joseph Byrum describes an intelligent enterprise as one that embraces AI to guide all its functions and decisions, small or large. However, this business is not run by the all-knowing, utopian artificial general intelligence (AGI) that science fiction writers and some commentators envision, which is a distant dream. Rather, it is an enterprise run by augmented intelligence — humans using AI and decision support tools that are enriched to the extent that is currently realistic and feasible. Byrum discusses the advantages of enterprises embracing augmented intelligence but cautions that making the entire enterprise "intelligent" requires concerted effort.

In our next article, Namratha Rao and Jagdish Bhandarkar outline the concept of intelligent automation using AI, ML, and RPA. A case study from the financial sector highlights the benefits gained through RPA. The authors explain how an intelligent bot can be trained and deployed over a period of a few months, and they emphasize establishing a roadmap, applying the right security measures, and setting up robust governance as three key tenets for scaling automation.

Currently, in most enterprises, business processes are automated in isolation, creating "automation silos" — a major barrier to realizing the fuller potential of enterprise-wide integrated automation. In their article, Aravind Ajad Yarra and Danesh Zaki address this issue. They differentiate between first- and second-generation smart automation and identify key imperatives to ensure desired integration across an entire business process. Furthermore, they present a detailed architecture for, and a pathway toward, smart automation 2.0, which enterprises can adopt to enable their automation bots to cooperate across the value chain.

Our next article discusses an interesting paradigm: human-machine hybrid intelligence. In their article, Tad Gonsalves and Cutter Consortium Senior Consultant Bhuvan Unhelkar argue that while machine intelligence facilitates smart automation and autonomous operations, yielding benefits, it cannot handle decisions that need to account for subjective factors, such as satisfaction, perceived quality, or joy, which cannot be parameterized in an ML algorithm. The authors recommend judicious superimposition of human natural intelligence (NI) on machine intelligence as a better way to facilitate business decisions that factor in customer value. In their discussion of how to achieve this goal, they also present a few use cases that embrace this hybrid intelligence.

Our concluding article focuses on another important issue facing enterprises and society: the governance and regulation of intelligent automation. Daniel J. Power, Ciara Heavin, and Shashidhar Kaparthi argue that a better governance mechanism is necessary to minimize the dangers of rushing to adopt AI and automation without due consideration of the risks. They present a governance framework for intelligent automation that includes all key stakeholders and offer policy prescriptions and guidelines for successful intelligent automation.

In addition to these articles, we recommend you also look at last year's *CBTJ* issue on automation,[5] which features eight articles that cover the technologies that drive and support new frontiers in automation, such as blockchain, AI, and security, and automation strategies and design considerations.

A New Automation Mindset

As we start to implement smart technologies to automate enterprise processes and activities, we must look at automation with a new mindset, holistically, broadening our vision. To gain dramatic benefits across an organization's activities, we need to move from traditional opportunistic automation of processes, which offers incremental benefits, to systematic, organization-wide automation of processes. However, we shouldn't blindly pursue and embrace automation without first examining the need for it, its relevance, and its consequences. Furthermore, to work along with machines, humans have to know the capabilities and limitations of machine intelligence and automation and change their own mindset and behavior to be compatible with machine behavior and activities.

To raise intelligent automation to higher heights, we need to pursue further development in the following areas: collaborative automation, where two or more systems work collaboratively, sharing insights and working toward a higher-level shared objective; human-machine interaction and transfer of roles when the situation warrants it; security and reliability; standards that facilitate integration and coordination between different processes; and governance and regulation.

We hope the articles in this issue present perspectives and ideas on intelligent automation that you'll find insightful. We also hope these articles inspire and encourage you to harness advanced automation in your domain of interest.

References

[1]Coughanour, Mariesa. "How Intelligent Automation Is Solving Unforeseen COVID-19 Challenges." Cognizant, 6 April 2020.

[2]Coughanour (see 1).

[3]Knight, Will. "The Pandemic Is Propelling a New Wave of Automation." *Wired*, 12 June 2020.

[4]Mithas, Sunil, San Murugesan, and Priya Seetharaman. "What Is Your Artificial Intelligence Strategy?" *IT Professional*, Vol. 22, No. 2, March-April 2020.

[5]Murugesan, San (ed.). "The Next Frontier in Automation: Opportunities, Challenges, and Impact." *Cutter Business Technology Journal*, Vol. 32, No. 2, 2019.

San Murugesan (BE [Hons], MTech, PhD; FACS) is a Senior Consultant with Cutter Consortium's Data Analytics & Digital Technologies practice, Director of BRITE Professional Services, and an Adjunct Professor in the School of Computing and Mathematics at Western Sydney University, Australia. He is former Editor-in-Chief of the IEEE's IT Professional. *Dr. Murugesan has four decades of experience in both industry and academia, and his expertise and interests include AI, the Internet of Everything, cloud computing, green computing, and IT applications. He offers certificate training programs on key emerging topics and keynotes. Dr. Murugesan is coeditor of a few books, including* Encyclopedia of Cloud Computing *and* Harnessing Green IT: Principles and Practices. *He is a member of the COMPSAC Standing Committee, a fellow of the Australian Computer Society, and Golden Core Member of IEEE CS. Dr. Murugesan held various senior positions at Southern Cross University, Australia; Western Sydney University; the Indian Space Research Organization, Bangalore, India; and also served as Senior Research Fellow of the US National Research Council at the NASA Ames Research Center. He can be reached at smurugesan@cutter.com.*

The Intelligent Enterprise Defines the Future of Business

by Joseph Byrum

The fact that businesses are able to adapt and serve the needs of hundreds, millions, or even billions of customers is a marvel of complexity management. Getting this right is no easy task. Most individuals who try their hand at starting a business fail within the first 10 years; only 38% of startups survive beyond that.[1] Even the companies that make it to the top cannot rest on past achievements lest they join the ranks of forgotten giants that have disappeared due to the bad choices of their executives.

Former multibillion-dollar giants like Blockbuster, Tower Records, and Toys "R" Us famously missed the signs that the public wanted the convenience of online options and that brick-and-mortar operations would have to adjust to this new reality to survive. Each of these companies had ample opportunity to follow the leads of other companies that transitioned to online-only operation or found a way to mix the best of both worlds in a way that satisfied customer needs. Their failure to act on the information available to them at the time proved fatal.

Through the use of artificial intelligence (AI) tools, companies today are boosting their decision-making skills with custom-designed systems that work on a task-by-task basis.[2] For example, a doctor might use an AI program to assist with diagnosing a particular disease. A stock analyst can use an AI tool to spot a hot investment overlooked by others. Shipping companies can use AI to plot more efficient routes to save on fuel and transportation expenses.

Several industries are increasingly using AI tools for logistics and production purposes. The benefit of optimizing core functions is well established. The discipline of operations research has been lowering the cost of doing business for decades by using data analytics. Just one small slice of such projects — the ones selected as finalists for the Franz Edelman Award for Achievement in Advanced Analytics, Operations Research, and Management Science — have brought home the tidy sum of US $250 billion in cost savings.[3] Businesses are investing $30 billion a year on AI tools to optimize operations.[4] That might seem like a big number, but less so when you consider that US businesses also spend more than $40 billion a year on landscaping.[5] There's room to think bigger about what can be done with AI.

If you can succeed through optimization in one aspect of a business, it makes sense that extending the concept to other areas of the enterprise can produce similar results. The question then becomes: what would happen if businesses were built from the ground up to take advantage of AI? Every aspect of such a business would be optimized to take advantage of integral AI tools. Everything from the C-suite to HR, the general counsel's office, and even the landscapers, would bolster their skills with AI. Such a business would open a new frontier in business efficiency: the *intelligent enterprise*.

The phrase "intelligent enterprise" describes a business designed around the use of AI that guides all decisions, big and small. While this might sound like it is turning every human function over to an impersonal machine, in this case it would be the opposite. The intelligent enterprise is not a business run by the all-knowing artificial general intelligence (AGI) dreamed up by the authors of science fiction novels and movie screenplays. Rather, it is an enterprise run by humans using much-less-scary augmented intelligence decision support tools firmly rooted in reality.

Augmented Intelligence vs. AGI

AGI makes for some of the silver screen's greatest villains: *Space Odyssey's* Hal 9000, *The Matrix's* Agent Smith, and Arnold Schwarzenegger's homicidal Cyberdyne Systems Model 101 in *The Terminator*. Benevolent AI — the sort sometimes seen on *Star Trek* — rarely makes an appearance. That's why the public associates AI with the dangers of AGI.

The belief that AGI is possible — or even inevitable — is a product of linear reasoning. If one assumes that machines are growing smarter and smarter every year, and human abilities are finite, it follows that one day, machines must, of necessity, reach a status that far exceeds human capabilities. Once this happens, humanity will be obsolete. For an all-knowing AGI, there would be no logical reason to have humans stick around.

Fortunately, developing AGI is a distant dream. In fact, we're not even particularly skilled at simulating an AGI that's good enough to win the Turing test, which measures whether a machine can be crafted to present a conversation convincing enough to fool a panel of human judges.[6]

Augmented intelligence tools, as they already exist in early forms, are firmly rooted in what's possible. Rather than replacing humans, these tools take existing human employees and make them better at their jobs. Today's augmented intelligence tools are developed for one-off purposes, but they can be extended to serve more functions. Such systems can be more easily designed to fit within appropriate ethical guidelines[7] because humans are always in control, unlike with the use of "deep learning" forms of AI, where the rationale underlying the decisions is hidden from the end user.

Man vs. Machine

Augmented intelligence is built on the basic premise that humans and machines have their own strengths and weaknesses. Machines benefit from having effectively infinite memory capacity. They crunch numbers and perform every task, no matter how mundane, with precision. They never grow tired.

Humans, of course, don't do well with tedious tasks. We grow tired or are bored easily when forced to perform mindless or repetitive functions. On top of this, our memory capacity is limited and subject to influence from emotion and stress.[8] We can't perform computations 24/7, and we certainly can't do so without error. It's a mistake, though, to assume that this makes humans inferior.

What we lack in endurance and capacity, we make up for in creativity, insight, inspiration, and judgment. Machines by their very nature follow the rules under which they have been programmed. Venture beyond the rule set, or the training database, and a machine becomes quite dumb, failing to recognize things that a child instantly would know.

People can be funny; machines can't tell a joke. As surprise is one of the key elements of humor, it says a lot that no AI is sophisticated enough to come up with a joke on its own. Alexa can tell you a knock-knock joke, but only one recorded in advance by an Amazon engineer. A truly funny response requires inspiration not available to the current generation of AI.

Augmented intelligence takes into account the relative strengths and weaknesses of humans and machines, combining human judgment with the unerring, untiring processing power of machines. The result is something more powerful than either a human or a machine could ever create on its own. With augmented intelligence, the machine (1) takes in all the data, (2) processes and sorts out what's a priority, and (3) suggests possible courses of action based on a statistical analysis of possible outcomes.

Augmented intelligence takes into account the relative strengths and weaknesses of humans and machines, combining human judgment with the unerring, untiring processing power of machines.

This is an extremely powerful capability. AI can absorb the sum of human knowledge and make an analysis based on a review of every available scientific study on the subject in question. For example, medical AI tools can process every medical journal article so that a symptom of a rare disease mentioned in the footnote of an obscure foreign study, perhaps written in a language the doctor doesn't understand, can be flagged for the doctor to investigate. This is an example of how a level of knowledge once available only to the top specialists in a given field can be made available to everyone. The system sorts through the raw data, presenting only those facts most relevant to the task at hand.

With an augmented intelligence algorithm performing "information triage," the human mind is freed of the tedium of that work. No longer overwhelmed with raw data, the human has the mental capacity to focus on the details that matter. From this, an informed judgment can be made about what action to take. The judgment will be more effective because it's not just a human operating on a whim or a hunch. An objective, emotion-free, and scientific evaluation of the facts bolsters the decision.

Applying Augmented Intelligence to the Enterprise

Such powerful tools would be the foundation of the intelligent enterprise. At the top, the C-suite could use them to guide all major decisions in setting the company's direction. Computer models would test scenarios about all the typical top-line decisions about funding levels, adapting to economic conditions, and production. The human CEO would play an essential role, checking the impact of each potential choice against the likely impacts on, for example, employee morale. The most efficient choice on paper (or in an algorithm) is not always the best choice. The CEO also provides inspiration and motivation that machines simply cannot provide.

> *The intelligent enterprise at this point is still just a thought experiment, but it is achievable.*

The decisions humans make wouldn't simply be rubber-stamping the offerings of the "superior" machine intelligence. For instance, AI might calculate that moving a call center to a low-tax jurisdiction would reduce expenses. The entire cost of the move could be covered with savings over the next five years. The human executives would look at the option to determine the impact of the move on the personnel. Is that proposed location a good fit for the company? Would the move have a cost in terms of losing key staff? Would the best workers need to be replaced, and could you even find the right talent in that location?

More routine, day-to-day work decisions would also be bolstered by AI in the intelligent enterprise. At a bank, augmented intelligence systems would crunch the data for loan applications. The systems would look for evidence of fraud. They would track payments and assist with customer service. Many of these functions are already automated in this way. The difference: these augmented intelligence systems would be integrated and have the capability of communicating with one another across divisions and subdivisions of the enterprise.

The availability of data from subsidiaries and other business divisions expands the available data, enhancing the analysis. For example, a manufacturing division might want to reduce inventory if a financing division forecasts an economic slowdown or, conversely, it could ramp up production when given signs of boosted consumer spending.

Even the landscaping for the physical space of an intelligent enterprise would be optimized. The system would help choose which seeds to plant and when in order to keep all the plants looking their best under constantly changing weather and soil conditions. The system would optimize watering by having sprinklers operate only when needed and would order trimming at the most efficient times.

Employees might be happier working in an intelligent enterprise. Beyond the nicer landscaping at the headquarters and the fact that office supplies would never run out, employees in this new enterprise would be free from the most tedious tasks. Their contributions to the organization's objectives would be more obvious, and their performances could be better measured and appreciated because the intelligent enterprise tracks all decisions made and compares them against the expected results.

A Powerful Thought Experiment

The intelligent enterprise at this point is still just a thought experiment, but it is achievable. It takes years of concerted effort to build and, more importantly, to validate custom AI tools for a given, limited purpose. It's too time-consuming and resource-intensive to develop augmented intelligence systems for anything other than the core functions of a business. Until more general-purpose augmented intelligence tools are available that could be quickly adapted to each new function, we probably won't see a true intelligent enterprise.

Once shrink-wrapped, interoperable tools are available for executives to deploy and optimize any business task at hand, the intelligent enterprise would seem to be almost inevitable. Remove the cost of development from the equation, and optimization throughout the enterprise becomes pure profit.

Until that day comes, however, we can draw a few conclusions about the nature of AI from thinking through how an intelligent enterprise would be designed. The first is that, as long as AGI does not exist, humanity is not in danger of being replaced by machines any time soon. This is not to say that individuals won't find their particular profession rendered obsolete. It is inevitable that mundane tasks that can be automated will be automated. What's left for humans will be the role of taking the information provided by the machines and deciding what to do with it.

Far from a demotion for humanity, such roles fit our natures perfectly. Of course, just as every revolution in technology has caused economic dislocation, the transition to the intelligent enterprise won't be easy. Ultimately, however, it represents a bright future.

References

[1]"Table 7. Survival of Private Sector Establishments by Opening Year." Business US Bureau of Labor Statistics, 2020.

[2]Phillips-Wren, Gloria. "AI Tools in Decision Making Support Systems: A Review." *International Journal of Artificial Intelligence Tools*, Vol. 21, No. 2, April 2012.

[3]"Franz Edelman Award for Achievement in Advanced Analytics, Operations Research, and Management Science: 2019 Winners." The Institute for Operations Research and the Management Sciences (INFORMS), 2020.

[4]Bughin, Jacques, et al. "Artificial Intelligence: The Next Digital Frontier?" McKinsey Global Institute , June 2017.

[5]"Landscape Industry Statistics." National Association of Landscape Professionals, 2020.

[6]Biever, Celeste. "No Skynet: Turing Test 'Success' Isn't All It Seems." *New Scientist*, 9 June 2014.

[7]"Ethically Aligned Design, First Edition." IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, 2019.

[8]Pezdek, Kathy. "Fallible Eyewitness Memory and Identification." In *Conviction of the Innocent: Lessons from Psychological Research*, edited by B.L. Cutler, American Psychological Association, 2012.

Joseph Byrum is Chief Data Scientist at Principal, a global investment management group based in Des Moines, Iowa, USA. Previously, he was an R&D executive at Syngenta, an agricultural firm based in Basel, Switzerland. Dr. Byrum holds a PhD in quantitative genetics from Iowa State University and an MBA from the University of Michigan. He can be reached via Twitter @ByrumJoseph.

Intelligent Automation: An Alchemy of Technology and Human Intelligence

by Namratha Rao and Jagdish Bhandarkar

Evolution of Automation

One look inside Mercedes' famed "Factory 56" in Germany stops you in your tracks.[1] Production of its latest class of cars is wholly automated; from planning to assembly of parts to quality monitoring to the finishing touches, every aspect is controlled by digital systems, which are powered by artificial intelligence (AI). The team of people in the plant works in tandem with these systems, and the result is spectacular. The time needed to produce the cars has been drastically reduced, while the quality has risen exponentially.

Blending intelligence with bots provides an opportunity to increasingly automate the control aspect of tasks.

Automation has always been around us. Our collective intelligence has explored ways and means to achieve more with less, beginning with simple tools and evolving to complex machinery that has given us the advantage of efficiency. Given that mechanical automation started the current trend, it is only natural that the manufacturing industry (as in the Mercedes example) adopted automation before other industries. Manufacturing results have been highly visible, and more players have followed the automation trend. The software industry is now catching up, with the emergence of techniques to apply machine learning (ML) concepts with a combination of high computing capabilities. The intent is the same as in any other industry: to achieve higher levels of efficiency in the least amount of time and at the lowest possible cost. The creation of a software bot (i.e., a digital twin of the mechanical bot on the factory floor) has accelerated the embracing of software processes and systems automation.

Robotic process automation (RPA) is changing the way organizations operate. New metrics, key performance indicators (KPIs), and key result areas (KRAs) have evolved, and a new perspective has emerged to assess operations. In this article, we attempt to describe how an organization can adopt these new concepts.

Blending of Intelligence and Automation

Automation, by definition, is utilizing a tool or a device instead of a person to perform a task. The decision of what to automate has been driven mainly by the nature of possible tasks to be automated. A task usually consists of control and execution, with execution being the easiest to automate because it does not contain any decision-making points. RPA focuses on the execution aspects of a task. In contrast, the control aspect encompasses all those points where a person needs to make a decision using human intellect and an understanding of the task itself. Common bots do not have the capability to make decisions unless a decision is purely rule-based. Blending intelligence with bots provides an opportunity to increasingly automate the control aspect of tasks.

With the advancement of ML techniques, bots are now augmented with intelligence, taking automation a notch higher. This is typically referred to as *intelligent automation*. Intelligent automation makes use of AI and analytics to provide automation with some awareness of its environment and an ability to configure and heal itself to evolve in response to client requirements.

Figure 1 shows how intelligent automation moves from automating transactional-based processes to judgment-driven processes utilizing key concepts such as ML and AI. In the following section, we use a case study of a financial institution to describe how intelligent automation can help organizations.

Case Study: RPA and Intelligent Automation in a Bank's KYC Process

US financial institutions under the Bank Secrecy Act (BSA), the US's primary anti-money laundering (AML) regulation, are required to provide information on all

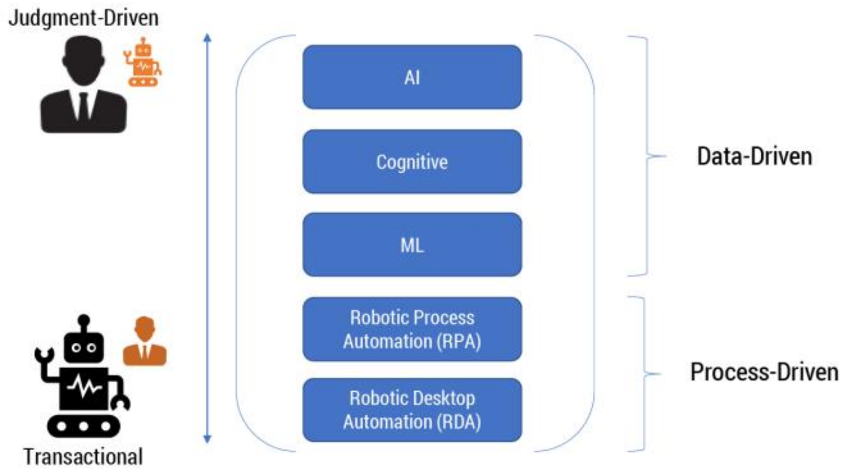

Figure 1 — Intelligent automation toolkit.

customers to US government agencies to detect and prevent money laundering. The cornerstone for preventing money laundering is to learn as much as possible about both new and old customers, which is the aim of the Know Your Customer (KYC) program. The foundation of an effective BSA/AML program includes comprehensive customer due diligence (CDD) policies, procedures, and processes.[2] Accounts that the CDD procedures flag as high risk are subject to enhanced due diligence (EDD). The overall increase in the regulatory requirement complexity of the KYC, CDD, and EDD policies hinder institutions' operating speed. However, due to the risk mitigation of these processes, KYC-CDD is core to the onboarding of new customers and has become the most strategic area of the compliance chain.

Our case study examines the implementation — utilizing intelligent automation — of KYC-CDD with the EDD process in a US super-regional bank. The CIO wanted to assess how the bank was handling the KYC process and regulatory compliance around this process. She assessed the process from three perspectives:

1. **From her colleagues in the AML line of business, she learned that the ever-increasing regulatory complexity and the broadening scope of the regulatory requirements were making analysts' jobs tougher.** The analysts had to monitor numerous sectors and information sources, and the lists generated were then screened against the names of the institution's clients. (These names numbered in the millions and ranged from politically exposed persons to state-owned companies; this detail signifies the level of risk with which intelligent automation models must deal). If the bank failed to comply with the regulations, potential fines could run into millions of dollars, in addition to exposure to high remediation costs.

2. **From the customer onboarding analysts' team, she understood that the onboarding and post-matching processes were time- and cost-sensitive.** The current process involved dealing with 70%-80% of alerts being false positives and with more than 20% needing review by more than three senior analysts. The resulting balancing act weighed the risk of huge penalties against the ever-increasing cost of manual investigations.

3. **Her final check was with the customer experience department to correlate the attrition of prospective customers with the KYC process.** She came to the realization that the time-consuming identity checks created a bitter experience for the bank's potential customers and yielded a very negative customer experience. Hence, the rate of attrition of prospective customers was high.

The overall scenario in the bank was concerning, and the fines that the bank was paying made it more so. Many repetitive and manually intensive processes made evident a dire need to improve the efficiency of monitoring activities, which in turn would help reduce errors and speed up the whole process. The CIO proposed introducing RPA to implement a digital twin concept. The robots deployed alongside the analysts to perform the monitoring process would be well suited for spotting anomalies. They could flag anomalies to the analysts, who could then perform confidence scoring and due diligence.

Upon approval from the board to perform a pilot for automation, the newly established automation team started analyzing the KYC-EDD process for RPA

	Manual	RPA	Benefits Realized by Bank
Entering customer data	Analysts refer to scanned customer identification documents and populate the CRM and other systems with key customer information.	Bots automate the activity of populating the CRM and systems.	• Increased time for analysts to perform more value-added tasks
Gathering customer information	Analysts validate customer information by collecting social media information, merging data, extracting data from various documents, accessing various databases, and then populating the forms.	Bots perform the work of collecting customer data from several information sources and entering that data into the required systems.	• Improved operational speed • Enhanced quality due to greater accuracy and an updated audit trail with accurate information • Increased analyst productivity
Compiling customer information	Analysts compile information, such as brokerage or savings information, in accordance with services sought by customer.	Bots gather customer information from disparate systems to provide analysts with a holistic view of the customer.	• 360-degree view of customers • Increased analyst productivity

Table 1 — Manual activities to be automated with RPA.

implementation. Table 1 shows the assessment of traditional manual activities that were good candidates for RPA.

The RPA implementation quickly brought about a noticeable difference. Customer onboarding analysts saw the time required to collect data drastically reduced due to the bots taking over that task. Furthermore, the review process could now be done by two senior analysts rather than three. The customer experience improved within two quarters due to the ease of the process, shortened cycle time, and a reduction in the number of human-introduced errors.

It was now time to introduce intelligence into the process. Intelligent automation could be used to analyze the information captured, especially in the areas of negative news analysis and transaction monitoring of accounts. Intelligent automation would augment analysts' ability to identify accounts at risk for money laundering.

Unlike RPA, which was implemented quickly at minimal cost, integrating ML solutions with the current tools landscape at the bank would need a phased approach. The first step was to train the intelligent bots using historical data. Once the bots had been trained, an analyst would validate the bots' decisions. As time progressed, confidence in the decisions made by the bots would increase until it reached a level where the majority of the KYC processes could be handed over to them. Table 2 shows the areas of the KYC process where intelligent automation played a role.

The success of the approach depended on the quality of the historical data used to train the AI models. However, the phased approach was sufficient to convince regulators that this step-by-step approach yielded high accuracy within the KYC process. Intelligent automation further increased analysts' efficiency by allowing them to focus on the flagged cases. Intelligence blended with RPA increased both speed and accuracy.

The CIO announced the very-promising results to the board. In summary, the results showed that:

- Intelligent automation encompassing AI and ML models led to a reduction in false positives, from 70% to 15%.
- There was a 60% overall improvement in investigation efficiency.
- Efficiency provided by automation and by analysts' having a digital twin increased system efficiency by more than 80%.

Step	Applying Intelligent Automation	Benefits Realized by Bank
Information capture from customers	The intelligent OCR engine is used to identify details from scanned images, rather than the analyst reading the images manually.	More analysts' time can be invested on exception cases that need human judgment.
Risk profiling of customers	ML algorithms detect and flag sudden activity of certain types in dormant accounts. This enables continuous monitoring of customer accounts in real time. Bots search for any negative news or any changes affecting the watch list. A profile dossier is created for each case.	Analysts' digital twins provide analysts with completed dossiers, which include a confidence score for the created profile, assertions, and findings.
360-degree view of customers	Network analysis helps build customer profiles, which include transactions and counterparties, to provide a holistic understanding of the customer.	Accurate information helps analysts obtain approvals faster.
Periodic review of customer information	ML models institutionalize periodic reviews of the customer profiling.	Implementing such ML models helps reduce number of false positives compared to using traditional rule-based parameters for the analysis.

Table 2 — Areas of the KYC process that intelligent automation enhances.

- Relationships and risks that would have otherwise gone undetected were identified.
- Planning and decision-making time was reduced, with a 76% increase in employee satisfaction.
- Customer retention went up by 48%.

Implementing Intelligent Automation

How does an organization decide to take up the route of intelligent automation? Are there benefits beyond cost savings? Yes, in addition to the cost benefits, advantages of improved quality, speed, governance, security, and business continuity are realizable due to a high degree of automation.

A few important questions may provide direction to an organization that is undecided about taking the first step. Figure 2 shows a sample list of questions. A higher number of "yes" answers to these questions would indicate that automation is the right decision. Almost any project an organization undertakes will advance through the program lifecycle shown in Figure 3.

With intelligent automation projects, certain key elements must be considered within each stage of the lifecycle, including:

- **Discovery.** The initiative needs to be well scoped, with the main focus on identifying and improving the processes to be automated, since automation will not just amplify existing efficiencies but will also amplify any inefficiencies. For instance, if the process under consideration for automation has flaws, then automation without first improving the process will prove costly in later stages. "Garbage in, garbage out" is applicable to automation.
- **Pilot and scope.** Proof of value should be carried out after the proof of concept to provide an indication of the initiative's ROI.
- **Software selection.** An important decision involves the standardization of the tool stack to ensure the selection of the right intelligent capabilities and RPA features.
- **Build, implement, and deploy.** The business does the process design, in contrast to other projects where this task is carried out by the IT team. Having the business do the design ensures that the translation

	Questions	Yes	No
1	Are the tasks in your organization structured, rule-based, repeatable, and computer-based?	○	○
2	Do the tasks in your organization involve searching, collating, and updating information?	○	○
3	In your organization, in order to complete a process, do you access one or more systems?	○	○
4	Are there tasks in your organization that perform simple or complex decision routines?	○	○
5	Does your organization fall under a highly regulated industry (e.g., banking, financial services, healthcare)?	○	○
6	Does your organization have fluctuating volumes when it comes to seasonality, client transitions, and production rollouts?	○	○
7	Finally, do you have a CFO who is looking for geography and standardization models and cost reductions?	○	○

Figure 2 — Sample questionnaire.

Figure 3 — Lifecyle of a typical program.

of requirements to design is nearly flawless, as the business has the knowledge experts. Operationally, an automation center of excellence should be set up that focuses exclusively on automation-related R&D.

- **Monitoring and benefit realization.** Continuous monitoring of KPIs is crucial to realize the primary objectives of the program.

3 Key Tenets of Automation: Roadmap, Security & Governance

Since intelligent bots learn on the job, efficiency might not be realized right at the beginning. Typically, a time period of around two to three months is needed for an intelligent bot to fully take over from its human counterpart.

For instance, consider a typical case of a team of employees working to capture data from emails, social media, and images/scans to spreadsheets that is then sent to back-end processing. Simple RPA bots can be programmed to work on back-end processing that is rule-based (see Figure 4).

However, to help reduce the amount of effort the team spends in capturing data, an intelligent model needs to be deployed. This ML model is trained on historical data. After an extended period of training, the intelligent bot can take over part of the task (see Figure 5).

This simple example illustrates how an intelligent bot can be trained and deployed over a period of a few months. Once a few initial automation projects have shown success, an enterprise's next step is to scale this success. Establishing a roadmap, applying the right security measures, and setting up robust governance are key tenets for scaling automation.

Establishing a Roadmap

An extensive process-discovery exercise will yield a process heat map that shows the types of processes that can be considered for automation. Assessing these processes reveals their complexities. Figure 6 shows a typical automation suitability assessment for creating a roadmap. The roadmap should slot the processes, based on ROI, into a timeline, which should be reviewed periodically. This artifact will drive the implementation in a manner that brings maximum returns to the enterprise.

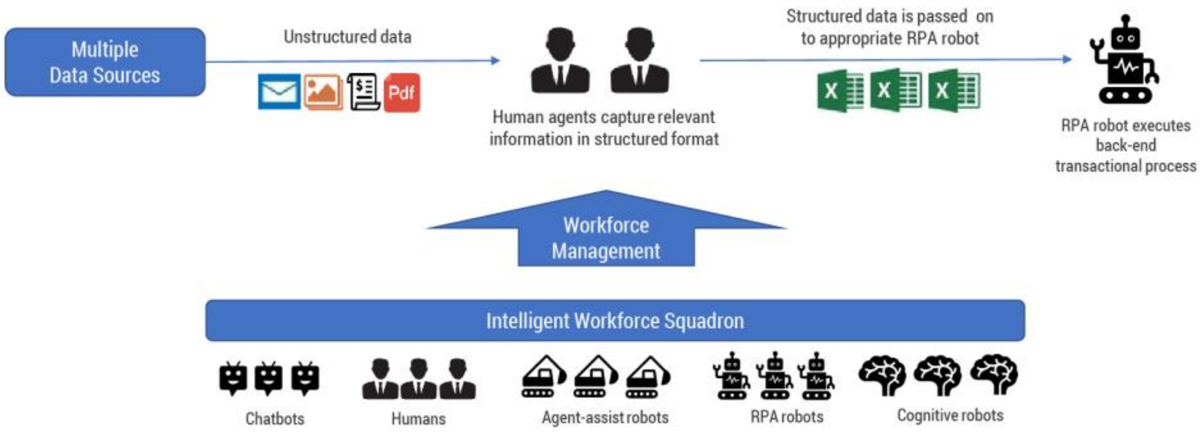

Figure 4 — Capturing unstructured data without intelligent automation.

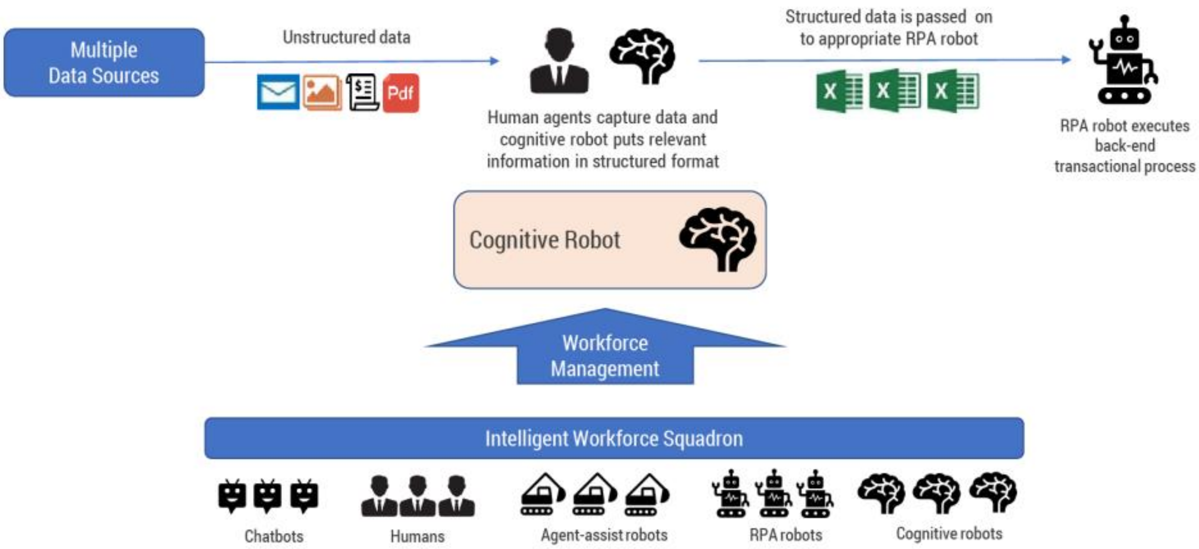

Figure 5 — Capturing data with a trained intelligent bot.

Applying Security

Security is a key consideration when implementing automation. The most vital aspect is to secure the process. Elements of applying security include the following:

- When we automate a process, it is essential to validate that personally identifiable information (PII) encryption is in place for data at rest and for data in motion.

- Bot developers need to be trained on risks and controls so that when bots are developed, bot credentials are considered (e.g., the inclusion of strong, unique passwords that comply with standard company policies or the disabling of an interactive login from any source other than the virtual machine that runs it).

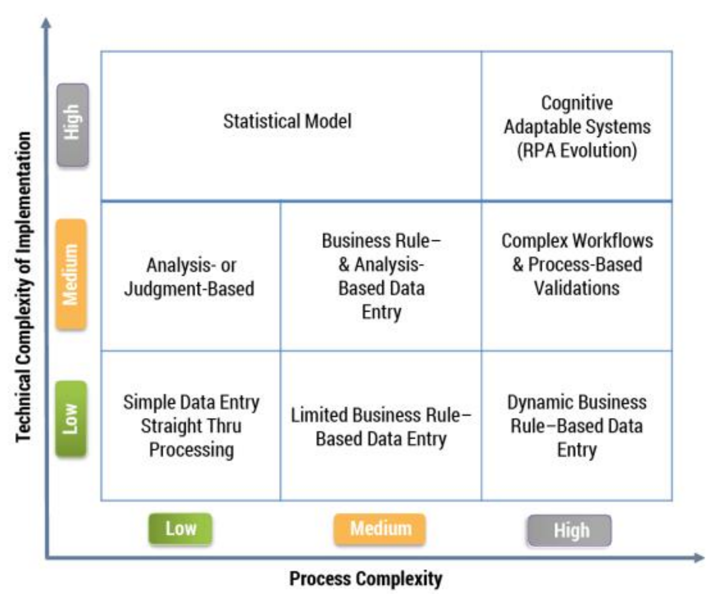

Figure 6 — Automation assessment and suitability.

- The environment in which the bots run needs audit logging of account activities, the credentials must be stored in vaults, servers need to be hardened, and all applications must have role-based access control.
- To determine any weaknesses on the technology front, vulnerability scanning and threat modeling need to be performed periodically.

> *The right blend — alchemy — of human intelligence and AI technologies has the power to bring about enormous changes in the way we run our organizations.*

Establishing Governance

As with the roadmap and security, governance also needs a distinct approach when it comes to automation programs. The very nature of automation, with bots coexisting with human counterparts and overlapping the IT and business domains, creates security- and privacy-related concerns that require new attention to governance. Figure 7 shows a few key governance principles.

Change management needs to incorporate an additional area for scaling automation implementation. Change in this case involves not only people and processes, but also robots (the new entities participating in the day-to-day business). Incorporating robot maintenance into the standard change management framework is imperative.

Apart from the three key tenets discussed above, it is important to ensure the consideration of the softer aspects of such a program, including:

- **The people factor.** It is crucial that an organization determines impacts and mitigates any risks associated with productivity and behavioral issues.
- **Leadership.** Alignment of leaders at various levels of an organization assists in successful implementation of advanced automation.
- **Communication.** Communication is key to ensure that automation is accepted at all levels of the organizational hierarchy.

The Promise of Intelligent Automation

There is no dearth of case studies in the world today that highlight human beings' innovative abilities to make life more comfortable — the holy grail we are always seeking.

The right blend — alchemy — of human intelligence and AI technologies has the power to bring about enormous changes in the way we run our organizations. Nearly nine in 10 (i.e., 87.3%) manufacturing organizations have already adopted some form of smart automation; in the banking and finance sector, 84.7% of financial institutions already report RPA and AI implementations. About 60% of retailers have adopted some form of smart automation.[3] Most of these organizations, we must assume, are rearchitecting their systems to align with the new normal.

While the promise of such technologies knows no bounds, applying the concepts with a clear objective

Should Do This	Must Do This
Strategy Alignment & Measurement: • All automation aligns to enterprise strategy and can be measured for meaningful outcomes. **Technology:** • Platforms/products are utilized; development standards are defined and continuously improved. **People:** • There is motivation and leadership of the teams involved in automation.	**Safe:** • Automation deployed uses standards that ensure SLAs do not exceed the organizational risk appetite. **Secure:** • Deployments support accountability and ensure security by not trespassing critical controls (e.g., unified access management/dual control). **Compliant:** • Deployments do not result in bypassing regulations and comply with regulations for the relevant domain.

Figure 7 — Governance in automation: "should do this" and "must do this."

and a roadmap is imperative to attain sustained progress; or else, in the words of Aristotle, "Well begun is half done." As with all emerging technologies, the promise of AI could fade away if implementations are not well thought out and supported during the entire lifecycle. Organizations that crack this challenge will be the winners in the race.

References

[1]Padeanu, Adrian. "Mercedes' New Factory Looks Straight from a Sci-Fi Movie." Motor1.com, 16 November 2018.

[2]Vasudevan, Jagannathan. "KYC — Enhanced Due-Diligence." White paper, Association of Certified Anti-Money Laundering Specialists (ACAMS), 2020.

[3]Automation Anywhere. "New Research Shows That Businesses Have Passed the Tipping Point Towards Universal Intelligent Automation Adoption." Press release, Cision PR Newswire, 4 November 2019.

Namratha Rao is a techno-change management expert with nearly two decades of industry experience. She is AVP of Product Engineering Development at SLK Software Services, where she heads the organizational change management, product development, and applied innovation areas. A thought leader and practitioner in "blue ocean" strategy, Ms. Rao has led many initiatives in SLK's Value-Based Advisory Services to create a differentiation for SLK and its clients. She earned a bachelor's degree in engineering and is currently pursuing her master's degree in the disruptive innovation space. She can be reached at namratha.rao@slkgroup.com.

Jagdish Bhandarkar is a technology geek with 20+ years of industry experience with various Fortune *500 companies. He is the CTO at SLK Software Services, responsible for driving the company's technology vision of continuously enhancing innovation on behalf of customers at a global scale. Dr. Bhandarkar has authored various articles for journals and conferences. He earned a master's degree in engineering and a PhD from MIT. Dr. Bhandarkar is an alumnus of the Indian Institute of Management Calcutta, India, and graduated from Mangalore University, India. He can be reached at jagdish.bhandarkar@slkgroup.com.*

Breaking Automation Silos: An Integration Approach for Smart Automation 2.0

by Aravind Ajad Yarra and Danesh Zaki

With initial implementations of automation now showing promise across industries, smart automation has become an important priority. Whether building conversational assistants at a bank or automating goods allocations to warehouses at a retail firm, businesses in all sectors are investing significantly in their automation capability. Smart automation initiatives are no longer simply "nice to have" but have now become a necessity, an important element of business today. Smart automation has direct impacts on businesses. A recommendation bot guiding customers with product selection, a chatbot supporting returns of items, a handwriting recognition bot verifying supporting documents during a claims process — all are great examples of how smart automation can help businesses with innovation, the cost of operations, and the ability to be nimble.

> *Given the narrow scope within which first-generation smart automation bots operate, they have proved to be highly effective and practical for use in business.*

First-Generation Smart Automation

Current implementations of smart automation represent the first generation of the technology, which has the following characteristics:

- Employed for point tasks in a business process, such as identifying the next best action in response to a customer complaint or answering frequently asked questions (FAQs) about products via a bot
- Often modeled as bots, which are independent and highly focused on their scope of operation (which is often a very narrow domain)
- Applied in a domain that is narrow enough to keep complexity in check; data within such a narrow domain is analyzed and contextualized to build and use the machine learning (ML) models needed for smart automation
- Often employs human-in-the-loop design to ensure smart automation works effectively for a given task

Effectiveness of First-Generation Smart Automation

Given the narrow scope within which first-generation smart automation bots operate, they have proved to be highly effective and practical for use in business. Since these bots often have human oversight (human-in-the-loop design), it is easy to identify exception scenarios and fall back onto alternative methods, such as manual handling of tasks. Nearly all business processes in modern enterprises now use some form of smart automation bot. However, first-generation automation bots have also led to widespread limitations, or *automation silos* (see Figure 1).

Building bots with a narrow scope results in silos of automation within enterprises. These silos ensure that the scope and impact of automation is often very limited, confined within a single business process. Based on our experience with various businesses in the financial services, retail, and technology sectors, the effectiveness of first-generation smart automation on the value chain of any given business is low, considering less than ~5% of the total manual tasks that currently happen in a value chain can be automated.

With smart automation's ultimate goal being fully autonomous processes across a value chain, with bots working independently while also working with each other to achieve full autonomy, these silos lead to a lack of sharing of data and context across process steps. These automation silos mean the promise of smart automation is unfulfilled.

Since first-generation smart automation bots work independently, they do not cooperate across a business

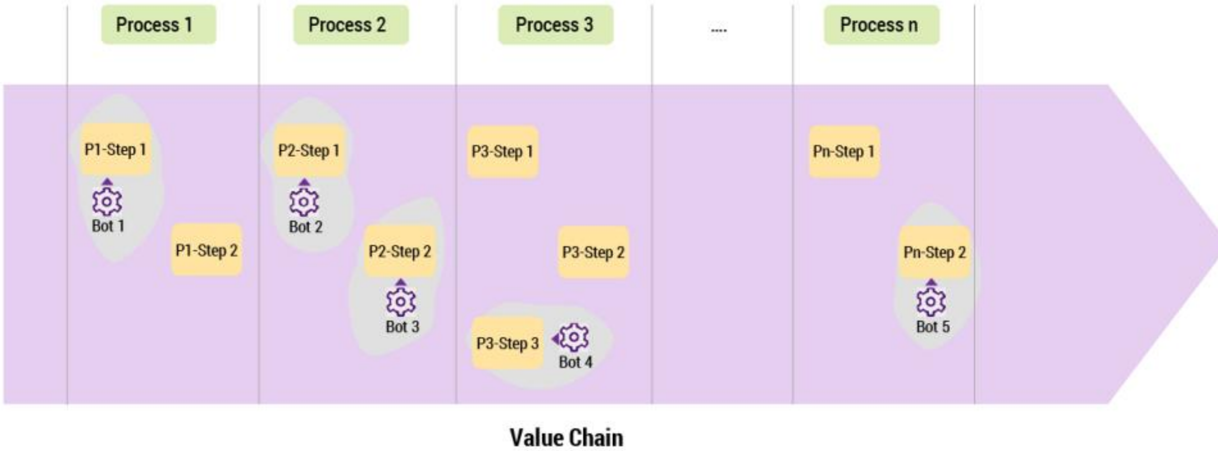

Figure 1 — First-generation smart automation: silos of automation, with each bot operating within a process step across the value chain.

process or across the value chain. For example, an automation bot working on conversation automation during a customer service process step often cannot access the intelligence gathered by another bot that is learning product preferences in a product catalog process step. Automation bots must be better integrated to cooperate with each other in real time.

Automation silos also mean best practices around data wrangling, cleansing, and model building are not being shared. While automation helps with agility and produces faster release cycles, automation within silos creates high amounts of duplication of effort when developing the automation and in the technology used. Bot development teams in different silos often cannot share their knowledge, as each team has different technology stacks, architecture, and governance practices.

One could argue that overcoming these challenges is simply a matter of integrating these automation bots using classic integration technologies and patterns. However, certain limitations make this impossible. Those limitations include:

- While automation works best when the scope of a bot is within one domain, extending the scope across domains adds to an impedance mismatch across domains, making automation modeling extremely difficult.
- Integration of multiple automation bots involves building shared knowledge repositories, which are extremely difficult to model and build across domains, from a design, operations, and governance point of view.
- Integration is extremely difficult in the context of shared learning (as in how ML models learn). Techniques such as transfer learning, for example, cannot work across domains when the bots are very different from each other.
- Orchestrating the integration of bots across a value chain externally is not desirable, as it leads to rigid architectures, creating lower resilience and reduced agility.

Second-Generation Smart Automation

To address the challenges with smart automation, we need to build on and improve the current generation. The second generation of smart automation will address the integration limitations. The effectiveness of the second generation, however, depends upon breaking the automation silos (see Figure 2).

Smart automation 2.0 is about amplifying automation effectiveness across business process steps and achieving automation across the value chain. Achieving this goal requires addressing the challenges posed by current automation silos and the limitations created by the architectural approaches currently used. Table 1 lists some of the changes that would occur between smart automation 1.0 and smart automation 2.0.

A Supply Chain Scenario

In an end-to-end supply chain process, humans will handle some parts of the process, while bots will perform others. The bots work in coordination with humans and with other bots. Some bots, for example, might help with invoice matching and in other order-related areas and would collaborate with bots from freight management used to measure and manage containers (see Figure 3).

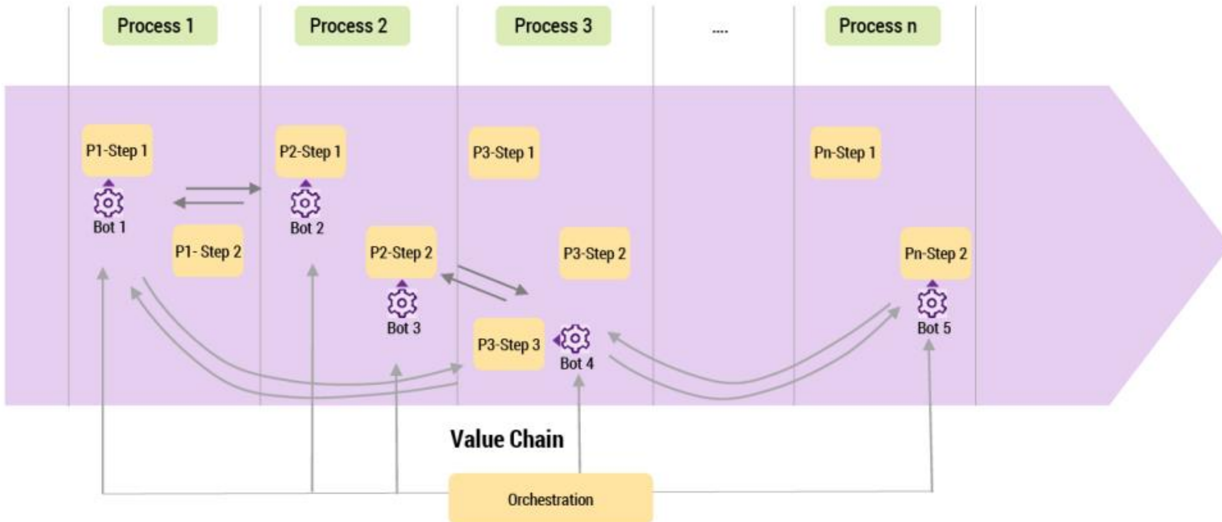

Figure 2 — Automation bots sharing and cooperating with each other in a peer-to-peer manner, orchestrated across the value chain.

Smart Automation 1.0	Smart Automation 2.0
Bots are autonomous and work independently.	Bots are autonomous but also cooperate with each other.
Knowledgebases and learning are specific to a bot within a process step.	Knowledgebases are shared across the bots and across domains throughout the value chain.
No integration exists between bots.	Bots can coordinate and collaborate by communicating with each other synchronously or asynchronously.
There is no process and value chain view of automation.	A composite view of automation across the value chain amplifies automation effectiveness across the value chain.

Table 1 — Moving from smart automation 1.0 to smart automation 2.0.

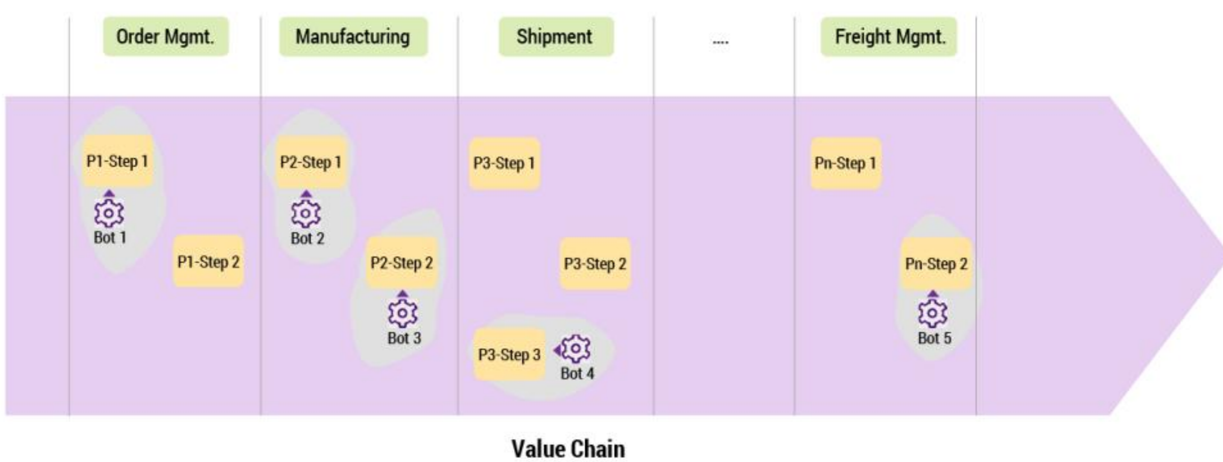

Figure 3 — Automation bots in a supply chain process.

In this scenario, first-generation smart automation can be manifested as follows:

- Bots are employed for point tasks (e.g., order matching, intelligent shipment scheduling, and freight container optimization).
- Bots are independent and highly focused on their scope of operation (e.g., the order management domain or the shipment domain).
- A human-in-the-loop design is used (e.g., as liaison with shippers to enable effective work on a given task).

As we have noted, first-generation smart automation bots operate in specific domain silos. Inadequate cooperation among bots results in issues such as shipment schedules not matching the container optimization recommended by a freight bot, an issue that might necessitate rework or manual intervention. Employing smart automation 2.0 can help improve the scenario in the following ways:

- Bots across the steps of order matching, shipment scheduling, freight container loading, and so on, cooperate to share intelligence gleaned through analysis or "conversations."
- Bots collaborate with each other throughout the process by sharing knowledge and coordinating actions.

Bot cooperation means that a freight bot can design a container optimization strategy that perfectly matches the shipment schedules. Any change in shipment schedules can be accommodated in real time with the container allocation readjusted accordingly.

An Integration Approach to Smart Automation 2.0

From an architecture perspective, the following imperatives are key to ensure that integration works at the scale of an entire business process and meets the goals of smart automation 2.0:

- **Preserve autonomous behavior.** It is essential that automation bots be developed and delivered autonomously and in a loosely coupled manner as required for each step within a business process. This is extremely important to allow for agility and the changing nature of business process automation.
- **Share data and knowledge.** Sharing of data and knowledge plays an important role in ensuring that automation bots can understand each other and work in a cooperative manner. The challenge, however, is to have sharing happen without building centralized and global data and knowledge models and stores. Creating centralized, enterprise-wide models and data stores is an extremely complex endeavor, and such models and data stores are often not practical to implement.
- **Design the cooperation of bots.** Externally orchestrating automation bots across a business process is not pragmatic and is one of the limitations of smart automation 1.0. To attain smart automation 2.0 goals across a process, a mechanism of designing cooperation to ensure automation bots work together is imperative.
- **Integrate heterogenous technology landscapes.** Enterprise business processes across a value chain typically use multiple technology stacks, making them heterogenous. An approach to achieve integration of such heterogenous technology landscapes is extremely important.

Reference Architecture

We can think of smart automation 2.0 architecture as designing cooperation among autonomous bots to realize a value chain–scoped automation goal. In other words, these autonomous bots cooperate with each other to accomplish this shared automation goal across the process. Such an architecture needs to address conflicting goals (see Figure 4):

- Addressing process-wide automation goals, yet retaining the autonomous behavior of bots
- Having bots leverage knowledge across the process and learn from each other, yet not hamper themselves with rigid integrations

These goals require a finely balanced architectural approach with capabilities that go beyond those of traditional integration architecture. Table 2 describes the key capabilities such an architecture should have.

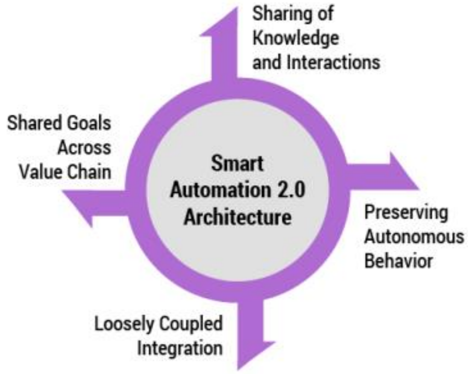

Figure 4 — Balancing conflicting architecture goals.

Putting together a smart automation 2.0 reference architecture with the key capabilities in Table 2 requires a combination of architectural styles. Table 3 describes the relevance of each architectural style to smart automation 2.0.

Figure 5 (on page 24) illustrates how these architectural styles are used to create the necessary building blocks of a reference architecture for smart automation 2.0.

The *bot registry* maintains a catalog of bots that perform automation tasks across domains. The registry is used by each bot individually for cooperation or for composition using an external component. For example, a bot doing data collection in the order management domain, by using character recognition in scanned PDFs, uses the registry to find and cooperate with a bot that can intelligently validate the address in the shipment domain to complete the entire business process. Distributed data stores, such as etcd, Consul, or the equivalent, are options to build a bot registry. (Table 4 [on page 24] describes possible technology choices for each building block.)

Events are chosen as the preferred mechanism of communication. Events enable the architecture to be loosely coupled and facilitate communication through a well-defined interface abstracting the various implementations of the bots. This helps avoid the impedance mismatch–like situations that occur in databases. A common standard for events, such as CloudEvents or AsyncAPI, goes a long way in helping bots converse and understand each other well. Events

Capability	Description
Context beyond the domain	Bots, although autonomous, also need to consider the overall context of the value chain when they need to cooperate with other bots to achieve automation goals. Context includes the overall process structure, process metadata, and the overall automation goals.
Knowledge graph	Bots maintain a knowledge graph with predefined rules as well as the knowledge gained during training and ongoing operations. However, such a knowledge graph is limited to local automation. Value chain–wide automation requires a broader knowledge graph that encompasses the scope of the entire value chain.
Bot cooperation plan	Automation goals need a higher-level plan for bot cooperation. Instead of defining a rigid orchestration, the bot cooperation plan describes the cooperation goals and constraints so that bots can come together in a loosely coupled but structured way.
Bot catalog	When bots need to cooperate, a list of bots, their metadata, and their specific automation capabilities and potential end points (if they support external inputs) needs to be available.
Bot communication	Bots need to communicate when working toward an overall automation goal. This communication must be asynchronous and, as much as possible, loosely coupled.
Bot negotiation	Value chain–wide automation goals at times require bots to negotiate and exchange information in order to make intelligent automation decisions based on the environment and process context.

Table 2 — Capabilities of a smart automation 2.0 reference architecture.

defined in adherence to the standard create a vocabulary that is known to all.

Events are published to a *channel of event mesh* to which bots subscribe based on their interests. When bots receive an event, they identify the next course of action; they check if a task is completed, identify the system where it is completed, and find the next task to be performed. This approach helps bots work without the need of a central component to facilitate the next step in the process flow.

While the channel of event mesh serves as the running medium for information-sharing on specific events of interest across the business process, the knowledge acquired by the bots within each domain, and other information about process context, is needed for process-level automation. Storing such knowledge in a data store is realized using data virtualization, combining multiple knowledge stores as a process-wide *knowledge graph* (a collection of nodes that help in traversing and making sense of the overall knowledge within the system). A knowledge graph could also contain state-specific-to-process instances, which can help in resolving conflicts across bots and with negotiation among bots, as well as directing cooperation among bots. A knowledge graph also serves as the store for analytics related to automation across the process, which will help with further improvements in automation. *ML models* that work along with the knowledge graph, while being specific to each bot, can be shared as well.

Architectural Style	Relevance to Smart Automation 2.0
Multi-agent systems (MAS) architecture	MAS architecture guides autonomous agents (e.g., bots) as they carry out automation in their respective domains. MAS architecture also helps bots work with their environment and cooperate with other bots. MAS architecture has traditionally been used, with success, in controlled environments and recently in domains such as autonomous vehicles and robotics.
Service-oriented architecture (SOA)	SOA plays an important role in how bots in a multi-agent system work together, whether by communicating and integrating with each other as peers or through an information bus. SOA helps in building a catalog of bots, which aids bots in discovering each other's capabilities.
Microservices architecture	Using MAS architecture in business processes creates certain constraints because of its homogeneity. Using it across domains requires that bots be modeled as independently deployed services with heterogeneous technology stacks. Microservices architecture greatly assists with autonomous deployment and heterogeneity. Combining MAS architecture with microservices architecture is a good way to balance the goals of autonomy, loose coupling, and heterogeneity.
Event-driven architecture (EDA)	EDA brings an integration capability that will help bots publish and communicate information as events in such a way that it is not necessary to know which other bots will be interested in the information. Events can be subscribed to as needed.
Data virtualization	While a knowledgebase for a bot often is not useful outside a given domain, such knowledgebases have an important role to play in process-level automation. Bringing together domain-specific knowledgebases into a single knowledge graph is one approach that can be considered, although it is not practical, as the impedance between domains makes centralization hard to implement. A data virtualization approach, in which a semantic layer can connect multiple knowledgebases to create a process-level knowledge graph, is pragmatic.

Table 3 — Architectural styles relevant to smart automation 2.0.

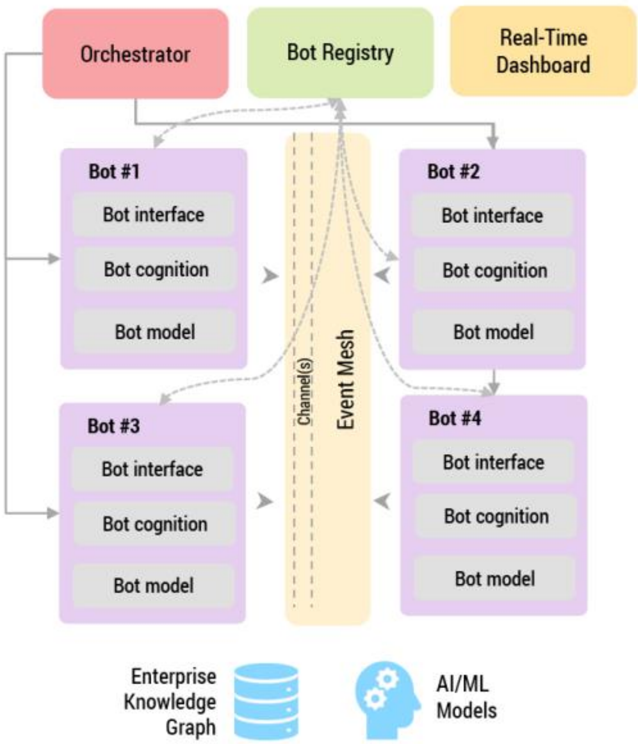

Figure 5 — Architectural building blocks for integration in smart automation 2.0.

An *orchestrator* in the architecture oversees the automation being executed by cooperative bots. It offers a plan for bots based on process goals and objectives. It also assists in real-time bot-to-bot communication by invoking a requested bot. Inter-bot communication may be required to enhance accuracy through specific feedback to a given bot. The orchestrator also analyzes the performance of bots and looks for areas of optimization by improving accuracy and/or reducing handoffs or instruction complexity.

Bot activity and events are visualized in real time using the *dashboard*. The dashboard not only shows visualizations of real-time information but can also be configured with rules to identify exceptions. It works with a preexisting alerting mechanism as required to trigger human intervention.

A Path to Smart Automation 2.0

Businesses with the intention of breaking automation silos to mature to smart automation 2.0 must take a more strategic approach to automation. Automation initiatives within domains also will have to consider broader automation goals. Key considerations that should go into planning for smart automation 2.0 include:

- **Platform-driven automation.** Integrating automation bots that cooperate across a value chain is best done through a platform. A platform with capabilities such as a bot registry, event mesh, and a knowledge graph can make its capabilities available to all the bots within a value chain. Such a platform should be designed to be open, so that bots built on heterogeneous technologies can participate and leverage the platform's services.

- **Value chain automation strategy.** A comprehensive automation strategy across the value chain will go a long way in delivering smart automation 2.0. An automation blueprint serves as a reference for bot

Architecture Block	Possible Technology Choices
Bots	Bots built and packaged as containers leveraging ML platforms, including those from cloud providers (e.g., AWS, Google, IBM Watson, or MS Azure)
Bot registry	Distributed data stores (e.g., etcd, Consul, Zookeeper)
Event middleware	NATS, Kafka, or event mesh technologies such as Solace
Event format	CloudEvents, AsyncAPI
Knowledge graph	Graph databases (e.g., Neo4j, Amazon Neptune, or document stores such as MongoDB)
Orchestrator	Lightweight workflow engines (e.g., Camunda BPM, Netflix Conductor, Zeebe)
Real-time dashboard	Elastic Kibana Dashboard
AI/ML tools	PyTorch, TensorFlow

Table 4 — Technology choices for architecture building blocks.

design. However, it is important to understand that enterprise value chains can be complex, making it difficult to devise a comprehensive automation strategy in one go. An evolutionary strategy is the best option.

- **Event architectures.** While the reference architecture incorporates event architecture, event architectures for bot communication are not easy to implement. Even though event architectures are becoming much easier to implement, consider synchronous communication where needed, such as for negotiations between bots and any real-time queries to knowledge graphs.

- **Business process management versus automation.** Most enterprises already have business process management initiatives running and have the necessary platforms to support them. Most often, process management efforts focus on process efficiencies beyond automation. It is important that automation initiatives lean into process management. While there are synergies, do not confuse automation initiatives with process management. Process management architectures are highly structured and rigid, while smart automation 2.0 initiatives should maintain autonomy and loose coupling between bots as key goals.

Smart automation 2.0 is a great intermediate step as we move toward complete autonomy, where all the processes within a value chain are intelligent and the lines between process management and automation have disappeared.

Aravind Ajad Yarra is a Wipro Fellow and Chief Architect at Wipro Digital. His areas of focus are emerging technologies and digital architectures. With over 24 years' experience in the technology services industry, Mr. Yarra helps enterprises adopt emerging technologies to build smart applications leveraging artificial intelligence, machine learning, and cloud computing. His current areas of research include quantum computing and decentralized applications. Previously, Mr. Yarra worked as a Solution Architect for several complex, transformational programs across the banking, capital markets, insurance, and telecom industries. He can be reached at aravind.ajad@wipro.com.

Danesh Zaki is a Senior Distinguished Member of the technical staff at Wipro Ltd. He has over 22 years' experience in consulting, architecture, implementation, and development. Mr. Zaki's core expertise is on enterprise integration covering API management, open source middleware, integration platforms as a service (iPaaS) and service-oriented architecture. He has published white papers and has been a speaker on integration and architecture. He can be reached at danesh.zaki@wipro.com.

Superimposing Natural Intelligence on Artificial Intelligence: Optimizing Value

by Tad Gonsalves and Bhuvan Unhelkar

> Time you enjoy wasting is not wasted time.
>
> — Attributed to T.S. Eliot (among others)

Automation, with the help of machine learning (ML), aligns processes and technologies and offers mechanisms to *keep people out of the loop*, facilitating autonomous operations. People, however, base their decisions on *values*, which are subjective in nature. For instance, customer value depends on subjective factors that are difficult to parameterize in an ML algorithm. So ML, by its very nature, cannot fulfill the total value proposition for a customer. Quality business decisions are arrived at *when humans are kept in the loop*.

Rather than favoring the optimization of business decision making through pure artificial intelligence (AI) automation, we present a model that benefits businesses through collaboration between humans and AI. We argue that judiciously superimposing human natural intelligence (NI) on AI is a better way to enhance customer value than using AI alone to arrive at business decisions.

Optimizing the Business

AI, on its own, automates business decision making and business operations, but is not always successful in accounting for and improving customer value. Customer value is a subjective entity based on myriad and ever-changing personal and environmental factors. "Satisfaction," "quality," and "joy" denote factors that provide customer value, but an ML algorithm cannot easily parameterize these factors. Customers' needs, moods, and context dynamically change. What might be insignificant for a machine, for example, may be of immense value to a person. Consider the thrill and excitement of watching a tennis match between two tennis stars; they produce occasional brilliant shots as well as unforced errors and mistakes. Spectators derive immense value from the aesthetics of the game, which include those unforced errors. The value parameters (e.g., joy and satisfaction) justify high ticket prices, as compared to a tennis match between robots; the robots play perfectly — with no unforced errors — but the result is a boring match with little value for spectators.

AI/ML conducts analytics over a vast suite of data and establishes correlations between data points. However, AI/ML, limited to this "black box" of correlations, cannot explain causation. Machines — good at plowing through data (big and small) beyond human capabilities — sift through a massive collection of data, including a huge number of possibly relevant variables, and quantify the variables' relevance to predicting a target. Humans, on the other hand, can identify small sets of relevant aspects of an issue and enquire into their cause. Explainable AI[1] is an attempt to understand the aforementioned black box, but that explanation usually occurs after the event, is about the correlations, and does not provide any understanding of why the event occurred or why the machine made certain predictions. Causation (dealing with the "why" of an event), subjectivity in customer value, and the corresponding business decisions — all important ingredients in an optimized business — remain within the purview of humans.

Thus, we propose using NI in combination with AI in the automation and optimization of business. NI usage was prevalent throughout the industrial and information eras, during which machine-based computations were used as a tool in human decision making. Here, we explore the opportunity to incorporate human decisions and human feedback in an AI-based decision engine. NI provides the appropriate subjectivity to enable value-producing business decisions. The key points to consider while optimizing a business based on AI capabilities are as follows:

- NI has the potential to provide judicious support to AI and the ability to make good use of AI in decision making.
- NI and AI need to be treated as complementary aspects of business decision making; AI must not entirely replace NI.

- AI must be continuously improved as it learns from and simulates decisions made by humans, but this improvement should not aim to take over all human-made decisions.

- Evaluating the consequences of a decision and investigating the "why" of an event is a NI function that need not be automated.

- Decision engines must facilitate the incorporation of feedback about the consequences of a decision in order to enhance customer value.

- Optimizing business performance is a distant goal that the intermediate goal of engineering collaboration between intelligent humans and intelligent machines supports.

The question we seek to address is: how can an advanced automation system and humans interact and work together in an intelligent way to capitalize on the strengths of each and maximize the value proposition for the customer? We now move on to our arguments for superimposing NI on AI to provide authentic and holistic customer value.

Relevance of AI/ML in Automation

ML algorithms in supervised, unsupervised, and reinforced formats boost decision making, since they handle, in a short time frame, the myriad complexities that the human mind cannot easily comprehend. A massive increase in GPU-led computing power and the exponential growth in data available for training, testing, and analysis have produced decisions that are more accurate and faster than those that humans are capable of making. Computer vision, natural language processing, and games excel in performance because ML models enable excellence in automation. AI transforms businesses and enterprises to an extent beyond that made possible by traditional industrial automation and advanced robots.

ML models ingest relevant historical data and undertake analytics on it. Decision trees, random forests, support vector machines, regression, and neural networks are some of the readily available models for performing ML.[2] Data is fed to the ML model in several batches and the learning algorithm is executed until the desired prediction accuracy is achieved at the learning output. Training a learning model using historical data results in the learned model. The learned model extracts patterns found in the historical data while learning and is ready to make predictions when new and fresh data is fed into it. The detailed steps involved in ML are automated in order to form a single pipeline from data collection to prediction.

The manufacturing and construction sectors are more automated than, say, the health and sports sectors because manufacturing and construction processes are more mature and well defined. Processes in manufacturing and construction, unlike those in health and sports, are primarily linear in nature. This means they have less need for feedback and adjustments, making them ideal for automation. Incorporating agility into processes, however, brings in iterations and increments that require the handling of uncertainty and change.[3] ML algorithms can mimic human behavior fairly well, as long as that behavior is relatively stable. When human behavior changes because the context and values have changed, as with agility, then such behavior is difficult to mimic. Therefore, it is imperative that AI be combined with NI, which accommodates change.

Limitations of AI/ML in Practice

AI systems exhibit better performance than humans do when it comes to data crunching at high speed. Nonetheless, these AI systems still have significant limitations. AI models can only be as good as the data fed to them and as well as the algorithms are coded. Biases in models can crop up based on skewed observations, wrongly recorded data, and developer viewpoint. AI learning models can potentially (1) develop biases from the data on which they are trained and/or (2) incorporate the biases into their learning algorithms, at the same time that they (3) lack the inherent ability to make ethical decisions or understand shifting contexts, which are subjective and continuously changing. Next, we explore these limitations.

Biases in Data Used for AI Learning

While data is usually considered objective, a collection of data can still be biased. Data is a record of observations that involve human decisions; it incorporates the beliefs, purposes, biases, and pragmatics of those who designed the data collection systems. Data is not a singular record but a collection of many records of observations. Therefore, the potential exists that the beliefs of the observers have colored the meaning of the data.[4] Sample bias, prejudicial bias, exclusion bias, measurement bias, noise bias, and accidental bias are all examples of data-specific biases.[5] These biases influence the models built upon the data.

Biases in the Learning Algorithms

Humans use reasoning to make decisions, and they have rationales for their decisions. They can explain the causal relationships between the data at hand (input) and the decision they have arrived at (output). ML, on the other hand, does not proceed by deciphering the cause-effect relationship. Rather, it learns a function that maps inputs to outputs. This nonlinear function, however, is not in a closed analytic form but instead lies distributed in millions of parameters scattered throughout the learning network. Programmers and network designers cannot pinpoint which data patterns or features the learning network extracted to produce its final prediction. Thus, the learning model is a black box.[6]

> *People, processes, and technology make up business. ML attempts to automate business processes by using technology. Automation aligns processes and technology while keeping people out of the loop.*

The opaqueness inherent in ML architectures severely limits their applications and reduces the overall trust people put in such machine-made predictions and decisions. Biases in algorithms are a result of the developer's viewpoint and the system's ability to learn from previous decisions (i.e., learnability). Structure bias (i.e., number of layers, neuronal units in each layer and their synapses), hyperparameters bias (i.e., millions of hyperparameters controlling the learning mechanism), and train-test bias (i.e., the training data set and testing data set used to solve a certain problem may follow totally different distributions) are examples of algorithm-specific biases.

AI models by their very nature are based on achieving a certain level of performance. Users and business leaders all aim to maximize their performance based on the performance of the AI systems. The ML code embedded in AI models is agnostic to a user's specific situation. AI systems are data-driven (correlations); human systems are knowledge-driven (causation). AI systems are performance-driven, not value-driven, but business customers are looking for value along with performance. ML code cannot produce value on its own.

Inability of AI Models to Make Decisions of Subjective Value

Full automation with AI, from data collection to decision making, implies that humans are out of the loop and machines have full control. But accurate predictions do not necessarily mean ethically correct decisions, and machines cannot be trained to respond to ethical and moral questions when arriving at decisions because ethical and moral values are inherently subjective and not codable. Machines are not capable of self-evaluating the consequences of their decisions. Without human involvement in the decision-making processes, greater ethical and moral challenges involving the use of AI/ML are likely.

Automation, Optimization & Subjectivity in AI-Based Decision Making

People, processes, and technology make up business. ML attempts to automate business processes by using technology. Automation aligns processes and technology while keeping people out of the loop. In an optimized business, the people-process-technology elements are aligned with each other on an ongoing basis with the aim of maximizing customer value. When these elements are not aligned, slack, wastage, and errors occur. While automation is the direct application of AI to business processes, optimization, in practice, also requires natural (human) intelligence (NI). As shown in Figure 1, NI handles the subjectivity involved in providing customer value; that subjectivity does not lend itself to being coded as an ML algorithm.

ML generates increasingly accurate business predictions based on the patterns and relations inherent in large data sets, thereby creating a learned model. Decision makers then make their decisions based on the machine-learned predictions. While the process of machine-learned prediction can be automated, the actual decision making still requires human input. Technology automation continues to optimize business processes without fully comprehending shifting contexts. Automation based on correlations (i.e., automated ML) can provide many results but can neither explain those results nor ascertain and make use of the subjectivity inherent in customer value.

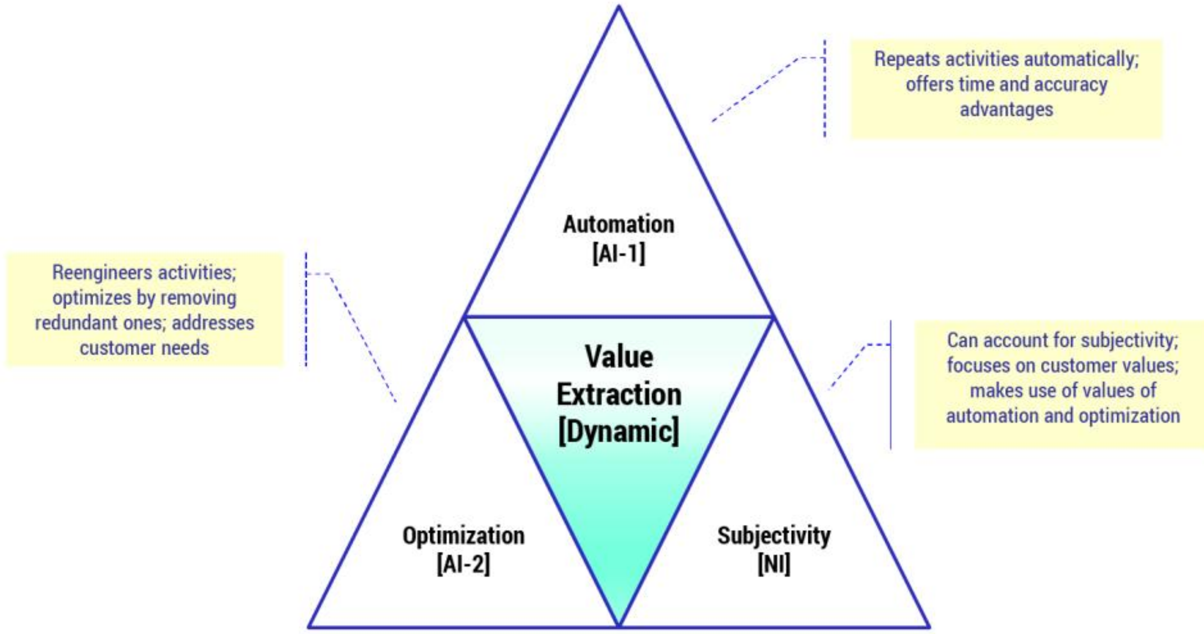

Figure 1 — Automation, optimization & subjectivity in "value extraction."

The ideal decision-making process is a judicious combination of NI and AI. Thus, as we asserted in our opening, quality business decisions are arrived at *when humans are kept in the loop*. Rather than optimizing business decision making through pure AI automation, a model that benefits businesses through optimizing collaboration between humans (NI) and AI is most appropriate.

Intelligent Automation: Learning-Correction-Relearning Cycle

Implementing collaboration between NI and AI yields intelligent automation. Figure 2 shows the AI pipeline, with four phases: data collection, ML, prediction, and decision making. The first three phases are relatively easy to automate based on current AI technologies; the fourth phase cannot (yet) be successfully automated. The limitations of AI, as we have discussed, need to be handled by superimposing NI through the design, development, and implementation of the solution. The following outlines the role of NI in each phase:

1. **Data collection** — choosing the right kind of data for a given ML problem and filtering the varied types of possible biases from the data

2. **ML** — allocating the right kind of ML algorithm

3. **Prediction —** opening the ML black box to explain causal relationships among inputs and prediction

4. **Decision making** — fully engaging in decision making

Normally, decisions are based on the predictions obtained from the automated ML modules.[7] With the ever-increasing cutting-edge performance demonstrated by recent ML, the final decision-making process is at risk of being entrusted entirely to full automation. Fully automated decisions — leaving humans completely out of the loop — may not provide the necessary customer value due to the subjective interpretation and changing nature of that value. At times, such fully automated decisions may even be detrimental to business goals and to society in general due to ethical challenges.

Quality decisions, which are also ethical decisions, include humans in the decision-making loop. Humans are capable of considering the consequences of decisions vis-à-vis their quality and ethical ramifications. NI provides invaluable insights, after inspecting the consequences of decisions, by considering ethics and values. These NI-based insights are superimposed on the learning algorithm (as shown in Figure 2). The feedback loop illustrated in Figure 2 then tweaks the historical data, learning model, and new data to filter possible sources of error and bias and retrains the model. The learning-correction-relearning cycle is repeated multiple times to enable the system to continue to learn and improve its performance. Eventually, after multiple iterations, the model shown in Figure 2 arrives at ethically sound decisions that produce adequate customer value. The caveat to keep in mind is that in earlier iterations of this model, NI makes the

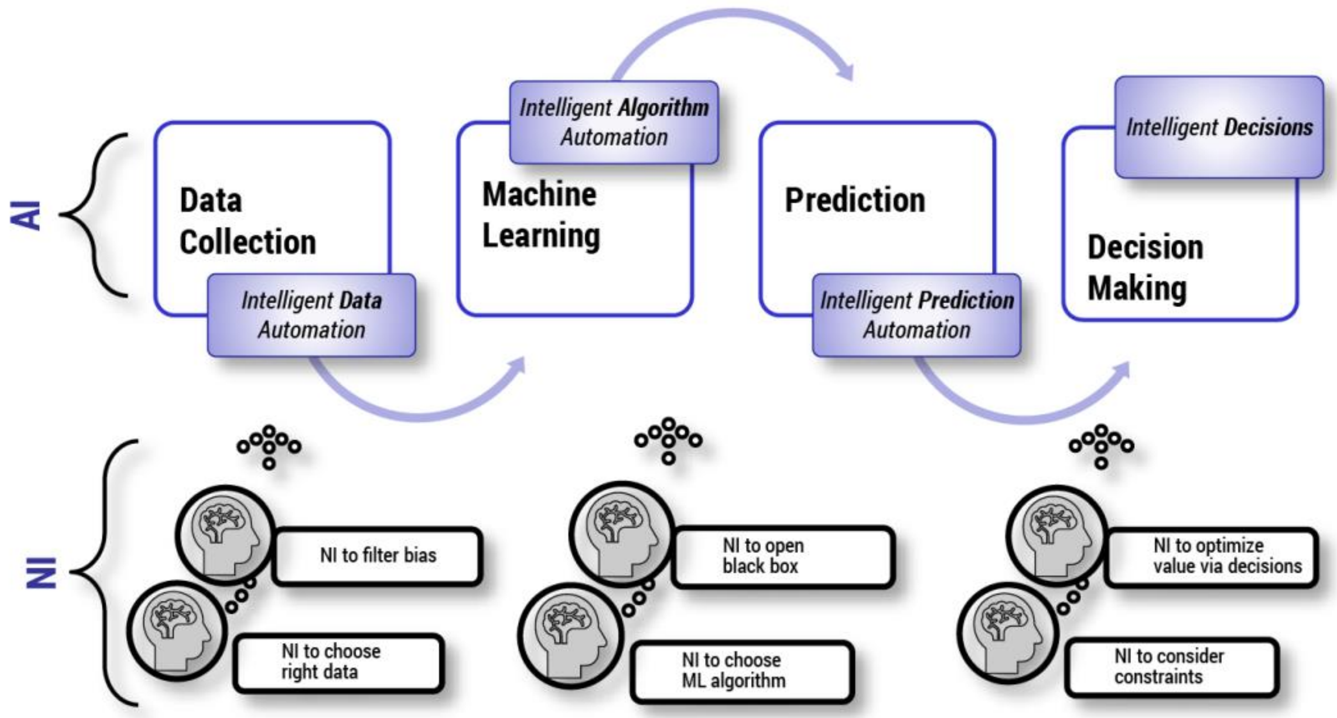

Figure 2 — Superimposing NI on AI for intelligent automation.

actual decision, whereas in later iterations, AI learns from NI and stores those insights.

Optimizing "Subjectives" Under Subjective Constraints

Optimization can be defined as minimizing or maximizing an objective under certain constraints, where the objective to be minimized or maximized is cast in the form of an objective function of the decision variables.[8] An optimization problem is defined as follows:[9]

Find X,
such that f(X) is minimum/maximum
subject to the constraints:
$g(X) = 0$ & $h(X) < 0$

where X is a vector of decision variables,
usually bounded as:
$X_{min} \leq X \leq X_{max}$

f(X) is the objective function, g(X) and h(X) are the constraints, and
X_{min} and X_{max} are the bounds on the decision variables in vector X.

The decision variables, the objective function, and the constraints are all quantified, making them "objective" (as opposed to unquantified concepts, which are "subjective"). The decision variables serve as inputs to many AI optimization algorithms, such as evolutionary computation[10] or swarm intelligence,[11] which optimize even large-scale or NP-hard[12] (nondeterministic polynomial time) problems in a reasonable amount of time.

For example, let's say the manager of a software development company plans to develop a software package in the minimum possible time. She hires teams of analysts, designers, and programmers. The constraints are that she can hire not more than five members for each team and cannot exceed the hiring and development cost decided upon by the company. In this optimization problem, the software development time is the *objective function*, development cost is the *constraint*, and the number of professionals to hire in each team is the *decision variable*. The important thing to note here is that the objective function, constraints, and the decision variable are all quantified because any AI optimization algorithm will soon yield the optimum solution without any human intervention.

Now let's review some examples of nonquantifiable problems. Consider, for instance, a qualitatively different optimization problem from the one above. Let's say a couple is out on a date at a very fine dining establishment. Imagine the restaurant's kitchen and serving processes are highly optimized. Orders for food and drinks are taken by a robot and promptly delivered via conveyor belt, all at the ring of a bell from the couple's table. The situation is fully AI-automated and optimized. But would dining in such an environment be an enjoyable and romantic experience for the couple? Perhaps as a novelty for a technocratic-oriented couple — but unlikely. The tennis match described previously is another AI-automated and optimized experience, but one not likely to be appreciated by its spectators.

In these subjective examples, humans do not cherish the experience because they do not derive any value (joy, aesthetic satisfaction) from it. Since customer value is not quantifiable, it cannot be objective. Value is a subjective phenomenon based on customers' emotions, judgments, impressions, and opinions. Therefore, we propose to optimize not an objective problem, but a subjective (i.e., value) one by making decisions based on our human intelligence (NI), which comprises emotions, intuition, experience, knowledge, and expertise. The NI decision-making phase follows the ML phases shown in Figure 2.

The maximization of customer value is evaluated from the following four points of view:

1. **Right decision.** This is a decision taken at the right time, in the right context.
2. **Tradeoff.** This occurs between enterprise profits and societal well-being.
3. **Ethics.** A perfectly good "objective" decision may not be an ethical one.
4. **Legality.** A perfectly good "objective" decision may not be a legal one.

These four aspects of evaluating a decision are also "subjective" in the sense that they cannot be quantified. They are the subjective constraints against which the subjective customer value is to be optimized. NI, with its experience, intuition, knowledge, and expertise, can judge the value of the decision, its potential bias, and the legal and ethical impacts of the decision on society. Finally, let's examine three use cases to briefly describe the impact of NI on AI models.

Insurance

Even if we overcome all possible biases and a prediction is accurate, that prediction can still pose ethical and moral challenges. For example, a prediction may be accurate in indicating the possibility of heart disease in a certain cross-section of society based on age, gender, ethnicity, and financial well-being. Acting on the prediction can lead to differences in insurance coverage amounting to discrimination. NI must override the "logical" action resulting from the prediction to ensure equality in insurance coverage.

Crime Spotting

Crimes may be accurately predicted through unbiased data and unbiased algorithms. Preventive and preemptive action by the regulatory authority may lead to discriminatory and unfair actions even before a crime has been committed. What would be viewed as unfair justification of preemptive action is an ethical and moral challenge that fully automated systems are not capable of handling. Better decisions are taken when NI is superimposed on ML predictions.

Healthcare

Healthcare is one sector in which personalization is urgently needed. Most patients trying to come to terms with a serious illness expect emotional support from their doctors. However, in the current medical setup, most doctors, nurses, and other healthcare providers are overworked and cannot take additional time to deal with every patient on an individual level. Providers can train robots to take over routine functions (e.g., delivery of medications by bedside or movement of patients in a wheelchair), making it possible for medical staff to aptly provide the human empathy and caring that patients need. Human-robot collaboration is an important part of superimposing NI on AI;[13] in this case because AI is used to free up NI for important "subjective" functions.

Conclusion

Computers are number-crunching machines. Humans cannot — and, perhaps, *need* not — compete with computers when it comes to handling big data volume and velocity. Machines are good at data-driven predictions that are difficult for humans to make. Humans are better at handling smaller sets of data, inspecting the relevant aspects (contexts), and arriving at decisions appropriate to specific situations. As opposed to computers, humans are fuzzy, knowledge-driven analyzers. Their NI (human or natural intelligence) includes not only knowledge and experience, but also intuition and insights. Values and their dynamics are beyond the grasp of algorithms. Perhaps machines may never be fully learned, not because of a shortcoming in ML systems but because values in decision making are inherently unlearnable.

This unlearnable aspect poses a challenge to automation and to the optimization of ML. AI models are extremely complex because they incorporate a wide variety of data with high volume, velocity, and variety. Added to the complexity of the data is the complexity of ML algorithms. Finally, AI models are not singular, stand-alone systems, but rather a combination of multiple, collaborative systems typically interacting with each other in the cloud. Thus, when the systems are totally automated and are, in effect, black boxes, any break-down in a system may result in total chaos across all business functions. The results can be uncontrolled and dangerous, as it is typically very difficult to determine the source of the fault and fix it.

Therefore, for a business decision-making or analytic task, human experience and an intuition-based, fuzzy way of thinking should complement the data-driven techniques of machines. In this article, we have argued for the need to superimpose NI on the predictions made by AI and have examined the likely risks without such superimposition as well as the advantages resulting from it.

References

[1]Goebel Randy, et al. "Explainable AI: The New 42?" Proceedings of the *2nd International Cross-Domain Conference for Machine Learning and Knowledge Extraction* (*CD-MAKE*), Hamburg, Germany, 27-30 August 2018.

[2]Finlay, Steven. *Artificial Intelligence and Machine Learning for Business: A No-Nonsense Guide to Data Driven Technologies*. Relativistic, 2018.

[3]Unhelkar, Bhuvan. *Big Data Strategies for Agile Business*. Auerbach Publications, 2018.

[4]Provost, Foster, and Tom Fawcett. *Data Science for Business: What You Need to Know About Data Mining and Data-Analytic Thinking*. O'Reilly Media, 2013.

[5]Jones, M. Tim. "Machine Learning and Bias: Look at the Impact of Bias and Explore Ways of Eliminating Bias from Machine Learning Models." *IBM Developer*, 27 August 2019.

[6]Doran, Derek, Sarah Schulz, and Tarek R. Besold. "What Does Explainable AI Really Mean? A New Conceptualization of Perspectives." arXiv e-print, Cornell University, October 2017.

[7]Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press, 2018.

[8]Gonsalves, Tad. *Artificial Intelligence: A Non-Technical Introduction*. Sophia University Press, 2017.

[9]Gonsalves (see 8).

[10]Simon, Dan. *Evolutionary Optimization Algorithms*. Wiley, 2013.

[11]Kennedy, James, Russell C. Eberhart, and Yuhui Shi. *Swarm Intelligence (The Morgan Kaufmann Series in Evolutionary Computation)*. Morgan Kaufmann, 2001.

[12]"NP-hard." National Institute of Standards and Technology (NIST), 2020.

[13]Unhelkar, Vaibhav V., Shen Li, and Julie A. Shah. "Semi-Supervised Learning of Decision-Making Models for Human-Robot Collaboration." Proceedings of the *Conference on Robot Learning* (*CoRL*), Osaka, Japan, 30 October-01 November 2019.

Tad Gonsalves is a Professor in the Department of Information & Communication Sciences at Sophia University, Japan. His research interests include computational intelligence and ML algorithms as well as bio-inspired optimization techniques for engineering and business applications. Dr. Gonsalves is also actively engaged in designing image recognition and low-cost autonomous driving systems. As an educator, his passion is teaching programing to students without a scientific background. Dr. Gonsalves is the author of Artificial Intelligence: A Non-Technical Introduction, *an attempt to introduce the latest developments in AI to the general public. He holds a bachelor of science degree in theoretical physics, a master's of science degree in astrophysics, and a PhD in information science and systems engineering. He can be reached at t-gonsal@sophia.ac.jp.*

Bhuvan Unhelkar (BE, MDBA, MSc, PhD; FACS) is a Senior Consultant with Cutter Consortium's Business Technology & Digital Transformation Strategies practice. He has more than two decades of strategic as well as hands-on professional experience in the information and communications technologies (ICT) industry. Dr. Unhelkar is a full Professor at the University of South Florida, Sarasota-Manatee. As a founder of MethodScience and PLATiFi, he has demonstrated consulting and training expertise in big data (strategies), business analysis (use cases, BPMN), software engineering (object modeling, Agile processes, and quality), collaborative Web services, green IT, and mobile business. His domain experience includes banking, financial, insurance, government, and telecommunications. Dr. Unhelkar has authored/edited 20 books in the areas of collaborative business, globalization, mobile business, software quality, business analysis/processes, the UML, and green ICT, and he has extensively presented and published papers and case studies. He has designed multiple industrial and master's degree courses in such areas as global information systems, Agile method engineering, object-oriented analysis and design, and business process reengineering and new technology alignment, which he has delivered to universities in Australia, US, China, and India. Dr. Unhelkar holds a Certificate-IV in TAA and TAE and is a Certified Business Analysis Professional (CBAP of the IIBA). He is an engaging and sought-after speaker; a Fellow of the Australian Computer Society (for distinguished contribution to the field of ICT); a life member of the Computer Society of India; President of Rotary Club of Sarasota Sunrise (Florida) — Multiple Paul Harris Fellow, AG; a Discovery volunteer at NSW parks and wildlife; a member of the Society for Design and Process Science; and a former TiE Mentor. Dr. Unhelkar earned his PhD in the area of object orientation from the University of Technology, Sydney. He can be reached at bunhelkar@cutter.com.

IT'S THE RIGHT THING TO DO

Governing Intelligent Automation

by Daniel J. Power, Ciara Heavin, and Shashidhar Kaparthi

For more than 70 years, automation has been expanding in scope and capabilities. Today, intelligent automation (IA) is being quickly deployed in both the product manufacturing and service sectors of developed economies. Current crisis conditions may lead to even more rapid adoption and greater acceptance of smart machines. In crisis situations, companies and organizations are more likely to rush development and adoption of innovative, "magic bullet" technology such as IA. The rush to market and the rush to adopt IA may, however, lead to increased risks to civil liberties, health, and safety; discrimination and fairness issues; overreliance on artificial intelligence (AI) applications; limited or ineffective human decision-maker oversight; and other negative consequences.

Society is in the midst of the fourth wave of automation, or what some call the "Fourth Industrial Revolution." Over the years, a number of related terms have been used to describe and explain the automation of tasks, including cybernetics, robotics, and adaptive control. The term "automation" has traditionally been associated with manufacturing; specifically, the application of technology, as part of the manufacturing process, to routinize repetitive tasks, with people replaced by machines in the execution of tasks, processes, and activities. Today's IA technologies and applications are extremely complex software systems that are deployed to perform a wide variety of tasks, and the systems should be monitored and regulated by both internal and external stakeholders.

For many years, formal and informal governance mechanisms dealt with the automation issues of data collection and use, intellectual property (IP), labor displacement, product liability, and warranties. Governance was ad hoc and involved unsystematic and partial solutions. Moving forward, governance of automation should be more proactive and coordinated with thoughtful policies, structures, and roles. Improving the governance of process automation, especially of systems incorporating data that is assessed using AI technologies, should be a proactive management goal.

Information technologies — for example, machine learning (ML), motion sensors, optimization algorithms, and robotic vision — are disrupting many traditional manufacturing, service, and retail industries. IA can be used in a variety of settings, from the tangible — such as factories with assembly lines, continuous process production lines (e.g., an oil refinery), and small-batch job shops — to the less tangible (e.g., computer-based processes such as call centers with voice bot advice, tax advisor and accounting services, and automated insurance claims processing). Implementing more sophisticated IA, especially for advice, fraud detection, security monitoring, autonomous factories/manufacturing, and automated warehousing, among other applications, can increase the resiliency, cost effectiveness, customer satisfaction, and scalability of many productive endeavors.

This article explores the need for IA and considers possible policies for governing smarter factories, processes, and supply chains. In the following section, we examine the meaning of IA. We then investigate the need for governing automated processes, especially IA using AI. Based on our analysis, we propose a governance framework and eight major policy prescriptions associated with governing IA.

> *Society is in the midst of the fourth wave of automation, or what some call the "Fourth Industrial Revolution."*

Defining Intelligent Automation

The convergence of AI and operations innovation creates a paradigm shift. Smart automation is a vision for how organizations can and should enhance productivity. IA is a technology frontier. IA is intended to assist human workers by taking over repetitive, routine, and manual tasks. Potentially, IA can process most invoices, reconcile accounts, conduct background reviews, pick products in a warehouse, build cars and trucks, and perform many other tasks. IA includes decision automation using AI with sensors to control and manage complete manufacturing systems and production processes. It also includes specific products such as vision systems, production robots, robotic floor scrubbers, and chat and voice bots that provide advice. Some IA systems go beyond applying rules and procedures to include ML and cognitive technologies. IA systems are becoming more mature and are increasing in capabilities and sophistication. Let's briefly examine three companies on the leading edge of IA — Amazon, Tesla, and Progressive Casualty Insurance Company:

1. **Amazon** has more than 100,000 robots in its warehouses. In addition, the company has installed machines in several of its US warehouses that scan and box items to be sent to customers. At Amazon, robotic automation has taken over certain duties, such as carrying pods of inventory and transporting pallets through buildings.[1] Using robots and people in the same environment can cause problems, so Amazon is providing warehouse workers with utility belts that signal robots to avoid fatal encounters and increase employee safety.[2]

2. **Tesla's** factory in Fremont, California, USA, is one of the world's most advanced automotive plants. The Tesla assembly line is built with a combination of AI software and automation. The manufacturing process for the Tesla Model S uses more than 160 specialist robots.[3] IA systems in the Tesla factory sense and produce very large amounts of data used to automate entire processes, make decisions, and guide robots.

> *IA applications are reaching the point where both organizations and governments must create policies regulating the use of and the liability associated with using smart technologies.*

3. **Progressive**, a technology and automation innovator, uses a mobile app (Snapshot) for an optional insurance discount program.[4] Usage-based insurance involves collecting data directly from a person's car and then using algorithms to infer patterns in behavior and adjust insurance rates. Snapshot creates a personalized insurance rate based on a driver's actual habits. The Progressive app has many features, including a claims center for filing and tracking claims.

IA is transformative and strategically and tactically important for organizations. It is likely to change how organizations operate and how people work. IA or "intelligent process automation" (IPA) integrates AI, ML, natural language processing, and automation to redesign business processes. Robotic process automation (RPA) encompasses many technologies, including workflow orchestration, mobile data capture, sensors, analytics, AI bots, more traditional robots, and e-signature technologies.

The possibilities for automation of both tangible goods and delivery of less tangible services have increased greatly in recent years. Now, most companies have some automation, ranging from supporting isolated production tasks or semi-independent work processes to more comprehensive systems.

Need for Governance

Change is happening so quickly that some stakeholders have real concerns that adopting IA, RPA, and IPA will result in job losses, problems with invasive bots, loss of privacy, and poor ROI because of technology obsolescence. These concerns help justify the need for better governance mechanisms.

IA applications are reaching the point where both organizations and governments must create policies regulating the use of and the liability associated with using smart technologies. Many outstanding issues exist; for example, who is liable if IA applications make incorrect decisions or recommendations? How and when should IA applications be tested and validated? Should an organization disclose that a customer is interacting with an AI-enhanced automated process?

In 2018, Google announced that its new virtual AI assistant, Duplex, could make phone calls for you using a human-like voice that sounded like a real person.[5] People immediately objected to Duplex, concerned that a bot could be confused with a human. Well-thought-out governance policies could have avoided this implementation mistake. Instead, Google had to react to the negative feedback and changed Duplex's functioning so that it would identify itself as a robot when making calls. Government regulations will likely need to address this disclosure issue.

Also in 2018, the Singapore government's Infocomm Media Development Authority (IMDA) announced the establishment of an "Advisory Council on the Ethical Use of AI and Data" that will assist the government in developing ethics standards and reference governance frameworks and that will publish advisory guidelines, practical guidance, and/or codes of practice for voluntary adoption by industry.[6] Singapore provides an example that other governments should follow. The current council discussion paper recommends two key principles that are especially relevant for IA governance:[7]

1. "Decisions made by or with the assistance of AI should be explainable, transparent, and fair to consumers."

2. "AI systems, robots, and decisions should be human-centric."

A potential problem and ethical issue with the adoption and use of AI and intelligent systems is that ML is often a "black box" with seemingly unexplainable results. Examination of the data used for ML and testing can identify problems and correct them or disclose them. Testing and transparency encourage confidence and trust in a machine-generated algorithm.

AI and IA are expanding in use, but technology companies are often secretive about how the software makes decisions. IA systems may be poorly designed and, as in the case of Duplex, need redesign. There may be inadvertent or malicious coding errors, data collection errors, or process changes. Application functionality may need to change rapidly. Government regulation and company policies should address the need for transparency and validation of AI, ML, and algorithms. A trusted third party may be needed to certify some or all decision automation software.

Security issues also create a need for governance mechanisms and policies. Reasonable, effective governance of IA and advanced AI applications is both necessary and possible.

Governance Framework and Policy Prescriptions

IA governance includes (1) governance of the data collected and used in IA and (2) governance of the decision automation software and tools incorporated in systems. The entire automated system — with components including robots, automated handling, workflow software, and people — requires governance.

A governance framework should specify who has the responsibility and accountability for making specific categories of decisions (decision rights) and establish an accountability framework to ensure appropriate behavior. A governance framework should specify and evolve processes, roles, policies, standards, and metrics to ensure the appropriate and effective use of IA. Managers and regulators must use current knowledge and anticipate change. Given the novelty of IA, governance will need to evolve over time as more is learned. Because IA governance is complex, managers must do more than establish policies, they must solicit feedback from all stakeholders, including customers and suppliers. Managers must also have mechanisms for feedback and for monitoring policies and performance. Governance stakeholders include the following:

- Adopters/owners — senior management, shareholders

- IA vendors — creators/producers, distributors

- Government and regulatory bodies

- Industry competitors

- Suppliers and customers of products and services

- Citizenry and community

Figure 1 tries to capture the broad IA governance domain, including stakeholders' interests and responsibilities and the interactions among stakeholders.

The central actor in governance is the adopter/owner, but government at all levels has a role to play, especially in relationship to IP and liability laws. The testing and monitoring of problems by IA vendors — the creators and providers of technology solutions — is especially important.

Because the adopter/owners of automated factories or automated workflow processes are at the center of the governance ecosystem, their governance responsibility is greater. Adopters must keep records, have proactive policies, and ensure that managers make ethical governance choices. A central component of governance is thoughtful, well-defined policies.

Governing a technology or process is about control and risk minimization. Suggesting generalized policies for the wide variety of IA systems that exist and that have been suggested is a daunting task. We tried to apply a number of lenses to the task, including examining ethical concerns, examining traditional topics in IT governance, and using a stakeholder framework to examine governance from a broader context. In an adopter organization, the right amount of governance depends upon factors such as risk tolerance, the types

Figure 1 — Governing intelligent automation.

of IA in use and contemplated for development, and the maturity of the current IA systems. We tried to anticipate what could go wrong with IA and how policies could reduce those risks. The following suggestions provide a starting point for defining IA governance policies. The suggestions are categorized as "public" or "company" policies.

Government with Appropriate Authority and Jurisdiction

Public Policy 1

The owner of an IA application cannot limit her/his/its liability related to the use of the IA in any way, even with a disclaimer limiting liability. Liability is a legislative issue, and current laws should be reviewed. Public policy should resolve three fundamental questions: (1) Who is responsible for any harm arising from the use of IA? (2) What does it mean to own a smart machine? and (3) Who owns the IP?

Public Policy 2

IA is especially useful when the technology replaces routine, repetitive, low-value jobs. Organizations should be incentivized to upskill employees that IA displaces and to prepare them for knowledge-intensive, data-driven work that requires empathy and other complex human characteristics that are difficult for IA to provide. In other situations, IA applications that replace human employees should be discouraged, except when the task to be automated is dangerous or creates other potential harm to a human. If smart automation is an appropriate way to engineer jobs and reduce costs, then taxation may be necessary to compensate for any social disutility.

Public Policy 3

Developers should provide complete disclosures about IA algorithms and decision automation software that include: (1) evidence that an algorithm does not create bias or discrimination in the decisions that will be made using it; (2) identification of the frequency of errors in classification when the algorithms were developed; and (3) acknowledgment of any possibility that IA interactions could be mistaken for human interaction (e.g., when using chatbots).

Adopters, Suppliers, and Vendors

Company Policy 1

In assessing IA opportunities, management must collaborate with stakeholders to assess the impact on the workforce. Adopters should commit to funding the upskilling of displaced employees to assume higher-value roles. Anxiety continues to grow around the possibility that smart robots will take over the jobs of people. In general, it is important to provide training so workers can adapt to new, more complex, and technologically advanced operating environments.

Company Policy 2

IA applications should only make autonomous decisions in routine, recurring, well-understood decision situations. Even in those situations, knowledgeable humans should regularly monitor the decisions and their consequences to ensure the application is performing satisfactorily. It is especially important to limit the domain of discretion for IA applications in specialized settings such as hospitals.

Company Policy 3

There should be a commitment to transparency and disclosure. An adopter should disclose when an IA application is making a decision, and the reasoning behind the decision should be understood and transparent to anyone impacted by the decision.

Company Policy 4

Each IA application and system should be tested to ensure it is operating as intended. IA decisions should be logged. Customers and owners of an IA application must understand why a machine model took a specific action, reached a specific conclusion, or made a specific recommendation. A person with governance authority should monitor logs and be able to determine what information was analyzed to reach a conclusion and be able to selectively review that information. An expert should be convinced that the data used is unbiased, and testing should confirm at a satisfactory confidence level that the choices made by an IA application are appropriate.

Company Policy 5

IA application code should be restricted and secured. IA adopters should have copies of application code. Security is especially important to prevent IP theft of IA applications and AI algorithms. IA governance must include the traditional data governance issues, including backup, recovery, maintenance of change logs, security, and data privacy.

> ***Industry associations, vendors, and other public sources are likely to share and disseminate best practices.***

Many other issues and concerns will likely need to be addressed, including: (1) an IA usage policy that clearly articulates the business processes and data that management considers appropriate for automation; (2) a policy that identifies who is authorized to procure IA products and services; (3) the required use of business process management (BPM) to analyze, improve, optimize, and automate business processes prior to implementing IA; and (4) standards for testing, installation, maintenance, and performance monitoring.

Recently, some vendors are promoting the "democratization" of process automation using low code, cloud-based development environments. These new environments can speed up development for some process automation projects, but organizations may end up with new automation silos. In general, governance policies should specify a more strategic approach to project selection and development.

Policies that were relevant at one point in time may become dated or irrelevant. Industry associations, vendors, and other public sources are likely to share and disseminate best practices. The above suggestions are general and generic. They serve as a starting point for thinking about and creating a more customized governance document and framework. Lack of controls for an IA program can prevent an organization from meeting security, privacy, and compliance requirements.

Conclusion

IA will solve some problems but may create new ones, such as biased decision making or worker alienation. More intelligent business processes and production automation may lead to higher productivity, more flexible production systems, and, potentially, more stable profitability. Effective governance helps ensure the benefits of IA are realized. IA is an experiment in advanced production concepts, ML, and digital transformation of work. The automated plant of just a few years ago is no longer an advanced production facility. Technology obsolescence is occurring rapidly, and that phenomenon will continue.

> *IA will solve some problems but may create new ones, such as biased decision making or worker alienation.*

The comprehensiveness and sophistication of a governance structure for an IA project and for specialized applications should be determined by the scale of the project, its importance and perceived risk, and the complexity of the integration of emerging technologies such as AI, Internet of Things, blockchain, robotics, and RPA. The prior experiences of regulators, vendors, and managers also impact the sophistication of governance. A smart city or factory should have a governance framework that differs from that for traditional, legacy IT systems.

There are a wide variety of IA implementations in organizations. Emerging, innovative technologies for IA will be disruptive and there are many "unknowns," so caution and "pilot" testing are important, but managers should use controlled experimentation rather than a "try it and see what you think" approach. For large-scale, urgent, perceived high-risk IA projects that integrate many emerging technologies, managers should identify needs and design a system, anticipate consequences, build and pilot test components, and then observe the operation of the system. Once enough operational data has been collected, it is important to analyze the data and formalize lessons learned. In most situations, engineers and technologists will then improve, refine, and expand the IA systems, and managers will improve governance policies before data collection and observation and monitoring of the systems are repeated. Technology governance should not be a bureaucratic, static process but should be envisioned as a dynamic framework with continuous improvement. Managers should embed "test and learn" processes in an IA governance structure.

Finally, managers can and should anticipate the production processes of the future in many industries, including retail, insurance, logistics, food processing, traditional manufacturing (e.g., automobile assembly), and life sciences. Adoption of IA, however, requires more sophisticated governance and control than earlier RPA. Developers of IA products and systems should be proactive about issues such as testing, quality control, and compliance. Senior managers must be proactive in defining policies and procedures and monitoring changes. Managers must create a broad, adaptive governance program for IA that assesses, develops, enforces, and monitors policies and procedures prior to implementation of complex, integrated IA systems. After implementation, a governance program must be agile, ethical, and responsive to social, economic, regulatory, and technological change.

References

[1]About Amazon Staff. "What Robots Do (and Don't Do) at Amazon Fulfillment Centers." Amazon, 2020.

[2]Vincent, James. "Amazon Warehouse Workers Are Getting Utility Belts That Ward Off Robots." The Verge, 21 January 2019.

[3]See Wikipedia's "Tesla Factory."

[4]"Get Snapshot from Progressive." Progressive Casualty Insurance Company, 2020.

[5]Leviathan, Yaniv. "Google Duplex: An AI System for Accomplishing Real-World Tasks over the Phone." *Google AI Blog*, 8 May 2018.

[6]Infocomm Media Development Authority (IMDA). "Singapore Implements Artificial Intelligence Governance and Ethics Initiatives." Cision PR Newswire, 5 June 2018.

[7]IMDA (see 6).

Daniel J. Power is a Professor of Information Systems and Strategy at the College of Business Administration, University of Northern Iowa. He is also the Editor of DSSResources.COM. Dr. Power has published more than 50 refereed journal articles and book chapters. His articles have appeared in leading academic journals, including Decision Support Systems, Decision Sciences, *and* Journal of Decision Systems. *Dr. Power has authored or coauthored six books on computerized decision support. His two most recent books, coauthored with Ciara Heavin, are* Decision Support, Analytics, and Business Intelligence *and* Data-Based Decision-Making and Digital Transformation. *Dr. Power earned a PhD in business administration from the University of Wisconsin-Madison. He can be reached at daniel.power@uni.edu.*

Ciara Heavin is a lecturer/researcher in the Department of Business Information Systems at Cork University Business School, University College Cork (UCC), Ireland. Her research focuses on opportunities for using information systems in the global healthcare ecosystem and in digital transformation. As Co-Director of the Health Information Systems Research Centre, Dr. Heavin has directed funded research in the investigation, development, and implementation of innovative technology solutions in the health/healthcare domain. She has published articles in several top international IS journals and conference proceedings. Dr. Heavin coauthored two books with Daniel J. Power: Decision Support, Analytics, and Business Intelligence *and* Data-Based Decision-Making and Digital Transformation*. She earned a PhD from UCC. She can be reached at C.Heavin@ucc.ie.*

Shashidhar "Shashi" Kaparthi is an Associate Professor of Management Information Systems at the University of Northern Iowa. He specializes in decision support systems; Internet technologies; business data analytics; supply chain and operations management; IT strategy and innovation; and artificial intelligence, including neural network models and applications in management research. Dr. Kaparthi was the inaugural CIO at the University of Northern Iowa. In that role, he led the university community and the IT services unit in delivering high-quality, cost-effective, and secure services; provided the university's IT vision; and was responsible for the alignment of IT infrastructure and capabilities with the university's mission and strategic goals. Dr. Kaparthi is the author of Macromedia ColdFusion *and has been published in numerous journals, such as* Decision Sciences, Journal of Decision Systems, International Journal of Production Research, European Journal of Operational Research, and Journal of Intelligent Manufacturing*. He earned a PhD in management systems from the University at Buffalo School of Management at State University of New York Buffalo and an undergraduate degree in mechanical engineering from the Indian Institute of Technology Madras. He can be reached at shashi.kaparthi@uni.edu.*

CUTTER Business Technology Journal

About Cutter Consortium

Cutter Consortium is a unique, global business technology advisory firm dedicated to helping organizations leverage emerging technologies and the latest business management thinking to achieve competitive advantage and mission success. Through its research, training, executive education, and consulting, Cutter Consortium enables digital transformation.

Cutter Consortium helps clients address the spectrum of challenges technology change brings — from disruption of business models and the sustainable innovation, change management, and leadership a new order demands, to the creation, implementation, and optimization of software and systems that power newly holistic enterprise and business unit strategies.

Cutter Consortium pushes the thinking in the field by fostering debate and collaboration among its global community of thought leaders. Coupled with its famously objective "no ties to vendors" policy, Cutter Consortium's *Access to the Experts* approach delivers cutting-edge, objective information and innovative solutions to its clients worldwide.

For more information, visit www.cutter.com or call us at +1 781 648 8700.