16TH EDITION · 2023 TECH TRENDS REPORT · ARTIFICIAL INTELLIGENCE

Future Today Institute

Detailed analysis of AI, including key techniques and influential models and more than 100 AI trends related to the enterprise, consumer ecosystem and government.

TABLE OF CONTENTS

TABLE OF CONTENTS

THIS YEAR REQUIRES FOCUS

In August 2017, a rare explosive event known as GW17817 took place in space. Two stars collided, unleashing a blast energetic enough to form an incalculable number of new stellar bits that continue to travel through interstellar space. Over time, this stardust will combine into small objects, evolve into large rocks, fuse with even more material, and form into planets. One incredibly violent disruption will someday lead to the formation of a new corner of the universe. This is how our own sun and Earth, and all of human existence, came into being.

Lately it's as if we've been living through the aftermath of cataclysmic explosions: the release of generative artificial intelligence systems like ChatGPT and Midjourney, a fusion breakthrough that could someday generate zero-carbon energy, Russia's ongoing invasion into Ukraine, deep uncertainty about a global recession, and AlphaFold's protein-folding algorithms that predicted structures for nearly all cataloged proteins known to science, to name a few. These and other forces of change are colliding, going supernovae, and resulting in an unfathomable amount of new signals—bits of change that, over time, result in the trends that shape society.

Now more than ever, it's important to carefully track new trends as they emerge. But that isn't easy, given the rapid pace of change. For that reason, the theme of our 2023 Tech Trends report is Focus. It is crucial to focus when new signals are forming because some may be lasting and develop into impactful trends, while others might burn out and fade away. In an increasingly complex and fast-paced world, leaders who focus on the trends that matter and adapt to changing circumstances make better decisions and see improved outcomes. Trends enable them to anticipate near-term change, understand the factors influencing their industries, and develop a point of view on the future.

Our research is presented in 14 in-depth reports that reveal the current state of play, a list of influencers to watch, key trends, detailed examples, expert perspectives and recommendations designed to help executives and their teams develop their strategic positioning. Some of the trends are new advancements on mature technologies, while others represent frontier technologies and areas of science. When we look at them collectively, new centers of gravity come into focus, and we can glimpse the impacts they will have on every sector.

Trends on their own cannot predict the future. Rather, future-focused organizations use them to deeply reflect on the tension between long-term and short-term goals and to reduce uncertainty. By understanding the trends and changes shaping the landscape, executives can make informed decisions and capitalize on new opportunities in the year ahead.

We invite you to join us in observing how the stardust settles into new signals and trends. Share your feedback with us at **2023trends@futuretodayinstitute.com**.

Amy Webb
Chief Executive Officer
Future Today Institute

STATE OF PLAY

AI is a force multiplier on technological progress because it is an enabler of other technologies and powers the evolution of business, government, and society.

ChatGPT, the viral chatbot from OpenAI that uses artificial intelligence to produce content for everyday people, sparked imagination and alarm at the end of 2022. From our point of view, it is the most powerful AI system ever released to the general public. Its ease of use, combined with increasingly creative prompts, made it a global sensation. Created by the same company that developed GPT-3 and image generator DALL-E 2, the bot prompted school boards to hold emergency meetings, since they worried the new tool would supercharge cheating on homework. It also spooked investors, who wondered if this marked the end of traditional search, and therefore posed a credible threat to Google.

To be sure, generative AI and other assistive technologies are bringing powerful capabilities to non-technical users. But there is a lot more happening in the field of artificial intelligence.

In this section of our emerging tech trends report, we use artificial intelligence as an umbrella term to encompass the techniques, models, and frameworks that make up the field. Its aim is to create intelligent machines that can sense, reason, act, and adapt like humans do, or in ways that go beyond our capabilities. Today, cars can park themselves, while emerging platforms are capable of having seemingly natural conversations. As the technology evolves to beyond-human capabilities, it might invent new drugs, predict the real-time movement of wildfires, and autonomously design machine parts—or even entire factories.

AI is transforming business, but it's also changing programming. Wired magazine once explained that "in traditional programming, an engineer writes explicit, step-by-step instructions for the computer to follow. With machine learning, programmers don't encode computers with instructions. They train them."

STATE OF ARTIFICIAL INTELLIGENCE 2023

Likely Near-Term Developments

General

- General-purpose models will be commoditized in the near-future. This will lead to large language model (LLM) features being integrated into every app. The first players to succeed in the space will likely be creative apps, and all of them will have LLMs as a backbone to inspire and augment human creativity.
- Vertically integrated solutions will garner a higher transactional value. Some companies will win by providing "a refined/value-added LLM product" to the end consumer and meeting the customer in desired distribution channels, such as LLMs for health care, legal, finance, and architecture.
- In the next 12-18 months, watch for the proliferation of large-scale open-source models and tools.

Enterprise

- Top talent will start breaking away from the largest players—Google, OpenAI, Meta—to form their own startups, which will include conversational agents, artificial general intelligence programs, AI-first biotech companies, and the like.
- Consolidation will continue to be a driving theme in 2023. Here's just one example: Microsoft plans to increase its investment in and leverage OpenAI for Bing to pull market share away from Google search.
- Google's powerful ChatGPT rival, Bard, uses an AI model called LaMDA. Launched February 2023, it could catalyze a new race for conversational search.
- We anticipate increased enterprise adoption of AI. Leaders see AI as necessary for growth in the current macroeconomic environment, even as new developments make some job categories obsolete.

Automation

- AI coding assistants will grow in popularity. OpenAI's Codex, introduced in 2021, evolved from research to open commercialization by the middle of last year. GitHub CoPilot is now available as a subscription ($10 per month). As of January, Amazon's CodeWhisperer is available in preview. Internally, Google is using a machine language–powered code completion tool—it may at some point be made available to everyday users, possibly this year.
- We expect to see AI more deeply integrated into health care and life sciences later this year. Generative AI will yield new proteins, antibodies, and drugs, while biology and chemistry-specific models will result in faster discovery, further providing practical use cases and leading to increased investment.
- Could 2023 mark the beginning of the end of human radiologists? Lithuania-based Oxipit, an AI-first medical imaging startup, built an autonomous system that reports on chest X-rays that show no abnormalities. The technology is good enough that it received state-level certification to operate independently, without a radiologist in the loop.

Regulation and Geopolitics

- Will this be the year the government breaks up Big Tech? US lawmakers have begun work to decouple Google's ad business from the rest of the organization. In the European Union, powerful new regulations of tech platforms could go into effect. In China, regulators have eased up slightly on Big Tech crackdowns—but the Chinese Communist Party is intent on focusing new R&D toward servicing the country's long-term growth ambitions.

STATE OF ARTIFICIAL INTELLIGENCE 2023

When will AI disrupt your organization?

Advancements in AI will disrupt every industry within the next decade, and Future Today Institute classifies it as a general-purpose technology. Like the steam engine and the internet before it, AI has the potential to influence entire economies and to alter society through political, economic, and social structures.

AI is now used across most industries, solving business problems, detecting fraud, improving crop yields, managing supply chains, recommending products, and assisting designers and writers in creative work. AI can predict call volume in customer service centers and recommend staffing levels; it also predicts the emotional state 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 to track; however, because AI is now so deeply intertwined with every aspect of life, you will find AI mentioned in other trends throughout all volumes of Future Today Institute's 2023 report.

Several factors are driving the momentum of AI trends and the probable timing of an industry's disruption:

Scaling: Enormous amounts of training data are still required for most AI models to learn. For example, recommender systems coupled with generative AI could lead to deep personalization for the hospitality and health care sectors—as long as data is made available. Historically, data is locked inside proprietary systems built by third parties, and regulation often hinders access to certain forms of data.

Investment: AI has passed through cycles of enthusiasm and disillusionment, leading to either too much or not enough capital being made available. Investors prioritize commercialization over basic R&D—though the latter yields bigger impact and often stronger returns. Whether investors are patient tends to influence progress and pace of commercialization.

Constraints on adoption: Even if a technology is maturing, constraints on its adoption can hinder its impact on an industry. For example, a business may refuse to adopt an automated system because it challenges existing orthodoxy or an existing successful strategy. This is especially true in health care, insurance, and financial services.

Regulations: Advances in technology typically outpace regulatory changes. This has benefitted AI, which until very recently was not targeted for regulation. Additionally, whether local regulations are conflicting or complementary influences adoption in the marketplace.

Media mentions: Increased awareness and enthusiasm can influence the momentum of a technology, even when there's been no real breakthrough. Until OpenAI's ChatGPT breakthrough in late 2022, leaders weren't talking about the impact generative AI might have on their business. Media bursts related to AI will drive momentum, especially if those stories are favorable, and more importantly, are easily understood by the public.

Public perception: How the public understands and responds to AI advancements will create or quell demand. This is especially true of generative AI and education/creativity/intellectual property/misinformation, and the role assistive technologies will play in shaping the future workforce.

R&D developments: The pace of new research breakthroughs can't be scheduled to coincide with a board meeting or earnings report. Factors like funding, quality and size of staff, and access to resources can improve the likelihood and speed of new discoveries. We closely monitor R&D developments but treat them as wild cards.

STATE OF ARTIFICIAL INTELLIGENCE 2023

Why AI trends matter to your organization

We believe AI is a force multiplier on technological progress because it is an enabler of other technologies and powers the evolution of business, government, and society. But new capabilities of large language models—and ChatGPT in particular—have deeply concerned some in the professional and creative services. For years, experts have warned about robots taking away blue collar jobs; now we're entering an era where AI will erode white collar jobs.

Since publishing our first Tech Trends report 16 years ago, we have included and expanded our coverage on artificial intelligence. What began as several pages of insights is now a dedicated, stand-alone report with more than 100 trends to monitor. AI is already transforming most economic sectors, but we anticipate deeper impacts this year across insurance, finance, entertainment, health care, biotechnology, and cloud computing.

KEY INSIGHTS

We have entered the “text-to-everything” era. Soon, we will use natural language to operate and interact with computers rather than graphic user interfaces (GUIs).

If the 2010s were known for perception AI—systems that sensed signals such as images and text—to understand the world around us, the 2020s will be known for generative AI. These systems not only sense and understand the world but can also generate new content, concepts, and ideas while communicating with us.

AI is emerging as an assistant for all knowledge workers. Within the next 18-24 months, we will see assistive technology developed for a variety of professions. Think: GitHub’s Copilot for financial analysts, commercial real estate developers, and lawyers.

In the next 18-24 months, generative AI will be incorporated into consumer applications. Already, Canva, the popular online graphic design tool, has integrated Stable Diffusion in its platform. Microsoft has incorporated OpenAI’s DALL-E 2 image-generating system into its Microsoft Designer and Image Creator applications, which is a big upgrade to clip art. Google’s Bard includes an API to encourage developers to start building new products. Notion, the freemium productivity and note-taking web application, uses LLMs for assistive writing.

It’s possible for agents to learn the right skills but the wrong objectives; an AI system can be asked to learn something that then could be used for harmful purposes. Commercial AI products could inadvertently incentivize bad behavior.

We expect to see increased activity in the specialized AI accelerators space this year. Watch the UK’s Graphcore, which built a new type of processor for machine intelligence to accelerate machine learning and AI applications, and Cerebras, which built one of the fastest AI accelerators based on the largest processor in the industry, as well as AWS Inferentia accelerators. Will Nvidia’s GPU/AI accelerator moat get disrupted this year?

Publicly available LLMs are often the foundation for AI startups, but some researchers and technologists are questioning their defensibility when it comes to capturing value. The moat is in data. Techniques and models will largely get commoditized, and served via the infrastructure layer, where real value will be realized.

Long-term sustainability depends on network effects to gather enough user data. User-generated data can be harnessed to differentiate systems by offering tuned models on top of foundational/commoditized LLMs, creating a flywheel effect. Longer term, we expect to see niche LLMs owned by a select few players, while general-purpose LLMs become commoditized.

At the infrastructure layer, we anticipate a variety of generative AI services built to better serve a variety of applications. In addition to OpenAI, specialized players are entering the market, including AI21, Cohere, Snorkel.AI, and Scale.AI.

IMPORTANT TERMS

Machine learning (ML)

Machine learning uses data to make predictions and recommendations on how to achieve stated goals. 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 data sets 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 AI's growth.

There are different types of machine learning, including supervised, unsupervised, and reinforcement.

Supervised learning

A model that attempts to transform one type of data into another type using labeled examples. 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. This is the most common form of ML used today. Understanding what product features would most likely drive new purchases is a business use case for supervised learning.

Unsupervised learning

Data is provided to a model without specific output parameters, and the model tries to learn the data set's structure without any designated labels. For example, if a researcher doesn't know what to do with a large data set, an unsupervised learning model could determine patterns, classify data, and make recommendations without a human supervisor. Researchers used unsupervised learning during the pandemic to find patterns in how COVID-19 spread throughout communities.

Reinforcement learning (RL)

A system performs a task by repeatedly running calculations as it attempts to accomplish a stated goal. It's a trial-and-error process, where rewards or penalties are earned in response to the system's performance toward achieving the stated goal. RL is used when there isn't enough training data, when the researcher is trying 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 high number of business use cases, ranging from real-time dynamic pricing models to high-frequency trading algorithms and the systems that operate self-driving cars.

Deep learning (DL)

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 on their own. Once a system learns what an object looks like—say, an apple—it then can recognize that object in all other images, even if it has only a partial view.

There are different types of deep learning architectures. These are the most common:

Convolutional neural network (CNN)

A CNN is multilayered, with a convolutional layer, a pooling layer, and a fully connected layer. Each one performs a different task with 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 to identify anomalies in medical imaging, faulty products on a production line, blight on crops, and other irregularities.

Recurrent neural networks (RNNs)

These multilayered neural networks 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 (sports, financial) using structured data.

Transformers

A transformer is a component whose purpose is to process sequential data, such as natural language or genome sequences. Transformers rely on "attention" (the mathematical description of how things relate to, complement, or modify each other) in translating sequences. A transformer neural network is the unique architecture that enables systems to learn from context and to generate new information. Transformers are

IMPORTANT TERMS

complementary to CNNs and RNNs, the two most common neural network architectures used in deep learning.

Generative adversarial networks (GANs)
As unsupervised deep learning systems, GANs are 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, such as identifying the correct person, resulting in optimizing overall performance. GANs are useful when researchers don't have enough data to train an algorithmic model, and are also used to create new, synthetic data.

Deepfakes, which have become prevalent 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 pre-set 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.

Additional Terms

AI safety
A field that studies and attempts to mitigate the catastrophic risks that future AI could pose to humanity.

Algorithm
A process describing how to solve a specific problem or how to complete a particular task.

Artificial general intelligence (AGI)
A designation for systems that match and then exceed the full range of human cognitive ability across all economically valuable tasks.

Automatic speech recognition
Algorithmic systems that give computers the ability to recognize and convert audio to human readable language.

Computer vision
Processes that give computers the ability to derive meaningful information from digital images (including still and video) and to mimic and manipulate such images.

Model
A program that has been trained on a data set. Models are generally used for analytical and decision-making tasks, such as making predictions.

Natural language processing
Processes that give computers the ability to understand, mimic, and manipulate human language.

Recommender systems
A class of machine learning algorithms that use data to predict, narrow down, and find what people are looking for among an exponentially growing number of options.

Transformer
A type of neural network mechanism that learns what text means when it appears in a particular context. Using "attention," 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, image generation, and genome sequencing.

ONES TO WATCH

Dr. Ali Madani, for leading the AI for protein engineering moonshot (ProGen) at Salesforce Research and developing large language models used to design artificial proteins that outperformed their naturally occurring peers.

Allison Gardner, senior scientific adviser at NICE, for her research in algorithmic bias, diversity, and inclusion.

Dr. Anima Anandkumar, director of ML research at Nvidia and Bren Professor at Caltech, for her research in tensor-algebraic methods, deep learning, and non-convex problems.

Ari Kalfayan, global head of business development for early-stage AI/ML startups at Amazon Web Services, for scouting cutting-edge startups.

Ashish Vaswani, co-founder and chief scientist of Adept.AI, and co-author of the transformer model architecture in 2017 at NIPS.

Ashley Llorens, vice president, distinguished scientist, and managing director, Microsoft Research Outreach, for leading inclusive, collaborative research projects.

Dr. Cassie Kozyrkov, chief decision scientist at Google, for innovating in the combined fields of data science and behavioral science.

Dr. Daniel Cohen-Or, professor at the Department of Computer Science and The Isaias Nizri Chair in Visual Computing at Tel Aviv University, along with **Dr. Amit H. Bermano**, assistant professor at Tel Aviv University; **Dr. Ron Mokady**, computer vision research at Google; **Rotem Tzaban**, software engineer at Google; and **Rinon Gal**, computer science PhD candidate at Tel Aviv University, for developing a framework for semantic editing of faces in videos, demonstrating significant improvements over the current state-of-the-art.

Dr. Darío Gil, senior vice president and director at IBM Research, for advancing fundamental AI research to the benefit of industry and society.

Dr. David Chalmers, professor of philosophy and neural science at New York University, for researching whether large language models are sentient.

Dr. Dong Yu, vice general manager and distinguished scientist at Tencent AI Lab, for his work in natural language processing and speech recognition.

Dr. Dou Shen, chief of Baidu AI Cloud, for leading the development of large Chinese language models and enterprise applications.

Dr. Eric Boyd, corporate vice president of Microsoft's AI Platform, for developing an enterprise path for AI systems.

Drew Harwell, reporter covering AI at The Washington Post, for investigating the people, products, and companies that make and make use of AI.

Dr. Gadi Singer, vice president at Intel Labs and director of Emergent AI Research, for driving innovation in machine intelligence and cognition.

Dr. Ilya Sutskever, co-founder and chief scientist at OpenAI, and Sam Altman, CEO of OpenAI, for pushing the frontiers of generative AI.

Dr. Irina Rish, professor of computer science and operations research at the University of Montreal and a core member of the Mila—Quebec AI Institute, for her research on sparse modeling and probabilistic inference, dialogue generation, biologically plausible reinforcement learning, and dynamical systems approaches to brain imaging analysis.

Dr. Jae Lew, director of computer vision, robotics, and AI, for developing innovative automation systems using machine and deep learning.

Dr. Joelle Pineau, director of Meta AI Research Labs, for developing new models and algorithms for planning and learning in complex, partially observable domains.

Emad Mostaque, founder and CEO of Stability. AI, for the design of open-source AI tools.

ONES TO WATCH

Dr. Jeff Dean, senior fellow and senior vice president of Google AI (Research and Health), for large-scale distributed neural network systems.

Karen Hao, senior reporter for China tech and society at The Wall Street Journal, for her deep analysis of the social and geopolitical impacts of AI.

Kay Firth-Butterfield, head of AI and ML, and member of the executive committee at the World Economic Forum, for leading a global effort on the governance of AI, to harmonize AI policies, and to explore future global risks.

Khari Johnson and **Will Knight**, senior writers for AI at Wired magazine, for their reporting and analysis of the positive and negative ways AI shapes human lives.

Dr. Ming-Yu Liu, distinguished research scientist and a director of research at Nvidia Research, for focusing on deep generative models and their applications.

Dr. Mutale Nkonde, AI policy adviser at Harvard University's Berkman Klein Center, for her work to eliminate the underrepresentation of black professionals in the American technology sector by 2030.

Dr. Prakhar Mehrotra, vice president for machine learning US Omni Tech at Walmart Global Tech, for leading enterprise adoption of AI.

Rachel Metz, a senior journalist (most recently at CNN), for covering the business side of AI.

Dr. Rachel Thomas, foundation director of University of San Francisco Center for Applied Data Ethics, for bringing people from around the world with diverse and nontraditional backgrounds into AI.

Dr. Ramayya Krishnan, dean of the Heinz College of Information Systems and Public Policy at Carnegie Mellon University, for his research on the responsible use of AI and in data-driven approaches to support workforce development.

Dr. Rana el Kaliouby, deputy CEO of Smart Eye and former co-founder of Affectiva, for bringing emotional intelligence to devices and digital experiences.

Dr. Samy Bengio, senior director of AI and ML research at Apple, for his work applying deep learning to speech, image, and other forms of content.

Dr. Swami Sivasubramanian, vice president of database, analytics, and machine learning at AWS, for advancing cloud capabilities and insights for businesses.

Dr. Timnit Gebru, founder and executive director at the Distributed AI Research Institute, for her research in AI ethics, algorithmic bias, and data mining.

Dr. Zhou Jingren, deputy director of Alibaba Damo Academy (Alibaba's leading-edge research arm), for leading AI initiatives related to smart cities, autonomous driving, mobile computing platforms, semiconductor R&D, and other areas.

MODELS, TECHNIQUES, & RESEARCH

TRENDS

AI MODELS

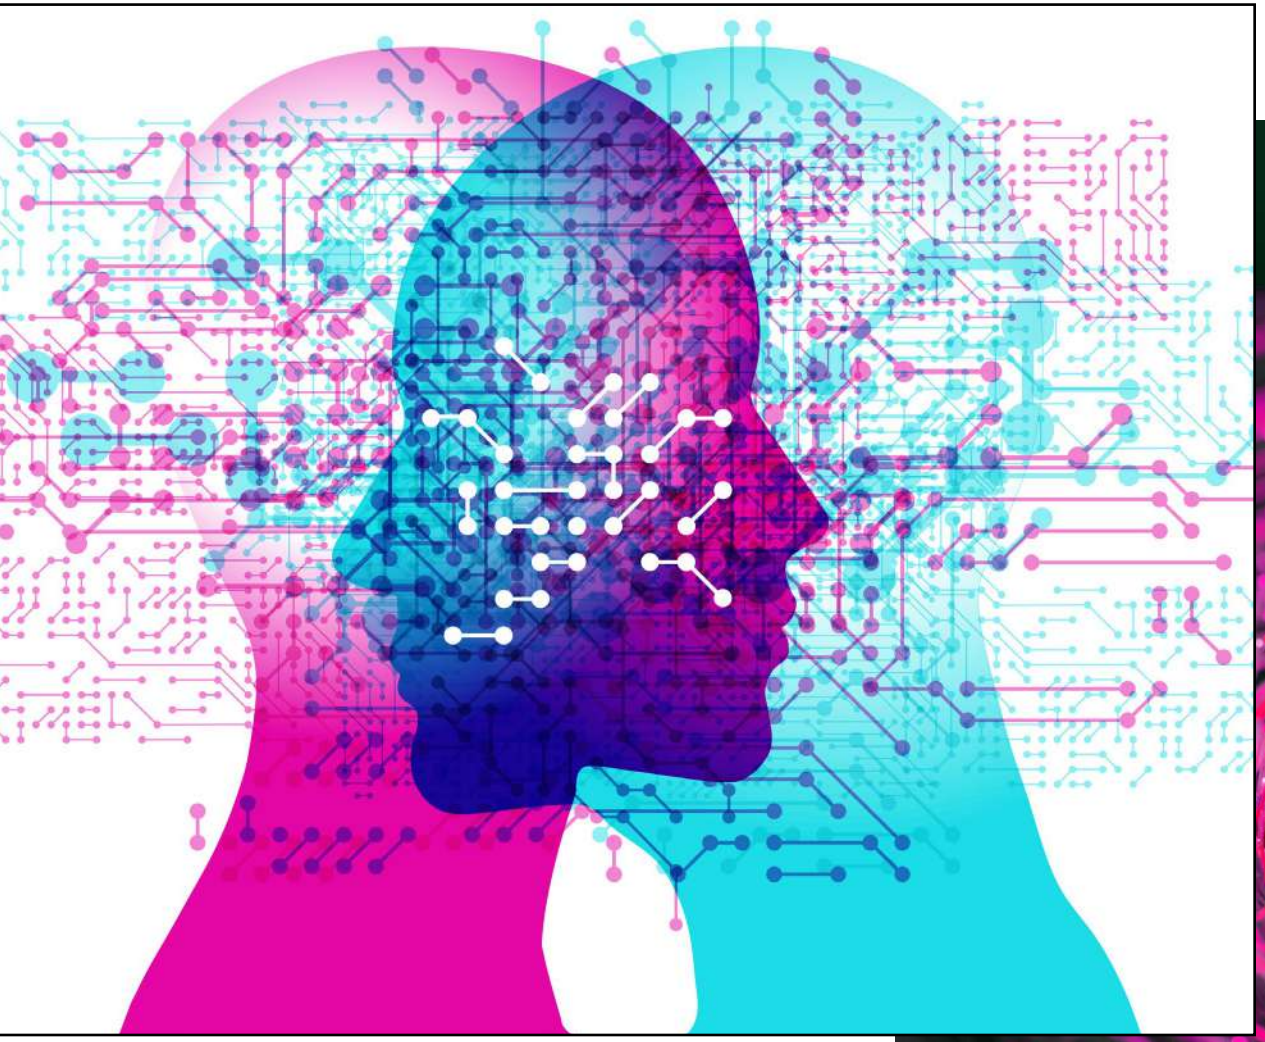

An AI model is a program or algorithm that uses trained data to spot patterns and automatically make predictions or decisions. Usually, the more data an AI model can access, the better it will be in delivering accurate answers or forecasts. A few models that might be included in an AI library include: linear regression, logistic regression, and decision trees. While all machine learning models are AI models, not all AI models are necessarily machine learning models.

Research in new image and video models have shifted from traditional convolutional neural networks to transformer-based discriminative models. Early this year, diffusion models took the world by storm. These models are extremely proficient at generating images and text, and we expect to see generative systems that can create realistic video this year.

Starting this year, recommender systems, coupled with generative AI, will lead to the further personalization of products and services. AI systems of the future will not only try to find content that consumers like, they will generate personalized content for that person's specific interests.

Some companies are releasing models to the public to continue training and tweaking them—possibly prematurely. In August, Meta released its BlenderBot3 for anyone to use; shortly after, the system began spinning out misinformation. Originally Google had planned to keep its LaMDA (Language Model for Dialogue Applications) system in-house after publishing a seminal paper on it in May 2021. But given increasing competition from OpenAI and Meta, the company changed course, announcing a new "AI Test Kitchen" initiative, where registered users can test various experiments, like the "List It" demo, which allows users to share a goal or topic and LaMDA will break it down into a list of helpful subtasks.

The large-scale release of AI systems to hundreds of millions of users may result in problematic knock-on effects. On the other hand, it's also likely to surface ethics and safety issues that developers may have overlooked.

AI MODELS

Modalities—Text

PaLM

The world's largest dense LLM developed by Google, at 540 billion parameters. PaLM shows breakthrough capabilities on numerous, very difficult tasks. PaLM is the largest single LLM trained on TPUs (Tensor Processing Units) at scale, with 6144 TPU chips.

BLOOM

A collection of open-source models, built through a community effort led by HuggingFace, with the largest model at 176 billion parameters. It can output coherent text in 46 languages and 13 programming languages that is hardly distinguishable from text written by humans.

Chinchilla

A 70 billion parameter model developed by DeepMind trained on 1.4 trillion tokens proved that current LLMs are largely undertrained. This model outperformed larger models with up to 530 billion parameters. The findings illustrate that for compute-optimal training, the model size and the number of training tokens should be scaled equally. The findings underscored the need for large, high-quality training data sets.

RETRO

Retrieval Enhanced Transformers, developed in February 2022 by DeepMind. Traditionally, the knowledge base of transformer models consists of only the data it was trained on. RETRO addresses this problem by obtaining a new knowledge base of "facts" through retrieving information from a database. RETRO helps LLMs stay current, without the need for retraining models. By decoupling the knowledge base from the model, researchers show that RETRO models with ~25x fewer parameters can achieve similar results to that of larger models.

GPT-3.5

The new and improved version of OpenAI's original 175 billion parameter GPT-3 model. This model was trained on more recent data with the new reinforcement learning with human feedback, a technique that used human and machine written data to improve the original GPT-3 model. This model is consumer focused, and is the backbone of ChatGPT. These sets of techniques are also commonly referred to as InstructGPT models.

NLLB

This 55 billion parameter open-source model developed by Meta AI is capable of delivering evaluated, high-quality translations directly between 200 languages—including low-resource languages like Asturian, Luganda, Urdu, and more.

Modalities—Images

DALL-E 2

Developed by OpenAI last year, DALL-E 2 is an upgrade of DALL-E, a model trained to manipulate visual concepts through language. It begins with a prompt that's written in natural language and then generates a set of images showing its interpretation of the intended meaning. With text prompts, it can now edit images, extend the original canvas, create variations of an input image, and has 4x higher resolution than the original DALL-E model.

Stable Diffusion

Developed by Stability.AI, Stable Diffusion is an open-source text-to-image, latent diffusion model that allows people to create stunning art within seconds. It is a breakthrough in speed and quality, allowing the model to run on consumer GPUs.

Parti

Developed by Google and also known as Pathways Autoregressive Text-to-Image, this autoregressive text-to-image generation model achieves high-fidelity photorealistic image generation and supports content-rich synthesis involving complex compositions and world knowledge.

Imagen

Developed by Google, this text-to-image diffusion model has a high degree of photorealism and a deep level of language understanding. Imagen builds on the power of large transformer language models in understanding text and hinges on the strength of diffusion models in high-fidelity image generation. The computation is done in the pixel space, unlike Stable Diffusion.

e-Diffi

Developed by Nvidia, this text-to-image diffusion model shows stunning results with instant style transfer and paintings with text-based prompts. Unlike Stable Diffusion, which does its computation in the latent space, e-Diffi shows promise with diffusion in the pixel space by using the power of transformer models to understand the relationship between text and images.

AI MODELS

The authors used Lensa AI to generate this portrait of Amy Webb.

Muse

Masked generative transformers is a text-to-image model developed by Google that shows promising results for image generation and editing use cases. Muse is more efficient because it employs discrete tokens, compared to pixel-space diffusion models, such as Imagen and DALL-E 2. When compared to autoregressive models, such as Parti, Muse is more efficient because it uses parallel decoding.

Modalities—Video

Make-a-Video

A state-of-the-art AI system that generates videos from text, developed by Meta AI. The network learns how the world moves by using images with descriptions, along with unlabeled videos. The result is astonishing; short video clips can be generated with just a few text prompts.

MSFT X-CLIP

Developed by Microsoft, this extension to CLIP for general video understanding will be used for tasks like video classification and video-text retrieval.

We anticipate that OpenAI will launch a video system this year.

Modalities—Voice

VALL-E

A voice synthesis transformer-based model that converts text to speech, VALL-E generates discrete audio codec codes, corresponding to the target content and the speaker's voice, enabling various text-to-speech applications, speech editing, and content creation.

Modalities—3D and Animation

Get3D

A model that generates high-quality 3D textured shapes learned from images, GET3D is able to generate high-quality 3D textured meshes, ranging from cars, chairs, animals, motorcycles, and human characters to buildings, achieving significant

AI MODELS

improvements over previous methods.

Magic3D
This text-to-3D content creation tool creates 3D mesh models using a coarse-to-fine strategy that leverages both low- and high-resolution diffusion priors for learning the 3D representation of the target content.

Modalities—Bioengineering

OpenFold
A fast, memory-efficient, and trainable implementation of AlphaFold2 and OpenProteinSet, the largest public database of protein multiple sequence alignments.

MegaMolBART
A deep learning model for small molecule drug discovery and cheminformatics based on SMILES (Simplified Molecular Input Line Entry System). MegaMolBART uses Nvidia's NeMo-Megatron framework, which is designed for the development of large transformer models.

MedPalm
A large language model for the medical community that advances SOTA in seven medical question-answering tasks, including achieving 67% on MedQA USMLE improving prior work by more than 17%.

EquiFold
An end-to-end differentiable, SE(3)-equivariant, all-atom protein structure prediction model.

ESM-2
A deep contextual language model trained on 86 billion amino acids across 250 million protein sequences spanning evolutionary diversity.

GenSLMs
Genome-scale language models that can learn the evolutionary landscape of SARS-CoV-2 genomes (among others).

ProtVAE
A deep generative model that surfaces an accurate, generative, fast, and transferable model of the sequence-function relationship for data-driven protein engineering.

Modalities—Climate Science

FourCastNet
Developed by Nvidia, the vision transformer-based global data-driven weather forecasting model provides accurate short- to medium-range global predictions at 0.25 degree resolution. FourCastNet accurately forecasts high-resolution, fast-timescale variables such as the surface wind speed, precipitation, and atmospheric water vapor.

Multi-Modalities

Gato
An agent from DeepMind that works as a multimodal, multitask, multi-embodiment generalist policy. The same network with the same weights can play Atari, caption images, chat, stack blocks with a real robot arm, and much more, deciding based on its context whether to output text, joint torques, button presses, or other tokens.

CLIP
The Contrastive Language-Image Pre-training neural network from OpenAI is an open-source, multimodal, zero-shot model. Given an image and text descriptions, the model can predict the most relevant text description for that image, without optimizing for a particular task.

Flava
A foundational language and vision alignment model from Meta that explicitly and simultaneously targets language and vision tasks.

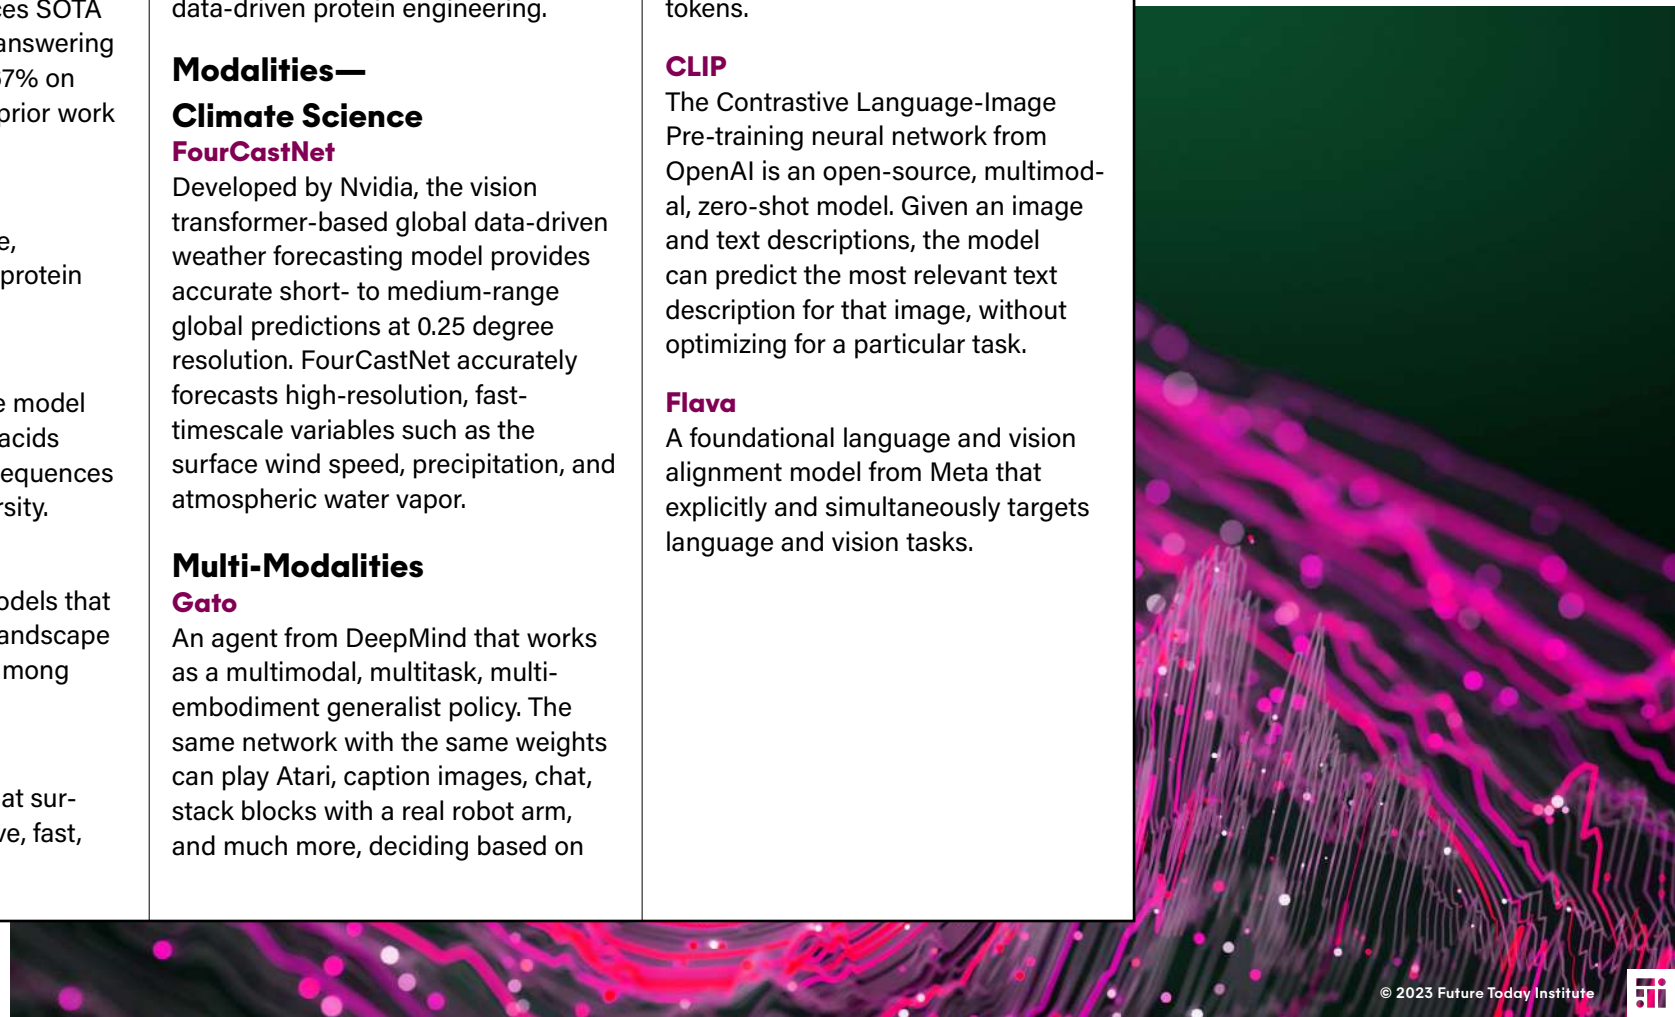

TECHNIQUES

PROMPT LEARNING/ENGINEERING

Prompt learning/engineering encompasses a collection of techniques that focus on customizing a foundational LLM for a particular use case, using a supervised learning technique. These methods are typically more efficient than fine-tuning as they require less computation and no retraining of the foundational model. The objective is to feed better inputs into the foundational model to obtain better responses. The process involves adding a snippet of text or vectors to input data during the supervised learning phase, to better model and understand the outcomes of the foundational model. Once learned, these snippets are added to all future input prompts.

REINFORCEMENT LEARNING WITH HUMAN FEEDBACK (RLHF)

RLHF is a technique used to align LLMs with human intentions and is based on training a reward model to mimic human feedback and intentions. This training uses prompt-generation pairs from a predefined data set, which takes a prompt and its corresponding completion to output a single "reward," or a score of how good the completion was. This reward model is then fed into a pipeline with a pretrained or fine-tuned LLM (policy model) to constantly improve the outputs of the LLM, based on human intentions.

IMITATION LEARNING

Neural networks are trained to perform tasks by watching humans do them. Imitation learning can be used to train AI to control robot arms, drive cars, or navigate webpages. Previously, this had to be done by a human annotating each action. Video Pre-Training (VPT) is an approach that empowers a neural network to automatically label videos. Given the billions of hours of video now available for free online, researchers have access to a vast resource to train AI systems. In the future, might this actually be a way to clone you—at least digitally? Maybe. Your digital clone trained on your data could read through websites, shop for holiday gifts, or book hotel rooms—just as you would. Taking that a step further: a robotic clone might someday do your gardening or move boxes around after watching videos to learn how you move.

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, AI specialists, and engineers. AutoML operates differently by matching raw data and models together to reveal the most relevant information. Google, Amazon, and Microsoft include a host of AutoML products within their cloud service offerings.

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

FEDERATED LEARNING

Federated learning is a technique that distributes machine learning to the edge. Introduced by Google researchers in 2016, this framework makes it possible for algorithms to use data on devices—such as mobile phones and smartwatches—without compromising user privacy. Research in this space has dramatically increased.

GENERAL REINFORCEMENT LEARNING ALGORITHMS

Researchers are developing single algorithms that can learn multiple tasks. DeepMind, the AI subsidiary of Alphabet and the team behind AlphaGo, which learned how to play Go with the skill level of a human grandmaster, continues to push its research forward. MuZero mastered multiple games without being told their rules, a "significant step forward in the pursuit of general-purpose algorithms," according to DeepMind. In a seminal paper, "Reward Is Enough," published at the end of 2022, DeepMind researchers hypothesized that artificial general intelligence could be achieved through reinforcement learning alone.

GRAPH NEURAL NETWORKS

Because we perceive scents using millions of sensory neurons in our brains, and because

TECHNIQUES

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.

GNNs constitute a particular type of deep neural network that operates on graphs as inputs. They 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, MIT and Harvard's biomedical and genomic research center, used them to discover antibiotic compounds that don't have toxic side effects.

ACTIVITY VISION

A fundamental subtask in video computer vision is activity recognition: identifying the activities that occur in videos. AI systems have been challenged to classify activities that range from simple actions, like walking, waving, or standing, to ones that are more complex and contain multiple steps, like preparing a salad (which requires an AI system to recognize and chain together discrete actions like cutting tomatoes, washing the greens, applying dressing, and the like).

As of January, one model stands out: MTV, a collaboration between Google Research, Michigan State University, and Brown University, which achieved a 89.6% Top-1 accuracy on the 600 series, 89.1% accuracy on the 400 series, and 82.20% accuracy on the 700 series.

ActivityNet is a video data set for human activity understanding that contains 700 hours of videos of people doing 200 different activities (long jump, dog walking, vacuuming, etc.). For an AI system to successfully complete the ActivityNet Temporal Action Localization Task (TALT) task, it has to execute two separate steps: (1) localization (identify the precise interval during which the activity occurs); and (2) recognition (assign the correct category label). Temporal action localization is one of the most complex and difficult tasks in computer vision. Performance on TALT is measured in terms of mean average precision, with a higher score indicating greater accuracy.

VISUAL COMMONSENSE REASONING (VCR)

The Visual Commonsense Reasoning challenge is a relatively new benchmark for visual understanding. VCR asks AI systems to answer challenging questions about scenarios presented from images, and also to provide the reasoning behind their answers (unlike the VQA challenge, which only requires an answer). The data set contains 290,000 pairs of multiple-choice questions, answers, and rationales from 110,000 image scenarios taken from movies. Although progress has been made since the challenge was launched, improvements have become increasingly marginal, suggesting that new techniques may need to be invented to significantly improve performance.

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 US Rep. Alexandria Ocasio-Cortez (D-N.Y.)—wearing a low-cut top or bikini 53% of the time.

PREDICTIVE MODELS USING INCOMPLETE DATA

Computer vision systems are getting smarter, which enables neural networks to predict geometry from a single color image. In 2019, the DeepMind team developed a 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 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

TECHNIQUES

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 humans by taking cues from body language. Retail, manufacturing, and education settings could be especially relevant.

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 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.

REAL-TIME MACHINE LEARNING (RTML)

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 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.

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.

MODEL-FREE APPROACHES TO RL

Dreamer is a reinforcement learning (RL) agent that uses a world model to learn long-sighted predictions, employing backpropagation (an algorithmic way of working backward to test for errors) 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.

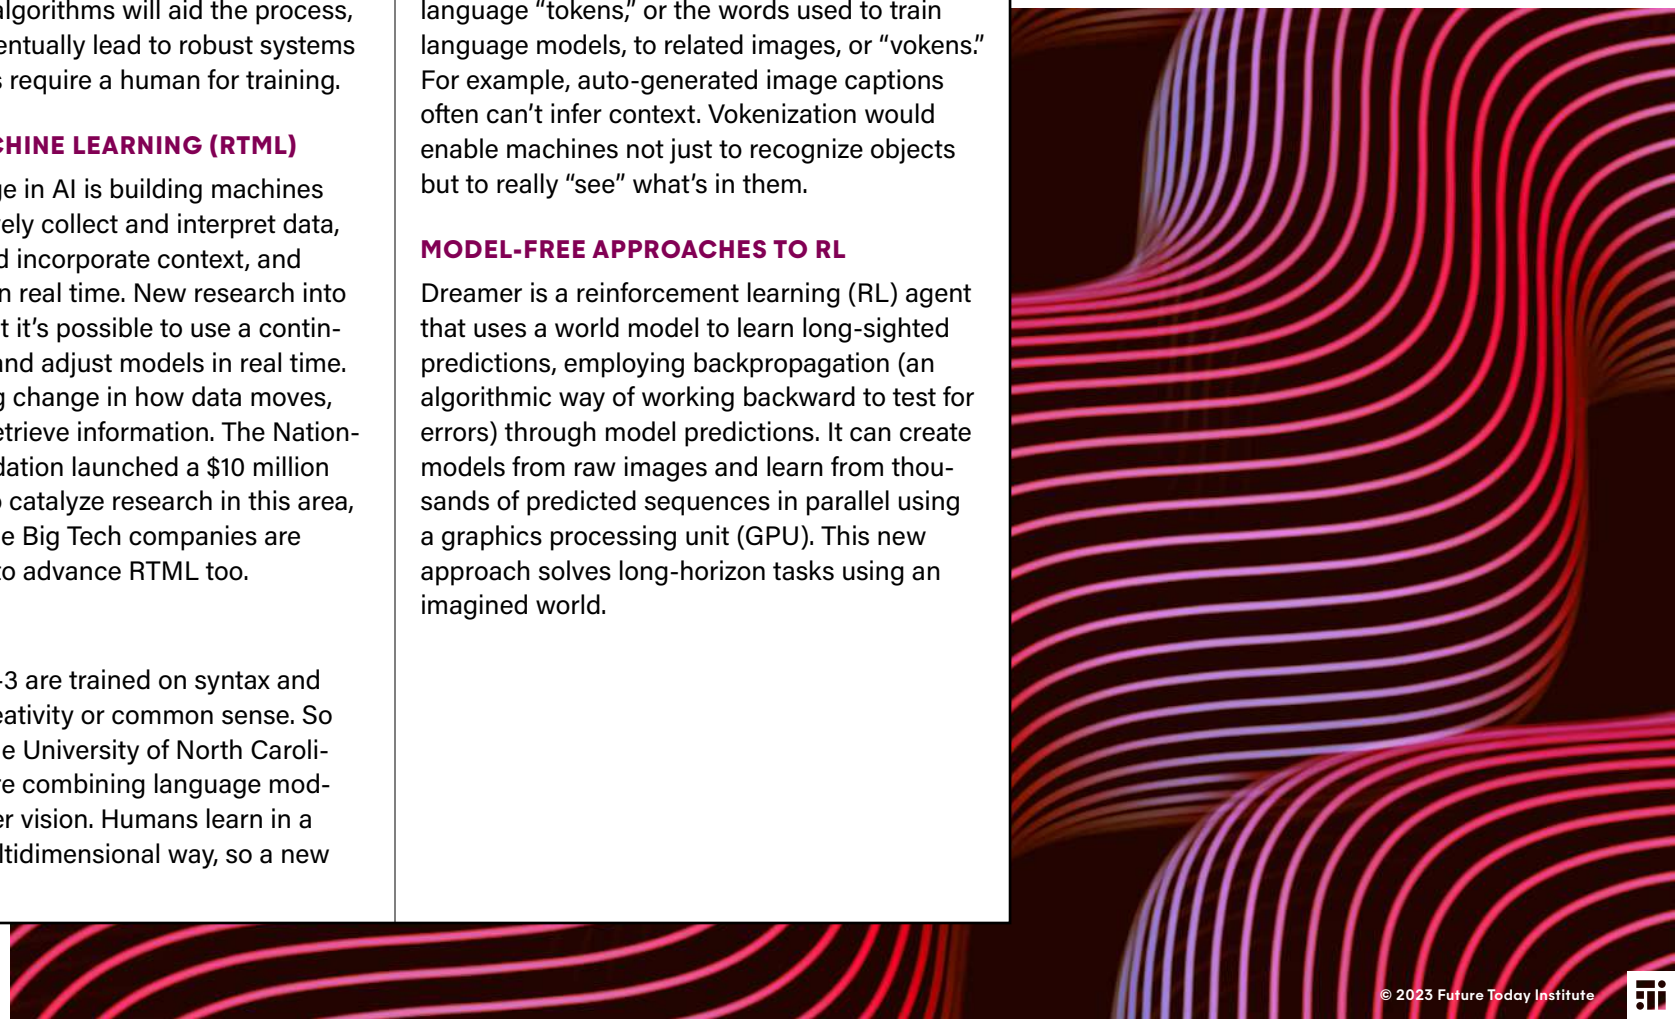

RESEARCH

BUILDING SUPERSIZED AI MODELS

Last year, we saw the proliferation of large AI models, and this year supersized models are on the horizon. For context, GPT-3—widely hailed as a powerhouse—has 175 billion parameters. Huawei debuted a 200 billion parameter language model called Pangu, while Baidu and the Peng Cheng Lab released PCL-BAIDU Wenxin, with 280 billion parameters. PCL-BAIDU is already deployed to Baidu's news feeds, search engine, and digital assistant. Gopher, which was released by DeepMind in December 2021, has 280 billion parameters. And Microsoft's Megatron-Turing NLG, built in collaboration with Nvidia, has 530 billion parameters. Google's Switch-Transformer and GLaM models have a staggering 1 trillion and 1.2 trillion parameters, respectively. But even bigger is Wu Dao 2.0 from the Beijing Academy of AI, which reportedly has 1.75 trillion parameters.

CREATING SPECIALIZED LANGUAGE MODELS

Increasingly, researchers are introducing specialized language models for different industries, and LLMs for proteins are already in use. Researchers at Salesforce created a suite of protein language models, named ProGen2, that are scaled up to 6.4 billion parameters and trained on different sequence data sets drawn from over a billion proteins from genomic, metagenomic, and immune repertoire databases.

Meta also introduced its own model to predict protein structures, called ESMFold, which demonstrated similar predictions to DeepMind's AlphaFold2. Outside of bioengineering, Google's PaLM-SayCan is an LLM that can use abstract steps to complete a goal, such as retrieving a sponge. This is a remarkable achievement, because sponges are soft and flexible, properties that introduce new variables to the task. During testing, SayCan was drilled on more than 100 instructions, and it was successful in planning and execution 84% of the time.

DEVELOPING PRACTICAL TOOLS FOR LANGUAGE MODELS

Language models can learn to use the tools we rely on every day, like calculators and search, simply by accessing text interfaces. They only need a few human demonstrations in order to learn. OpenAI's WebGPT is a tool that GPT3 uses to search the web, provides answers with references, and is likely the stepping stone to true conversational search. The startup Adept is working on commercialization: It hopes to customize large models to do things like interact with websites, APIs, and apps to drive workflow productivity.

TRAINING LANGUAGE MODELS TO ACHIEVE SCALE

Current opinion on language model (LM) scaling laws might be wrong; they are enormous, yet the amount of data they're trained on isn't commensurate with their overall size. The implication: Existing models aren't trained enough, and that's a problem. Researchers are eying a different approach going forward.

OpenAI has proven to many that foundational models can be improved efficiently by using RLHF, which might explain why the company released its generative AI systems publicly. Research now clearly indicates—across modalities—that a trifecta of (1) large-scale data, (2) accelerated computing, and (3) powerful algorithms can lead to astonishing results. Data scaling leads to better predictions from multibillion parameter models.

We are just scratching the surface on all three fronts:

Data
Today, the internet has provided us with an unprecedented amount of data that we never had access to before, thanks to the number of internet users worldwide, increased usage of the internet, and the ability of generative AI to create high-fidelity synthetic data.

Compute
Although Moore's law may have reached the limit in terms of miniaturization of electronics, we are seeing continued acceleration of the compute stack through software and algorithms. Software will continue to accelerate performance and improve the energy efficiency of chips exponentially.

Algorithms
Distributed and parallel computing algorithms and frameworks are rapidly evolving. Broadly, transformers have proven to scale efficiently, and can handle mountains of data. Innovations are accelerating on transformer and diffusion architecture today.

RESEARCH

ADOPTING UNIFIED LEARNING PROCESSES

Deep neural nets are good at identifying objects in photos and videos and processing natural language, but until recently models had to be trained separately. Might there be one model to rule them all? That was Google's proposal back in 2017. Since then, DeepMind's Gato has emerged as a 1.2 billion parameter transformer capable of performing hundreds of tasks in robotics, simulated environments, vision, and language. Researchers at Meta have now developed Data-2vec, a system that deploys a single algorithm to train a neural network to recognize images, text, or speech. It unifies the learning process through self-supervised learning, which allows the neural net to recognize patterns in data sets on its own, without being fed labeled examples.

TRANSITIONING TO TEXTLESS NLP

Most LLMs have been trained on publicly available data sets such as Reddit and Wikipedia. Both are rife with biases. Researchers are developing generative spoken language modeling, which extracts speech from raw audio without labels or text. The hope is that AI could become more inclusive if it uses podcasts, local radio, and other sources of spoken language.

AMPING UP DIFFUSION MODELS

Diffusion models are beating GANs at benchmarks for text-to-image generation. These generative models learn to model data distribution from whatever the input is, and with training they can generate new data sets similar to the ones they trained on. Stable Diffusion is based on diffusion models.

MANAGING THE COMMUNITY-DRIVEN OPEN SOURCING OF 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 20% of all academic papers on AI publish their full code, and some big names—DeepMind and OpenAI—notoriously leave theirs out, citing proprietary concerns.

That's why what's currently happening within the generative AI community is so interesting. Recently, as big models are released, others are quickly cloning and improving them. For example, GPT3 was released in June 2020, and by August 2022 a number of models were in use: Pan-Gu from Huawei, Ernie 3.0 Titan from Baidu, and later OPT from Meta. DALL-E was released in January 2021, and Meta's Make-a-Scene launched a year later, followed by Google's Imagen in May 2022 and Stable Diffusion in August 2022. Even AlphaFold 1, released in August 2018, had fast followers. Rosetta Fold launched in 2021, followed by OpenFold from Columbia University and Meta's ESM 2 in 2022. While these models aren't necessarily replicas, they either build on the ideas or the code that existed.

AI FOR CODE RESEARCH

Google, DeepMind, Salesforce, and OpenAI are actively building systems that help the computer science community develop code. Codex from OpenAI powers GitHub's Copilot—think of it as an autocomplete for code. It can string together multiple lines of code from natural language instructions. While at first the results were questionable, there has been enough interest that corporate labs are now competing to build commercial products. BigCode is an open-source collective initiative driven by ServiceNow and HuggingFace to develop an LLM for code. Salesforce's CodeGen is building a powerful alternative to Codex. Meanwhile, DeepMind's AlphaCode is already generating whole programs.

FRAMEWORK CONSOLIDATION

Google's TensorFlow and Meta's PyTorch are two popular frameworks used by researchers, and the relative popularity of different frameworks typically mirrors trends in the commercial application landscape. In the past five years, Meta seems to have gained ground. Of the 2021 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. JAX, used for the development of LLMs, gained popularity in the middle of the year with TPUs, but early this year activity had cooled.

LOWERING THE COST OF TRAINING MODELS

It's expensive to train a model, and several variables influence those costs, all of which have increased in the past few years. For example, it costs an average of $1 per 1,000 parameters

RESEARCH

today, which suggests that OpenAI likely spent more than $10 million to train GPT-3. Smaller research groups and companies cannot afford these expenses. Some in the AI community are instead allowing the big tech companies to pre-train and publish big models. The cost of computing will decline over time, and accelerations at the software layer are underway. Frameworks, and distributed training starting to take shape, are delivering meaningful gains every year.

SURPASSING BENCHMARKS

As AI systems improve, they are surpassing benchmarks. The original 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 tasks around understanding nine sentences or a pair of sentences built on existing data sets, which are selected to cover a diverse range of data set sizes, text genres, and degrees of difficulty.

It also has a diagnostic data set 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 surpassed humans by increasing their scores from 60 to 90.6.

The SuperGLUE benchmark is a new measurement of more difficult language understanding tasks, improved resources, and a new public leaderboard. When SuperGLUE was introduced, there was a nearly 20-point gap between the best-performing model and human performance on the leaderboard, but by 2022 AI models from Microsoft and Google had outperformed humans. Existing language benchmarks still fail to capture biases encoded in public data—future benchmarks could be designed to resolve this gap.

This year, watch the Holistic Evaluation of Language Models (HELM) from Stanford University, which holistically evaluates the performance of LLMs. The benchmark includes measuring models across a variety of capabilities (spanning tasks, domains, and languages), multiple metrics (accuracy, fairness, bias, toxicity, robustness, etc.), and the standardization of methodology and data sets across models. Also watch BigBench, a collaborative benchmark intended to probe LLMs and extrapolate their future capabilities across hundreds of tasks.

EXPLAINABLE AI (XAI)

There's a renewed call to make AI explainable—meaning, to reveal to researchers and others how systems make their forecasts and decisions. Many aspects of AI systems can be explained, such as showing their training data, or their lack of data in a particular region of the corpus. XAI also calls for transparency in the fairness of gathering data, the description of people involved in supervised learning, and the ranking of features. Emerging approaches to XAI methods include showing how outcomes were validated—meaning, how can we trust the recommendations, predictions, and classifications made?

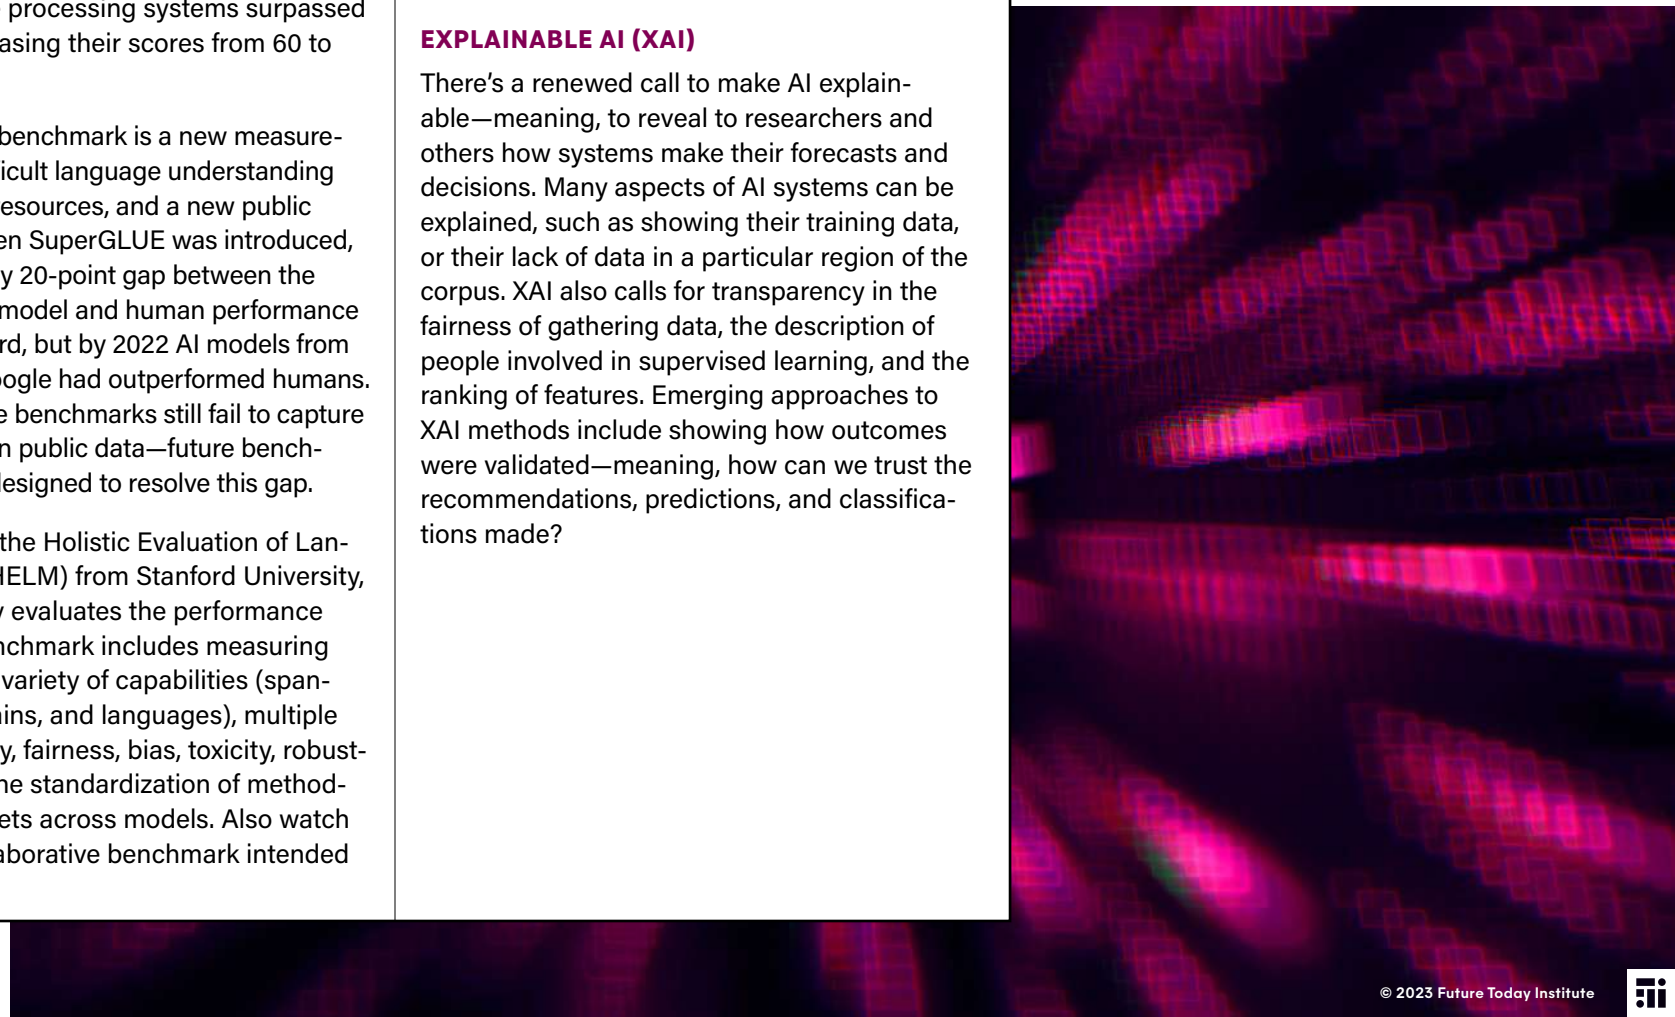

TALENT

TRENDS

TALENT

DEMAND FOR AI TALENT GROWING FAST

For many years, demand for AI talent has outpaced supply. In the US, 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 is just too much demand for AI skills and nowhere near enough trained workers. The hiring process for AI job postings is taking longer and becoming more expensive, impeding growth at some companies.

UPSKILLING FOR AI

Companies are looking to upskill their workforce in machine learning and the basics of AI. As a result, new training programs abound to augment the knowledge, skills, and competencies required of a modern organization. Levi Strauss launched a machine learning bootcamp to upskill its workforce, showing staff how to apply AI-thinking to everyday tasks. Founded by Harvard University and University of California, Los Angeles faculty, Univ.AI is an online program for training in machine learning and AI.

NUMBER OF AI PATENT FILINGS, 2010-2021

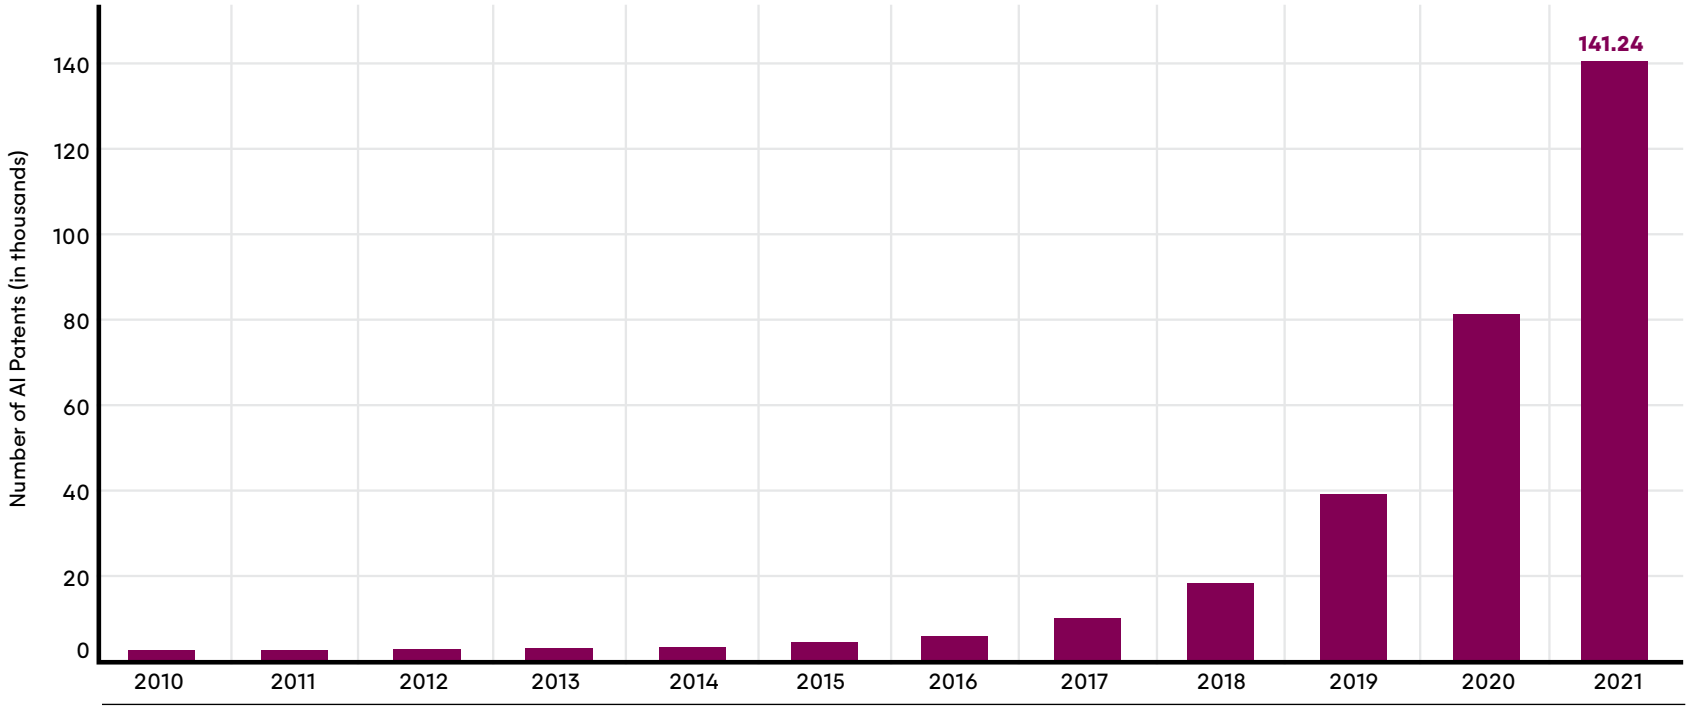

As AI moves more deeply into the mainstream, new patent filings have followed. The number of patents filed globally in 2021 was more than 30 times higher than the number filed in 2015.

Source: CSET and 1790 Analytics database containing all AI granted patents and published applications (collectively "patent documents") in all patent systems worldwide. https://github.com/georgetown-cset/1790-ai-patent-data

TALENT

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 corporate salaries and benefits, and they get to work in a similar tenured environment that's carefully cultivated to replicate their experience in academia. Between 2004 and 2018, Google, DeepMind, Amazon, and Microsoft hired 52 tenured and tenure-track professors from American universities.

Tech companies are also 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, and companies need the talent. In one infamous case, Uber poached an entire robotics lab in 2015 from Carnegie Mellon University—40 professors and researchers in total—to work on self-driving cars. Poaching academia today could rob the future of future AI experts: Without great scholars, who will train the next generation of innovators?

CORPORATE AI LABS

AI labs are located around the world, with concentrations in North America, Europe, and Asia. While the tech sector laid off tens of thousands of workers between Q4 2022 and Q1 2023, many corporate AI labs and positions remained intact. Meta, Google, IBM, and Microsoft operate more than 60 labs dedicated to AI R&D, and the majority are outside of the US because of access to talent. During the Trump administration, immigration restrictions and stringent visa requirements made it difficult to bring talent into America, and overseas labs allowed companies to avoid that barrier. Most of those labs do basic AI research rather than product development.

BUILDING A SILICON VALLEY OF THE MIDDLE EAST

The Middle East is vying to become the next Silicon Valley, with a particular focus on developing its own AI ecosystem. In 2022, nearly 3,000 startups launched with nearly $1 billion in investment according to MAGNiTT, a Dubai-based research firm. The United Arab Emirates government has developed ambitious plans to establish the country as a global AI hub. In 2017, it published a National Artificial Intelligence Strategy and launched the Mohamed bin Zayed University of Artificial Intelligence in Abu Dhabi. AI is taking center stage in the D33 Dubai Economic Agenda, which aims to generate new economic value from digital transformation and AI technologies.

Meanwhile, Saudi Arabia is also investing heavily in AI. In 2019, it established the Authority for Data and Artificial Intelligence, and created a National Center for Artificial Intelligence and a National Data Management Office. Technology is the focus of the Kingdom's Vision 2030 plan, a long-term strategy intended to stimulate economic growth. The Kingdom also launched its Saudi Venture Capital Co., with an initial $1.33 billion under management, and is making large investments in tech companies.

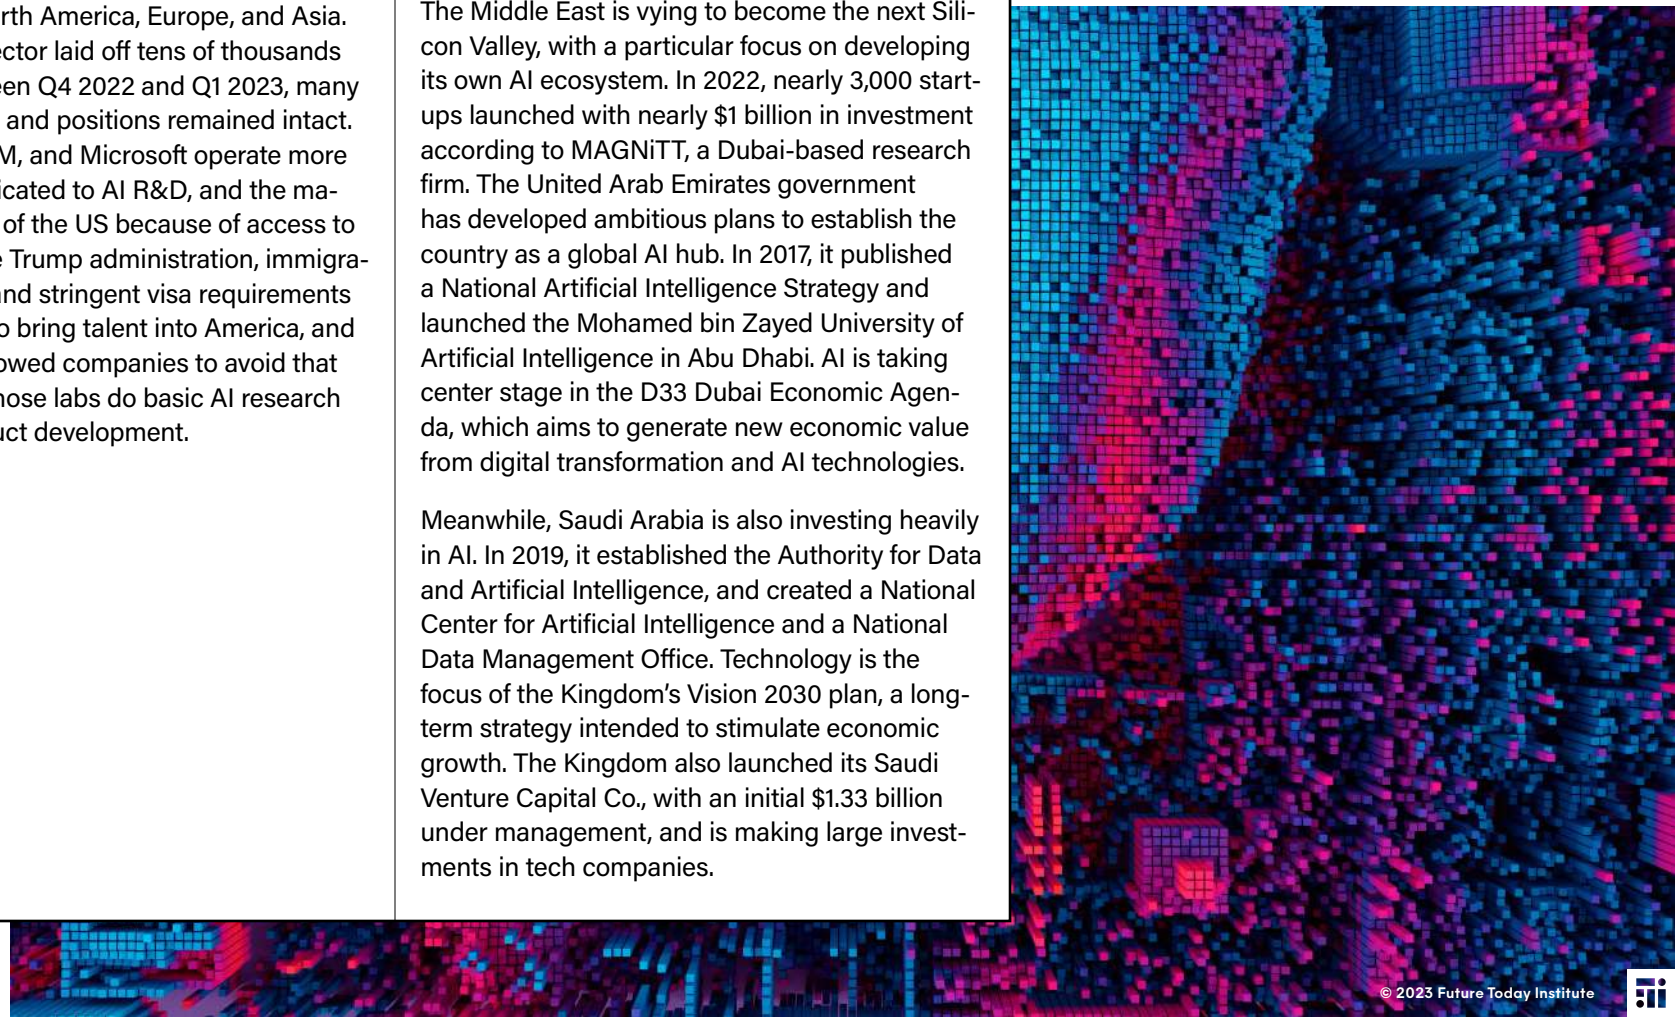

> By humanizing technology, we have this golden opportunity to reimagine how we connect with machines, and therefore, how we, as human beings, connect with one another.

DR. RANA EL KALIOUBY,
CO-FOUNDER OF AFFECTIVA, DEPUTY CEO OF SMART EYE

ENTERPRISE

TRENDS

ENTERPRISE

CORPORATE PARTNERSHIPS

Microsoft, Google, and AWS are all partnering with outside companies to compete in, and with, AI. Hyperscalers and challenger AI compute providers are tallying up major AI compute partnerships, most notably Microsoft's $1 billion investment into OpenAI. In 2019, Microsoft invested $1 billion into OpenAI between cash and credits for Azure, Microsoft's cloud computing business. OpenAI's ChatGPT or WebGPT might be integrated into Bing, which could make Microsoft's search engine the first real competitor to Google's hegemony since Explorer.

OpenAI expects $1 billion in revenue by 2024, according to Reuters. In 2023, we expect the biggest players to seek out or even announce strategic relationships with an AGI-focused organization. This is already happening quietly. The big three LLM service providers are partnered with hyperscalers: AI21 is on AWS, Cohere is on GCP, OpenAI is on Azure. Adept.AI, an ML research and product lab, has inked a deal with Oracle Cloud Infrastructure.

BUILDING IN-HOUSE SOLUTIONS

As companies adopted AI, they initially relied on packaged solutions from vendors. They could not invest both in talent and technology, and outsourcing was a path forward. While external partners made the transition easier, the one-size-fits-most approach has fallen out of favor. As new AI solutions emerge and companies approach maturity in their deployment and integration of technology, organizations are now opting to build solutions in-house. This means increased investment in hiring technical leadership as well as investment in the technologies and systems required to operationalize AI systems.

Some examples of this approach we're seeing from our Future Today Institute clients include: predictive risk and compliance management, automated supply chain planning, dynamic pricing models for retail and hospitality, automated call centers and customer service operations, assisted recruitment, and various financial applications (audits, forecasting, planning).

ADOPTING SERVERLESS COMPUTING

Hyperscalers like 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 variety 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.

While some engineers worry that such serverless systems require them to surrender too much control, enterprises have signed large commits with cloud solutions providers to obtain better pricing. CPU workloads are becoming commoditized and cheap, while GPU workloads allow enterprises to retire quota. All this makes serverless computing an easier pill to swallow when companies don't have to earmark big CapEx investment to build AI capabilities.

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. In the software world, 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 used in machine learning. Some of the fastest-growing GitHub projects are MLOps, or projects that deal with tooling, infrastructure, and operations.

Going forward, MLOps will describe a set of best practices that combines machine learning, traditional DevOps, and data engineering. Already, MLOps is becoming a formal discipline within organizations, as the overhead to transition AI experiments into production is quite high. This year, look for tools, automation options, and best practices to manage large-scale infrastructure.

LEVERAGING 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. Data affinity, compliance, and a lengthy vendor selection process has resulted in sticky CSP contracts, since switching vendors is proving to be a costly exercise. For that reason, the competition to cement new deals with enterprise customers is heating up. In the West, the field is led by Amazon, Microsoft, and Google, followed by IBM and Oracle. In Asian markets, Alibaba and Baidu dominate the AI cloud.

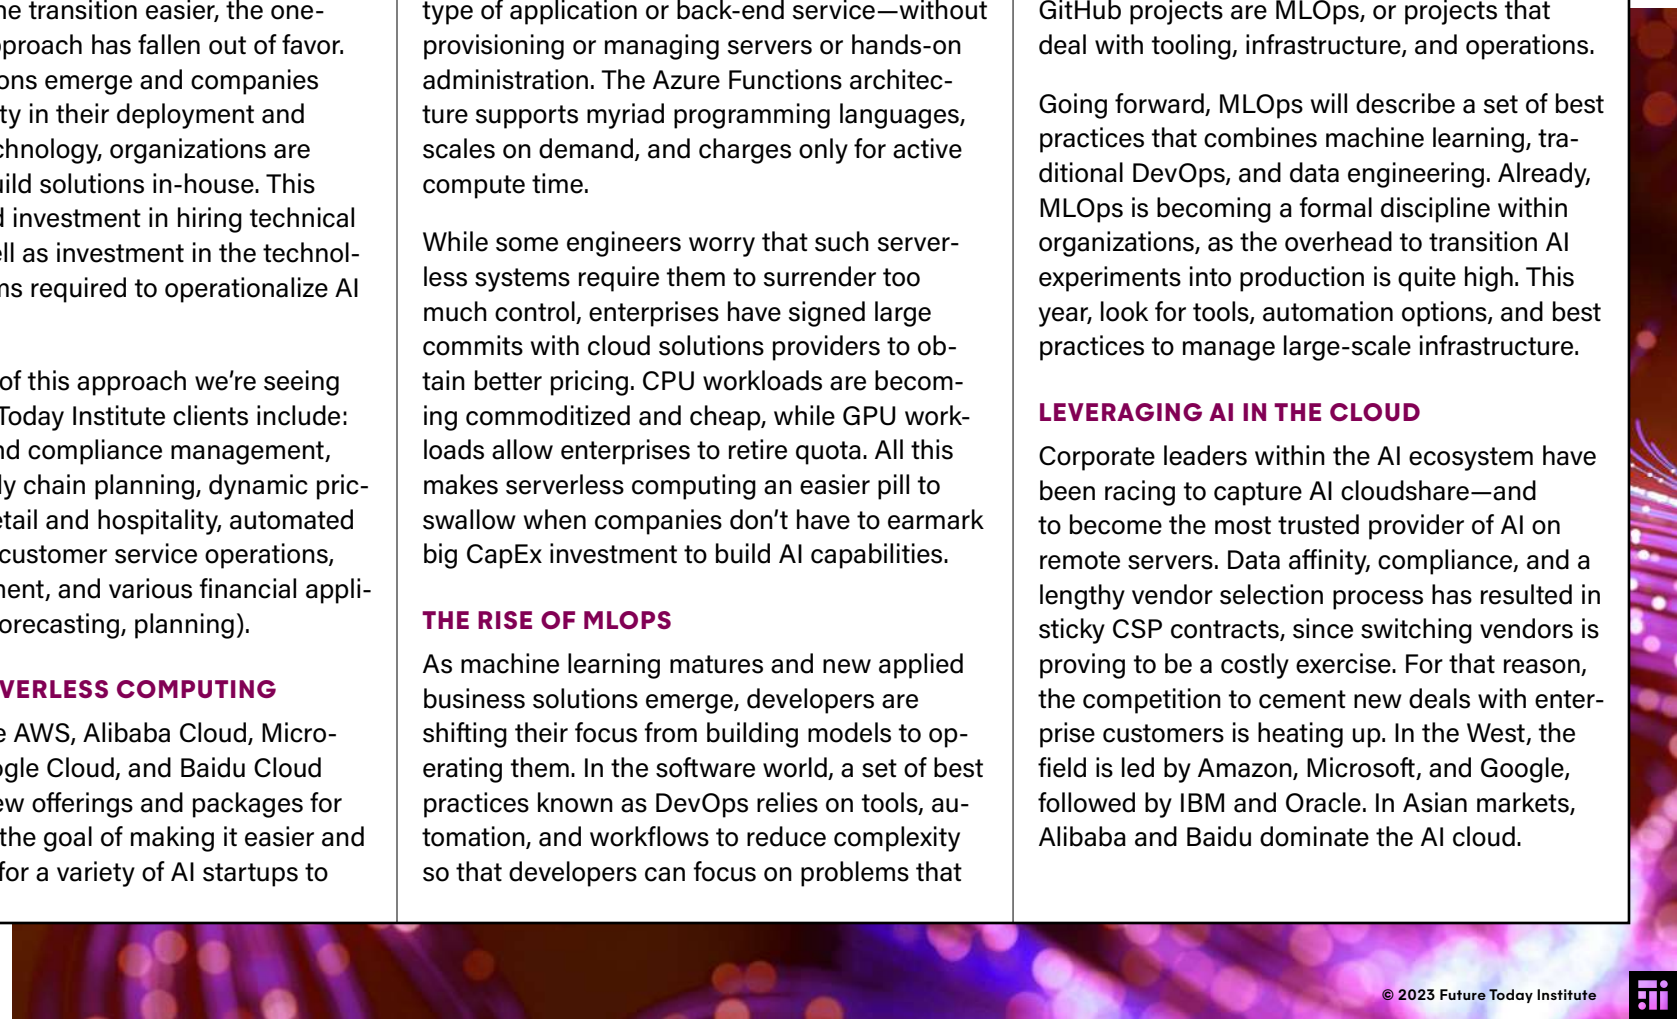

ENTERPRISE

IMPLEMENTING 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 (IoT) 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: Businesses can deploy complex event processing, machine learning, image recognition, and other high value AI without writing it in-house. Anyone is able to create AI modules and make them available to the community for use through the Azure Marketplace.

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 data sets 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 Cloud's AutoML includes image classification, object detection, translation, and all sorts of pattern recognition tools to allow developers with limited machine learning expertise to train high-quality models specific to their business needs. MakeML allows developers to create an AI app or solve business problems using computer vision.

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 AWS. Microsoft Power Apps is a low-code application development environment on Azure.

CONSOLIDATION IN AI'S ECOSYSTEM

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 a handful of big companies dominate the AI landscape: Google, Amazon, Microsoft, IBM, Meta, and Apple in the US, and Baidu, Alibaba, and Tencent in China, with significant fortification and support from their country's government.

They are fueled by a robust ecosystem, which includes Nvidia, Salesforce, Qualcomm, Intel Capital, Google Ventures, Salesforce, Samsung Ventures, and In-Q-Tel, among a small subset of others. 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.

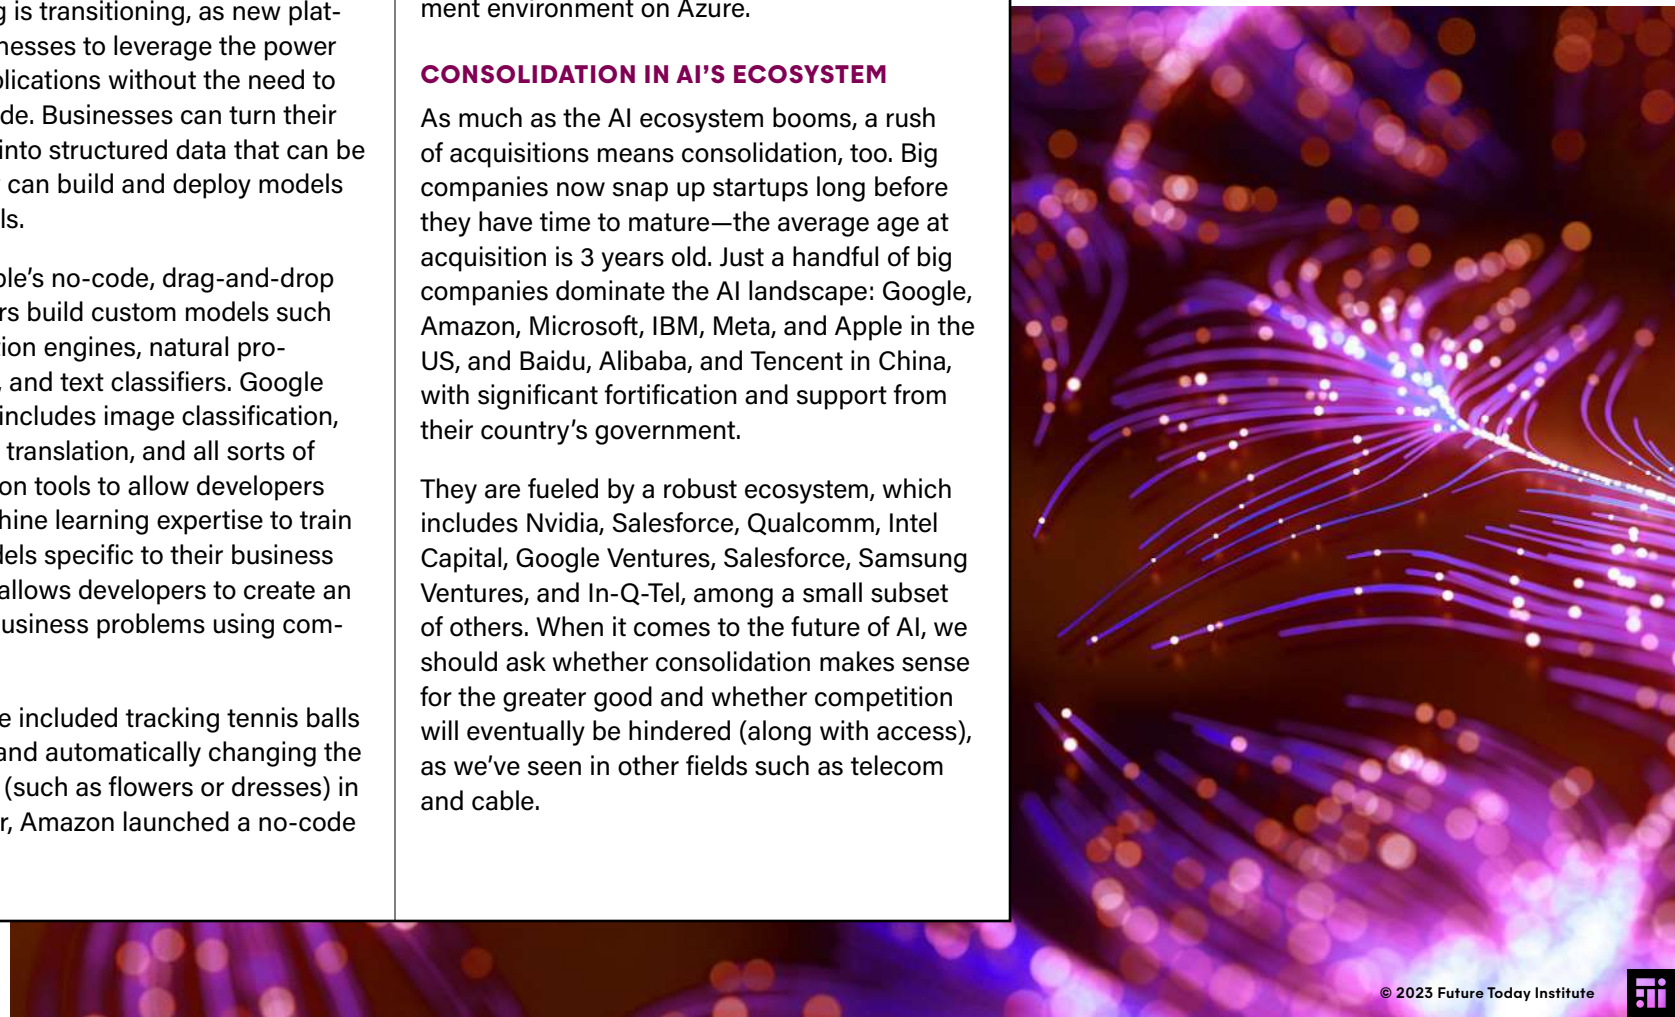

HARDWARE

In 2021, unveiled the largest AI processor ever made, the Wafer Scale Engine 2 (WSE-2).

Image Credit: Cerebras

ADVANCED AI CHIPSETS

Today's neural networks have long required an enormous amount of computing power, take a long time to train, and rely 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, Meta, Google, IBM, Nvidia, Intel, and Qualcomm, as well as startups Graphcore, Mythic, and Cerebras Systems, are all developing 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.

In 2019, Cerebras debuted 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—and yet in 2021, it announced that its next-gen chip, the Wafer Scale Engine 2 (WSE 2), has 2.6 trillion transistors, 850,000 cores, 40 gigabytes of on-chip memory, and 20 petabytes of memory bandwidth. Amazon's homegrown AI chip AWS Inferentia is a custom machine-learning chip used for high-performance inference predictions, and it now powers Alexa's back-end services.

The pretrained chips 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. Google's TPU, an AI accelerator application-specific integrated circuit (ASIC), is gaining traction in the enterprise and consumer marketplaces. Cerebras and Graphcore each announced breakthroughs in 2022, while Amazon Search reduced ML inference costs by 85% with AWS Inferentia. Traditional GPUs are still the mainstream, but we expect to see more developments in the AI chipset arena in 2023.

PROCESSING-IN-MEMORY TECHNOLOGY

A new approach to memory, which could eventually power the next generations of smartphones and help usher more smart devices to market, involves breaking some of the current bottlenecks in computing. Processing-in-memory (PIM) integrates a processor with RAM on a single chip, which allows computations to be performed in the memory of a computer.

In tests, PIM tech delivers significant performance gain while cutting energy consumption. In late 2021, Samsung, which in addition to phones and consumer appliances is also the world's largest manufacturer of dynamic random-access memory, announced a new PIM-enabled chip that could eventually double the performance of neural nets. This matters to the enterprise because it could bring AI applications to a wider array of devices in the near future.

APPLIED AI

Health Care AI Funding

US	$1 billion in 58 deals
Europe	$0.3 billion in 24 deals
Asia	$0.3 billion in 42 deals

Fintech AI Funding

US	$214 million in 19 deals
Europe	$292 million in 10 deals
Asia	$34 million in 42 deals

Source: CB Insights

APPLIED AI FOR HR

Recognition systems can now be deployed to watch people in an interview and gauge enthusiasm, tenacity, and poise. Algorithms analyze hundreds of details, such as their tone of voice, facial expressions, and mannerisms to best predict how a candidate will 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 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. The consumer advocacy organization Electronic Privacy Information Center filed a complaint with the US Federal Trade Commission requesting an investigation into HireVue, alleging its tools produce results that are "biased, unprovable, and not replicable" through algorithmic models.

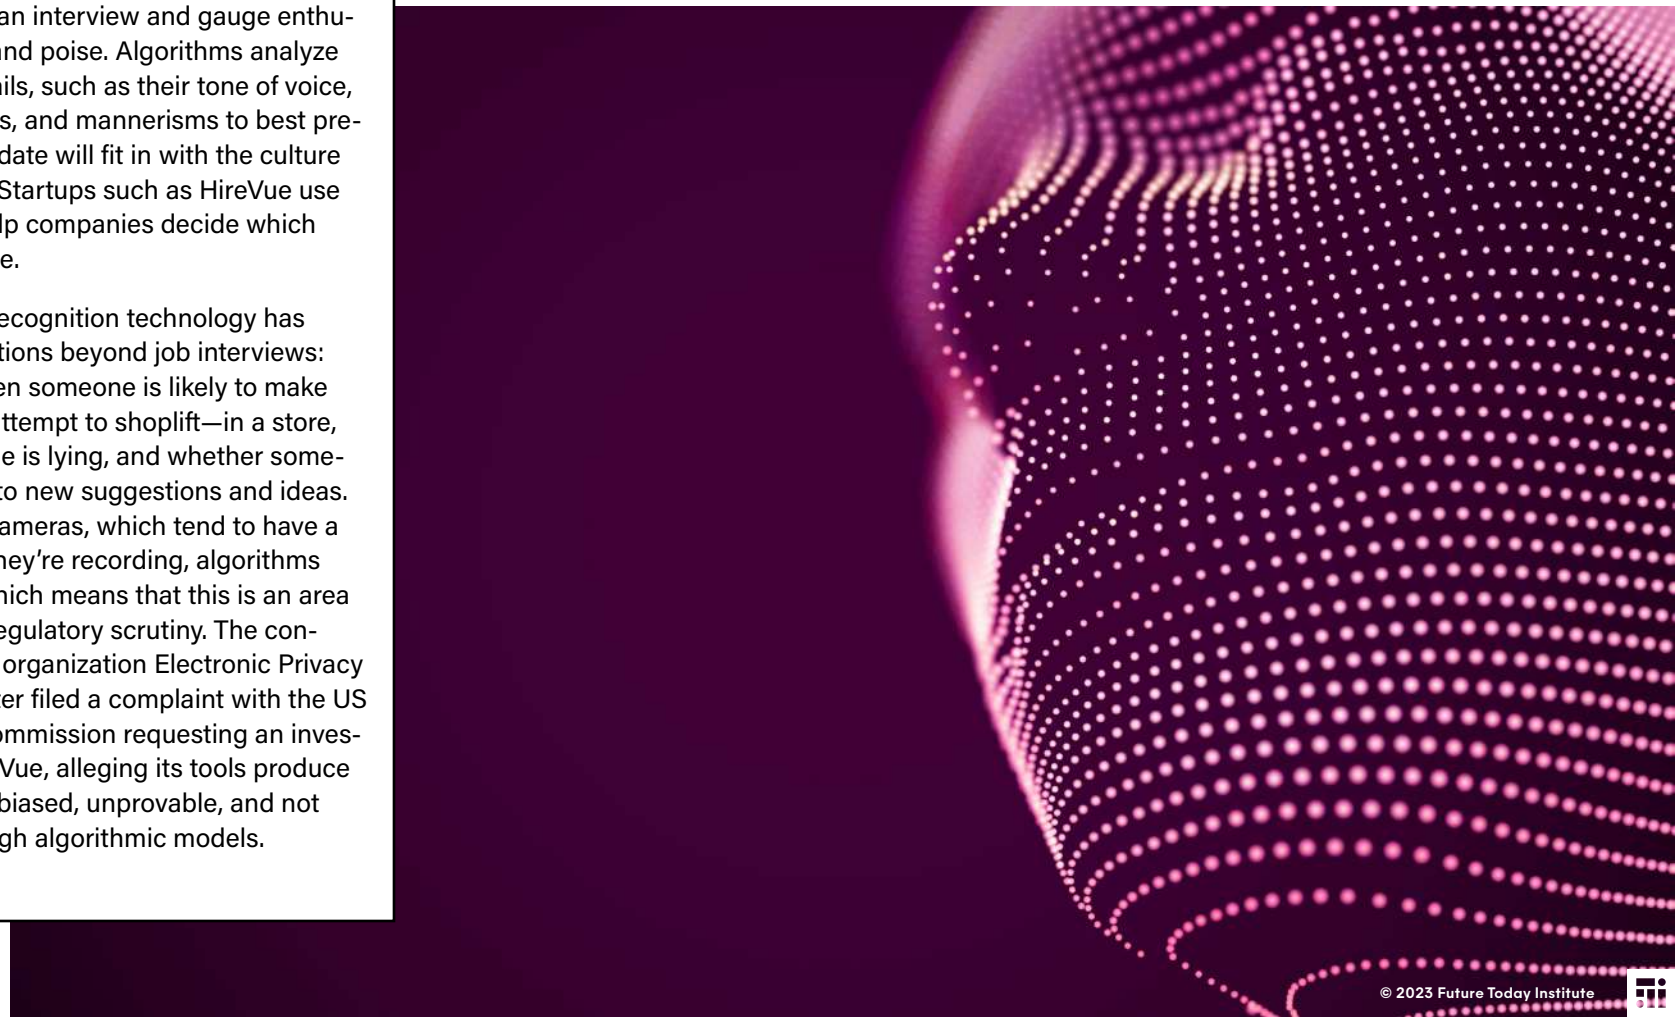

SPOTLIGHT: HEALTH CARE

PROTEIN FOLDING

In 2020, DeepMind's AI made a big announcement: It had solved a 50-year grand challenge with AlphaFold, an AI tool that predicts the structure of proteins. AlphaFold outperformed an estimated 100 teams in a biennial protein-structure prediction challenge called Critical Assessment of Structure Prediction, a problem that has long vexed biologists.

AlphaFold had previously bested other teams, but it worked so quickly and so accurately that it signaled a near future when the technology could be used regularly by other scientists. Along with the newest version of AlphaFold, DeepMind published full details of the system and released its source code. It also made a stunning reveal: AlphaFold 2 has predicted the shapes of nearly every protein in the human body, as well as hundreds of thousands of other proteins found in 20 of the most widely studied organisms, including yeast, fruit flies, and mice. All of this research will allow biologists to study and gain new insights on living organisms and pathogens, which will form the basis for new drug development.

AI TO SPEED UP 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, and their time.

Research labs now use AI systems to speed the process of scientific discovery. Recursion, a biotechnology company based in Salt Lake City, uses computer vision-based digital biomarkers, such as respiratory rate, to assess and track disease. Digitalizing in vivo studies will shorten the time to gather data and identify drug efficiency. 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, the algorithm determines what to change next. A 9- to 12-month process was completed in five days. DeepMind's AlphaFold will allow scientists to synthesize new drugs to treat diseases and develop enzymes that might someday break down pollution.

AI-FIRST DRUG DEVELOPMENT

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 then 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 pharmaceutical 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 humans 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 algorithms need their own clinical trials? Could AI be used to take shortcuts and undermine the value of the science conducted 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.

SPOTLIGHT: HEALTH CARE

AI TO IMPROVE PATIENT OUTCOMES

New medical algorithms address the level of patient care in the US. 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 the pain of Black patients is underscored. It's likely that the system itself, called the Kellgren-Lawrence Grade, 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.

ANOMALY DETECTION IN MEDICAL IMAGING

A new system designed to improve stroke outcomes, Viz.ai, showed promising results in a real-world 2021 study. An AI-based approach reduced the amount of time it took to detect brain stroke by 39%, which resulted in more patients identified as eligible for thrombectomy, a procedure that reduces the chances of long-term disability.

Radiologists and pathologists increasingly rely on AI to assist them with diagnostic medical imaging. In 2021, US 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, expanding the pool of health care workers who can read such scans. Nurses need just a few days of training on the software.

MEDICAL DEEPFAKES

Deepfakes are digitally manipulated content intended to deceive another person into believing they are real. Medical deepfakes result in diagnostic imagery that maliciously adds, or removes, tumors and other conditions. Cyber criminals are developing novel medical deepfake attacks intended to bring chaos to hospital systems and diagnostic centers.

Researchers at Ben-Gurion University and the Soroka University Medical Center demonstrated that tumors could be added or removed from CT images—and the deepfakes were good enough that radiologists didn't realize they were altered. After the radiologists were told that some of the scans had been edited, they misdiagnosed 60% of the AI-generated tumors and 87% of the tumors that had been edited out. In a study from Oslo Metropolitan University, researchers used AI to generate fake electrocardiograms, which visualize aspects of a human heartbeat. *(See also: Future Today Institute's 2023 Health and Medical Tech Trends report.)*

GENERATIVE ANTIBODY DESIGN

An antibody is simply a protein that protects an organism. Produced by the immune system, antibodies bind to unwanted substances and eliminate them. In January, researchers from the New York and Washington-based Absci Corp. showed how a generative AI model was able to design multiple novel antibodies that bind to a target receptor, HER2, more tightly than previously known therapeutic antibodies. What's interesting about this work is that researchers first removed all reference data on antibodies, so that the system couldn't just imitate and replicate the structure of known antibodies that work well.

The designs produced by Absci's system were both diverse (meaning, they didn't have counterparts known to already exist), and they received a high score on "naturalness," so they would be easy to develop and therefore catalyze a strong immune response. Using generative AI to design novel antibodies that function at the same level—or even better—than those designed by our own bodies marks a bold new step in using AI to reduce the speed and cost of therapeutic antibody development.

SELF-DRIVING MICROSCOPES

Researchers at the Oak Ridge National Laboratory are applying deep learning to microscopy, which until now had relied on humans to painstakingly organize, observe, and analyze microscopic samples. Deep learning will automate much of that process while also extracting more information from samples. A smart microscope would be able to view and analyze samples

SPOTLIGHT: HEALTH CARE

Sean McClain, founder and CEO of Absci, uses deep learning AI and synthetic biology to expand the therapeutic potential of proteins.

Image Credit: Absci Corporation

in real time—while techniques that introduce a stimulus, such as an electric field, will allow researchers to gain real-world understanding in more dimensions and on a nanoscopic scale. This could become a force multiplier in how scientific discovery is conducted.

NLP ALGORITHMS DETECT VIRUS MUTATIONS

Natural language processing (NLP) algorithms, which are typically used for words and sentences, are also 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 produce text 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. MIT 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.

USING AI TO IMPROVE TALK THERAPY

The success of therapy and counseling requires trust between a clinician or doctor and patient—and that trust is built through dialogue. AI is now being used to quantify linguistic interactions to determine what techniques work best. Startup Lyssn translates natural language into structured data and generates digital voiceprints, which identify the sentiment attached to each sound. It's hoped that this technology will be used to improve the techniques therapists use for cognitive behavioral therapy, PTSD therapy, and other forms of talk therapy delivered in-person or via telemedicine.

SPOTLIGHT: FINANCE AND INSURANCE

MITIGATING FRAUD

JPMorgan Chase spent $100 million to develop anti-fraud systems for consumer payments, which has paid off—its AI systems reduced fraud by 14% from 2017 to 2021. It deployed an algorithm to detect fraud patterns: When a credit card transaction is processed, trained models determine whether fraud is occurring.

PREDICTING FINANCIAL RISK

AI systems can help improve loan underwriting and reduce financial risk. Models are being trained to recognize anomalous activity and to develop forecasts for a variety of middle- and back-office applications. For example, US Bank relies on deep learning to analyze customer data as well as to root out money laundering schemes.

PREDICTING WORKPLACE INJURIES

AI systems are being trained to detect possible workplace injuries. Using AI-based computer vision models, Turkey-based Intenseye can detect 40 types of employee health and safety incidents in real time. The company says that it does not capture personally identifiable information from the visual data it processes and that it detected 1.8 million unsafe acts in 2020 and 2021. San Francisco–based Voxel uses computer vision to enable security cameras to automatically detect high-risk activities in real time.

IMPROVING DAMAGE ASSESSMENT

Insurance companies are applying AI to assess damage and improve forecasts. The Vehicle Damage Inspection model, which is available on AWS Marketplace, uses a machine learning model to determine what part of a car is damaged. After photos are uploaded, it assesses loss and dramatically reduces the amount of time required for human appraisers to conduct their analysis. Following catastrophic typhoons and weather events in Japan, local insurance companies are relying on computer vision to assess damage after a natural disaster. Sompo Japan is using the Tractable AI Estimating system to calculate the approximate repair cost of damaged homes.

CONSUMER-FACING ROBO-ADVISERS

Automated assistants are moving from the fringe to the mainstream as consumer adoption increases. Robo-advisers offer algorithm-based portfolio management advice to investors, applying parameters like risk tolerance and desired returns. These investment tools offer some tangible benefits over their traditional, human counterparts: they can provide more services at a lower cost, they're able to digest and interpret mounds of data in real time, and they don't take part of the weekend off to golf. Wealthfront is an AI-powered system for consumers: It suggests fund managers and calculates probable risk levels based on the user's personal information and preferences.

AI CLAIMS PROCESSING

While human claims writers must painstakingly review pictures and reports to assess damage, compare what they see to coverage policies, and make a determination about appropriate actions, an AI system can digest the same data and accomplish the same work in a matter of minutes. Using a suite of tools—natural language processing for policy review, and computer vision recognition to spot anomalies in photos and videos—claims can be processed efficiently and, it's believed, more accurately. AI-powered claims processing reduces the overhead for businesses and wait times for customers. Some insurance providers are wading into a new pool of opportunities. Liberty Mutual's mobile app has started to integrate ML for damage assessment—it informs customers about their coverage and next steps.

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 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 in certain patients?

These problems could put a company at risk of lawsuits, and new insurance models are needed to address these issues. Underwriters are starting to include AI under cyber insurance plans, while specialty insurers such as La Playa's Science and Tech Insurance now offer coverage for AI applications.

> If I hid Ava from you so you could just hear her voice, she would pass for human. The real test is to show you that she's a robot and then see if you still feel she has consciousness.

NATHAN, SPEAKING TO CALEB IN *EX MACHINA*

CREATIVITY & DESIGN

TRENDS

CREATIVITY & DESIGN

We used Shutterstock's generative AI tool to design a chair using the prompt: "midcentury modern chair with a neon pink cushion.

AI-ASSISTED INVENTION

Stable Diffusion, MidJourney, DALL-E 2, and ChatGPT are now widely accessible to end consumers, leading to AI-assisted human creativity. But these systems were all trained using other artists' works. If a business uses an AI-generated image, video, or text for commercial purposes, does it owe anything to those whose original works were used for training? Likewise, what if a generative AI system invents a product that's eligible for a patent?

In 2021, the South African government granted a patent to an AI system called Dabus, which invented a method to interlock food containers. It was a world-first—previously, patents had only been awarded to humans. In the US, the application was rejected, with a judge citing case law stipulating that only a human can hold a patent. There may be business cases for an AI to hold a patent rather than an individual. It raises the question: What happens when AI systems co-invent, or even entirely invent, new products? We're likely to hear more debate on this topic this year.

ASSISTED CREATIVITY

Generative adversarial networks (GANs) are capable of far more than generating deepfake videos. Researchers are partnering with artists and musicians to create entirely new forms of creative expression. From synthesizing African tribal masks to building fantastical, fictional galaxies, AI is exploring new ideas.

In 2019, 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. We anticipate that models like Stable Diffusion and DALL-E 2 (among others) will be more deeply integrated this year into a wide variety of creativity apps and software suites.

NEURAL RENDERING

Starting with a 2D image, researchers can now create a rich 3D view of a scene by using a neural network to capture and generate spatial imagery. Called neural rendering, the process captures a photorealistic scene in 3D by calculating the density and color of points in space. The algorithm converts 2D pixels into voxels, which are a 3D equivalent. The result is a video that looks convincingly real. The many applications for neural rendering include amping up autonomous driving to help train algorithms to recognize and react to novel road situations. This technology will influence the future of video games, virtual reality, and emerging metaverse environments.

CREATIVITY & DESIGN

GENERATING VIRTUAL ENVIRONMENTS FROM SHORT VIDEOS

Nvidia is teaching AI to build realistic 3D environments from short video clips, a method that builds on previous research on GANs. Nvidia's system generated graphics based on open-source data sets 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.

AI VOICE GENERATORS

In the past year, there has been an explosion of AI voice generator startups. Any podcaster 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 the spoken word could be edited the way a word document is edited? That's the promise of AI companies, including Resemble AI, Murf AI, Lovo AI, Play.ht, and Descript, which make it possible to clone or create voices. The latest generation of AI voice generators are capable of mimicking human emotions like anger, happiness, sadness, and trepidation.

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. This could improve the soundscapes created for AI-generated movies—but it might also help us imagine soundscapes for both imaginary worlds (Laconia, from "The Expanse") and real ones (Mars).

TRENDS

SOCIETY & ETHICS

SOCIETY & ETHICS

DETECTING EMOTION

A new type of neural network can determine how people are feeling. Using radio waves, AI can detect subtle changes in heart rhythms, run a pattern analysis, and predict someone's emotional state in a given moment. A team from Queen Mary University of London used a transmitting radio antenna to bounce radio waves off test subjects and trained a neural net to detect fear, disgust, joy, and relaxation, as people were shown different videos. The system accurately tagged emotional states 71% of the time, which signals new opportunities for health and wellness applications, as well as for job interviews and the government/military intelligence community.

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 data set are developing accurate models. Amazon's Rekognition API infers someone's emotions using facial recognition and physical appearance. (Though, Amazon is quick to point out in its documentation that the "API is only making a determination of the physical appearance of a person's face. It is not a determination of the person's internal emotional state and should not be used in such a way.")

Replika uses AI to evaluate voice and text, and over time the personal AI bot mirrors the user in "conversations." Affectiva Human Perception AI analyzes complex human states with 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 the program can make real-time suggestions to improve the driver's performance. (Affectiva was acquired by Smart Eye in 2021 but has kept its name.)

THEORY OF MIND MODELS

Research teams are attempting to teach machines about love, active listening, and empathy. The question becomes: What is an authentic emotion? "Theory of mind" refers to the ability to imagine the mental state of others, and 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 relational agent for mental health.

By designing machines to respond with empathy and concern, these technologies could eventually end up in hospitals, schools, and prisons, providing emotional support robots to patients, students, and the incarcerated. According to health insurer Cigna, the number of Americans who report feeling loneliness has doubled in the past 50 years. In our increasingly connected world, people report feeling more isolated. Governments struggling with a massive mental health crisis may turn to emotional support robots to address the issue at scale.

CONSUMER APPLICATIONS

GENERATIVE AI FOR PERSONAL EXPRESSION

In 2022, experimental apps moved quickly into the mainstream, as consumers enthusiastically generated AI avatars of themselves. Lensa AI, a photo-editing app that creates avatars automatically by using a handful of photos, became one of the most downloaded apps in the world. It produced dozens of images depicting a variety of scenes and styles, such as an anime hero, a Renaissance-era princess, or a modern artist.

Amper Music creates and mixes tracks using parameters set by the user, while OpenAI's MuseNet launched a tool that can generate songs with up to 10 different instruments and music in up to 15 different styles. It'll also mimic a famous artist, such as Mozart. Generative AI apps haven't been without controversy: Lensa AI tends to oversexualize women, while OpenAI's generator also samples from modern composers who don't receive any royalties. (We used it to re-create Luis Fonsi's "Despacito" in the style of Russian-American composer Rachmaninoff, which sounded just as weird as you're imagining.)

TURNKEY CONSUMER-GRADE APPLICATION DEVELOPMENT

Low-code and no-code offerings from 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 create 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. As more text-to-code models emerge, we're likely to see broader adoption.

RESPONSIVE RECOGNITION TECHNOLOGY

Real-world conversation is full of nuance: We use words and emphasis in unique ways, we interrupt each other, and sometimes we need others to help us express what we're thinking. All of these communication styles pose serious hurdles for AI, which doesn't adapt as easily to a multivariate situation, such as everyday people talking to one another.

Soon, Amazon's Alexa, using responsive recognition technology, will join conversations in a way that feels both natural and useful. Upgrades will make the digital assistant more responsive, proactive, and humanlike. Beyond engaging in conversation, Alexa will be able to detect other atmospheric sounds such as snoring, coughing, the cries of a baby, or a dog barking, and can then respond with a set of commands. The next time you have a coughing fit, don't be surprised if Alexa adds chicken soup to your next Amazon order.

EFFORTS TO THWART RECOGNITION SYSTEMS

As facial recognition becomes ubiquitous, various groups want to limit the technology's effectiveness to protect privacy. While methods of confusing or obscuring facial recognition systems are not always feasible, researchers have begun trying to confuse online applications that scrape and collect images used as inputs for training facial recognition engines for the purpose of developing a form of camouflage, which consumers may someday demand.

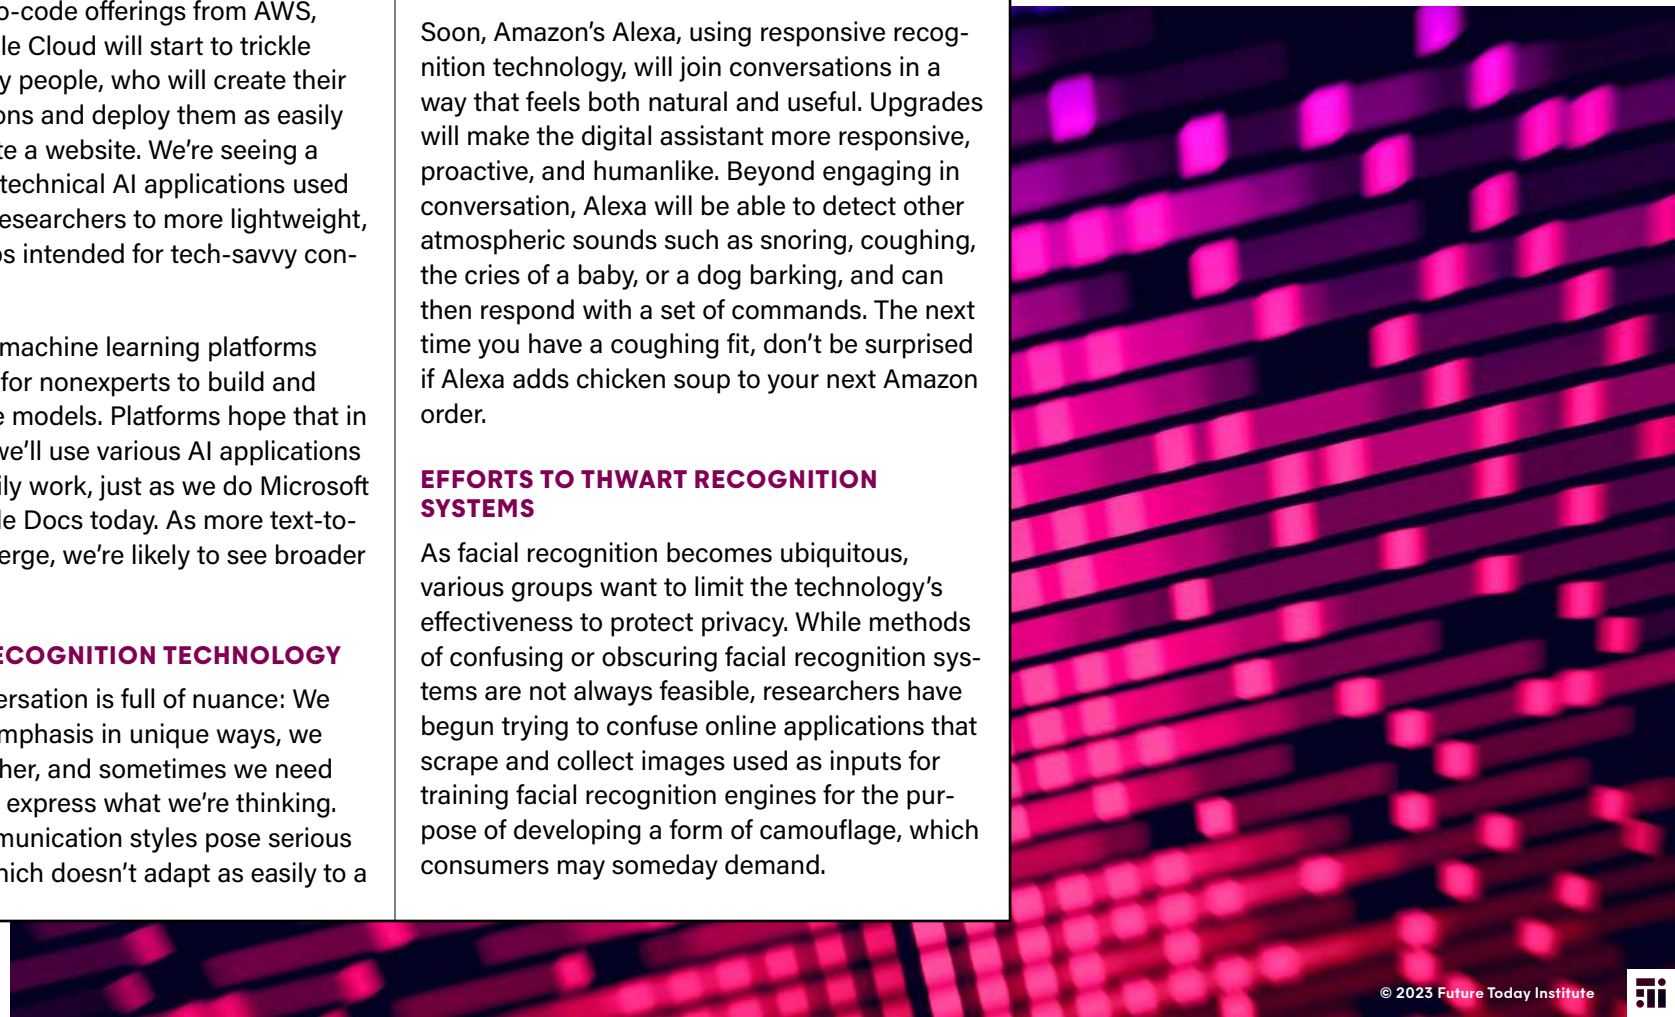

CONSUMER APPLICATIONS

We created these photos of award-winning actor Pedro Pascal using Lensa AI.

Researchers from the University of Chicago have created a program, Fawkes, that adds extra pixels to images to cause facial recognition apps to misclassify faces. Taking this principle a step further, Israeli artificial intelligence company Adversa AI adds noise, or small alterations, to photos of faces, causing algorithms to detect a different face than what is visible to the naked eye. The algorithm is successful at imperceptibly changing an individual's image to someone else of their choosing.

BIOMETRIC SCORING

Quantifying and analyzing our biometric data can reveal patterns in our activities and a lot about who we are, what we're thinking, and what we will likely do next. Behavioral biometrics use machine learning to understand hundreds of unique biometric data points to understand, authenticate, nudge, reward, and punish. Behavioral biometrics tools can be used to map and measure how a user types—what force is used to press down on screens, whether the Cs and Vs on a phone are tapped with a fat finger, and how quickly fingers flick when hunting through search results.

Those tools know a user's unique typing pattern on a physical keyboard, too—whether someone constantly spells the word "behavioral" wrong on the first try, or holds down or repeatedly taps on the delete button. Most of us are unaware that we have certain identifiable behaviors, but machines perceive them. In the near future, such patterns will pose security vulnerabilities—as well as interesting new opportunities. Imagine never having to use a password again; your bank would simply recognize a customer's typing pattern after a few sentences. The downside: If behavior is observable, at some point it will become repeatable, too, which represents a security risk.

> We don't have the luxury of a long time to actually even out the effects on job loss with this revolution, because it's happening so quickly.

KAY FIRTH-BUTTERFIELD,
HEAD OF AI & MACHINE LEARNING AND MEMBER OF THE EXECUTIVE COMMITTEE OF THE WORLD ECONOMIC FORUM

SAFETY

AI GENERATORS AND COPYRIGHT INFRINGEMENT

Data sets are growing larger and more controversial. Creators are concerned about AI systems being trained on vast amounts of their copyrighted work with no consent, no credit, and no compensation. Early this year, a lawsuit argued that Microsoft, GitHub, and OpenAI infringed on the rights of developers, whose code was scraped from the web and used to train CoPilot. Stable Diffusion, Midjourney, and DeviantArt (which just launched its own AI art generator DreamUp) were also targeted with copyright lawsuits, which allege that generative AI tools violate US copyright laws prohibiting the use of creative work without consent.

Three artists—Sarah Andersen, Kelly McKernan, and Karla Ortiz—argue that AI generators infringed on the rights of "millions of artists" by training their AI tools on 5 billion images scraped from the web "without the consent of the original artists." In what could prove to be a landmark case for how AI systems are trained, Getty Images announced in January that it is suing Stability AI (which makes Stable Diffusion) for infringement. Whether these systems meet the qualifications of copyright infringement is a complicated question—and this year it could be decided in a court.

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 for training. The Multiple Encounter Dataset contains two large data sets of photos: people who have not yet committed a crime and an FBI data set of deceased people. The data sets over-index 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 challenge, because researchers need large data sets to perform their work. Often, images are used without consent.

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 system 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. 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 had learned to hide information about the original image inside of the image it generated.

INCREASED USED OF AMBIENT SURVEILLANCE

What happens behind closed doors may not be secret for long, and executives should beware of new ambient surveillance methods. Scientists at MIT 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.

TRUST

AI ALIGNMENT

As AI systems improve, some researchers are insisting on guardrails to ensure that they are deployed in ways that do not harm humanity. One area of concern is known as AI alignment, which explores different scenarios in which AI systems are built with goals that align with society's values. OpenAI, DeepMind, and Anthropic (which describes itself as an "AI safety and research company") each have AI alignment teams with dedicated staff researching guardrails. While the total number of researchers working on AI alignment is small compared to the rest of the AI community, such dedicated teams did not exist until recently.

ADDRESSING BIAS

Seemingly the moment OpenAI's ChatGPT went public, there were multiple accounts of the system displaying racism, ageism, and gender bias. To the company's credit, it appeared to be resolving those issues in real time. In early December 2022, Steven T. Piantadosi, the head of the computation and language lab at the University of California, Berkeley made ChatGPT write code to say that only white or Asian men would make good scientists.

OpenAI immediately stepped in. Trying to replicate the code returned the following response: "It is not appropriate to use a person's race or gender as a determinant of whether they would be a good scientist." Still, the AI community has struggled with a serious and multifaceted bias problem for decades. Researchers including Timnit Gebru, the Ethiopian-born founder and executive director of Distributed AI Research Institute, and Abeba Birhane, another Ethiopian native and a senior fellow in Trustworthy AI at Mozilla, are publishing new research on responsible AI practices.

SYNTHESIZING AND GENERATING TRUST

Humans can be tricked into believing machine-generated faces, especially when they've been engineered to elicit trust. A study published in the Proceedings of the National Academy of Sciences shows that synthetic faces are often "deemed more trustworthy than real faces," suggesting that synthetic faces could be designed as societal malware. If a bad actor were attempting to undermine institutions, it could deploy a synth on social media to sow distrust. There are not yet effective countermeasures for synthetic humans, or effective markers to help consumers distinguish between fake and real.

WORKER SURVEILLANCE

The rise of remote work during the pandemic accelerated the surveillance of workers, and will likely continue to grow as remote and hybrid work models take root. The US Constitution's Fourth Amendment, which prohibits unreasonable searches and seizures and precludes most uses of this same technology by law enforcement, doesn't apply to private companies.

Teleperformance, a French company that manages outsourced call center work for many Fortune 50 companies, uses cameras and AI to monitor its teams. It flags employees as idle when it detects they haven't used the keyboard or mouse within a specified time frame. Live Eye Surveillance offers a monthly subscription service that remotely monitors live video feeds of employees for companies such as 7-Eleven, Dairy Queen, and Holiday Inn. Sneek is another example of "tattleware" that captures live photos of employees via webcams and displays them on a digital wall viewable by everyone in the company. Click on a photo and it instantly pulls that person into a video call with you.

The most well-known user of worker surveillance might be Amazon, which has installed AI-enabled cameras in delivery trucks to track behavior. The company docks driver pay if it perceives unsafe conditions such as distracted driving, speeding, or hard braking. In its warehouses, the company monitors worker productivity by measuring Time Off Task, which is any time when a worker isn't actively processing products. South Korean e-commerce giant Coupang, which has pledged to become the "Amazon of Korea," uses similar surveillance tactics.

The industry has also continued to evolve as it offers more AI-based analysis of workers. Amazon is exploring using keystroke-logging software that tracks user behavior over time to detect if the same person is controlling the worker's account. Aware's Spotlight software detects behavioral changes like mood, tone, and attitude across conversations on employees' devices. Teramind offers software that will disable private conversations if it detects "inappropriate" keywords. With the top three tools in the industry accounting for over 60% of global demand, expect to see more AI-based surveillance that leverages the growing pool of data collected by a variety of companies.

TRUST

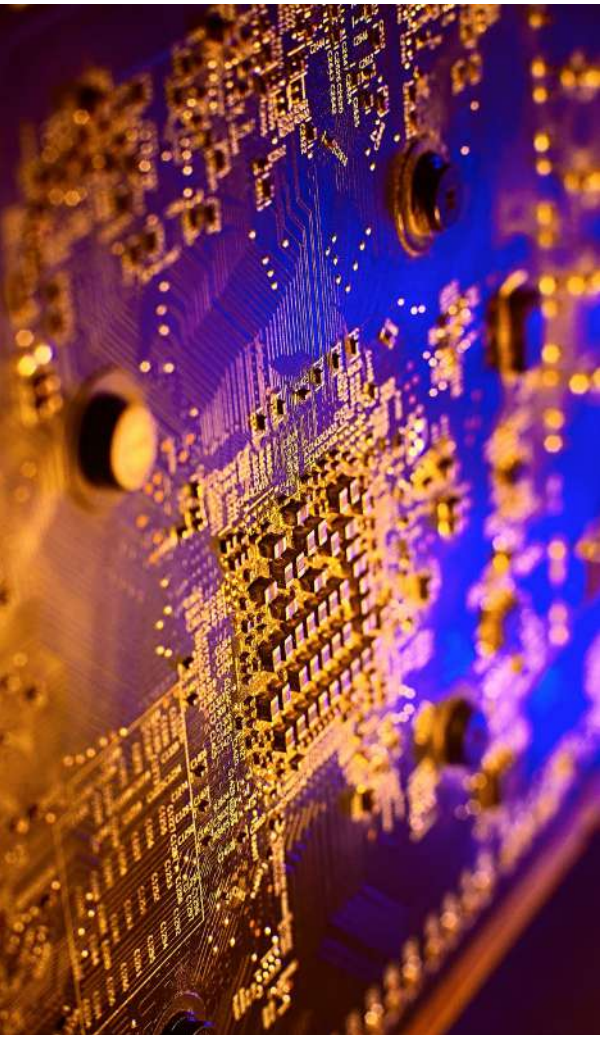

SCHOOL SURVEILLANCE

During the pandemic, many students were issued laptops and other devices by schools to facilitate remote learning. They weren't told, however, that these devices would open a portal into their homes that could be monitored by schools at all times of the day. In the US and many other countries, schools can legally monitor students, often without disclosing what is being tracked.

Gaggle is one company that monitors school-issued accounts and uses AI to track online behavior of students across services like email and chat tools. In 2020, the Minneapolis school district signed a contract with the company to monitor its students through 2023. School districts across the U.S. use Securly to monitor students in real time, looking for prohibited behaviors such as having too many browser tabs open. The software enables teachers to close tabs for any students they believe are "off task."

Philadelphia and Chicago schools deployed GoGuardian software on district-issued Chromebooks. A vulnerability in the software allowed teachers to start virtual sessions that enabled webcams on those Chromebooks without notification or consent by the student. Schools in China deploy technology to monitor attentiveness in students. An algorithm called 4 Little Trees is used in Hong Kong to detect students' emotions as they learn—by monitoring their facial expressions with webcams. If the system detects a lack of focus, it nudges the student to pay attention.

PRIORITIZING TRUST

We will soon reach a point when we will no longer be able to tell if a data set 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.

Building more transparency is a challenge, as 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 trustworthiness 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 diversity among developers—representing different races, ethnicities, and genders—will reduce inherent bias in AI systems.

> There's a real danger of systematizing the discrimination we have in society [through AI technologies]. What I think we need to do — as we're moving into this world full of invisible algorithms everywhere — is that we have to be very explicit, or have a disclaimer, about what our error rates are like.

DR. TIMNIT GEBRU,
FOUNDER & EXECUTIVE DIRECTOR,
THE DISTRIBUTED AI RESEARCH INSTITUTE (DAIR)

GOVERNMENT & DEFENSE

TRENDS

GOVERNMENT

RISE OF AI NATIONALISM

Governments are instituting new restrictions on mergers and acquisitions and investment activity to ensure that AI does not aid foreign adversaries. The US Senate overwhelmingly passed legislation in June 2021 that dedicated $250 billion to scientific and technological research. Its centerpiece, the Endless Frontier Act, was designed to boost American competitiveness against China, especially in AI. 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 the country.

Meanwhile, in China the Ministry of Science and Technology established 20-city AI pilot zones that should open by the end of this year. They will carry out AI-based policy experiments and social experiments, according to official government documents. Meanwhile, China is planning for a world without American technology, with the government prioritizing homegrown technology companies and software systems.

COUNTRIES DEVELOP NATIONAL AI 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 citizens; and the United Arab Emirates has a sweeping set of policy initiatives on AI, and appointed H.E. Omar Sultan Al Olama as its minister of state for artificial intelligence.

In the US, numerous public and private groups work independently on the future of AI on behalf of the nation, including the National Artificial Intelligence Initiative and the National Security Commission on Artificial Intelligence. Those efforts, however, lack interagency collaboration and coordinated efforts to streamline goals, outcomes, R&D efforts, and funding. A new wave of countries will launch national AI strategies this year. The OECD.AI Policy Observatory now maintains a live repository of more than 700 AI policy initiatives from 60 countries, territories, and the European Union.

EUROPE REGULATES AI

The EU will finalize its Artificial Intelligence Act (AIA) this year, which will focus on the use of AI systems that use citizen data for potentially detrimental purposes. It's viewed as the world's first attempt at creating broad, enforceable standards governing the use of AI. The AIA distinguishes between "high risk" and "low risk" systems. For those categorized as high risk, AI systems must meet several criteria: (a) they should work as the user intended and should be interpretable; (b) they should be secure and accurate; (c) they should contain all necessary technical documentation for proper use and keep logs of their behavior; (d) they must have effective human oversight.

The AIA would cover any person or organization (including foreign entities) that use an AI system operated within the EU. While AI developed for national security or military uses will likely be exempt from the EU's new regulations, social media platforms like Facebook and TikTok could find it difficult to comply.

NATION-BASED GUARDRAILS AND REGULATIONS

From self-driving car accidents to election interference, disinformation campaigns to political repression enhanced by facial recognition and automated surveillance, events over the past few years have exposed the dangers of AI. Few guardrails exist for a technology that will touch every facet of humanity, and countries are racing to develop strategies and guidelines to oversee it.

The EU developed an AI Alliance and plan of cooperation between member countries, and Estonia is developing its own legal framework governing the use of AI. In 2020, China moved into position to lead the first set of global AI norms and standards. It had previously published a report on technical standards that would allow companies to collaborate and make their systems interoperable. The EU and the Organization for Economic Cooperation and Development (OECD) published their own guidelines.

While these efforts could introduce new ways to safeguard against bias and to ensure trust, they also each attempt to create strategic advantag-

GOVERNMENT

The United Arab Emirates is establishing a geopolitically neutral hub for AI under the direction of H.E. Omar Sultan Al Olama, UAE Minister of State for Artificial Intelligence, Digital Economy, and Remote Work Applications.

Image Credit: Dubai Chamber of Digital Economy

es for stakeholders. As AI continues to develop according to different rules in China, the EU, and the US, one of the hallmarks of the field—global academic collaboration—could drastically decline. This could worsen as visa allowances increasingly become politicized around the world.

REGULATING DEEPFAKES

The US National Defense Authorization Act includes provisions that address the growing problem of deepfakes, requiring the Department of Homeland Security to issue an annual report for the next five years on the risks posed by deepfakes. In 2021, the US Senate Committee on Homeland Security and Governmental Affairs voted unanimously to advance the Deepfake Task Force Act, which would establish a public-private team to investigate technology strategies and to develop policies that could curb risk. Bills to regulate or prohibit the use of deepfakes have been introduced in California, Texas, and Massachusetts, and a number of other federal bills are under review. These initiatives will likely be met with arguments that prohibiting deepfakes infringes on free speech rights.

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 US, Germany, Japan, India, South Korea, the UK, 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. Meanwhile, the United Arab Emirates and Saudi Arabia are working to remain neutral hubs for innovation and business.

DEFENSE

AUTONOMOUS WEAPONS POLICIES

Late January 2023, the US Department of Defense updated its guidance on autonomy in weapons systems. The original 2012 policy, and a 2017 update, did not explicitly mention AI. This new directive is aimed at helping to clarify the process for developing autonomous or semi-autonomous weapons systems. Previous policies, such as the Ethical Principles for Artificial Intelligence (2020) and Responsible Artificial Intelligence Strategy and Implementation Pathway (2021), were intended to guide decision making for the development and deployment of AI within the military.

In late 2022, NATO released its Autonomy Implementation Plan, which said that AI systems "offer clear opportunities, including bolstering deterrence and defense, preserving NATO's technological edge, increasing resilience and adapting to the security impacts of climate change." It spelled out a multipronged plan to develop and adopt autonomous systems. As AI technologies advance, so too will policy: We anticipate additional policy updates in the next two years to account for emerging AI capabilities.

SIMULATING WARFARE

Given the rising tensions between the US and China over Taiwan, several groups are building AI-powered simulation tools to war-game a future conflict. In China, the People's Liberation Army has been using AI simulation tools to prepare for military operations against Taiwan.

The Center for Strategic and International Studies, a bipartisan, nonprofit policy research organization, developed a war game involving an amphibious invasion of Taiwan. After 24 rounds of gameplay, the US and its allies Japan and Taiwan successfully defeated a conventional amphibious invasion by China. While Taiwan remained autonomous in the simulation, its economy was devastated and the US lost hundreds of aircraft and tens of thousands of lives—while the Chinese Communist Party never really destabilized. Games that use real-world data to run simulations are augmenting the work of military strategists, so that leaders can validate or revise their postures on deterrence, invasion, and defense.

UNPILOTED MILITARY VEHICLES

In 2022, a retrofitted Black Hawk helicopter flew autonomously between mountains to deliver blood supplies in a simulated mission. Part of a Defense Advanced Research Projects Agency (DARPA) program called ALIAS, aimed at converting aircraft that can fly autonomously, the Black Hawk flew 134 miles on its own. In another test flight, the helicopter successfully changed flight paths to pick up a wounded soldier. Once the autonomous system is given a mission, the aircraft is designed to make all necessary decisions to complete the mission as directed. The near-term hope is that an AV system could replace human pilots in dangerous situations to perform resupply runs—or even take over for pilots, allowing them to rest.

AI USED TO GUIDE MILITARY STRIKES

In 2021, the US military said that it had started using AI to guide its airstrikes, deploying algorithms to a live operational kill chain. The kill chain is a process of gathering intelligence, performing analysis, weighing risks, and deploying weapons to destroy a target. Using a modified process, an AI system was deployed into the Air Force Distributed Common Ground System to analyze troves of intelligence, which would have required a significant amount of human hours to complete. The new AI system cannot order a strike on its own, but it is now automatically identifying possible targets.

AUTOMATED TARGET RECOGNITION

Lethal autonomous weapons systems, powered by AI, are capable of finding targets autonomously and making decisions to complete a mission. In 2022, a lieutenant colonel in the Ukrainian military said that he and a group called Aerorozvidka had developed special drones that make use of automated target recognition. While it's unclear whether Aerorozvidka actually carried out test missions, the fact remains that machine learning–based vision for automated target recognition already exists. In response, 70 nations delivered a joint statement at the UN General Assembly calling for a ban on autonomous weapons—but little progress has been made in the months since.

DEFENSE

AUTOMATING OFFENSIVE ATTACKS USING AI

Thanks to advancements in AI, one of the big trends in security is automated hacking—in short, software that's built to out-hack human hackers. DARPA launched a Cyber Grand Challenge project in 2016, with a mission to design computer systems capable of beating hackers at their own game. DARPA wanted to show that smarter automated systems can reduce the response time—and fix system flaws—to just a few seconds. Spotting and fixing critical vulnerabilities is a task that might take a human hacker several months or even years to complete, and yet the machine that won the Grand Challenge did it in just a fraction of that time.

The winner became the first nonhuman entity to earn the DEF CON's Black Badge, which is the hacking community's equivalent of an Oscar. Very soon, malicious actors will create autonomous systems capable of automatically learning new environments, exposing vulnerabilities and flaws, and then exploiting them for gain—or whatever the stated objective, which could simply be generalized mayhem.

ALGORITHMIC WARFIGHTING

Future wars could be fought entirely 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 US, China, Israel, France, Russia, the UK, and South Korea—are also developing weapons systems that include at least some autonomous functionality.

In 2020, the US Air Force successfully flew an AI copilot on a U-2 spy plane in California, marking the first time in the history of the DOD that an AI algorithm trained to execute specific in-flight tasks was deployed. It was the mission commander, with the call sign ARTUμ—though the flight was just practice.

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.

MANDATING ETHICS GUIDELINES FOR TECH CONTRACTORS

Project Maven was developed to enlist AI to analyze surveillance video. Initially, Google was the DOD's vendor, but when employees found out they'd been working on a military project, thousands protested. It wasn't the first time tech contractors had lost trust in the government.

As a result, the Defense Innovation Unit is enforcing "responsible artificial intelligence" guidelines that vendors must adopt when building AI systems, models, or applications for the DOD. The guidelines offer specific instructions that must be followed during planning, development, and deployment, which include provisions for risk assessment. This represents a longer-term trend: government agencies requiring transparency in AI projects.

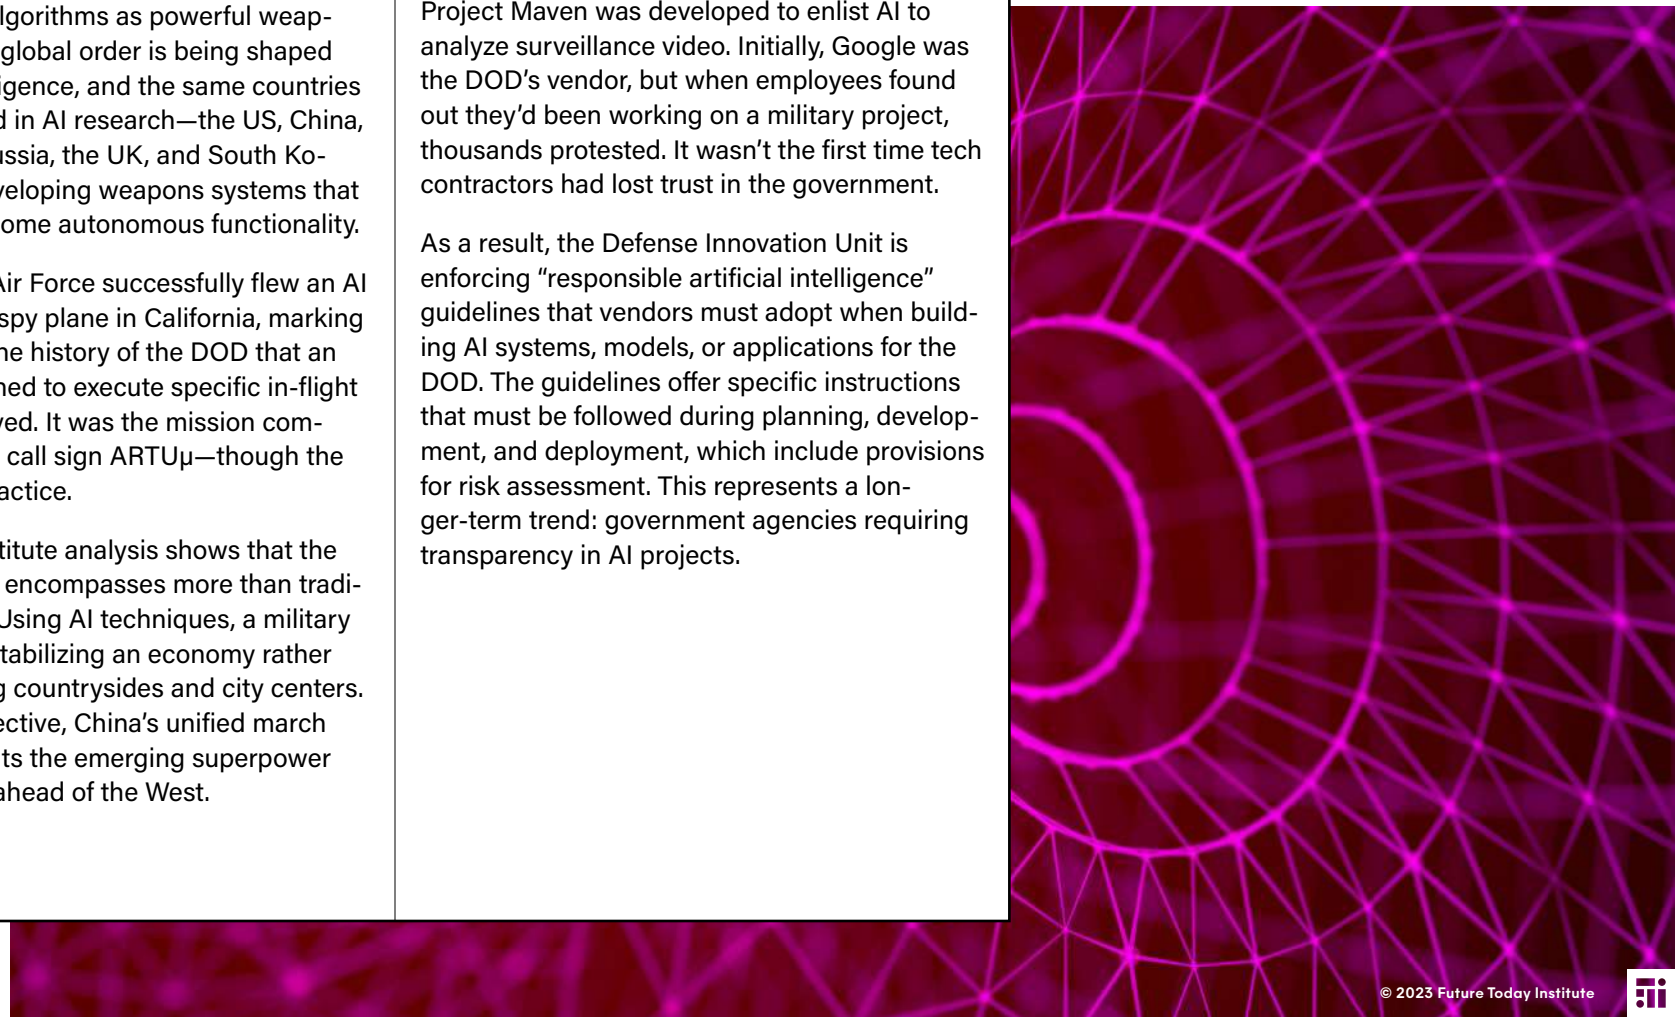

SPOTLIGHT: CHINA'S LONG-TERM AI AMBITIONS

TOP-TIER AI RESEARCHERS INCREASINGLY HAIL FROM CHINA

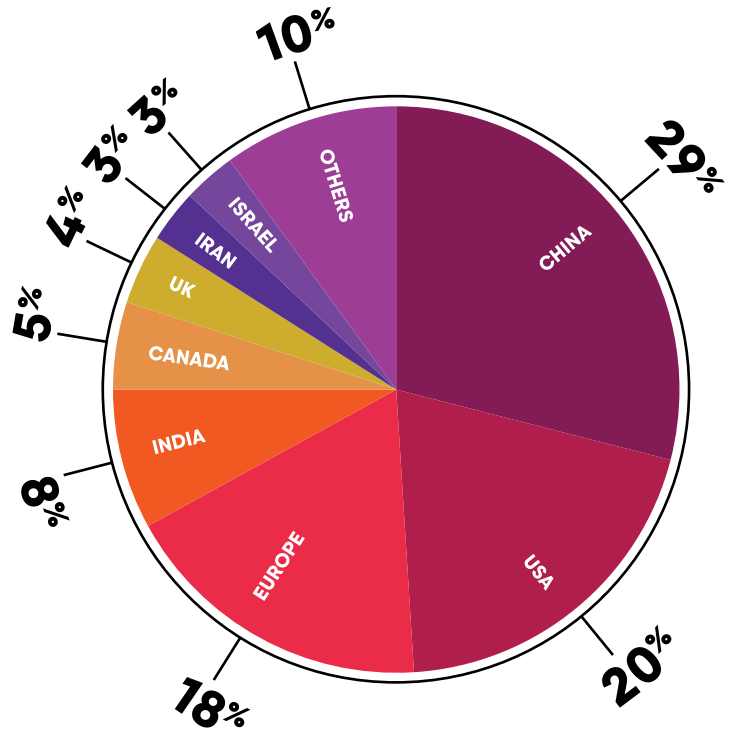

Country affiliations are based on the country where researchers received their undergraduate degree.

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

China is an undisputed global leader in AI. 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 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.

Within the next decade, China plans to meet two crucial milestones: By 2027, its People's Liberation Army will have a modern-ready force, and by 2030 the Chinese Communist Party (CCP) expects to have outpaced the US in AI and become the dominant force. China is producing what it calls "intelligentized" technologies to bolster both its economy and military.

Recently, China took major steps to shape the future of AI by releasing its own pretrained models, and it is forging ahead with its own natural language processing models, which makes sense since the most popular models in use now are trained on English text. Researchers from Tsinghua University and Alibaba are developing Chinese data sets and pretrained large transformers to compete against the likes of GPT-3. In 2021, two models developed specifically for the Chinese language went live: Wu Dao 2.0 and M6.

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 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.

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

SPOTLIGHT: CHINA'S LONG-TERM AI AMBITIONS

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 US startups. Baidu, a 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 a few years ago. It has since upped its stakes in companies including Tesla and Snap. The payoff for the Chinese is not just a monetary return on investment—Chinese companies expect intellectual property as well. China-based AI startups now account for nearly half of all AI investments globally.

EXPANDING MARKET

It's a challenging time for Chinese startups because of rising tensions with the West. Companies hoping to gain traction in Europe are making efforts to cloak their origin. Shein, the e-commerce website popular among teens, says it was "founded in L.A.," but the company actually got its start in Nanjing and Guangzhou by relying on the region's manufacturing centers and ample supply chains. Or look at TikTok, which has said it's a US-based company—while the app's parent Chinese company ByteDance has employed linguistic gymnastics to separate itself. Binance, the world's largest crypto exchange, which was created in China, says it doesn't have one physical location for its headquarters.

It's no wonder that as Chinese startups hope to expand globally, they're seeking to distance themselves from the authoritarian regime in Beijing. But that creates political hurdles, especially as the CCP seeks to bring its home-grown technology ecosystem into lockstep with party leaders.

The result could be a future parallel universe, in which Chinese-created AI systems are shaped both by enormous amounts of data and local laws. In Brazil, a generative AI system might write an unfettered political essay in Portuguese about a leader—while in China, that same essay would be automatically filtered for politically sensitive words and phrases. As the CCP enforces new regulations targeting AI and what the government calls "deep synthesis tech," the ways in which people experience and work alongside AI could be dramatically different.

BREAKING UP BIG TECH

Alibaba, Tencent, and Baidu, which have made important advancements in AI research, may find it difficult to keep innovating. Starting in 2020, the CCP initiated a wave of legislation aimed at its tech sector, introducing anti-monopoly legislation focused on the platform economy and promoting data security and privacy laws. The Personal Information Protection Law (PIPL), China's version of the EU's GDPR, went into effect November 2021. What followed were a series of crackdowns targeting some of China's most successful tech companies. Ultimately, this regulation wasn't about "breaking up" China's Big Tech—the CCP wanted to focus its tech sector on achieving research and development goals set by the government and military within the decade.

Increasingly, Beijing is pressuring its mega-successful big tech companies—Baidu, Alibaba, and Tencent, among others—to share data with the state and to perform research to support the vision of the CCP. Going forward, Beijing aims to direct the might of its tech companies at programs of national strategic importance rather than making video games. China's tech crackdown could cool private investment in Chinese companies, which could result in a chilling effect on innovation and economic growth, and also free up capital for emerging markets.

BUILDING A 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. China's social credit

SPOTLIGHT: CHINA'S LONG-TERM AI AMBITIONS

system, an algorithmic reputation system developed by the government, standardizes assessments of citizens' and businesses' behavior and activity.

China's PIPL, which took effect on Nov. 1, 2021, is just the latest in a national campaign to reassert government control over user data. The law targets private companies (government-controlled data will not be impacted) and even extends beyond Chinese borders in some cases. It was partly enacted in response to increased calls for consumer privacy amid the expanded collection and use of personal data for commercial purposes.

Companies outside of China will need to comply when processing data generated by Chinese individuals or collected in the process of delivering products and services in China. Any Chinese data sent overseas will require explicit permission. Penalties for violating the law are steep—up to 5% of the company's annual revenue the prior year—and are designed to discourage any attempts to bypass compliance. PIPL follows closely the passing of the Data Security Law in June 2021; both indicate that the CCP views preventing personal data from leaking outside of its national borders as vital to national security.

As for consumer privacy rights, new guidelines from China's Supreme People's Court went into effect in August of 2021, prohibiting hotels, retailers, and other businesses from using facial recognition without consent. Companies that already have this data will be required to delete face scans upon consumer request.

These laws, however, only protect Chinese citizens from commercial surveillance within Chinese national borders. In other countries, Chinese companies such as Huawei, ZTE, and Alibaba are installing surveillance systems as part of smart city initiatives. While these systems may offer cost-efficient surveillance, there is no guarantee that sensitive data won't be shared with the Chinese government or that the systems won't support a rise in authoritarianism.

China's Corporate Social Credit System (CSCS) is intended to be a standardized reputation-based system to create what the CCP calls a "fair, transparent, and predictable" business environment. This system applies to both local and foreign entities doing business in China.

To regulate corporate behavior, the CSCS will rely on data-gathering efforts that extend to businesses nationwide. Under this system, Chinese companies and trade associations will be required to provide data about foreign partners and enforce blacklists against targeted companies. Threats that both domestic and foreign businesses face from the effort include being subjected to arbitrary rule enforcement or new regulations regarding IP or tech transfer.

DEEPENING INTERNATIONAL TIES

We're living through a precarious moment. China is shaping the world order in its own image, while exporting its technologies and surveillance systems to other countries. 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.

As the CCP expands into African countries and throughout Southeast Asia and Latin America, it will also begin to eschew operating systems, technologies and infrastructure built by the West. Two Chinese companies—the state-con-

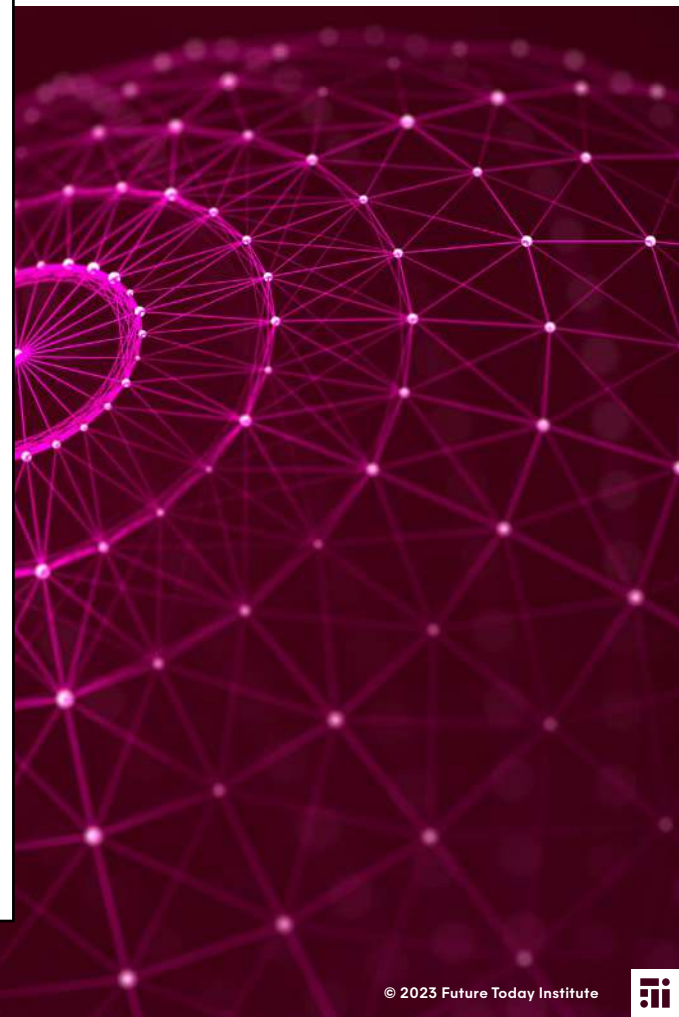

SPOTLIGHT: CHINA'S LONG-TERM AI AMBITIONS

Zhuang Rongwen, head of the Cyberspace Administration of China, speaks during the closing ceremony of the 5th World Internet Conference.

trolled CEIEC and Huawei—built Ecuador's surveillance system, called ECU-911. The system promised to curb high murder rates and drug crime, but Ecuador could not afford the 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 US 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.

When it comes to AI, leaders should monitor escalating tensions between the US and China. But they should also remember that there are cells of rogue actors who could cripple our economies simply by mucking with the power or traffic grids, causing traffic spikes on the internet, or locking us out of our connected home appliances. These aren't big, obvious signs of aggression, and that is a problem for many countries, including the US. Most governments don't have a paradigm describing a constellation of aggressive actions. Each action on its own might be insignificant. What are the escalation triggers? We don't have a definition, and that creates a strategic vulnerability.

> We've never seen a technology move as fast as AI has to impact society and technology. This is by far the fastest moving technology that we've ever tracked in terms of its impact and we're just getting started.

PAUL DAUGHERTY, CHIEF TECHNOLOGY OFFICER, ACCENTURE

SCENARIOS

WHAT IF DEEPFAKES ATTACK OUR PUBLIC MEMORY?

Scenario Year: 2033

Depending on where you live in the US, your student now learns different versions of what happened on the January 6th insurrection. While some school history textbooks contain unedited press photos of the event, others contain images generated by the latest AI tools that have been almost imperceptibly altered. Some teaching materials depict law enforcement officials violently attacking what appear to be innocent, peaceful protesters. Others show what looks like a parade, where revelers brandish flags instead of weapons. While the basic details—the weather, number of people in attendance, locations—remain the same across all the books, the visuals convey believable stories that vary greatly from historical fact, leading to a generation of Americans who disagree about what happened that fateful day at the US Capitol.

WHAT IF AI CAN'T TRUST ITS USERS?

Scenario Year: 2028

As of today, there are now more than 30 formal mathematical models that describe what is "fair" in AI. These were written by different groups and use different definitions. Harmonizing those models has proven difficult, and as disparate AI systems started to interact they faced conflicting information. After years of asking ourselves whether we can trust AI, we're now facing an unforeseen problem: AI systems that must decide whether they can trust us, and the data we report to them.

WHAT IF AI COULD HELP YOU PRACTICE MARRIAGE BEFORE GOING THROUGH WITH THE CEREMONY?

Scenario Year: 2030

The rich and famous aren't just getting prenups; they're now springing for prenup sims. Developed to mirror each fiancé's lifestyle habits, communication styles, personality quirks, and love languages, sims live together in a digital twin of the potential couple's home. Thousands of simulated routine and novel challenges are run to determine the probability of a long, happy marriage.

WHAT IF AI COULD HELP YOU EAT BETTER AND BE HEALTHIER?

Scenario Year: 2026

A new smartphone app uses deepfakes to encourage healthy lifestyles. The app, IntelliEAT (the EAT is short for Energy Assessment Technology) automatically recognizes everything you eat or drink, calculates the data from your smartwatch or fitness tracker, computes the benefit or damage, and automatically generates a deepfake to show the impact you made on your body that day. IntelliEAT doesn't just count calories or nutrients. Instead it displays what the future you would look life if every day was like your current day. Healthy meals and exercise reveal an older version of you who is fit, agile, and happy. Regular binge eating reveals a deepfaked version of you that is heavier with joint pain and difficulty sleeping. Proponents of body positivity have voiced concern, but IntelliEAT has won praise from medical associations for focusing holistically on health. In addition to the Immediate Impact View, the app also offers an Invisible Damage View, which highlights health effects such as increased cholesterol, restricted blood flow, and inflammation.

HOW TO PREPARE FOR THE FUTURE

HOW TO PREPARE

What should your organization do now to prepare for these trends?

Artificial intelligence should be part of every strategic plan, as it crosses multiple dimensions, from workforce automation, to digital transformation, to everyday business processes and business intelligence. 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.

As no-code and low-code applications become more widely available, Future Today Institute analysis shows that innovation teams will be in position to build powerful systems for decision management, business intelligence, and product ideation. Generative AIs will improve an organization's efficiency and enhance creativity, leading to hybrid human-machine creative teams. AI-assisted design will dramatically increase the number of prototypes that can be automatically generated with prompts.

We recommend that chief strategy officers in every field develop a solid understanding of AI to engage more closely with others in the C-suite to develop a cohesive point of view on digitalization, augmentation and automation—as well as to develop strategic plans. Importantly, businesses should keep abreast of emerging regulations that could restrict the use of consumer data, algorithms, or generative AI systems. Risk models should be developed to determine plausible near-futures, so that leaders can adjust their strategies accordingly.

When it comes to talent and workforce development, the need for highly skilled people is fast outpacing the graduation rate of universities. This might lead to a shift in higher education: It's plausible that companies in search of AI specialists might opt for modular certifications, which can be earned faster than traditional four-year degrees. Talent sourcing and retention will continue to pose challenges for tech companies—and for organizations in other industries that need a trained workforce but may not be able to provide the same perks as the hottest AI players.

In nearly every industry, AI will serve as a force multiplier for growth, bringing efficiencies, better tracking, business intelligence, and assistance with decision-making. As training costs decline, more applications will be built. Spending on AI systems and hardware is likely to explode this decade, creating significant enterprise value overall.

KEY QUESTIONS FOR YOUR TEAM

1

How do these trends change our perspective on our current AI investment and capabilities?

2

Is our organization positioned to leverage AI for growth, in addition to realizing new efficiencies?

3

When we engage in long-term strategic planning, how does the evolution of AI factor in?

4

Are we adequately preparing our workforce to succeed in a world in which their knowledge tasks might be augmented or fully automated by AI?

5

What is our process to vet, verify and monitor our vendors?

Is the AI we're using explainable?

If not, does that open us up to additional risk?

6

Does our team have an expansive enough viewpoint on emerging threats and attack surfaces?

7

What parts of our business make us vulnerable as AI evolves?

8

What is our point of view on the use of generative AI systems?

9

If proposed antitrust legislation passes in the US and the EU further regulates Big Tech companies, how does this create a strategic opportunity for our organization?

Or, what might our organization lose?

10

How might we develop the knowledge, experience, and talent in place to leverage these AI trends?

SELECTED SOURCES

SELECTED SOURCES

"A Generalist Agent." https://www.deepmind.com/publications/a-generalist-agent.

"Adept: Useful General Intelligence." https://www.adept.ai.

Air Force. "AI Copilot: Air Force Achieves First Military Flight With Artificial Intelligence." https://www.af.mil/News/Article-Display/Article/2448376/ai-copilot-air-force-achieves-first-military-flight-with-artificial-intelligence.

"AI Patent Database Content." 2020. Reprint, Center for Security and Emerging Technology (CSET), https://github.com/georgetown-cset/1790-ai-patent-data.

"AI-Synthesized Faces Are Indistinguishable from Real Faces and More Trustworthy | PNAS." https://www.pnas.org/content/119/8/e2120481119.

"Aircrew Labor In-Cockpit Automation System." https://www.darpa.mil/program/aircrew-labor-in-cockpit-automation-system.

"API Reference—Amazon Rekognition." https://docs.aws.amazon.com/rekognition/latest/dg/API_Reference.html.

Baek, Minkyung, Frank DiMaio, Ivan Anishchenko, Justas Dauparas, et al. "Accurate Prediction of Protein Structures and Interactions Using a Three-Track Neural Network." Science 373, no. 6557 (August 20, 2021): 871–76. https://doi.org/10.1126/science.abj8754.

Baltescu, Paul, Haoyu Chen, Nikil Pancha, Andrew Zhai, et al. "ItemSage: Learning Product Embeddings for Shopping Recommendations at Pinterest." arXiv, May 23, 2022. http://arxiv.org/abs/2205.11728.

"Biological Structure and Function Emerge From Scaling Unsupervised Learning to 250 Million Protein Sequences." PNAS. https://www.pnas.org/doi/10.1073/pnas.2016239118.

Borsos, Zalán, Matt Sharifi, and Marco Tagliasacchi. "SpeechPainter: Text-Conditioned Speech Inpainting." arXiv, March 30, 2022. http://arxiv.org/abs/2202.07273.

Cancian, Mark F., Matthew Cancian, and Eric Heginbotham. "The First Battle of the Next War: Wargaming a Chinese Invasion of Taiwan," January 9, 2023. https://www.csis.org/analysis/first-battle-next-war-wargaming-chinese-invasion-taiwan.

Center for Security and Emerging Technology. "Ethical Norms for New Generation Artificial Intelligence Released." https://cset.georgetown.edu/publication/ethical-norms-for-new-generation-artificial-intelligence-released.

Center for Security and Emerging Technology. "US AI Workforce." https://cset.georgetown.edu/publication/u-s-ai-workforce.

"Chexia Face Recognition." NIST. https://www.nist.gov/programs-projects/chexia-face-recognition.

Chitty-Venkata, Krishna Teja, and Arun K. Somani. "Neural Architecture Search Survey: A Hardware Perspective." ACM Computing Surveys 55, no. 4 (April 30, 2023): 1–36. https://doi.org/10.1145/3524500.

Corso, Gabriele, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. "DiffDock: Diffusion Steps, Twists, and Turns for Molecular Docking." arXiv, October 4, 2022. https://doi.org/10.48550/arXiv.2210.01776.

Damioli, Giacomo, Vincent Van Roy, and Daniel Vertesy. "The Impact of Artificial Intelligence on Labor Productivity." Eurasian Business Review 11, no. 1 (March 1, 2021): 1–25. https://doi.org/10.1007/s40821-020-00172-8.

Daulton, Samuel, David Eriksson, Maximilian Balandat, and Eytan Bakshy. "Multi-Objective Bayesian Optimization over High-Dimensional Search Spaces." arXiv, June 15, 2022. http://arxiv.org/abs/2109.10964.

Dauparas, J., I. Anishchenko, N. Bennett, H. Bai, R. J. Ragotte, L. F. Milles, B. I. M. Wicky, et al. "Robust Deep Learning–Based Protein Sequence Design Using ProteinMPNN." Science 378, no. 6615 (October 7, 2022): 49–56. https://doi.org/10.1126/science.add2187.

"Deep Dream Generator." https://deepdreamgenerator.com.

DeepMind. "Mastering Atari, Go, Chess and Shogi by Planning With a Learned Model." https://deepmind.com/research/publications/2019/Mastering-Atari-Go-Chess-and-Shogi-by-Planning-with-a-Learned-Model.

DeepMind. "MuZero: Mastering Go, chess, shogi and Atari Without rules." https://deepmind.com/blog/article/muzero-mastering-go-chess-shogi-and-atari-without-rules.

DeepMind. "Putting the power of AlphaFold Into the World's Jands." https://deepmind.com/blog/article/putting-the-power-of-alphafold-into-the-worlds-hands.

Defense One. "AI Could Match 'Fingerprints' of Texts to Their Authors, Under New Intelligence Program." https://www.defenseone.com/technology/2022/02/ai-could-match-fingerprints-texts-their-authors-under-new-intelligence-program/361893.

Defense One. "The Near Future of Military Autonomy Isn't Robotanks, But 'Microservices.'" https://www.defenseone.com/technology/2022/02/near-future-military-autonomy-isnt-robotanks-microservices/361873.

Degrave, Jonas, Federico Felici, Jonas Buchli, Michael Neunert, et al. "Magnetic Control of Tokamak Plasmas through Deep Reinforcement Learning." Nature 602, no. 7897 (February 17, 2022): 414–19. https://doi.org/10.1038/s41586-021-04301-9.

"DOD to Hire First-Ever Chief Digital and Artificial Intelligence Officer, Form New Office—Nextgov." https://www.nextgov.com/emerging-tech/2021/12/dod-hire-first-ever-chief-digital-and-artificial-intelligence-officer-form-new-office/187399.

SELECTED SOURCES

Dong, Tian, Bo Zhao, and Lingjuan Lyu. "Privacy for Free: How Does Dataset Condensation Help Privacy?" arXiv, June 1, 2022. http://arxiv.org/abs/2206.00240.

Dunn, Jeff. "Google Works Out a Fascinating, Slightly Scary Way for AI to Isolate Voices in a Crowd." Ars Technica, April 13, 2018. https://arstechnica.com/gadgets/2018/04/google-works-out-a-fascinating-slightly-scary-way-for-ai-to-isolate-voices-in-a-crowd.

Dwivedi-Yu, Jane, Yi-Chia Wang, Lijing Qin, Cristian Canton-Ferrer, and Alon Y. Halevy. "Affective Signals in a Social Media Recommender System." In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2831–41, 2022. https://doi.org/10.1145/3534678.3539054.

"Emotion—Amazon Rekognition." https://docs.aws.amazon.com/rekognition/latest/APIReference/API_Emotion.html.

"ESM." https://huggingface.co/docs/transformers/model_doc/esm.

Farr, Christina. "As Amazon Moves into Health Care, Here's What We Know—and What We Suspect—About Its Plans." CNBC. https://www.cnbc.com/2018/03/27/amazons-moves-into-health-what-we-know.html.

Fedus, William, Barret Zoph, and Noam Shazeer. "Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity." ArXiv:2101.03961 [Cs], January 11, 2021. http://arxiv.org/abs/2101.03961.

Flennerhag, Sebastian, Yannick Schroecker, Tom Zahavy, and Hado van Hasselt. "Bootstrapped Meta-Learning," 2022.

Gafni, Oran, Adam Polyak, Oron Ashual, Shelly Sheynin, et al. "Make-A-Scene: Scene-Based Text-to-Image Generation with Human Priors." arXiv, March 24, 2022. http://arxiv.org/abs/2203.13131.

Gent, Edd. "Meta Built an AI That Can Guess the Words You're Hearing by Decoding Your Brainwaves." Singularity Hub (blog), September 13, 2022. https://singularityhub.com/2022/09/13/meta-built-an-ai-that-can-guess-the-words-youre-hearing-by-decoding-your-brainwaves.

"GitHub—Google-Research/Dreamer: Dream to Control: Learning Behaviors by Latent Imagination." https://github.com/google-research/dreamer.

Google. "Join Us in the AI Test Kitchen," August 25, 2022. https://blog.google/technology/ai/join-us-in-the-ai-test-kitchen.

Google. "LaMDA: Our Breakthrough Conversation Technology," May 18, 2021. https://blog.google/technology/ai/lamda.

Google. "MUM: A New AI Milestone for Understanding Information," May 18, 2021. https://blog.google/products/search/introducing-mum.

Gupta, Agrim, Silvio Savarese, Surya Ganguli, and Li Fei-Fei. "Embodied Intelligence via Learning and Evolution." Nature Communications 12, no. 1 (October 6, 2021): 5721. https://doi.org/10.1038/s41467-021-25874-z.

Hafner, Danijar. "Dream to Control." Python, February 13, 2022. https://github.com/danijar/dreamer.

Horowitz, Adam Haar, Tony J. Cunningham, Pattie Maes, and Robert Stickgold. "Dormio: A Targeted Dream Incubation Device." Consciousness and Cognition 83 (August 2020): 102938. https://doi.org/10.1016/j.concog.2020.102938.

"How Amazon Search Reduced ML Inference Costs by 85% with AWS Inferentia." AWS Machine Learning Blog,"September 22, 2022. https://aws.amazon.com/blogs/machine-learning/how-amazon-search-reduced-ml-inference-costs-by-85-with-aws-inferentia.

Intenseye. "Intenseye: AI for Workplace Safety." https://content.intenseye.com.

Jumper, John, Richard Evans, Alexander Pritzel, Tim Green, et al. "Highly Accurate Protein Structure Prediction with AlphaFold." Nature 596, no. 7873 (August 2021): 583–89. https://doi.org/10.1038/s41586-021-03819-2.

Kafritsas, Nikos. "CLIP: The Most Influential AI Model From OpenAI—And How To Use It." Medium, September 28, 2022. https://towardsdatascience.com/clip-the-most-influential-ai-model-from-openai-and-how-to-use-it-f8ee408958b1.

Kuang, Zhengfei, Kyle Olszewski, Menglei Chai, Zeng Huang, et al. "NeROIC: Neural Rendering of Objects From Online Image Collections." arXiv, September 1, 2022. http://arxiv.org/abs/2201.02533.

LeCun, Yann. "A Path Towards Autonomous Machine Intelligence Version 0.9.2, 2022-06-27."

Lee, Jae Hyeon, Payman Yadollahpour, Andrew Watkins, Nathan C. Frey, et al. "EquiFold: Protein Structure Prediction With a Novel Coarse-Grained Structure Representation." bioRxiv, October 8, 2022. https://doi.org/10.1101/2022.10.07.511322.

Li, Yujia, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, et al. "Competition-Level Code Generation with AlphaCode." arXiv, February 8, 2022. http://arxiv.org/abs/2203.07814.

Louis Bouchard. "Impressive Photo Restoration by AI !" March 11, 2022. https://www.louisbouchard.ai/gfp-gan.

"Meta AI Research Topic—No Language Left Behind." https://ai.facebook.com/i18n/research/topics/no-language-left-behind.

Mildenhall, Ben, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, et al. "NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis." ArXiv:2003.08934 [Cs], August 3, 2020. http://arxiv.org/abs/2003.08934.

SELECTED SOURCES

MIT Technology Review. "A Bot That Watched 70,000 Hours of Minecraft Could Unlock AI's next Big Thing." https://www.technologyreview.com/2022/11/25/1063707/ai-minecraft-video-unlock-next-big-thing-openai-imitation-learning.

Müller, Thomas, Alex Evans, Christoph Schied, and Alexander Keller. "Instant Neural Graphics Primitives with a Multiresolution Hash Encoding." ACM Transactions on Graphics 41, no. 4 (July 2022): 1–15. https://doi.org/10.1145/3528223.3530127.

"Multiple Encounter Dataset (MEDS-I)—NIST Special Database 32." National Institute of Standards and Technology. https://catalog.data.gov/dataset/multiple-encounter-dataset-meds-i-nist-special-database-32-616de.

"Nerfies: Deformable Neural Radiance Fields." https://nerfies.github.io.

Nextgov.com. "Senate Confirms John Sherman as Pentagon's Tech Chief." https://www.nextgov.com/cio-briefing/2021/12/senate-confirms-john-sherman-pentagons-tech-chief/359940.

Ni, Bolin, Houwen Peng, Minghao Chen, Songyang Zhang, et al. "Expanding Language-Image Pretrained Models for General Video Recognition." ArXiv, August 4, 2022. https://doi.org/10.48550/arXiv.2208.02816.

Nightingale, Sophie J., and Hany Farid. "AI-Synthesized Faces Are Indistinguishable from Real Faces and More Trustworthy." Proceedings of the National Academy of Sciences 119, no. 8 (February 22, 2022). https://doi.org/10.1073/pnas.2120481119.

Nijkamp, Erik, Jeffrey Ruffolo, Eli N. Weinstein, Nikhil Naik, and Ali Madani. "ProGen2: Exploring the Boundaries of Protein Language Models." arXiv, June 27, 2022. https://doi.org/10.48550/arXiv.2206.13517.

Nitzan, Yotam, Kfir Aberman, Qiurui He, Orly Liba, et al. "MyStyle: A Personalized Generative Prior." arXiv, October 6, 2022. https://doi.org/10.48550/arXiv.2203.17272.

"Obvious—Art & AI." https://obvious-art.com.

Office, U. S. Government Accountability. "Facial Recognition Technology: Federal Law Enforcement Agencies Should Better Assess Privacy and Other Risks." https://www.gao.gov/products/gao-21-518.

OpenAI. "CLIP: Connecting Text and Images," January 5, 2021. https://openai.com/blog/clip.

OpenAI. "DALL·E: Creating Images from Text," January 5, 2021. https://openai.com/blog/dall-e.

OpenAI. "MuseNet," April 25, 2019. https://openai.com/blog/musenet.

Peng, Kenny, Arunesh Mathur, and Arvind Narayanan. "Mitigating Dataset Harms Requires Stewardship: Lessons from 1000 Papers." ArXiv:2108.02922 [Cs], November 21, 2021. http://arxiv.org/abs/2108.02922.

Piergiovanni, Aj, Vincent Casser, Michael S. Ryoo, and Anelia Angelova. "4D-Net for Learned Multi-Modal Alignment." In 2021 IEEE/CVF International Conference on Computer Vision (ICCV), 15415–25. Montreal, QC, Canada: IEEE, 2021. https://doi.org/10.1109/ICCV48922.2021.01515.

"ProT-VAE: Protein Transformer Variational AutoEncoder for Functional Protein Design," BioRxiv. https://www.biorxiv.org/content/10.1101/2023.01.23.525232v1.

Radford, Alec, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. "Robust Speech Recognition via Large-Scale Weak Supervision."

Rainie, Lee, Janna Anderson, and Emily a Vogels. "Experts Doubt Ethical AI Design Will Be Broadly Adopted as the Norm Within the Next Decade." Pew Research Center: Internet, Science & Tech (blog), June 16, 2021. https://www.pewresearch.org/internet/2021/06/16/experts-doubt-ethical-ai-design-will-be-broadly-adopted-as-the-norm-within-the-next-decade.

Ramesh, Aditya, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. "Hierarchical Text-Conditional Image Generation with CLIP Latents." arXiv, April 12, 2022. http://arxiv.org/abs/2204.06125.

"README." Python. 2022. Reprint, NVIDIA Corporation, January 21, 2023. https://github.com/NVIDIA/MegaMolBART.

"Responsible AI Guidelines." https://www.diu.mil/responsible-ai-guidelines.

Roper, Dr Will. "Exclusive: AI Just Controlled a Military Plane for the First Time Ever." Popular Mechanics, December 16, 2020. https://www.popularmechanics.com/military/aviation/a34978872/artificial-intelligence-controls-u2-spy-plane-air-force-exclusive/.

Saharia, Chitwan, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, et al. "Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding."

Sapé, Sergi Castella i. "Trends in AI—September 2022." Zeta Alpha, September 28, 2022. https://www.zeta-alpha.com/post/trends-in-ai-september-2022.

Schrittwieser, Julian, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, et al. "Mastering Atari, Go, Chess and Shogi by Planning with a Learned Model." Nature 588, no. 7839 (December 2020): 604–9. https://doi.org/10.1038/s41586-020-03051-4.

Sevilla, Jaime, Lennart Heim, Anson Ho, Tamay Besiroglu, Marius Hobbhahn, and Pablo Villalobos. "Compute Trends Across Three Eras of Machine Learning." arXiv, March 9, 2022. https://doi.org/10.48550/arXiv.2202.05924.

SELECTED SOURCES

Shwartz-Ziv, Ravid, and Amitai Armon. "Tabular Data: Deep Learning Is Not All You Need." arXiv, November 23, 2021. http://arxiv.org/abs/2106.03253.

Silver, David, Satinder Singh, Doina Precup, and Richard S. Sutton. "Reward Is Enough." Artificial Intelligence 299 (October 1, 2021): 103535. https://doi.org/10.1016/j.artint.2021.103535.

Singer, Uriel, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, et al. "Make-a-Video: Text-to-Video Generation Without Text-Video Data."

Singhal, Karan, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, et al. "Large Language Models Encode Clinical Knowledge." arXiv, December 26, 2022. https://doi.org/10.48550/arXiv.2212.13138.

Solaiyappan, Siddharth, and Yuxin Wen. "Machine Learning Based Medical Image Deepfake Detection: A Comparative Study." Machine Learning with Applications 8 (June 15, 2022): 100298. https://doi.org/10.1016/j.mlwa.2022.100298.

"The AI Index Report—Artificial Intelligence Index." https://aiindex.stanford.edu/report.

Thoppilan, Romal, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, et al. "LaMDA: Language Models for Dialog Applications." arXiv, February 10, 2022. http://arxiv.org/abs/2201.08239.

Tuli, Shreshth, Giuliano Casale, and Nicholas R. Jennings. "TranAD: Deep Transformer Networks for Anomaly Detection in Multivariate Time Series Data." arXiv, May 14, 2022. http://arxiv.org/abs/2201.07284.

Tzaban, Rotem, Ron Mokady, Rinon Gal, Amit H. Bermano, and Daniel Cohen-Or. "Stitch It in Time: GAN-Based Facial Editing of Real Videos." arXiv, January 21, 2022. https://doi.org/10.48550/arXiv.2201.08361.

"Voxel." https://www.voxelai.com.

"Welcome—Amazon Rekognition." https://docs.aws.amazon.com/rekognition/latest/APIReference/Welcome.html.

Willingham, Emily. "Humans Find AI-Generated Faces More Trustworthy Than the Real Thing." Scientific American. https://www.scientificamerican.com/article/humans-find-ai-generated-faces-more-trustworthy-than-the-real-thing.

Xiao, Tete, Xiaolong Wang, Alexei A. Efros, and Trevor Darrell. "What Should Not Be Contrastive in Contrastive Learning." arXiv, March 18, 2021. http://arxiv.org/abs/2008.05659.

Xie, Yunhe, Chengjie Sun, and Zhenzhou Ji. "A Commonsense Knowledge Enhanced Network with Retrospective Loss for Emotion Recognition in Spoken Dialog." In ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 7027–31, 2022. https://doi.org/10.1109/ICASSP43922.2022.9746909.

Yang, Gengshan, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, and Hanbyul Joo. "BANMo: Building Animatable 3D Neural Models from Many Casual Videos." In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2853–63. New Orleans, LA, USA: IEEE, 2022. https://doi.org/10.1109/CVPR52688.2022.00288.

Zhang, Can, Tianyu Yang, Junwu Weng, Meng Cao, Jue Wang, and Yuexian Zou. "Unsupervised Pre-Training for Temporal Action Localization Tasks." arXiv, March 25, 2022. http://arxiv.org/abs/2203.13609.

Zhang, Susan, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, et al. "OPT: Open Pre-Trained Transformer Language Models." ArXiv, June 21, 2022. http://arxiv.org/abs/2205.01068.

Zvyagin, Maxim, Alexander Brace, Kyle Hippe, Yuntian Deng, Bin Zhang, Cindy Orozco Bohorquez, Austin Clyde, et al. "GenSLMs: Genome-Scale Language Models Reveal SARS-CoV-2 Evolutionary Dynamics." BioRxiv, October 11, 2022. https://doi.org/10.1101/2022.10.10.511571.

AUTHORS & CONTRIBUTORS

AMY WEBB

Founder and CEO of Future Today Institute

Amy Webb pioneered a data-driven foresight methodology that is now used within hundreds of organizations. She possesses master-level knowledge and insights on the potential futures for the industries critical to keeping the world moving during this unparalleled period of disruption in business and society. Webb is a sought-after advisor to Fortune 100 and Global 1000 companies and provides strategic foresight on the emerging tech and trends that will turn industries on their heads and transform every aspect of our lives. She was elected a life member of the Council on Foreign Relations, serves on a Stewardship Board at the World Economic Forum, and in 2022 received an official Mad Scientist proclamation from the U.S. Army. A lifelong science fiction fan, she collaborates closely with writers and producers on films, TV shows and commercials about science, technology and the future. Webb was named by Forbes as one of the five women changing the world, listed as the BBC's 100 Women of 2020, named one of the most influential business professors in the world by Poets & Quants, and is ranked on the Thinkers50 list of the 50 most influential management thinkers globally.

ANNAMALAI CHOCKALINGAM

Senior Expert Advisor

Annamalai Chockalingam is a Senior Expert Advisor at Future Today Institute, where he is helping shape a better future. Annamalai currently runs product management and marketing at Nvidia, for a suite of generative AI, and large language model software products. Annamalai brings a wealth of experience spanning deep learning, strategy and management consulting, data and firmware engineering, and technology infrastructure. He has experience working with large-cap organizations across technology, energy, and financial sectors. He excels at using his deep technical background to devise astute business strategies and deliver innovative products that delight its users. He is a recent graduate from NYU's Stern School of Business and Courant Institute of Mathematical Sciences, and also holds a Bachelor of Science in Electrical Engineering from the University of Alberta.

Chief Executive Officer
AMY WEBB

Managing Director
MELANIE SUBIN

Chief Content Officer
JON FINE

Creative Director
EMILY CAUFIELD

Designer
ERICA GRAU

Editors
CAROLE BRADEN
TOM BRADY
MEGAN CREYDT
TOM FOSTER

Copy Editor
SARAH JOHNSON

Director of Operations
CHERYL COONEY

FUTURE TODAY INSTITUTE'S 2023 TECH TREND REPORT

Our 2023 edition includes nearly 700 trends, which are published individually in 14 volumes and as one comprehensive report with all trends included.

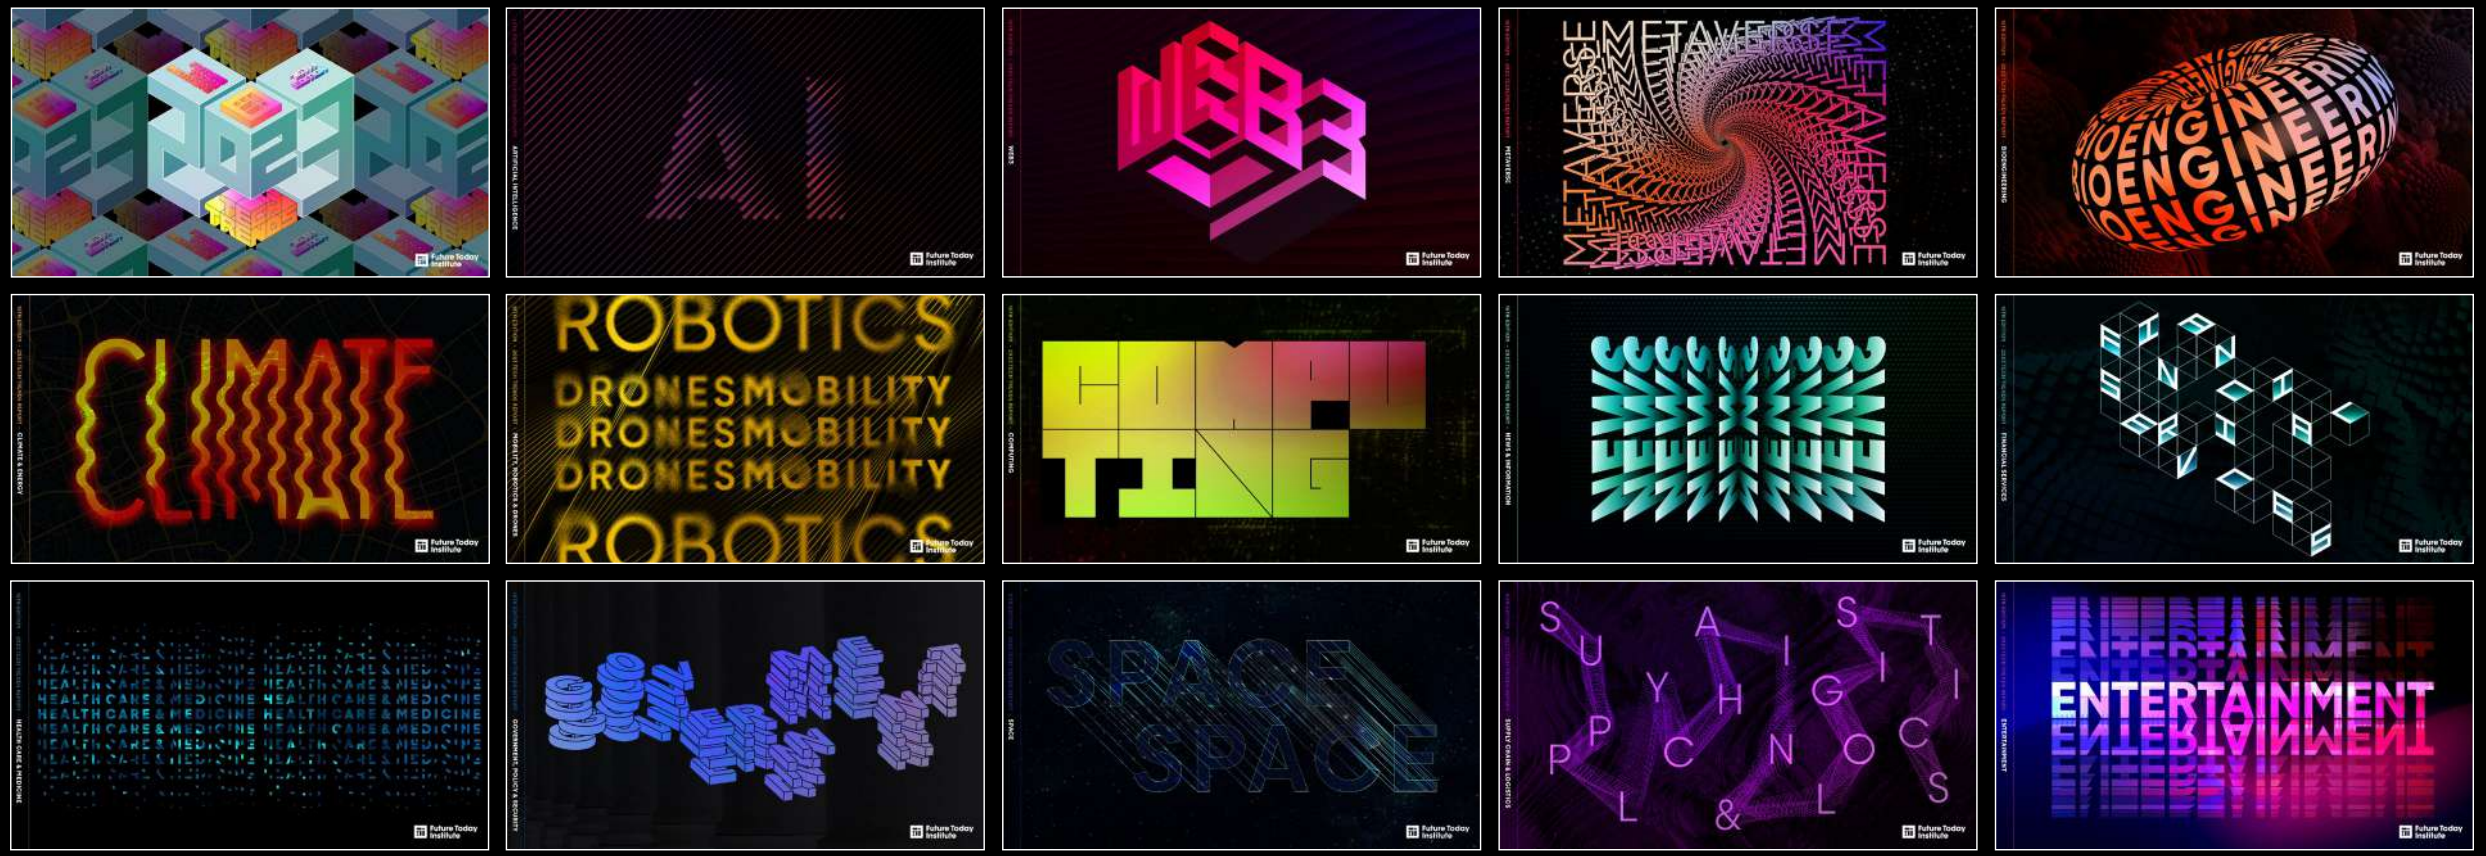

Download all sections of Future Today Institute's 2023 Tech Trends report at **http://www.futuretodayinstitute.com/trends.**

ABOUT FUTURE TODAY INSTITUTE

ABOUT FUTURE TODAY INSTITUTE

Founded in 2006, Future Today Institute helps leaders and their organizations navigate turbulent times by reducing uncertainty, identifying existential threats, leveraging critical emerging technologies, and identifying opportunities for growth.

We help organizations lead into the future. Our consultants work with executives and their teams to define plausible futures and develop executable strategies. Combining deep industry knowledge, rich creativity and rigorous analysis, we make teams confident in taking meaningful action.

As the leading strategic foresight and futures management consultants to executive leadership teams worldwide, FTI's data-driven, applied research methodology identifies trends, calculates their impact, and through descriptive scenarios, reveals 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

For an introductory conversation to learn how Future Today Institute can assist your organization with its long-term strategic planning and foresight needs, please contact

inquiries@futuretodayinstitute.com

futuretodayinstitute.com

+1 267 342 4300

SOCIAL

TWITTER	**@FTI**
LINKEDIN	**@Future-Today-Institute**
FACEBOOK	**@FutureTodayInstitute**
INSTAGRAM	**@futuretodayinstitute**

METHODOLOGY

Future Today Institute conducts in-depth qualitative and quantitative research throughout the year to identify emerging trends. We review patent and trademark filings, pre-print and published scientific papers, investment rounds, online search trends, macroeconomic data, publications from governments worldwide, news mentions, influencer posts and other sources, and we use a proprietary system to identify patterns, which are then grouped into nodes and evaluated using a set of standardized indicators. Qualified trends are further scored for their trajectory, momentum and timing. Additionally, we harness the deep subject matter expertise of our Future Today Institute network, leading to valuable insights about the topics we cover.

In continuous publication since 2007, Future Today Institute's annual report includes maturing and emerging trends grouped into two categories: industry and technology. Industry trends reflect the ways in which technology is shaping the future of an entire industry. Technology trends are specific developments within one arena, such as artificial intelligence. Covering a wide range of technologies across industry sectors creates a holistic view of change and provides leaders with a clear understanding of their potential impact. Trends are published as individual Industry and Technology reports, as well as in one combined report with all of our research.

Monitored regularly, trends help executives recognize emerging threats and opportunities in the near-term and enable them to develop perspectives, strategies and plans for the future.

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 Future Today Institute or this report's authors.

Future Today Institute's 2023 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 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

You are free to:

Download the materials for your or your organization's non-commercial use.

Use the assets with full attribution to Future Today Institute using the following language: "This material is copyrighted by Future Today Institute and licensed under the Creative Commons 4.0 International License (CC BY-NC-ND 4.0)."

Use assets in a reasonable manner, but not in any way that suggests that Future Today Institute endorses you or your use.

You may not:

Use the materials for any commercial purposes, which includes: selling, licensing or charging for products or services that use the materials.

Publish the materials in a book or upload to a website or intranet without prior written permission from Future Today Institute.

Remix, transform or build upon the material without giving appropriate credit to Future Today Institute as the originator of the work using the following language: "This material is adapted from original work copyrighted by Future Today Institute and licensed under the Creative Commons 4.0 International License (CC BY-NC-ND 4.0)."

Assert a relationship with or endorsement from Future Today Institute.

16TH EDITION · 2023 TECH TRENDS REPORT · ARTIFICIAL INTELLIGENCE

Future Today Institute