RESEARCH BRIEF 17 • DECEMBER 2020

# ARTIFICIAL INTELLIGENCE AND THE FUTURE OF WORK

THOMAS W. MALONE, DIRECTOR, MIT CENTER FOR COLLECTIVE INTELLIGENCE
PATRICK J. MCGOVERN PROFESSOR OF MANAGEMENT, MIT SLOAN SCHOOL OF MANAGEMENT

DANIELA RUS, DIRECTOR, COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL)
ANDREW & ERNA VITERBI PROFESSOR OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE
MIT TASK FORCE ON THE WORK OF THE FUTURE MEMBER

ROBERT LAUBACHER, ASSOCIATE DIRECTOR, MIT CENTER FOR COLLECTIVE INTELLIGENCE

# Artificial Intelligence and the Future of Work

Thomas W. Malone, Daniela Rus, and Robert Laubacher

## Executive Summary

Artificial intelligence (AI) is one of the most important technologies in the world today. The United States and China compete for dominance in its development (Lee, 2018). CEOs believe it will significantly change the way they do business (Rao et al., 2019). And it has helped companies such as Facebook, Google, and Apple to become among the largest in the world.

But how will this technology affect work in the future? Will it lead to a permanent underclass of people who are no longer employable because their jobs are being done by computers? Will superintelligent computers someday take over the world, finding little use for the humans who created them? Or will robotic servants usher in a golden age of human leisure and prosperity?

In this report, we conclude that recent fears about AI leading to mass unemployment are unlikely to be realized. Instead, we believe that—like all previous labor-saving technologies—AI will enable new industries to emerge, creating more new jobs than are lost to the technology. But we see a significant need for governments and other parts of society to help smooth this transition, especially for the individuals whose old jobs are disrupted and who cannot easily find new ones.

In addition, even though AI is advancing rapidly, we believe that we are at least many decades away from the day when computers have complete, human-level artificial intelligence. For the foreseeable future, therefore, the most promising uses of AI will not involve computers replacing people, but rather, people and computers working together—as "superminds"—to do both cognitive and physical tasks that could not be done before.

## RECOMMENDATIONS

Specifically, we recommend the following actions for the following key stakeholders, which range from schools to businesses and government. In some instances, we call for stakeholders to provide direct assistance to people whose work life is disrupted by AI. In other cases, we call for stakeholders to pay for services delivered by other entities.

- *K-12 schools, colleges, and universities*
  - Ensure that every high school has a computer science teacher with appropriate subject training.
  - Enhance secondary-school programs by expanding current computer literacy curricula to include computational thinking.
  - Build community college programs that include reskilling tracks and online micro-degree offerings matched to the needs of local employers. Offerings should include apprenticeships that provide on-the-job training.
  - Expand post-secondary-school enrollment—through traditional, online, and hybrid programs—to better educate the population as a whole, enhancing not only general cognitive skills but also social skills.
- *Businesses*
  - Focus on applying AI to work with people in achieving new or better outcomes—rather than to replace people.
  - Offer training for employees whose positions will be eliminated or transformed by AI to prepare them for other jobs.
  - Provide appropriate continuing education to decision-makers, developers, and users of AI; this should include training in testing and evaluation practices.
  - Create nontraditional training initiatives focused on enhancing the skills of partners and customers.
- *Worker organizations*
  - To meet the needs of an increasingly dynamic world of work, current worker organizations (such as labor unions and professional associations) or new ones (perhaps called "guilds") should expand their roles to provide benefits previously tied to formal employment (such as insurance and pensions, career development, social connections, a sense of identity, and income security).
- *Community organizations*
  - Use strong local ties and deep understanding of the nature of the specific challenges faced by community members to help workers deal with disruptions caused by AI.
  - Provide aggressive communication programs that inform displaced workers about their reskilling and placement opportunities, endeavoring to inspire them to imagine new career pathways.

- *Government*
  - Increase federal and state government investment in post-secondary education and reskilling/training programs to make the American workforce once again the best educated in the world.
  - Reshape the legal and regulatory framework that governs work to encourage job creation and to adapt to other disruptions created by AI.

## What Is Artificial Intelligence?

There are many ways to define AI, but one simple definition is "intelligence demonstrated by machines" (Wikipedia, 2020). A leading textbook in the field defines it more precisely as machines thinking and/or acting humanly and/or rationally (Russell and Norvig, 2020).

In general, the field of artificial intelligence seeks to advance the science and engineering of intelligence, with the goal of creating machines with human-like characteristics. This includes developing machines with a wide range of human-inspired capabilities, including communication, perception, planning, reasoning, knowledge representation, the ability to move and manipulate objects, and learning. AI approaches problems using tools and techniques from a wide variety of other fields, including probability and statistics, symbolic computation, search and optimization, game theory, information engineering, mathematics, psychology, linguistics, and philosophy.

### BRIEF HISTORY OF AI

Artificial intelligence traces its beginnings to Alan Turing, who in a 1950 paper imagined a machine that could communicate—via an exchange of typed messages—so capably that people conversing with it could not tell whether they were interacting with a machine or another person (Turing 1950). The term artificial intelligence was coined in mid-1955, in a proposal by a group of computer scientists, including Marvin Minsky of MIT, to hold a workshop at Dartmouth College the following summer. The goal of the envisioned workshop was "to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves" (McCarthy et al., 1955). The workshop was held between June and August 1956; many consider the event to mark the birth of the field of AI.

In the years following that seminal workshop, labs were established at MIT, Stanford, and Carnegie Mellon to pursue research in AI. Other labs soon followed. Some took highly focused, narrow approaches. MIT's Minsky, in contrast, advocated combining a broad, multifaceted range of methods—an approach outlined in his influential 1961 paper "Steps Toward Artificial Intelligence" (Minsky, 1961). MIT's Patrick Winston nicely characterized the contributions of two of the pioneers of the field of AI: “Turing told us that we could do this … Minsky told us what to do” (MIT Sloan and CSAIL, 2017, Module 1).

The next several decades saw a broad range of advances in AI. The first wave of breakthroughs relied on such approaches as dividing complex problems into sets of small, constituent tasks that machines could do (Slagle, 1961). A later wave involved encoding specialized human expertise into rules for the machine to follow (Buchanan and Shortliffe, 1984).

By the mid-1980s, however, advances slowed, and the field experienced what came to be called the “AI winter.” Some new discoveries occurred over the next few decades, but the field’s previously confident projections that rapid progress would continue were not realized.

Progress in the field has since accelerated, notably over the past decade or so. This has been enabled by the convergence of advances in three areas: the explosion of data, ever-increasing computing power, and new algorithms (especially algorithms called “neural networks” and “deep learning”). Together, these advances have created a technological tidal wave. The first sign appeared in 2011, when IBM’s Watson program beat the best human players of the TV game show *Jeopardy*. Other advances followed in quick succession. In 2015, for example, Google’s AlphaGo beat a grand master of the game of Go, and in 2017, AlphaGo beat the number one–ranked human player of the game, a feat that had previously been considered impossible, since Go is far more complex than chess.

## SPECIALIZED VS. GENERAL INTELLIGENCE

Many people don’t realize that even the most advanced examples of AI today involve machines that can only perform very specific tasks. This specialized AI, sometimes called narrow or weak AI, is only able to tackle a limited number of problems within a narrow scope. Examples include Google’s AlphaGo, IBM’s Watson, and numerous speech transcription programs.

Narrow AI systems solve problems by working in predefined ways, and they often rely on complex pattern recognition. Typically, they look at vast amounts of data, extract patterns, make predictions, and act based on those predictions. To play Go, for example, the computer can study every game ever recorded and model the likely outcome of every possible move. Narrow AI solutions exist for a wide range of specific problems and can do a lot to improve efficiency and productivity in the work world. For example, in addition to the version of IBM’s Watson system that plays *Jeopardy!*, a substantially different version of the system can help doctors make better decisions.

*Specialized intelligence* can be contrasted with *general intelligence*, which humans possess: the ability to undertake a wide variety of tasks in a broad range of settings (Malone, 2018, 24; Russell & Norvig, 2020). A long-term objective of the AI field is to develop systems that can exhibit artificial general intelligence (AGI, also called strong AI). A recent step toward AGI was taken by OpenAI in 2020, with the introduction of the Generative Pre-trained Transformer (GPT-3), a natural language model capable of performing a wide range of tasks for which it was not explicitly trained (Bussler, 2020). But, as we'll explain below, predictions that general AI is just around the corner should be regarded with caution.

## FIELDS WITHIN AI

AI systems aim to perform complex, problem-solving tasks in a way that is similar to what humans do to solve problems. Efforts include developing and implementing algorithms for playing games, planning and executing movement in space, representing knowledge and reasoning, making decisions, adapting actions, perceiving the world, communicating using natural languages, and learning. Below we summarize the objectives of four subfields of AI with recent advances that impact the workforce (for more on AI and its subfields, see Russell and Norvig, 2020):

- Machine learning
- Robotics
- Computer vision
- Natural language processing

### Machine learning

Arguably the most difficult work in computing is telling the machines, in painstaking detail, exactly what they need to do. This is usually done by professional programmers who write software programs with these very detailed instructions.

Machine learning is an alternative, very powerful approach. With machine learning, human programmers don't need to write detailed instructions for solving every different kind of problem. Instead, they can write very general programs with instructions that enable machines to learn from experience, often by analyzing large amounts of data.

More specifically, machine learning refers to a process that starts with a body of data and then tries to derive rules or procedures to explain the data or predict future data. The function of a machine learning system can be *descriptive*, meaning that the system uses the data to explain what happened; *predictive*, meaning the system uses the data to predict what will happen; or *prescriptive*, meaning the system will use the data to make suggestions about what action to take.

The output of a machine learning system is a model that can be thought of as an algorithm for future computations. The more data the system is presented with, the more refined the model. The quality of the learned model is also dependent on the quality of the data used to train it. If the data is biased, the output of the model will also be biased.

Currently, much research effort is focused on getting robust and trustworthy behavior from machine learning systems. Some of today's greatest successes in the field are due to a technique called "deep learning." Deep learning uses data—often millions of human-labeled examples—to determine the weight (or strength of connection) for each link in special kinds of many-layered "neural networks." These weights are adjusted mathematically so that when the network is presented with new inputs, it will produce the correct outputs. For example, in an image recognition system, some layers of the neural network may be designed to detect individual features of a face, such as eyes, nose, or mouth. Another layer then detects whether these features appear in an arrangement that means they are part of a face (Antonio Torralba in MIT Sloan and CSAIL 2019, Module 1).

Three important categories of machine learning are (a) *supervised learning*, where manually labeled data is used to train the system, (b) *unsupervised learning*, where the system is able to use naturally occurring data without any specially added labels, and (c) *reinforcement learning*, where the objective is to train a system to take the best action by establishing a reward system (e.g., programming the system to maximize points and then rewarding points for optimizing a task). Drawing on these and other techniques, machine learning has been used to create a wide range of successful applications in most industries where data can be collected and labeled, including many areas of scientific discovery.

For example, PathAI is a young company that combines data from a variety of sources to diagnose diseases and identify appropriate therapies. The primary data inputs are images from diagnostic tests, details on the treatment protocols that patients underwent, and the outcomes of treatment. Using these inputs, PathAI's machine learning system makes predictions about which treatments will be best for future patients and what kind of drugs pharmaceutical companies could develop to address the disease being treated (Andrew Beck in MIT Sloan and CSAIL, 2019).

In another application, a large bank took a sample of data from its credit card customers and used a machine learning system to discern which variables could predict payment delinquencies. Banks traditionally use credit scores to make this prediction. But the machine learning system found that another variable—an interruption in direct deposits to a customer's checking account—was far more predictive of delinquencies (Andrew Lo in MIT Sloan and CSAIL 2019, Modules 1 and 5).

### Natural language processing

Since very early in their development, computers have been able to understand highly structured computer programming languages (such as today's C, Python, and Java) and commands such as those used in spreadsheets. But humans typically don't use these highly structured languages; they use "natural languages" like English, Spanish, and Chinese. And it is surprisingly difficult to get computers to understand and generate these unrestricted, natural human languages with anything like the facility of a normal, human five-year-old. The area of AI focused on understanding and generating these human languages is called "natural language processing."

Natural language processing (NLP) includes automatic text parsing and understanding, speech recognition, machine translation between human languages, text generation, text summarization, and question answering.

NLP's roots go back to the 1950s, to Alan Turing's paper (Turing, 1950), which proposed a task centered on automatic text generation and interpretation as a criterion for intelligence: the Turing test. Since then, the NLP research community has proposed many algorithms and systems for understanding human language that address syntactic and semantic aspects of speech and written language and use techniques from statistics, symbolic reasoning, signal processing, and machine learning.

The availability of large bodies of text (such as are available on the Web) has made it possible for NLP to benefit from the most recent advances in machine learning, leading to unprecedented performance for NLP tasks in a wide range of applications. For example, Microsoft's speech recognition system has reduced its error rate to 5.1 percent, matching the error rate of multiple human transcribers in a widely recognized accuracy test (Xiong et al., 2018). Google Translate has made it possible to instantaneously translate text for scientists, financial analysts, and anyone else who needs information written in a foreign language. NLP is impacting many fields, including medicine, finance, and education. One of the most promising applications of NLP is in the creation of chatbots, software applications capable of conducting online conversations with people in specific domains.

Another major application of natural language processing that has emerged over the past decade is its use by attorneys to speed the process of pretrial discovery and to aid contract management during premerger due diligence. Before, junior associates at law firms had to pore over hundreds, even thousands of documents, to find ones that were relevant. NLP software can now analyze scanned versions of the documents and assign a probability to each as to whether it is relevant to the case at hand or not. An attorney will further screen a document when the software cannot definitively ascertain relevance (Remus and Levy, 2017).

Other NLP systems are used in call centers, where they are employed to translate words spoken by customers over the phone line into written text. This text gets parsed by the NLP tool to ascertain the nature of the customer's request and in many cases to deliver a standardized response, either via a recording or through a call center worker, who reads it off a monitor. More complex requests, for which a standardized response isn't appropriate, get referred to human customer service agents (Frank Levy in MIT Sloan and CSAIL 2017, Module 3).

## Robotics

Robotics is the study of robots, machines that move and act independently in the physical world with the purpose of automating physical tasks like those done by humans. Robots are programmable mechanical devices that can exert physical force in the world; they take input from sensors, reason internally about the input to determine actions, and then execute those actions.

Researchers in the field of robotics work to develop the science and engineering of autonomous physical action by designing novel machines capable of interacting with the physical world in an intelligent and controlled way. Tasks include developing robot locomotion and robot manipulation, designing multi-robot systems, and facilitating human-robot interaction. The work requires both novel designs for the physical robots and new algorithms for perception, planning, reasoning, learning, and control.

Robotics has transformed industrial manufacturing and proved helpful in many other applications that have dull, repetitive, dangerous, and dirty tasks. Early deployments of industrial robots date to the 1960s and were successful because the robots operated under very precise conditions, in environments that did not change. Today, through hardware, algorithmic, and software advancements, robots are becoming increasingly more adaptive to operating in unfamiliar and human-centered environments.

Thanks to such advances, robotics is enabling powerful solutions in many fields, including manufacturing, transportation, agriculture, defense, medicine, environmental monitoring, and in-home activities. A particularly promising area of robotic application is autonomous transportation, including both the robotic taxi systems currently pursued by many companies, along with self-driving mobile systems such as trucks, boats, agricultural equipment, golf carts, delivery vehicles, and wheelchairs. Robotics is also contributing to scientific discovery by traveling to places humans cannot go, including outer space, deep oceans, and the interior of active volcanos.

By now, most people are used to seeing robots work in manufacturing plants; they're even familiar with driverless cars. Another application that has had significant impact in recent years is the use of robots in warehouses. For example, the robots built by Kiva Systems, a company acquired by Amazon, have transformed warehouse processes.

Each Kiva robot is a small box on wheels, a bit over a foot tall and 2-by-2½-feet wide. When a customer buys something from Amazon, the robot retrieves the ordered item from a moveable pod. The robot drives under the pod, lifts it, and delivers it to the worker responsible for boxing up the order. Once the worker has removed the item from the pod, the robot returns the pod to its appropriate location.

Human workers used to have to move through the aisles of Amazon's warehouses to get the items they needed to complete an order. Now the robots do the moving. And items no longer need to be stored in a fixed location; items projected to be in high demand can be kept closer to the workers, with other items farther away, thus minimizing travel time for the robots (Mountz, 2011).

### Computer vision

Automating tasks done by the human visual system is the goal of computer vision. Developing computers that can understand the physical world based on visual information such as images and video streams involves developing algorithms for capabilities such as object recognition, scene reconstruction, 3D-scene modeling, vison-based robot control, motion analysis, and image restoration.

Computer vision algorithms take input from single images or from sequences of images, and endeavor to extract, analyze, and understand the salient information in a task-specific way. Computer vision has seen significant advancements thanks to the creation of an open-source repository of millions of labeled images, ImageNet (Deng et al., 2009), which opened the door to using machine learning for object recognition and other visual tasks. Computer vision, powered by machine learning, is providing powerful solutions to applications that necessitate object recognition, including automatic inspection for manufacturing; map-making, localization, and place recognition for autonomous driving; detecting events for visual surveillance; and information extraction for medical diagnosis.

For example, the startup Tellus Labs uses computer vision to analyze satellite images and make predictions about crop yields or the movement of oil tankers and ocean shipping. Its customers include commodity traders.

## What Is Today's AI Frontier?

Despite major recent advances, artificial intelligence is nowhere close to matching the breadth and depth of perception, reasoning, communication, and creativity of people. AI systems remain quite limited in their ability to reason, make decisions, or interact reliably with people and objects in the physical world. And they usually can't act creatively or put together unfamiliar combinations of ideas. They also typically lack social intelligence—though researchers at MIT, including Rosalind Picard and Cynthia Breazeal, are working to build systems that can read human emotion and even mirror it back (see, for example, Picard, 1997).

Current limitations hold AI back from outthinking and outdoing humanity in many domains, and this means that there are significant constraints on how much and in what ways machines will be able to displace workers on the job today.

Other challenges associated with AI stem from some of the qualities inherent to today's systems: they need to be trained with very large data sets; they can be brittle (which means they can fail due to small changes in their inputs); they require massive amounts of computing power; and they are not able to explain how they reach their decisions (in large part because they cannot reason like people). Researchers are working to address these shortcomings, but solutions will not arrive overnight.

### *Data availability*

Deep learning systems are often trained using vast amounts of data. This is why companies sometimes refer to data as "the new oil." But data typically needs to be manually labeled, which can time-consuming and expensive (Metz, 2019). When the data has bias, the system propagates that bias. And poor-quality training data results in poor system performance. The AI community is currently developing new approaches to reduce training needs and assess data quality.

### Robustness

Deep learning systems achieve impressive performance, but their predictions can be brittle (Marcus, 2018). That means that errors can result from carefully constructed but indiscernible perturbations to input data. These carefully constructed examples are sometimes referred to as adversarial examples, because they are specifically created to "fool" the deep learning systems. Among other things, this means that criminals or others could use such examples to gain control of the predictions. And even benign noise can trigger errors. Consider the two images of dogs in Figure 1. The right image is obtained by slightly perturbing the one on the left with additional pixels in the middle. Those small changes are enough to trick the system into incorrectly identifying the righthand picture as an ostrich (Szegedy et al., 2014)

**Figure 1: Lack of Robustness in Image Recognition**

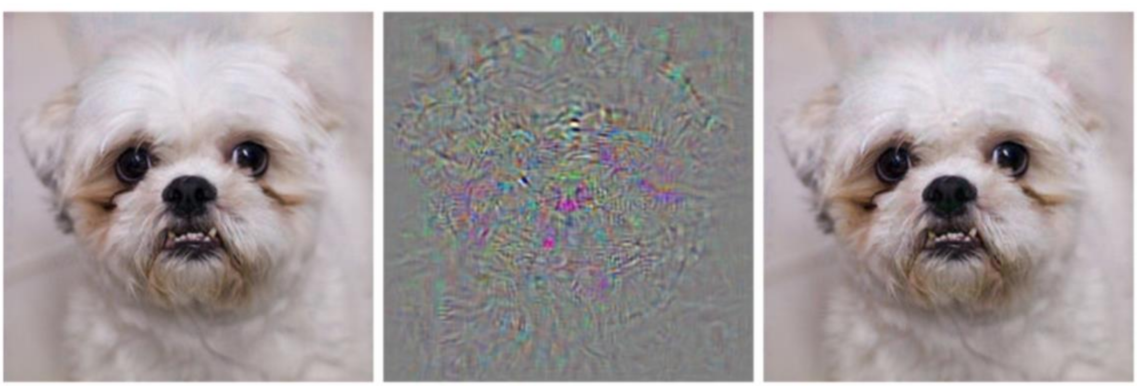

**Dog** **Image Perturbation** **Ostrich**

Because of their lack of robustness, many deep neural nets work merely "most of the time," which is not acceptable in critical applications. Lack of robustness can also reduce trust, since users cannot be assured their system is delivering correct answers. The AI community is developing new training methods and data sensitization techniques to significantly reduce or even eliminate this kind of brittleness.

### Computational load

Recent advances in computing technology have allowed deep learning solutions to scale up to address real-world problems. Such successes have encouraged the creation of larger and larger models, trained on more and more data. Deep learning programs are computationally hungry, and the costs can be staggering (Peng, 2019). For example, the cost of the electricity alone that powered the servers and graphical processing units (GPUs) during training of the GPT-3 language model is estimated at $4.6 million. According to OpenAI, accelerating demand has caused power requirements for training large models to skyrocket 300,000-fold in 2018, and they are currently doubling every 3½ months.

In the mid-1960s, Gordon Moore of Intel projected that the number of transistors on a computer chip would double every 18 months for the foreseeable future, which meant each new generation of chips could deliver twice the processing power for the same price. This projection, which came to be known as Moore's law, has held true ever since and is a major reason the cost of computing has declined significantly in recent decades. But transistors have now become so small they are close to the size of atoms, which means we are approaching the point where Moore's law will no longer hold true. The computing community is working on alternative, algorithmic approaches to make computation more efficient. In the meantime, training deep neural networks will remain a costly process.

### Interpretation

While decisions made by traditional, rule-based software can be traced through the code, it's often hard to know how today's AI systems return their results. They provide only a black-box, input-output functionality. As a result, users cannot learn from the systems, and it's hard to detect abnormal behavior that could pose security or safety risks. This means that rare and unexpected inputs could potentially lead to failures with catastrophic consequences.

Researchers are working to make AI models more transparent by incorporating explanation as a central design goal rather than an afterthought. Approaches for providing greater interpretability include highlighting regions of the input that maximally contribute to output; building additional models that can explain results using natural language; building models that are inherently more interpretable and visualizable; and developing external algorithms that can probe a model's underlying mechanism.

### Reasoning

The “learning” aspects of deep learning systems primarily involve pattern recognition. A system than can translate from English to French has no idea what is being said, and an object classification system that can recognize an image of a beach does so by solely by comparing the pixels of the image to pixels previously labeled by humans as representing a beach—the system has no understanding of what a beach is.

Machine learning systems can go well beyond human capabilities in identifying and recognizing patterns in data. But they cannot do the kind of thinking people do easily: infer symbols, abstractions, and functional connections. For machines to navigate the real world, they will need to move beyond pattern recognition and toward representing and understanding objects and how they relate to one another.

The AI community is working to increase the ability of AI systems to reason—giving machines the power to connect facts, observations, and learned concepts. One promising direction for machine reasoning involves using the mathematical language of logic and finite-state machines to represent concepts and relationships. One system employing this approach learns the rules of the road by observing drivers (Araki et al., 2020).

## AI OPPORTUNITIES AROUND THE CORNER

To deal with the current limitations of AI systems, the research community is pushing the AI frontier forward in of the following areas.

### Reducing computational load

One way to address the computational load of machine learning is to use more efficient models and training data. Promising possibilities for doing this include using a special type of algorithm called “the Transformer” and pruning training data and models.

> *Transformers*. The Transformer is a deep learning approach introduced in 2017 that is primarily used for tasks such as translation and text summarization (Vaswani et al., 2017). Transformers predict what words should come next in a sequence, but they do so without requiring that language data be processed in the order in which it appears in the text. Instead, the processing is done in parallel. This allows the model to consider the relationships between words regardless of how far apart they are and to determine which words and phrases are most important. If the input is a natural language sentence, for example, the Transformer does not need to process the first words before the last ones, and that reduces training time.
>
> Parallelization makes Transformers much more computationally efficient. The GPT-3 Transformer, introduced in the summer of 2020, was trained on roughly 500 billion words and consists of 175

billion parameters. GPT-3 is fueling a revolution in the field of chatbots, allowing, for example, the insurance startup Lemonade to process routine claims wholly automatically.

*Pruning training data and models*. Modern deep learning networks with millions of parameters require large amounts of memory and processing power. To increase the efficiency of deep learning systems, AI researchers are developing ways to reduce the amount of training data required and the size of the networks.

One technique is to curate training sets to capture salient features. One method called few-shot learning uses a very small amount of data. For example, the MNIST database is a large data set of handwritten digits. By carefully engineering and optimizing a sample set for each digit image, researchers have trained the system to recognize digits using only 10 data points of comparison. This system achieves the same accuracy as one trained with the entire MNIST data set of 60,000 examples.

The size of neural networks can be reduced using pruning algorithms that remove redundant parameters. Pruning redundant filters can generate network architectures with 80 percent to 90 percent fewer parameters while retaining the same accuracy as the original network (Liebenwein, 2019).

### New models for interpretability, robustness, and reasoning

A central goal of artificial intelligence research is to design algorithms that represent the world coherently while simultaneously reflecting an understanding of its dynamics. Surprisingly, animals as small as the nematode (which is only a few millimeters long) can do this thanks to their near-optimal nervous system structure and harmonious neural information-processing mechanisms, which enable locomotion, motor control, and navigation capabilities.

A large body of active research focuses on designing methodologies to explain how a large neural network arrives at its output/decision and alleviate some of the limitations of machine learning to function in the real world. For example, neural circuit policies (NCPs) are intended to improve our understanding of how neural information processing systems work in continuous-time environments. NCPs are a new AI approach inspired by nature that involves giving each artificial neuron increased computational capabilities relative to what's provided in contemporary deep learning models (Lechner, 2020)—replacing the simple thresholding function in today's models with a differential equation. The result is a significantly smaller network that is more interpretable. This approach has been applied to keep autonomous vehicles in their lanes using small networks of neural models inspired by the brain that comprise just 19 neurons. These NCPs have outperformed state-of-the-art deep learning models with more than 100,000 nodes.

Other new models for machine learning address the lack of robustness demonstrated in Figure 1 by measuring how certain pixels in training images can influence the system's predictions (Madry et al., 2017). The models can be tweaked to focus on pixel clusters closely correlated with identifiable features—such as detecting an animal's paws, ears, and tail, yielding models capable of symbolic reasoning with increased robustness. Significant efforts are also directed at understanding the origins of common sense in humans and how this understanding can be captured in machines and mapped into more extensive algorithms for causal inference and symbolic reasoning (Lake et al., 2017; Mao et al., 2019; Smith et al., 2020).

### Reducing problems with data availability

Two important ways to reduce the need for manually labeled data are unsupervised learning and reinforcement learning. In addition, federated learning can enable the use of private data in training models.

*Unsupervised learning*. Most AI applications today use supervised learning. The term "supervised" means humans must manually label the data used to train the network. This is a labor-intensive task often performed through services such as Amazon Mechanical Turk or companies such as iMerit. In contrast, unsupervised learning (also called self-supervised learning) does not need labeled data. Instead, the algorithms cluster data and discover natural groupings. If the data is an unlabeled mix of images of animals and cars, for example, an unsupervised system can determine common features and group the cars and animals into separate categories. So far, this approach has been used effectively in natural-language applications and robotics.

*Reinforcement learning*. Reinforcement learning (RL) allows robots to learn complex behaviors through interaction with the world. If successful, this approach could enable deployment of systems that operate much more autonomously. In RL, the robot explores options for performing a given task, determines its success according to a specific reward metric, and learns from experience.

To date, applying reinforcement learning has proven challenging due to the complexity of real-world environments; the number of options robots can pursue is huge. Researchers are therefore seeking to develop simpler representations of the environment and create methods that will allow robots to explore options more efficiently. If these advances are achieved, RL could power robotic systems that truly operate on their own.

*Federated learning*. Most real-world deployments of AI involve the use of valuable and sensitive data. Ensuring the data stays private is a big challenge. Today, a deep learning model is typically trained on a large data corpus stored in a single repository. If the training data includes medical records, personal photos, correspondence, records of user behavior, or other confidential information, having the data aggregated in one place creates the risk that a single breach could expose sensitive information.

To address this concern, federated learning leaves the data distributed across devices and servers close to where the data originated, instead of bringing all the data to a single place. Versions of the model get sent to each device that has training data, and they are trained locally on each subset of data. The resulting model parameters, but not the data itself, are then sent back to the cloud. The "minimodels" are all aggregated to create a unified model that works just as though it had been trained on a single data set.

### Artificial general intelligence

An obvious longer-term question is: How soon will these and other approaches overcome the limitations and challenges seen at today's AI frontier? When, for example, will we have human-level artificial intelligence (also known as artificial general intelligence)? One review of predictions made about AI since the field began showed that the average prediction by both experts and nonexperts for when human-level general AI would be available has been about 20 years in the future (Armstrong and Sotola, 2015). In other words, human-level AI has seemed 20 years away for the last 60 years. While it's theoretically possible that this time the prediction could be right, we believe that is very unlikely. Barring some kind of major societal disaster, we believe it's very likely that human-level AI will be achieved someday, but that day is probably still many decades in the future.

## How Will Work Change with AI?

Even if human-level general AI is still far in the future, AI is likely to have great economic importance in the meantime because it will steadily increase the number of tasks that computers can do faster, more cheaply, or better than people. The question is, how can we predict what these tasks will be?

### WHAT TASKS WILL COMPUTERS DO?

One simple way of predicting what tasks computers will do in the future is to take note of what today's computers are often better at doing than people now: routine data processing (e.g., summing up a retail store's transactions) and predictable physical work (e.g., assembly line production).

Figure 2 shows another simple framework, this one designed to predict where machine learning can be useful. Since most recent advances in AI are based on various forms of machine learning, this framework applies to many kinds of current AI.

**Figure 2: Framework for Predicting where Machine Learning will be useful**

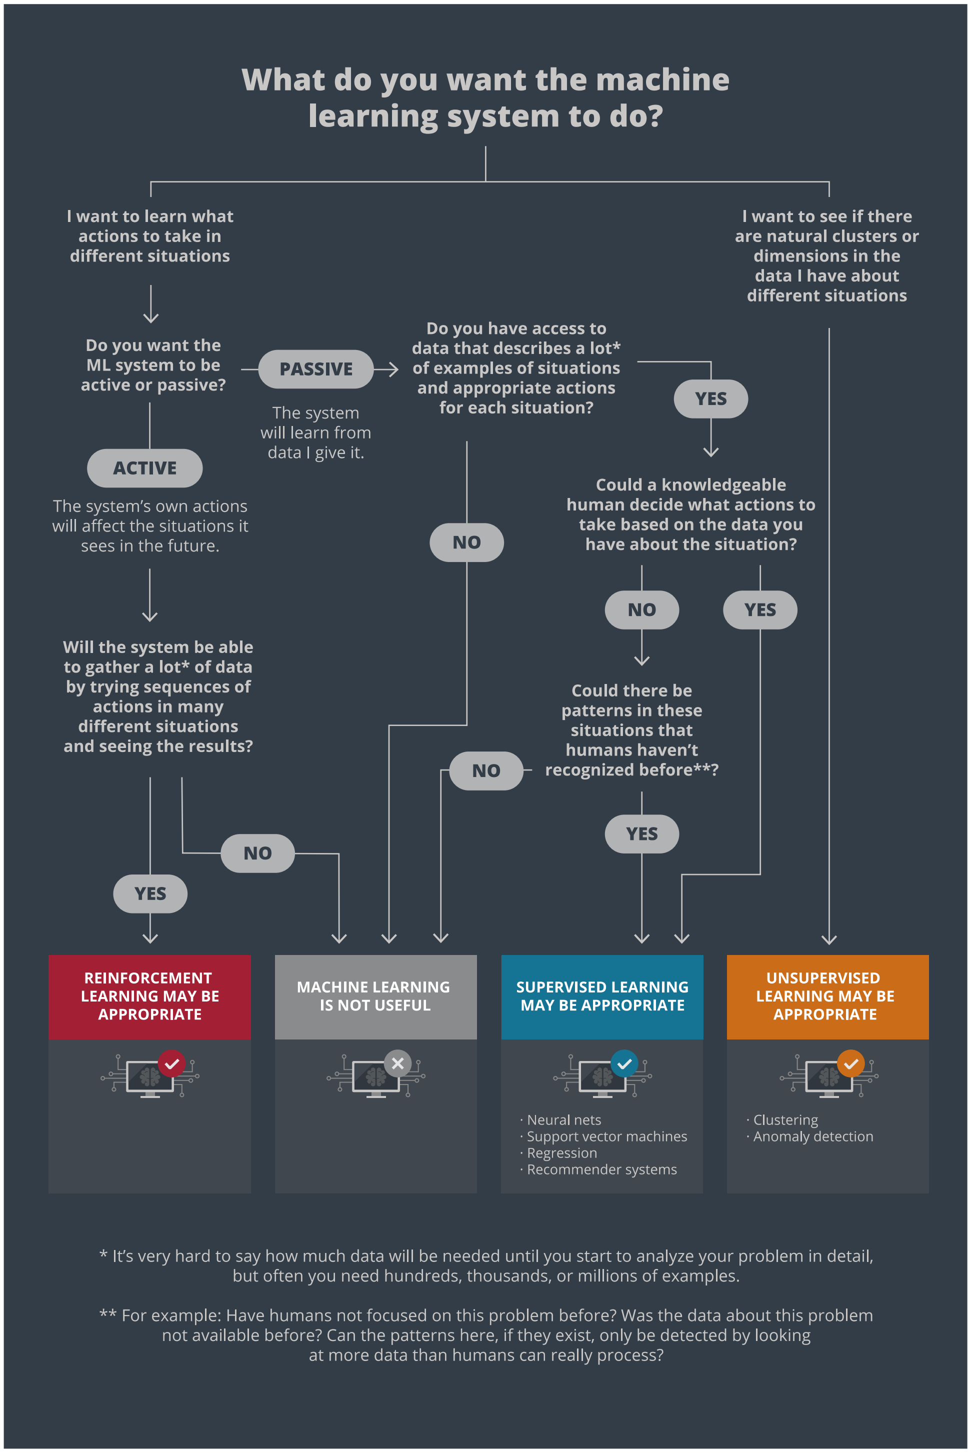


Source: Thomas Malone in MIT Sloan and CSAIL 2019, Module 2

As this figure shows, two important factors are usually needed for today's machine learning algorithms to be useful. First, as noted earlier, these algorithms usually require large amounts of data, often hundreds, thousands, or millions of examples. When such data is available—such as from recordings of previous conversations with customers or from sensor logs of previous machine operations—machine learning can be very useful. For instance, the success of Google Translate was enabled by the availability of vast amounts of textual data in many languages on the web.

Second, machine learning only works if there is an actual pattern to be recognized. For example, machines can often make credit-risk assessments as well as people can when using information such as a borrower's current assets, income, and debt. More interestingly, machine learning programs can sometimes be trained to recognize patterns in pools of data that are far larger than humans would ever be able to analyze. For instance, online retailers such as Amazon use AI techniques to detect subtle patterns in vast amounts of data related to which kinds of customers buy which kinds of products (Hardesty, 2019). Still, no amount of machine learning is likely to be able to predict the exact path an individual shopper will take through a shopping mall.

### Difficulties in predicting when machines can perform better than people

In general, it's very hard to specify in any simple, precise way what machines will or won't be able to do. The above guidelines simply provide ways to make rough judgments about what specific tasks can be done by today's AI; many subtle details can also affect whether today's AI is actually appropriate for a given task.

Unfortunately, even though it's easy to assert a tautology such as "computers can do tasks that are precisely specifiable in ways computers can execute," it's very difficult to provide general rules for what is and isn't precisely specifiable in this way. And it's even harder to predict the effects of future progress in AI (Darwiche, 2018).

Many people, for instance, make statements such as "computers will never be really creative" or "computers will never be able to understand human emotions," but we don't know of any such categorical statements that we believe are accurate. In the distant future, computers may be able to do all these things as well as people do. And in the meantime, computers will be able to do more and more types of things every year. For example, computers have already been used to generate strikingly artistic images (Gayford, 2016), including a painting that sold at auction for more than $400,000 (Vincent, 2018). They have also proved able to read human emotions in certain situations (Picard, 1997).

One useful way of understanding the difficulty in predicting what machines will be able to do in the future is to recognize a key difference between classical AI and modern machine learning (Darwiche, 2018; Marcus and Davis, 2020b). Classical AI (especially common in the 1950s to 1980s) uses an approach that can be described as "represent and reason." This approach explicitly represents the knowledge needed (e.g., "birds lay eggs" and "chickens are birds") and then reasons logically using this knowledge (e.g., concluding that "chickens lay eggs").

Machine learning, on the other hand, uses a "function fitting" approach. It tries to find the parameters for complex mathematical functions that allow the system to predict the input-output relationships observed in the training data. For example, a machine learning model with millions of parameters might be able to predict which configurations of pixels in an image correspond to pictures of cats and which correspond to pictures of dogs.

The problem is that even though both of these approaches can do useful things, neither can yet "think" in the ways humans do. The classical approach is closer to how humans think consciously, but it has not yet been successful at replicating many kinds of real-world human cognitive activities. The machine learning approach may be closer to unconscious human mental processes (such as visual perception), but it has difficulties related to interpretability and reasoning, as described earlier in this report.

One promising direction is to combine the classical and more recent approaches (e.g., Lake et al., 2017; Mao et al., 2019), but no one knows the exact limitations of either approach separately or what might be possible when the two are combined.

### Frontier capabilities

Even though it's hard to predict precisely what tasks computers will be able to do, it is still possible to outline some promising possibilities. In the foreseeable future, for example, we see important opportunities for AI in tasks such as the following:

- *Sensing*. Today's AI systems are already good at analyzing many kinds of images, sounds, motions, and other inputs. For example, AI systems are now sometimes better than human physicians at interpreting X-rays and similar medical images (Rajpurkar et al., 2018). Amazon Go stores allow people to pay for their purchases without ever standing in a checkout line, because cameras and sensors track what items they take and charge them when they leave (Palmer, 2020). Airbus is using various kinds of maintenance records, motion sensor logs, and other data to predict when preventive maintenance is needed on their aircraft (Matthew Evans in MIT Sloan and CSAIL, 2019, Module 1). We believe these examples just hint at the many innovative ways AI can be used to sense the world.
- *Deciding*. For decisions where the inputs and the desired outcomes can all be codified precisely, and where substantial data about relevant past cases is available, machines can usually make decisions at least as well as and often better than people (Kahnemann, 2011, chapter 21). For instance, many credit risk decisions are now done by machines (Andrew Lo in MIT Sloan and CSAIL,

2019, Module 5), and as large databases of medical cases becomes available, many kinds of automated medical diagnoses are likely to become feasible as well (e.g., Yala et al 2017).

- *Creating*. Many kinds of creation involve patterns that machines can learn. For instance, machines have successfully generated news articles about Little League baseball games from simple statistics about the game (Wright, 2015), and it seems likely that they will soon be able to generate customized versions of many legal contracts, sales proposals, and other documents. The GPT-3 system, for example, can flexibly generate remarkably human-like text. Even though it doesn't really "understand" anything it generates, it may substantially enhance other kinds of text-generation software (Brown et al., 2020, Marcus and Davis, 2020a; Bussler, 2020). And, in the physical world, machines are beginning to be useful in generating detailed, innovative designs for objects such as heat exchangers (McAfee & Brynjolfsson, 2017, 111).

For tasks like these, machines are likely to perform work currently done by people. But for the foreseeable future, human work is likely to become even more valuable in tasks that humans usually do better than machines (McKinsey, 2017 and Malone, 2018, 273-281), including those that require

- *social skills* (such as providing emotional support to patients diagnosed with cancer),
- *unpredictable physical skills* (such as repairing plumbing in the complex environment under a kitchen sink),
- *common sense* (such as knowing that someone calling from out of town is not likely to want a local restaurant reservation in five minutes), or
- the *general intelligence* needed to deal with new, non-routine situations (such as what to do when you see flames from a forest fire across the freeway ahead of you).

## NOT PEOPLE OR COMPUTERS, BUT PEOPLE AND COMPUTERS

As these examples illustrate, reframing the question of AI and the future of work around activities suggests that a useful strategy is to begin with tasks that comprise a job and imagine the computers doing the ones they can do best and people doing the ones they can do best. Taking such an approach means thinking less about people OR computers and more about people AND computers.

Adopting this perspective shifts the emphasis from what intelligent computers can do to what intelligent human-computer groups can do. Another term for these human-computer groups is *superminds*, which Malone (2018) defines as "groups of individuals acting together in ways that seem intelligent." Such combinations of many individual minds—often including both people and computers—can do things that couldn't be done by people alone or computers alone. In other words, a supermind—like a superorganism such as a beehive, ant colony, or coral reef—can do many things that its individual members can't.

By focusing on human-computer groups—superminds—we can move away from thinking of AI as a tool for replacing humans by automating tasks, to thinking of AI as a tool for augmenting humans by collaborating with them more effectively. As we've just seen, AI systems are better than humans at some tasks such as crunching numbers, finding patterns, and remembering information. Humans are better than AI systems at tasks that require general intelligence—including non-routine reasoning and defining abstractions—and interpersonal and physical skills that machines haven't yet mastered. By working together, AI systems and humans can augment and complement each other's skills.

The possibilities here go far beyond what most people usually think of when they hear a phrase like "putting humans in the loop." Instead of AI technologies just being tools to augment individual humans, we believe that many of their most important uses will occur in the context of groups of humans. As the Internet has already demonstrated, another very important use of information technology—in addition to AI—will be providing hyperconnectivity: connecting people to other people, and often to computers, at much larger scales and in rich new ways that were never possible before (Malone, 2018, 3).

Hyperconnectivity already enables people to communicate across the entire world, practically for free, and has transformed how we work and live. It has also had downsides, such as the spread of misinformation and extremist beliefs on social media. Some researchers are now seeking to combat the problems caused by hyperconnectivity using hyperconnectivity (Pennycook and Rand, 2019).

In hyperconnected groups, artificially intelligent computers will sometimes be *tools* or *assistants* to individual humans, but they will also sometimes act as *peers* or *managers* for humans. More interestingly, computers will increasingly work with humans—buying and selling in markets or interacting in other ways. These arrangements will enable computers to use their specialized intelligence to do what they do best while people use their general intelligence and other skills to do what they do best. Together, these human-computer groups or superminds will be able to do things that were never possible before. That's why we need to move from thinking about *putting humans in the loop to putting computers in the group* (Malone, 2018, 50-57, 75).

## DESIGNING ORGANIZATIONS USING A SUPERMIND PERSPECTIVE

In implementing AI and other information technology tools today, many organizations simply look for ways that technology can substitute for human labor while continuing to do the same activities they did in the past. Examples include automated menus in telephone customer-service systems and self-service check-out kiosks in stores.

The economists Daron Acemoglu and Pascual Restrepo call these kinds of systems "so-so technologies" (Acemoglu and Restrepo, 2019). They don't increase productivity much—in many instances, they just take part of the work formerly done by employees and pass it off to customers—and they also often lead to reductions in quality. Thus, focusing on AI primarily as a means of cost-cutting is not a panacea.

Better outcomes can often be achieved by focusing on creative new ways to configure groups of people and computers to accomplish valuable tasks (Guszcza and Schwartz, 2020). We believe there is substantial opportunity—both in research and in practice—for shifting focus in this way from designing computers that replace people to designing human-computer groups that are better than anything we ever had before.

For example, MIT's Center for Collective Intelligence is developing a new approach called *supermind design* that includes systematic techniques for triggering creative insights about how to configure human-computer groups (Malone, 2018; MIT CCI, 2020). In one case, this approach helped a large pharmaceutical company design a novel system for detecting and treating adult depression in Japan using AI, social media, and human facilitators (Ammarito et al., 2019; Laubacher et al., 2020).

### A VISION FOR AN AI-ENABLED FUTURE

AI systems are already partnering with humans in industrial and domestic settings. They work side by side with people in factories and operating rooms. Businesses use AI to understand what customers are saying on social media in near-real time, and systems automatically alert companies about impending supply chain bottlenecks. The credit division of Ford has been testing new software that uses machine learning to help underwriters better assess the loan applications they receive, and the developers of this system believe it will make a big difference for people without credit histories who today struggle to get loans. Speech recognition systems are approaching the accuracy of human transcribers and are making instantaneous translations possible. These advances promise to greatly improve our ability to give clear instructions to people and machines and to communicate easily with those who speak other languages.

Positive impacts can also be expected for manufacturing. AI support for new design and simulation tools can speed the validation of new product designs. Sensors on manufacturing lines can feed data into real-time systems that can adjust the machines to minimize defects. Robots can manage dangerous work and help workers avoid safety risks. Computational and AI support for fabrication will enable customization. The potential changes are much bigger than marginal improvements. In the future, we may be working with entirely new molecules and materials. Businesses such as Cambridge-based Kebotix are already using AI and robotics to explore, discover, and produce new materials that can help address such massive challenges as climate change and water pollution (Kebotix, 2018). Given this potential, McKinsey estimates that the United States could boost annual manufacturing-value-added by up to $530 billion over current trends by 2025—and add up to 2.4 million jobs (Ramaswamy et al., 2017).

In a few years, AI will touch even more areas of our lives. Though there has been turbulence along the way, the natural evolution of human society over the centuries has generally been to develop and adopt tools that have made life better. AI is another new tool, and our goal should be to give people ways to use this new tool effectively. But for AI to reach its full potential, we'll need to think hard about what actions can ensure that the benefits it provides are broadly distributed.

## Policy Implications

In the half decade after 2010, a series of influential academic books and articles noted that information technology—artificial intelligence in particular—could supplant large numbers of human workers in the coming decades (Brynjolfsson and McAfee, 2011 and 2014; Frey and Osborne, 2013; Ford, 2013 and 2015; for a review, see Wolters, 2020). These studies generated a plethora of articles by journalists on this theme. The studies and articles emphasized the potential for AI to substitute for human cognitive effort, just as machines substituted for physical labor during the First Industrial Revolution.

Many of the early 2010s predictions presented the risk posed to the workforce by AI as unprecedented, but technology causing job losses is an old story. In 1800, 90 percent of Americans worked in agriculture. By 1900, thanks to such innovations as the mechanized reaper and steam-driven tractor, that number had dropped to 40 percent. With further advances in the 20th century, the share of the American workforce employed in agriculture had fallen to less than 2 percent by 2000 (Acemoglu, 2009, chapter 20). The reduction in the agricultural workforce was followed by a rapid increase in employment in manufacturing and services. And, ever since manufacturing employment peaked around 1970, the United States has seen continued growth in service jobs.

Vast increases in productivity, enabled by technology, made these shifts in sectoral employment possible. During the second half of the 20th century, output per agricultural worker grew nearly fifteen-fold (Fuglie et al., 2007). Productivity growth of this kind, made possible by technology, is what allows us to enjoy a standard of living today more than twenty times greater than that of our preindustrial ancestors (Morris, 2010).

It's worth noting, however, that it took several generations before working people experienced real wage growth during the Industrial Revolution in Britain (Feinstein, 1998). The shift from agriculture to industry generated wrenching social and cultural dislocation. Technological change provides long-term benefits, but the transitions can be rocky.

One reason the impact of technology on jobs creates alarm is that it's easy to see existing jobs being eliminated, but it's far more difficult to imagine the new ones that will be created. It would have been nearly impossible for a person in 1970 to imagine some of the jobs listed by the U.S. Bureau of Labor Statistics as being as among today's fastest-growing occupations: wind turbine service technician, information security analyst, and genetic counselor (U.S. BLS, 2019). Even two decades ago, when the dot-com boom was under way, few foresaw the emergence of social media, smart devices, and cloud computing—or the millions of jobs that have been created in connection with those new technologies.

One reason we can expect the number of jobs to continue to grow is that humans have proven remarkably consistent in their craving for new goods and services. This principle—the insatiability of human demand (Autor, 2017)—has driven the creation of new kinds of jobs to meet emerging wants. The iconoclastic American economist Thorstein Veblen captured this idea in his quip: "Invention is the mother of necessity" (Veblen, 1914).

Unless there is a sharp break with the past, it's reasonable to assume that while some jobs will be eliminated in the future by AI, new jobs will also be created. In recent years, most of the job growth in advanced economies has been at either the high- or the low-skilled ends of the spectrum (Autor et al., 2006), hollowing out the middle class. A key question for the future will be: where on the spectrum of low-medium-high skill will newly created jobs fall?

There has also been a growing move to focus not on jobs, but on the tasks or activities that comprise jobs. Employing that perspective leads to the insight that many jobs will be transformed as technology takes over tasks formerly performed by people, and people assume responsibility for new activities in the time that has been freed up (Levy et al., 2003; Autor, 2013).

Thus, the coming years and decades will likely see AI eliminating some jobs, creating many new ones, and transforming others. McKinsey & Company projects advances in technology over the next decade will require 14 percent of workers globally to change occupations or significantly change how they do their jobs (Manyika et al., 2017). One of the biggest policy challenges presented by AI will be to smooth the transition for people whose jobs are significantly impacted.

Many of the policy issues associated with AI also arise from other kinds of technologies and are discussed in other reports in this series. Here we therefore focus on issues most closely linked to AI.

## WHAT CAN DIFFERENT PARTS OF SOCIETY DO?

A broad range of institutions can help workers whose lives are disrupted by AI. These include K-12 schools, colleges, and universities; businesses; worker organizations; community-based organizations; and government. Each case is outlined below.

### K-12 schools, colleges, and universities

In the early 20th century, mechanization was displacing U.S. agricultural workers even as new jobs were opening in factories. But the technologies being used in manufacturing required workers with better skills in arithmetic, reading, and writing than were needed for farming. In response, states built new high schools and began to require children to attend them until age 16. This effort, which came to be known as the High School Movement, quickly spread nationwide; by midcentury, the United States had the best-educated population in the world (Goldin and Katz, 2008).

We believe that a comparable, 21st-century push to boost post-secondary education—at both two- and four-year institutions—could help workers gain the skills they need to thrive in an AI-enabled world.

It's important to realize, however, that what's needed is not just the cognitive skills that have been the traditional focus of schools. As AI systems take over some of the routine cognitive tasks that have historically provided employment for high school and college graduates, there is likely to be an increasing need for people to use their social skills for tasks that still can't be done well by computers. For example, when banks introduced ATMs, they continued to employ about the same number of the tellers as before. But the tellers became more like salespeople who built relationships with customers (Bessen, 2015). In fact, research suggests that interpersonal skills are critical to the effective functioning of face-to-face groups, and they may be equally important in online groups whose members never meet in person (Wooley et al., 2010; Engel et al., 2014).

These findings were reinforced by research done at LinkedIn, which found that 57 percent of senior leaders believe soft skills (such as creativity, an ability to work well with others, and adaptability) are more important than technical skills. Many companies are looking for people with creativity, persuasion, and collaboration skills (Petrone, 2018).

Schools at all levels can help prepare the workforce of the future for the changes that AI will bring by adding programs that not only train students to search for information online and use word processing and spreadsheet applications, but also teach them to think computationally. Computational thinking is not just about learning to code; it's also designed to build broad problem-solving skills. For example, students can learn to break down complex challenges into manageable parts that could be addressed by an "information processing agent," be it a computer or a person doing calculations (Wing, 2006).

Workers who want to gain new skills don't have to be restricted to traditional education. Innovations in online education, undertaken by legacy universities, for-profit companies such as Udacity and Coursera, and nonprofits like Khan Academy make training more affordable and accessible to workers across many different sectors (Sarma and Bonvillian, 2020).

One particularly interesting possibility is that online (or in-person) coursework could be combined with activities on web-based crowdsourcing platforms to give students a version of real work experience in a low-risk way (Malone, 2018, chapter 19). In a traditional organizational setting, low-quality work can be a serious problem for the employer. But many crowdsourcing platforms are structured so that participants' contributions are averaged, or the best entries are selected. In this setting, students could submit work and see how it stacks up against work by professionals in the field, gaining experience without hurting the performance of the overall organization.

### Businesses

Workers have long obtained job-specific (as opposed to general) skills through training delivered by businesses, either in formal programs or informally on the job. For instance, a nationally representative survey of nearly 4,000 American workers by MIT's Paul Osterman found that in the prior year, nearly 60 percent of respondents had participated in a formal training program offered by their employers and nearly 50 percent had received informal training on the job (Osterman, 2020). Many large companies today envision undertaking workforce reskilling in the face of AI (Illanes, 2018), and AI tools are even being used to help identify reskilling needs for some workers (Weise, 2020).

Some company-sponsored programs, however, look different than their mid-20th century counterparts. For example, Grow with Google is a reskilling program focused not on Google employees, but on increasing the technical skills of people in the company's broad network of partners and customers. By enhancing the tech savvy of the entire ecosystem it serves, Google hopes to enable such partners to use company offerings in more sophisticated ways.

Another way that businesses can help is by matching jobs to job seekers. Traditionally, employers advertise open positions, prospective employees provide information about their credentials and experience, and the employers screen for promising candidates. Now AI has begun to play a role in this process. AI algorithms are used to screen candidates and AI-enabled tools like chatbots and automation software can streamline interactions between the candidate and the company (Albert, 2019).

In recent decades, online marketplaces have also emerged with the potential to dramatically improve the matching of jobs to job seekers, both for permanent work and short-term assignments (Malone and Laubacher, 1998). The kinds of assignments on offer on these platforms vary widely, ranging from freelance software development on Upwork, to permanent medical assistant jobs on Monster.com, to giving rides via Uber and Lyft or completing small tasks through Amazon's Mechanical Turk. These platforms offer the promise of better matching an employer needs with a job seeker's skills.

### Worker organizations

Worker organizations can also help to meet the challenges AI will present. Novel kinds of worker associations, modeled on mutual aid societies or even medieval guilds, are seeking to invent new models (sometimes also called guilds) (Laubacher and Malone, 2003). One example is Freelancers Union, a New York City–based nonprofit that provides insurance plans and represents the political interests of independent workers. Modern guilds can step in to provide many of the features linked to traditional employment: access to health insurance and pension plans, pathways for career development, social connections, a sense of professional identity, and even a form of unemployment insurance.

For instance, members might pay a percentage of their income to the guild in good times in return for a guaranteed minimum income in bad times. This system would give guilds a natural incentive to provide training and job placement services to members to prevent them from being unemployed. One recently launched startup, Trupo, offers such income-smoothing insurance for freelancers.

A natural way new kinds of guilds could form is for people who work in the same occupation to band together, the way they did in medieval guilds and in the labor unions of the early industrial era. People might also join guilds based on common educational or work experience—alumni of the same school or the same large company. Or, they might join together based on a common geographic or ethnic background. For example, immigrants from Taiwan and South Asia have formed strong informal ties that have helped propel their shared professional interests in the Silicon Valley region of California (Saxenian, 2007).

Guilds may be particularly suited to fields where new technologies are creating fluid work arrangements, such as those prevalent in today's tech sector. And they may be an attractive option for workers who possess in-demand skills and power in the labor market.

Many lesser skilled workers are also already being affected by AI: for example, piece workers who perform micro tasks such as content moderation though crowdworking platforms. These workers provide the "last mile" of many AI-driven systems, the human touch that steps in when AI algorithms fall short (Gray and Suri, 2017). But these workers are typically poorly paid and thus have little market power. If they are mistreated, a caste of "ghost workers" could arise, toiling away in "digital sweatshops" (Zittrain, 2009) behind the shiny veneer of AI.

For those with lower skills and less power, labor unions, which band workers together to negotiate with employers collectively, still have a role to play. Their power has waned in the face of decades of globalization, mobile capital, and the embrace of market-oriented policies by national governments. To protect "ghost workers," policies could be adjusted to make organizing easier, and labor standards in place for traditional employment can also be applied to virtual work.

For example, on Amazon's Mechanical Turk program, which is used by many social science researchers, payment rates are typically pegged informally to the minimum wage (though this is not required). An increase in the minimum wage could make online experiments more expensive but would help the piece workers whose labor ensures that AI systems work.

### Community-based organizations

Organizations embedded in local communities will be needed to provide opportunities for workers displaced by AI. People who have lost their jobs, or have seen their jobs transformed, are in a vulnerable position. Organizations with deep roots in areas where displaced workers live will be in the best position to reach them effectively.

Bit Source, a small business based in Eastern Kentucky, provides an example of how an organization with strong community ties can reach workers rooted in a particular locale. It was launched in 2014 by Rusty Justice and Lynn Parrish, who formerly worked in the coal industry. The company's mission was to teach former coal miners to code computers and to apply their newly gained skills to providing technology services for clients.

In 2015, a cohort of ten prospective coders earned $15 an hour to participate in a twenty-two-week training program. After the training ended, Bit Source began paying its employees locally competitive salaries, working toward what they had been making before in the mines (Smiley, 2015).

A key insight behind the launch of Bit Source was its founders' recognition that miners, from their experience working in coal, had a technical orientation and problem-solving skills that could readily be applied to software development. As Justice put it, "a coal miner is just a tech worker that gets dirty" (Grow with Google, 2019). Nevertheless, miners needed to learn domain-specific skills to make the transition.

Bit Source was embraced by the Appalachian community. But its journey has not followed a straight line. Teaching the miners to code went quickly but gaining credibility in the tech services marketplace has taken longer. Bit Source was designed to help workers displaced by clean energy technology, but its experience offers lessons that will apply to other community-based programs designed to help workers displaced by AI.

### Government

An important traditional role of government has been to support the development of basic skills as a public good. In the U.S. government, there is a growing awareness that reskilling initiatives are needed to address the challenges AI will present. For instance, both the President's Council of Advisors on Science and Technology and the National Science Foundation have recently supported such programs (Reardon, 2019).

Last year, NSF launched a program to identity creative new ways to upskill/reskill workers (NSF 2019). And PCAST featured skill development prominently in its recent report on the industries of the future (PCAST 2020). As Osterman (2020) emphasizes, there is an important need for the government to fund worker reskilling programs.

In some instances, the government has even created jobs directly, usually in response to economic downturns. In the United States, the best-known example is the Works Progress Administration (WPA), a government agency founded during Franklin Roosevelt's presidency with the mission of giving work to the unemployed (Hopkins 1937). In the face of the economic slowdown caused by the COVID-19 pandemic, a "Digital WPA" (Malone, 2020) could provide work for people who have lost their jobs and give those who lack technical skills a chance to develop them. In similar manner, government-provided jobs could help workers displaced by AI to pay their bills and to reskill.

Government has also traditionally played a role in helping people whose employment is in transition, whether as a result of technical change or other reasons, primarily through unemployment insurance. In response to the challenges expected to come with AI, many observers have called for a universal basic income (UBI) as a way to provide a financial foundation for workers facing a potentially turbulent future (see, for example, Yang, 2018). UBI has been embraced by some, especially in Silicon Valley, but criticized by others who cite the expense of providing such income and who argue it could rob people of the sense of identity and meaning they now find in work.

Another area in which government can support the workforce transition to AI centers on legal and regulatory frameworks. Internet technology has already brought to the fore many issues related to employment practices. Notably, the rise of the gig economy has blurred even further the already hazy line between people who work as formal employees and those who are independent contractors. Since AI can create new, distributed forms of work that combine people and machines in increasingly novel ways, governments may need to step in to create new rules.

Government may also need to assume a larger role in confronting the societal implications of AI. For example, unrecognized biases embedded in the underlying data or algorithms may exacerbate discriminatory practices that already exist due to broadly shared and unexamined social attitudes.

While government can certainly play a key role in addressing such issues, they are properly the concern of other key stakeholders as well: educators, business executives, and officials at worker associations, nonprofits, and community-based organizations.

## Conclusion

Artificial intelligence already pervades many parts of our daily lives—providing services ranging from instant credit card authorizations to effortless speech-to-text transcriptions—and the ways in which AI might improve our lives in the future are almost limitless. However, ensuring that the benefits of the changes wrought by AI are distributed widely across the workforce and across the globe will require a concerted effort to help people whose jobs are eliminated—by facilitating the creation of new jobs, matching jobs to job seekers, and providing education, training, and sometimes financial support to people as they transition from old jobs to new ones.

In the long term, AI-enabled human-computer superminds will facilitate new forms of societal endeavor in business, science, and art that are as different from what we know today as today's computers are from the adding machines of half a century ago. We have not attempted in this report to identify all the issues that will face us when that time comes, but guiding this ongoing process well will require all the collective wisdom our species can muster.

# References


Acemoglu, Daron. 2009. *Introduction to Modern Economic Growth*. Princeton, New Jersey: Princeton University Press.

Acemoglu, Daron, Andrea Manera, and Pascual Restrepo. 2020. Taxes, Automation, and the Future of Labor. MIT Task Force on the Work of the Future Research Brief. https://workofthefuture.mit.edu/wp-content/uploads/2020/10/2020-Research-Brief-Acemoglu-Manera-Restrepo.pdf.

Acemoglu, Daron and Pascual Restrepo. 2017. "Robots and Jobs: Evidence from US Labor Markets." NBER Working Paper No. 23285 https://www.nber.org/papers/w23285.

Acemoglu, Daron and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33 (2): 3-30. https://www.aeaweb.org/articles?id=10.1257/jep.33.2.3.

Albert, Edward Tristam. 2019. "AI in Talent Acquisition: A Review of AI-applications Used in Recruitment and Selection." *Strategic HR Review* 18 (5): 215-221. https://www.emerald.com/insight/content/doi/10.1108/SHR-04-2019-0024/full/html.

Ammirato, Wendy, David Anderton-Yang, Hiroaki Atsuji, Annalyn Bachmann, Olympia Banerjee, Sujit Basu, Linda Calandra, et al. 2019. "How Can the Japanese Healthcare System Deal More Effectively with Adult Depression? Harnessing the Power of Superminds." MIT Center for Collective Intelligence and Takeda Pharmaceutical Working Paper. https://www.dropbox.com/s/et4k24wih62f8fz/Takeda%20Report-Harnessing%20the%20Power%20of%20a%20Supermind%20FINAL.pdf?dl=0.

Araki, Brandon, Kiran Vodrahalli, Thomas Leech, Cristian Ioan Vasile, Mark Donahue, Daniela Rus. 2020. "Deep Bayesian Nonparametric Learning of Rules and Plans from Demonstrations with a Learned Automaton Prior." *Proceedings of the AAAI Conference on Artificial Intelligence*. https://aaai.org/ojs/index.php/AAAI/article/view/6559.

Armstrong, Stuart and Kaj Sotala. 2015. "How We're Predicting AI—or Failing To." In *Beyond Artificial Intelligence: The Disappearing Human-Machine Divide*, edited by Jan Romport, Eva Zackova, and Jozef Kelemen. New York: Springer. https://link.springer.com/chapter/10.1007/978-3-319-09668-1_2.

Autor, David H. 2013. "The 'Task Approach' to Labor Markets: An Overview." *Journal of Labour Market Research*. https://labourmarketresearch.springeropen.com/articles/10.1007/s12651-013-0128-z.

Autor, David. 2017. "Will Automation Take Away All Our Jobs?" Ideas.TED.com. https://ideas.ted.com/will-automation-take-away-all-our-jobs/.

Autor, David H., Frank Levy, Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics* 118 (4): 1279–1333. https://academic.oup.com/qje/article-abstract/118/4/1279/1925105.

Autor, David, H., Lawrence F. Katz, and Melissa S. Kearney. 2006. "The Polarization of the U.S. Labor Market." *American Economic Review* 96 (2): 189-194. DOI: 10.1257/000282806777212620.

Bessen, James. 2015. *Learning by Doing: The Real Connection Between Innovation, Wages, and Wealth*. New Haven: Yale University Press.

Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. 2020. "Language Models Are Few-Shot Learners." July 22, 2020. https://arxiv.org/abs/2005.14165.

Brynjolfssson, Erik and Andrew McAfee. 2011. *Race Against the Machine: How the Digital Revolution Is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy*. Lexington, Massachusetts: Digital Frontier Press.

Brynjolfssson, Erik and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: Norton.

Brooks, Rodney A. 2017. "The Seven Deadly Sins of AI Predictions." *MIT Technology Review*, October 6, 2017. https://www.technologyreview.com/2017/10/06/241837/the-seven-deadly-sins-of-ai-predictions.

Buchanan, Bruce G. and Edward H. Shortliffe, editors. 1984. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Reading, Massachusetts: Addison-Wesley.

Bussler, Frederik. 2020. "Will the Latest AI Kill Coding?" *Towards Data Science*, July 20, 2020. https://towardsdatascience.com/will-gpt-3-kill-coding-630e4518c04d.

Dahik, Adriana, Deborah Lovich, Caroline Kreafle, Allison Bailey, Julie Kilmann, Derek Kennedy, Prateek Roongta, Felix Schuler, Leo Tomlin, and John Wenstrup. 2020. "What 12,000 Employees Have to Say About the Future of Remote Work." Boston Consulting Group. https://www.bcg.com/publications/2020/valuable-productivity-gains-covid-19.

Darwiche, Adrian. 2018. "Human-Level Intelligence or Animal-Like Abilities?" *Communications of the ACM* 61 (10): 56-67. https://doi.org/10.1145/3271625.

Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li and Li Fei-Fei. 2009. "ImageNet: A Large-Scale Hierarchical Image Database." Department of Computer Science, Princeton University. http://www.image-net.org/papers/imagenet_cvpr09.pdf.

Engel, David, Anita Williams Woolley, Lisa X. Jing, Christopher F. Chabris, Thomas W. Malone. 2014. "Reading the Mind in the Eyes or Reading between the Lines? Theory of Mind Predicts Collective Intelligence Equally Well Online and Face-to-Face." *PLOS One* 9 (12): e115212. https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0115212.

Feinstein, Charles H. 1998. "Pessimism Perpetuated: Real Wages and the Standard of Living in Britain during and after the Industrial Revolution." *Journal of Economic History* 58 (3): 625-658. https://ideas.repec.org/a/cup/jechis/v58y1998i03p625-658_02.html.

Ford, Martin. 2013. "Could Artificial Intelligence Create an Unemployment Crisis?" *Communications of the ACM* 56: 7 (July), 37–39. http://cacm.acm.org/magazines/2013/7/165475-could-artificial-intelligence-create-an-unemployment-crisis/abstract.

Ford, Martin. 2015. *Rise of the Robots: Technology and the Threat of a Jobless Future*. New York: Basic.

Frey, Carl Benedikt, and Michael A. Osborne. 2013. "The Future of Employment: How Susceptible are Jobs to Computerisation?" Oxford Martin School Programme on the Impacts of Future Technology Working Paper. http://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf.

Fuglie, Keith O., James M. MacDonald, and Eldon Ball. 2007. Productivity Growth in U.S. Agriculture. U.S. Department of Agriculture, Economic Research Service, Economic Brief Number 9. https://www.ers.usda.gov/webdocs/publications/42924/11854_eb9_1_.pdf?v=37.7.

Gayford, Martin. 2016. "Robot Art Raises Questions about Human Creativity." *MIT Technology Review*, February 15, 2016. https://www.technologyreview.com/2016/02/15/162067/robot-art-raises-questions-about-human-creativity/.

Goldin, Claudia, and Lawrence F. Katz. 2008. *The Race between Education and Technology*. Cambridge MA: Belknap for Harvard University Press.

Gray, Mary L. and Siddharth Suri. 2017. "The Humans Working Behind the AI Curtain." *Harvard Business Review*, January 9, 2017. https://hbr.org/2017/01/the-humans-working-behind-the-ai-curtain.

Grow with Google. 2019. "Bit Source: Appalachian Ingenuity." YouTube. https://www.youtube.com/watch?time_continue=17&v=Bj3f1ap4Lfs&feature=emb_logo.

Guszcza, James and Jeff Schwartz. 2020. "Superminds, not Substitutes: Designing Human-Machine Collaborations for a Better Future of Work." *Deloitte Review* 27: 26-41. https://www2.deloitte.com/content/dam/insights/us/articles/6672_superminds-not-substitutes/DI_DR27-Superminds.pdf.

Hardesty, Larry. 2019. "The History of Amazon's Recommendation Algorithm." *Amazon Science*, November 22, 2019. https://www.amazon.science/the-history-of-amazons-recommendation-algorithm.

Hopkins, Harry. 1937. *Inventory: An Appraisal of Results of the Works Progress Administration*. Washington, D.C.: Government Printing Office.

Illanes, Pablo, Susan Lund, Mona Mourshed, Scott Rutherford, and Magnus Tyreman. 2018. "Retraining and Reskilling Workers in the Age of Automation." McKinsey Global Institute. https://www.mckinsey.com/featured-insights/future-of-work/retraining-and-reskilling-workers-in-the-age-of-automation.

Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. New York: Macmillan.

Kebotix. 2018. "Harvard Scientists Launch Kebotix, a Breakthrough AI and Robotics Tech Company for Rapid Innovation of Materials." BusinessWire. https://www.kebotix.com/press/2018/11/7/harvard-scientists-launch-breakthrough-ai-and-robotics-tech-company-for-rapid-innovation-of-materials.

Lake, Brenden M., Tomer D. Ullman, Joshua B. Tenenbaum, and Samuel J. Gershman. 2017. "Building Machines that Learn and Think Like People." *Behavioral and Brain Sciences* 40, e253. http://gershmanlab.webfactional.com/pubs/Lake17.pdf.

Laubacher, Robert, Gianni Giacomelli, Kathleen Kennedy, David Sun Kong, Annalyn Bachmann, Katharina Kramer, Paul Schlag, and Thomas W. Malone. 2020. "Using a Supermind to Design a Supermind: A Case Study of University Researchers and Eorporate Executives Co-designing an Innovative Healthcare Concept." MIT Sloan Working Paper No. 6109-20. https://cci.mit.edu/wp-content/uploads/2020/07/03.31.20-MIT-Sloan-CCI-Working-Paper-6109-20.pdf.

Laubacher, Robert and Thomas W. Malone. 2003. "The Retreat of Firm and the Rise of Guilds: The Employment Relationship in an Age of Virtual Business." In *Inventing the Organizations of the 21st Century*, edited by Thomas W. Malone, Robert Laubacher, and Michael S. Scott Morton. Cambridge, Massachusetts: MIT Press.

Lechner, Mathias, Ramin Hasani, Alexander Amini, Tom Henderson, Daniela Rus, and Radu Grosu. 2020. "Neural Circuit Policies Enabling Auditable Autonomy." *Nature Machine Intelligence* 2, 642–652. https://www.nature.com/articles/s42256-020-00237-3.

Lee, Kai-Fu. 2018 *AI Superpowers: China, Silicon Valley, and the New World Order*. New York: Houghton Mifflin Harcourt.

Liebenwein, Lucas, Cenk Baykal, Harry Lang, Dan Feldman, and Daniela Rus. 2019. "Provable Filter Pruning for Efficient Neural Networks." *Proceedings of the International Conference on Learning Representations*. https://openreview.net/forum?id=BJxkOlSYDH.

Madry, Aleksander, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, Adrian Vladu. 2017. "Towards Deep Learning Models Resistant to Adversarial Attacks. *Proceedings of the International Conference on Representation Learning*. https://arxiv.org/abs/1706.06083.

Malone, Thomas W. 2018. *Superminds: The Surprising Power of People and Computers Thinking Together*. New York: Little, Brown and Company.

Malone, Thomas W. 2020. "Why We Need a 'Digital WPA' Similar to the Depression-era Works Progress Administration." *The Hill*, May 4, 2020. https://thehill.com/opinion/technology/495943-why-we-need-a-digital-wpa-similar-to-the-depression-era-works-progress.

Malone, Thomas W. and Robert Laubacher. 1998. "The Dawn of the E-Lance Economy." *Harvard Business Review*, September-October. https://hbr.org/1998/09/the-dawn-of-the-e-lance-economy.

Manyika, James, Susan Lund, Michael Chui, Jacques Bughin, Jonathan Woetzel, Parul Batra, Ryan Ko, and Saurabh Sanghvi. 2017. "Jobs Lost, Jobs Gained: What the Future of Work Will Mean for Jobs, Skills, and Wages." McKinsey Global Institute. https://www.mckinsey.com/featured-insights/future-of-work/jobs-lost-jobs-gained-what-the-future-of-work-will-mean-for-jobs-skills-and-wages.

Mao, Jiayua, Chuang Gan, Pushmeet Kohli, Joshua B. Tenenebaum, and Jiajun Wu. 2019. "The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words, and Sentences from Natural Supervision." ArXiv.org. https://arxiv.org/abs/1904.12584.

Marcus, Gary. 2018. "Deep Learning: A Critical Appraisal." ArXiv.org. https://arxiv.org/abs/1801.00631.

Marcus, Gary and Ernest Davis. 2020a. "GPT-3, Bloviator: OpenAI's Language Generator Has No Idea What It's Talking About." *MIT Technology Review*, August 22, 2020. https://www.technologyreview.com/2020/08/22/1007539/gpt3-openai-language-generator-artificial-intelligence-ai-opinion.

Marcus, Gary and Ernest Davis. 2020b. *Rebooting AI: Building Artificial Intelligence We Can Trust*. New York: Penguin.

McCarthy, J., M. L. Minsky, N. Rochester, and C. E. Shannon. 1955. "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence." http://raysolomonoff.com/dartmouth/boxa/dart564props.pdf.

Metz, Cade. 2019. “A.I. is Learning From Humans. Many Humans.” *New York Times*, August 16, 2019. https://www.nytimes.com/2019/08/16/technology/ai-humans.html.

Minsky, Marvin. 1961. “Steps Toward Artificial Intelligence.” *Proceedings of the IRE* 49 (1): 8-30.

MIT Center for Collective Intelligence. 2020. *Supermind Design Primer*. https://cci.mit.edu/primer.

MIT Sloan School of Management and MIT Computer Science and Artificial Intelligence Lab. 2017. Artificial Intelligence: Implications for Business Strategy. Online short course. https://executive.mit.edu/openenrollment/program/artificial-intelligence-implications-for-business-strategy-self-paced-online/.

MIT Sloan School of Management and MIT Computer Science and Artificial Intelligence Lab. 2019. Machine Learning in Business. Online short course. https://executive.mit.edu/openenrollment/program/machine-learning-implementation-in-business-self-paced-online/.

Molteni, Megan. 2017. “Want a Diagnosis Tomorrow, Not Next Year? Turn to AI.” *Wired*, August 10, 2017. https://www.wired.com/story/ai-that-will-crowdsource-your-next-diagnosis/.

Morris, Ian. 2010. *Why the West Rules—for Now: The Patterns of History, and What They Reveal About the Future*. New York: Farrar Strauss & Giroux.

Mountz, Mick. 2011. “A Day in the Life of a Kiva Robot.” Wired Business Conference. https://www.youtube.com/watch?v=6KRjuuEVEZs.

National Science Foundation. 2019. Careen Compass Challenge winner announced, August 30, 2019. https://www.nsf.gov/news/special_reports/announcements/083019.jsp.

Osterman, Paul. 2020. “Skill Training for Adults.” MIT Task Force on the Work of the Future Research Brief. https://workofthefuture.mit.edu/research-post/skill-training-for-adults/.

Palmer, Annie. 2020. “Amazon Starts Selling Cashierless Checkout Technology to Other Retailers.” CNBC. https://www.cnbc.com/2020/03/09/amazon-selling-cashierless-checkout-technology-to-other-retailers.html.

Peng, Tony. 2019. “The Staggering Cost of Training SOTA AI Models.” *Synced Review*, June 27, 2019. https://medium.com/syncedreview/the-staggering-cost-of-training-sota-ai-models-e329e80fa82

Pennycook, Gordon and David G. Rand. 2019. “Fighting Misinformation on Social Media Using Crowdsourced Judgments of News Source Quality.” *PNAS* 116 (7) 2521-2526. https://doi.org/10.1073/pnas.1806781116.

Petrone, Paul. 2018. "The Skills Companies Need Most in 2019—And How to Learn Them." LinkedIn Learning Blog. https://www.linkedin.com/business/learning/blog/top-skills-and-courses/the-skills-companies-need-most-in-2019-and-how-to-learn-them.

Picard, Rosiland. 1997. *Affective Computing*. Cambridge, Massachusetts: MIT Press.

President's Council of Advisors on Science and Technology. 2020. Recommendations for Strengthening American Leadership in Industries of the Future. June 2020. https://science.osti.gov/-/media/_/pdf/about/pcast/202006/PCAST_June_2020_Report.pdf?la=en&hash=019A4F17C79FDEE5005C51D3D6CAC81FB31E3ABC.

Rajpurkar, Pranav, Jeremy Irvin , Robyn L. Ball, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, et al. 2018. "Deep Learning for Chest Radiograph Diagnosis: A Retrospective Comparison of the CheXNeXt Algorithm to Practicing Radiologists." *PLOS Medicine*, November 20, 2018. https://doi.org/10.1371/journal.pmed.1002686.

Ramaswamy, Sree, James Manyika, Gary Pinkus, Katy George, Jonathan Law, Tony Gambell, and Andrea Serafinol. 2017. "Making it in America: Revitalizing US Manufacturing." McKinsey Global Institute. https://www.mckinsey.com/featured-insights/americas/making-it-in-america-revitalizing-us-manufacturing.

Rao, Anand, Flavio Palaci, and Wilson Chow. 2019. "This Is What the World's CEOs Really Think of AI." World Economic Forum, Meeting of the New Champions, Dalian, People's Republic of China. https://www.weforum.org/agenda/2019/06/ai-will-shape-our-future-how-do-we-ensure-it-is-used-responsibly/.

Reardon, Sara. 2019. "Trump's Science Adviser on Research Ethics, Immigration and Presidential Tweets." *Nature News Q&A*. https://www.nature.com/articles/d41586-019-01396-z.

Remus, Dana and Frank Levy. 2017. "Can Robots Be Lawyers? Computers, Lawyers, and the Practice of Law." *Georgetown Journal of Legal Ethics* 30, no. 3 (summer). https://go.gale.com/ps/anonymous?id=GALE%7CA514460996&sid=googleScholar&v=2.1&it=r&linkaccess=abs&issn=10415548&p=AONE&sw=w.

Russell, Stuart J. and Peter Norvig. 2020. *Artificial Intelligence: A Modern Approach*, 4th ed. New York: Pearson.

Sarma, Sanjay and William Boone Bonvillian. 2020. "Applying New Education Technologies to Meet Workforce Education Needs." MIT Task Force on the Work of the Future Research Brief. https://workofthefuture.mit.edu/wp-content/uploads/2020/10/2020-Research-Brief-Sarma-Bonvillian.pdf.

Saxenian, Annalee. 2007. *The New Argonauts: Regional Advantage in a Global Economy*. Cambridge, Massachusetts: Harvard University Press.

Slagle, James Robert. 1961. "A Heuristic Program that Solves Symbolic Integration Problems in Freshman Calculus: Symbolic Automatic Integrator (SAINT)." PhD diss., MIT Department of Mathematics. https://dspace.mit.edu/bitstream/handle/1721.1/11997/31225400-MIT.pdf?sequence=2.

Smiley, Lauren. 2015. "Canary in the Code Mine." *Backchannel*. https://medium.com/backchannel/canary-in-the-code-mine-903884eca853.

Smith, Kevin, Eliza Kosoy, Alison Gopnik, Deepak Pathak, Alan Fern, Josh Tenenbaum, Tomer D. Ullman. 2020. "The Origins of Common Sense in Humans and Machines." CogSci 2020. https://cognitivesciencesociety.org/wp-content/uploads/2020/06/cogsci20a-W3-doc.pdf.

Szegedy, Christian, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2014. "Intriguing Properties of Neural Networks." International Conference on Learning Representations. https://arxiv.org/abs/1312.6199.

Turing, A. M. 1950. "Computing Machinery and Intelligence." *Mind* 59 (236): 433-460.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. "Attention Is All You Need." 31st Conference on Neural Information Processing Systems. https://arxiv.org/pdf/1706.03762.pdf.

Vebelen, Thorstein. 1918. *The Instinct of Workmanship and the State of the Industrial Arts*. New York: Macmillan.

Vincent, James. 2018. "Christie's Sells Its First AI Portrait for $432,500, Beating Estimates of $10,000." *The Verge*, October 15, 2018. https://www.theverge.com/2018/10/25/18023266/ai-art-portrait-christies-obvious-sold.

Volcovici, Valerie. 2015. "In Kentucky, Talk of a 'War on Coal' Gives Way to Hope for New Economy." *Reuters*, June 4, 2015. https://www.reuters.com/article/us-coal-usa-kentucky-insight/in-kentucky-talk-of-a-war-on-coal-gives-way-to-hope-for-new-economy-idUSKBN0OK0CE20150604.

US Bureau of Labor Statistics. 2019. *Occupational Outlook Handbook: Fastest Growing Occupations*. https://www.bls.gov/ooh/fastest-growing.htm.

Weise, Michelle R. 2020. "Research: How Workers Shift from One Industry to Another." *Harvard Business Review*, July 7, 2020. https://hbr.org/2020/07/research-how-workers-shift-from-one-industry-to-another.

Wikipeida. 2020. “Artificial Intelligence.” Last modified September 25, 2020. https://en.wikipedia.org/wiki/Artificial_intelligence.

Wing, Jeannette. 2006. “Computational Thinking.” *Communications of the ACM* 49 (3): 33-35. https://dl.acm.org/doi/10.1145/1118178.1118215.

Wolters, Lukas. 2020. “Robots, Automation, and Employment: Where We Are.” MIT Work of the Future Working Paper 05-2020. https://workofthefuture.mit.edu/wp-content/uploads/2020/08/WotF-Working-Paper-05-2020.pdf.

Wooley, Anita Williams, Christopher F. Chabris, Alex Pentland, Nada Hashmi, and Thomas W. Malone. 2010. “Evidence of a Collective Intelligence Factor in the Performance of Human Groups.” *Science* 330 (6004): 686-688. https://science.sciencemag.org/content/330/6004/686.

Wright, Alex. 2015. “Algorithmic Authors.” *Communications of the ACM* 58 (11): 12-14. https://dl.acm.org/doi/10.1145/2820421.

Xiong, Wayne, Lingfeng Wu, Jasha Droppo, Xuedong Huang, and Andreas Stolcke. 2018. “The Microsoft 2017 Conversational Speech Recognition System.” *Proceedings IEEE International Conference on Acoustics, Speech and Signal Processing*. https://arxiv.org/abs/1708.06073.

Yala, Adam, Regina Barzilay, Laura Salama, Molly Griffin, Grace Sollender, Aditya Bardia, Constance Lehman, et al. 2017. “Using Machine Learning to Parse Breast Pathology Reports.” *Breast Cancer Research and Treatment* 161: 203–211. https://doi.org/10.1007/s10549-016-4035-1.

Yang, Andrew. 2018. *The War on Normal People: The Truth About America’s Disappearing Jobs and Why Universal Basic Income Is our Future*. New York: Hachette.

Zittrain, Jonathan. 2009. “The Internet Creates a New Kind of Sweatshop.” *Newsweek*, December 7, 2009. https://www.newsweek.com/internet-creates-new-kind-sweatshop-75751.