Continue

Intelligenza artificiale pdf

 What is AI? Artificial Intelligence (AI) refers to the development of computer systems of performing tasks that require human intelligence. AI aids, in processing amounts of data identifying patterns and making decisions based on the collected information. This can be achieved through techniques like Machine Learning, Natural Language Processing, Computer Vision and Robotics. AI encompasses a range of abilities including learning, reasoning, perception, problem solving, data analysis and language comprehension. The ultimate goal of AI is to create machines that can emulate capabilities and carry out diverse tasks, with enhanced efficiency and precision. The field of AI holds potential to revolutionize aspects of our daily lives.In this article, we will know about What is Artificial Intelligence?, its evolution, various types of AI, training models, benefits.What is Artificial Intelligence ?Examples of AIArtificial Intelligence (AI) has become increasingly integrated into various aspects of our lives, revolutionizing industries and impacting daily routines. Here are some examples illustrating the diverse applications of AI:Virtual Personal Assistants: Popular examples like Siri, Google Assistant, and Amazon Alexa utilize AI to understand and respond to user commands. These assistants employ natural language processing (NLP) and machine learning algorithms to improve their accuracy and provide more personalized responses over time.Autonomous Vehicles: AI powers the development of self-driving cars, trucks, and drones. Companies like Tesla, Waymo, and Uber are at the forefront of this technology, using AI algorithms to analyze sensory data from cameras, radar, and lidar to make real-time driving decisions.Healthcare Diagnosis and Treatment: AI algorithms are used to analyze medical data, including patient records, imaging scans, and genetic information, to assist healthcare professionals in diagnosing diseases and planning treatments. IBM's Watson for Health and Google's DeepMind are examples of AI platforms employed in healthcare.Recommendation Systems: Online platforms like Netflix, Amazon, and Spotify utilize AI to analyze user behaviour and preferences, providing personalized recommendations for movies, products, and music. These systems employ collaborative filtering and content-based filtering techniques to enhance user experience and increase engagement.Fraud Detection: AI algorithms are employed by financial institutions to detect fraudulent activities in real-time. These systems analyze.AI has the potential to revolutionize many industries and fields, such as healthcare, finance, transportation, and education. However, it also raises important ethical and societal questions, such as the impact on employment and privacy, and the responsible development and use of AI technology.Importance of AIToday, the amount of data in the world is so humongous that humans fall short of absorbing, interpreting, and making decisions of the entire data. This complex decision-making requires higher cognitive skills than human beings. This is why we're trying to build machines better than us, in these task. Another major characteristic that AI machines possess but we don't is repetitive learning. Let consider an example of how Artificial Intelligence is important to us. Data that is fed into the machines could be real-life incidents. How people interact, behave and react ? etc. So, in other words, machines learn to think like humans, by observing and learning from humans. That's precisely what is called Machine Learning which is a subfield of AI. Humans are observed to find repetitive tasks highly boring. Accuracy is another factor in which we humans lack. Machines have extremely high accuracy in the tasks that they perform. Machines can also take risks instead of human beings. AI is used in various fields like: Health CareRetailManufacturingBanking etc.Types of AIAI can be broadly classified into two major categories: Based on Capabilities:1. Narrow AI: Narrow AI, also known as Weak AI, refers to artificial intelligence systems that are designed and trained to perform a specific task or a narrow range of tasks. These systems excel at their designated tasks but lack the broad cognitive abilities and understanding of human intelligence. Narrow AI is the most common form of AI currently in use and has found widespread application across various industries and domains.Characteristics of Narrow AI include:Specialized capabilities: Narrow AI systems are highly specialized and focused on performing a specific function or solving a particular problem.Limited scope: These systems are not capable of generalizing their knowledge or skills to other domains outside of their designated tasks.Lack of consciousness: Narrow AI lacks self-awareness and consciousness. It operates based on predefined algorithms and data inputs without understanding the context or implications of its actions.Examples: Virtual personal assistants like Siri and Alexa, recommendation systems, image recognition software, chatbots, and autonomous vehicles are all examples of Narrow AI.2. General AI: Also referred to as "General AI". Here is where there is no difference between a machine and a human being. This is the kind of AI we see in the movies, the robots. A close example (not the perfect example) would be the world's first citizen robot, Sophia. She was introduced to the world on October 11, 2017. Sophia talks like she has emotions.General AI, also known as Strong AI or Artificial General Intelligence (AGI), refers to artificial intelligence systems that replicate human-like cognitive abilities and understanding across a wide range of tasks and domains. Unlike Narrow AI, which is task-specific, General AI aims to simulate human intelligence comprehensively, including reasoning, problem-solving, learning, and adaptation to new situations.Characteristics of General AI include:Human-like cognitive abilities: General AI systems can understand, learn, and apply knowledge across various tasks and domains, similar to human intelligence.Adaptability: These systems have the ability to generalize their knowledge and skills to new situations, tasks, and environments.Consciousness: General AI is theorized to possess self-awareness, consciousness, and subjective experiences, although achieving this level of intelligence remains a theoretical challenge.Examples: General AI remains largely theoretical and speculative, as researchers have not yet achieved human-level artificial intelligence. Development in this area continues to be a subject of ongoing research and exploration.3. Super Intelligent AI: Super intelligent AI refers to artificial intelligence systems that surpass human intelligence in virtually every aspect. This type of AI, also known as Artificial Superintelligence (ASI), represents the highest level of AI capabilities and poses significant implications for society and the future of humanity.Characteristics of Super intelligent AI include:Cognitive superiority: Super intelligent AI outperforms humans in terms of cognitive abilities, including problem-solving, creativity, and strategic planning.Rapid learning and adaptation: These systems can acquire and process information at an unprecedented speed and scale, leading to rapid advancements in various fields.Ethical and existential risks: The development of Super intelligent AI raises concerns about its potential impact on society, including risks related to control, alignment with human values, and existential threats to humanity.Examples: Super intelligent AI remains theoretical, and no concrete examples exist as of now. However, researchers are actively exploring the implications and challenges associated with its development.Based on Functionalities : 1. Reactive machines: These are the most basic type of AI and are purely reactive as the name suggests. They neither can form memories nor can use past experiences to form decisions. An example would be IBM's Deep Blue chess-playing supercomputer which is mentioned above. Deep Blue beat the international grandmaster Garry Kasparov in 1997. It can choose the most optimal of the chess moves and beat the opponent. Apart from a rarely used chess-specific rule against repeating the same move three times, Deep Blue ignores everything before the present moment, thus not storing any memories. This type of AI just perceives the world, the chess game in the case of Deep Blue, and acts on it.2. Limited memory: These machines can look into the past. Not the ability to predict what happened in the past, but the usage of memories to form decisions. A common example could include self-driving cars. For example, they observe other cars' speed and directions and act accordingly. This requires monitoring of how a car is driven for a specific amount of time. Just like how humans observe and learn the specifics. These pieces of information are not stored in the library of experiences of the machines, unlike humans. We humans automatically save everything in the library of our experiences and can learn from it, but limited memory machines can't.3. Theory of mind: These are types of machines that can understand that people have beliefs, emotions, expectations, etc., and have some of their own. A "theory of mind" machine can think emotionally and can respond with emotions. Even though there are close examples of this kind of AI like Sophia, the research is not complete yet. In other words, these machines have a notion of not just the world, but also the existing entities of the world, like human beings, animals, etc. These machines will be capable of answering simple "what if" questions. They'll have a sense of empathy.4. Self-Awareness: These types of machines can be called human equivalents. Of course, no such machines exist and the invention of them would be a milestone in the field of AI. These basically will have a sense of consciousness of who they are. The sense of "I" or "me". Here's a basic example of the difference between "theory of mind" and "self-awareness" AI. The feeling of I want to play is different from the feeling of I know I want to play. In the latter, if you notice, there is a sense of consciousness and is a characteristic of a self-aware machine, while the former feeling is a characteristic of a theory-of-mind machine. Self-aware machines will have the ability to predict others' feelings. Let's hope the invention is not so far away.How Does AI Work ?Artificial Intelligence (AI) uses a wide range of techniques and approaches that enable machines to simulate human-like intelligence and perform tasks that traditionally require human assistance. AI systems work through a combination of algorithms, data, and computational power. Here's an overview of how AI works:Data Collection: AI systems rely on vast amounts of data to learn and make decisions. Data can be collected from various sources, including sensors, digital devices, databases, the internet, and user interactions. The quality and quantity of data are crucial for training accurate and reliable AI models.Data Pre-processing: Once data is collected, it needs to be pre-processed to ensure it's clean, structured, and suitable for analysis. This pre-processing stage may involve tasks such as cleaning noisy data, handling missing values, standardizing formats, and encoding categorical variables.Algorithm Selection: AI algorithms are chosen based on the specific task or problem the AI system aims to solve. Different algorithms are suited for different types of tasks, such as classification, regression, clustering, and pattern recognition. Common AI algorithms include neural networks, decision trees, support vector machines, and k-nearest neighbours.Model Training: In the training phase, AI models are fed with labelled data (supervised learning) or unlabelled data (unsupervised learning) to learn patterns and relationships. During training, the model adjusts its parameters iteratively to minimize errors and improve its performance on the given task. This process involves optimization techniques like gradient descent and backpropagation in neural networks.Model Evaluation: After training, the AI model is evaluated using separate validation data to assess its performance and generalization ability. Performance metrics such as accuracy, precision, recall, F1-score, and area under the curve (AUC) are used to quantify the model's effectiveness in making predictions or decisions.Model Deployment: Once the AI model meets the desired performance criteria, it can be deployed into production environments to perform real-world tasks. Deployment involves integrating the model into existing systems, such as mobile apps, web services, or embedded devices, to provide AI-driven functionalities.Continuous Learning and Improvement: AI systems can adapt and improve over time through continuous learning. They can be updated with new data and retrained periodically to stay relevant and accurate in dynamic environments. Techniques like online learning, transfer learning, and reinforcement learning enable AI models to learn from new experiences and feedback.Inference and Decision-Making: During inference, the trained AI model applies its learned knowledge to make predictions or decisions on new, unseen data. Inference involves feeding input data into the model and obtaining output predictions or classifications based on the model's learned patterns and representations.Overall, AI systems work by leveraging data, algorithms, and computational power to learn from experience, make decisions, and perform tasks autonomously. The specific workings of an AI system depend on its architecture, algorithms, and the nature of the tasks it's designed to accomplish.Applications of AIArtificial Intelligence (AI) has a wide range of applications and has been adopted in many industries to improve efficiency, accuracy, and productivity. Some of the most common uses of AI are:Healthcare: AI is used in healthcare for various purposes such as diagnosing diseases, predicting patient outcomes, drug discovery, and personalized treatment plans.Finance: AI is used in the finance industry for tasks such as credit scoring, fraud detection, portfolio management, and financial forecasting.Retail: AI is used in the retail industry for applications such as customer service, demand forecasting, and personalized marketing.Manufacturing: AI is used in manufacturing for tasks such as quality control, predictive maintenance, and supply chain optimization.Transportation: AI is used in transportation for optimizing routes, improving traffic flow, and reducing fuel consumption.Education: AI is used in education for personalizing learning experiences, improving student engagement, and providing educational resources.Marketing: AI is used in marketing for tasks such as customer segmentation, personalized recommendations, and real-time audience analysis.Gaming: AI is used in gaming for developing intelligent game characters and providing personalized gaming experiences.Security: AI is used in security for tasks such as facial recognition, intrusion detection, and cyber threat analysis.Natural Language Processing (NLP): AI is used in NLP for tasks such as speech recognition, machine translation, and sentiment analysis.These are some of the most common uses of AI, but the Applications of AI are constantly expanding, evolving, and it is likely that new uses will emerge in the future.Future of AIThe future of AI is likely to involve continued advancements in machine learning, natural language processing, and computer vision, which will enable AI systems to become increasingly capable and integrated into a wide range of applications and industries. Some potential areas of growth for AI include healthcare, finance, transportation, and customer service. Additionally, there may be increasing use of AI in more sensitive areas such as decision making in criminal justice, hiring and education, which will raise ethical and societal implications that need to be addressed. It is also expected that there will be more research and development in areas such as explainable AI, trustworthy AI and AI safety to ensure that AI systems are transparent, reliable and safe to use.   The time may have finally come for artificial intelligence (AI) after periods of hype followed by several "AI winters" over the past 60 years. AI now powers so many real-world applications, ranging from facial recognition to language translators and assistants like Siri and Alexa, that we barely notice it. Along with these consumer applications, companies across sectors are increasingly harnessing AI's power in their operations. Embracing AI promises considerable benefits for businesses and economies through its contributions to productivity growth and innovation. At the same time, AI's impact on work is likely to be profound. Some occupations as well as demand for some skills will decline, while others grow and many change as people work alongside ever-evolving and increasingly capable machines. This briefing pulls together various strands of research by the McKinsey Global Institute into AI technologies and their uses, limitations, and impact. It was compiled for the Tallinn Digital Summit that took place in October 2018. The briefing concludes with a set of issues that policy makers and business leaders will need to address to soften the disruptive transitions likely to accompany its adoption. AI's time may have finally come, but more progress is needed The term "artificial intelligence" was popularized at a conference at Dartmouth College in the United States in 1956 that brought together researchers on a broad range of topics, from language simulation to learning machines. Despite periods of significant scientific advances in the six decades since, AI has often failed to live up to the hype that surrounded it. Decades were spent trying to describe human intelligence precisely, and the progress made did not deliver on the earlier excitement. Since the late 1990s, however, technological progress has gathered pace, especially in the past decade. Machine-learning algorithms have progressed, especially through the development of deep learning and reinforcement-learning techniques based on neural networks. Several other factors have contributed to the recent progress. Exponentially more computing capacity has become available to train larger and more complex models; this has come through silicon-level innovation including the use of graphics processing units and tensor processing units, with more on the way. This capacity is being aggregated in hyperscale clusters, increasingly being made accessible to users through the cloud. Another key factor is the massive amounts of data being generated and now available to train AI algorithms. Some of the progress in AI has been the result of system-level innovations. Autonomous vehicles are a good illustration of this: they take advantage of innovations in sensors, LIDAR, machine vision, mapping and satellite technology, navigation algorithms, and robotics all brought together in integrated systems. Despite the progress, many hard problems remain that will require more scientific breakthroughs. So far, most of the progress has been in what is often referred to as "narrow AI"—where machine-learning techniques are being developed to solve specific problems, for example, in natural language processing. The harder issues are in what is usually referred to as "artificial general intelligence," where the challenge is to develop AI that can tackle general problems in much the same way that humans can. Many researchers consider this to be decades away from becoming reality. Deep learning and machine-learning techniques are driving AI Much of the recent excitement about AI has been the result of advances in the field known as deep learning, a set of techniques to implement machine learning that is based on artificial neural networks. These AI systems loosely model the way that neurons interact in the brain. Neural networks have many ("deep") layers of simulated interconnected neurons, hence the term "deep learning." Whereas earlier neural networks had only three to five layers and dozens of neurons, deep learning networks can have ten or more layers, with simulated neurons numbering in the millions. There are several types of machine learning: supervised learning, unsupervised learning, and reinforcement learning, with each best suited to certain use cases. Most current practical examples of AI are applications of supervised learning. In supervised learning, often used when labeled data are available and the preferred output variables are known, training data are used to help a system learn the relationship of given inputs to a given output—for example, to recognize objects in an image or to transcribe human speech. Unsupervised learning is a set of techniques used without labeled training data—for example, to detect clusters or patterns, such as images of buildings that have similar architectural styles, in a set of existing data. In reinforcement learning, systems are trained by receiving virtual "rewards" or "punishments," often through a scoring system, essentially learning by trial and error. Through ongoing work, these techniques are evolving. Limitations remain, although new techniques show promise AI still faces many practical challenges, though new techniques are emerging to address them. Machine learning can require large amounts of human effort to label the training data necessary for supervised learning. In-stream supervision, in which data can be labeled in the course of natural usage, and other techniques could help alleviate this issue. Obtaining data sets that are sufficiently large and comprehensive to be used for training—for example, creating or obtaining sufficient clinical-trial data to predict healthcare treatment outcomes more accurately—is also often challenging. The "black box" complexity of deep learning techniques creates the challenge of "explainability," or showing which factors led to a decision or prediction, and how. This is particularly important in applications where trust matters and predictions carry societal implications, as in criminal justice applications or financial lending. Some nascent approaches, including local interpretable model-agnostic explanations (LIME), aim to increase model transparency. Another challenge is that of building generalized learning techniques, since AI techniques continue to have difficulties in carrying their experiences from one set of circumstances to another. Transfer learning, in which an AI model is trained to accomplish a certain task and then quickly applies that learning to a similar but distinct activity, is one promising response to this challenge. Businesses stand to benefit from AI While AI is increasingly pervasive in consumer applications, businesses are beginning to adopt it across their operations, at times with striking results. AI's potential cuts across industries and functions AI can be used to improve business performance in areas including predictive maintenance, where deep learning's ability to analyze large amounts of high-dimensional data from audio and images can effectively detect anomalies in factory assembly lines or aircraft engines. In logistics, AI can optimize routing of delivery traffic, improving fuel efficiency and reducing delivery times. In customer service management, AI has become a valuable tool in call centers, thanks to improved speech recognition. In sales, combining customer demographic and past transaction data with social media monitoring can help generate individualized "next product to buy" recommendations, which many retailers now use routinely. Such practical AI use cases and applications can be found across all sectors of the economy and multiple business functions, from marketing to supply chain operations. In many of these use cases, deep learning techniques primarily add value by improving on traditional analytics techniques. Our analysis of more than 400 use cases across 19 industries and nine business functions found that AI improved on traditional analytics techniques in 69 percent of potential use cases (Exhibit 1). In only 16 percent of AI use cases did we find a "greenfield" AI solution that was applicable where other analytics methods would not be effective. Our research estimated that deep learning techniques based on artificial neural networks could generate as much as 40 percent of the total potential value that all analytics techniques could provide by 2030. Further, we estimate that several of the deep learning techniques could enable up to $6 trillion in value annually. So far, adoption is uneven across companies and sectors Although many organizations have begun to adopt AI, the pace and extent of adoption has been uneven. Nearly half of respondents in a 2018 McKinsey survey on AI adoption say their companies have embedded at least one AI capability in their business processes, and another 30 percent are piloting AI. Still, only 21 percent say their organizations have embedded AI in several parts of the business, and barely 3 percent of large firms have integrated AI across their full enterprise workflows. Other surveys show that early AI adopters tend to think about these technologies more expansively, to grow their markets or increase market share, while companies with less experience focus more narrowly on reducing costs. Highly digitized companies tend to invest more in AI and derive greater value from its use. At the sector level, the gap between digitized early adopters and others is widening. Sectors highly ranked in MGI's Industry Digitization Index, such as high tech and telecommunications, and financial services are leading AI adopters and have the most ambitious AI investment plans (Exhibit 2). As these firms expand AI adoption and acquire more data and AI capabilities, laggards may find it harder to catch up. Several challenges to adoption persist Many companies and sectors lag in AI adoption. Developing an AI strategy with clearly defined benefits, finding talent with the appropriate skill sets, overcoming functional silos that constrain end-to-end deployment, and lacking ownership and commitment to AI on the part of leaders are among the barriers to adoption most often cited by executives. On the strategy side, companies will need to develop an enterprise-wide view of compelling AI opportunities, potentially transforming parts of their current business processes. Organizations will need robust data capture and governance processes as well as modern digital capabilities, and be able to build or access the requisite infrastructure. Even more challenging will be overcoming the "last mile" problem of making sure that the superior insights provided by AI are inculcated into the behavior of the people and processes of an enterprise. On the talent front, much of the construction and optimization of deep neural networks remains an art requiring real expertise. Demand for these skills far outstrips supply; according to some estimates, fewer than 10,000 people have the skills necessary to tackle serious AI problems, and competition for them is fierce. Companies considering the option of building their own AI solutions will need to consider whether they have the capacity to attract and retain workers with these specialized skills. Economies also stand to benefit from AI, through increased productivity and innovation Deployment of AI and automation technologies can do much to lift the global economy and increase global prosperity. At a time of aging and falling birth rates, productivity growth becomes critical for long-term economic growth. Even in the near term, productivity growth has been sluggish in developed economies, dropping to an average of 0.5 percent in 2010–14 from 2.4 percent a decade earlier in the United States and major European economies. Much like previous general-purpose technologies, AI has the potential to contribute to productivity growth. AI could contribute to economic impact through a variety of channels The largest economic impacts of AI will likely be on productivity growth through labor market effects including substitution, augmentation, and contributions to labor productivity. Our research suggests that labor substitution could account for less than half of the total benefit. AI will augment human capabilities, freeing up workers to engage in more productive and higher-value tasks, and increase demand for jobs associated with AI technologies. AI can also boost innovation, enabling companies to improve their top line by reaching underserved markets more effectively with existing products, and over the longer term, creating entirely new products and services. AI will also create positive externalities, facilitating more efficient cross-border commerce and enabling expanded use of valuable cross-border data flows. Such increases in economic activity and incomes can be reinvested into the economy, contributing to further growth. The deployment of AI will also bring some negative externalities that could lower, although not eliminate, the positive economic impacts. On the economic front, these include increased competition that shifts market share from nonadopters to front-runners, the costs associated with managing labor market transitions, and potential loss of consumption for citizens during periods of unemployment, as well the transition and implementation costs of deploying AI systems. All in all, these various channels net out to significant positive economic growth, assuming businesses and governments proactively manage the transition. One simulation we conducted using McKinsey survey data suggests that AI adoption could raise global GDP by as much as $13 trillion by 2030, about 1.2 percent additional GDP growth per year. This effect will build up only through time, however, given that most of the implementation costs of AI may be ahead of the revenue potential. The AI readiness of countries varies considerably The leading enablers of potential AI-driven economic growth, such as investment and research activity, digital absorption, connectedness, and labor market structure and flexibility, vary by country. Our research suggests that the ability to innovate and acquire the necessary human capital skills will be among the most important enablers—and that AI competitiveness will likely be an important factor influencing future GDP growth. Countries leading the race to supply AI have unique strengths that set them apart. Scale effects enable more significant investment, and network effects enable these economies to attract the talent needed to make the most of AI. For now, China and the United States are responsible for most AI-related research activities and investment. A second group of countries that includes Germany, Japan, Canada, and the United Kingdom have a history of driving innovation on a major scale and may accelerate the commercialization of AI solutions. Smaller, globally connected economies such as Belgium, Singapore, South Korea, and Sweden also score highly on their ability to foster productive environments where novel business models thrive. Countries in a third group, including but not limited to Brazil, India, Italy, and Malaysia, are in a relatively weaker starting position, but they exhibit comparative strengths in specific areas on which they may be able to build. India, for instance, produces around 1.7 million graduates a year with STEM degrees—more than the total of STEM graduates produced by all G-7 countries. Other countries, with relatively underdeveloped digital infrastructure, innovation and investment capacity, and digital skills, risk falling behind their peers. AI and automation will have a profound impact on work Even as AI and automation bring benefits to business and the economy, major disruptions to work can be expected. About half of current work activities (not jobs) are technically automatable Our analysis of the impact of automation and AI on work shows that certain categories of activities are technically more easily automatable than others. They include physical activities in highly predictable and structured environments, as well as data collection and data processing, which together account for roughly half of the activities that people do across all sectors in most economies. The least susceptible categories include managing others, providing expertise, and interfacing with stakeholders. The density of highly automatable activities varies across occupations, sectors, and, to a lesser extent, countries. Our research finds that about 30 percent of the activities in 60 percent of all occupations could be automated—but that in only about 5 percent of occupations are nearly all activities automatable. In other words, more occupations will be partially automated than wholly automated. Three simultaneous effects on work: Jobs lost, jobs gained, jobs changed The pace at and extent to which automation will be adopted and impact actual jobs will depend on several factors besides technical feasibility. Among these are the cost of deployment and adoption, and the labor market dynamics, including labor supply quantity, quality, and associated wages. The labor factor leads to wide differences across developed and developing economies. The business benefits beyond labor substitution—often involving use of AI for beyond-human capabilities—which contribute to business cases for adoption are another factor. Social norms, social acceptance, and various regulatory factors will also determine the timing. How all these factors play out across sectors and countries will vary, and for countries will largely be driven by labor market dynamics. For example, in advanced economies with relatively high wage levels, such as France, Japan, and the United States, jobs affected by automation could be more than double that in India, as a percentage of the total. Given the interplay of all these factors, it is difficult to make predictions but possible to develop various scenarios. First, on jobs lost: our midpoint adoption scenario for 2016 to 2030 suggests that about 15 percent of the global workforce (400 million workers) could be displaced by automation (Exhibit 3). Second, jobs gained: we developed scenarios for labor demand to 2030 based on anticipated economic growth through productivity and by considering several drivers of demand for work. These included rising incomes, especially in emerging economies, as well as increased spending on healthcare for aging populations, investment in infrastructure and buildings, energy transition spending, and spending on technology development and deployment. The number of jobs gained through these and other catalysts could range from 555 million to 890 million, or 21 to 33 percent of the global workforce. This suggests that the growth in demand for work, barring extreme scenarios, would more than offset the number of jobs lost to automation. However, it is important to note that in many emerging economies with young populations, there will already be a challenging need to provide jobs to workers entering the workforce and that, in developed economies, the approximate balance between jobs lost and those created in our scenarios is also a consequence of aging, and thus fewer people entering the workforce. No less significant are the jobs that will change as machines increasingly complement human labor in the workplace. Jobs will change as a result of the partial automation described above, and jobs changed will affect many more occupations than jobs lost. Skills for workers complemented by machines, as well as work design, will need to adapt to keep up with rapidly evolving and increasingly capable machines. Four workforce transitions will be significant Even if there will be enough work for people in 2030, as most of our scenarios suggest, the transitions that will accompany automation and AI adoption will be significant. First, millions of workers will likely need to change occupations. Some of these shifts will happen within companies and sectors, but many will occur across sectors and even geographies. While occupations requiring physical activities in highly structured environments and in data processing will decline, others that are difficult to automate will grow. These could include managers, teachers, nursing aides, and tech and other professionals, but also gardeners and plumbers, who work in unpredictable physical environments. These changes may not be smooth and could lead to temporary spikes in unemployment (Exhibit 4). Second, workers will need different skills to thrive in the workplace of the future. Demand for social and emotional skills such as communication and empathy will grow almost as fast as demand for many advanced technological skills. Basic digital skills have been increasing in all jobs. Automation will also spur growth in the need for higher cognitive skills, particularly critical thinking, creativity, and complex information processing. Demand for physical and manual skills will decline, but these will remain the single largest category of workforce skills in 2030 in many countries. The pace of skill shifts has been accelerating, and it may lead to excess demand for some skills and excess supply for others. Third, workplaces and workflows will change as more people work alongside machines. As self-checkout machines are introduced in stores, for example, cashiers will shift from scanning merchandise themselves to helping answer questions or troubleshoot the machines. Finally, automation will likely put pressure on average wages in advanced economies. Many of the current middle-wage jobs in advanced economies are dominated by highly automatable activities, in fields such as manufacturing and accounting, which are likely to decline. High-wage jobs will grow significantly, especially for high-skill medical and tech or other professionals. However, a large portion of jobs expected to be created, such as teachers and nursing aides, typically have lower wage structures. In tackling these transitions, many economies, especially in the OECD, start in a hole, given the existing skill shortages and challenged educational systems, as well as the trends toward declining expenditures on on-the-job training and worker transition support. Many economies are already experiencing income inequality and wage polarization. AI will also bring both societal benefits and challenges Alongside the economic benefits and challenges, AI will impact society in a positive way, as it helps tackle societal challenges ranging from health and nutrition to equality and inclusion. However, it is also creating pitfalls that will need to be addressed, including unintended consequences and misuse. AI could help tackle some of society's most pressing challenges By automating routine or unsafe activities and those prone to human error, AI could allow humans to be more productive and to work and live more safely. One study looking at the United States estimates that replacing human drivers with more accurate autonomous vehicles could save thousands of lives per year by reducing accidents. AI can also reduce the need for humans to work in unsafe environments such as offshore oil rigs and coal mines. DARPA, for example, is testing small robots that could be deployed in disaster areas to reduce the need for humans to be put in harm's way. Several AI capabilities are especially relevant. Image classification performed on photos of skin taken via a mobile phone app could evaluate whether moles are cancerous, facilitating early-stage diagnosis for individuals with limited access to dermatologists. Object detection can help visually impaired people navigate and interact with their environment by identifying obstacles such as cars and lamp posts. Natural language processing could be used to track disease outbreaks by monitoring and analyzing text messages in local languages. Our work and that of others has highlighted numerous use cases across many domains where AI could be applied for social good. For these AI-enabled interventions to be effectively applied, several barriers must be overcome. These include the usual challenges of data, computing, and talent availability faced by any organization trying to apply AI, as well as more basic challenges of access, infrastructure, and financial resources that are particularly acute in remote or economically challenged places and communities. AI will need to address societal concerns including unintended consequences, misuse, algorithmic bias, and questions about data privacy In economic terms, difficult questions will need to be addressed about the widening economic gaps across individuals, firms, sectors, and even countries that might emerge as an unintended consequence of deployment. Other areas of concern include the use and misuse of AI. These range from use in surveillance and military applications to use in social media and politics, and where the impact has social consequences such as in criminal justice systems. We must also consider the potential for users with malicious intent, including in areas of cybersecurity. Multiple research efforts are currently under way to identify best practices and address such issues in academic, nonprofit, and private-sector research. Some concerns are directly related to the way algorithms and the data used to train them may introduce new biases or perpetuate and institutionalize existing social and procedural biases. For example, facial recognition models trained on a population of faces corresponding to the demographics of artificial intelligence developers may not reflect the broader population. Data privacy and use of personal information are also critical issues to address if AI is to realize its potential. Europe has led the way in this area with the General Data Protection Regulation, which introduced more stringent consent requirements for data collection, gives users the right to be forgotten and the right to object, and strengthens supervision of organizations that gather, control, and process data, with significant fines for failures to comply. Cybersecurity and "deep fakes" that could manipulate election results or perpetrate large-scale fraud are also a concern. Three priorities for achieving good outcomes The potential benefits of AI to business and the economy, and the way the technology addresses some of the societal challenges, should encourage business leaders and policy makers to embrace and adopt AI. At the same time, the potential challenges to adoption, including workforce impacts, and other social concerns cannot be ignored. Key challenges to be addressed include: The deployment challenge We have an interest in embracing AI, given its likely contributions to business value, economic growth, and social good, at a time when many economies need to boost productivity. Businesses and countries have a strong incentive to keep up with global leaders such as the United States and China. Increased and broad deployment will require accelerating the progress being made on the technical challenges, as well making sure that all potential users have access to AI and can benefit from it. Among measures that may be needed: Investing in and continuing to advance AI research and innovation in a manner that ensures that the benefits can be shared by all. Expanding available data sets, especially in areas where their use would drive wider benefits for the economy and society. Investing in AI-relevant human capital and infrastructure to broaden the talent base capable of creating and executing AI solutions to keep pace with global AI leaders. Encouraging increased AI literacy among business leaders and policy makers to guide informed decision making. Supporting existing digitization efforts that form the foundation for eventual AI deployment for both organizations and countries. The future of work challenge A starting point for addressing the potential disruptive impacts of automation will be to ensure robust economic and productivity growth, which is a prerequisite for job growth and increasing prosperity. Governments will also need to foster business dynamism, since entrepreneurship and more rapid new business formation will not only boost productivity, but also drive job creation. Addressing the issues related to skills, jobs, and wages will require more focused measures. These include: Evolving education systems and learning for a changed workplace by focusing on STEM skills as well as creativity, critical thinking, and lifelong learning. Stepping up private- and public-sector investments in human capital, perhaps aided by incentives and credits analogous to those available for R&D investments. Improving labor market dynamism by supporting better credentialing and matching, as well as enabling diverse forms of work, including the gig economy. Rethinking incomes by considering and experimenting with programs that would provide not only income for work, but also meaning and dignity. Rethinking transition support and safety nets for workers affected, by drawing on best practices from around the world and considering new approaches. The responsible AI challenge AI will not live up to its promise if the public loses confidence in it as a result of privacy violations, bias, or malicious use, or if much of the world comes to blame it for exacerbating inequality. Establishing confidence in its abilities to do good, at the same time as addressing misuses, will be critical. Approaches could include: Strengthening consumer, data, and privacy and security protections. Establishing a generally shared framework and set of principles for the beneficial and safe use of AI. Best practice sharing and ongoing innovation to address issues such as safety, bias, and explainability. Striking the right balance between the business and national competitive race to lead in AI to ensure that the benefits of AI are widely available and shared.