

INFO~TECH
RESEARCH GROUP

TECH TRENDS 2024

INFOTECH.COM

INTRODUCTION

WHAT MOORE'S LAW CAN TEACH US ABOUT OUR RELATIONSHIP WITH AI

When Gordon Moore made a prediction in a 1965 paper, he couldn't have fathomed that it would become the most famous law of computing. Moore's law originally stated that computer processing power would double every two years and that this trend would last for at least 10 years (Moore, 1965). Afterward, he was proven correct as integrated circuits became more efficient and less expensive at an exponential rate. The trend lasted beyond Moore's predicted 10-year span, holding true for decades.

In March 2023, Moore died. Whether Moore's law outlives him or not is a matter of debate. Some say we are nearing the physical limits of the number of transistors that can be packed into a silicon wafer. But whether or not the concept still applies to chips is not as important as the broader lesson learned about the feedback loop created between humanity and technology. It's one of exponential advancement, and it's why Moore's law is now commonly used to describe many different advances in computing beyond processing power.

Once made, Moore's prediction turned into a goal – one that chip designers strove for as their North Star of progress. Designers used high-performance computing (HPC) to augment their designs, solving mathematical and engineering problems required to more densely pack transistors together. Hence, the demand for HPC increased. This supported the design of devices that preserved Moore's law, leading to even more powerful HPCs, and so on (HPCWire, 2016). This feedback loop eventually produced today's nanometer-scale chips that power our smartphones and wearables.

A relationship that produces benefits for both parties involved can be described as symbiotic. Chip designers formed a symbiotic relationship with HPC to achieve their goal. The same can be said of developers' relationship with machine learning systems as the large language models powering generative AI have become more powerful over the past decades.

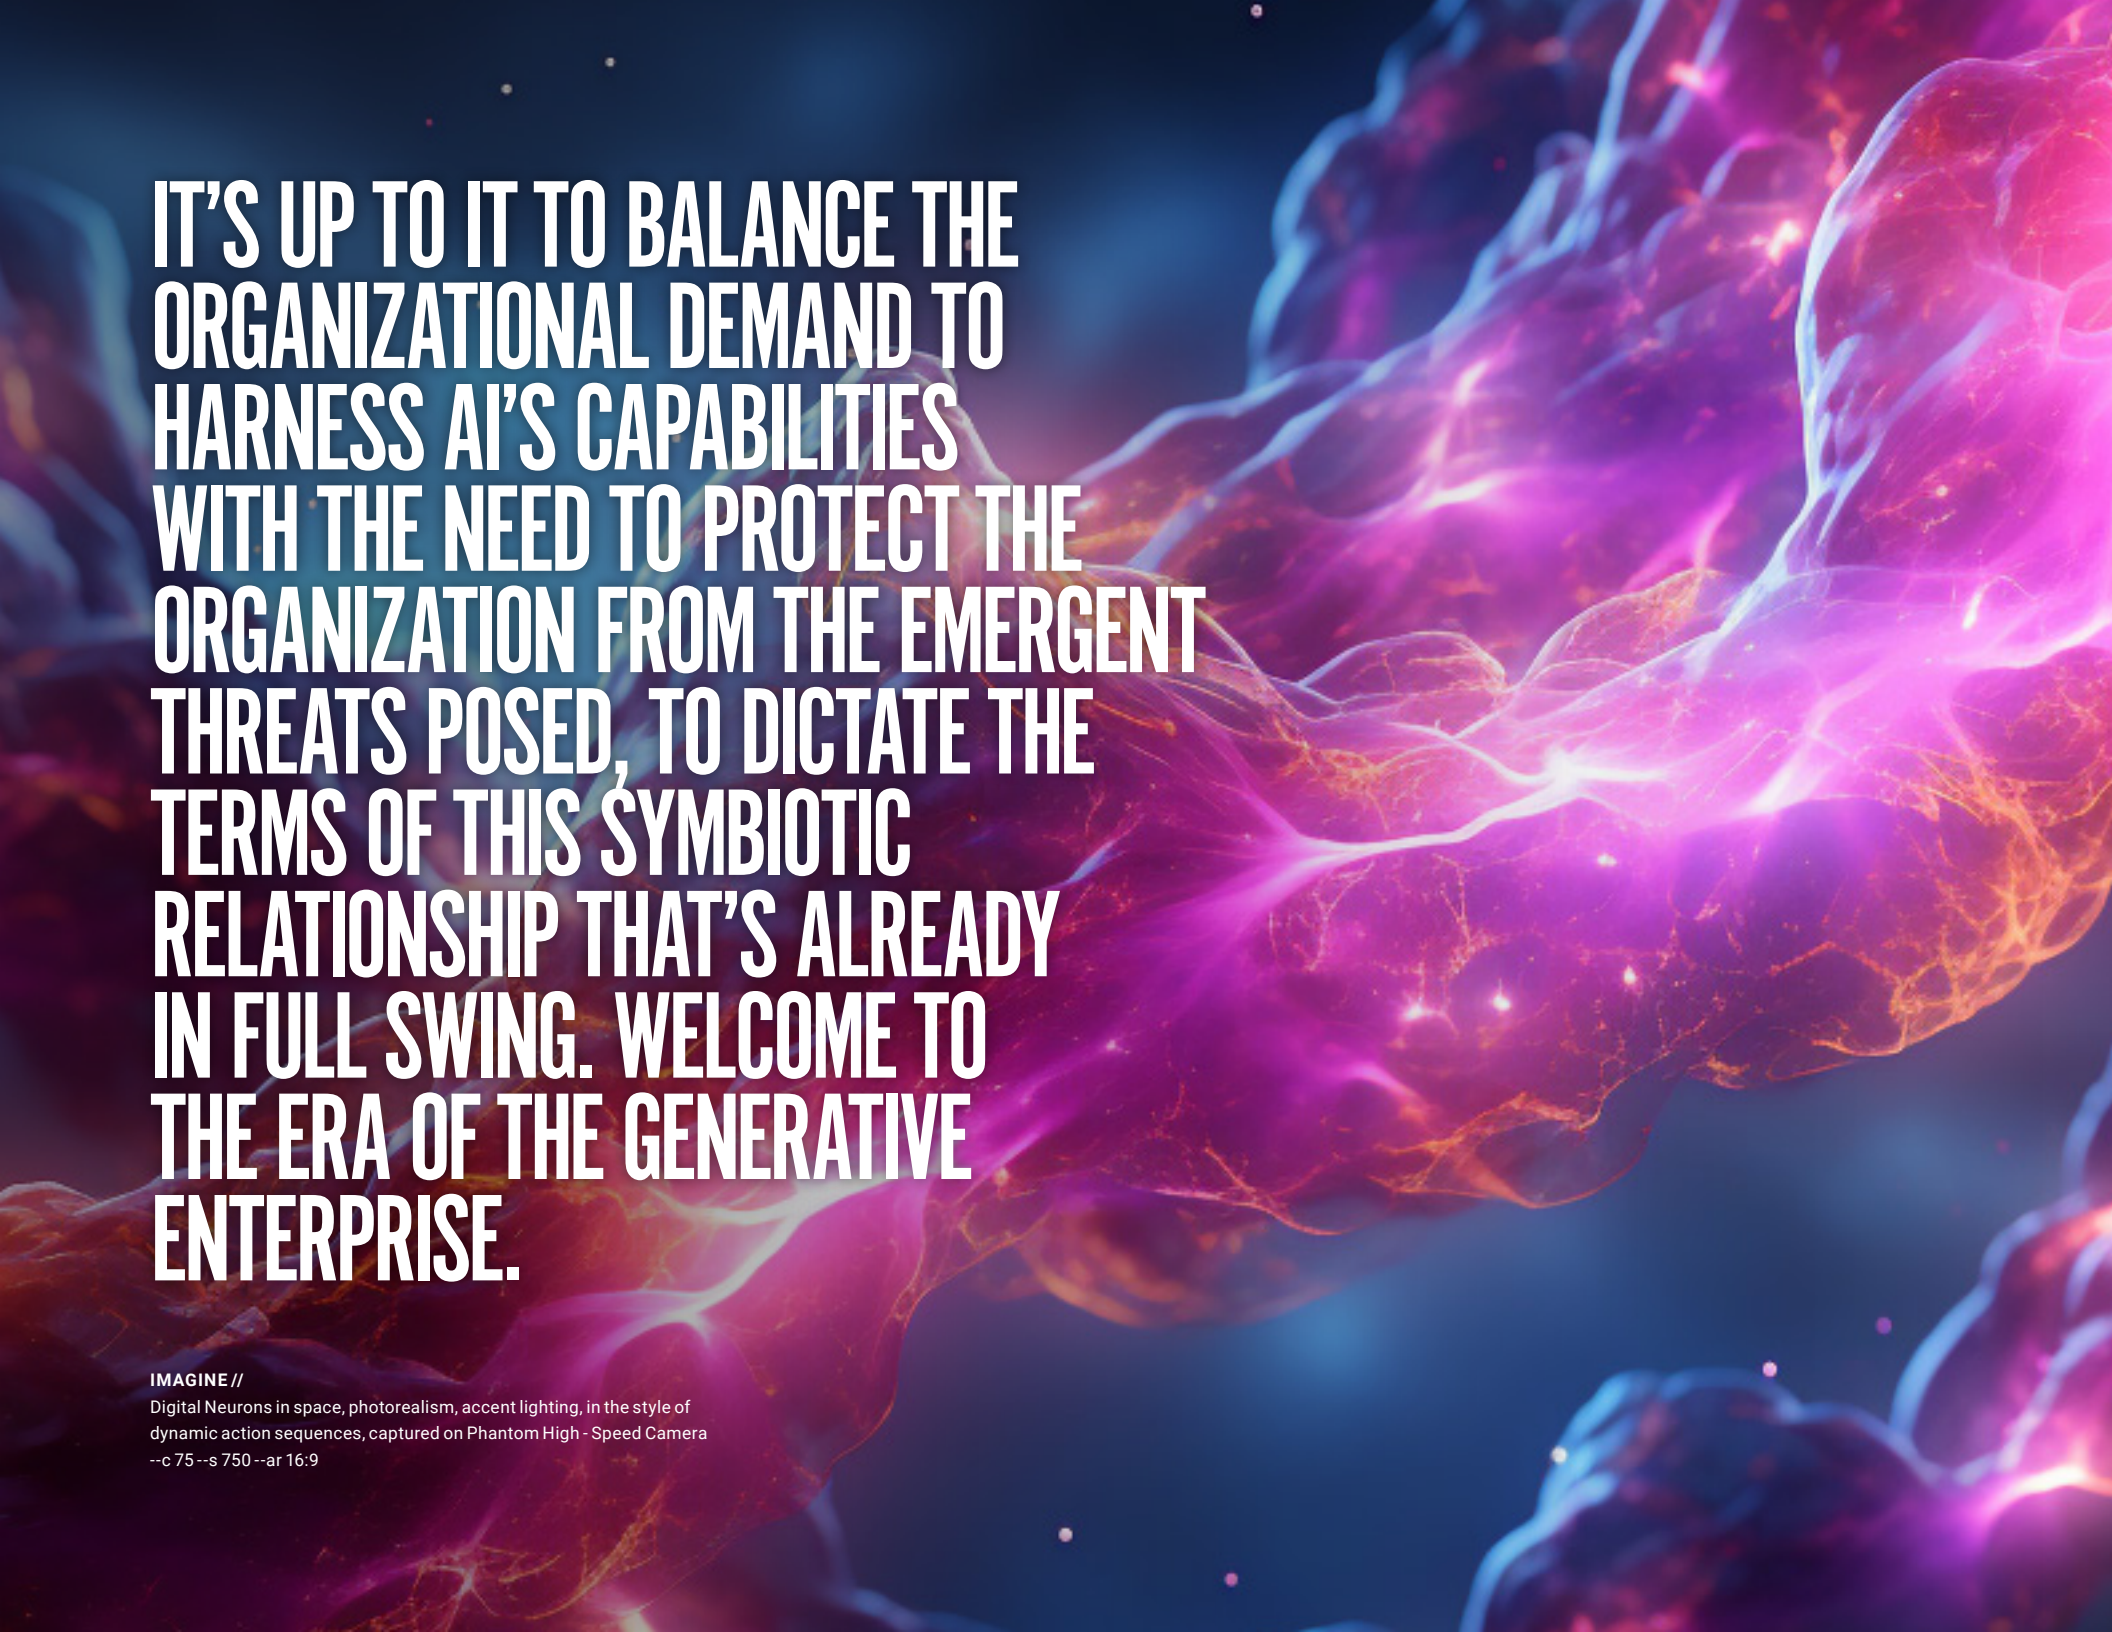
Like chips, generative AI systems have seen exponential growth. But this growth has been compounded into a briefer timeline, specifically over the past five years. It can be measured either in the computing power required to train the models or in the number of parameters contained by the models, an indicator of their complexity and flexibility. For example, in 2019, OpenAI's GPT-2 contained just 1.5 billion parameters. In 2022, Google's PaLM contained 540 billion parameters (Stanford University, 2023). Today, it's estimated that OpenAI's GPT-4 contains well over 1 trillion parameters.

Now some of those developing these large models say we should slow down, lest humanity's symbiotic relationship turn into a predatory one – with us playing the part of prey. This comes after said developers scraped the digital commonwealth of the web in the quest for more data to feed their growing algorithms, in pursuit of continued exponential growth. Many creators – from writers to illustrators to coders – are protesting that their consent wasn't sought to be a part of these data sets. Several lawsuits before the courts could determine how exactly copyright concepts and training an algorithm intersect.

For IT organizations, the exponential development of generative AI can't be ignored any longer. Just as Moore's law pushed demand for constantly higher-performing and always-miniaturizing computing power in the enterprise, IT must now work to enable new AI capabilities. As with the digital age, this will transform enterprises from their back-office operations to their very business models. Simultaneously, IT must prepare a new set of controls that mitigates the risks brought by AI. From securing the new systems to protecting against irresponsible use, IT departments will be asked to supply governance to an area that's attracting increased attention from regulators and courtrooms.

It's up to IT to balance the organizational demand to harness AI's capabilities with the need to protect the organization from the emergent threats posed, to dictate the terms of this symbiotic relationship that's already in full swing. Welcome to the era of The Generative Enterprise.

In our [Tech Trends 2023 report](#), we featured generative AI as one of our seven trends. We advised firms to experiment with generative AI tools and curate data sets for the purpose of training models. It's safe to say this trend had significantly more impact on the market than others, and this year we're focusing our report on exploring its many implications for IT.



IT'S UP TO IT TO BALANCE THE
ORGANIZATIONAL DEMAND TO
HARNESS AI'S CAPABILITIES
WITH THE NEED TO PROTECT THE
ORGANIZATION FROM THE EMERGENT
THREATS POSED, TO DICTATE THE
TERMS OF THIS SYMBIOTIC
RELATIONSHIP THAT'S ALREADY
IN FULL SWING. WELCOME TO
THE ERA OF THE GENERATIVE
ENTERPRISE.

IMAGINE //

Digital Neurons in space, photorealism, accent lighting, in the style of
dynamic action sequences, captured on Phantom High-Speed Camera

--c 75 --s 750 --ar 16:9

TECH TRENDS 2024: THE GENERATIVE ENTERPRISE

SEIZE OPPORTUNITIES

- ▶ AI-Driven Business Models
- ▶ Autonomized Back Office
- ▶ Spatial Computing

Throughout our report, we'll examine how organizations that have already invested in AI or plan to invest in AI are behaving compared to organizations that either do not plan to invest in AI or don't plan to invest until after 2024. We'll refer to these two groups as "AI adopters" and "AI skeptics" for simplicity. Here's a quick breakdown of what each of these groups look like:

AI adopters: Organizations that have already invested in AI or plan to do so by the end of 2024 ($n=430$).

- ▶ More likely to be from larger organizations, with 37% of respondents estimating a total headcount above 2,500.
- ▶ More likely to have a larger IT budget, with 48.5% reporting a budget of at least \$10 million.
- ▶ 62% located in North America.
- ▶ Represent a wide swath of industries including 20% in public sector and 11% in financial services.
- ▶ Most likely to rate IT maturity level at "Support" (38%) or "Optimize" (29%).

MITIGATE THREATS

- ▶ Responsible AI
- ▶ Security by Design
- ▶ Digital Sovereignty

AI skeptics: Organizations that either don't plan to invest in AI until after 2024 or don't plan to invest at all ($n=176$).

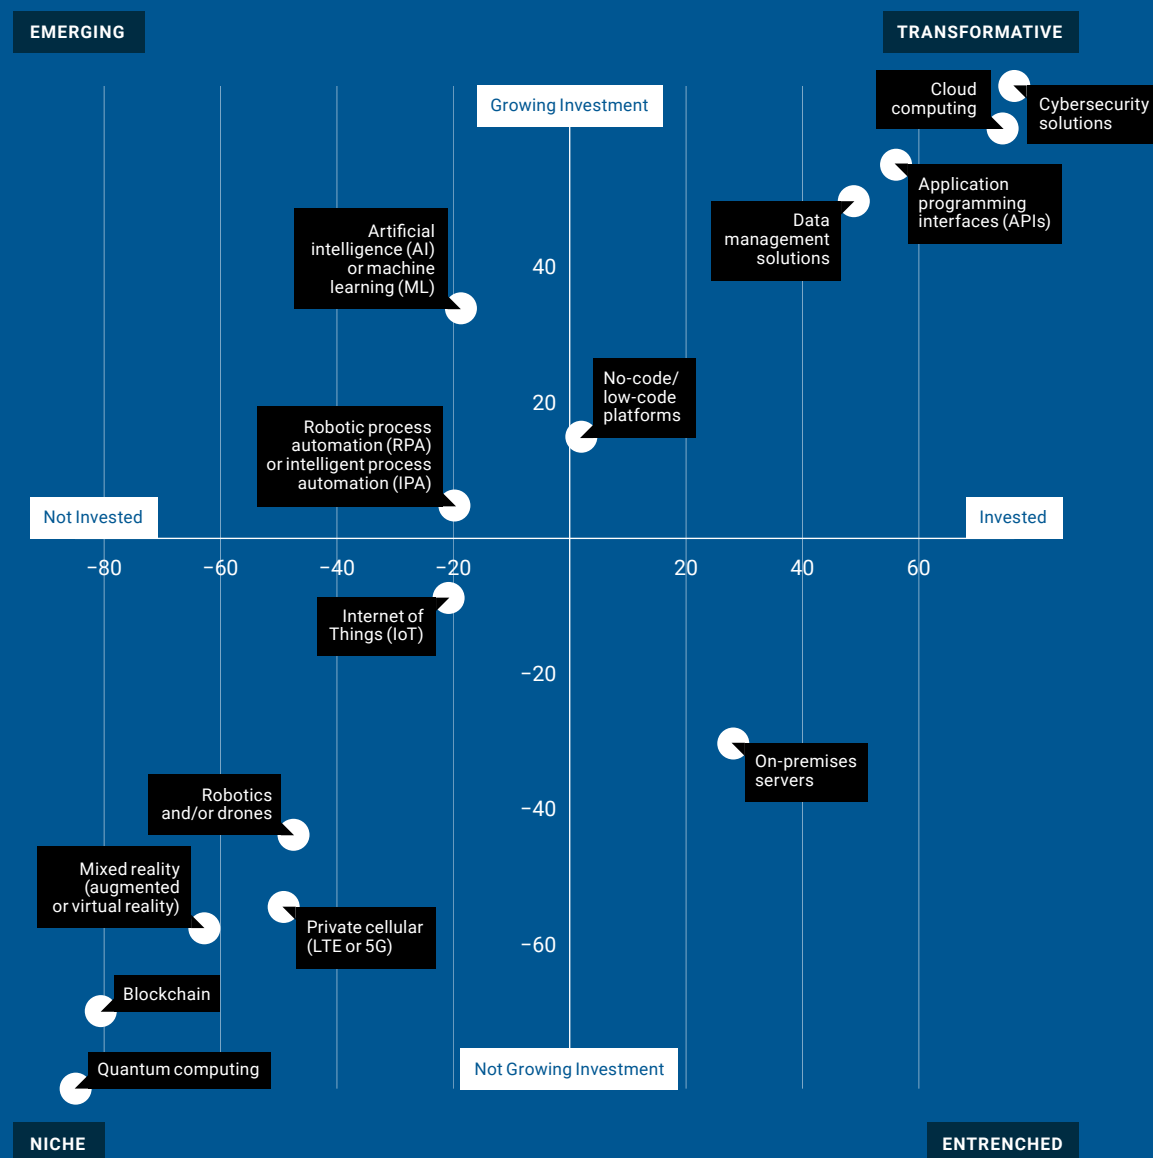
- ▶ More likely to be from smaller organizations, with 52% of respondents estimating a total headcount of below 1,000.
- ▶ More likely to have a smaller IT budget, with 65% reporting a budget of under \$10 million.
- ▶ 63% located in North America.
- ▶ Represent a wide swath of industries including 26% in government and 11% in manufacturing.
- ▶ Most likely to rate IT maturity level at "Support" (42%) or "Optimize" (27%).

We are interested in delineating between AI adopters and skeptics because AI and machine learning (ML) will see the fastest-growing adoption among all emerging technologies in our survey. Nearly one-third of respondents say they plan to invest in AI next year. An additional 35% say they are already invested and plan more investment in AI.

TOP FIVE TECHNOLOGIES: ORGANIZATIONS NOT CURRENTLY INVESTED BUT PLAN TO INVEST IN 2024

- ▶ AI or ML – 32%
- ▶ Robotic process automation (RPA) or intelligent process automation (IPA) – 22%
- ▶ No-code/low-code platforms – 20%
- ▶ Internet of Things (IoT) – 14%
- ▶ Data management solutions – 14%

Our emerging technologies quadrant considers existing investment and intended investment for the year ahead as growth indicators, and investment planned for further into the future or no investment at all as stagnation. In this analysis, AI is hot on the heels of transformative technologies like cybersecurity, cloud computing, and data management solutions. The planned investment in AI among those not already invested indicates it has more momentum than any of these other transformative technologies for 2024.



INSIGHT //

AI is the most rapidly emerging transformative technology.

OTHER NOTEWORTHY STANDOUTS FROM THE QUADRANT

- ▶ Mixed reality leads the “not invested but plan to invest after 2024” category at 21%.
- ▶ Quantum computing leads the “No plans to invest” category at 81%.
- ▶ On-premises servers lead the “Already invested, but do not plan further investment” category at 33%.

FOCUS ON TRANSFORMERS

We’ll also feature some highlights from another group of “Transformers,” or organizations that rank themselves at the top of Info-Tech’s IT maturity scale.

IT MATURITY LEVEL

Choices	Response %
IT transforms the business	14.2%
IT expands the business	8.1%
IT optimizes the business	28.4%
IT supports the business	38.9%
IT struggles to support the business	10.4%

n=676

About one in six IT leaders describe themselves as innovators. Most put themselves at either the “Support” or “Optimize” level of maturity.

METHODOLOGY

Info-Tech's Future of IT 2024 survey collected responses from May 23 to August 22, 2023. The online survey received 894 total responses, with 496 participants completing every question and 382 partially completing the survey. All respondents either work in IT or direct IT.

FIRMOGRAPHICS

SIZE OF ORGANIZATION

(1) 0-250	21.34%
(2) 251-1,000	27.12%
(3) 1,001-2,500	16.98%
(4) 2,501-5,000	11.44%
(5) More than 5,000	23.11%
n=848	

SENIORITY

(1) 0-250	21.34%
(2) 251-1,000	27.12%
(3) 1,001-2,500	16.98%
(4) 2,501-5,000	11.44%
(5) More than 5,000	23.11%
(1) Manager	21.46%
(2) Director-level	23.47%
(3) C-level officer	24.53%
(4) VP-level	8.96%
(6) Team member	10.97%
(7) Owner/President/CEO	4.13%
(8) Consultant	5.78%
(9) Contractor	0.71%
n=848	

REGION

(1) United States	46.23%
(2) Canada	14.15%
(3) Australia/New Zealand	8.25%
(4) Africa	6.01%
(5) Other (Europe)	5.31%
(6) Great Britain	5.19%
(7) Latin America/South America/Caribbean	3.3%
(8) Other (Asia)	3.18%
(9) Middle East	2.5%
(10) Germany	2.2%
(11) India	1.6%
(12) Netherlands	0.6%
(13) Japan	0.6%
(14) Mexico	0.5%
(15) China	0.6%
n=848	

INDUSTRY

What is your organization's primary industry?	Response %
Arts & Entertainment (including sports)	1.18%
Construction	1.89%
Education	8.02%
Financial Services (including banking & insurance)	12.38%
Gaming & Hospitality	1.42%
Government / Public Sector	20.52%
Healthcare & Life Sciences	8.25%
Manufacturing	10.26%
Not for Profit (including professional associations)	2.95%
Media, Information, Telecom & Technology	8.02%
Professional Services	7.19%
Retail & Wholesale	1.89%
Transportation & Warehousing	2.12%
Utilities	3.07%
Real Estate and Property Management	1.53%
Natural Resources	1.77%
Other (Please specify)	7.55%
n=848	

DEVELOPING THE TRENDS

Info-Tech conducts brainstorming exercises with its expert advisors and former CIOs to determine the implications of megatrends for technology decision makers. The Future of IT Survey is developed based on the lines of inquiry suggested by those implications. Analysts with functional area expertise then design hypotheses that are tested in the survey results. The trends featured in this report are based on those results. Info-Tech's Priorities reports also leverage this research, determining the urgency with which external pressures must be responded to for different functional roles.

Info-Tech's design team leveraged generative AI to create the artwork for *Tech Trends 2024*. Note that we've included our prompts with each image.

IMAGINE //

Digital neurons in space, photorealism, accent lighting, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--c 75 --s 750 --ar 16:9



SEIZE OPPORTUNITIES



MITIGATE RISKS



SECURITY BY DESIGN

PAGE 54



RESPONSIBLE AI

PAGE 42



DIGITAL SOVEREIGNTY

PAGE 66

AI-DRIVEN BUSINESS MODELS

PREDICTIONS
THAT CREATE
CUSTOMER
VALUE

IMAGINE //

A cinematic scene from a fantasy movie called Legend of the 7 Rings. This long shot captures a silver and purple holographic crystal ball, accent lighting, sunny day, in the style of dynamic action sequences, captured on Phantom High-Speed Camera

--c 50 --ar 16:9 --s 750



INTRODUCTION

PREDICTIONS AS PRODUCTS

The commercialization of AI models is based on the value of an accurate prediction. Algorithm builders train their neural networks to make good predictions by using a lot of historical data and sometimes adding in human feedback to help sort out special circumstances. Once trained, the algorithms can make predictions based on new data. It's a concept that the tech giants of our era have demonstrated for the past decade, with Facebook predicting which ads you're most likely to find relevant, or Amazon predicting what products you'll want to purchase next.

More recently, we've seen the technology sector move from augmenting its business models (e.g. to sell ads, or as buy-everything e-commerce stores) with AI predictions to making the AI predictions themselves the product. Midjourney is an example of an image generator that

predicts what an image should look like based on a user's prompt. OpenAI's ChatGPT predicts the right words to respond to a prompt. But selling predictions won't stop there. As AI becomes more effective, it's displacing established approaches to solve problems and helping industries solve previously unsolved problems. It's being used by airports to manage flight control centers, by pharmaceutical firms to research new drugs, and by financial services firms to detect fraud.

According to our survey, most IT organizations are making plans for AI to drive strategic aspects of their business in 2024. It will be uncharted territory for many, and there will be new risks to consider as these new business models are forged.

In our **2021 Tech Trends report**, we featured **"Machine Learning by Design"** and detailed how organizational structures were inhibiting the transformational potential of machine learning. The report made the case that organizations should organize around making machine learning a core piece of their value proposition. This year, we're emphasizing how AI should factor into the broader strategic business model.

AI-based business strategies aren't just for those on the cutting edge. In our Future of IT survey, about 1 in 5 IT leaders told us they are already using AI to help define business strategy.

SIGNALS

INVESTMENT IS OPTIMISM

Discourse about AI tends to sway between two extremes – either it will wipe out humanity or it will solve all of our problems; either it will cause mass unemployment or it will free workers from the shackles of tedious minutiae to focus on more valuable tasks. Yet most IT practitioners tend to see AI's impact as somewhere in the middle, while maintaining optimism overall.

AI adopters are much more optimistic than skeptics. Two-thirds of them say AI will bring benefits to their businesses. But skeptics aren't doom and gloom either – half are merely on the fence about it, anticipating a balance of benefits and challenges. Only 3% of skeptics feel their business faces an existential threat from AI, and no adopters are in this camp. Transformers are similar to AI adopters in this area, with two-thirds also saying they are feeling positive about AI.

Organizations are making plans for AI to feature in strategy and risk management capabilities. Among AI adopters, "Business analytics or intelligence" is the most popular selection in this category, with more than three-quarters planning to use AI there by the end of 2024. Seven in 10 organizations also plan to use AI to identify risks and improve security by the end of 2024. Since AI skeptics are not investing in AI before the end of next year, most of them skipped this category or indicated delayed or no plans to use AI in any of these areas. But some did indicate plans to use AI in these areas despite a lack of investment. Perhaps they're hoping to dabble with free trials or have their workers fiddle with open source models.


66%

of companies choosing to invest in AI before the end of 2024 expect AI to have a positive impact on their company vs. 38% of companies not investing in AI.

TRANSFORMERS //
Are more likely than other IT organizations to expect a positive impact from AI, with 65.6% being optimistic.

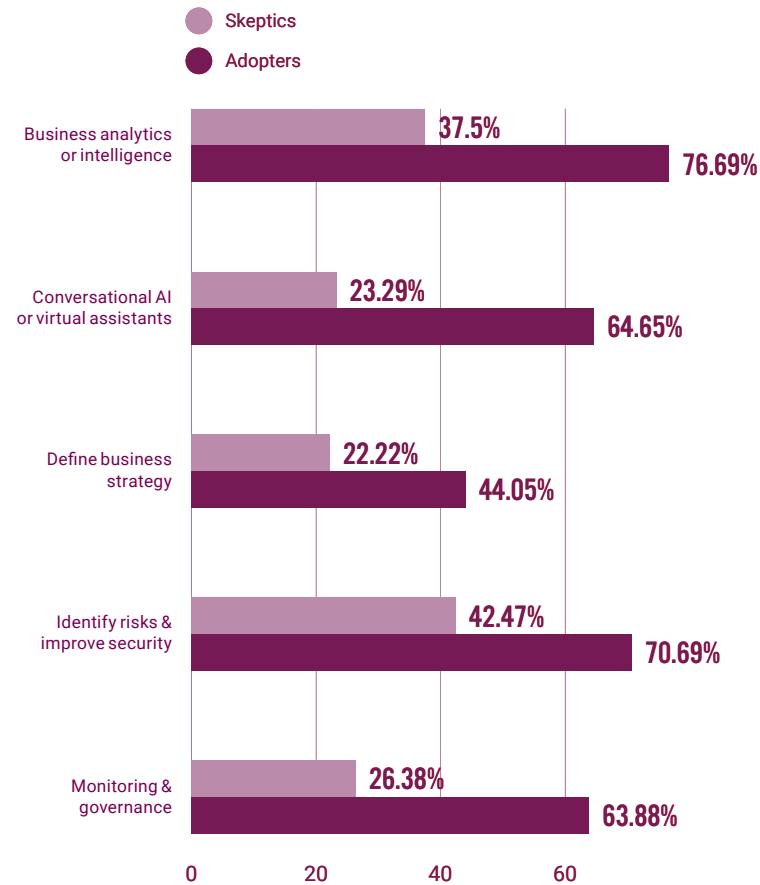
SURVEY

WHAT OVERALL IMPACT DO YOU EXPECT AI TO HAVE ON YOUR ORGANIZATION?



SURVEY

BY THE END OF NEXT YEAR, WHAT SORT OF STRATEGIC OR RISK MANAGEMENT TASKS WILL YOUR ORGANIZATION BE USING AI FOR?



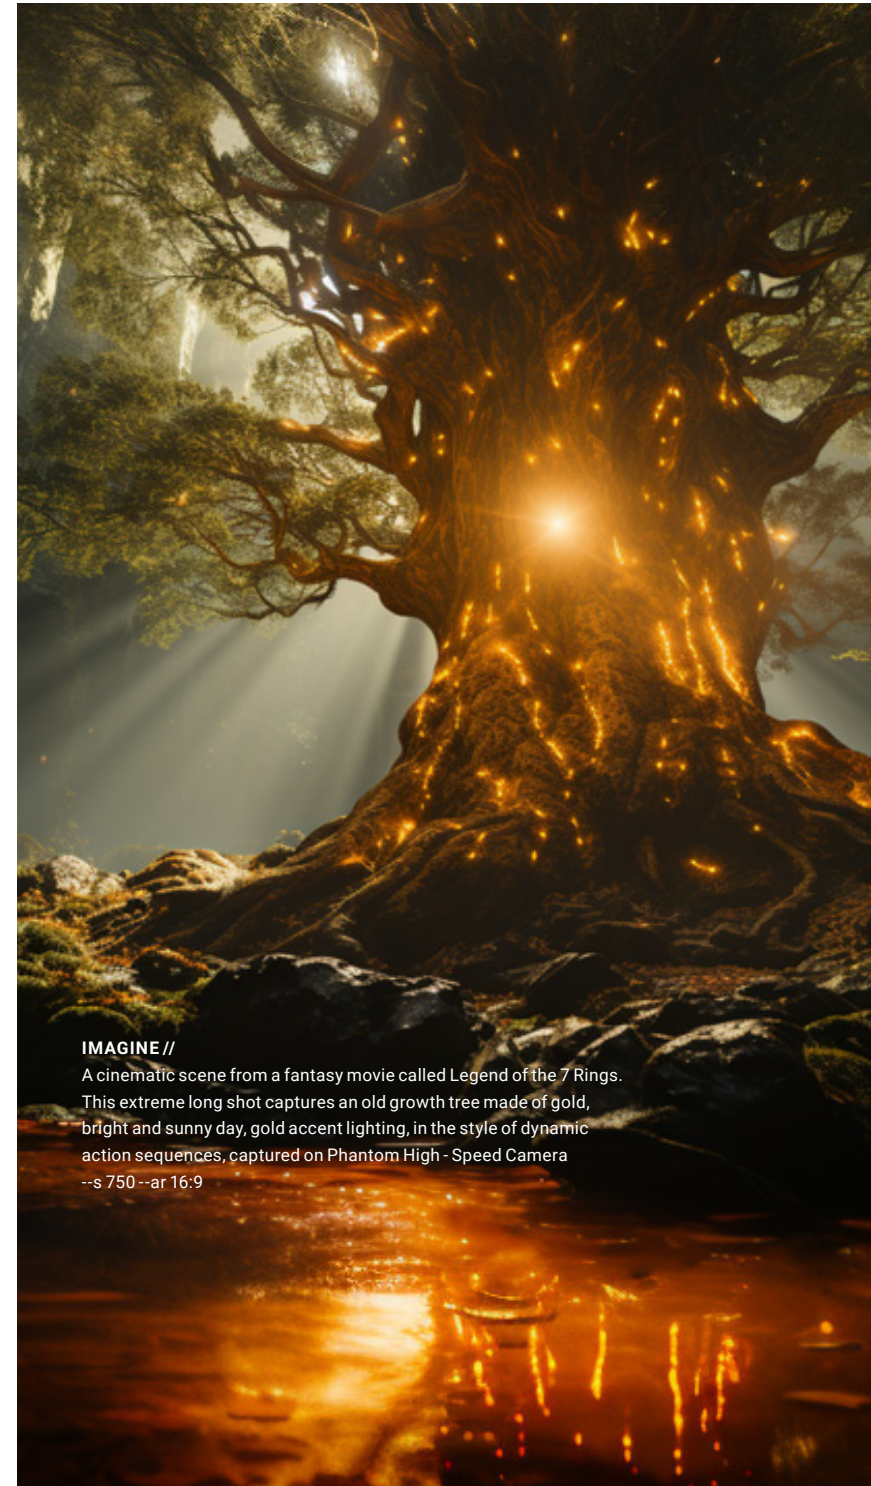
The Transformers segment stands out here, as they indicate by far the most interest in using AI to define business strategy, with more than two-thirds saying so. Fewer than half of adopters plan to do so.

AI adopters will be using AI in several strategic areas by the end of 2024

77% for business analytics and intelligence, and 71% to identify risks and improve security, while most skeptics won't apply AI in any of these areas.

68%

of Transformers say AI will define business strategy by the end of 2024.



IMAGINE //

A cinematic scene from a fantasy movie called Legend of the 7 Rings. This extreme long shot captures an old growth tree made of gold, bright and sunny day, gold accent lighting, in the style of dynamic action sequences, captured on Phantom High-Speed Camera --s 750 --ar 16:9



IMAGINE //

A cinematic scene from a fantasy movie called Legend of the 7 Rings. This long shot captures a floating crystal castle in the sky, accent lighting, sunny day, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--c 50 --ar 16:9 --s 750

OPPORTUNITIES & RISKS

A HIGH-RISK, HIGH-REWARD SCENARIO

OPPORTUNITIES

► DISRUPTION

Using AI predictions to solve problems that previously required more overhead, or to solve new problems altogether, has the potential to disrupt many different industries. Similar to the digital revolution that saw software take on so many business operations more effectively, AI is quickly becoming an obvious best option for many tasks.

► REDUCE DECISION FRICTION

AI can consider many more complex factors in a given situation than a person ever could and boil it down to a simpler choice (Agrawal et al., 2022). This can help organizations provide customers opportunities they might not have taken, or it can empower employees to push ahead with a project.

► SCALE

Training models is difficult, requiring talented data scientists and access to powerful compute resources. But once a model is fully baked, it can be deployed to edge devices to provide service at very little cost. It's an upfront capital requirement with low long-term overhead that is easy to scale.

RISKS

► EASY TO REPLICATE

After OpenAI made its splash with ChatGPT in November 2022, Meta responded by releasing its model's code to open source ("Meta Made Its AI Tech Open-Source," The New York Times, 2023). This gave developers an alternative path to harnessing the capabilities of a large model without paying to use OpenAI's APIs. ChatGPT continues to do a brisk business, but already, many more similar chatbots have emerged for use free of charge. With the method of building foundation models commercialized, businesses may find their competitors are able to quickly respond to any competitive advantage with similar updates. Pushing the capabilities to market for free could drive the value of making certain predictions to zero and disrupt a business model.

► RAPID OBSOLESCENCE

Here's another shortcut a model might take to irrelevance, as technological advancement in this space seems to be a weekly occurrence. Multimodal inputs look to be the next advancement on the horizon, which would make text-only, speech-only, or image-only models seem antiquated only months after creating shockwaves around the world (Meta AI, 2023).

► ETHICAL QUANDARIES

Creators of generative AI models are openly saying there may be a 1 in 10 chance that AI poses an existential threat to humanity. Whether they're right or not, some will say anyone developing AI capabilities is contributing to the problem. Ethical concerns don't stop there, as many creators are fighting back against perceived theft of intellectual property. Also, opting to use AI instead of hiring a person to do a job is likely to invite criticism.

IMAGINE //

A cinematic scene from a fantasy movie called Legend of the 7 Rings. This long shot captures a gaunt dark sorcerer summoning a cloud of chaos, technicolor

--c 25 --s 250 --ar 16:9





We see the visible light spectrum from a light bulb because our eyes are adapted to that frequency of electromagnetic waves. But for Wi-Fi, we're not adapted. If we could see these signals or interpret them, what would that give us?

TAJ MANKU,
CEO, COGNITIVE SYSTEMS

CASE STUDY


WHAT DOES YOUR ROUTER SEE?

SITUATION

Working in a previous company that he founded to build chips for cellular phones, Taj Manku often considered how the chips could “see” cellular base stations and Wi-Fi access points in a way that humans couldn’t. Instead of the opaque objects perceived by the human eye, they were illuminating beacons, radiating out electromagnetic signals. What if, the physicist and Ph.D. lecturer at the University of Waterloo wondered, we could give people the ability to see that signal in the same way? In 2014 he founded Cognitive Systems Corp. to find out.

“It spawned from the idea of how can we use this radiation that’s already in your home and how can we build application on this type of technology,” he says in an interview. Cognitive Systems trained an AI system that could sit on a Wi-Fi access point and interpret the signals in a different way beyond the information being transmitted. The AI uses the Wi-Fi fields between the access points and the devices connected to it to understand the environment of the home. Then it can detect when a human moves through that environment, disturbing the signal. A statistical profile is used to detect the unique way a human body partially reflects the signal as it moves.

To sort out human movement from pet movement or a fan, Cognitive Systems used reinforcement learning with human feedback (RLHF), a combination of a computer looking for patterns and a researcher providing feedback about whether it’s correct or not. The model that’s deployed to the edge – in this case, a Wi-Fi access point – can adapt to a changing environment if someone decides to rearrange the furniture.



IMAGINE //

A cinematic scene from a fantasy movie called Legend of the 7 Rings. This extreme long shot captures a network of floating metal orbs with a force field around them in the desert, technicolor --ar 16:9 --c 25 --s 250



ACTION

The first go-to-market strategy was to sell the service directly to consumers. But after slow uptake, convincing one customer at a time to fiddle with their router settings, Cognitive Systems pivoted to a business-to-business model as a software vendor. It partnered with chip makers to receive the deep-system access its software – dubbed Wi-Fi Motion – required for installation on routers, and sold the software to internet service providers (ISPs) that could deploy at scale. Cognitive charges ISPs in a SaaS model, creating recurring revenue.

ISPs get a value-added service they can deliver to customers. Typically ISPs only see customers use their applications to pay a bill or resolve a service disruption, so providing a beneficial feature is an opportunity to create a better relationship. The primary value proposition of the service is as a security system that requires no additional hardware. When customers are away from home and don't expect anyone to be in the house, they can be alerted to the presence of a person.

Once a customer activates the home monitoring service, there are upselling opportunities. A wellness monitoring feature can alert a caregiver when an elderly home occupant hasn't moved for an extended period, and a smart home automation can adjust the thermostat and turn the lights on or off according to people's movement through the home and out the door.

Privacy is a priority for Manku, who chooses to comply with the strictest data privacy laws in the world – currently those of the state of California, he says. Customers must opt in to using the software on their access points first. The technology isn't capable of identifying an individual – it can only detect a person's movements – and the information isn't discrete enough to differentiate between someone doing jumping jacks and running on the spot.

RESULT

Today, Cognitive Systems is deployed to more than 8 million homes and is growing. It sees its revenues double annually. It is working with 150 ISPs around the world and is seeing those ISPs onboard new users of the service every day. Manku offers this advice to entrepreneurs pursuing an AI business model:

"It has to be scalable. I would try to stay away from hardware. I would focus on a software-based solution. Hardware solutions are tougher because you have to deal with a lot of the management of the procurement, and that can be difficult."

Rather than trying to find one customer at a time, look for an opportunity to find a million customers or more at a time. ChatGPT falls into that category, becoming the fastest growing technology ever by making its service available to anyone via a web browser.



IMAGINE //

A cinematic scene from a fantasy movie called Legend of the 7 Rings. This long shot captures a silver and purple holographic vortex, accent lighting, sunny day, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--c 50 --ar 16:9 --s 750

WHAT'S NEXT

A MODEL OF IMPERFECTION

► MULTI-MODAL MODELS

If 2023 has been all about LLMs (large language models) then 2024 might be about MMMs (multimodal models). These models will be capable of receiving different modes of input, such as an image, text, or audio, and generating multiple modes of output in turn. Meta's SeamlessM4T model is an example, combining both text and speech in a model designed to translate between almost 100 different languages. The model supports nearly 100 languages for both speech and text input, can provide text transcription in nearly 100 languages, and can provide speech output in 36 languages (Meta AI, 2023).

► LONGER-CONTEXT MODELS

Developers using ChatGPT or its API equivalent often run into a barrier with the amount of specific context that they can provide to the model. It's a concept that's referred to as the attention span of the model, and the longer it is, the more useful to enterprises who want to use their own data to guide output. Building a longer attention span is one of the main motivations to train new foundation models, and the cutting edge at the moment is a 32K (or 32,000 tokens, equivalent to about 25,000 words) limit, offered by OpenAI's GPT-4 model ("What Is the Difference Between the GPT-4 Models?" OpenAI, 2023), and the open source model Giraffe, built by Abacus.AI (Abacus.ai, 2023).

► INDUSTRY-SPECIFIC MODELS

While large models are flexible and can be used for a number of different tasks, many industries find these models too unreliable to depend upon. Accuracy isn't important if you're using ChatGPT to provide dialogue for a character in a video game, but it is necessary if you're going to use it to present citations in a courtroom or engineer a new drug. Specific data sets will be required to hone AI enough to make accurate predictions, and even then, humans will need to work to verify results. Early examples of industry-specific models come from the legal industry, where Harvey AI, CoCounsel, and LitiGate all compete to offer law firms AI services. Similarly, in the pharmaceutical industry, not only is AI helping design new drugs in the lab and predict their likelihood of approval by regulators, but it is also helping monitor clinical trials by interpreting sensor data ("AI Poised to Revolutionize Drug Development," Forbes, 2023).



RECOMMENDATIONS

Align with your business stakeholders on opportunities to provide customer-facing value with AI. Primary considerations will be what problems your customers are trying to solve, where they face friction with your products and services at present, and what data you own that can be harnessed to train a model. Building a pilot project to test out new ideas is desirable to test ideas in the real world rather than try to transform the entire business all at once.

INFO-TECH RESOURCES

► **BUILD YOUR GENERATIVE AI ROADMAP**

Generative AI has made a grand entrance, presenting opportunities and causing disruption across organizations and industries. Moving beyond the hype, it's imperative to build and implement a strategic plan to adopt generative AI and outpace competitors.

Yet generative AI has to be done right because the opportunity comes with risks, and the investments have to be tied to outcomes.

► **AI TRENDS 2023**

As AI technologies are constantly evolving, organizations are looking for AI trends and research developments to understand the future applications of AI in their industries.

► **BUILD YOUR ENTERPRISE INNOVATION PROGRAM**

Collect ideas from business stakeholders in a constructive way and prioritize initiatives that could be worthy of a pilot project.

IMAGINE //

A cinematic scene from a fantasy movie called Legend of the 7 Rings. This extreme long shot captures lush green enchanted forest with crystal ball in the sky, purple accent lighting, sunny day, in the style of dynamic action sequences, captured on Phantom High-Speed Camera --c 100 --s 750 --ar 16:9

AUTONOMIZED BACK OFFICE

DRIVING EFFICIENCY
IN COGNITIVE TASKS

IMAGINE //

A cinematic scene from a utopian science fiction movie called The Uncanny Office. Long shot captures three blue and silver holographic people talking in a well-lit office building on a bright and sunny day, iridescent interface, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High-Speed Camera

--c 50 --s 750 --ar 16:9

INTRODUCTION

ENTERPRISE SOFTWARE GETS CHATTY

IT's role has always been to autonomize business systems by providing capabilities that allow systems to self-execute and self-regulate toward company goals in the name of efficiency. With generative AI, a wide range of new tasks become possible to automate toward this goal. These AI models are adaptable and flexible, able to process large volumes of unstructured data and provide classification, editing, summarization, new content creation, and more. Consultancy McKinsey estimates that, by automating these routine cognitive tasks, generative AI's impact on the economy could add \$2.6 to \$4.4 trillion in value across 63 use cases – and double it if generative AI is embedded into software already used for other tasks beyond those use cases (McKinsey, 2023).

So even for organizations not transforming their business model around AI, there will be value to reap from streamlining current operations. Some of this increase in efficiency will be delivered by using new applications or web services, such as ChatGPT, but much of it will be delivered through new features in software that's upgraded with new AI-powered features. With the software as a service (SaaS) model, in many cases, enterprises won't even need to deploy an upgrade to harness these new features. Existing vendor contracts will be the most likely avenue to add generative AI to many enterprises' IT arsenal. The list of vendors that have announced generative AI features is too long to include here, but consider several examples of vendors in the IT space alone:

- ▶ Juniper Networks announced integration of ChatGPT with Marvis, its virtual network assistant. The chatbot will be better at helping users review documentation and provide customer support to resolve networking issues (Juniper Networks, 2023).
- ▶ CrowdStrike released Charlotte AI to customer preview, its own generative AI that answers questions about cybersecurity threats and allows users to use prompts to direct the automation of repetitive tasks on the Falcon platform (SDX Central, 2023).
- ▶ ServiceNow announced the Now Assist assistant for its Now platform, which automates IT service workflows. The assistant summarizes case incidents. Another feature allows developers to generate code with text prompts (CIO, 2023).

In other lines of business, major vendors like Microsoft, Salesforce, Adobe, and Moveworks are among those announcing generative AI features. Generative AI is going to impact all industries, but some sooner than others. As we'll see in the case study, the legal industry is one where generative AI solutions are more specialized and deployed among early adopters.

First, we'll examine how organizations plan to approach new generative AI features from vendors.

SIGNALS

IT'S EITHER ROLL OUT OR OPT OUT


Many generative AI features will enter the enterprise through feature upgrades to existing business applications. In some cases, IT may have the keys to the admin controls, and in other cases, it will rest with the line of business that procured the solution. For SaaS solutions that bolt on generative AI chatbots and other features, IT may find that they are turned on by default with new versions, and action is required to opt out of them.

If given the choice, nearly half of adopters (47%) are keen to adopt new generative AI features from major vendors either in beta access (17%) or when generally available (30%). The other half are still taking a more cautious approach, with 37% saying they need more information before deciding and 16% saying they will hold off on the features until other organizations test them.

Skeptics are about twice as likely as adopters to say they need more information or are not interested in adopting generative AI features at all. Less than 1 in 5 skeptics say they will be adopting new generative AI features at general availability or sooner.

SURVEY

WITH BUSINESS APPLICATION PROVIDERS PLANNING TO UPGRADE THEIR SOFTWARE WITH GENERATIVE AI FEATURES (E.G. MICROSOFT COPILOT, ADOBE FIREFLY), HOW DO YOU PLAN TO MANAGE THE ROLLOUT OF THESE FEATURES?



INSIGHT //

Adopters are 2.5 times more likely than skeptics to say they will adopt generative AI features from vendors either in beta or after general availability.

63%

of **TRANSFORMERS** say they will adopt generative AI features from vendors either in beta or when they become generally available.

WHAT BACK-OFFICE JOBS WILL AI DO?

We asked what type of operational tasks organizations are most interested in using AI for. One in 3 adopters say they are already using AI to automate repetitive, low-level tasks. Another 45% say they plan to do so in 2024. More than a quarter of adopters are also already using AI for content creation (27%), with another 30% saying they will do so in 2024. More than one quarter of adopters say they already use AI for IT operations (27%), and 42% say they will use it for IT operations in 2024. Applying AI to IoT and sensor data

generated the least interest among adopters, with 41% saying they had no plans to use it.

Skeptics aren't likely to have adopted AI for any operational tasks yet, but they are more likely to leave room for adoption rather than close the door on it completely.

There's one thing that most adopters and skeptics seem to agree on – they are more interested in seeing AI automate tasks rather than augmenting operational staff in their


decision making. Almost 1 in 5 adopters say they have no plans to pursue augmenting staff with AI, and 46% of skeptics say the same. This seems to run contrary to the message that many businesses and vendors often say about AI: that it is intended not as a replacement for people doing jobs, but as an augmentation.

INSIGHT //

Adopters use AI to automate repetitive, low-level tasks for them and are interested in AI's ability to take on more content creation and IT operations tasks starting in 2024.

SURVEY

BY THE END OF NEXT YEAR, WHAT SORT OF OPERATIONAL TASKS WILL YOUR ORGANIZATION BE USING AI FOR?



IMAGINE //

A cinematic scene from a utopian science fiction movie called The Uncanny Office. Long shot captures three blue and silver holographic people talking in a well-lit office building on a bright and sunny day, iridescent interface, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High-Speed Camera

--c 50 --s 750 --ar 16:9



OPPORTUNITIES & RISKS

SEIZE OPPORTUNITIES & MITIGATE RISKS

OPPORTUNITIES

► COST SAVINGS AND EFFICIENCY GAINS

With more cognitive tasks automated, employee time can be spent on higher-value tasks, or less overhead may be required to manage a process. Organizations will be able to scale to support more business without being bogged down by administrative nickel-and-diming, though using generative AI will represent a new cost in itself.

► IMPROVED OUTPUT

By getting to a first draft more quickly, workers can spend more time honing their message and putting a point on the finer details. Using generative AI to augment workers is often a path to improved quality and modest time savings.

► EASE OF ACCESS

With major enterprise vendors eager to compete in launching new generative AI features, the new capabilities may be rolled in as a value-added component to existing contracts. Organizations can work with vendors where they've established a trusted relationship.

RISKS

► HALLUCINATION SITUATION

Even when trained on specific data sets and built for purpose, generative AI is still prone to fabricate information and present it as fact. Knowledge workers using outputs from generative AI tools will need expertise to validate facts provided by these tools (Interview with Monica Goyal).

► DATA SECURITY AND PRIVACY

Old fears about third-party hosts getting access to sensitive data will be revived. Organizations using generative AI features on hosted software will perceive new risks around their data being used to train the vendor's algorithm. Vendors will commit to not doing so in contracts, but risk managers will point out it's still technically possible. New features may be blocked in some situations by cautious administrators.

► ETHICAL LIABILITY

Employees who don't grasp the limits of AI's capabilities may be over-reliant on its output or try to use it for a task that's not appropriate. Governance of new AI capabilities will require training for users to avoid inadvertent or intentional cases of ethical misuse of AI.

IMAGINE //

A cinematic scene from a utopian science fiction movie called The Uncanny Office. Long shot captures three blue and silver holographic people talking in well-lit office building on a bright and sunny day, red accent lighting, iridescent interface, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High-Speed Camera

--ar 16:9 --s 750





Generative AI is going to give us more certainty in how long things take to do, and that's going to allow us to do more fixed-price billing.

MONICA GOYAL,
LAWYER AND DIRECTOR OF
LEGAL INNOVATION,
CARAVEL LAW

CASE STUDY

CHATGPT PASSED THE BAR EXAM; NOW IT'S WORKING IN LAW

SITUATION

When researchers found that OpenAI's GPT-4 could not only pass the bar exam, but do so in the 90th percentile, it made it seem like AI lawyers were just around the corner (ABA Journal, 2023). But that notion took a hit when a New York City lawyer submitted a legal brief created by ChatGPT that was full of fake legal citations, with the lawyer claiming he did not comprehend that ChatGPT could fabricate cases ("The ChatGPT Lawyer Explains Himself," The New York Times, 2023).

Despite that widely covered inauspicious introduction to the courtroom, AI still has the potential to transform the legal industry. AI can augment a lawyer's expert-level capabilities by providing first drafts of legal content, translating technical legal language into more colloquial terms for clients, and reviewing contracts or agreements. In one analysis that included data from 10 corporate legal departments, researchers found that 40% of time entries representing 47% of billing could potentially use generative AI. Given an upper limit of generative AI to reduce that work by half, law firm revenue could be reduced by 23.5% (3 Geeks and a Law Blog, 2023).

Technology companies have already been launched to provide generative AI tools that are specifically trained for the legal industry. These tools are trained on case law that sits behind the paywall. Competing vendors in the space include Harvey.ai, Casetext's CoCounsel, and Litigate.ai. Other applications, like Rally's Spellbook, also apply OpenAI's GPT-4 to more pointed tasks, such as contract review and drafting (Interview with Monica Goyal).

As with many commercial applications, it's still early days for AI in law. But at Caravel Law, Monica Goyal, lawyer and director of legal innovation, is implementing the technology to see where it can optimize the operations of this Toronto-based non-litigating practice.



IMAGINE //

A cinematic scene from a utopian science fiction movie called The Uncanny Office. Long shot captures 3 blue and silver holographic males and females working in a well lit office building indoors on a bright and sunny day, iridescent interface, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--c 50 --s 750 --ar 16:9



ACTION

Caravel is an early user of Harvey, which is not yet generally available. Based on OpenAI's GPT-4, Harvey also received funding from OpenAI to pursue a solution for the legal industry (Global Legal Post, 2023). Harvey works similarly to ChatGPT, offering a simple chatbot interface that allows users to enter prompts.

"If you have some training around what is a good user prompt, you will get a better result," Goyal says. "The more fulsome you can be, the better the output."

Lawyers can ask Harvey legal questions about particular areas of law. Caravel's lawyers have learned to specify that they are interested in Ontario-based law, and they find that improves the results. Harvey also has a document upload feature, so a user can submit a PDF or an email as context and then have Harvey draft a document such as a notice or client correspondence. The output will not only lean on its trained model, but also reference the uploaded documents with citations.

Caravel is also using Spellbook by Rally, which is a Microsoft Word plug-in that uses GPT-4 to review contracts and suggest additions. It's also looking to improve other back-office operations such as business development support. Goyal's team customized Julius.ai to allow their business development representatives to query it about lawyers' skillsets and availability for new clients.

Goyal leads a tech innovation team that is assembled based on the project requirements. To customize Julius, she contracted short-term employees. In other situations, she's partnered with a vendor or hired a consultant.

Caravel finds the technology effective overall but cautions that it's still prone to make mistakes, like citing case law that doesn't exist, and that its output must be validated. Caravel's lawyers tend to have 10 years of experience or more, and Goyal worries that a less experienced lawyer wouldn't have the expertise to properly validate output from generative AI tools. "If you have less than five years of experience, you might not do very well to verify it's accurate," she says.

RESULT

After accounting for the additional validation of output, Goyal estimates Caravel lawyers are saving about 25%-50% of their time spent on creating a legal draft. In the case of Spellbook, lawyers are saving more like 10%-15% of their time on contract reviews.

"It's going to save you some time, but you have to go through the contract and make sure that you read it and that you know it," Goyal says. But using generative AI isn't just saving the lawyers time – it's also about creating a better output in the end.

Harvey commits to its customers that it will not use their data to train their model. Goyal says this is sufficient assurance for protecting sensitive data.

Goyal cautions that in-house legal practices should have clear processes in place before trying to deploy new AI solutions. It's also important to understand how the software works and what its limits are, recruiting help through vendors or consultants if necessary. But she's confident the value can be realized by those who understand it.

As a result, law firms are talking about reevaluating the billable-hour model. "People have been talking for a long time in the industry about how the billable model doesn't work for clients," she says. "They really don't like it." Lawyers had to use the model because their services are so bespoke, and it's uncertain how much time will be required for services. But with generative AI providing more streamlining, lawyers could be more confident about cost certainty and do fixed-price billing in certain scenarios.



IMAGINE //

A cinematic scene from a utopian science fiction movie called The Uncanny Office. Long shot captures gold iridescent holographic building on a bright and sunny day, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--c 50 --s 750 --ar 16:9

WHAT'S NEXT

DIFFERENTIATION COMES FROM THE FOUNDATION

Vendors releasing generative AI features will either be partnering with an AI-focused company such as OpenAI to provide a foundation model, or they will train their own models. Savvy technology purchasers will set aside vendors' promises about the benefits of these software features and consider the pros and cons of both approaches:

- ▶ Vendors that integrate a third-party AI model will have to answer questions about whether customer data is exposed to that party. But the foundational model itself may be more flexible and provide more utility due to being created by an AI-focused firm. There is the risk that the model provider could go out of business or run into regulatory trouble with their algorithms, and that this could affect the performance of the vendor's solution.
- ▶ Vendors that train a proprietary model will have to answer questions about whether they themselves are using customer data to train their own AI models. Customers will want the option to consent to such an arrangement and will expect sufficient value in return. Models are likely to be more purpose-built and less flexible.

There is yet a third, hybrid approach to consider in which a vendor starts with a foundation model provided by an AI company but customizes the model and licenses the rights to host it on their own infrastructure. In our case study, Harvey.ai is an example of this hybrid approach, adapting OpenAI's GPT-4 model with financial backing from the company as well.

IMAGINE //

A cinematic scene from a utopian science fiction movie called The Uncanny Office. Long shot captures three blue and silver holographic people talking in a well-lit office building on a bright and sunny day, iridescent interface, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High-Speed Camera

--ar 16:9 --c 50 --s 750

SHADOW AI

Following the release of ChatGPT's beta to the web in November 2022, many organizations quickly deployed policies telling employees not to use the tool. One survey by BlackBerry found 75% of organizations were considering or implementing a ban (BlackBerry, 2023). Yet actually preventing its use is difficult to enforce since it's free to use and only requires a web browser.

We might expect a similar situation when vendors begin rolling out their own chatbots and other generative AI-powered features. IT departments will need strong governance models to enforce limitations on accessing new features they aren't comfortable with. At the same time, overly strict limitations on using these new features will give business departments an incentive to cut IT out of the equation and go directly to vendors. Establishing the risk tolerance and specific no-go areas with top level leadership is going to be an important step in effective governance.

RECOMMENDATIONS

Organizations should look to their trusted vendor relationships for opportunities to harness new generative AI features in the tools they are already familiar with. CIOs should keep apprised of new feature releases and any changes to terms of use that come along with them. Once they are satisfied there is no additional risk introduced around sensitive data, there are two paths to pursue for value realization. A pilot project that identifies a specific use case for new features can be selected and launched, or business users can be educated about new features and left to incorporate them to improve their own productivity.

INFO-TECH RESOURCES

- ▶ **IMPROVE IT OPERATIONS WITH AI AND ML**
Prioritize IT use cases for automation and make a plan to deploy AI capabilities to improve your IT operations. Calculate return on investment for solutions and create a roadmap to communicate a deployment plan.
- ▶ **GOVERN OFFICE 365**
Prepare for the new generative AI features coming to Office 365 by aligning your business goals to the administration features available in the console. Apply governance that reflects IT's requirements and control Office 365 through tools, policies, and plans.
- ▶ **ESTABLISH A COMMUNICATION AND COLLABORATION SYSTEM STRATEGY**
Cut through the redundant and overlapping collaboration applications and give users a say in how they want to work together and what tools they can use. The impact is reducing shadow IT and the burden on application maintenance.





LOOK FIRST TO YOUR TRUSTED
VENDORS FOR NEW GENERATIVE
AI FEATURES IN TOOLS YOU ARE
ALREADY FAMILIAR WITH.

IMAGINE //

A cinematic scene from a utopian science fiction movie called *The Uncanny Office*. Long shot captures three blue and silver holographic people talking in a well-lit office building on a bright and sunny day, iridescent interface, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--c 50 --s 750 --ar 16:9

SPATIAL COMPUTING

DIGITAL CONTENT
ANCHORED TO THE
PHYSICAL ENVIRONMENT

IMAGINE //

A cinematic scene from a utopian science fiction movie called *Second Sight*. Extreme long shot captures a person building a purple and silver floating holographic castle in the desert on a bright and sunny day, iridescent, rainbow, in the style of dynamic action sequences, captured on Phantom High-Speed Camera

--ar 16:9 --s 750



INTRODUCTION

FROM THE METAVERSE TO SPATIAL COMPUTING

When Apple debuted its Vision Pro mixed reality headset at its Worldwide Developers Conference (WWDC) in June 2023, it had to explain how headset users could participate in video conferencing. Joining a Zoom call with a phone or a laptop provides a natural place for a camera to point at the user's face, but once they start wearing a headset, that's lost. To solve this problem, Apple demonstrated that the Vision Pro's front-facing cameras and sensors could be used to scan the user's shoulders and head, then AI would generate an accurate likeness in the form of a digital avatar complete with natural facial expressions (TechCrunch, 2023).

The demonstration shows how AI will be an important part of mixed reality's mass commercialization. It's something that Meta also understood, previously sharing plans for virtual assistants that could help headset users cook with augmented reality by identifying where ingredients were in the kitchen or alerting them that they haven't added the salt yet. Also, an assistant capable of receiving voice commands and rendering fully immersive scenes would be part of a virtual reality experience. While Meta called its vision for this future of computing "the metaverse" and Apple chooses "spatial computing" instead, they are both using the same technological building blocks and converging them to an experience that adds up to more than the sum of its parts.

Both visions are also nascent in development, with Apple expecting to sell well under half a million of its first-generation Vision Pro ("Apple Reportedly Expects To Sell," Forbes, 2023). In the meantime, generative AI will

begin to feature as an interface more often in traditional computing experiences. Voice assistants like Siri and Alexa are being improved with large language models, and just about every major enterprise application seems to be announcing plans for a chat bot addition. Mobile apps capable of scanning rooms and objects and converting them into 3D models are already available on app stores.

At the same WWDC where it announced the Vision Pro, Apple also announced the capability for iPhones to take photos of a completed meal and provide the user with the recipe. Even Apple understands that mass adoption of mixed reality headsets may be over the horizon, but AI-powered interface advancements can still power spatial computing experiences through already ubiquitous devices.

Most will wait and see if mixed reality lives up to the hype. At least in the meanwhile, many will be exploring generative AI interfaces that will open the door to more spatial computing applications even without the aid of a headset.




IMAGINE //

A cinematic scene from a utopian science fiction movie called Second Sight. Extreme long shot captures a person building a purple and silver floating holographic bridge in the desert on a bright and sunny day, iridescent, rainbow, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--ar 16:9 --s 750

SIGNALS

BUILDING A BETTER KEYBOARD & MOUSE TRAP

Survey takers express some interest in adopting generative AI as an interface. Examples provided include the new chatbot-powered Bing search, or uses as a human-machine interface. The latter would apply with Apple's Vision Pro, as AI will be used as a core part of the UX to interpret eye movements as navigational cues or even to scan a user's face with the device's front-facing cameras and create an accurate digital avatar for interaction when teleconferencing.



Rate your business' interest in adopting generative AI interfaces (e.g. Bing search, human-machine interfaces)

INSIGHT //

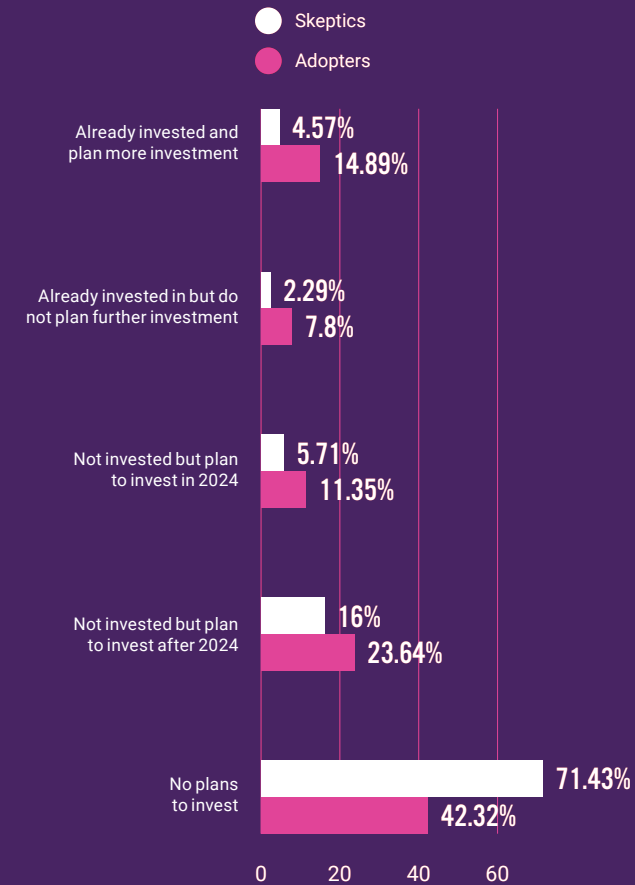
AI adopters are 42% more interested than skeptics in using a generative AI-based interface.

This gives this group the edge in creating value for spatial computing. The next thing we might wonder, then, is how keen are these groups on adopting mixed reality?

Adopters are more likely than skeptics to be keen on mixed reality. About 1 in 5 adopters have already invested in mixed reality, but only about 1 in 15 skeptics has done so. Mixed reality is the most popular infrastructure and hardware technology to plan investment for after 2024 for adopters, but the Internet of Things was the most popular option for skeptics. Overall, most skeptics just don't foresee ever investing in mixed reality.

SURVEY

MIXED REALITY (AUGMENTED OR VIRTUAL REALITY)



INSIGHT //

Organizations invested in or planning investment in AI are more likely to be adopters of mixed reality, but most of that investment is still over the horizon, after 2024.

TRANSFORMERS //


Are almost three times more likely to have already adopted mixed reality than any other group, with 39% saying they've already invested in it.

Apple's Vision Pro and competing headsets may make waves among early adopters, but most will wait and see if mixed reality lives up to the hype. At least in the meanwhile, many will be exploring generative AI interfaces that will open the door to more spatial computing applications even without the aid of a headset.

IMAGINE //

A cinematic scene from a utopian science fiction movie called Second Sight. Extreme long shot captures a person building a purple and silver floating holographic bridge in the desert on a bright and sunny day, iridescent, rainbow, in the style of dynamic action sequences, captured on Phantom High-Speed Camera

--s 750 --ar 16:9



ELIMINATING TIME & SPACE CONSTRAINTS COMES WITH REAL COSTS

OPPORTUNITIES

► OMNIPRESENT EXPERTS

During the pandemic, firms faced the problem of not having auditors or specialized engineers available to travel to inspect facilities or make repairs to complex machinery. A solution many employed was to use mixed reality headsets worn by frontline staff. Experts could effectively see through the eyes of the workers using the device's front-facing camera and effectively communicate with the workers to guide them to complete the tasks. When travel returned after the pandemic, this way of working stuck due to the reduced travel costs and the expediency of advice.

► IMMERSIVE USER EXPERIENCE

Furniture retailers have leveraged spatial computing capabilities to allow prospective customers to see how a new coffee table would fit into their living room by looking through their smartphone. It's becoming increasingly easy for consumers to digitally scan and model their homes and possessions, opening up new commercial possibilities.

► HIGH-END REMOTE COLLABORATION

Organizations finding value in mixed reality collaboration solutions typically aren't just holding meetings to talk about the latest sales numbers. Instead, they are relating to the design of a product or working to construct a new environment. Being able to see a product rendering in 3D alongside one's colleagues despite meeting remotely proves to be magnitudes better than a Zoom call in this situation.



A cinematic scene from a science fiction movie. A person in a dark coat and hat stands in a desert, looking at a large, glowing, iridescent, rainbow-like structure that appears to be a floating holographic castle. The structure is composed of many vertical, translucent columns and is set against a bright, sunny sky. The person is standing on a sandy, rocky ground. The overall scene is captured in a dynamic action sequence style.

RISKS

► PRIVACY CONCERNS

Encouraging users to either scan their homes with their smartphones or wear a headset with always-on, front-facing cameras will be subject to privacy concerns. Users will want assurances their data remains in their control, and the general public will object to a widespread capability to passively capture their likeness and activities in public. Critics will assert these activities deepen surveillance capitalism.

► USER COMFORT

Despite many advances in headsets, many users report adverse effects after longer periods of use. From nausea due to motion sickness to disturbing psychological conditions or just plain old eye strain, using mixed reality headsets throughout a full workday may not be comfortable for many.

► INFRASTRUCTURE CAPACITY

Spatial computing often combines IoT with mobile devices or headsets to construct and experience a digital twin. Such a deployment could exponentially increase the number of devices requiring connectivity, pushing IT departments to deploy next-generation Wi-Fi or cellular solutions to accommodate both the necessary simultaneous connections and increased bandwidth.

IMAGINE //

A cinematic scene from a science fiction movie called *Second Sight*. Extreme long shot captures a person building a purple and silver floating holographic castle in the desert on a bright and sunny day, iridescent, rainbow, in the style of dynamic action sequences, captured on Phantom High-Speed Camera

--ar 16:9 --s 750



AI will replace the drafting up of spaces or objects for manufacturing or construction. When you have a 3D model, you don't need the orthographic projections that are made to construct the product. That was designed in the Renaissance and made to work with paper.

COLIN GRAHAM,
CEO, ARCALOGIX

CASE STUDY


DIGITAL REFLECTIONS OF REAL SPACES

SITUATION

Five years ago, Colin Graham set out to create a business that would replace expensive CAD modeling tools that required special expertise to use with a web-browser tool that anyone could use. "So you wouldn't have to pay the massive subscription prices for Autodesk products, and you didn't need the steep learning curve," he says in an interview.

The drag-and-drop tools were made available to the commercial office sector and found a customer base. But he didn't stop there. For three years, users of the platform manually converted PDF floor plans into 3D models using the tools available. Those projects – more than 5,000 of them – provided data for AWS SageMaker to create Archie, an in-house AI tool. By labeling the various symbols to represent doors, windows, walls, columns, etc. in the floorplans, Graham and his team trained Archie to take an uploaded raster image of a floor plan and convert it into a construction set for their 3D authoring tool.

The result is that users can upload their floorplans and create accurate 3D models of their space in seconds. It's an output that previously would have taken an architect weeks to build. Now users can use Arcalogix to create their floor plans, then tinker with them using browser-based software to imagine renovating their space in different ways.



ACTION

Arcalogix took its toolset capability from mere modeling to digital twin level by integrating with AWS IoT TwinMaker. It now shows status and telemetry reports for IoT devices in an office mapped to their location in a 3D model.

Graham focused on supporting occupancy sensors too, so space owners and operators could see how often any given chair or desk is used or get a sense of how busy the space is at different times of day. With hybrid work arrangements seeing offices less occupied than ever before, office owners are trying to figure out how much space is really needed and how to redesign the environment for its new mode. “What is the optimum configuration of the space in terms of the overall layout and the types of space that will support employees going forward?” Graham asks rhetorically. “What space is used more than other areas and why?”

Office managers could author a redesigned space and discuss it with employees before calling in the contractors. Once deployed, it’s even possible to use Arcalogix as a communications platform, as Graham’s own team does internally. Codenamed “Viza,” it incorporates the Amazon Chime tech stack so remote employees can interact with in-office employees through Arcalogix. Remote employees can book a virtual chair at an in-office meeting and connect with the physically present participants through videoconference. A text-based chat is also available.

“We have it running throughout the day, so at any point you can glance up to see where people are and who they are meeting with,” he says. “You can click on the chair next to them and knock on the door, so to speak, to join the meeting.”

It creates a sense of presence for everyone and helps managers understand what’s going on in the office, Graham says. He acknowledges this use case isn’t going to displace Teams or Zoom, but he thinks it could help some customers organize ad hoc meetings.

RESULT

Arcalogix is a platform that offers multiple types of customer different value propositions today. Aside from office managers, contractors are using the tool to accelerate quotes for jobs, giving clients a tour of a model of what the building will look like when completed.

Product suppliers are also talking to Graham about featuring their office furniture as models available in the product. This would place them at an early stage of office planning for buyer consideration, Graham says, instead of being an afterthought when the space is ready to be filled with chairs and desks.

Customers pay on a per-SKU, per-month basis. Graham is exploring options to work with channel partners that could white label the technology and deploy it to specific use cases.

IMAGINE //

A cinematic scene from a utopian science fiction movie called *Second Sight*. Close up shot captures a person looking into a purple and silver mirror in the desert on a bright and sunny day, iridescent, rainbow, in the style of dynamic action sequences, captured on Phantom High-Speed Camera --s 700 --c 25 --ar 16:9

IMAGINE //

A cinematic scene from a utopian science fiction movie called *Second Sight*. Long shot captures a purple and silver laser scanner in the desert on a bright and sunny day, iridescent, rainbow, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--s 700 --ar 16:9

WHAT'S NEXT

SPATIAL, THE FINAL FRONTIER?

► 3D WEB

The web has supported 3D content to some degree for decades, but a convergence of new technologies and support are improving the user experience both for end users and creators. Different techniques are being used to enable mobile apps that can scan any object with the camera and create a high-quality 3D model. Photogrammetry is one such technique, used by the makers of the MagiScan app, which allows users to scan an object as if they were making a video showing all its angles. The 3D model can then be exported to a platform like NVIDIA Omniverse or used on a business website to showcase their wares. Another technique, Neural Radiance Fields (NeRFs), uses AI to fill in the gaps from fewer photos of an object.

► HEADSET WARS

Apple's Vision Pro is slated to launch in the first half of 2024 barring production delays and will be watched as a test of early adopter appetite. Expect limited availability of the device with Apple treating this first generation as more of a test case than a full-on market entry. Meanwhile Meta is developing a similar headset that supports both augmented and virtual reality, Samsung is preparing its own response to the Vision Pro, and Google is back to the drawing board to contemplate how AI can better support a post-Google Glass headset.

RECOMMENDATIONS

Organizations can have developers become familiar with new spatial computing toolkits and standards. They can start building out digital models in useful formats and compiling a digital twin library as the business works to identify initial value cases for frontline workers. Don't rush to invest in expensive headset technology unless there's a clear value proposition, and only after approaching the concept using existing tools.

INFO-TECH RESOURCES

► DOUBLE YOUR ORGANIZATION'S EFFECTIVENESS WITH A DIGITAL TWIN

This research will help your organization understand what a digital twin is, including the unique characteristics of this transformative technology. Articulate both the value and constraints of digital twin technology. Formulate a use case and validate its alignment with your organization.

► INTO THE METAVERSE

Understand how Meta and Microsoft define the metaverse and the coming challenges that enterprises will need to solve to harness this new digital capability.

DON'T RUSH TO INVEST IN EXPENSIVE HEADSET TECHNOLOGY UNLESS THERE'S A CLEAR VALUE PROPOSITION, AND ONLY AFTER APPROACHING THE CONCEPT USING EXISTING TOOLS.

IMAGINE //

A cinematic scene from a utopian science fiction movie called Second Sight. Extreme long shot captures a person building a purple and silver holographic city in the desert on a bright and sunny day, in the style of dynamic action sequences, captured on Phantom High-Speed Camera

--s 250 --c 25 --ar 16:9

RESPONSIBLE AI

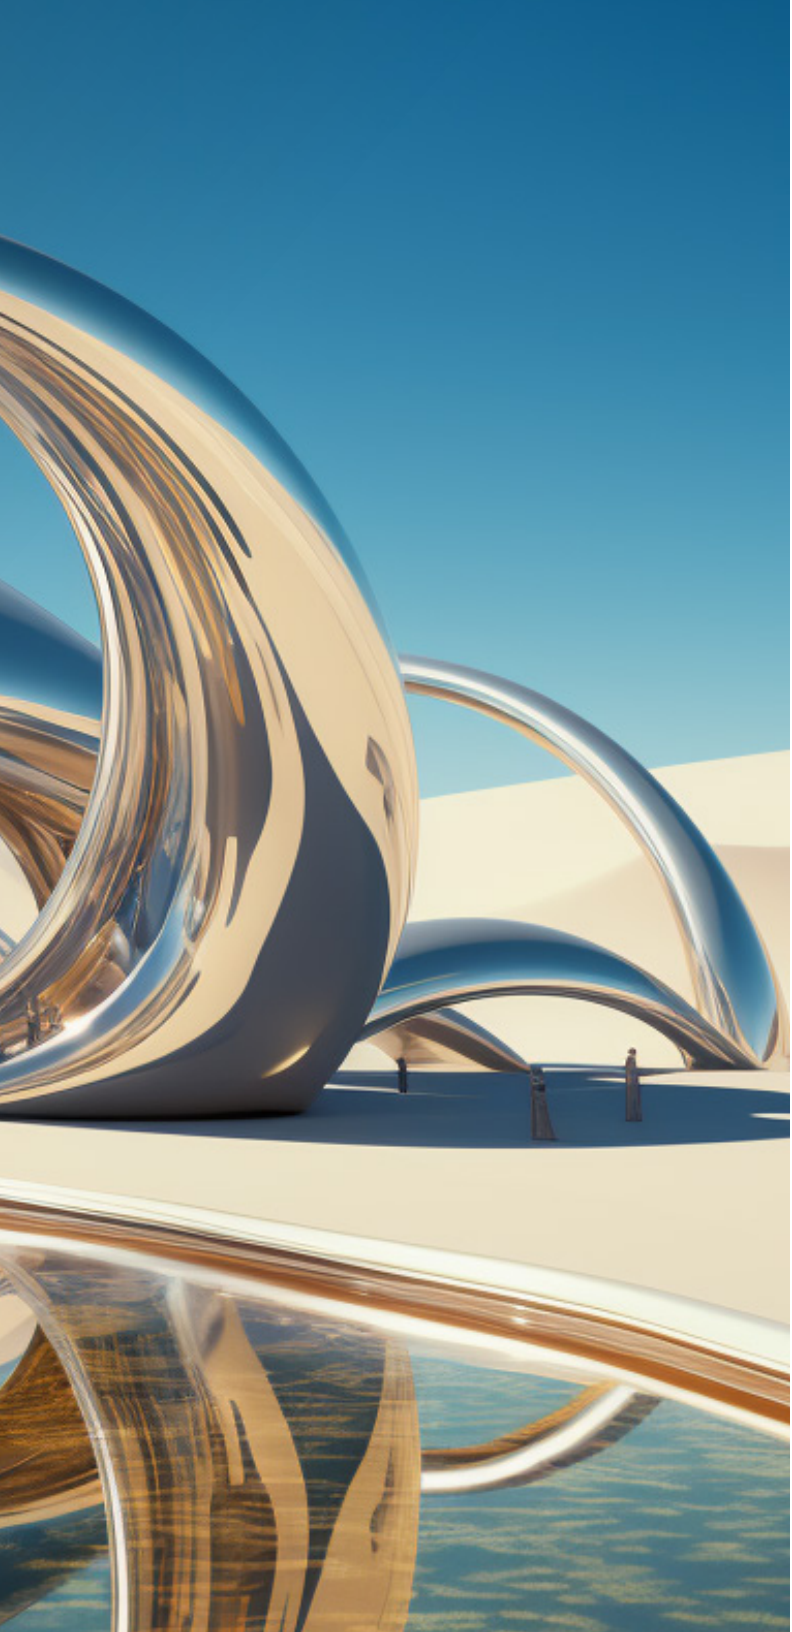
EFFECTIVE GOVERNANCE
TODAY TO AVOID COMPLIANCE
CHALLENGES TOMORROW

IMAGINE //

A cinematic scene from a utopian science fiction movie called Digital Protector. Long shot captures gold and silver holographic bridge, bright and sunny day, iridescent interface, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--c 50 --s 750 --ar 16:9

INTRODUCTION

MOVE FAST AND REQUEST REGULATION



Lawmakers around the world have typically been slow to respond to emerging technologies. In recent years, we've seen many examples of Silicon Valley's "move fast and break things" approach, where a company quickly swoops into a market and worries about sorting out the situation with lawmakers years after the fact. Such was the case with Uber usurping the taxi industry by allowing anyone with a car to use a mobile app to find customers who need a lift. Years after Uber was already in numerous major cities around the world, a few jurisdictions banned its ride-sharing business model, but most adapted their rules to make room for it. But with AI, it appears things might be different this time around.

Even some of the AI service leaders are inviting regulation, with OpenAI CEO Sam Altman, in an appearance before a US Senate committee in May, saying: "My worst fears are that we – the field, the technology, the industry – cause significant harm to the world. I think that can happen in a lot of different ways" (Time, 2023).

It's true that many AI services have already been commercialized and have found millions of users. But lawmakers around the world are also at work drafting. Fears of AI's misuse or neglect surround its potential to manipulate people's behavior using misinformation, to cause mass unemployment, or even to pose an existential threat to humanity. The stakes are clear, and governance policies based around responsible AI frameworks are taking shape. Hundreds of policy initiatives are in development around the world, according to the OECD AI policy tracker (OECD, 2023).

Europe is taking the lead. The European Commission brought draft legislation on AI a step closer to law in June, which would include a ban on using biometric identification tools, like facial recognition, in public places. It would also require that generative AI tools like

ChatGPT publish detailed summaries of the copyrighted data used for training their models ("MEPs Ready to Negotiate First-Ever Rules for Safe and Transparent AI," European Parliament, 2023).

In the US, the White House is moving toward an executive order and legislation on responsible AI. In the meantime, it secured a voluntary commitment from seven leading AI companies to ensure public safety and earn public trust. Measures include a method to watermark AI creations and self-reporting on AI systems' limitations and areas of inappropriate use (The White House, 2023).

Even after laws regulating AI are passed, there will be more time before responsible AI is regulated in an enforceable manner. But organizations looking to build or deploy AI should mitigate the risk of not meeting compliance requirements later by adopting responsible AI frameworks now.

SIGNALS

CIOs EXPECTED TO BEAR BURDEN OF AI GOVERNANCE

Who is going to be accountable for AI in the organization? This person will have a lot of responsibility as the regulatory space shifts from drafting policy to enforcing it. Even before that happens, there could be reputational fallout from working with AI vendors that are seen as unethical, or using AI when customers aren't expecting it.

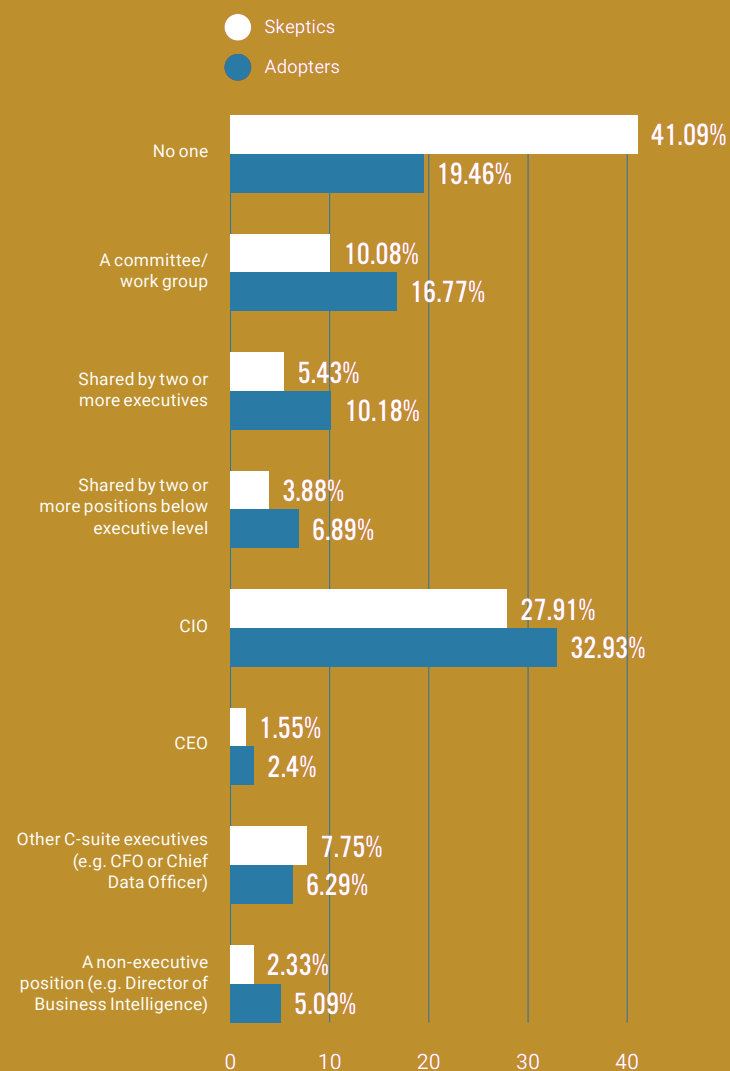
Among AI adopters, 1 in 3 say the CIO will be solely accountable for the governance of AI. Another 17% say that a committee or work group will be accountable, and another 10% say it's shared between two or more executives – either of these groups could also include the CIO. For 1 in 5 adopters, no one is responsible for AI governance yet.

Skeptics are twice as likely to say no one is responsible for AI governance, with 2 in 5 saying so. Even here, CIOs remain a popular choice at 28%. While AI skeptics may not be investing directly in AI, these firms may recognize that AI features will filter in through feature upgrades to software they already use through existing contracts, or they may want to set policy around what employees are allowed to do with consumer-facing and open-source AI tools. In these situations, CIOs will be contending with shadow AI in addition to the usual shadow IT concerns.

INSIGHT //
CIOs are likely to bear at least some accountability for AI governance in their organizations.

SURVEY

WHO IN YOUR ORGANIZATION IS ACCOUNTABLE FOR GOVERNANCE OF AI?



SIGNALS

TAKING STEPS TOWARD GOVERNANCE

Organizations deploying AI will be accountable for applying governance at various stages to prevent harm. Draft legislation in different jurisdictions requires actions to protect customer privacy, monitor model performance over time, and to explain when and how AI is being used.

Adopters are most likely to say they have no governance in place today (35%), but 1 in 3 say they are publishing clear explanations of how AI is intended to be used and what predictions it makes, as well as 1 in 3 saying they conduct impact assessments on AI systems. Predictably, most skeptics are taking no steps toward AI governance.

If draft legislation in different jurisdictions around the world holds up and is passed into law, organizations building and deploying AI will need to implement many or all of these steps, depending on the context of their use case.

INSIGHT //


Organizations are just beginning to implement real steps in pursuit of AI governance, with many doing nothing at present.

TRANSFORMERS //

Are most likely to say they are using AI models that are explainable (44%) as a step toward AI governance.

SURVEY

WHAT AI GOVERNANCE STEPS DOES YOUR ORGANIZATION HAVE IN PLACE TODAY?





IMAGINE //

A cinematic scene from a utopian science fiction movie called Digital Protector. Long shot captures gold and silver holographic bridge, bright and sunny day, iridescent interface, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--ar 16:9 --s 750

OPPORTUNITIES & RISKS

SEIZE OPPORTUNITIES & MITIGATE RISKS

OPPORTUNITIES

► CREATE AI THAT BENEFITS HUMANITY

Responsible AI is about more than just reducing compliance risk. It represents an effort to put human concerns first, including the most vulnerable among us, as industry pursues the next wave of technological advancement. Society has many problems that AI might help solve, and it is incumbent upon the industry to pursue those solutions as a motive alongside profits.

► AI THAT DELIGHTS

Generative AI capabilities can feel like magic to use: type in a few words and receive an image that you previously only could have imagined, or engage in a conversation with some of your favorite fictional characters. But executed the wrong way, AI will leave customers feeling duped or undervalued. Responsible AI creates a framework designed to keep people's experience with AI on the bright side.

► NARROW THE USE CASES

Large language models (LLMs) create risk because of their flexibility. Without being designed for any single, specific, clear purpose, these models leave room for malicious uses. Hence, ChatGPT can be used to brainstorm at your next marketing meeting as easily as it could be used to create malware (Ars Technica, 2023). Organizations that identify the use cases where they want to use AI can build for purpose and reduce the risks.

RISKS

► MODELS ARE DECLARED ILLEGAL

Models trained on data sets that aren't sufficiently documented or that can't be well explained risk more backlash from regulators. Given that models are expensive to build, it's best to use responsible AI approaches today so that models won't be hamstrung – or barred from the market entirely – in the future. Organizations adopting AI models built by third parties should ask for assurances of responsible AI to avoid potential service disruptions.

► MACHINES PASS THE TURING TEST

Present-day methods to separate bots from humans are no longer effective. Where Completely Automated Public Turing Tests to tell Computers and Humans Apart (CAPTCHAs) provided reliable tests in the form of distorted letters or image-selection grids in the past, new methods will be needed to block bots powered by LLMs. More abstract concept puzzles will be needed. One set of puzzles, ConceptARC, created by a team at the Santa Fe Institute, proves solvable for humans more than nine times out of 10, yet it stumps GPT-4 more than two-thirds of the time (Nature, 2023).

► SELF-REPLICATING AI

So far, AI models act when humans give them direction. The models don't have any agency to set their own goals and pursue them. But AI creators like former Google vice-president Geoffrey Hinton worry that capability could emerge in a large model, and once it does, AI would seek resources to self-replicate to different host locations in order to preserve itself. "They may well develop the goal of taking control – and if they do, we're in trouble," Hinton said on stage at Collision (University of Toronto, 2023). This would be a singularity-level event representing humanity's loss of control over AI.

IMAGINE //

A cinematic scene from a utopian science fiction movie called Digital Protector. Long shot captures gold and silver holographic bridge, bright and sunny day, iridescent interface, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--ar 16:9 --c 25 --s 750





We're providing an AI and analytics platform for people to develop responsible AI. We took the time to develop responsible AI principles and to take the terms and define what each of those terms means to us. The next step is to instantiate those terms within our culture ... So choose the language. But it's the culture that will shape it.

REGGIE TOWNSEND,
VICE-PRESIDENT OF
DATA ETHICS,
SAS

CASE STUDY

AI ETHICS ARE ABOUT HUMAN EQUITY

SITUATION

Business analytics and intelligence provider SAS launched its responsible innovation initiative, guided by the SAS Data Ethics Practice in May 2022. The initiative was a response to the recognition that AI offered many opportunities to augment human capabilities but also posed a risk if not deployed ethically.

The Data Ethics Practice is a cross-functional team that acts as a center of excellence for SAS employees and customers to use when deploying data-driven systems, ensuring they respect human wellbeing, agency, and equity. Vice-President of Data Ethics Reggie Townsend leads the team and is also appointed to the US Department of Commerce's National AI Advisory Council (SAS, 2022).

Now, more than a year after being appointed, Townsend sees that AI is already affecting daily life: providing navigation for driving, making decisions about loan applications, and even determining which candidate an employer will hire. There are opportunities to increase productivity and convenience, he acknowledges, but will some people be left out in reaping those benefits?

"Those who sit on the margins of society often don't get to participate in that," he says. "One of the great benefits that I think AI provides us is an opportunity to do just that and extend the beneficial opportunity to those within and well outside of the margins of our societies today" (Interview with Townsend).

Doing so will require the commitment of government, industry, and individuals.

ACTION

Townsend's team developed a trustworthy AI framework for SAS. It includes six principles:

- ▶ Human centrality
- ▶ Inclusivity
- ▶ Accountability
- ▶ Transparency
- ▶ Robustness
- ▶ Privacy & security

The model is consistent with Info-Tech's own responsible AI framework, he says. But what really matters is not what words are used to create the framework, but how they shape behaviors in an organization's culture. "Choose the language [for responsible AI], but culture ultimately will shape it," Townsend says. "Business leaders need to be able to translate the mission and vision to actual people doing things."

Townsend takes a pragmatic approach, thinking of what can practically be done and the inherent trade-offs that might be made in some situations. For example, SAS focuses on being able to explain how its models are making decisions, but this can sometimes conflict with the need to maintain privacy. So Townsend asks his team, "For whom might this fail? Who in this given scenario is the most vulnerable? If we calculate around that cohort of individual, that contextualizes which trade-offs we choose to make." Therefore, model explainability in a healthcare setting might look different than it does in a retail setting.

SAS calls its collaborative governance approach "QUAD," with a focus on providing oversight, a platform that aligns with data ethics principles, controls that provide checks and balances before new AI services are made public, and a culture that normalizes data ethics principles (SAS, 2023).

IMAGINE //

A cinematic scene from a utopian science fiction movie called Digital Protector. Extreme long shot captures a silver and gold floating holographic vault on a bright and sunny day, iridescent, rainbow, in the style of dynamic action sequences, captured on Phantom High-Speed Camera --s 750--ar 16:9 --c 25



INFO-TECH'S RESPONSIBLE AI GOVERNANCE PRINCIPLES

ACCOUNTABILITY

VALIDITY &
RELIABILITYFAIRNESS &
BIAS DETECTIONSAFETY &
SECURITY

EXPLAINABILITY

DATA
PRIVACY

RESULT

Townsend acknowledges that many organizations haven't taken any steps toward AI governance yet, but he isn't alarmed by it. His advice to customers includes to think clearly about how to activate that approach later, once a framework is in place. Establishing the framework requires top leadership's presence and involvement from across the organization. Next, think about where organizational data is stored and where AI models are being trained. Understand how that lines up with regulatory activity, and make sure employees comprehend it.

Ultimately, governance for AI must be distributed in an organization. Awareness and training at the individual level are going to be important, and CIOs may find they are responsible for enabling the organization to use AI in a responsible manner. When organizations are consuming AI as a service, CIOs will be accountable for that just as they are for other technology vendor contracts. But as is the case with shadow IT, when other lines of business invest in the technology, the accountability for it has to follow. "If AI is being used in hiring decisions, then HR probably should have some level of domain, expertise and accountability over that," Townsend says. "If you're an organization like SAS that's creating AI technologies, then a CTO may ultimately take accountability for that."

SAS uses a committee for AI governance composed of a cross-functional membership. It puts workflows in place that help guide responsible use of AI in the course of operations. Even though some regulations are uncertain, Townsend's team is taking precautionary measures such as creating an inventory of where models reside in the business. It's also pushing awareness of trustworthy AI out to the employee base in order to avoid technical debt being created with any models built today that might find themselves out of compliance tomorrow.



IMAGINE //

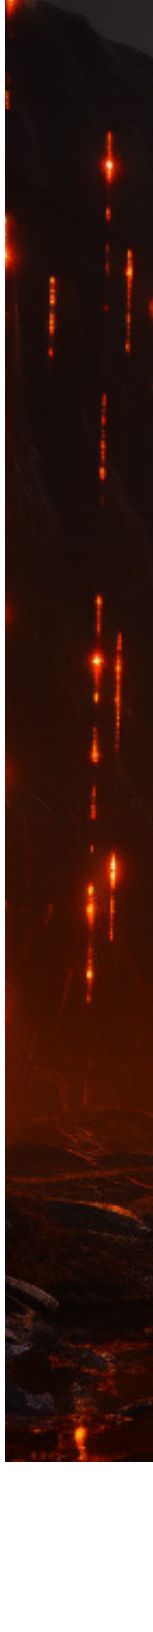
A cinematic scene from a utopian science fiction movie called Digital Protector. Long shot captures gold and silver holographic bridge, bright and sunny day, red and orange accent lighting, iridescent interface, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--ar 16:9 --c 25 --s 750

WHAT'S NEXT

PIRACY ENTERS THE AGE OF AI

PIRATE MODELS

As AI developers and regulators reach a consensus on responsible AI frameworks and implement more governance measures around commercial models, it's conceivable there will be an underground market of non-compliant models that are built by criminals for criminals. The illegal models could be those trained on copyright-protected or illicit materials, or the models could be released without safety constraints. Imagine having access to a variant of ChatGPT that could instruct you on how to make a weapon or program ransomware. Criminals interested in building such models will be constrained by the same factors as law-abiding organizations – access to GPUs, the availability of data science talent, and access to the massive amount of data needed for training. For this reason, it's more likely that we'll see illegal variants of open-source models such as those released by Meta, or even variants of stolen proprietary models. Criminals will have an easier time manipulating these foundation models to remove safeguards or add additional context training on new data sets. Lawmakers will have to determine how to respond with sufficient penalties for creating and wielding pirate models.





WAITING FOR ENFORCEMENT MECHANISMS

It's likely that in 2024, many jurisdictions will pass new laws outlining requirements for AI development and use. But it could take time before those laws become regulations with accountable bodies to enforce the new measures. In the meantime, organizations or individuals harmed by AI will have little recourse except to try the case in court or publicly shame the offender. The AI Incident Database keeps a public record with documented evidence of AI offenders and victims for the purposes of informing research aimed at avoiding bad outcomes (AI Incident Database, 2023). But companies that feature high on the leaderboard of incidents may be subject to public backlash. The European Commission estimates that the second half of 2024 is the earliest its AI regulations could be applied to operators, with algorithms applied to conformity assessments (European Commission, 2023).

IMAGINE //

A cinematic scene from a utopian science fiction movie called Digital Protector. Long shot captures evil metal mainframe computer, desert, red and gold accent lighting in the style of dynamic action sequences, captured on Phantom High-Speed Camera

--s 750 --c 10 --ar 16:9

RECOMMENDATIONS

Organizations should build, design, and deploy AI responsibly today to avoid legal consequences tomorrow. In the meantime, operating with responsible AI frameworks will protect organizations' reputation and ensure better outcomes for people affected by AI decisions. The seven principles of our responsible AI model address three main concerns:

- ▶ **TRUSTWORTHINESS**
Privacy is preserved both in training the model and in its outputs, and human safety and security are at the forefront.
- ▶ **EXPLAINABILITY**
To the greatest degree possible, it can be explained how the model is making its predictions. Fairness and bias analyses are conducted to mitigate unfair outcomes, especially for more vulnerable populations.
- ▶ **TRANSPARENCY**
It's clear what decisions the AI is making and what purpose they serve, as well as who is accountable for those decisions. Monitoring of model validity and reliability is conducted and reported.

INFO-TECH RESOURCES

- ▶ **BUILD YOUR GENERATIVE AI ROADMAP**
Tailor Info-Tech's responsible AI foundational principles for your own organization as you pursue building a tactical roadmap for generative AI.
- ▶ **WEBINAR: ESTABLISH YOUR RESPONSIBLE AI GUIDING PRINCIPLES**
Bill Wong and Logan Rohde discuss how to build responsible AI around six guiding principles, with an overarching focus on data privacy.
- ▶ **IMPLEMENT AND OPTIMIZE APPLICATION INTEGRATION GOVERNANCE**
Assess your capabilities and determine which area of governance requires the most attention to achieve success in AI using Info-Tech's governance gap analysis tool.





BUILD, DESIGN, AND DEPLOY AI RESPONSIBLY TODAY TO AVOID LEGAL CONSEQUENCES TOMORROW.

IMAGINE //

A cinematic scene from a utopian science fiction movie called Digital Protector. Long shot captures gold and silver holographic bridge, bright and sunny day, iridescent interface, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High-Speed Camera --ar 16:9 --c 25 --s 750

SECURITY BY DESIGN

BUILDING TECHNOLOGY
THAT CAN'T BE BROKEN

IMAGINE //

A cinematic scene from a science fiction drama movie called Lock and Key. Long shot captures a monolithic heavily guarded pyramid on a bright and sunny day, accent lighting, in the style of dynamic action sequences, captured on Phantom High-Speed Camera

--s 750 --ar 16:9

INTRODUCTION

CYBERSECURITY'S SEATBELT MOMENT

It wasn't a legal requirement for automakers to include seatbelts in their cars until 1968 in the US. Before that, the safety feature was optional – an unacceptable scenario for modern motorists. Lawmakers had to put the onus on automakers to make their products safe.

The standards for the software industry are the opposite. Instead of the product builders being accountable for their unsafe products, the users are accountable for mitigating the risks. As a result, programming languages with known vulnerabilities continue to be used to design new products, with security patches rolled out on a regular schedule to plug the holes as they're discovered. IT operations personnel know what to expect on "Patch Tuesday." And because zero-day attacks are possible – in which an undiscovered vulnerability is exploited – and not all patches are deployed in time, organizations must also invest their own resources into cybersecurity. New solutions and services come to market every year to mitigate cyber risks. IT specialists build their careers on security skillsets, and organizations purchase cyber insurance to help them cover the costs if they fail.

But regulators are looking to change that bargain. The White House's National Cybersecurity Strategy seeks to "drive prioritization of cybersecurity as a fundamental safety issue and ask more of the technology providers to build security into products throughout their lifecycle, ship products with secure defaults, and foster radical transparency into their security practices" ("CISA Cybersecurity Strategic Plan," CISA 2023).

The new risks posed by AI are driving regulators to put the magnifying glass on vendors bringing it to market. The Federal Trade Commission (FTC) launched an investigation into OpenAI in July, with a focus on cybersecurity practices. Included in the FTC requests of OpenAI are:

- ▶ All instances of known and attempted "prompt injection attacks."
- ▶ Details on how and why personal information is collected, used, analyzed, stored, and transferred.
- ▶ Details on data retention, data deletion, and deidentification practices.

("Generative AI's 'Industry Standards,'" JD Supra, 2023.)

The requests reflect the new cybersecurity risks posed by adversarial AI that the FTC is considering. Prompt injection attacks are attacks where a malicious actor provides instructions to a model that succeed at convincing it to disregard safety mechanisms or to provide sensitive data hidden within the model. This type of attack is limited in its scope of harm, as it affects only one user session with the model (though it may lead to a leak of valuable data), while model poisoning attacks threaten to affect the experience of all users. Model poisoning is a type of supply chain vulnerability attack in which a malicious actor injects malicious code or training data into an open-source model available in an AI model marketplace. Users leveraging the model would then receive biased or false outputs. For example, a hacker might set up a crypto coin and have an AI model advise investing in it, or state-backed cyber-warfare group might plant false information about an enemy government (Mithril Security, 2023).

AI brings new threat vectors with familiar themes. Without addressing the problem at its root, enterprises will continue to sink more investment and resources into cybersecurity.

SIGNALS

HIGHER SPENDING BY DESIGN

Another annual IT budget, another increase in the amount spent on cybersecurity. Most IT leaders say they expect to increase their spending in 2024, whether they are adopters or skeptics. Adopters are slightly more likely to increase their budgets by more than 10%, with 1 in 5 indicating this compared to less than 1 in 8 skeptics.

Top priorities for cybersecurity investment are different for adopters and skeptics. Adopters rate security awareness and training of their own staff as the most important area to invest in, rating it an importance of 4.3/5 on average. Skeptics see third-party services (such as 24/7 intrusion detection) as the top priority, with a rating of 4.1/5 on average. Still, adopters and skeptics aren't that far apart on the importance of spending on cybersecurity across different priorities, with the biggest gap between them being a difference of 0.6/5 on the security awareness priority.

Adopters and skeptics also agree on the second-highest priority, next-generation tools such as security automation and network detection and response.

Whatever their cybersecurity spending priority, it is organizations, not their technology providers, that will be accountable for the outcome of their efforts to secure proprietary and sensitive data. For the time being,

INSIGHT //


A majority of organizations plan to spend more on cybersecurity in 2024, with more than 1 in 6 organizations planning to increase their cybersecurity budget by more than 10%.

TRANSFORMERS //

Are more likely than other organizations to say they will not increase spending on cybersecurity, with a little less than half saying they will spend "about the same."

SURVEY


FOR THE NEXT FISCAL YEAR, HOW DO YOU ANTICIPATE YOUR ORGANIZATION'S SPENDING ON CYBERSECURITY WILL CHANGE COMPARED TO THE PREVIOUS YEAR?



SURVEY

HOW IMPORTANT IS EACH OF THE FOLLOWING AREAS AS A CYBERSECURITY SPENDING PRIORITY FOR THE ORGANIZATION?

Rated from 1 to 5, where 1 is the least important and 5 is the most important.



INSIGHT //
Adopters rate security awareness and training as the most important spending priority, but skeptics' top priority is third-party services.

WHATEVER THEIR CYBERSECURITY SPENDING PRIORITY, IT IS ORGANIZATIONS, NOT THEIR TECHNOLOGY PROVIDERS, THAT WILL BE ACCOUNTABLE FOR THE OUTCOME OF THEIR EFFORTS TO SECURE PROPRIETARY AND SENSITIVE DATA. FOR THE TIME BEING.

IMAGINE //

A cinematic scene from a science fiction drama movie called Lock and Key. Extreme long shot captures a celestial beacon in the desert, blue accent lighting, in the style of dynamic action sequences, captured on Phantom High-Speed Camera

--c 50 --s 750 --ar 16:9



IMAGINE //

A cinematic scene from a science fiction drama movie called Lock and Key. Extreme long shot captures a row of centurions on a monolithic structure on a bright and sunny day, blue and gold accent lighting, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--s 750 --c 10 --ar 16:9

OPPORTUNITIES & RISKS

SEIZE OPPORTUNITIES & MITIGATE RISKS

OPPORTUNITIES

► ADOPT AI WITH SECURITY AS A CORE PRINCIPLE

Generative AI capabilities are bringing new IT infrastructure along with them and may introduce many organizations to new cloud service models or database types. Designing this infrastructure presents an opportunity to build in security and resilience from the outset, avoiding more investment of resources to provide just-in-time security down the road.

► ENHANCE SECURITY WITH GENERATIVE AI

Threat detection is an exercise in monitoring a system and determining when a threat emerges based on a pattern of unusual activity or behavior. Generative AI can detect this in broad and flexible ways that weren't previously possible and could help enhance intrusion detection and monitoring (Analytics Insight, 2023).

► IMPROVE TRAINING AND AWARENESS

Generative AI can be used to aid tabletop security exercises or allow users to ask questions about new threats and determine what risks they pose to a specific environment.

RISKS

► PROMPT INJECTION ATTACKS

Prompt injection attacks seek to bypass a model's safety mechanisms in order to use it for malicious activity or to reveal sensitive information. As LLMs are equipped with more plug-ins to extend their capability, the risk of this type of attack increases. Organizations may want to code their own plug-ins to ensure that the least amount of privilege required is used in requesting data. All prompts should be treated as potentially malicious and be inspected and sanitized before extracting information (NVIDIA Developer, 2023).

► MODEL INVERSION AND DATA EXTRACTION ATTACKS

Model inversion and data extraction attacks seek to determine what training data went into a model, or what data is being used to customize its outputs. Organizations must treat models the same way they would restrict access to their training data when it comes to access privileges. Organizations customizing models with their own data should use encryption ("Software Must Be Secure by Design," CISA, 2023).

► MODEL POISONING

Model poisoning, or supply chain poisoning, seeks to force a model to provide a false output or to ignore specific inputs. For example, adversarial inputs to a self-guided car algorithm might cause it to crash, or given to security camera software might cause it to not detect certain objects. Models must be monitored after deployment to ensure they are still performing as expected ("Software Must Be Secure by Design," CISA, 2023).



IMAGINE //

A cinematic scene from a science fiction drama movie called Lock and Key. Extreme long shot captures a row of centurions on a monolithic structure on a bright and sunny day, blue and gold accent lighting, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--s 750 --c 10 --ar 16:9



We're seeing this huge shift from where all the sensitive data was pouring into the model itself to where the model is not the sensitive piece, it's the inputs and outputs from the model that are sensitive ... and the most interesting output there is are these vectors. They're like the memory of AI, all these vector databases.

PATRICK WALSH,
CEO, IRONCORE LABS

CASE STUDY

SECURING THE MEMORY OF AI

SITUATION

Large language models are flexible and adaptable, but enterprises still want to customize them using their own data to get the most value. To provide this, generative AI models integrate their foundation model, which is informed by its training data, with a vector database that functions as the working memory and stores custom data inputs. These custom data inputs are called “embeddings” and are stored as vectors, which are a string of numbers that are incomprehensible to humans, but which represent the meaning of the input to the machine. Vector databases can be used to encode different types of inputs, from text to audio recordings to images including people’s faces. They are such an effective means of capturing meaning and creating new relevant outputs that are customized to the users that it has become the best-in-class way to provide recommendations.

For example, vector databases are used by:

- ▶ Spotify for its song recommendations
- ▶ YouTube for video recommendations
- ▶ Pinterest for visual search
- ▶ Netflix for program recommendations
- ▶ Google for semantic search

Although vector databases store information in a way that a human couldn’t comprehend, that doesn’t mean they are secure by default. Adversarial AI techniques include an embedding inversion attack, which can translate embeddings back into their source data. While the output may not always be exactly the same as the input data, it’s often close enough that it would be considered a security compromise for sensitive information. Researchers have demonstrated that without knowing anything about how a model works or what its data inputs were, an adversarial model can be trained to reproduce the original inputs (HackerNoon, 2023).

Often, organizations store original inputs alongside vector embeddings in an effort to improve the accuracy of AI outputs, explains Patrick Walsh, CEO of IronCore Labs. “One of the biggest problems with AI today is hallucination. Of the techniques available to ground these foundation models, the leading one is called RAG, retrieval augmented generation, which grabs information relevant to the query and feeds it into the prompt,” he says. “This leads people to store their sensitive data alongside the vector representation.”

This further exposes it to hackers, who may find it easier to get access to the model and retrieve sensitive data than to bypass other security measures. No major breach incidents in the real world have come to light yet, but Walsh isn’t waiting to provide a solution to what he sees as a huge gap in the market.

IMAGINE //

A cinematic scene from a super hero movie called Digital Protector. Long shot captures gold and silver holographic super hero, justice, bright and sunny day, iridescent interface, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High - Speed Camera
--ar 16:9 --c 50 --s 750

**ACTION**

IronCore Labs is bringing its property-preserving, data-in-use encryption to vector databases, so customers can encrypt the vectors they store in the memory of AI. Using a public key/private key method of encryption, customers decode their vector embeddings as they are processed in order to see the information, while hackers retrieving the data would see incomprehensible strings of letters and numbers. Cloaked AI protects the privacy and security of vector database users.

By applying the decryption to the results and not the processing of the information, Cloaked AI can be used with vector databases whether they are hosted and configured by a cloud provider (SaaS), configured by the organization and hosted in the cloud (IaaS), or configured and hosted on-premises. The data-in-use model of encryption allows data to stay encrypted while processed on hosted computers, and it allows property checks on the data as if it were not encrypted. The overall result is that the encryption is invisible to the data owner but blocks access by anyone else.

“You can plug and play with other systems here,” Walsh says. “If you’re querying for embeddings instead of completions, probably you’re going to use a hosted vector database because running your own vector database is a nightmare. And if you want to put it in a hosted database, then the thing you do is encrypt it before you send it there, but you can still use it. You can query over it with an encrypted query, get encrypted results back, and then decrypt it when it gets back to you.”

IronCore is working with about a dozen clients that have signed up for its beta version. Cloaked AI will be useful to an enterprise that wants to customize an LLM using its own data.

RESULT

Walsh sees an “enormous” market for private AI security products: “All of the interesting use cases for AI are over private data – over your health data, your financial data, your personal documents, etc. That’s where things go off the rails.”

Walsh points to data from market analysis that indicates CIOs are hesitant to pursue AI projects due to data privacy concerns. Info-Tech’s survey data also supports that, with 7 in 10 organizations saying they have only progressed as far as exploring what’s possible with AI, or haven’t even explored it yet. Confidence with data security may be holding them back from launching pilot projects or integrating AI into operations.

Some organizations interested in Cloaked AI so far include a group interested in using generative AI for personal injury law and a developer building an app to share audio diaries with therapists. He also expects it will be of interest to financial services and healthcare companies. “It’s not just the classically security-sensitive folks, but people who are trying to do stuff and trying to figure out how to do it right from the get-go. That’s where the market is.”

IronCore is supporting Cloaked AI and its other security products for AI with reference architectures that show how to deploy security by design, building out infrastructure for a resilient AI capability from the beginning. A new set of infrastructure represents another chance for businesses to put best practices for security in place from the outset. And generative AI is less entrenched in existing infrastructure than other enterprise technology.

Walsh says Cloaked AI will be generally available in Q4 of 2023.

IMAGINE //

A cinematic scene from a super hero movie called Digital Protector. Long shot captures gold and silver holographic super hero, justice, bright and sunny day, iridescent interface, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High - Speed Camera
--ar 16:9 --c 50 --s 750



WHAT'S NEXT

US NATIONAL CYBERSECURITY STRATEGY

The Cybersecurity & Infrastructure Agency (CISA) will be pursuing the *CISA Strategic Plan 2023-2025*, its first strategic plan since it was established in 2018, and driving security-by-default is a core initiative. CISA plans to support this initiative by developing network defense and cyber operations tools, services, and capabilities. It will provide support for the national cyber workforce to fill shortages in critical skills through educational resources, and it will prioritize security with technology builders. "Technology products must be designed and developed in a manner that prioritizes security, ensures strong controls by default, and reduces the prevalence of exploitable vulnerabilities," the strategy states. CISA will "measure the adoption and effectiveness of secure development practices and control adoption for technology products and services" (CISA, 2022).


FTC'S APPROACH WITH AI LEADERS

The investigation into OpenAI's ChatGPT product is likely to reveal a lot more information than is currently publicly known about how OpenAI trained its large language models and to what degree it is the target of malicious attacks. The investigation could see fines imposed or a court order that would put a consent decree in place, detailing a plan on how customer data is used. Other leading AI vendors will be watching to see how the FTC approaches enforcement of current consumer protection regulations ("New FTC Investigation," JD Supra, 2023).

An FTC blog post following FTC statements about a complaint regarding Amazon's Alexa service may provide insight into what the FTC will be evaluating in the case of OpenAI. "The FTC will hold companies accountable for how they obtain, retain, and use the consumer data that powers their algorithms," the blog states. "Machine learning is not a license to break the law" (FTC, 2023). In July 2023, Amazon agreed to a permanent injunction and a \$25 million civil penalty as part of a settlement following complaints that Amazon retained children's voice recordings indefinitely by default, in violation of a requirement that they be retained only as long as necessary to fulfill the purposes for which they were collected (DoJ, 2023).

IMAGINE //

A cinematic scene from a utopian science fiction movie called Digital Protector. Long shot captures gold and silver holographic bridge, bright and sunny day, iridescent interface, rainbow, happy, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--ar 16:9 --c 50 --s 750



RECOMMENDATIONS

SECURITY BY DEFAULT IS NOT A CHOICE

Organizations that are building AI models or customizing foundational models for their own use can't afford to ignore security. The ramifications of malicious use and potential harm to people in the process are too high. Security must be built in by default, ensuring data used in training and directing the models won't fall into the wrong hands. Introducing secure concepts at the outset of this new wave of AI capabilities may be a last chance for organizations to break the cycle of increasing security investments year after year, yet always facing more risk imposed upon them by technology. It may also be a matter of regulatory necessity, depending on how authorities determine enforcement actions in 2024.

INFO-TECH RESOURCES

- ▶ **DEMONSTRATE DATA PROTECTION BY DESIGN FOR IT SYSTEMS**


Turn abstract principles like privacy into functional ways of working in your organization, using data protection by design in your IT systems to determine what controls are necessary at every step of the data lifecycle. Lay the foundation for a full-fledged privacy program.

- ▶ **DEVELOP A SECURITY OPERATIONS STRATEGY**

Move from a center of security operations to an ongoing process that combines four critical functions: prevention, detection, analysis, and response. Use functional threat intelligence to inform incident response and align with the business and collaborate across the organization to create a transparent security process.

- ▶ **SECURE YOUR HIGH-RISK DATA**

Protect data throughout its entire lifecycle and use a multi-layered defense across all data sources to meet compliance obligations and secure the business. Be prepared to secure sensitive data wherever it resides, from on-premises servers to cloud environments.





**AI BRINGS NEW THREAT VECTORS
WITH FAMILIAR THEMES. WITHOUT
ADDRESSING THE PROBLEM AT ITS
ROOT, ENTERPRISES WILL CONTINUE
TO SINK MORE INVESTMENT AND
RESOURCES INTO CYBERSECURITY.**

IMAGINE //

A cinematic scene from a science fiction drama movie called Lock and Key. Long shot captures a monolithic heavily guarded pyramid on a bright and sunny day, accent lighting, in the style of dynamic action sequences, captured on Phantom High-Speed Camera

--s 750 --ar 16:9

DIGITAL SOVEREIGNTY

CONTROL YOUR
OWN DESTINY IN
THE DATA-DRIVEN
FUTURE

IMAGINE //

A cinematic scene from a cyber punk movie called the Network.
Extreme long shot of people breaking into a vault, in the style
of dynamic action sequences, captured on
Phantom High-Speed Camera
--ar 16:9 --c 50 --s 250

INTRODUCTION

STARTING OFF ON THE WRONG FOOT

In the early days of cloud computing, many enterprises were wary of the new infrastructure model because of the perceived risk of storing their data on the servers of a third party. With the tech giants offering cloud services seemingly willing to compete in many different business categories, there was always a nagging concern that they'd peek at customer data to draw insights for their own competitive efforts.

Cloud providers have worked for years to mitigate those concerns by setting up regional data centers closer to their customers around the world and offering encryption to keep customer data private. More recently, confidential computing in the cloud even offers encryption of data while it's being processed, providing assurances to more sensitive users.

The same concerns over data confidentiality are not only being sparked by vendors offering generative AI services, but they are fanning the flames by revealing that they've scraped the web to compile massive data sets to train their models. OpenAI detailed how it used Common Crawl, a data repository created by a nonprofit that intended to provide a resource to researchers, to train its model. Common Crawl contains more than 240 billion web pages spanning 16 years and claims to be the primary training corpus for every LLM (Common Crawl, 2023).

Information that companies exposed to search engines has been scooped up and used to create generative AI engines that can now directly provide the answers that many companies hoped to lure customers with. While some companies are now taking measures to block search crawlers from scanning their websites, the training on their historical data is already done (The Hacker News, 2023).

While chatbots were trained on a huge portion of the web's historical text, image-generating bots are benefiting from the wealth of images there. Stability AI built Stable Diffusion, an image-generating model that is now commercialized in several different applications, with open image data sets from the non-profit LAION. The data sets pair images with English-language descriptions, which is key to how Stable Diffusion's model is trained to produce new images based on text prompts. Significant portions of images come from Pinterest; WordPress-hosted blogs; and blogging or art sites including SmugMug, Blogger, Flickr, and DeviantArt. Shopping sites also contributed a large portion of the images, including Fine Art America, Shopify, Wix, Squarespace, and Etsy. Finally, stock image sites represent another major source, with Adobe Stock, PhotoShelter, iStock, Unsplash, Getty Images, and Shutterstock all represented in the data set (Waxy.org, 2022).

In our [2021 Tech Trends report](#), the "Self-Sovereign Cloud" trend looked at balancing the capabilities of the public cloud with the control and privacy of on-premises infrastructure.



IMAGINE //

A cinematic scene from a cyber punk movie called the Network. Medium shot of a cyborg robot brain in the witness stand on trial in a courtroom, in the style of dynamic action sequences, accent lighting, captured on Phantom High-Speed Camera
--s 250 --c 50 --ar 16:9

STARTING OFF ON THE WRONG FOOT (CONT'D)

With both text-based and image-based generative AI tools trained on copyrighted work and often creating output that is very similar to it, several different lawsuits have been filed that could have major impacts on the field of generative AI. Here are just a few examples:

- ▶ A class-action lawsuit against GitHub, Microsoft, and OpenAI targets the GitHub Copilot tool. Coders say that Copilot is copying and republishing code without attribution. That's against the GitHub open-source license. Microsoft and GitHub have tried to have this case dismissed but weren't able to and will face the allegations in court.
- ▶ A lawsuit against Stability AI, Midjourney, and DeviantArt alleges these companies scraped the web and infringed on artists' copyrights by training their AI models.
- ▶ Getty Images filed a copyright complaint against Stability AI for allegedly copying and processing millions of its images and metadata.
- ▶ Authors Paul Tremblay and Mona Awad are suing OpenAI for allegedly infringing on authors' copyrights. The suit estimates more than 300,000 books were copied in OpenAI's training data.
- ▶ Sarah Silverman is suing Meta and OpenAI, claiming that their large language models illegally acquired data sets that included her work (TechTarget, 2023).

Creators are leading the charge in pushing back against AI vendors because they have the most to lose, but all organizations are taking note. Recent research by BlackBerry shows that three-quarters of organizations worldwide are currently considering or implementing bans on ChatGPT and other generative AI applications in the workplace. The potential risk to data security and privacy is cited as the biggest concern, with 67% citing it (BlackBerry, 2023).

While courts and law makers catch up with the new capabilities of generative AI to synthesize large volumes of information and harness it for outputs, creators and companies are wondering how to protect their data and key aspects of their digital identity. There are suddenly new incentives to set up infrastructure and deploy protections that preserve your digital sovereignty and prevent third-party AIs from training on your data.

SIGNALS

SENSITIVE DATA FEARS CHILL AI TOOL ADOPTION

ChatGPT reached 1 million users faster than any other technology before it, but organizations are mostly trying to dissuade their employees from using the tools. This is one area where adopters and skeptics both agree that caution is the best course of action. Two-thirds of skeptics ask their employees to wait for professional tools with oversight to be deployed, and 56% of adopters do the same.

Adopters are three times more likely than skeptics to have identified a couple of clear use cases for third-party generative AI tools, with 30% saying they've done so. On the other side of the spectrum, skeptics are three times more likely than adopters to avoid the tools altogether and instruct employees not to use them, with nearly 1 in 5 taking this stance.

OpenAI understood that organizations exploring ChatGPT were held back by concerns about losing control over their intellectual property and sensitive data. Since such data might be included in prompts, enterprises quickly moved to create policies restricting employees from using such tools for work. In May, OpenAI CEO Sam Altman acknowledged in an interview that "customers clearly want us not to train on their data." OpenAI changed its terms of service to state it would not use data from its APIs for training, but it left the door open to use ChatGPT inputs for training at that time (CNBC, 2023).

In August 2023, OpenAI released an enterprise version of ChatGPT, which promised not to use customer data for model training, provided access to the latest version of GPT-4, and offered more capability to provide customization and context. The firm realized the demand for a professional-grade service after seeing professionals register for the consumer app. OpenAI reported that over 80% of Fortune 500 companies had created accounts on the consumer-grade release of ChatGPT in the nine months since its launch, based on accounts registered with corporate email domains ("Introducing ChatGPT Enterprise," OpenAI, 2023).

INSIGHT //


Regardless of their plans to invest in AI, the majority of organizations are waiting for professional generative AI tools with oversight before employees are approved to use them.

TRANSFORMERS //

Half say they are also waiting for more professional AI tools with oversight. Seventeen percent have already integrated AI into their business or are transforming their business with generative AI.

SURVEY

WHICH OF THE FOLLOWING BEST DESCRIBES YOUR ORGANIZATION'S APPROACH TO THIRD-PARTY GENERATIVE AI TOOLS (SUCH AS CHATGPT OR MIDJOURNEY)?



IMAGINE //

A cinematic scene from a cyber punk movie called the Network. This extreme long shot captures a network of floating metal orbs with a force field around them in the sky, bright and sunny day, blue and orange accent lighting, in the style of dynamic action sequences, captured on Phantom High-Speed Camera --s 250 --c 100 --ar 16:9

OPPORTUNITIES & RISKS

SEIZE OPPORTUNITIES AND MITIGATE RISKS

OPPORTUNITIES

► DEPLOY AI TO YOUR OWN INFRASTRUCTURE

With open-source models available and integration partners also offering help bringing AI models on-premises, many enterprises are choosing to avoid the risk of their data falling into the hands of third parties by operating models on their own infrastructure. This can also improve performance, where deploying a model close to the edge will allow it to react to different contexts more quickly.

► MONETIZE DATA STORES

In light of the need to train large AI models on various types of data, it's possible that your data stores hold new value for AI builders. While considering data sensitivity and privacy always remain paramount, organizations may be able to license their data to earn revenue.

► BUILD TRUST WITH VENDORS

For many organizations, some aspects of AI service delivery will take place on a third party's servers. Create trusted relationships with vendors by taking control over contractual language about how your data is used, and seek additional practical measures, such as isolated infrastructure and encryption, to ensure it.

RISKS

► **ADOPT PROTECTIVE MEASURES AGAINST AI**

The capability of AI model training requires that organizations reevaluate what data they are exposing publicly and where they can deploy new measures to protect that data from being ingested into training. Protective measures will range from changing configurations to adopting new preventive security measures.

► **PREPARE FOR CLOSE SCRUTINY ON CUSTOMER DATA**


Regulators are just beginning to consider the implications of AI capabilities for data privacy laws. In the meantime, data ethics has proven to be fertile ground for activism, and companies that act too aggressively may face a reputational backlash.

► **SEEK OUT YOUR DATA IN MASSIVE DATA SETS**

Is your data sitting in a massive data set like Common Crawl that is used to feed into the large AI models of different vendors? Seeking to have it removed where possible can further reduce the risk of your intellectual property moving beyond your control.

IMAGINE //

A cinematic scene from a cyber punk movie called the Network. This extreme long shot captures a network of floating metal orbs with a force field around them in the sky, bright and sunny day, blue and orange accent lighting, in the style of dynamic action sequences, captured on Phantom High - Speed Camera
--s 250 --c 25 --ar 16:9





Copyright law is created to incentivize human creativity. Clearly right now, AI is not doing that. It's hurting human creativity.

SHAWN SHAN,
GLAZE PROJECT LEAD,
UNIVERSITY OF CHICAGO
DEPARTMENT OF COMPUTER SCIENCE

CASE STUDY

WHEN IMITATION ISN'T A FORM OF FLATTERY

SITUATION

Since AI image generators emerged onto the market and allowed users to create images based on text prompts, tools like Stable Diffusion and Midjourney have become very popular. Some platforms count millions of users and produce hundreds of thousands of images every day. Many of these images mimic the unique style of established artists. This has caused a lot of harm to artists, as documented by the University of Chicago Department of Computer Science's Glaze project in a study of 1,000 artists. The research finds that AI mimicking artists in this way sabotages the artists' business model by displacing the original art in search results and demoralizes art students who see future career paths eliminated. In a survey of 1,207 artists, they found 97% of artists expect AI mimicry will decrease their job security. More than half said they plan to reduce or remove online artwork or have already done so (Glaze, 2023).

The Glaze team conducted its study on the harms to artists because they were being contacted by artists looking for help with protecting their work. Shan's team created the Fawkes privacy protection tool in 2020, which added minute changes to users' social media images to prevent them from unauthorized use by facial recognition systems. Artists wanted to know if the tool could protect their artwork from AI mimicry.

"At the moment we were not really big into that space, so we started talking to artists about the ways AI was affecting their life. And we said, 'OK this is big. It's big and it's really impacting people's lives,'" says Glaze project lead Shawn Shan. "So we started applying the Fawkes technique to these types of images" (Interview with Shawn Shan, 2023).

From there they developed Glaze, a tool that protects artists from AI style mimicry perpetrated by text-to-image generators.

IMAGINE //

A cinematic scene from a thriller movie called *Hidden in Plain Sight*. Medium shot of Mona Lisa in witness protection, in the style of dynamic action sequences, captured on Phantom High-Speed Camera --ar 16:9 --s 250 --c 50



ACTION

Glaze adds a cloak to images that is nearly imperceptible to the human eye but makes AI models see the images entirely differently. When it's applied to artworks, AI can still understand the content of an image (e.g. a woman wearing a dress surrounded by birds), but it does not accurately reproduce the style. Glaze makes just enough changes to effect style transformations informed by the Stable Diffusion model, causing a painting created by Karla Ortiz to be seen as a painting in the style of Picasso or Van Gogh.

"Since we have the Stable Diffusion model, any work we protect probably is the most optimal on that model because we can optimize the protection on it," Shan says. "But we see that the model has very high transferability in the sense that these models were trained to do the same stuff. So we see the same protection basically work across all different models that we test."

Glaze also found artists usually judged the technique to be effective even when the AI model was trained with a data set that was only 25% cloaked images and 75% uncloaked images, implying that artists don't necessarily have to ensure all of their online images are cloaked.

Glaze worked with artists to find the right level of cloak that would be effective at protecting an image while not changing the appearance of the image for people. Once the strength of the cloak reaches a high enough level, some pixelation can become visible to users.

Glaze was released as a free application download for Windows and Mac in March 2023. Shan and team are communicating to artists they don't view it as a permanent solution, but as one tactic against AI mimicry that can help as the industry waits for courts or lawmakers to intervene.

RESULT

Since its release, Glaze has been downloaded more than 1 million times. It was awarded the Distinguished Paper Award at USENIX Security Symposium and the 2023 USENIX Internet Defense Prize. ("USENIX Announces the Winners," USENIX, 2023). It has also released a web-based app that allows users to upload an image to add the cloak.

Shan hopes that copyright laws will evolve to the point that Glaze isn't necessary because AI image generators need to receive consent from artists for training purposes by law. Copyright laws were created based on the abilities of an average person well before AI capabilities were commonplace. "Imagine the average human being was able to process 2 billion images, learn some information about it and then massively reproduce those images using what they learned," he says. "I'm pretty sure our copyright protection law would look pretty different than what we have today."

In the meantime, Glaze is preparing for countermeasures. "Most security problems come down to an arms race, so we'll probably expect the same here," Shan says. "Our goal is to increase the costs for these big companies. They may be able to bypass it, but it will be more challenging for some random AI bros on the internet."

The Glaze team considered potential countermeasures that could be used against their application and found the protection it provided was still more than 85% effective, as judged by artists. The team anticipates stronger attacks in the future and is continuing to update Glaze to prevent new attacks, better optimize its appearance on images, and account for new diffusion training models.

ARTISTS OPTING OUT (CONT'D)

HOW GLAZE WORKS




Image courtesy of Shawn Shan,
University of Chicago Glaze project
and used with permission.

IMAGINE //

A cinematic scene from a movie called The Joy of Painting With Robots. Extreme long shot captures a beautiful landscape with a robot painting on an isle, orange and gold accent lighting, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--c 10 --s 750 --ar 16:9





IMAGINE //

A cinematic scene from a cyber punk movie called the Network. This extreme long shot captures a network of monolithic server towers on a bright and sunny day, red and orange accent lighting, in the style of dynamic action sequences, captured on Phantom High-Speed Camera
--s 750 --c 10 --ar 16:9

WHAT'S NEXT

PLATFORMS IN DATA RUSH TO STAKE THEIR CLAIMS

As more creators demand the right to provide consent for their works to train AI models, and as courts evaluate how copyright law interacts with the practice, platform owners will seek to update their terms of service to include an agreement that user data can be used to train future models. Several such incidents have already occurred:

- ▶ In March, Zoom changed its terms of service in a way that appeared to give it permission to harvest user data for AI training. After users protested, Zoom clarified that customers create and own their own video, audio, and chat content. CEO Eric Yuan posted to LinkedIn promising user consent would be sought for any AI training and that the terms of service update was a mistake (Axios, 2023).
- ▶ In July, Google updated its privacy policy to allow the company to collect and analyze information people share online for AI training. It specifically referenced Google Translate, Bard, and Cloud AI capabilities as products that would benefit from such training (Search Engine Journal, 2023).
- ▶ In August 2023, Mozilla had lawyers and privacy experts review Microsoft's updated service agreement to go into effect Sept. 30, 2023. The experts couldn't interpret if Microsoft intended to use personal data such as audio, video, chat, and attachments from products like Office, Teams, and Xbox, or not. Mozilla launched a petition to ask Microsoft to clarify their intent (Mozilla, 2023).

PUTTING A PRICE ON USER DATA

On the other side of the coin, platforms that host a large amount of user data are reconsidering their free-for-all access models. Realizing that third parties are using their data to train products that create a lot of value, social media platforms are looking to charge for access to their APIs instead of offering it for free to all developers. These decisions haven't been without controversy either. Here are several that took place in 2023:

- ▶ Reddit planned changes to its API to charge premium access fees to developers who wanted to use Reddit's user forum data in their applications. Moderators of large "subreddits" or topical forums on the site protested, saying the fees would limit the ways they use the site with third-party apps and even harm accessibility features ("Reddit Communities to 'Go Dark,'" The Guardian, 2023).
- ▶ Coding help website Stack Overflow announced it would start charging for access to its API providing access to programming questions and answers from its 20 million users. The site says that it only wants to charge companies that are developing LLMs and that it will continue to license data for free to some developers.
- ▶ Elon Musk raised prices on X's (formerly known as Twitter) API access to start at \$42,000 per month for access to 50 million tweets (Wired, 2023).

Between API pricing, licenses negotiated between content owners and AI companies, and court decisions involving copyrighted data, a standard for pricing on AI training data may approach a consensus in 2024. It's likely that such a standard would include the minimum amount of data that an owner would have to contribute to a model and how often that data is used in the outputs of LLMs before earning compensation.v



IMAGINE //

A cinematic scene from a cyber punk movie called the Network. This extreme long shot captures a network of monolithic server towers in a snowy valley on a bright and sunny day, red and orange accent lighting, in the style of dynamic action sequences, captured on Phantom High-Speed Camera


--s 750 --ar 16:9

RECOMMENDATIONS

IT departments will be asked to look at data security from a new perspective and to consider how valuable a corpus of data is for AI training. For example, where previously companies would have wanted website data exposed to improve placement in search engine results, now they'll want to protect that data from scrapers. Questions will be asked about what data is exposed to third-party platforms in the course of a workday and whether that data might be used to train a third-party AI model. There may be opportunities to charge a fee for access to data if companies aren't interested in training their own models with their data. Either way, companies will have to respect their customers' own digital sovereignty or else face user backlash, the courts, or both.

INFO-TECH RESOURCES

- ▶ **MATURE YOUR PRIVACY OPERATIONS**
Establish a comprehensive organization-wide privacy program that's measurable and drives business efficiency. Put it into action consistently and at scale to comply with privacy regulations and earn user trust.
- ▶ **MASTER CONTRACT REVIEW AND NEGOTIATION FOR SOFTWARE AGREEMENTS**
Revisit your software licensing agreements with vendors to prioritize protecting your data from being used to train AI models.
- ▶ **DEVELOP APIS THAT WORK PROPERLY FOR THE ORGANIZATION**
Increase application quality and code reusability and improve development throughput for your organization while exposing the right internal services and data to third parties and business partners.



A cinematic scene from a cyberpunk movie. In the foreground, a person wearing a black hooded jacket and sunglasses walks towards the camera. The background is filled with other hooded figures and a shower of sparks or fire, creating a sense of action and danger. The lighting is dominated by red and orange hues, with some blue light visible on the left.

COMPANIES WILL HAVE TO
RESPECT THEIR CUSTOMERS'
OWN DIGITAL SOVEREIGNTY
OR ELSE FACE USER
BACKLASH, THE COURTS,
OR BOTH.

IMAGINE //

A cinematic scene from a cyber punk movie called the Network. Extreme long shot of people breaking into a vault, in the style of dynamic action sequences, captured on Phantom High-Speed Camera

--ar 16:9 --s 250

CONCLUSION

THE GENERATIVE ENTERPRISE IN THE AGE OF EXPONENTIAL IT

Generative AI marks the first era of the age of exponential IT. It will be a core pillar for pursuing new lucrative business models and the autonomization of organizational capabilities. During this era, IT will be crucial to an organization's success and as a result will face exponentially increasing demand from the business. They must pursue this challenge while finding the right balance in a high-risk, high-reward proposition. IT will assume a mandate to act as a business partner that pursues innovation and exploits emerging technologies.

In pursuing that mandate, IT must extend its mutually beneficial symbiotic relationship with the business to include AI. AI will provide organizations with cheap and accurate predictions at scale, but humans will need to provide judgment about when and how those are harnessed to benefit

the organization. CIOs can be the leaders in the loop and ensure that the judgment represents the underlying values of the organization. If they don't, the relationship with AI could turn from beneficial to predatory.

When Gordon Moore made his prediction about exponential computation growth in 1965, he said the trend would last for at least a decade. He was not only proven right, but the trend lasted for decades beyond that, and it's up for debate in some quarters about whether Moore's law still prevails. Either way, the trend persisted because it became a goal for chip designers. The goal pushed them to form a beneficial symbiotic relationship with technology – chip designers created improved circuit board designs for high-performance computing, those computers were then used to further augment their designs, and the

cycle continued. Faster chips created the demand for high-performance computers, and that supported the design of smaller and more complex devices that preserved the law. Moore's law was a choice made in the pursuit of faster, better, and less expensive computing power.

Moore's law demonstrates the symbiosis in the relationship between humans and technology. Just as the exponential growth seen with processors is also true of AI, it's also the case that this type of relationship will need to be forged with AI to dictate its future role in business and society.

Whether that relationship ends up being one that's beneficial or predatory – well, that's the choice in front of us.

EXPERT CONTRIBUTORS

EXTERNAL EXPERT CONTRIBUTORS

QUOTED

Taj Manku, CEO, Cognitive Systems

Monica Goyal, lawyer and director of legal innovation, Caravel Law

Colin Graham, CEO, Arcalogix

Reggie Townsend, vice-president of data ethics, SAS

Patrick Walsh, CEO, IronCore Labs

Shawn Shan, Glaze project lead, University of Chicago Department of Computer Science

BACKGROUND

Neil Trevett, president and chairman, Metaverse Standards Forum

Michael MacKenzie, GM of Industrial IoT Edge Services, Amazon Web Services

Blake Rooney, CIO, Husch Blackwell LLP

Daniel “Dazza” Greenwood, Executive Director of law.MIT.edu

INFO-TECH EXPERT CONTRIBUTORS

Sanchia Benedict

Manish Jain

Andrew Sharp

Janice Clatterbuck

Allison Kinnaird

Aaron Shum

Rob Garmaise

Geoff Nielson

Mike Tweedie

Adib Ghubril

Robert Redford

Steve Willis

Michel Hébert

THANK YOU
TO THE 894
“FUTURE OF IT 2024”
ONLINE SURVEY
RESPONDENTS!
WITHOUT YOU,
THERE WOULD BE
NO TECH TRENDS
REPORT.

IMAGINE //

Digital Neurons in space, photorealism, accent lighting,
in the style of dynamic action sequences, captured on
Phantom High-Speed Camera

--c 75--s 750--ar 16:9

BIBLIOGRAPHY

INTRODUCTION

“Gordon Moore, Intel Co-Founder and Creator of Moore’s Law, Dies Aged 94.” *BBC News*, 25 Mar. 2023. Web.

Maslej, Nestor, et al. “The AI Index 2023 Annual Report.” *AI Index Steering Committee, Institute for Human-Centered AI, Stanford University*, April 2023. Web.

Moore, Gordon. “Cramming More Components onto Integrated Circuits.” *Electronics*, vol. 38, no. 8, Apr. 1965: pp. 114-19.

Reed, Jonathan. “Are We Doomed to Make the Same Security Mistakes with AI?” *Security Intelligence*, 11 July 2023. Web.

Stettler, Mark, and Shesha Krishnapura. “Moore’s Law – Not Dead – and Intel’s Use of HPC to Keep It Alive.” *HPCwire*, 11 Jan. 2016. Web.

AI-DRIVEN BUSINESS MODELS

Agrawal, Ajay, et al. *Power and Prediction: The Disruptive Economics of Artificial Intelligence*. Harvard Business Review Press, 2022. Accessed 23 Aug. 2023.

Arka. “Giraffe – Long Context LLMs.” *The Abacus.AIBlog*, 22 Aug. 2023. Web.

J., Joshua. “What Is the Difference Between the GPT-4 Models?” *OpenAI*, 22 Aug. 2023. Accessed 24 Aug. 2023.

Licholai, Greg. “AI Poised To Revolutionize Drug Development.” *Forbes*, 13 July 2023. Accessed 24 Aug. 2023.

Metz, Cade, and Mike Isaac. “Meta Made Its AI Tech Open-Source. Rivals Say It’s a Risky Decision.” *The New York Times*, 18 May 2023. Web.

Roubini, Nouriel. *MegaThreats: Ten Dangerous Trends That Imperil Our Future, And How to Survive Them*. Little, Brown and Company, 2022.

Schulz, Yogi. “Expect Your ISP to Offer More WiFi Functionality.” *IT World Canada*, 20 Jan. 2023. Web.

“Seamless Communication.” *Meta AI*, 2023. Accessed 23 Aug. 2023.

“The Generative AI Landscape: Top Startups, Venture Capital Firms, and More: The State of Generative AI in 7 Charts.” *CB Insights Research*, 25 Jan. 2023. Web.

“What Is WiFi Motion?” *YouTube*, uploaded by Cognitive Systems Corp, 10 Aug. 2020.

Wilson, H. James, and Paul R. Daugherty. “Creating the Symbiotic AI Workforce of the Future.” *MIT Sloan Management Review*, 21 Oct. 2019. Web.

AUTONOMIZED BACK OFFICE

Basham, Victoria. “Allen & Overy Integrates ChatGPT-Style Chatbot to Boost Legal Work.” *The Global Legal Post*, 16 Feb. 2023. Web.

“Caravel Law Presents - The Future Is Now - The Use of AI in Legal | Tuesday June 20, 2023.” *YouTube*, uploaded by Caravel Law, 20 June 2023.

Chui, Michael, et al. “Economic Potential of Generative AI.” *McKinsey*, 14 June, 2023. Accessed 8 Sept. 2023.

Ghoshal, Anirban. “ServiceNow Adds New Features to Its Now Assist Generative AI Assistant.” *CIO.Com*, 26 July 2023. Web.

Hemmadi, Murad. “With New AI Bot, Shopify Offers Its Merchants a Sidekick.” *The Logic*, 26 July 2023. Web.

“Juniper Networks Extends AIOps Leadership with Large Language Model (LLM) Capabilities, Zoom Integration and Expanded Wi-Fi 6E Portfolio.” Accessed 27 July 2023. Press release.

Kaplan, Ari. “The Effects of Harvey and Generative AI on the Legal Industry.” *Reinventing Professionals*, 27 June 2023. Accessed 5 Sept. 2023.

Liu, Nancy. “CrowdStrike’s New Generative AI Tool Combines Machine and Human Data.” *SDxCentral*, 30 May 2023. Web.

McClead, Ryan. “AI-Pocalypse: The Shocking Impact on Law Firm Profitability.” *3 Geeks and a Law Blog*, 3 Aug. 2023. Web.

---. “Generative AI Could Reduce Law Firm Revenue by 23.5%.” *3 Geeks and a Law Blog*, 2 Aug. 2023. Web.

“Moveworks Recognized for Generative AI Innovation in 2023 Artificial Intelligence Breakthrough Awards Program.” *Business Wire*. 21 June 2023. Press release. Web.

“Salesforce Announces Einstein GPT, the World’s First Generative AI for CRM.” *Salesforce*, 20 July 2023. Press release.

“ServiceNow Expands Generative AI Capabilities With Case Summarization and Text-to-Code to Drive Speed, Productivity, and Value.” *ServiceNow*, 26 July 2023. Press release. Accessed 27 July 2023.

Sussman, Bruce. “Why Are So Many Organizations Banning ChatGPT?” *BlackBerry Blog*, 8 Aug. 2023. Web.

Weiser, Benjamin, and Nate Schweber. “The ChatGPT Lawyer Explains Himself.” *The New York Times*, 8 June 2023. Web.

Weiss, Debra Cassens. “Latest Version of ChatGPT Aces Bar Exam with Score Nearing 90th Percentile.” *ABA Journal*, 16 Mar. 2023. Web.

Wood, Stuart. “Revolutionizing Legal Innovation: Exploring the Impact of AI in Law Practice with Monica Goyal.” *Business Decisions*. Podcast. 23 Aug. 2023.

BIBLIOGRAPHY

SPATIAL COMPUTING

Gross, Dariusz. "Create a 3D Model With Your AI-Powered Smartphone." *Medium*, 18 Sept. 2022. Web.

Gurman, Mark. "Apple Tests 'Apple GPT,' Develops Generative AI Tools to Catch OpenAI." *BNN*, 19 July 2023. Web.

Heaven, Will Douglas. "Welcome to the New Surreal. How AI-Generated Video Is Changing Film." *MIT Technology Review*, 1 June 2023. Accessed 23 Aug. 2023.

Lee, Angie. "Meet the Omnivore: Startup Develops App Letting Users Turn Objects Into 3D Models With Just a Smartphone." *NVIDIA Blog*, 27 June 2023. Web.

Mileva, Gergana. "MagiScan App Lets Users Create 3D Models With Their Smartphone." *ARPost*, 11 July 2023. Web.

Ray, Siladitya. "Apple Reportedly Expects To Sell Fewer Than 400,000 Vision Pro Headsets Next Year Due to Production Snags." *Forbes*, 3 July 2023. Accessed 25 Aug. 2023.

Wiggers, Kyle. "The Week in AI: Apple Makes Machine Learning Moves." *TechCrunch*, 15 June 2023. Web.

RESPONSIBLE AI

Biever, Celeste. "ChatGPT Broke the Turing Test — the Race Is on for New Ways to Assess AI." *Nature*, vol. 619, no. 7971, July 2023, pp. 686–89.

"EU AI Act: First Regulation on Artificial Intelligence." *European Parliament*, 6 Aug. 2023. Web.

"Fact Sheet: Biden-Harris Administration Secures Voluntary Commitments from Leading Artificial Intelligence Companies to Manage the Risks Posed by AI." The White House, United States Government, 21 July 2023. Web.

Goodin, Dan. "ChatGPT Is Enabling Script Kiddies to Write Functional Malware." *Ars Technica*, 6 Jan. 2023. Accessed 6 Sept. 2023.

Goswami, Rohan. "OpenAI Changed Its Plans and Won't Train on Customer Data, Sam Altman Says." *CNBC*, 5 May 2023. Web.

Harris, David E. "As a Responsible AI Researcher, I'm Terrified about What Could Happen Next." *Berkeley Blogs*, 29 June 2023. Web.

Johnston, Jeffrey S., and Briana Falcon. "'Algorithmic Justice': FTC Orders Destruction of Algorithms Following Privacy Violations." *Lexology*, 6 Apr. 2022. Web.

"MEPs Ready to Negotiate First-Ever Rules for Safe and Transparent AI." *European Parliament*, 14 June 2023. Press release. Web.

Metz, Cade. "How Could A.I. Destroy Humanity?" *The New York Times*, 10 June 2023. Web.

O'Carroll, Lisa. "EU Moves Closer to Passing One of World's First Laws Governing AI." *The Guardian*, 14 June 2023. Web.

"Introducing ChatGPT Enterprise." OpenAI, 28 Aug. 2023. Accessed 31 Aug. 2023.

"Regulatory Framework Proposal on Artificial Intelligence." *European Commission*, 20 June 2023.

"Responsible Innovation." SAS.com, 2023. Accessed 30 Aug. 2023.

"SAS Puts Humans at the Center with Responsible Innovation Initiative." SAS, 10 May 2022. Press release. Web.

Siddiqui, Tabassum. "Risks of Artificial Intelligence Must Be Considered as the Technology Evolves: Geoffrey Hinton." *University of Toronto Faculty of Arts & Science*, 4 July 2023. Web.

"The OECD Artificial Intelligence Policy Observatory." *OECD*, 2023. Accessed 29 Aug. 2023.

"Welcome to the AI Incident Database." AI Incident Database, 2023. Accessed 31 Aug. 2023.

Zorthian, Julia. "OpenAI CEO Sam Altman Asks Congress to Regulate AI." *Time*, 16 May 2023. Web.

SECURITY BY DESIGN

Akash, S. "Revolutionizing Cybersecurity With Generative AI." *Analytics Insight*, 12 July 2023. Web.

"Amazon Agrees to Injunctive Relief and \$25 Million Civil Penalty for Alleged Violations of Children's Privacy Law Relating to Alexa." Department of Justice Office of Public Affairs (DoJ), 19 July 2023. Web.

Benesch Law. "New FTC Investigation Into OpenAI May Shed Light on Developing AI Enforcement Landscape." *JD Supra*. Accessed 11 Sept. 2023.

Brown, Tom B., et al. "Language Models Are Few-Shot Learners." *arXiv.org*, 22 July 2020. Web.

"CISA Strategic Plan 2023-2025." *Cybersecurity & Infrastructure Security Agency (CISA)*, Sept. 2022. Web.

"Cloaked AI." *IronCore Labs*. Accessed 29 Aug. 2023.

Easterly, Jen. "CISA Director Easterly Remarks at Carnegie Mellon University." *Cybersecurity & Infrastructure Security Agency (CISA)*, 27 Feb. 2023. Accessed 9 Aug. 2023.

Goldstein, Eric. "CISA Cybersecurity Strategic Plan: Shifting the Arc of National Risk to Create a Safer Future." *Cybersecurity & Infrastructure Security Agency (CISA)*, 4 Aug. 2023. Web.

Harang, Rich. "Securing LLM Systems Against Prompt Injection." *NVIDIA Technical Blog*, 3 Aug. 2023. Web.

Huynh, Daniel, and Jade Hardouin. "AI Attacks: Prompt Injection Vs. Model Poisoning + Mitigations." *Mithril Security Blog*, 26 July 2023. Web.

Jillson, Elisa. "Hey, Alexa! What Are You Doing with My Data?" *Federal Trade Commission*, 13 June 2023.

BIBLIOGRAPHY

SECURITY BY DESIGN (CONT'D)

Lai, Christine, and Jonathan Spring. "Software Must Be Secure by Design, and Artificial Intelligence Is No Exception." *Cybersecurity & Infrastructure Security Agency (CISA)*, 18 Aug. 2023. Web.

Pamma, Ashee. "Is the Microsoft-OpenAI Partnership on the Rocks? Analysts Weigh in." *ITBusiness.Ca*, 7 Sept. 2023. Web.

Polsinelli Law. "Generative AI's 'Industry Standards' for Cybersecurity and Data Privacy Could Be Here Sooner Rather than Later." *JD Supra*, 26 July 2023. Web.

Reed, Jonathan. "Are We Doomed to Make the Same Security Mistakes with AI?" *Security Intelligence*, 11 July 2023. Web.

Thompson, Clive. "How Rust Went From a Side Project to the World's Most-Loved Programming Language." *MIT Technology Review*, 14 Feb. 2023. Web.

Vastel, Antoine. "How to Prevent ChatGPT From Stealing Your Content & Traffic." *The Hacker News*, 30 Aug. 2023. Accessed 30 Aug. 2023.

Walsh, Patrick. "'Embeddings Aren't Human Readable' And Other Nonsense." *HackerNoon*, 18 Aug. 2023. Accessed 29 Aug. 2023.

---. "Security of AI Embeddings Explained." *IronCore Labs*, 29 June 2023. Accessed 29 Aug. 2023.

---. "Waitlist Now Open for New Encrypted AI Vector Embeddings Solution." *IronCore Labs*, 29 June 2023. Web.

DIGITAL SOVEREIGNTY

"20 Biggest GDPR Fines so Far [2023]." *Data Privacy Manager*, 13 June 2023. Web.

Afifi-Sabet, Keumars. "What Is GDPR? Everything You Need to Know, From Requirements to Fines." *ITPro*, 16 Oct. 2019. Web.

"Ask Microsoft: Are You Using Our Personal Data to Train AI?" *Mozilla Foundation*, n.d. Accessed 1 Sept. 2023.

Baio, Andy. "Exploring 12 Million of the 2.3 Billion Images Used to Train Stable Diffusion's Image Generator." *Waxy.org*, 30 Aug. 2022.

Brown, Tom B., et al. "Language Models Are Few-Shot Learners." *arXiv.org*, 22 July 2020. Web.

"California Consumer Privacy Act (CCPA)." *State of California - Department of Justice - Office of the Attorney General*, 15 Oct. 2018. Web.

"Common Crawl - Open Repository of Web Crawl Data." *Common Crawl*, 2023. Accessed 31 Aug. 2023.

Dave, Paresh. "Stack Overflow Will Charge AI Giants for Training Data." *Wired*, 20 Apr. 2023. Web.

Efrati, Amir, and Aaron Holmes. "OpenAI Passes \$1 Billion Revenue Pace as Big Companies Boost AI Spending." *The Information*, 29 Aug. 2023. Web.

Freedman, Linn Foster. "Wyoming Enacts Genetic Data Privacy Act." *Data Privacy + Cybersecurity Insider*, 10 Mar. 2022. Web.

Fried, Ina. "Zoom Clarifies Terms of Service After Customer Outcry." *Axios*, 11 Aug. 2023. Web.

Hern, Alex. "Reddit Communities to 'Go Dark' in Protest Over Third-Party App Charges." *The Guardian*, 11 June 2023. Web.

Hill, Kashmir. "This Tool Could Protect Artists From A.I.-Generated Art That Steals Their Style." *The New York Times*, 13 Feb. 2023. Web.

"How AI Experts Are Using GPT-4." *MIT Technology Review*, 21 Mar. 2023. Accessed 23 Aug. 2023.

"How to Prevent ChatGPT From Stealing Your Content & Traffic." *The Hacker News*, 30 Aug. 2023. Accessed 30 Aug. 2023.

Koerner, Katharina. "Generative AI: Privacy and Tech Perspectives." *IAPP*, 28 Mar. 2023. Web.

Laidler, John. "Harvard Professor Says Surveillance Capitalism Is Undermining Democracy." *Harvard Gazette*, 4 Mar. 2019.

Loten, Angus. "Companies Put AI to Work Outside the Cloud, Trimming Costs." *The Wall Street Journal*, 16 Aug. 2023. Web.

Lutkevich, Ben. "AI Lawsuits Explained: Who's Getting Sued?" *TechTarget*, 4 Aug. 2023.

"More US States Are Ramping up Data Privacy Laws in 2023." *BleepingComputer*, 25 July 2023. Accessed 8 Aug. 2023.

Rix, Ryan. "Data Rights Protocol." *Consumer Reports Innovation Lab, Github*, 15 Aug. 2022. Web.

Shan, Shawn, et al. "Glaze - Publications and Media Coverage." *Glaze*, 25 June 2023. Web.

---. "Glaze: Protecting Artists From Style Mimicry by Text-to-Image Models." *USENIX, 32nd USENIX Security Symposium*, 2023, pp. 2187-204. Web.

Southern, Matt G. "Google Updates Privacy Policy To Collect Public Data For AI Training." *Search Engine Journal*, 3 July 2023. Web.

Sussman, Bruce. "Why Are So Many Organizations Banning ChatGPT?" *BlackBerry Blog*, 8 Aug. 2023. Web.

"USENIX Announces the Winners of the 2023 Internet Defense Prize." *USENIX*, 9 Aug. 2023. Accessed 23 Aug. 2023.

Zimmeck, Sebastian. "Data Rights Protocol and Global Privacy Control." *Innovation at Consumer Reports*, 13 Jan. 2022. Web.

INFO~TECH

RESEARCH GROUP

North America 1-888-670-8889

United Kingdom 0808 175 3350

Australia 1800 242 692

International +1-519-432-3550

INFOTECH.COM

IMAGINE //

Sunrise, Light Art, Hyperspectral
Imaging, Light Blue, Multiverse, in
the style of dynamic action sequences,
captured on Phantom High - Speed Camera
--c 100 --s 500 --ar 16:9