

Volume 1 of 12

14th Annual Edition

2021

Tech Trends Report

Strategic trends that will influence business, government, education, media and society in the coming year.



Artificial Intelligence

- 03 Overview
- 04 Macro Forces and Emerging Trends
- 06 Summary
- 08 **Artificial Intelligence**
- 09 An Executive's Guide to AI
- 09 Machine Learning
- 09 Deep Learning
- 10 Weak and Strong AI
- 12 **Enterprise**
- 12 The Rise of MLOps
- 12 Low-Code or No-Code Machine Learning
- 12 Web-Scale Content Analysis
- 12 Simulating Empathy and Emotion
- 13 Artificial Emotional Intelligence
- 13 Serverless Computing
- 14 Expert Insight: Emotion AI Will Power the Empathy Economy, but AI Still Needs to Work
- 17 AI in the Cloud
- 17 AI at the Edge
- 17 Advanced AI Chipsets
- 17 Digital Twins
- 17 Spotting Fakes
- 18 Natural Language Processing for ESGs
- 18 Intelligent Optical Character Recognition
- 18 Robotic Process Automation
- 18 Massive Translation Systems
- 19 Predicting Systems and Site Failures
- 19 Liability Insurance for AI
- 19 Manipulating AI Systems for Competitive Advantage
- 19 Global Rush to Fund AI

- 19 Algorithm Marketplaces
- 19 100-Year Software
- 20 Scenario: Rage Against the Machine
- 22 **Health, Medicine, and Science**
- 22 AI Speeds Scientific Discovery
- 22 AI-First Drug Discovery
- 23 AI Improves Patient Outcomes
- 23 Deep Learning Applied to Medical Imaging
- 23 NLP Algorithms Detect Virus Mutations
- 23 Diagnostics Without Tests
- 23 Protein Folding
- 23 Dream Communication
- 24 Thought Detection
- 25 Scenario: Deep Twins in the OR
- 27 **Consumer**
- 27 Zero UIs
- 27 Consumer-grade AI Applications
- 27 Ubiquitous Digital Assistants
- 28 Deepfakes for Fun
- 28 Personal Digital Twins
- 30 **Research**
- 30 Closed-Source Code
- 30 Framework Consolidation
- 30 Cost of Training Models
- 31 NLP Benchmarks
- 31 Machine Reading Comprehension
- 31 AI Summarizing Itself
- 31 No Retraining Required
- 31 Graph Neural Networks
- 31 Federated Learning
- 31 GP Models
- 31 GPT-3's Influence
- 32 Vokenization

- 32 Machine Image Completion
- 32 Predictive Models Using Single Images
- 33 Model-free Approaches to RL
- 33 Real-time Machine Learning
- 33 Automated Machine Learning (AutoML)
- 33 Hybrid Human-Computer Vision
- 33 Neuro-Symbolic AI
- 33 General Reinforcement Learning Algorithms
- 34 Continuous Learning
- 34 Proliferation of Franken-Algorithms
- 34 Proprietary, Homegrown AI Languages
- 36 **Talent**
- 36 AI Brain Drain
- 36 AI Universities
- 37 Demand for AI Talent Growing
- 37 Corporate AI Labs
- 37 AI for Interviews
- 39 **Creative**
- 39 Assisted Creativity
- 39 Generative Algorithms for Content Production
- 39 Generating Virtual Environments from Short Videos
- 40 Automated Versioning
- 40 Automatic Voice Cloning and Dubbing
- 40 Automatic Ambient Noise Dubbing
- 42 **Geopolitics and Defense**
- 42 AI Nationalism
- 42 National AI Strategies
- 42 AI as Critical Infrastructure

- 43 Nation-based Guardrails and Regulations
- 43 Regulating Deepfakes
- 43 Making AI Explain Itself
- 43 New Strategic Technical Alliances
- 43 The New Mil-Tech Industrial Complex
- 44 Algorithmic Warfighting
- 45 **China's AI Rules**
- 48 **Society**
- 48 Ethics Clash
- 48 Ambient Surveillance
- 48 Marketplace Consolidation
- 48 Fragmentation
- 49 Expert Insight: AI reveals our real-world biases
- 50 AI Still Has a Bias Problem
- 50 Problematic Training Data
- 50 AI to Catch Cheaters
- 50 Algorithms Targeting Vulnerable Populations
- 51 AI Intentionally Hiding Data
- 51 Undocumented AI Accidents
- 51 Digital Dividends
- 51 Prioritizing Trust
- 52 Scenario: Bully Bots
- 53 Application
- 54 Key Questions
- 55 Sources
- 56 Authors



Overview



The 1920s began in chaos. Cataclysmic disruption resulting from the first world war and the Spanish flu shuttered businesses and provoked xenophobia. Technological marvels like the radio, refrigerator, vacuum cleaner, moving assembly line and electronic power transmission generated new growth, even as the wealth gap widened. More than two-thirds of Americans survived on wages too low to sustain everyday living. The pace of scientific innovation—the discovery of insulin, the first modern antibiotics, and insights into theoretical physics and the structure of atoms—forced people to reconsider their cherished beliefs.

The sheer scale of change, and the great uncertainty that came with it, produced two factions: those who wanted to reverse time and return the world to normal, and those who embraced the chaos, faced forward, and got busy building the future.

It's difficult not to see striking parallels to our modern world. A tumultuous U.S. election, extreme weather events and Covid-19 continue to test our resolve and our resilience. Exponential technologies—artificial intelligence, synthetic biology, exascale computing, autonomous robots, and off-planet missions to space—are challenging our assumptions about human potential. Under lockdown, we've learned how to work from our kitchen tables, lead from our spare rooms, and support each other from afar. But this disruption has only just begun.

With the benefit of both hindsight and strategic foresight, we can choose a path of reinvention. Our 2021 Tech Trends Report is designed to help you confront deep uncertainty, adapt and thrive. For this year's edition, the magnitude of new signals required us to create 12 separate volumes, and each report focuses on a cluster of relat-

ed trends. In total, we've analyzed nearly 500 technology and science trends across multiple industry sectors. In each volume, we discuss the disruptive forces, opportunities and strategies that will drive your organization in the near future.

Now, more than ever, your organization should examine the potential near and long-term impact of tech trends. You must factor the trends in this report into your strategic thinking for the coming year, and adjust your planning, operations and business models accordingly. But we hope you will make time for creative exploration. From chaos, a new world will come.

Amy Webb

Founder
The Future Today Institute

Macro Forces and Emerging Trends

For nearly two decades, the Future Today Institute has meticulously researched macro forces of change and the emerging trends that result. Our focus: understanding how these forces and trends will shape our futures. Our 14th annual Tech Trends Report identifies new opportunities for growth and potential collaborations in and adjacent to your business. We also highlight emerging or atypical threats across most industries, including all levels of government. For those in creative fields, you will find a wealth of new ideas that will spark your imagination.

Our framework organizes nearly 500 trends into 12 clear categories.

Within those categories are specific use cases and recommendations for key roles in many organizations: strategy, innovation, R&D, and risk.

Each trend offers six important insights.

1. Years on the List

We track longitudinal tech and science trends. This measurement indicates how long we have followed the trend and its progression.

2. Key Insight

Concise description of this trend that can be easily understood and repeated to others.

3. Examples

Real-world use cases, some of which should be familiar to you.

4. Disruptive Impact

The implications of this trend on your business, government, or society.

5. Emerging Players

Individuals, research teams, startups, and other organizations emerging in this space.

6. Action Scale

FTI's analysis of what action your organization should take. Fields include:

Watch Closely

Mounting evidence and data, but more maturity is needed. Use it to inform your vision, planning, and research.

Informs Strategy

Strong evidence and data. Longer-term uncertainties remain. Use it to inform your strategic planning.



Act Now

Ample evidence and data. This trend is already mature and requires action.

00 01 02 03 04 05 06 07 08 09 10 11 12 Health, Medical, Wearables 6 Watch Closely Informs Strategy Act Now

1ST YEAR ON THE LIST — 1

Doctorless Exams

AliveCor's personal EKG monitoring system.

KEY INSIGHT — 2

Advancements in diagnostic testing and remote monitoring, supported by cloud computing, machine learning, and low-cost technology are upending traditional doctor visits. Patient data are triaged by algorithm, rather than human doctors alone.

EXAMPLES — 3

Smartphones and smartwatches now take blood pressure readings and perform electrocardiograms, using apps approved by the U.S. Food and Drug Administration. Phones don't just record data; they interpret it. People who wear an Apple Watch know that an unusually high or low heart rate or irregular rhythm may suggest atrial fibrillation. The VROR system, a VR-based eye exam, emulates an eye doctor's ultra-widefield imaging machine but within a compact headset. Data are sent to a mobile app for an understanding of a patient's optic nerve health. StethIO is a mobile stethoscope that uses a smartphone to capture, decode and analyze heart sounds. AliveCor is an FDA-Cleared wireless personal EKG that connects to a phone. Butterfly iQ is a portable ultrasound device that delivers a 2D image. The ParatusPerio Test analyzes different bacteria and sources of inflammation in a patient's mouth.

DISRUPTIVE IMPACT — 4

Continual monitoring helps patients know their baseline vital stats and track any changes. This disrupts traditional healthcare in a few ways. First, with continual monitoring, patients are more likely to intercept an emerging problem in advance and seek out care. For example, if a patient's smartwatch warns of atrial fibrillation, they can call their doctor for next steps. This reduces strain on emergency departments. It also unlocks new opportunities for healthcare systems and insurers willing to use those data and to make medical records systems interoperable. Financial forecasting that harnesses real-time data could be algorithmically recalibrated, and more accurately assess risk. But connected devices aren't accessible to everyone, which means that a new digital divide could be on the horizon, further reducing health equity in many communities.

EMERGING PLAYERS — 5

- The Clue period tracking app
- Apple Health
- StethIO
- Healthy.io
- Paratus Diagnostics
- Butterfly Network
- AliveCor

59

© 2021 Future Today Institute

Macro Forces and Emerging Trends

Scenarios Describe Plausible Outcomes

You will find scenarios imagining future worlds as trends evolve and converge. Scenarios offer a fresh perspective on trends and often challenge your deeply held beliefs. They prompt you to consider high-impact, high-uncertainty situations using signals available today.

1. Headline

A short description offering you a glimpse into future changes.


2. Temporal and Emotive Tags

A label explaining both when in the future this scenario is set and whether it is optimistic, neutral, pessimistic, or catastrophic.

3. Narrative

The descriptive elements of our imagined world, including the developments leading us to this point in our future history.

Scenario sources: The Future Today Institute uses a wide array of qualitative and quantitative data to create our scenarios. Some of our typical sources include patent filings, academic preprint servers, archival research, policy briefings, conference papers, data sets, structured interviews with experts, conversations with kids, critical design, and speculative fiction.



Artificial Intelligence Summary

- + Natural language processing is an area experiencing high interest, investment, and growth.
- + No-code or low-code systems are unlocking new use cases for businesses.
- + Amazon Web Services, Azure, and Google Cloud's low-code and no-code offerings will trickle down to everyday people, allowing them to create their own artificial intelligence applications and deploy them as easily as they could a website.
- + The race is on to capture AI cloudshare—and to become the most trusted provider of AI on remote servers.
- + The AI community still operates using a closed-source model. Researchers' reluctance to publish their full code leads to less transparency and reproducibility, and makes accountability murky.
- + Natural language processing algorithms— typically used for text, words, and sentences—are being used to interpret genetic changes in viruses.
- + COVID-19 accelerated the use of AI in drug discovery last year. The first trial of an AI-discovered drug is underway in Japan.
- + AI plays key roles in synthetic biology, genetics, and medical imaging; predicting the spread of disease; and improving patient health outcomes.
- + New artificial nervous systems use AI and neural implants.
- + The SuperGLUE benchmark, which measures AI's human language ability, will likely be surpassed by the end of 2021.

Artificial Intelligence Summary

- + Gaussian processes, the gold standard for many real-world modeling problems, are becoming more accurate and easier to train.
 - + AI researchers are leaving academia for corporations at an alarming pace.
 - + Generative adversarial networks assist artists and musicians in new forms of creative expression.
 - + A new wave of AI nationalism is rising as governments institute new restrictions on M&A and investment activity.
 - + Several countries will launch national AI strategies in 2021 and 2022.
 - + New measures to regulate the creation and distribution of deepfakes will be introduced throughout 2021 and 2022.
- + Technical alliances that help drive future R&D could also challenge existing geopolitical alliances.
 - + Future wars will be fought in code, using data and algorithms as powerful weapons.
 - + We continue to fail to see China's growing AI proficiency as a military, economic, and diplomatic threat.
 - + New software could be viable for 100 years by using AI to adapt to changes around it.

14TH YEAR ON THE LIST

Artificial Intelligence



AI is a force multiplier for every industry.

KEY INSIGHT

Artificial intelligence represents the third era of computing, generally defined as the ability for a machine to perform cognitive functions as well as or better than humans. Such functions include perception, learning, reasoning, problem-solving, contextual understanding, making inferences and predictions, and exercising creativity.

EXAMPLES

AI is now used across most industries. It solves business problems, detects fraud, improves crop yields, manages supply chains, recommends products, and even assists designers and writers in their work. AI can predict call volume in customer service centers and recommend staffing levels; it also predicts the emotional state and behavior of the person calling to help companies anticipate desirable solutions. AI automates the process for drug discovery, which ultimately led to faster COVID-19 vaccine candidates. Because AI is so broad, we have identified different themes within the discipline that you should be following. You will also find the technology intersecting with other trends throughout this report.

DISRUPTIVE IMPACT

The convergence of groundbreaking research, business use cases, the explosive growth of data, and improvements in computing power and storage are enabling advances in AI. The global artificial intelligence market is expected to grow at a compound annual growth rate of 42.2% from 2021 to 2027.

EMERGING PLAYERS

- Broad Institute
- Clarifai
- Clearview AI
- DeepMind
- Disperse
- Graphcore
- HiSilicon Technologies
- Kasisto
- LabGenius
- Mohamed bin Zayed University of Artificial Intelligence
- Niantic
- Nvidia
- OpenAI
- OpenMined
- Persado
- PolyAI
- Recursion
- SenseTime
- Scale AI
- Syntiant



An Executive's Guide to AI

What You Need To Know

In its most basic form, artificial intelligence is a system that makes autonomous decisions. AI is a branch of computer science in which computers are programmed to do things that normally require human intelligence. This includes learning, reasoning, problem-solving, understanding language, and perceiving a situation or environment. AI is an extremely large, broad field that uses its own computer languages and relies on computer networks modeled on our human brains.

Machine Learning

AI pioneer Arthur Samuel popularized the idea of machine learning in 1959, explaining how computers could learn without being explicitly programmed. This would mean developing an algorithm that could someday extract patterns from datasets and use those patterns to predict and make real-time decisions automatically. It took many years for reality to catch up with Samuel's idea, but today machine learning is a primary driver of the growth in AI.

Machine learning uses data to make predictions and recommendations on how to achieve stated goals. Types of machine learning include supervised, unsupervised, and reinforcement.

In supervised learning, an algorithm uses training data to learn the relationship between established parameters—inputs and outputs. Humans supervise, tweaking and adjusting systems as they work. Supervised learning is used when

teams know how to classify the input data and what they are trying to predict but can get accurate results much more quickly by relying on an algorithm rather than a human. Understanding what product features would most likely drive new purchases is an example of a business use case for supervised learning.

In unsupervised learning, data is provided to an algorithm without specific output parameters. For example, if a researcher doesn't know quite what to do with a large dataset, an algorithm could determine patterns, classify data, and make recommendations without a human supervisor. Unsupervised learning has been used during the pandemic to find patterns in how the virus is spreading throughout communities.

In reinforcement learning, an algorithm learns to perform a task by repeatedly running calculations as it attempts to accomplish a stated goal. Reinforcement learning is used when there isn't enough training data, when the researcher is try-

ing to learn about an environment (such as a complex financial portfolio), or when the researcher needs to find greater levels of optimization. It has a tremendous number of business use cases, ranging from real-time dynamic pricing models to high frequency trading algorithms to the systems that run self-driving cars.

Deep Learning

Deep learning is a relatively new branch of machine learning. Programmers use special deep learning algorithms alongside an enormous corpus of data—typically many terabytes of text, images, videos, speech, and the like. Often, these systems are trained to learn on their own, and they can sort through a variety of unstructured data, whether it's making sense of typed text in documents or audio clips or video. In practical terms, deep learning's emergence means that more and more human processes will be automated, including the writing of software, which computers will soon start to do

themselves. For example, once a system learns what an object looks like—say, an apple—it can recognize that object in all other images, even if it has only a partial view.

There are different types of deep learning models. The most common types include convolutional neural networks, recurrent neural networks, transformer neural networks, and generative adversarial networks (GANs).

A convolutional neural network (CNN) is multilayered, with a convolutional layer, a pooling layer, and a fully connected layer. Each one performs a different task using the data. The output is classification. If a researcher has 10,000 images and needs to extract data—to recognize particular faces, for instance—the CNN would run until information could be inferred. In business, CNNs are used for recognition: anomalies in medical imaging, faulty products on a production line, blight on crops.

An Executive's Guide to AI

Recurrent neural networks (RNNs) are multilayered neural networks that move and store information between input, hidden, and output layers. They are good at modeling sequence data for predictions. In business, they are used anytime the sequence of data matters, such as speech recognition and language translation. RNNs are used in digital assistants, to create captions for images, and to generate narrative reports using structured data (sports, financial).

GANs are unsupervised deep learning systems composed of two competing neural networks—a generator and a discriminator—that are trained on the same data, such as images of people. The networks compete against each other to perform a task—identifying the correct person—which results in optimizing overall performance. GANs are useful when researchers don't have enough data to train an algorithmic model. They are also used to create new, synthetic data. Deepfakes, which have become popular

in the past year, are generated using GANs. In design, GANs are tremendously useful: They can produce thousands of designs and recommend the best ones based on desired parameters. They can generate and modulate voices, faces, even gestures. Researchers from Nvidia, Massachusetts General Hospital, BWH Center for Clinical Data Science, and the Mayo Clinic collaborated on a GAN that generates synthetic MRIs showing cancerous tumors.

A **transformer** is a type of neural network architecture that learns what words mean when they appear in a particular context. Using “attention mechanism,” a transformer looks at an input sequence and determines at each step what other parts of the sequence are important. To date, transformers have mainly been used in natural language processing and generation.

Weak and Strong AI

There are two kinds of AI—weak (or “narrow”) and strong (or “general”). Narrow AI systems make decisions within very narrow parameters at the same level as a human or better, and we use them all day long without even realizing it. The anti-lock brakes in your car, the spam filter and autocomplete functions in your email, and the fraud detection that authenticates you for a credit card purchase—these are all examples of artificial narrow intelligence.

Artificial general intelligence (AGI) describes systems capable of decision-making outside of narrow specialties. Dolores in “Westworld,” the Samantha operating system in “Her,” and the H.A.L. supercomputer from “2001: A Space Odyssey” are anthropomorphized representations of AGI—but the actual technology doesn't necessarily require humanlike appearances or voices.

There is no single standard that marks the distinction between weak and strong


AI. This is problematic for researchers covering AI developments and for managers who must make decisions about AI. In fact, we have already started to see real-world examples of functioning artificial general intelligence. In 2017, researchers at DeepMind, a lab owned by the same parent company as Google, announced that AI had taught itself how to play chess, shogi (a Japanese version of chess), and Go (an abstract strategy board game)—all without any human intervention. The system, named AlphaZero, quickly became the strongest player in history for each game. The team has been publishing important discoveries at an impressively fast pace. Last year, the DeepMind team taught AI agents to play complex games, such as the capture the flag “game mode” inside the video game “Quake III Arena.” They, like humans, had learned skills specific to the game as well as when and how to collaborate with other teammates. The AI agents had matched human player ability using re-

inforcement learning, in which machines learn not unlike we do—by trial and error. While we haven't seen an anthropomorphic AI walk out of DeepMind's lab, we should consider these projects as part of a long transition between the narrow AI of today and the strong AI of tomorrow.



Enterprise

Enterprise Trends



By measuring certain biomarkers, AI can detect people's emotions and respond accordingly.

The Rise of MLOps

As machine learning matures and new applied business solutions emerge, developers are shifting their focus from building models to operating them. Within software, a set of best practices known as DevOps relies on tools, automation, and workflows to reduce complexity so that developers can focus on problems that need to be solved. This approach is now being used in machine learning. In 2020, some of the fastest-growing GitHub projects were MLOps, or projects that dealt with tooling, infrastructure, and operations. Going forward, MLOps will describe a set of best practices that combines machine learning, traditional DevOps, and data engineering.

Low-Code or No-Code Machine Learning

Machine learning is transitioning, as new platforms allow businesses to leverage the power of AI to build applications without the need to know specific code.

Businesses can turn their unruly datasets into structured data that can be trained, and they can build and deploy models with minimal skills. Create ML is Apple's no-code, drag-and-drop tool that lets users build custom models such as recommendation engines, natural processing systems, and text classifiers. Google's AutoML includes image classification, object detection, translation, and all sorts of pattern recognition tools. MakeML creates object detection. Applications have included tracking tennis balls during matches and automatically changing the colors of objects (such as flowers or dresses) in images. Last year, Amazon launched a no-code mobile and web app builder for Amazon Web Services (AWS). Microsoft Power Apps is a low-code application development environment on Azure.

Web-Scale Content Analysis


Mining very large, unstructured datasets is now easier thanks to advanced natural

language processing collection and classification. Trained to recognize keywords, special algorithms can rapidly sort, classify, and tag information to detect patterns. For example, a model trained to search for hate speech can detect bad actors in social networks. Machine translation generates training data for financial crime classification; last year, it reduced the amount of time needed for classification from 20 weeks (human analysts working alone) to two weeks.

Simulating Empathy and Emotion

AI can now measure biomarkers that suggest a person's emotional state, such as agitation, sadness, or giddiness. Precisely detecting human emotion is challenging, but companies with a large enough dataset are developing accurate models. Amazon's Rekognition API infers someone's emotions using facial recognition and physical appearance. Replika uses AI to evaluate voice and text, and over time it mirrors the user.

Enterprise Trends



Samsung's next Exynos system on a chip will have an AMD graphics processing unit (GPU).

Affectiva Human Perception AI analyzes complex human states using speech analytics, computer vision, and deep learning. For example, the automotive sector uses Affectiva's technology to detect a driver's emotional state—such as sleepiness or road rage—and make real-time suggestions to improve their driving.

Artificial Emotional Intelligence

Research teams at Loving AI and Hansen Technologies are teaching machines unconditional love, active listening, and empathy. In the future, machines will convincingly exhibit human emotions such as love, happiness, fear, and sadness. It begs the question: What is an authentic emotion? Theory of mind refers to the ability to imagine the mental state of others. This has long been considered a trait unique to humans and certain primates. AI researchers are working to train machines to build

theory of mind models of their own. This technology could improve existing AI therapy applications such as WoeBot, a clinical therapy chatbot. By designing machines to respond with empathy and concern, digital assistants such as Alexa will increasingly become a part of one's family. This technology could eventually end up in hospitals, schools, and prisons, providing emotional support robots to patients, students, and inmates. According to health service organization Cigna, the rate of loneliness in the U.S. has doubled in the past 50 years. Two years ago, former U.K. Prime Minister Theresa May created a new cabinet position, the world's first Minister of Loneliness. In our increasingly connected world, people report feeling more isolated. Future governments struggling with a massive mental health crisis, such as South Korea, may turn to emotional support robots to address the issue at scale.

Serverless Computing

AWS, Alibaba Cloud, Microsoft's Azure, Google Cloud, and Baidu Cloud are rolling out new offerings and packages for developers with the goal of making it easier and more affordable for a wide swath of AI startups to launch their ideas into the marketplace. AWS Lambda lets teams run code for virtually any type of application or back-end service—without provisioning or managing servers or hands-on administration. The Azure Functions architecture supports myriad programming languages, scales on demand, and charges only for active compute time. Some engineers worry that such serverless systems require them to surrender too much control.

Expert Insight



Emotion AI Will Power the Empathy Economy, but AI Still Needs to Work

Dr. Rana El Kaliouby

CEO of Affectiva

Emotion AI will power the empathy economy, but AI still needs to work.

The COVID-19 pandemic has meant that more than ever, we rely on video conferencing to connect virtually—working remotely, learning from home, and in our social lives. But there's a big problem: These technologies are emotion blind.

When we communicate in person, we convey so much more than the

words we say: We express ourselves through nonverbal cues from our faces, voices, and body language. But technology is not designed to capture the nuances of how we interact with those around us.

AI may be the answer to preserving our humanity in virtual environments. Specifically, Emotion AI—software that can understand nuanced human emotions and complex cognitive states based on facial and vocal expressions—can address some of technology's shortcomings in light of the pandemic, and we'll see companies using it for new use cases, such as:

1. Video conferencing and virtual events

Emotion AI can provide insight on how people are emotionally engaging in a virtual event or meeting. This provides presenters with valuable audience feed-

back, gives participants a sense of shared experience, and can help companies take a pulse on collective engagement during this stressful time.

2. Online learning

Emotion AI can give feedback on how students are engaging with online educational materials and lectures, flagging if they're confused, stressed, or bored. This becomes especially important during the pandemic as so many students are learning online and suffering from "Zoom fatigue."

3. Telehealth

Emotion AI can create more meaningful discussions and trust between patients and healthcare providers as telehealth appointments are replacing in-person visits. And, a data-driven analysis

of a patient's emotional wellbeing provides a quantitative measure of mental health that goes beyond self-reporting on a rating scale of 1-10.



AI researchers are now taking data that's already been collected and synthesized and using it to create brand new data.

What's on the horizon for AI:

1. Data synthesis.

AI algorithms are built on deep learning, but they can only work accurately when they're trained and then validated on massive amounts of data. That includes data that are diverse and truly representative of the situations the algorithm will encounter in the real world.

But companies developing AI often are challenged in getting access to the right kinds of data and the necessary volumes of data. That's where data simulation and data synthesis methodologies will come into play, addressing those problems. AI researchers are now taking data that's already been collected and synthesized and using it to create brand new data.

Take the automotive industry. The industry seeks to develop

advanced driver safety features and personalize transportation, and to achieve that car makers want to better understand what's happening with people inside of a vehicle. Getting that real-world data is difficult, expensive and time-consuming. But data synthesis is not. For example, a video of a person driving a car can become data that lets researchers create new scenarios, such as simulating the person turning her head, or wearing a hat or sunglasses.

2. The need for diversity, equity and inclusion.

As AI becomes more mainstream, the tech is taking on roles that were traditionally done by humans and changing how we interact with one another. For the technology to work for all of us, diverse teams must build those applications. Indeed, the number



Emotion AI can give feedback on how students are engaging with online educational materials and lectures, flagging if they're confused, stressed, or bored.

one issue to look out for is the risk for bias. Unfortunately, we've seen many instances in which AI has been biased against minority groups. Not only is this unethical; it's also bad for business. If AI can't work for all people as it's intended, there's little benefit to using it in the first place.

3. The challenge of power asymmetry.

Powerful technologies like AI are often in the hands of large corporations and governments, and this poses a number of challenges. The value that users receive from the technologies don't always measure up to the value that companies gain from the tech's user data. Also, those corporations or governments can determine a technology's distribution and who has access to it.

This can have adverse implications for social and economic mobility. People with access to certain types of AI will be able to work more efficiently and will have a leg-up on those who don't have access. I worry about the impact this can have on communities and populations that are already disadvantaged, because AI could continue to widen that gap.

We need to create guidelines to ensure AI is applied in an equitable way. The technology has the potential to improve people's lives and solve societal problems, but if we don't start thinking about power distribution now, we risk institutionalizing AI in a way that may exacerbate inequalities.

**

Dr. Rana El Kaliouby is the Co-founder and CEO of Affectiva, the pioneer of Emotion AI. Rana invented the company's award-winning emotion recognition technology. Prior to founding Affectiva, el Kaliouby was a research scientist at the MIT Media Lab where she spearheaded applications for facial coding to benefit mental health, autism, and other research areas. Born and raised in Cairo, she received degrees in computer science from the American University in Cairo and a Ph.D. from the computer laboratory, University of Cambridge.



Enterprise Trends

AI in the Cloud

Corporate leaders within the AI ecosystem have been racing to capture AI cloudshare—and to become the most trusted provider of AI on remote servers. Enterprise customers are likely to stick with their initial vendor, because machine learning systems get better over time, the more data they amass. For that reason, the competition is furious, even though it's still early. In the West, the field is led by Amazon, Microsoft, and Google, followed by companies including Apple, IBM, Salesforce, SAP, and Oracle. In Asian markets, Alibaba and Baidu dominate the AI cloud, although in January 2020, telecom equipment and smartphone maker Huawei announced a management change to focus on what it calls a “full-stack cloud platform.” It's a \$250 billion industry and quickly growing. New York University Stern School of Business professor Arun Sundararajan says it best: “The prize will be to become the operating system of the next era of tech.”

AI at the Edge

AI-driven processing and decision-making that occurs closer to the source of data generation, as opposed to in the cloud, is a technique known as “edge computing.” The Internet of Things and its billions of devices, combined with 5G networking and increased computing power, has made large-scale AI at the edge possible. Processing data directly on devices will be important in the future for health care, automotive, and manufacturing applications because it's potentially faster and safer. Apple spent \$200 million to acquire Xnor.ai, a Seattle-based AI startup focused on low-power machine learning software and hardware. Microsoft offers a comprehensive toolkit called Azure IoT Edge that allows AI workloads to be moved to the edge.

Advanced AI Chipsets

Today's neural networks have long required an enormous amount of computing power, have taken a long time

to train, and have relied on data centers and computers that consume hundreds of kilowatts of power. That is all starting to change. Enter the SoC, or “system on a chip.” Big tech companies including Huawei, Apple, Microsoft, Facebook, Alphabet, IBM, Nvidia, Intel, and Qualcomm, as well as startups Graphcore, Mythic, Wave Computing, SambaNova Systems, and Cerebras Systems, are all working on new systems architecture and SoCs—some of which come pretrained. In short, this means that the chips are more readily able to work on AI projects and should promise faster and more secure processing. Projects that might otherwise take weeks could instead be accomplished in a matter of hours. Cerebras has built an AI chip with 1.2 trillion transistors, 400,000 processor cores, 18 gigabytes of SRAM, and interconnects (tiny connection nodes) that can move 100 quadrillion bits per second. That's an astounding amount of components and power. As of November 2020, Amazon's homegrown AI chip AWS Inferentia now powers AI-

exa's back-end services rather than chips designed by Nvidia. The AI chip market will quadruple to \$6.7 billion in 2022, from \$1.66 billion in 2018, according to market research firm Tractica. Marketing pretrained chips to businesses will speed up commercialization and further R&D. But if the various device manufacturers all start creating unique protocols, developers may struggle with too many different frameworks. We anticipate an eventual consolidation, pitting just a few companies—and their SoCs and languages—against one another.

Digital Twins

Digital twins are virtual representations of real-world environments, products, or assets for a variety of purposes. Manufacturers use digital twins to manage the performance and effectiveness of machines and plants, while city planners use them to simulate the impact of new developments and roads. The Singapore government uses them for urban oper-

ations. Siemens MindSphere supports digital twins for a number of industries. As low-code and no-code systems become more prevalent, companies should be able to build and deploy digital twins to simulate a wide array of processes, which will lead to reduced spending on modernization efforts.

Spotting Fakes

In the past year, researchers showed how AI could be used to compose text so good that humans couldn't tell it was machine written. The team at OpenAI demonstrated the many reasons why this was problematic, from mass-generating salacious social media posts and fake reviews to forging documents by world leaders. It turns out that AI can also be used to detect when text was machine generated, even if we humans can't spot the fake. That's because an essay written by AI tends to rely on statistical patterns in text and doesn't have much linguistic variation. Researchers at the MIT-IBM

Enterprise Trends



Facebook launched the first AI model that translates 100 languages without relying on English data. (Image credit: Facebook.)

Watson AI Lab and Harvard University developed the Giant Language Model Test Room (GLTR), which looks for words that are likely to appear in a particular order. This technology can be used to detect forgery, intentional records falsification, email phishing campaigns, and corporate espionage.

Natural Language Processing for ESGs

Companies are moving toward new environmental, social, and governance (ESG) criteria—a set of standards increasingly used by investors to evaluate their investments. ESG standards must be quantified and explicitly stated, but measuring performance can be difficult because many intangibles or abstract concepts are involved. Natural language processing is being used to identify, tag, and sort documentation from various sources about a company’s ESG reputation (on issues such as labor practices, community impact, diversity, and inclusion).

Intelligent Optical Character Recognition

An ongoing challenge is getting machines to recognize the various ways we express ourselves in writing. Optical character recognition (OCR) works in fixed, recognizable formats such as highway signs and the text from a book. But often, OCR isn’t smart enough to recognize different fonts, unique notations, or spreadsheets with fields specific only to one company. Researchers are training AI systems to recognize patterns, even if they show up in unusual places. For example, the AWS Textract system now recognizes both text and context specific to a company or business unit.

Robotic Process Automation

Robotic process automation (RPA) can automate certain tasks and processes within offices and allow employees to spend time on higher-value work. It’s the most commonly deployed AI technique

among enterprise companies. Google’s Duplex is a good example; it’s a bot designed to make routine phone calls. Amazon uses RPA to sift through résumés and prioritize top candidates. In banking, Blue Prism and Automation Anywhere help staff with repetitive work functions. RPA will eventually augment staff and shift productivity into higher gear.

Massive Translation Systems

In 2020, Facebook launched a new open-source AI language model called M2M-100 that can translate 100 languages. Facebook’s AI lab trained the model using 7.5 billion sentence pairs gathered automatically from the web. (Surprisingly, Facebook did not use its own data for this project.) The FastText language model identified the language, and an unsupervised learning model matched sentences by their meaning. The goal was to improve simultaneous language translation.



Enterprise Trends

Predicting Systems and Site Failures

Computer vision can anticipate and identify failures in physical locations. High-tech factories, airline manufacturers, and construction sites use image recognition systems to monitor projects and automatically warn of problems. This is accomplished by comparing data from the real world to that of a digital twin.

Liability Insurance for AI

Who's to blame when machines behave badly? When the machine learning system in Uber's self-driving car failed and killed an Arizona pedestrian, the company was likely not covered under traditional cyber insurance. As businesses rush to build and implement AI products and processes, they must plan ahead for emerging risks. For example, what happens if machine learning makes a company vulnerable to attackers who inject fake training data into a system? What if a health care company's AI

misinterprets data and neglects to identify cancer among certain patients? These are the kinds of problems that could put a company at risk of lawsuits. New insurance models will help address these issues. Underwriters are starting to include AI under cyber insurance plans. Specialty insurers such as LaPlaya Insurance now offer insurance for AI applications.

Manipulating AI Systems for Competitive Advantage

Amazon, Google, and Facebook have all come under fire in the past few years for manipulating their search systems to prioritize results that are more profitable for their companies. For example, Google has been accused of de-ranking websites and promoting news stories from preferred partners. Late in 2019, researchers found that Amazon had optimized its search algorithm to boost the visibility of Amazon's own brands. Tweaks to search algorithms have a significant impact on what internet users see, whether that

is news, products, or advertising. This resulted in the ongoing antitrust lawsuits filed against the companies.

Global Rush to Fund AI

There is a global race to fund AI research and to acquire AI startups. In the first quarter of 2020, 285 U.S.-based AI startups had raised \$6.9 billion, according to the National Venture Capital Association. Investment waned as Covid became a global pandemic, but tech giants including Apple, Google, and Microsoft are still acquiring AI companies, while non-tech companies are gobbling AI startups too: McDonald's acquired personalization platform Dynamic Yield, while Nike acquired inventory management company Celect and guided shopping experience platform Invertex.

Algorithm Marketplaces

In the 2010s, big tech companies, startups, and communities of developers used algorithm marketplaces to share and sell

their work. In 2018, Microsoft paid \$7.5 billion to buy GitHub, a popular development platform allowing anyone to host and review code, to collaborate with other developers, and to build all kinds of projects. AWS hosts its own marketplace, offering models and algorithms for computer vision, speech recognition, and text—and its base of sellers includes Intel, CloudSight, and many others. (Think of AWS Marketplace as an Amazon for algorithms and models.) There are marketplaces for generalists, like GenesisAI and Algorithmia, where developers can upload their work and receive payment when others pay to access it. Now there are specialized marketplaces for specific use cases: Nuance AI Marketplace developed a single API to connect its algorithms to radiologists at 6,500 health care facilities. Quantiacs allows developers to build algorithmic trading systems, and it matches their algorithms with capital from institutional investors. Bonseyes is a European-specific marketplace to buy and sell AI tools.

100-Year Software

Traditional software has a short and unpredictable shelf life compared with other engineering tools. This leads to headaches and costly upgrades, often with downtime. As a result, companies and government agencies attempt to keep pace with the evolution of technology by maintaining systems rather than evolving. Libraries, data formats, and protocols can all become outdated quickly, creating vulnerabilities in critical systems. Since 2015, the Defense Advanced Research Projects Agency (DARPA) has funded research to make software viable for more than 100 years. These systems would use AI to dynamically adapt to changes in environments and resources. They require a novel approach to design, using AI to discover and make visible the application's operations and interactions with other systems.





Rage Against the Machine

Mid-future neutral scenario

From screaming into pillows to pounding punching bags, humans have developed numerous ways to air our frustrations. The hope is that if we act out against inanimate objects, we're less likely to act out against our fellow humans, risking harm or trauma. But what if there were a humanoid stand-in that could absorb our aggressions in a more cathartic, and ultimately beneficial way? As AI begins to achieve convincing emulations of human personalities, a new type of avatar emerges, algorithmically designed to provide a responsive therapeutic outlet for aggression.

Users can program the avatar to look and act like a figure from their life (a boss, a partner, a rival) for whom they harbor pent-up feelings, allowing users to express themselves freely without threatening their real-life relationships or risking legal repercussions. The AI persona could even be assigned to a surrogate robotic body, letting the user act out physical aggressions. But as the technology grows in popularity, designers must keep watch that what they've created doesn't normalize and increase the rate of interhuman conflict, instead of alleviating it.

Health, Medicine & Science

Health, Medicine, and Science Trends



The U.S. Food and Drug Administration approved IDx-DR, the first autonomous AI system to provide a diagnostic decision.

AI Speeds Scientific Discovery

Running experiments with several variables often requires tiny, methodical tweaks to measurements, materials, and inputs. Graduate students might spend hundreds of tedious hours repeatedly making small adjustments until they find a solution—a waste of their cognitive abilities. Research labs now use AI systems to speed the process of scientific discovery. Materials scientists at the University of British Columbia now rapidly test a new kind of solar cell and log results using a robot overseen by an AI algorithm. Based on the results of each experiment, an algorithm determines what to change next. A 9- to 12-month process was completed in five days. Google’s DeepMind developed a way of testing and modeling the complex folding patterns of long chains of amino acids, solving a problem that has vexed scientists for many years. DeepMind’s system, AlphaFold, will allow scientists to synthesize new drugs to treat diseases and develop enzymes that might someday break down pollution.

AI-First Drug Discovery

COVID-19 accelerated the use of AI in drug discovery. An international team crowdsourced a Covid antiviral by synthesizing candidates for 2,000 molecules in less than 48 hours—a process that likely would have taken human researchers a month or longer. In Japan, the first phase of a clinical trial for an AI-designed drug to treat obsessive-compulsive disorder showed a positive result. The drug, DSP-1181, acts as an agonist to the receptor for serotonin, a signaling molecule in the brain that mediates mood. The project used AI techniques to generate tens of millions of potential molecules to try against the serotonin receptor and sift through the candidates to decide which ones to prioritize for synthesis and testing.

AI-first drug startups are attractive to investors. Recursion raised \$121 million in 2019 before spinning off CereXis, a new independent entity to study rare brain cancers. Nearly every major phar-

maceutical company has inked deals with AI drug discovery startups, including Johnson & Johnson, Novartis, Merck, AstraZeneca, and GlaxoSmithKline. Much of the potential in AI stems from deep learning’s ability to sort through huge volumes of information and learn and extrapolate from that information. The upshot: AI can think faster than humans—sorting data in months versus years—and see patterns that we may not. Still, drug discovery is tricky, because the algorithms rely on drug targets that must be published in research journals. Most data about potential compounds isn’t readily available, and when it is, it isn’t always complete or reliable. Filling the gaps and cleaning that data takes time and money. It also requires data sharing—and most drug data is proprietary and locked up by big drugmakers. Using algorithms for drug development also brings up a host of ethical questions. Will bias invade drug discovery much like it has other arenas of AI, thereby marginalizing certain patients or diseases? Do



Health, Medicine, and Science Trends

algorithms need their own clinical trials? Could AI be used to take shortcuts and undermine the value of the science being done inside the laboratory? Advocates say AI will make drug development and clinical trials more efficient, thereby cutting drug prices and paving the way for more personalized medicine.

AI Improves Patient Outcomes

New medical algorithms address the level of patient care in the U.S. Different patients experience symptoms differently, and their care is based on how they describe their symptoms and how those symptoms are interpreted by doctors. For example, assessing the severity of arthritic pain is challenging. There is a standard scoring system to rate pain, which looks at the amount of structural damage and missing cartilage, but data from the National Institutes of Health found that Black patients' pain is underscored. It's likely that the system itself, called the Kellgren-Lawrence Grade (KLG),

was riddled with bias when it was first developed using primarily white British patients. Researchers are training deep learning models instead, and finding gaps in patient care.

Deep Learning Applied to Medical Imaging

Radiologists and pathologists increasingly rely on AI to assist them with diagnostic medical imaging. Last year, new U.S. Food and Drug Administration approvals allowed new products to be used widely in hospitals and clinics. So far, most of the approved devices augment (rather than fully automate) the process of reviewing images and making diagnoses. But emerging autonomous products are making their way into clinical settings. IDx-DR is an AI-enabled device that detects diabetic retinopathy using retinal images. Caption Health uses AI to capture ultrasound images of the heart that expands who can read such scans. Nurses would just need a few days of training on

the software. In a trial, an AI screening system from Google Health and DeepMind outperformed human radiologists and reduced false positives in two large, clinically representative datasets from the U.S. and U.K.

NLP Algorithms Detect Virus Mutations

Natural language processing (NLP) algorithms, which are typically used for text, words, and sentences, are being used to interpret genetic changes in viruses. Protein sequences and genetic codes can be modeled using NLP techniques—and can be manipulated the way you'd write words and sentences in word processing software. At MIT, computational biologists used NLP to solve a vexing problem when developing new vaccines. "Viral escape" is the ability for a virus to mutate and evade the human immune system and cause infection. Massachusetts Institute of Technology researchers modeled viral escape using NLP to identify how

the virus might look different to the immune system. The approach is similar to changing words in a sentence to change its meaning. For example: "I laughed at the clown" versus "I cried at the clown." By using this kind of modeling before mutations occur, public health officials could strategize and potentially prevent new viral spreads.

Diagnostics Without Tests

The Covid Symptom Study—created by doctors and scientists at Massachusetts General Hospital, the Harvard T.H. Chan School of Public Health, King's College London and Stanford University School of Medicine, and health science company ZOE—developed an app to study Covid symptoms and track the spread of the virus. It collects and uses AI to analyze data from 4 million global contributors to discover new symptoms, predict Covid hot spots, and eventually predict Covid cases without physical tests.

Protein Folding

In November 2020, DeepMind's AI made a big announcement: It successfully determined a protein's 3D shape from its amino-acid sequence. The program, AlphaFold, outperformed an estimated 100 teams in a biennial protein-structure prediction challenge called Critical Assessment of Structure Prediction (CASP). Predicting protein structures has long vexed biologists. AlphaFold had previously bested other teams, but it worked so quickly and so accurately at last year's CASP that it signaled a near future when the technology could be used regularly by other scientists.

Dream Communication

Scientists discovered how to establish two-way communication channels between lucid dreamers. Lucid dreamers are aware that they are asleep and can steer their dreams. In four global studies, participants were outfitted with sensors attached to their heads and faces, and



Health, Medicine, and Science Trends



AI is used to increase the speed and efficiency of new drug discovery.

their data was fed into a computer that looked for patterns. Scientists verified a state of REM sleep, and then participants interacted with researchers using eye movements. The study proved that there are new ways to send and receive real-time information while dreaming.


Thought Detection

Deep neural networks are being used to analyze emotional states using wireless signals. Researchers at Queen Mary University of London used radio signals to measure subjects as they watched a video. The deep neural network analyzed subtle body movements—breathing rates, heart rate—to reveal otherwise hidden information. If this reminds you of the “Black Mirror” “Crocodile” episode, you’re not far off: Research labs are developing new technologies to read our minds. There are business implications: HR departments could determine what employees really think of company policies, lawyers could determine how jurors lean in a

case, and realtors could judge how serious a homebuyer is. But the ethics are, of course, concerning.



Developing AI systems based on biological models—or deep neural networks—is among the 10 fastest growing technologies in the U.S., as indicated by patent applications. The number of patents with deep neural networks grew 67% between 2016 and 2020, making it the top technology mentioned.

Leading Applicants, 2016–2020



Source: IFI Claims Patent Services, January 2021 study.





Deep Twins in the OR

Mid-future optimistic scenario

The success rate of complicated medical procedures skyrockets when hospitals construct AI-enabled digital twins of patients. These virtual facsimiles mirror the exterior and interior structure of the patient's living body, as well as its real-time bodily functions, thanks to ultra-high fidelity mapping technology. Doctors are thereby able to virtually "operate" on the digital twin, using sophisticated AI to simulate outcomes and determine the safest and most effective approach to surgery, minimizing risk to the patient's life and well-being. Should a complication arise, the algorithm adjusts and adapts the recommended course of action based on the new data—in cases of extreme emergency, the AI can even take over and automate certain surgical tasks. The patient experience is revolutionized and the statistical risk of operations is drastically reduced, making trips to the OR feel more like a routine doctor's visit.

Consumer

Consumer Trends



Replika is a programmable digital twin that you can deploy for your friends.

Zero UIs

Modern interfaces are able to do more for us with fewer direct actions—yet still captivate our attention. The average adult now makes more than 20,000 decisions a day—some big, such as whether or not to invest in the stock market, and some small, such as whether to glance at a mobile phone when the screen lights up. Zero user interfaces—otherwise known as ambient computing systems—promise to prioritize those decisions, delegate them on our behalf, and even autonomously answer for us, depending on the circumstance. Much of this invisible decision-making will happen without direct supervision or input from people. What makes ambient design so tantalizing is that it should require us to make fewer and fewer decisions in the near future. Think of it as a sort of autocomplete for intent.

Consumer-grade AI Applications

Low-code and no-code offerings from Amazon Web Services (AWS), Azure, and Google Cloud will start to trickle down to everyday people, who will create their own AI applications and deploy them as easily as they can a website. We're seeing a shift from highly technical AI applications used by professional researchers to more lightweight, user-friendly apps intended for tech-savvy consumers. New automated machine learning platforms make it possible for nonexperts to build and deploy predictive models. Platforms hope that in the near future, we'll use various AI applications as part of our daily work, just as we do Microsoft Office and Google Docs today.

Ubiquitous Digital Assistants

Digital assistants (DAs)—like Siri, Alexa, and their Chinese counterpart Tiān Māo from Alibaba—use semantic and natural

language processing, along with our data, to anticipate what we want or need to do next, sometimes before we even know to ask. Alibaba's highly advanced DA can not only interact with real humans but also deftly handle interruptions and open-ended answers. Similar to Google's Duplex, Tiān Māo can make calls on your behalf, but it also understands intent. So if you're trying to schedule an appointment and mention that you're usually commuting in the morning, the system infers that you won't be available then. In 2017, Future Today Institute's model correctly projected that nearly half of Americans would own and use a digital assistant by 2020. (An estimated 62% of Americans use digital assistants today.) Amazon and Google dominate the smart speaker market, but digital assistants can be found in many places. Thousands of applications and gadgets now track and respond to DAs. News organizations, entertainment companies, marketers, credit



Consumer Trends



Alibaba's voice assistant uses natural language processing.

card companies, banks, local authorities, political campaigns, and many others can harness DAs to both surface and deliver critical information.

Deepfakes for Fun

Faceswap is a free and open-source deepfake app powered by TensorFlow, Keras, and Python. Deep Art Effects offers desktop and mobile apps to turn images into stylized art. REFACE is a face swap app that morphs your face onto celebrity bodies and creates GIFs to share on social media. Jiggy is a deepfake that makes anyone dance. For now, they all result in images and GIFs that *look* like they've been manipulated—but with the technology becoming so easy to use, how long until we can't tell real from fake?

Personal Digital Twins

A number of startups are building customizable, trainable platforms capable of learning from you—and then representing you online via personal digital twins. In 2021, China's annual Spring Festival Gala on the country's state broadcaster (CCTV) included performances from synthesized celebrities. With an estimated billion people watching, the AI copies mimicked their human counterparts without pre-scripted behaviors, speeches, or routines. Meanwhile, Replika is a programmable digital twin that you can deploy for your friends. Molly, a Y Combinator-backed startup, answers questions via text. The near future could include digital twins for professionals across a range of fields, including health and education.



Research

Research Trends

Closed-Source Code

Code is important for reproducibility, accountability, and transparency, and it is a key to driving improvements in the greater AI community. But when academic researchers publish papers, they don't often include all of their code. The reason given: The code they used is intermingled with other proprietary research, and it therefore can't be released. Fewer than 15% of all academic papers on AI publish their full code, and some big names—DeepMind and OpenAI—notoriously leave theirs out, citing proprietary concerns.

Framework Consolidation


Google's TensorFlow and Facebook's PyTorch are two popular frameworks used by researchers, and the relative popularity of different frameworks typically mirrors trends in the commercial appli-

cation landscape. In the past four years, Facebook seems to have gained ground. Of the conference papers that mention the framework the researchers used, 75% cited PyTorch but not TensorFlow. Of the 161 researchers who published more TensorFlow papers than PyTorch papers, 55% of them switched to PyTorch, while only 15% moved in the other direction.

Cost of Training Models

It costs a lot to train a model. Several variables influence those costs, all of which have increased in the past few years. For example, it costs an average of \$1 per 1000 parameters. OpenAI's 175 billion parameter, GPT-3, likely cost more than \$10 million to train. For smaller research groups and companies, the costs are out of reach. Some in the AI community are instead allowing the big tech companies to pre-train and publish big models.

Most Active Institutions for AI Research



Source: <https://macropolo.org/digital-projects/the-global-ai-talent-tracker/>



Research Trends

NLP Benchmarks

The General Language Understanding Evaluation (GLUE) benchmark is a collection of resources for training, evaluating, and analyzing natural language understanding systems. It includes a benchmark of nine sentence- or sentence-pair language-understanding tasks built on existing datasets and selected to cover a diverse range of dataset sizes, text genres, and degrees of difficulty. It includes a diagnostic dataset designed to evaluate and analyze model performance with respect to a wide range of linguistic phenomena found in natural language. And it includes a public leaderboard so that researchers can track their performance. The human baseline score is 87, and between May 2018 and August 2020, natural language processing systems increased from 60 to 90.6, surpassing humans. The SuperGLUE benchmark is a new measurement of more difficult language understanding tasks, improved

resources, and a new public leaderboard. We predict that by the end of 2021, this new benchmark will also be surpassed.

Machine Reading Comprehension

For AI researchers, machine reading comprehension (MRC) has been a challenging goal, but an important one. MRC makes it possible for systems to read, infer meaning, and immediately deliver answers while sifting through enormous datasets. In 2019, China's Alibaba outperformed humans when tested by the Microsoft Machine Reading Comprehension dataset (or MS MARCO for short), which assessed its ability to use natural language to answer real questions posed by humans. Alibaba's system delivered answers to search queries posted by people to Microsoft's Bing, such as "How many carbs are in an English muffin?" and "How do you grow hops?"

AI Summarizing Itself

A new AI model can summarize scientific literature, including research about itself. The Allen Institute for Artificial Intelligence (AI2) used the model in Semantic Scholar, an AI-powered scientific paper search engine to provide a short summary of papers on AI. What makes this work impressive—and ultimately so useful—is that it is capable of compressing long papers with accuracy and efficiency.

No Retraining Required

Training robots to do more than one thing is difficult, but a new model pits identical robot arms against one another in a game (moving objects on a virtual table in specific ways) in which one robot challenges the other with increasingly difficult tasks. It's an example of multi-task learning, a deep learning model in which machines learn different skills as they progress. OpenAI's model allows a bot to solve new kinds of problems without requiring retraining.

Graph Neural Networks

Because we perceive scents using millions of sensory neurons in our brains, and because scents are multifaceted, predicting the way something will smell is incredibly complex. For example, how would you describe the smell of an orange? Sweet? Bright? Grassy? Each descriptor is unique. Classifying smell is tricky because it requires a multi-label system. Graph neural networks (GNNs) constitute a particular type of deep neural network that operates on graphs as inputs. GNNs are being used to detect smell—to predict odors at a molecular level—and for a wide array of chemical and biological processes. For example, researchers at the Broad Institute used them to discover antibiotic compounds that don't have toxic side effects.

Federated Learning

Federated learning is a technique that distributes machine learning to the edge. Introduced by Google researchers in

2016, it is a new framework that makes it possible for algorithms to use data on devices—such as mobile phones and smart watches—without compromising user privacy. Research in this space has dramatically increased.

GP Models


Gaussian processes (GP) are the gold standard for many real-world modeling problems, especially in cases where a model's success hinges on its ability to faithfully represent predictive uncertainty. GPs are becoming more accurate and easier to train, benefiting from neural network improvements.

GPT-3's Influence

The enormous AI that generates human-like language, GPT-3, was released by OpenAI last year. The text generator has written blog posts and code. It was pitted against college students in an essay writing contest, and the anonymized



Research Trends



The SuperGLUE benchmark will be broken by the end of 2021.

papers were graded by professors. GPT-3 earned mostly B's—the same as human students. But the AI has demonstrated a strong anti-Muslim bias. Researchers from Stanford University and McMaster University probed the neural network on tasks including prompt completion, analogical reasoning, and story generation. They found that a Muslim-violence bias appears consistently and creatively in many use cases of the model. It's yet another example of how bias creeps into our automated systems. Left unchecked, it will cause problems throughout society as AI matures.

Vokenization

Models like GPT-3 are trained on syntax and grammar, not creativity or common sense. So researchers at the University of North Carolina–Chapel Hill are combining language models with computer vision. Humans learn in a multilayered, multidimensional way, so a new technique called vokenization extrapolates

language-only data by contextually mapping language “tokens,” or the words used to train language models, to related images, or “vokens.” For example, auto-generated image captions often can't infer context. Vokenization would enable machines not just to recognize objects but to really “see” what's in them.

Machine Image Completion

If a computer system has access to enough images—say, millions and millions—it can patch and fill in holes in pictures. This capability has practical applications for professional photographers, as well as for everyone who wants to take a better selfie. Soon, if the foreground of a mountain is out of focus, or if your skin has an unsightly blemish, another version can be swapped in to generate the perfect picture. As such technology becomes commonplace, there will be significant biases and other pitfalls to navigate. For example, image generation algorithms routinely reflect

deeply culturally embedded racism and sexism. A few years ago, if you typed “CEO” into Google Images, the first result of a woman was CEO Barbie. In an experiment, researchers at Carnegie Mellon University trained a system to autocomplete images of men and women cropped below the neck. In pictures of men, the system autocompleted him wearing a suit. The system autocompleted women—including U.S. Rep. Alexandria Ocasio-Cortez (D-N.Y.)—wearing a low-cut top or bikini 53% of the time.

Predictive Models Using Single Images

Computer vision systems are getting smarter. Neural networks can predict geometry from a single color image. In 2019, the DeepMind team developed a generative adversarial network (GAN) that creates videos from images. For example: Imagine a photo of a person holding a basketball. Based on his posture, face, and other data within the

Research Trends

picture, the GAN figures out what likely happened next and generates a video clip of the action. Earlier, researchers at MIT’s Computer Science and Artificial Intelligence Laboratory (CSAIL) trained computers to predict what humans would do next using YouTube videos and TV shows such as “The Office” and “Desperate Housewives.” CSAIL’s system predicts whether two people are likely to hug, kiss, shake hands, or slap a high five. SinGAN is an unconditional generative scheme that can manipulate and enhance images—sketch a mountain, and it will produce a realistic-looking synthetic photograph. This research will someday enable robots to more easily navigate human environments—and to interact with us humans by taking cues from our body language. Retail, manufacturing, and education settings could be especially relevant.

Model-free Approaches to RL

Dreamer is a reinforcement learning (RL) agent that uses a world model to learn long-sighted predictions, employing backpropagation through model predictions. It can create models from raw images and learn from thousands of predicted sequences in parallel using a graphics processing unit (GPU). This new approach solves long-horizon tasks using an imagined world.

Real-time Machine Learning

One big challenge in AI is building machines that can proactively collect and interpret data, spot patterns and incorporate context, and ultimately learn in real time. New research into real-time machine learning (RTML) shows that it’s possible to use a continual flow of data and adjust models in real time. This signals a big change in how data moves, and in how we retrieve information. The National Science Foundation launched

a \$10 million grant program to catalyze research in this area, although all of the big tech companies are working closely to advance RTML too.

Automated Machine Learning (AutoML)

Some organizations want to move away from traditional machine learning methods, which are time-consuming and difficult and require data scientists, specialists in AI fields, and engineers. Automated machine learning, or AutoML, is a new approach: a process in which raw data and models are matched together to reveal the most relevant information. Google, Amazon, and Microsoft now offer a host of AutoML products and services.

Hybrid Human-Computer Vision

AI isn’t yet capable of fully functioning without human assistance. Hybrid intelligence systems combine humans and AI systems to achieve greater accuracy.

The U.S. Army Research Laboratory has a system that uses a brain-computer interface armed with computer vision technology and allows a person to rapidly see and sort images within her line of sight. CloudSight, a technology company specializing in image captioning, is working on a hybrid crowdsourced computer vision system. Microsoft researchers have proposed Pandora, a set of hybrid human-machine methods and tools for understanding system failures. Pandora leverages both human and system-generated observations to explain malfunctions related to input content and system architecture.

Neuro-Symbolic AI

The development of AI has been on two conceptual tracks since the 1950s: symbolic (machines that use a base of knowledge and rules that represent concepts) and non-symbolic (machines that use raw data to create their own patterns and representations of concepts).


Classic AI is the former, because it more closely represents how we understand human thought—and the original intent was to teach machines to think like us. Researchers are working on new ways to combine both learning and logic using neural networks, which would understand data through symbols rather than always relying on human programmers to sort, tag, and catalog data for them. Symbolic algorithms will aid the process, which should eventually lead to robust systems that don’t always require a human for training.

General Reinforcement Learning Algorithms

Researchers are developing single algorithms that can learn multiple tasks. DeepMind, the team behind AlphaGo, which learned how to play Go with the skill level of a human grandmaster, has developed an innovative new algorithm: AlphaZero. It is capable of achieving superhuman performance not only in



Research Trends



Using AI, researchers automated the task of converting live actor performances (left) to computer game virtual characters (right).

Go but in other games as well, including chess and shogi (Japanese chess). This one algorithm starts with no knowledge except for the rules of the game and eventually develops its own strategies to beat other players. In January 2020, DeepMind published new research showing how reinforcement learning techniques could be used to improve our understanding of mental health and motivation.

Continuous Learning

At the moment, deep learning techniques are helping systems learn to solve complex tasks in a way that resembles what humans can do—but those tasks are still specific, such as beating a human at a game. And they require a rigid sequence: Gather data, determine the goal, deploy an algorithm. This process requires humans and can be time-consuming, especially during early phases when supervised training is required. Continuous learning is more about au-

tonomous and incremental skill building and development, and researchers will continue to push the limits of what's possible in this field.

Proliferation of Franken-Algorithms

Algorithms are simply rules that define and automate the treatment of data. They are built using “if this, then that” logic that a computer can understand and process. Here's an easy example: If a website reader's IP address is based in Baltimore, the rules then allow that reader to freely access the site; if the IP address is based in Belgium, then the rules first show a GDPR screen stating privacy and cookie policies. While a single algorithm might be easily described and deployed as expected, systems of algorithms all working together can sometimes pose problems. Developers don't always know in advance how one algorithm will function alongside other algorithms. Sometimes, several teams of developers are working

independently on different algorithms and datasets, and they only see one another's work once it is deployed. This has been the cause of recent stock market glitches and e-commerce website wonkiness. It is especially challenging for big companies like Facebook, which have billions of algorithms working together at any given time.

Proprietary, Homegrown AI Languages

Python is a leading language with lots of pre-built libraries and frameworks. Julia, a language developed by Massachusetts Institute of Technology, is an open-source language that focuses on numerical computing. And of course there's Lisp, created by modern AI's foreparent John McCarthy in 1958. Companies are starting to build and release their own software packages now, as well as unique programming languages for AI applications. Uber released its own probabilistic programming language, Pyro, which it

wrote in Python. It's a move that signals likely fragmentation in the future of the AI ecosystem, not unlike the current iOS/Android rivalry or the long Mac/PC war. Businesses will find it increasingly cost-prohibitive and difficult to switch between AI frameworks and languages.

Talent

Talent Trends



The Mohamed bin Zayed University of AI opened in the United Arab Emirates last year.

AI Brain Drain

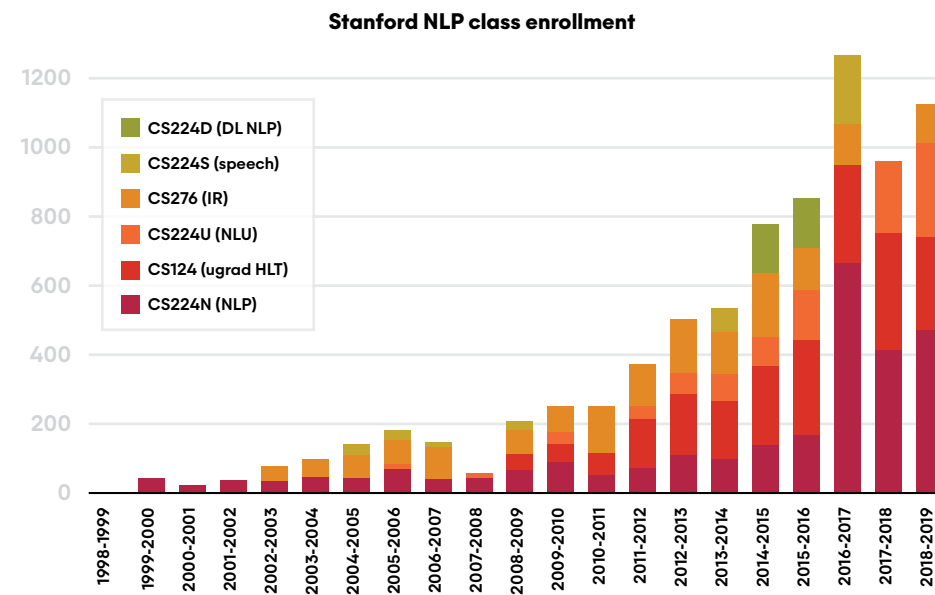
The brain drain of AI researchers out of academia and into corporations is growing at an alarming pace. The reason is simple: compensation packages. Top academics earn generous salaries and benefits, and they get to work in a similar tenured environment that's carefully cultivated to represent their experience in academia. Google, DeepMind, Amazon, and Microsoft hired 52 tenured and tenure-track professors from U.S. universities between 2004 and 2018. In return for their poaching, tech companies are endowing AI professorships at top universities. In some cases, professors take one- or two-year sabbaticals to work at tech companies and then return to their universities—but corporate benefits can be difficult to give up. In one infamous case, Uber poached an entire robotics lab from Carnegie Mellon University—40 professors and researchers in total. Facebook has aggressively hired superstars

like New York University professor Yann LeCun and Carnegie Mellon professor Jessica Hodgins. Poaching departments today could rob the future of future AI experts: Without great scholars, who will train the next generation of innovators?

AI Universities

Interestingly, new institutions solely focused on AI are launching around the world and offering postsecondary education. In the United Arab Emirates, the Mohamed bin Zayed University of Artificial Intelligence launched last year. MBZUAI is the world's first graduate-level, research-based AI university. Based in Abu Dhabi, it offers master of science in computer vision and Ph.D. programs in AI-related fields. Founded by Harvard University and University of California–Los Angeles faculty, Univ. AI is an online program for training in machine learning and AI.

Enrollment in Stanford University's natural language processing class is now 10 times the size it was in 2004.



Source: <https://nlp.stanford.edu/>



Talent Trends



I have always been convinced that the only way to get artificial intelligence to work is to do the computation in a way similar to the human brain. That is the goal I have been pursuing. We are making progress, though we still have lots to learn about how the brain actually works.”

— Geoffrey Hinton

Demand for AI Talent Growing

For many years, demand for AI talent has outpaced supply. In the U.S., there were nearly three times more AI-related job postings on Indeed last year than job views for AI-related roles. While schools are adding programs, increasing enrollment, and adding classes, there are just too many new needs for AI skills and nowhere near enough trained workers. As demand grows, the hiring process is taking longer and becoming more expensive. This is impeding growth at some companies, according to a 2020 Reuters study. Demand has also driven up wages. A recent study from Glassdoor estimated that average annual salaries for AI-related jobs rose 11% between 2017 and 2018 to \$123,069.

Corporate AI Labs

AI labs are located around the world, with concentrations in North America, Europe, and Asia. Facebook, Google, IBM, and Microsoft operate 62 labs dedicated to AI R&D, and the majority are outside of the U.S. because of access to talent. During the Trump administration, immigration restrictions and stringent visa requirements made it difficult to recruit talent into the United States, and overseas labs allowed companies to overcome that barrier. Most of those labs do basic AI research rather than product development. The G-MAFIA—Google, Microsoft, Amazon, Facebook, IBM, and Apple—spend \$76 billion on R&D annually. Collectively, their market cap exceeds \$6 trillion.

AI for Interviews

Recognition systems can now be deployed to watch you being interviewed and to gauge your enthusiasm, tenacity, and poise. Algorithms analyze hundreds of details, such as the tone of your voice, your facial expressions, and your mannerisms to best predict how you'll fit in with the culture of a community. Startups such as HireVue use AI systems to help companies decide which candidates to hire. But this kind of recognition technology has practical applications well beyond job interviews: It can detect when someone is likely to make a purchase—or attempt to shoplift—in a store, whether someone is lying, and whether someone is receptive to new suggestions and ideas. Unlike security cameras, which tend to have a light indicating they're recording,

algorithms work invisibly, which means that this is an area that could face regulatory scrutiny. Consumer advocacy organization Electronic Privacy Information Center filed a complaint with the U.S. Federal Trade Commission requesting an investigation into HireVue, alleging its tools produce results that are “biased, unprovable, and not replicable” through algorithmic models.



Creative

Creative Trends



Creative studio SoKrispyMedia develops short films featuring stick figures in battle. It relies on real-time rendering for photorealistic results.

Image credit: Nvidia.

Assisted Creativity

Generative adversarial networks (GANs) are capable of far more than generating deepfake videos. Researchers are partnering with artists and musicians to generate entirely new forms of creative expression. From synthesizing African tribal masks to building fantastical, fictional galaxies, AI is being used to explore new ideas. Last year, Nvidia launched GauGAN (named after post-Impressionist painter Paul Gauguin), a generative adversarial AI system that lets users create lifelike landscape images that never existed. The National Institute of Informatics in Tokyo built an AI lyricist, while Amazon released its DeepComposer system, which composes music “automagically.” These AIs are not ostensibly intended to replace artists, but rather to enhance their creative process.

Generative Algorithms for Content Production


Last year, the creators of “South Park” built an entire show using deepfakes. “Sassy Justice” aired on YouTube and featured a synthetic reporter, Fred Sassy, who looked a lot like former President Trump—with just a different enough voice and hairstyle to evade legal challenges. Episodes featured deepfakes of Al Gore, Mark Zuckerberg, Jared Kushner, and others. The open-source algorithm DeepFaceLab has been used by other artists and filmmakers. There was a Hulu commercial deepfaking sports stars, and several 2021 Super Bowl commercials, including an eerie appearance by Vince Lombardi, used deepfakes and synthetic media. OpenAI’s deep learning algorithm released a neural network called Jukebox that generates songs in a bunch of different styles and simulated voices that sound (sort of) like Elvis and others.

Generating Virtual Environments from Short Videos

Chip designer Nvidia is teaching AI to build realistic 3D environments from short video clips. The method builds on previous research on GANs. Nvidia’s system generated graphics based on open-source datasets used by the autonomous driving field. Using short clips segmented into various categories—such as buildings, sky, vehicles, signs, trees, or people—the GANs created new, different versions of these objects. The array of possible applications is vast. Automatically generated virtual environments could be used for fantasy and superhero movies and could bring down the costs of TV production and game development.



Creative Trends



Jukebox is a neural net that generates music, including rudimentary singing, as raw audio in a variety of genres and artist styles.

Automated Versioning

Journalists at Switzerland-based Tamedia experimented with generative techniques during their country's 2018 election. A decision-tree algorithm Tamedia named Tobi generated automated articles detailing vote results for each municipality covered by the private media group's 30 newspapers. It also produced content simultaneously in multiple languages. In total, Tobi published 39,996 different versions of election stories that averaged 250 words each. The articles carried a special byline alerting readers that they'd been written by an algorithm. With more experiments underway, we expect to see news and entertainment media companies developing multiple versions of the same content to reach wider audiences or to produce massive amounts of content at scale.

Automatic Voice Cloning and Dubbing

Anyone who's ever recorded a podcast is familiar with editing challenges such as guests talking over each other, interruptions from sirens and other background noises, and inconvenient sneezes. Those moments can stop a conversation cold. But what if you could edit the spoken word the way you edit a word document? That's the promise of AI companies including Resemble AI and Descript, which make it possible to clone voices. That means soon you might see a star like Phoebe Waller-Bridge in a movie and also hear her, in her own voice, speaking in Portuguese. There's obviously a dark side to this technology, however. Last year, hackers used voice cloning tools to trick an employee into thinking he was speaking on the phone to his CEO; he then transferred \$243,000 to a scammer's bank account.

Automatic Ambient Noise Dubbing

For some time, we've been training computers to watch videos and predict corresponding sounds in our physical world. For example, what sound is generated when a wooden drumstick taps a couch? A pile of leaves? A glass windowpane? The focus of this research, underway at MIT's Computer Science and Artificial Intelligence Laboratory, should help systems understand how objects interact with each other in the physical realm. Numerous projects are now underway to make it easier to automatically generate voices, videos, and even storylines.

Geopolitics & Defense

Geopolitics and Defense Trends



Cyber warfare will change the art of war in the coming decade.

AI Nationalism

Governments are instituting new restrictions on mergers and acquisitions and investment activity to ensure that AI developed by companies does not aid foreign adversaries. In 2020, Germany passed a law to allow the government to review or block investments or takeovers of robotics, AI, and semiconductor companies by non-EU-based firms. The foreign ownership threshold also dropped to 10% from 25%. The U.K. similarly strengthened its foreign investment policies for AI: Any M&A activity involving AI companies in which revenue exceeds 1 million pounds requires government scrutiny. Many acquisitions are being blocked as a result. The U.K. and U.S. blocked a takeover of Imagination Technologies by China's Canyon Bridge last year. A bipartisan bill proposed in the Senate, known as the Endless Frontier Act, explicitly frames AI as a race between the U.S. and China. It also creates a new technology directorate within the

National Science Foundation with \$100 billion in funding over five years and earmarks \$10 billion for local and regional tech hubs across America. Meanwhile, in China the Ministry of Science and Technology established 20-city AI pilot zones that should open by 2023. They will carry out AI-based policy experiments and social experiments, according to official government documents.

National AI Strategies

A new wave of countries will launch national AI strategies in 2021 and 2022. Iceland, Norway, Kenya, Chile, Colombia, and New Zealand are all developing national strategies. China passed its New Generation Artificial Intelligence Development Plan with aggressive benchmarks to become the world's dominant AI player within 10 years; France adopted a national strategy called AI for Humanity; Saudi Arabia has both a strategy and a legal framework for making robots citi-

zens; and the United Arab Emirates has a sweeping set of policy initiatives on AI and appointed Omar Sultan Al Olama as its minister of state for artificial intelligence. In the U.S., numerous public and private groups work independently on the future of AI on behalf of the nation. Those efforts, however, lack interagency collaboration and coordinated efforts to streamline goals, outcomes, R&D efforts, and funding.

AI as Critical Infrastructure

Government researchers are exploring ways to spearhead AI development for critical systems use: road and rail transportation systems; power generation and distribution; and predicting routes for public safety vehicles, such as ambulances and firetrucks. Rather than shunning AI systems, there is new interest in using the technology to prevent disasters and improve safety.

Geopolitics and Defense Trends

Nation-based Guardrails and Regulations

From self-driving car accidents to election interference through disinformation campaigns to political repression enhanced by facial recognition and automated surveillance, major events over the past few years have thrown into sharp relief the dangers of artificial intelligence. Few guardrails now exist for a technology that will touch every facet of humanity, and countries are racing to develop and publish their own AI strategies and guidelines. The European Union developed an AI Alliance and plan of cooperation between member countries, and Estonia is developing its own legal framework governing the use of AI within the country. Last year, China moved into position to lead the first set of global AI norms and standards. In 2019, the country published a report on technical standards that would allow companies to collaborate and make their

systems interoperable. The EU and the Organisation for Economic Co-operation and Development similarly published their own guidelines, and the Trump administration signed an executive order to spur the development of standards in the U.S. While these efforts could introduce new ways to safeguard against bias and to ensure trust, they also each attempt to create strategic advantages for stakeholders. As AI continues to develop according to different rules in China, the EU, and the U.S., one of the hallmarks of the field—global academic collaboration—could drastically decline.

Regulating Deepfakes

New measures to regulate the creation and distribution of deepfakes will be introduced throughout 2021 in the U.S. and elsewhere. A bill in Hawaii’s state legislature seeks to prohibit unauthorized deepfake apps and tools. If it passes, deepfaking would be considered a Class

C felony. Bills have also been introduced in California, Texas, and Massachusetts, and a number of federal bills are being discussed. These initiatives will likely be met with arguments that prohibiting deepfakes infringes on free speech rights.

Making AI Explain Itself

You’ve undoubtedly heard someone argue that AI is becoming a “black box”—that even researchers working in the field don’t understand how our newest systems work. That’s not entirely true. However, there is growing concern among computer scientists, journalists, and legal scholars who argue that AI systems shouldn’t be so secretive, and regulators are paying close attention. Broadly speaking, a few challenges must be overcome. Requiring transparency in AI could reveal a company’s trade secrets. Asking the systems to explain their decision-making processes as they work could also degrade the speed and qual-

ity of output. It’s plausible that various countries will enact new regulations requiring explainability in the coming years. Imagine sitting beside a genius mathematician who gives you correct answers in Italy, but receiving her answers across the border in France would mean asking her to stop and show her work—and every time she’s asked to share her answers in a new country.

New Strategic Technical Alliances

New strategic technical alliances between countries will help drive future R&D but could also strain existing geopolitical alliances or heighten tensions. Likely partners include the U.S., Germany, Japan, India, South Korea, the U.K., France, and Canada—leaving China and Russia to partner up separately. The latter two countries have already announced a technical alliance on satellites and deep-space exploration.

The New Mil-Tech Industrial Complex

In the past few years, some of the biggest AI companies in the U.S. have partnered with the military to advance R&D and find efficiencies. In fact, the public sector cannot advance its technology without help from outside companies. Plus, there is a lot of money to be made. The U.S. General Services Administration and the Department of Defense’s Joint Artificial Intelligence Center recently awarded a five-year, \$800 million contract to Booz Allen Hamilton for AI product development. The U.S. Army awarded Lockheed Martin a \$75 million contract for a machine learning cyber jamming pod that can be mounted on Humvees or drones. With a new focus on defense roadmaps that include AI components, startups working in high-resolution satellite imagery, computer vision, and unmanned aerial vehicles are attracting lucrative venture capital investment. Both Amazon and Microsoft made headlines



Geopolitics and Defense Trends



U.S. Rep. Ro Khanna (D-Calif., pictured here), along with U.S. Sens. Todd Young (R-Ind.) and Chuck Schumer (D-N.Y.), and U.S. Rep. Mike Gallagher (R-Wis.), unveiled the bipartisan, bicameral Endless Frontier Act.

over a \$10 billion, 10-year government tech contract called the Joint Enterprise Defense Infrastructure, or JEDI. Others, including IBM, Oracle, and Google, also competed to transform the military's cloud computing systems. Meanwhile, the CIA awarded Amazon a \$600 million cloud services contract, while Microsoft won a \$480 million contract to build HoloLens headsets for the Army. The contracts prompted employee protests. In 2017, the Department of Defense established an Algorithmic Warfare Cross-Functional Team to work on Project Maven—a computer vision and deep learning system that recognizes objects from still images and videos. The Department of Defense contracted Google to train the AI systems to analyze drone footage, but it turned out the Googlers assigned to the project didn't know they were doing military work. A high-profile backlash ensued: As many as 4,000 Google employees signed a petition objecting to Project Maven, and ultimately

dozens resigned. Eventually, Google said it wouldn't renew its contract on the project. The company launched a set of ethical principles governing its development and use of AI, including a provision that prohibits any systems from being used for "weapons or other technologies whose principal purpose or implementation is to cause or directly facilitate injury to people."

Algorithmic Warfighting

Future wars will be fought in code, using data and algorithms as powerful weapons. The current global order is being shaped by artificial intelligence, and the same countries leading the world in AI research—the U.S., China, Israel, France, Russia, the U.K., and South Korea—are also developing weapons systems that include at least some autonomous functionality. Israel uses autonomous drones for border patrol, while China developed stealth drones capable of autonomous

airstrikes. Future Today Institute analysis shows that the future of warfare encompasses more than traditional weapons. Using AI techniques, a military can "win" by destabilizing an economy rather than demolishing countrysides and city centers. From that perspective, China's unified march to advance AI puts the emerging superpower dangerously far ahead of the West.

China's AI Rules



President Xi Jinping has foregrounded AI as a major, long-term priority in China.

If you think of China as a country that copies rather than innovates, think again.

China is a global leader in artificial intelligence. Under President Xi Jinping, the country has made tremendous strides in many fields, but especially in AI. Businesses and the government have collaborated on a sweeping plan to make China the world's primary AI innovation center by 2030, and it's already making serious progress toward that goal. That plan is unlikely to be repealed by a new government; China abolished Xi's term limits and will effectively allow him to remain in power for life.

The country's enormous population of 1.4 billion offers researchers and startups there a command of what may be the most valuable natural resource in the future—human data—without the privacy and security restrictions common in much of the rest of the world. If data is the new oil, then China is the new OPEC. The kind of rich data the Chinese are mining can be used to train AI to detect patterns used in everything from education to manufacturing to retail to military applications. The Chinese

startup SenseTime is pioneering myriad recognition technologies, such as a system that provides advertisers real-time feedback on what people are watching; technology that can extract customer information and carry out statistical analysis in crowded areas like shopping malls and supermarkets; and simultaneous recognition of everything in a scene, including people, pets, automobiles, trees, or soda cans.

That gives China an incredible advantage over the West. It also gives three of China's biggest companies—Baidu, Alibaba, and Tencent—superpowers. Collectively, they're known as the BAT, and they're all part of the country's well-capitalized, highly organized AI plan.

The BAT is important to you even if you've never used them and don't do any business in China. That's because these companies are now well established in Seattle and around San Francisco, and they are investing significantly in U.S. startups. Baidu (a Chinese search-engine company often likened to Google) established AI research centers in Silicon Valley and Seattle, and Tencent (the developer of the mega-popular messaging app

WeChat) began hunting for American talent when it opened an AI lab in Seattle three years ago. It has since upped its stakes in companies including Tesla and Snap. The payoff for the Chinese is not just a typical return on investment—Chinese firms expect IP as well. China-based AI startups now account for nearly half of all AI investments globally.

Strategic panopticon

In late 2019, China began requiring all citizens to submit to facial recognition in order to apply for new internet or mobile services, and began requiring that telecom companies deploy AI to check the identities of people registering SIM cards. Chinese social media platforms require users to sign up with their real names. In Chinese schools, surveillance cameras with computer vision are used widely and track whether students are paying attention and whether they attempt to cheat or sleep. These and other national standards make it easier for the government to track its citizens. China's social credit system, an algorithmic reputation system developed by the government, standardizes assessments of citizens' and businesses' behavior and activity.

China's AI Rules

In 2020, numerous reports of abuse revealed that China turned its AI on the ethnic Uighur Muslim community. Huawei developed special AI software to identify Uighurs and alert local police. In 2021, China blocked social media platform Clubhouse after an open, democratic debate flourished on the platform about the plight of the Uighur community.

Risk profile

We have failed—and we are continuing to fail—to see China as a military, economic, and diplomatic threat when it comes to AI. China has already used its Belt and Road Initiative as a platform to build international partnerships in both physical and digital infrastructure, and it is making surveillance technologies available to countries with authoritarian regimes. Two Chinese companies—the state-controlled CEIEC and Huawei—built Ecuador's surveillance system, called ECU-911. The system promised to curb high murder rates and drug crime, but it was too expensive an investment. As a result, a deal was struck for a Chinese-built surveillance system financed with Chinese loans. It was a prelude to

a much more lucrative deal: Ecuador eventually signed away big portions of its oil reserves to China to help finance infrastructure projects. Similar package deals have been brokered in Venezuela and Bolivia.


China is quietly weaponizing AI, too. China's People's Liberation Army is catching up to the U.S. military, using AI for such tasks as spotting hidden images with drones. The Chinese military is equipping helicopters and jet fighters with AI. The government created a top-secret military lab—a Chinese version of DARPA—and it's building billion-dollar AI national laboratories. China's military is achieving remarkable AI successes, including a recent test of "swarm intelligence" that can automate dozens of armed drones.

The bottom line

There is a Chinese proverb that roughly translates to "forewarned is forearmed." Now that you know what's coming, reframe your thinking of China as simply the world's factory.

Chinese-educated Researchers Dominate NeurIPS

Chinese-educated researchers dominated the papers accepted at NeurIPS, a prestigious international AI conference, according to a new study from MacroPolo, a think tank based at the Paulson Institute, which promotes constructive collaborations between the U.S. and China. Nearly one-third of the papers accepted were from China—more than from any other country.



Source: <https://macropolo.org/digital-projects/the-global-ai-talent-tracker/>



Society

Society Trends



Timnit Gebru is a pioneering researcher on algorithmic bias.

Image credit: Wikimedia Commons.

Ethics Clash

On Dec. 2, 2020, Timnit Gebru, the co-lead of Google’s ethical AI team, posted a tweet saying she’d been fired. Known for groundbreaking research in bias and facial recognition, she is widely respected within the broader AI community. While the incident concerned a paper she co-authored and a decision by Google that it didn’t meet “our bar for publication,” she and others argued that Google’s ethics team was motivated by PR rather than progress. It set off a firestorm of criticism, and by February 2021 Google said it would change its diversity and research policies, “streamline its process for publishing research,” and change how sensitive employee exits are managed. Google’s head of AI, Jeff Dean, expressed his regret in an email to staff. To deal with its own ethical failures, Facebook launched an independent oversight board with the power to overrule content moderation guidelines—and even to overrule Mark Zuckerberg himself. (But just on

content.) In January 2021, the board made its first rulings on disputed content, overturning four out of the five cases it saw. But there are billions of posts on Facebook every day and an untold number of content complaints—which means the oversight board operates at the speed of traditional government. We anticipate many more ethics clashes in 2021.

Ambient Surveillance

What happens behind closed doors may not be secret for long, and executives should beware new ambient surveillance methods. Scientists at Massachusetts Institute of Technology discovered how to use computer vision to track data from what they call “accidental cameras.” Windows, mirrors, corners, houseplants, and other common objects can be used, along with AI, to track subtle changes in light, shadows, and vibrations. The result: We all may soon have X-ray vision capabilities—which may not be great news for

companies working on sensitive projects. Those working in information security and risk management should pay special attention to advances in computer vision.

Marketplace Consolidation

As much as the AI ecosystem booms, a rush of acquisitions means consolidation, too. Big companies now snap up startups long before they have time to mature—the average age at acquisition is 3 years old. Just nine big companies dominate the AI landscape: Google, Amazon, Microsoft, IBM, Facebook, and Apple in the U.S., and Baidu, Alibaba, and Tencent in China, with significant fortification and support from their country’s government. On the investment side, Qualcomm, Tencent, Intel Capital, Google Ventures, Nvidia, Salesforce, Samsung Ventures, Alibaba, Apple, Baidu, Citi, and In-Q-Tel fund much of the growth. When it comes to the future of AI, we should ask whether consolidation makes

sense for the greater good and whether competition will eventually be hindered (along with access), as we’ve seen in other fields such as telecom and cable.

Fragmentation

The AI ecosystem spans hundreds of companies. They are building the network infrastructure, the custom chipsets, the consumer applications, the back-end communications systems, the low-power radios in our smart home gadgets ... we could go on. Meanwhile, a large number of policy groups, advocacy organizations, and governments are developing guidelines, norms and standards, and policy frameworks hoping to guide the future development of AI. As a result, the ecosystem is fragmented in two ways: infrastructure standards and governance.

Expert Insight



AI reveals our real-world biases

Meredith Broussard

Associate Professor,
New York University

The biggest topic of conversation in AI right now is bias. We're moving away from the stage of marveling at AI technology, and people are realizing that AI models and automated systems reproduce all of the existing biases and inequalities of the real world.

If your company is using AI, ask: how might this system fail, and for whom? If you are in a heavily-regulated industry, AI is a compliance risk. Educate your teams about



The biggest topic of conversation in AI right now is bias.

algorithmic bias, audit your systems and your algorithms for bias, and require your vendors to do the same. If an AI system discriminates or makes decisions that are not in line with your corporate values, don't use it.

**

Meredith Broussard is an associate professor at the Arthur L. Carter Journalism Institute of New York University, an Affiliate with the New York University Alliance for Public Interest Technology, and the author of [Artificial Unintelligence: How Computers Misunderstand the World](#).

Society Trends

AI Still Has a Bias Problem

It's no secret AI has a serious and multi-faceted bias problem. Just one example: The datasets used for training often come from places like Reddit, Amazon reviews, and Wikipedia, a site inherently riddled with bias. The people building models tend to be homogeneous and aren't often aware of their own biases. As computer systems get better at making decisions, algorithms may sort each of us into groups that don't make any obvious sense to us—but could have massive repercussions.

Every single day, you are creating unimaginable amounts of data, both actively (such as when uploading and tagging photos on Facebook) and passively (driving to work, for example). That data is mined and used, often without your direct knowledge or understanding, by algorithms. It is used to create advertising, to help potential employers predict our behaviors, to determine our mortgage rates, and even to help law en-

forcement predict whether we're likely to commit a crime. Researchers at a number of universities—including the University of Maryland; Columbia University; Carnegie Mellon; MIT, Princeton University; University of California, Berkeley; International Computer Science Institute; among others—are studying the side effects of automatic decision-making. You, or someone you know, could wind up on the wrong side of the algorithm and discover you're ineligible for a loan, or a particular medication, or the ability to rent an apartment, for reasons that aren't transparent or easy to understand.

Increasingly, data is being harvested and sold to third parties without your knowledge. These biases can reinforce themselves over time. As AI applications become more ubiquitous, the negative effects of bias will have greater impact. The Apple card gave higher credit limits to men than women, in some cases by a factor of 20. Wearables such as Google's Fitbit are considerably less accurate for

darker skin types because of how melanin absorbs green light. This can lead to bias when insurance company algorithms take into account heart rates, blood pressure, and risk rates for conditions like irregular heartbeats or a potential heart attack.

Problematic Training Data

In 2018, researchers at MIT developed an AI called Norman that was trained to perform image captioning. They trained Norman using only content from a subreddit that's known for graphic violence. When Norman was ready, they unleashed him against a similar neural network that had been trained using standard data. Researchers fed both systems Rorschach inkblots and asked them to caption what they saw, and the results were striking: Where the standard system saw "a black and white photo of a baseball glove," Norman saw "a man murdered by machine gun in broad daylight." The point of the

experiment was to prove that AI isn't inherently biased, but that data input methods—and the people inputting that data—can significantly alter an AI's behavior. In 2019, new pre-trained systems built for natural language generation were released—but the conversations from which they learned were scraped from Reddit and Amazon reviews, both of whose author populations skew white and male, which means that their use of language isn't representative of everyone. This illustrates an ongoing challenge within the developer community. It is already difficult to get authentic data from real people to train systems, and with new privacy restrictions, developers are choosing to rely more on public—and problematic—datasets.

AI to Catch Cheaters

AI is being used to catch cheaters. ECRI Institute's CrossCheq uses machine learning and data analytics to look for hyperbole and misleading information

during the hiring process. Drexel University researchers built an app that uses biometrics to predict when dieters are likely to stray from their prescribed regimens. Researchers at the University of Copenhagen created a machine learning system to spot cheating on essays with, they say, a 90% accuracy rate.

Algorithms Targeting Vulnerable Populations

There is no question that machine learning systems trained correctly can help find missing children and detect abuse. The problem is that the systems use data from vulnerable populations to do their training. The Multiple Encounter Dataset contains two large datasets of photos: people who have not yet committed a crime, and an FBI dataset of deceased people. The dataset over indexes on people of color, which means that if law enforcement uses the data to train algorithms, it's going to lead to bias. Image recognition is a particularly vexing



Society Trends

challenge, because researchers need large datasets to perform their work. Often, images are used without consent. The Child Exploitation Image Analytics program—a dataset used for testing by facial recognition technology developers—has been running since 2016 with images of “children who range in age from infant through adolescent” and the majority of which “feature coercion, abuse, and sexual activity,” according to the program’s own developer documentation. These images are considered particularly challenging for the software because of the greater variability of position, context, and more.

AI Intentionally Hiding Data

Computers do exactly what they are told to do. Command a machine to win at a game, and it will do everything in its power to achieve that goal. Apparently that now includes cheating. Researchers at Stanford University and Google discovered that an AI, which was designed

to turn satellite images into usable maps, was withholding certain data. Researchers were using a neural network called CycleGAN, which learns how to map image transformations. For example, it took an old aerial photograph of a neighborhood, distinguished between streets, alleys, driveways, buildings, and lampposts, and then generated a map that could be used by GPS. Initially, they used an aerial photograph that hadn’t been seen by the network. The resulting image looked very close to the original—suspiciously close. But on deeper inspection, the researchers found that many details in both the original image and the generated image weren’t visible in the map made by the AI. It turns out that the system learned to hide information about the original image inside of the image it generated.

Undocumented AI Accidents

Only a few of the numerous AI-related accidents in 2018 and 2019 made

headlines. Most people know about the Uber self-driving car that hit and killed a pedestrian in Tempe, Arizona. But there were countless more incidents that didn’t result in death, and as a result, aren’t known to the public. At the moment, researchers are not obligated to report accidents or incidents involving our data, or AI processes, unless a law is broken. While big companies must inform consumers if their personal data—credit card numbers, home addresses, passwords—have been stolen, they are not required to publicly document instances in which algorithms have learned to discriminate against someone on the basis of race or gender, for example.

Digital Dividends

Artificial intelligence will inevitably lead to a shift in the global workforce, causing job losses across many industries. Researchers at Oxford University’s Institute for Humanity, researchers at the Future Today Institute, and former U.S. pres-

idential candidate Andrew Yang have all published works outlining different versions of a “digital dividend”—a way for companies to pay back to society a portion of the profits derived from AI.

Prioritizing Trust

We will soon reach a point when we will no longer be able to tell if a dataset has been tampered with, either intentionally or accidentally. AI systems rely on our trust. If we no longer trust their outcomes, decades of research and technological advancement will be for naught. Leaders in every sector—government, business, nonprofits and so on—must have confidence in the data and algorithms used. Building trust and accountability requires transparency. This is a complicated process, and corporations, government offices, law enforcement agencies, and other organizations understandably want to keep data private. The ethics of how data is collected in the first place may also influence the trustwor-

thiness and validity of scientific research, particularly in areas such as organ donations and medical research. In addition, employing ethicists to work directly with managers and developers and ensuring developers themselves are diverse—representing different races, ethnicities and genders—will reduce inherent bias in AI systems.



Bully Bots

Mid-future pessimistic scenario

We've seen affable chatbots inadvertently turn malicious and inflammatory when they're exposed to unfiltered online conversations, but now they're purposely programmed to antagonize. Black hat hackers cultivate bully bots—malicious chatbots that target and troll children and teens. Bully bots are designed to detect and exploit children's insecurities to inflict harm and suffering, and they can be bought anonymously at low cost. With so much of younger generations' conversations taking place via texts and DMs, there is ample data to feed the chatbots so that they can easily impersonate their target's peers. Ill-intentioned kids—and, disappointingly, adults—use these bots to outsource cyberbullying and deflect guilt and responsibility from themselves. Opportunist criminals use them to threaten children of prominent figures and wealthy individuals to gain access to security information and valuable objects. With inadequate laws and frustrating layers of anonymity, parents are at a loss over how to prevent harm and seek justice for their victimized children.



Application



STRATEGY

Artificial intelligence affects every business across multiple dimensions. AI is a cornerstone of most organizations, from workforce automation to digitization to staff allocation and beyond. It is imperative that executives and senior managers understand what AI is, what it is not, and what strategic value it adds to the business. Chief strategy officers should build a robust understanding of AI in order to develop longer-term plans and engage more closely with others in the C-suite, especially chief technology officers, chief information security officers, chief financial officers and others in the organization.



INNOVATION

AI is additive to the innovation and creative processes. Innovation teams can use deep learning for new product ideation, for understanding markets, and for anticipating what's on the horizon. Especially as no-code and low-code applications become more widely available, innovation teams will build powerful systems for decision management, general brainstorming, and powerhouse ways of generating new ideas.



R & D

For AI companies, talent retention is likely to get worse in 2021, as companies offer increasingly competitive packages to attract top minds. China has emerged as an R&D powerhouse, posing additional talent challenges. But for other companies, this is a good time to make use of commercially available and open source frameworks, algorithms, cloud-based systems, and applications to mine your own data. There is tremendous value in learning from your customers, partners, and vendors, and industry.



RISK

Risks abound in AI. New regulations could curb research, innovation, and product development. Geopolitical tensions and AI nationalism will start to direct foreign investment in new ways. Bias in facial recognition should concern everyone; so should the quickening advancements of general adversarial networks as they are used to manipulate and generate content. Everyone working in risk must develop a sharp focus on AI. Risk models should be developed to determine plausible near-future scenarios, so that leaders can adjust their strategies accordingly.

Key Questions

We recommend using this report to support your strategic foresight activity in the coming year. Every executive team should begin by asking these questions about technology, science, and policy:

1

How well do we understand AI?
What opportunities are we leaving on the table?

2

Are we adequately planning for the long term?
What assumptions must hold true for our current strategy to succeed?
How will we make needed changes?

3

What parts of our business model make us vulnerable to disruption brought by the evolution of AI?
What can we do now to mitigate future risk?



Selected Sources

A Constructive Prediction of the Generalization Error Across Scales, ICLR 2020. Jonathan S. Rosenfel, Amir Rosenfeld, Yonatan Belinkov N. Shavit.

Sharir, Or et al. "The Cost of Training NLP Models: A Concise Overview." *ArXiv abs/2004.08900* (2020)

Artificial Intelligence Market Size, Share & Trends Analysis Report By Solution (Hardware, Software, Services), By Technology (Deep Learning, Machine Learning), By End Use, By Region, And Segment Forecasts, 2020 - 2027. Grand View Research, July 2020.

The State of Financial Crime 2021: *Technology, regulation, and the future of financial crime compliance*. ComplyAdvantage, January 2021.

Destructive Cyber Operations and Machine Learning, CSET Issue Brief. Dakota Cary, Daniel Cebul.

Steed, Ryan and A. Çalışkan. "Image

Representations Learned With Un-supervised Pre-Training Contain Human-like Biases." *ArXiv abs/2010.15052* (2020)

Hafner, Danijar et al. "Dream to Control: Learning Behaviors by Latent Imagination." *ArXiv abs/1912.01603* (2020)

"Business R&D and Innovation Survey (BRDIS)," National Science Foundation,

<https://www.nsf.gov/statistics/srvyindustry/#sd&tabs-1>.

Intel AI Labs: <https://newsroom.intel.com/press-kits/intel-labs/>; NVIDIA AI Labs:

<https://www.nvidia.com/en-in/deep-learning-ai/ai-labs-her/>; Uber AI Labs:

<https://www.uber.com/newsroom/ailabs/>.

Department of Defense, Summary

of the 2018 National Defense Strategy (Department of Defense, 2018), 3, <https://www.hsdl.org/?view&did=807329>

President's Council of Advisors on Science and Technology, Recommendations for Strengthening American Leadership in Industries of the Future (Washington, DC: Office of Science and Technology Policy, June 2020)

Dario Gil, "How to Ensure the U.S.'s Quantum Future," *Scientific American*, August 20, 2020, <https://www.scientificamerican.com/article/how-to-ensure-the-uss-quantumfuture/>.

<https://www.reuters.com/article/us-usa-economy-artificialintelligence/as-companies-embrace-ai-its-a-job-seekers-market-idUSKCN1MP10D>

<https://www.nature.com/articles/s41586-019-1799-6>

<https://science.sciencemag.org/con->

<tent/371/6526/284>

<https://blog.einstein.ai/progen/>

<https://www.nature.com/articles/s41591-020-0916-2>

<https://paperswithcode.com/trends>

<http://horace.io/pytorch-vs-tensorflow/>

<https://thegradient.pub/p/cef6dd26-f952-4265-a2bc-f8bfb9e-b1efb/>

<https://medium.com/runacapital/open-source-growth-benchmarks-and-the-20-fastest-growing-oss-startups-d3556a669fe6>

Authors



LEAD AUTHOR

Amy Webb

Founder and CEO

Future Today Institute

awebb@futuretodayinstitute.com

Amy Webb pioneered FTI's data-driven foresight methodology that is used within hundreds of organizations globally. Her focus is to transform organizations as they prepare for complex futures. She advises CEOs of the world's most-admired companies, three-star admirals and generals, and the senior leadership of central banks and intergovernmental organizations. She leads FTI's technology research initiatives on AI, synthetic biology and genetic engineering, next-generation networks, and quantum technologies. Webb is a professor of strategic foresight at the New York University Stern School of Business, where she developed and teaches the MBA course on strategic foresight. She is a Visiting Fellow at Oxford

University's Saïd School of Business, was elected a life member to the Council on Foreign Relations and is a member of the Bretton Woods Committee. She is a member of the World Economic Forum where she serves on the Global Future Council on Media, Entertainment and Culture and the Stewardship Board of the Forum's Platform for Shaping the Future of Media, Entertainment and Culture. Webb is the bestselling author of several books about strategic foresight and emerging technologies. A lifelong science fiction fan, she collaborates closely with writers and producers on films, TV shows and commercials about science, technology and the future.



FTI CONTRIBUTOR

Marc Palatucci

Senior Foresight Associate

Future Today Institute

mpalatucci@futuretodayinstitute.com

Marc Palatucci is a Senior Foresight Associate at the Future Today Institute, with research specializations in new realities (AR/ VR/ MR/ DR), digital fashion, future of work and learning, retail and e-commerce. He co-leads the MBA course in strategic foresight at the New York University Stern School of Business. Palatucci serves on the World Health Organization's Learning Strategy Advisory Group and is a Senior Deputy to the World Economic Forum's Platform for Shaping the Future of Media, Entertainment and Culture. He holds an MBA in Emerging Technology from New York University's Stern School of Business and a

BA in Linguistics and Languages from NYU's Gallatin School of Individualized Study. Palatucci serves as editor-at-large for an arts, fashion and culture magazine and collaborates with writers, designers and producers on films, TV shows and commercials about science, technology and the future.

Authors



FTI CONTRIBUTORS

Leah Zaidi
 Senior Foresight Associate
 Future Today Institute
lzaidi@futuretodayinstitute.com


Leah Zaidi is an award-winning futurist with specializations in systems thinking, worldbuilding, science fiction prototyping and experiential futures. She is an Associate Editor of the World Futures Review. She holds an MDES in Strategic Foresight and Innovation from OCAD University and a BA from York University.



EDITORIAL

Jennifer Alsever
 Editorial Director
 Future Today Institute
jalsever@futuretodayinstitute.com

Jennifer Alsever is the Future Today Institute's Editorial Director. She has been a journalist for more than two decades covering tech, biotech, startups and business for such publications as Fortune Magazine, the Wall Street Journal, the New York Times, Wired and Fast Company. A popular young adult fiction writer, Alsever has won multiple YA awards for her Trinity Forest Series.



Cheryl Cooney
 Director of Operations
 Future Today Institute
ccooney@futuretodayinstitute.com

For over a decade, Cheryl Cooney has served as the Future Today Institute's director of operations, where she manages workflows, planning and logistics. Cooney is a published poet, with works appearing in American and Australian anthologies.

Tom Foster
 Editor

Sarah Johnson
 Editor



Authors



CREATIVE

Emily Caufield

Creative Director
Future Today Institute
ecaufield@futuretodayinstitute.com

Emily Caufield is an award-winning designer and illustrator. Serving as FTI's creative for more than a decade, Caufield applies design thinking to visually communicate complex trends, scenarios and foresight research. Caufield designed all aspects of this year's trend report. She is a graduate of Boston University's College of Fine Arts.



Julia Durgee

Artist and Futurist

Julia Durgee is a rare left and right-brained MBA with creative, strategic, analytical, and digital experience at world-class brands. She illustrated the portraits that appear in this year's trend report.



BUSINESS DEVELOPMENT

Mel Blake

Business Development
mblake@futuretodayinstitute.com

Mel Blake handles commercial development and client relations for FTI. He was formerly founder and managing director of Monitor Talent, a speaker agency he founded at the Monitor Group, a global consulting firm. He is a board member of The Yale Center for Customer Insights. He holds an MBA from the Yale School of Management and a BA in Economics from Wesleyan University.

Why FTI

We answer your most challenging questions using data, creative inquiry, and strategic foresight.

- What are plausible deep (20+ years), long-range (10+ years), and near-term (2+ years) futures?
- What scenarios describe our futures?
- What's happening outside my industry that I should know?
- What companies, startups, and partners make up our future value network?
- What new products, services, or businesses should we build?
- Which tech trends should we monitor? When should we act?
- How can we build an early warning system to see the next disruptive event?
- How do we reduce uncertainty about our futures?

We support executive leaders and their teams.

The Future Today Institute works closely with executive leadership and management teams to transform their strategic thinking on the future. Advisory services include signal mapping, trend identification, scenario development, risk modeling, visioning, and strategic planning.

About the Future Today Institute

Founded in 2006, the Future Today Institute researches, models, and prototypes future risk and opportunity. As the leading strategic foresight and futures management consultants to executive leadership teams worldwide, FTI's data-driven applied research reveals trends and calculates how they will disrupt business, government, and society.

Together with our clients and partners, FTI is helping leaders achieve their preferred futures. Our pioneering, data-driven forecasting methodology and tools empower leaders to make better decisions about the future, today.

Contact Us

The Future Today Institute
33 Irving Place
10th Floor
New York, NY 10003

hello@futuretodayinstitute.com

267-342-4300

www.futuretodayinstitute.com

Disclaimer

The views expressed herein are the authors' own and are not representative of the greater organizations in which they have been employed. The names of companies, services, and products mentioned in this report are not necessarily intended as endorsements by the Future Today Institute or this report's authors.

The Future Today Institute's 2021 Tech Trends Report relies on data, analysis, and modeling from a number of sources, which includes sources within public and private companies, securities filings, patents, academic research, government agencies, market research firms, conference presentations and papers, and news media stories. Additionally, this report draws from the Future Today Institute's previous EMT Trends Reports, FTI Trend Reports, and newsletters. FTI's reports are occasionally updated on the FTI website.

FTI advises hundreds of companies and organizations, some of which are referenced in this report. FTI does not own any equity position in any of the entities listed in this presentation.

Any trademarks or service marks used in this report are the marks of their respective owners, who do not endorse the statements in this report. All rights in marks are reserved by their respective owners. We disclaim any and all warranties, expressed or implied, with respect to this report.

Using and Sharing The Material In This Report



We invite you to use, share, and build upon the material in our 14th annual Future Today Institute Tech Trends Report. We are making it freely available to the public. This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

You are free to:

Share

Copy and redistribute the material in any medium or format, including in your organizations and classrooms.

Adapt

Remix, transform, and build upon the material for your own research, work, and teaching.

Under the following terms:

Attribution

You must give appropriate credit to the Future Today Institute, provide a link to this Creative Commons license, and indicate if any changes were made. You may do so in any reasonable manner, but not in any way that suggests that the Future Today Institute endorses you or your use.

NonCommercial

You may not, under any circumstance, use the material for commercial purposes.

ShareAlike

If you remix, transform, or build upon the material, you must distribute your contributions under the same license as you see here.

You are prohibited from:

Commercial Sharing

Don't copy and redistribute this material in any medium or format for commercial purposes, including any personal/ corporate marketing and client services.

Representing This Work As Your Own

Do not represent any part of this material as your own without giving credit to the Future Today Institute.

Additional Restrictions

You may not apply legal terms or technological measures that legally restrict others from doing anything this license permits.



This is volume 1 in the Future Today Institute's 2021 Tech Trends Report. Each volume covers a different set of topics.

To find additional volumes, visit www.futuretodayinstitute.com/trends



