Cognizant

Digital Business

Making AI Responsible – And Effective

To bring AI into the business mainstream, companies need to complement their technology advances with a focus on governance that drives ethics and trust. If they don't, their AI efforts will fall short of their expectations and lag the business results delivered by competitors that responsibly embrace machine intelligence, our latest research findings suggest.

October 2018

Executive Summary

Artificial intelligence (AI) has migrated from science fiction to lab experiments and, now, to actual business operations. Most people already encounter these technologies as they read product recommendations on Amazon, receive automated fraud alerts from credit card companies or ask their Google Assistant to play a certain song. But AI is also at work in less visible ways, as well – scoring bank customers for creditworthiness, analyzing warranty claims to uncover upstream production problems, grading college essays and even helping courts determine how to sentence criminals.

Nevertheless, businesses are in the early stages of AI adoption, and companies are still learning how to put the technology to work. The challenge today is less about understanding technical questions and technology capabilities, and more about crafting a strategy, determining the governance structures and practices needed for "responsible AI," and accelerating the move from experiments to full-scale AI adoption.

Earlier this year, we assessed the state of AI in business through the eyes of 975 executives across industries in the U.S. and Europe (see Methodology, page 22), asking them about their attitudes, expectations and plans. Key findings include:

- **Executives are enthusiastic about the importance and potential benefits of AI.** Roughly two-thirds of executives say that AI is extremely or very important to their companies today, and more than eight out of 10 expect this will be the case three years from now. About three-quarters foresee major or significant benefits from the technology, in terms of lower costs; increased revenues; and the ability to introduce new products or services, or to enter new businesses.
- **Faster-growing companies are more optimistic and aggressive when it comes to AI.** Executives at faster-growing companies are more likely than others to view AI as important, and more likely to expect major benefits in the coming years. Eighty-six percent of executives at these companies say AI is extremely or very important to their company's success, compared with 57% of those at their slower-growing competitors. These industry leaders say they plan to use AI to drive further growth, solidifying their leading positions and pulling even further away from the pack.
- **There is a major disconnect between executives' optimism about AI and actual AI implementation.** While two-thirds of executives said they knew about an AI project at their company, only 24% of that group – just 15% of all respondents – were aware of projects that were fully implemented., with the large majority of AI projects still in the early planning or pilot stages. This contrasts sharply with the fact that 84% of respondents believe AI will be critically important three years from now.
- **Many companies appear to lack a strategic focus for AI.** When identifying the challenges they face in implementing AI and the technologies they're targeting, executive responses were relatively undifferentiated in both areas – essentially adding up to "all of the above." This suggests that companies have yet to hone their strategic plans for

integrating AI into the business core. In addition, roughly 40% of respondents said that securing senior management commitment, buy-in by the business and even adequate budget were extremely or very challenging, indicating that many companies are not yet fully committed to AI's central role in advancing business objectives.

- **Many companies are not addressing the ethical dimensions of AI.** Only about half of the executives surveyed said their companies have policies and procedures in place to identify and address ethical considerations – either in the initial design of AI applications or in their behavior after the system is launched. We believe organizations need to focus on an ethics lens upfront in the strategy/value stream when deciding whether a problem/need has to be solved through AI.

  What's more, about two-thirds of that group said the ethical governance structures they have in place are extremely or very effective. That view may be overly optimistic, and an indication that many executives are underestimating or unfamiliar with the challenging ethical questions likely to emerge as AI becomes more sophisticated and pervasive both in its development and use.

Companies will need to take action on several fronts if they are to achieve significant business benefits with AI. The key agenda items include:

- **Formulate AI strategies.** These should focus on opportunities that promise measurable value – not only reduced costs and increased revenue, but also benefits such as improved customer service, entry to new lines of business and enhanced employee experiences. It's especially critical that strategies take a human-centric view of AI, so that machines can work successfully alongside and for people.
- **Develop governance structures.** Companies will need to work proactively to ensure that AI decision-making is transparent to those involved; that AI earns trust by avoiding errors and data-driven biases; and that AI is personalized and able to provide tailored, relevant and context-aware support as it interacts with humans.
- **Create and maintain responsible AI applications.** Because of AI's potential ubiquity and power, ethical concerns need to be interwoven into everything companies do with the technology. That means building AI systems ethically, and then providing oversight to ensure that those systems operate ethically over time, even as the AI applications learn and evolve. To be successful, companies will need to boost their ethics-related efforts upfront as AI touches more and more parts of business and society.

In many ways, these non-technical considerations – regarding trust, transparency, ethics and human-centric approaches – are more critical and complex than those related to developing and running the technology itself. But they are absolutely critical to the success of AI, because they will determine how well these technologies can be leveraged to advance business objectives.

AI operates in the real world. Companies that do not find solutions in these areas could see their AI efforts fall short, or fail completely. This, in turn, has the potential to irritate customers, alienate employees, drive up R&D and deployment costs, and undermine brand reputation as well as business opportunities – causing businesses to fall behind competitors in the race to unleash the vast business potential of AI.

AI enthusiasm outstrips deployment

Today, AI is clearly top of mind for many executives. Our study intended to demonstrate how this mindshare translates into how and where companies are using the technology and the issues they face. Our findings provide insight into the role that AI is playing and will play in companies, and how organizations will manage and use the technology in the coming years.

Our study found that (no surprise) executives are enthusiastic about AI. They see it as having a growing impact on their businesses, with 63% believing it is already extremely or very important to their company's success, and 84% saying that will be true in three years, including 48% who expect it to be *extremely* important. Business leaders also believe it can be applied throughout their organizations. Respondents reported on AI projects at their companies in a wide variety of functions, with customer service (30%), manufacturing/production (26%), operations (20%) and research and development (18%) cited most often.

Moreover, executives have high expectations for AI. About three-quarters of respondents foresee significant or major benefits from the technology, with one-third saying they expect major benefits. In particular, 79% of respondents expect AI to introduce major or significant benefits in increased efficiency and lower costs over the next three years (see Quick Take, page 5); 74% in increased revenues; and 73% in improved ability to introduce new products and services, or enter new businesses. More than one-third expect to see *major* benefits in each of these areas (see Figure 1).

Expected benefits of AI over next three years

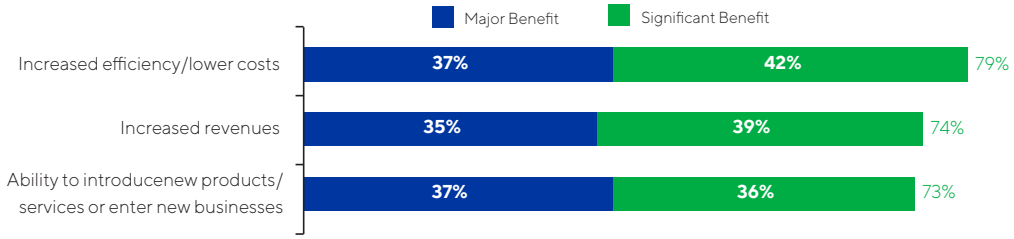

Response base: 975 executives in Europe and the U.S.
Source: Cognizant
Figure 1

Executives have high expectations for AI. About three-quarters of respondents foresee significant or major benefits from the technology, with one-third saying they expect major benefits.

Positive attitudes toward AI are particularly pronounced among respondents who said their companies are growing much faster than the average company in their industries. Executives at these faster-growing companies are more likely to view AI as important to their company's success and more likely to expect major benefits. If these leaders can unleash the potential of AI, their current lead could widen further (see Quick Take, page 6).

Quick Take

Accelerating fraud detection

While AI is still relatively new in business, some companies are already finding ways to use it to increase efficiency and keep costs down. For example, in spite of a great deal of automation in check processing operations, large banks still employ hundreds of people to review computer images of scanned checks to spot signs of fraud. Not only is this process time-consuming and inaccurate; it's also ineffective, with banks losing millions of dollars a year to counterfeiters.

Working with us, one global financial services organization took steps to use AI to augment and streamline its fraud detection process. We developed a machine-learning solution that essentially teaches itself to identify counterfeit checks. It analyzes scanned images of handwritten checks, automatically comparing a variety of factors on the checks to a growing database of checks that have been previously identified as fraudulent. The system then flags potential counterfeits in near real-time while deposit transactions are in process, helping to catch fraud early on.

When we tested our model on a historical portfolio of past transactions, it was able to process up to 20 million checks per day, delivering a fast, accurate confidence score for each check in less than 70 milliseconds. Fewer checks now require manual review, helping to keep processing costs low. And the system is forecast to provide a $20 million reduction in fraud losses annually, based on current models.[1]

Quick Take

The AI-growth connection

While executives in general are enthusiastic about AI, those from faster-growing companies are especially interested in the technology – and that is only increasing the urgency around AI implementations.

Among executives who said their company's revenue growth is far above the industry average, 45% said that AI is extremely important to the success of their companies today, compared with only 14% of executives from slower-growing companies. Looking three years ahead, that gap narrows, but it is still significant (63% vs. 43%) (see Figure 2, next page).

Executives at faster-growing companies are also more likely to expect major benefits from AI – and they expect to see them relatively soon, which suggests a more aggressive approach to implementing these technologies. For example, these executives are more likely than those at slower-growing companies to see major benefits in the next 12 months in terms of increased revenue (36% vs. 11%), increased efficiency/lower costs (33% vs. 9%) and their ability to introduce new products/services or enter new businesses (34% vs. 12%).

Executives from faster-growing companies also appear to be further along in terms of exploring a wider range of AI technologies. Executives at these companies reported on AI projects employing more key AI technologies: computer vision (64% vs. 47%), smart robotics/autonomous vehicles (63% vs. 43%) and analysis of natural language (67% vs. 42%). And faster-growing companies are more likely than others to have ethics policies in place for the initial design of AI applications (71% vs. 47%) and their behavior after they are launched (70% vs. 43%).

These differences may be due in part to faster-growing companies' ability to invest more in new technology. Or it may be that they tend to be more innovative in terms of technology usage – a key competitive quality today – and, therefore, more interested in the potential of AI. In any case, these findings are especially important since investment in technology is driving a trend in which larger companies are pulling further ahead of their smaller rivals.[2] AI promises to take efficiency and productivity to new levels – and that makes it critical for companies to develop and execute AI strategies sooner rather than later if they are not to be left behind.

AI embraced by fast-growers

Respondents were asked to rate the importance of AI to company success, today and in three years' time.

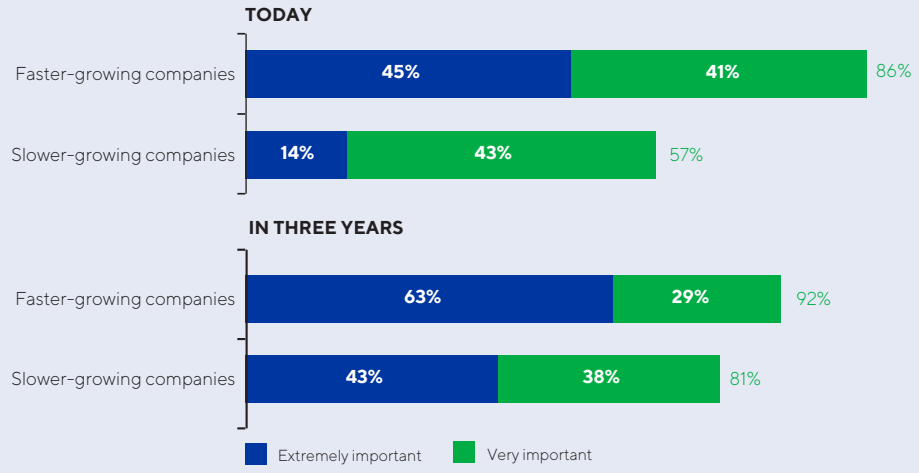

Note: Faster-growing companies were defined as those whose revenue growth is far above the industry average, as reported by respondents.

Response base: 975 executives in Europe and the U.S.
Source: Cognizant
Figure 2

There is a disconnect between executives' enthusiasm for AI and their actual deployment of AI applications. In reality, companies' real-world experience with AI is fairly limited.

Attitudes toward AI vary somewhat by industry: Financial services, technology and retail executives are most optimistic about both the importance and benefits of AI today, while insurance executives are less so. But these differences are not especially large, and looking ahead, they nearly disappear, with about eight out of 10 executives across industries saying AI will be extremely or very important to their company's success in three years. This broad, cross-industry appeal is fairly unusual for a new technology and is, perhaps, a reflection of the wide variety of roles envisioned for the technology, as well as the tremendous media attention it has received.

However, there is a disconnect between executives' enthusiasm for AI and their actual deployment of AI applications. In reality, companies' real-world experience with AI is fairly limited. Only about two-thirds of respondents were knowledgeable about an AI project at their company, and out of that group, only 24% were knowledgeable about AI projects that were fully implemented. Most were aware of projects that are still somewhat experimental: 21% reported on projects in the initial planning stage, 32% in the proof of concept stage, and 48% in the pilot stage (see Figure 3). In short, companies appear to be moving fairly slowly to implement actual AI projects.

This limited level of adoption is no doubt due in part to the relative newness of practical, business-ready AI technology. But executives cited a range of challenges that they see with AI. Leading the list are the need to secure executive talent, and attract and retain professionals with AI experience and skills, with 45% of executives considering these factors to be extremely or very challenging when employing AI (see Quick Take, next page). But not far behind are factors ranging from budget concerns (43%) to the interactions between different AI applications (42%) and securing senior management commitment (40%) – a list that may seem daunting to executives exploring the potential of AI.

Current status of AI projects

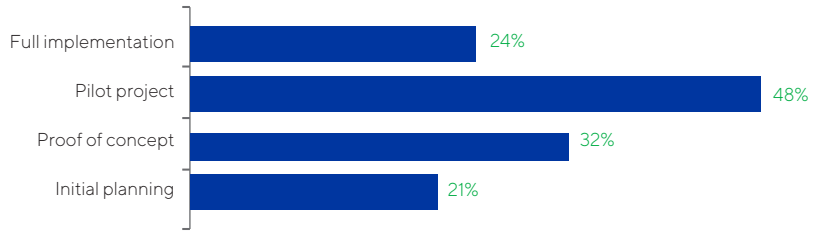

Note: Percentages total to more than 100% since respondents reported on multiple projects.

Response base: 975 executives in Europe and the U.S.
Source: Cognizant
Figure 3

Quick Take

Help still wanted

A number of observers today have expressed concern about the impact of AI on employment, with AI's potential to take on a variety of tasks traditionally performed by people. But just how AI will affect the labor market is not that easy to predict. Indeed, surveyed executives were divided on the question, with 37% saying AI will decrease the number of people they employ, and 44% saying it will increase their employee ranks. Retail and financial services industry executives were most likely to expect a boost to employment (see Figure 4).

We believe that although AI will eliminate some jobs, others will emerge. We estimate that over the next 10 to 15 years, nearly 21 million new jobs will be created by the advent of AI. In addition, many employees will spend less time doing repetitive tasks and more time doing tasks requiring judgment, empathy, creativity and human-specific skills.

AI will also create entirely new positions such as Chief Trust Officer or Ethical Sourcing Manager. (Our "21 Jobs of the Future" report[3] discusses these and other new jobs that are likely to be created over the next 10 years.) That is, as work changes, new roles will emerge

Expected impact of AI on company headcount over next three years

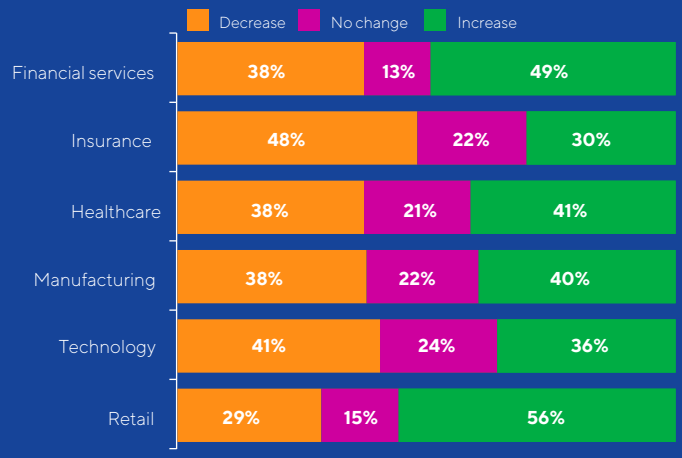

Percentages may not round to 100% due to rounding
Response base: 596 executives in Europe and the U.S. who reported one or more AI projects
Source: Cognizant
Figure 4

as others disappear. Surveyed executives seem to recognize this pending shift in work, as most cited the need for additional skills and the need for employees to interact with AI applications as important challenges.

To a great extent, tomorrow's jobs will use AI to augment and complement human skills, and use human-machine collaboration to bring together the best of each. Thus, we believe that many new jobs will draw on human abilities for tasks such as *coaching* – helping others do things well; *caring* – helping others with health and wellness; and *connecting* – helping to connect people with people, people with machines, the physical with the virtual world, and virtual work and commerce with ethics.[4]

Formulating an AI strategy

Companies will need to sharpen their focus and develop strategies for accelerating the move to AI. As always, the world is not standing still: The technology is evolving quickly, and some early adopters are already identifying AI opportunities and deploying AI technologies in the business.

The survey findings suggest that considerable uncertainty remains about how to proceed, and a lack of strategic focus when it comes to AI. For example, while talent and skills top the list of challenges, executives appear to give all challenges equal weight. Among the 13 issues listed in the survey, the difference between the most-cited and least-cited was just seven percentage points, and roughly 40% of executives considered each of the 13 issues to be extremely or very challenging. This "all of the above" finding suggests that companies have not developed clear strategies for pursuing AI. So too does the nature of some of their challenges, such as securing adequate budget (43%), senior management commitment (40%) and buy-in by the business (38%).

This lack of strategy is further underscored by the fact that roughly the same percentages of executives cited specific AI technologies as being used in their projects: virtual agents (55%), computer vision (51%), advice engines/machine learning (49%), smart robotics/autonomous vehicles (48%) and analysis of natural language (48%). As with the AI challenges, these responses once again essentially add up to "all of the above." This indicates that companies are not yet at a point of understanding where and how specific technologies can benefit the business.

To create rigorous AI strategies, companies need to look beyond technological capabilities. There is no recipe that companies can use to embed AI in the fabric of the company – each business challenge will require different tools, techniques and approaches. Rather than a sequential process, leveraging AI will demand extensive experimentation and the ability to apply learnings to the next stage of deployment. Companies need to factor that reality into their plans.

AI strategy should start with an emphasis on the creation of business value and ethical/responsible behavior. That may sound like common sense, but many AI initiatives today are focused more on the technology's capabilities and algorithms than on its impact and benefit for the business. That is, they tend to ask questions such as, "What can IBM's Watson or Amazon's machine learning do?" rather than, "What business problems or opportunities do we have that could be addressed with AI?"

Companies often think of AI as a tool for reducing costs, and that is certainly a key benefit – one that tops the list for executives in the survey. But there is a range of other potential business benefits to consider. If done right, AI can help improve product and service quality; reduce cycle time; create new and better employee experiences; and enhance safety, among other things.

Employing AI effectively requires a clear focus on applying intelligent technologies to solve tough operational challenges and deliver a lift to the business. For example, we worked with a utility to deploy voice-activated, AI-driven chatbots that assist executives, account managers and field-service technicians in conducting research into services and solutions using voice commands or typed queries, which helped the utility streamline customer interactions and enhance user experience. In another project, we helped an insurer use machine learning, along with geospatial analysis, to better understand the complex flood-insurance market – an effort that identified an overall financial opportunity of $3.3 billion and opened the door to a new line of business for the company.

Each company's situation is different, of course. But there are some general guidelines that companies can follow in the search for AI-enabled business value:

- **Look for opportunities to leverage data.** Companies are increasingly awash in data about everything from customer preferences and competitors to the details of Internet of Things (IoT)-equipped production facilities. This flood of data can quickly overwhelm humans trying to understand and interpret it. The irony: As more data is generated in the modern digital era, separating signal from noise becomes more difficult. AI, with its ability to apply human-like assessments and decisions to that data, rapidly and efficiently, provides a possible solution to that dilemma (see Quick Take, page 13).
- **Cast a wide net.** AI has the potential to touch many parts of the company, and it's important that underlying algorithms "understand" the larger context in which they operate. If they don't, they are likely to make recommendations for the right product but at the wrong time, for example, or struggle to make sound decisions involving cost and speed tradeoffs. With that in mind, organizations should establish cross-functional teams to identify AI opportunities, and arm them with a structured, systematic approach to uncovering potential improvements to processes and stakeholder relationships.
- **Solicit input and insight from external parties.** Partnering can be an important factor in AI strategy. In our survey, executives most often cited access to AI skills as a key challenge, and working with vendors can allow companies to quickly gain access to the required skills. Many companies appear to recognize this: 62% of executives said they rely on a mix of vendor and in-house resources for AI, while 13% said they always or usually secure AI applications from vendors. Companies should consider expanding these relationships into robust ecosystems that can work collaboratively to keep up with rapidly evolving technologies and use cases, and accelerate the journey to AI.

- **Encourage experimentation – and discipline.** While AI is rapidly moving into real-world business applications, it is still an emerging technology. To a great extent, companies will need to chart their own course forward – there are no ready-made, turnkey solutions, and there is still much to learn. That makes managed experimentation key. Companies should encourage a tolerance for risk-taking and innovation with AI, but balance that with rigorous testing and measurement of ROI and tangible business value. They should move on from failures quickly, and at the same time, be prepared to quickly scale up successful experiments and pilots into production.[5]

Perhaps most important, an effective strategy requires a human-centric approach to AI. The technology is most powerful when it collaborates with people and augments human activities and decisions – and a human-oriented approach is key to ensuring that AI fits into that real-world business context (see Quick Take, page 15). Thus, companies should consistently ask how AI can enable and support employees and customers, how it can best interact with people, how it can be designed to understand the behaviors and expectations of humans, and how it can have enough "emotional intelligence" to adjust its decisions to the people around it.

To help ensure a human-centric approach, companies can complement the input of business people with the insights of behavioral scientists (sociologists, anthropologists, etc.) to shape and train AI to be an effective, and valued, part of the organization.

AI technology is most powerful when it collaborates with people and augments human activities and decisions.

Quick Take

Working smarter to reduce compliance risk

AI brings a fundamental quality to the table – the ability to help companies leverage today's growing volumes of data, quickly and intelligently. That proved to be critical to an international professional services organization's efforts to reduce financial and compliance risk. The organization frequently has to conduct due diligence on the numerous vendor, partner and counterparty companies that are part of its business ecosystem. This typically involves going through more than 40,000 global sources, including media reports, corporate records, financial transactions and legal cases. This information had to be painstakingly reviewed for each entity before a report could be finalized – a laborious process that could take weeks.

The organization worked with us to apply advanced forms of AI to tackle the problem. It had adopted a third-party research tool to compile search results, but that system lacked "smart" search capabilities, and researchers were still spending a large amount of time reviewing irrelevant results. To improve the functionality of the system, we quickly developed a proof of concept for a machine-learning platform based on IBM's Watson Explorer 11 and Watson Content Analytics, along with advanced APIs to link the AI platform to the firm's existing systems.

With the solution, a contextual search model allows the organization to detect and deliver improved reports on news relating to risk and fraud. A machine-learning model understands linguistic nuances, meanings and relationships specific to an industry. This model extracts relevant information from enormous, diverse data sets, producing research reports with contextually relevant results. The solution can rank results by the likelihood of risk based on context, without researchers having to conduct multiple searches.

The solution automates more of the search process, integrates research workflow and reduces the time spent manually reviewing material by more than half – and 14% of reports can now be completed in one hour. As a result, the organization can generate up to 30% more due diligence reports a year, helping to reduce compliance and financial risks significantly.[6]

Quick Take

Bringing human-centric AI to healthcare

When approached from a human-centric perspective, AI can benefit individuals, businesses and society as a whole. For example, opioid dependency is devastating for patients and their families, and it can limit positive outcomes when addicted patients are being treated for other illnesses. It's also expensive: U.S. healthcare organizations spend more than $500 billion annually to care for patients suffering from opioid addiction, and treating addicted patients diverts much-needed resources from other patients.[7]

We worked with a leading healthcare services provider to develop an AI-driven solution that detects drug-seeking behavior and identifies patients who are at risk of becoming addicted. The solution uses an advanced machine-learning algorithm to mine information about patient behaviors and symptoms, and examines physicians' notes from patient interactions, captured using text analytics.

The system combines phase-based extraction, rule-filtering and advanced text clustering to mine highly variable data and ultimately identify patients who could become drug-seekers – before they turn into addicts. Pop-up alerts in the organization's electronic medical records system prompt physicians to take corrective actions at the point of care, interceding with patients in real time. The solution also learns continuously from its own results to verify the accuracy of its models and improve searches.

The tangible health and financial benefits from this solution include saving the healthcare organization as much as $60 million through the identification of some 85,000 at-risk patients. The intangible, but nevertheless invaluable, benefits include improved quality of life for thousands of patients, their families and society.

Overcoming strategic challenges

Even the most well-developed AI strategies won't succeed unless fundamental challenges are considered. One of these challenges is the ability to build a profitable and scalable business from AI. Many established businesses start with foundational AI centers of excellence, but these teams quickly confront questions that extend beyond the improvement of customer or user experience, or the use of AI to automate a process, such as:

- Can they trade in AI services in the future?
- Can they hive off a profitable business based on AI services on top of their existing product or service?

A large home appliance company recently partnered with a number of start-ups to understand how people cook, with the objective of influencing cooking and improving the dining experience. The company wanted to create a platform that will help it accomplish this goal using AI microservices embedded in its devices. Eventually, the company foresees an operating model where it monetizes this microservice by engaging retail partners across areas such as food, delivery and packaging. Other traditional companies are also partnering or buying AI start-ups in anticipation of realizing a "profitable and scalable" business from AI.

Another challenge is scalability of AI microservices. For example, each connected car or connected home appliance will have several different AI microservices and deep-learning models running on its units/devices to deliver a variety of services. Accommodating the millions of devices in use that want to make use of these microservices, globally, is non-trivial. For digital-native companies, this is slightly more manageable because they own more of their value chains and often have built their business around intelligent algorithms. However, the challenge for traditional companies like automobile manufacturers and home appliance makers is that their AI services scope will be incomplete unless they tap into their traditional value chain, including manufacturing, logistics, service partners, and warranty and fulfillment partners, to effect a wider AI impact.

AI governance: building transparency, trust and personalization

For companies to successfully unleash the potential of AI, people will need to see it as a reliable, dependable colleague. It will interact with customers, employees and partners, running operations and making important decisions. If an AI application is not well designed and managed, it may end up "misbehaving." The ramifications could be significant, ranging from damaged customer relationships, to missteps in factories that affect quality, to discriminatory decisions that elicit regulatory scrutiny.

Companies will need to put the right AI governance structures in place to guide its design and use. They'll need to develop policies for validating and operating the AI system, and work continuously to monitor and correct AI behavior. To that end, companies should consider establishing an AI council to monitor and oversee AI, and focus on ensuring three key factors: transparency, trust and personalization.[9]

Transparency: beyond the black box

Transparency boils down to allowing people to understand how an AI application makes decisions. Currently, AI systems tend to be "black boxes" whose operations are not well understood by humans, and that makes some people uncomfortable. We want to know why our colleagues and friends reach the conclusions they do, and we would like to do the same with intelligent machines. Companies will need to help stakeholders understand what AI applications know about people and why they behave the way they do. For example:

- Banking customers will want to know the criteria systems use for approving loans.
- Plant operators will want to know why AI decided to shut down a production line for safety reasons.
- On a more mundane level, customers will want to know why they are being offered, say, hiking boots rather than running shoes.

Transparency helps build trust, but trust is based on other factors, too. Can people feel confident that the system is working as it was designed to work, and that it understands their needs? Consistency, too, is important – people learn to trust over time, through experience.

Companies will need to put the right AI governance structures in place to guide its design and use. They'll need to develop policies for validating and operating the AI system, and work continuously to monitor and correct AI behavior.

Data needs to be not only accurate and free of "noise," but also free of bias. AI is not programmed in the traditional sense; rather, it learns from examples – and biased data is essentially a bad example.

Trust: it starts with the data

High-quality data is key to earning trust. Data needs to be not only accurate and free of "noise," but also free of bias. AI is not programmed in the traditional sense; rather, it learns from examples – and biased data is essentially a bad example. That was clearly illustrated by Microsoft's "Tay" chatbot, an AI application that was intended to learn how to behave well by interacting with Twitter users but ended up producing racist comments after a few hours of listening to online trolls.[10]

On another level, researchers have noted that lending and credit-scoring AI applications that learn from historical data can quickly pick up the gender and race biases built into that data. Companies will need to find ways to systematically mitigate biases and ensure that AI is learning from the right inputs – or risk implementing the kind of error-laden systems that a World Economic Forum report called "artificial stupidity."[11] (Such issues can also raise ethical questions; see next section, next page).

In using data with AI, companies will need to keep compliance in mind, as regulations around data privacy continue to evolve and expand. For example, the European Union's General Data Protection Regulation (GDPR) is a recently enacted set of rigorous data privacy rules designed to safeguard the privacy rights of data subjects – that is, customers and others whose data is used by a company. These rules are specific to AI, but will have important ramifications for the field. For example, GDPR requires companies to:

- Institutionalize the practice of asking data subjects for consent to use their data.
- Make the logic behind automation clear to data subjects who ask about it.
- Establish data protection checks and balance so that AI is not drawing on data that goes beyond its stated purpose.

While the GDPR is a European Union regulation, it applies to companies headquartered anywhere that do business in Europe, including U.S. companies.

Personalization: making it relevant and useful

Finally, personalization will be important in driving the acceptance and success of AI, because it is key to performing tasks and making decisions that are relevant and useful for the people they impact. Employees, for example, typically tailor their interactions with one another based on the individuals and situation involved. AI applications will need to do the same, whether it is suggesting an action to a banker or recommending a product to a customer.

This shaping of interactions not only increases the effectiveness of AI; it can also enhance the trust that people feel for the technology.

Do the right thing: responsible, ethical AI

In the end, the goal of good governance is "responsible AI." If AI is not responsible, it won't be embraced by employees or customers, not to mention shareholders who are put off by potential legal and reputational risk.

There is another, more fundamental thread at the heart of good AI governance and responsible AI: ethics. Increasingly, the development of AI ethics is recognized as critical to the success of AI, and it needs to be interwoven into everything companies do with the technology.

That means not only building ethical AI systems, but also running those systems ethically, in ways that are aligned with corporate and societal values. By its nature, AI is always adapting to and building on new information. Ethics will become particularly important as AI becomes more ubiquitous, and as AI systems increasingly learn from one another, not just from the inputs that humans provide – that is, when machine-learning AI applications "teach" other AI applications. An ethical foundation is needed to ensure that AI continues to operate for the benefit of humans over time. In a very real sense, ethics is at the heart of human-centric AI.

That reality is driving a growing number of programs that are focused on AI ethics. For example, the MIT Media Lab and Harvard University's Berkman-Klein Center for Internet and Society have launched an Ethics and Governance of AI Initiative to look at AI and the "social values of fairness, human autonomy and justice."[14] The University of Oxford is collaborating with other UK universities to find ways to make AI more transparent and accountable.[15] Facebook, Google, Microsoft and other tech companies have created The Partnership on AI to explore AI best practices and conduct "open research and dialogue on the ethical, social, economic and legal implications of AI."[14] And New York City recently established a task force to look at fairness and accountability in AI decisions made in various city departments.[15]

Companies will need to find ways to train AI to behave ethically. This is not as arcane a process as one might think. The concept is familiar to parents who provide feedback and guidance to raise their children to be good members of society, and there are well-understood tools and frameworks from the world of human sciences that can be used to instill ethics into the design and operation of AI.

While AI ethics is becoming an increasingly high-profile concern, its importance is not always widely recognized in business. Less than half of the executives (45%) surveyed said that ethical

Ethics will become particularly important as AI becomes more ubiquitous, and as AI systems increasingly learn from one another, not just from the inputs that humans provide – that is, when AI applications "teach" other AI applications.

Tools used to identify and address potential unethical behavior in AI applications

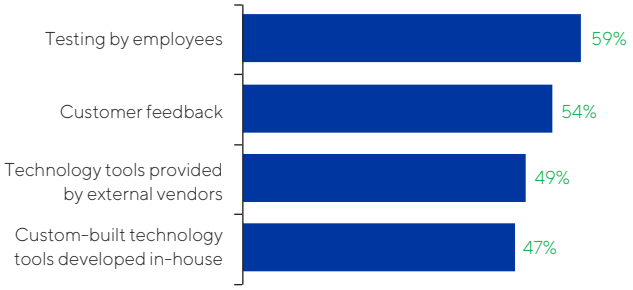

Multiple responses permitted
Response base: 975 executives in Europe and the U.S.
Source: Cognizant
Figure 5

considerations play a critical or significant role when their company develops and employs AI. In contrast, a majority (55%) said that ethical considerations play some, a minor, or little or no role.

Only about half of respondents reported that their companies have policies and procedures in place to address potential unethical behavior in the initial design of AI, although 35% plan to put these in place. Similar percentages of executives said the same about their approach to dealing with unethical behavior after AI applications are launched. In terms of how they uncover unethical behavior, executives most often cited testing by employees (59%), followed by customer feedback (54%), technology tools provided by external vendors (49%) and custom-built technology tools developed in-house (47%) (see Figure 5).

Approximately two-thirds of executives believe their companies' policies for ensuring ethical AI behavior are extremely or very effective. This somewhat optimistic view may be based in part on the relatively limited rollout of AI at many companies. As AI finds its way into more parts of the organization, and as it continues to learn and evolve in real-world business situations, maintaining an ethical foundation will become more complicated as well as more important. We believe that in time, many companies that are satisfied with their current AI ethics policies and procedures will realize they need to redouble their efforts in order to keep up with a growing range of ethical questions, such as:

- Who bears legal responsibility for a machine's harmful conduct?
- Should people be made aware when they are dealing with a very human-like AI system?
- What oversight is needed to keep AI from being used for criminal purposes?
- How can we manage the economic and labor disruptions that AI might bring?
- Should limits be placed on the intelligence and autonomy of AI applications?

To answer these questions, companies will need to maintain a relentless focus on ethics over time.

Looking forward

Unlike many established information technologies, it is often easy to get started with AI. New AI systems can be developed fairly quickly – an experimental version can typically be up and running in a month, and pilots can be rolled out in a month or two after that. The trick is to move rapidly from that starting point to grow and sustain full-scale, business-ready AI applications.

To help move projects forward, companies should start with an assessment of how prepared the organization is for AI, and what it needs to be put in place. What platforms will be used to run the scaled-up application? What skills and knowledge will be needed, and will these come from in-house personnel or be sourced from a partner? What data will be needed from traditional transaction systems, such as sales, marketing and finance; from newer digital systems, such as point-of-sale and social media applications; and from external, context-setting sources, such as governments, weather services, online review sites and so forth?

In terms of launching projects, companies may want to focus their initial AI efforts on targeted back-office applications, an approach that can provide an opportunity to gain experience with the technology without risking high-profile problems with customer-facing systems. They should also look for quick-win opportunities that allow them to chalk up successes and financial benefits to build AI momentum. And they need to understand and address AI's impact on employees – how their roles are shifting and what new skills they will need – and then employ change-management techniques to help people succeed with new AI-based tools.

As they nurture AI and move it from experimentation to the business core, companies should also keep an eye on the big picture. That will mean formulating plans for governance, and always keeping the need for responsible AI top of mind. Through it all, they should look at their efforts through a human-centric lens – because that will be key to designing and reaping business value from AI in the long run.

New AI systems can be developed fairly quickly. The trick is to move quickly from that starting point to grow and sustain full-scale, business-ready AI applications.

Methodology

We conducted a survey, through a combination of online and telephone interviews, with 975 respondents from companies in the U.S. and Europe in June and July 2018.

The companies represented a variety of industries and sizes, as measured by their annual revenues.

Fifty percent of the respondents were from companies headquartered in the U.S., 48% from companies headquartered in Europe, and 2% from companies in other locations. The respondents from European companies were located in the following countries: UK (11%), Germany (11%), France (10%), Netherlands (5%), Sweden (4%), Belgium (4%) and Norway (3%).

More than half the respondents had C-suite titles, with an additional 16% being an executive vice-president, senior vice-president or vice-president. Respondents work in the following functional areas: information technology (39%), management (19%), finance (13%), operations/manufacturing (9%), research & development (4%), sales & marketing (4%), human resources (4%), customer service (3%) and other functions (5%).

Respondent industry

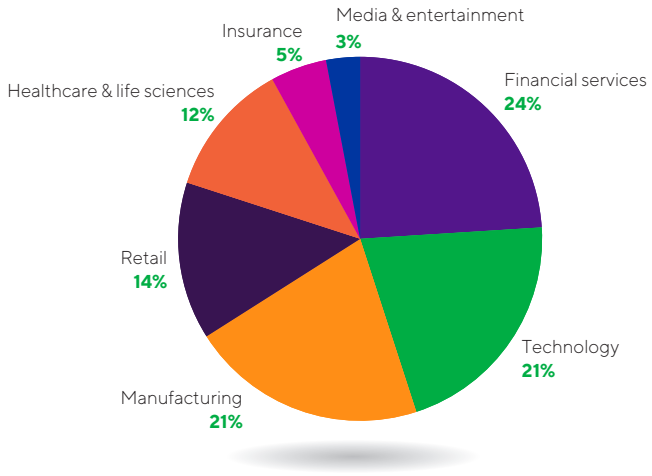

Annual revenues

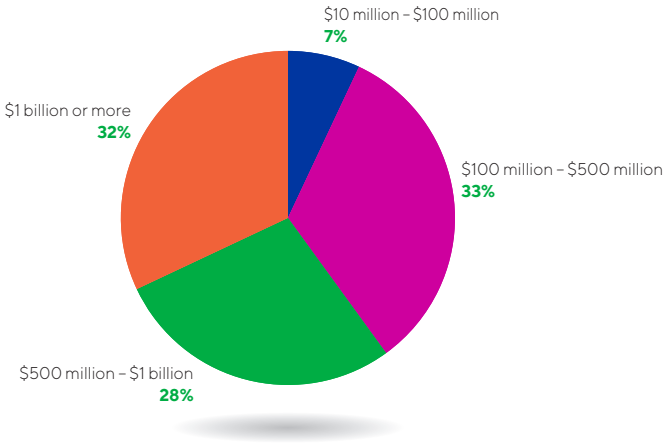

Respondent title

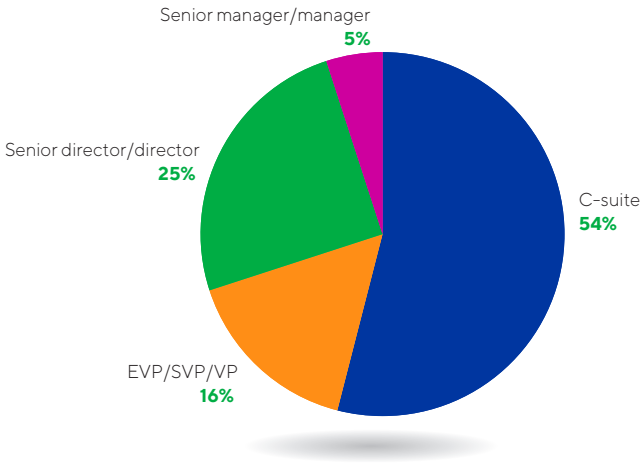

Endnotes

1 For more information on this and other case studies related to employing AI to deliver business value, see "Accelerate Business Growth and Outcomes with AI," Cognizant Technology Solutions, Sept. 7, 2017, https://www.cognizant.com/perspectives/accelerate-business-growth-and-outcomes-with-ai; or visit the "Featured Work" section of Cognizant's AI web pages: https://www.cognizant.com/artificial-intelligence.

2 Christopher Mims, "Why Do the Biggest Companies Keep Getting Bigger? It's How They Spend on Tech," *The Wall Street Journal*, July 26, 2018, https://www.wsj.com/articles/why-do-the-biggest-companies-keep-getting-bigger-its-how-they-spend-on-tech-1532610001.

3 "21 Jobs of the Future," Cognizant Technology Solutions, November 2017, https://www.cognizant.com/whitepapers/21-jobs-of-the-future-a-guide-to-getting-and-staying-employed-over-the-next-10-years-codex3049.pdf

4 Ibid.

5 For a discussion on developing an AI strategy, see "The Road to AI," Cognizant Technology Solutions, August 2018. https://www.cognizant.com/whitepapers/the-road-to-ai-codex3614.pdf.

6 For more information on this and other case studies related to employing AI to deliver business value, see "Accelerate Business Growth and Outcomes with AI," Cognizant Technology Solutions, Sept. 7, 2017, https://www.cognizant.com/perspectives/accelerate-business-growth-and-outcomes-with-ai, or visit the "Featured Work" section of Cognizant's AI web pages: https://www.cognizant.com/artificial-intelligence.

7 "The Underestimated Cost of the Opioid Crisis," The Council of Economic Advisers, November 2017, https://www.whitehouse.gov/sites/whitehouse.gov/files/images/The%20Underestimated%20Cost%20of%20the%20Opioid%20Crisis.pdf.

8 For more information on this and other case studies related to employing AI to deliver business value, see "Accelerate Business Growth and Outcomes with AI," Cognizant Technology Solutions, Sept. 7, 2017, https://www.cognizant.com/perspectives/accelerate-business-growth-and-outcomes-with-ai; or visit the "Featured Work" section of Cognizant's AI web pages: https://www.cognizant.com/artificial-intelligence.

9 For a discussion of AI governance, see "AI: Ready for Business," Cognizant Technology Solutions, July 2018, https://www.cognizant.com/whitepapers/ai-ready-for-business-codex3752.pdf .

10 "Bias In, Bias Out: How AI Can Become Racist," *Design News*, Sept. 29, 2017. https://www.designnews.com/content/bias-bias-out-how-ai-can-become-racist/176888957257555.

11 "Top 9 Ethical Issues in Artificial Intelligence" World Economic Forum, Oct. 21, 2016, https://www.weforum.org/agenda/2016/10/top-10-ethical-issues-in-artificial-intelligence/.

12 The Ethics and Governance of Artificial Intelligence Initiative website: https://aiethicsinitiative.org/.

13 "Oxford to Be Key Player in New AI Accountability Project," University of Oxford, May 9, 2018, http://www.ox.ac.uk/news/2018-05-09-oxford-be-key-player-new-ai-accountability-project.

14 Partnership on AI website: https://www.partnershiponai.org/tenets/.

15 "Mayor de Blasio Announces First-In-Nation Task Force To Examine Automated Decision Systems Used By The City," New York City press release, May 16, 2018. https://www1.nyc.gov/office-of-the-mayor/news/251-18/mayor-de-blasio-first-in-nation-task-force-examine-automated-decision-systems-used-by.


25%
2013
2012

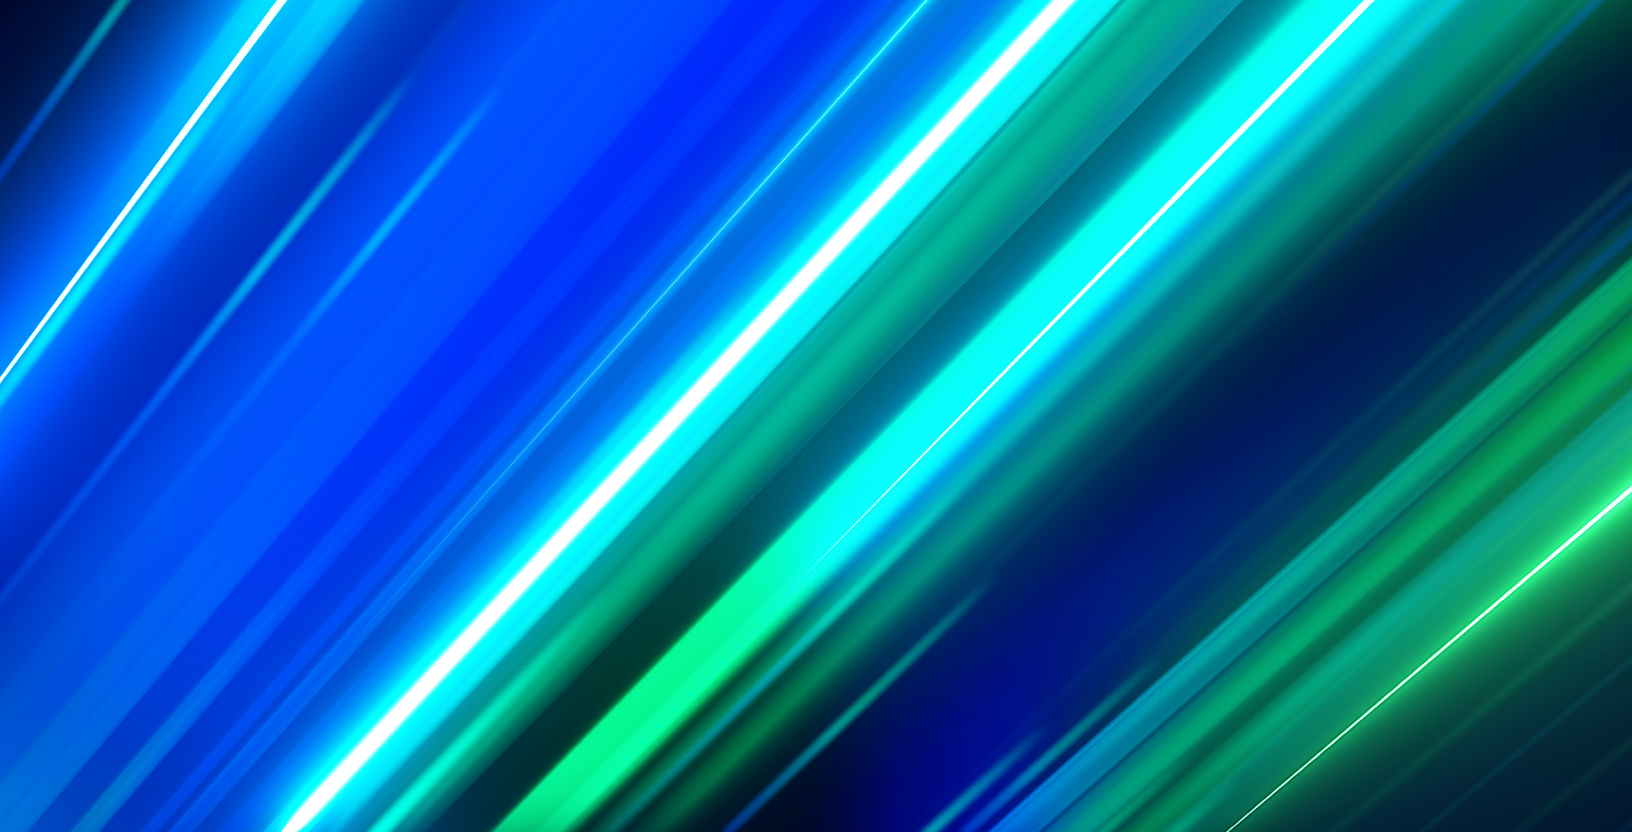

About Cognizant Digital Business

Cognizant Digital Business helps our clients imagine and build the Digital Economy. We do this by bringing together human insight, digital strategy, industry knowledge, design, and new technologies to create new experiences and launch new business models. For more information, please visit www.cognizant.com/digital or join the conversation on LinkedIn.

About Cognizant AI & Analytics Practice

As part of Cognizant Digital Business, Cognizant's AI & Analytics Practice provides advanced data collection and management expertise, as well as artificial intelligence and analytics capabilities that help clients create highly-personalized digital experiences, products and services at every touchpoint of the customer journey. We apply conversational AI and decision support solutions built on machine learning, deep learning and advanced analytics techniques to help our clients optimize their business/IT strategy, identify new growth areas and outperform the competition. Our offerings include AI to Insight, Customer Intelligence, Intelligent Automation, Product Intelligence, and Risk & Fraud Detection. To learn more, visit us at www.cognizant.com/cognizant-digital-business/applied-ai-analytics.

About Cognizant

Cognizant (Nasdaq-100: CTSH) is one of the world's leading professional services companies, transforming clients' business, operating and technology models for the digital era. Our unique industry-based, consultative approach helps clients envision, build and run more innovative and efficient businesses. Headquartered in the U.S., Cognizant is ranked 193 on the Fortune 500 and is consistently listed among the most admired companies in the world. Learn how Cognizant helps clients lead with digital at www.cognizant.com or follow us @Cognizant.

Cognizant

World Headquarters
500 Frank W. Burr Blvd.
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233
Fax: +1 201 801 0243
Toll Free: +1 888 937 3277

European Headquarters
1 Kingdom Street
Paddington Central
London W2 6BD England
Phone: +44 (0) 20 7297 7600
Fax: +44 (0) 20 7121 0102

India Operations Headquarters
#5/535 Old Mahabalipuram Road
Okkiyam Pettai, Thoraipakkam
Chennai, 600 096 India
Phone: +91 (0) 44 4209 6000
Fax: +91 (0) 44 4209 6060

APAC Headquarters
1 Changi Business Park Crescent,
Plaza 8@CBP # 07-04/05/06,
Tower A, Singapore 486025
Phone: + 65 6812 4051
Fax: + 65 6324 4051

Codex 3974.05