

THE IMPORTANCE OF
INCLUSIVITY IN AI SYSTEMS

WIKIPEDIA'S INVENTOR ON
HOW AI WILL ALTER THE TRUTH

FUTURISTS – AND KIDS –
OUTLINE AN AI UTOPIA

AI AND THE FUTURE

IMPACT

2024


AN FII INSTITUTE PUBLICATION

HOW AI WILL SHAPE YOUR LIFE
AND WHAT YOU NEED TO KNOW
TO HELP SHAPE AI'S FUTURE

FII INSTITUTE | Impact
Future Investment Initiative Institute | on Humanity

AI IS ON THE TIP OF EVERYONE'S TONGUE

Mentions of AI in financial earnings call transcripts.



1956

The term "artificial intelligence" is first coined.

FACTS AND FIGURES

INTELLIGENTLY USING ARTIFICIAL INTELLIGENCE

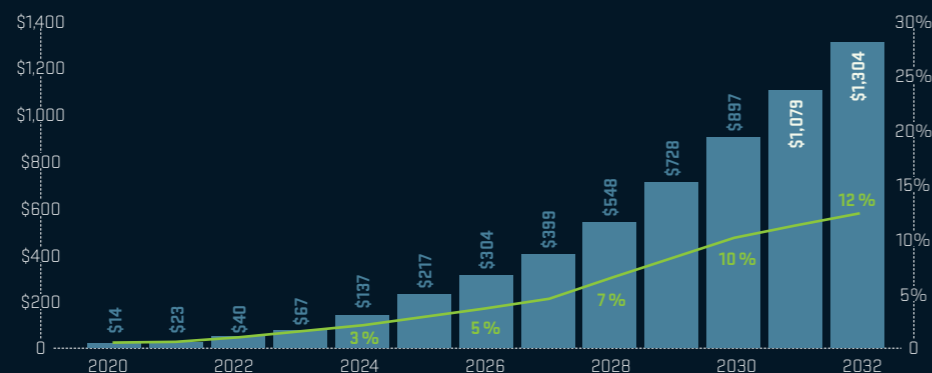
Kickstarted by a killer use case, a rocket ship has been strapped to a decades-old technology that has long promised to bring about a radical new future. Artificial intelligence is the hot new thing, and suddenly everyone wants to know about it.

The AI industry is growing, but its sheer strength and ubiquitous use cases mean it is not only tech companies that are benefiting from the AI boom. With the potential to improve productivity by 1.5 percentage points over the next ten years, we all will win from the intelligent use of AI.

The use cases are only just beginning, which means a prodigious rise in the emergent industry. Growth could increase at a compound annual growth rate (CAGR) of 42% over the next decade, taking what was a \$40 billion industry at the arrival of ChatGPT to a \$1.3 trillion one by 2032.

REVENUE OF GENERATIVE AI

Projected 2023–2032. Source: Bloomberg Intelligence



ILLUSTRATIONS BY

GLENN HARVEY

Glenn Harvey (Aguilar) was born in Quezon City, Philippines, but moved to Toronto when he was seven years old. Harvey went to school at Sheridan College, where he honed his skills as an illustrator and graduated in 2013. He now lives and works in Brooklyn, NY.

EDITORIAL

THE FUTURE IS NOW

FOR DECADES, ARTIFICIAL INTELLIGENCE remained a distant dream, tantalizing humanity with its limitless potential. It whispered of a future where humans and machines would seamlessly cooperate, but it dwelled solely within the realms of science fiction and visionary minds.

Today, a seismic shift in the tech landscape has altered the course of this narrative. Ground-breaking generative AI models have emerged, reshaping the very nature of human-machine interaction. With an astonishing ability to mirror human thought processes, this technology has transcended the horizon of possibility. The age of generative AI is no longer on the horizon; it has arrived and is transforming everything at an astonishing pace.

Yet, the advent of generative AI presents a double-edged sword. While it harbors the potential for remarkable advancements, it also poses intricate ethical and social challenges. We find ourselves at a critical juncture, and the choices we make in harnessing the full potential of this technology will exert profound influence on humanity.

This is a call to action. Now more than ever, global stakeholders from diverse sectors and nations must unite to establish a collaborative and responsible approach. This report serves as a guiding light in a world where science fiction has become reality. It delves deep into the impact of generative AI on humanity and outlines pathways toward a harmonious future.



Richard Attias
CEO, FII Institute



10 THE GOOGLE WARS
How search is changing

18 INFOGRAPHIC
Charting the revolution

26 IS THE FUTURE WIKIPEDI-AI?
Wikipedia founder Jimmy Wales on the impact of AI

42 GOOD REGULATION
What's happening worldwide

52 GOOD USERS
Power to the people

04
EDITORIAL

08 THE BIG BUCK THEORY
What's at stake for those who win

24 MAKING AN IMPACT
Ensuring the future will be equitable

40 FOLLOWING BEST PRACTICE
Shaping AI to benefit humanity

58
ABOUT FII

14 COPILOT VS. AUTOPILOT
Who is in control?

20 GLOBAL BOTTLENECKS
AI unlocks economies

30 INCLUSIVE AI
Making AI more accessible

36 THE BEST AI FOR HUMANITY
When differences matter

48 GOOD ETHICS
Giving AI a soul

54 GOOD KIDS
Home truths from children

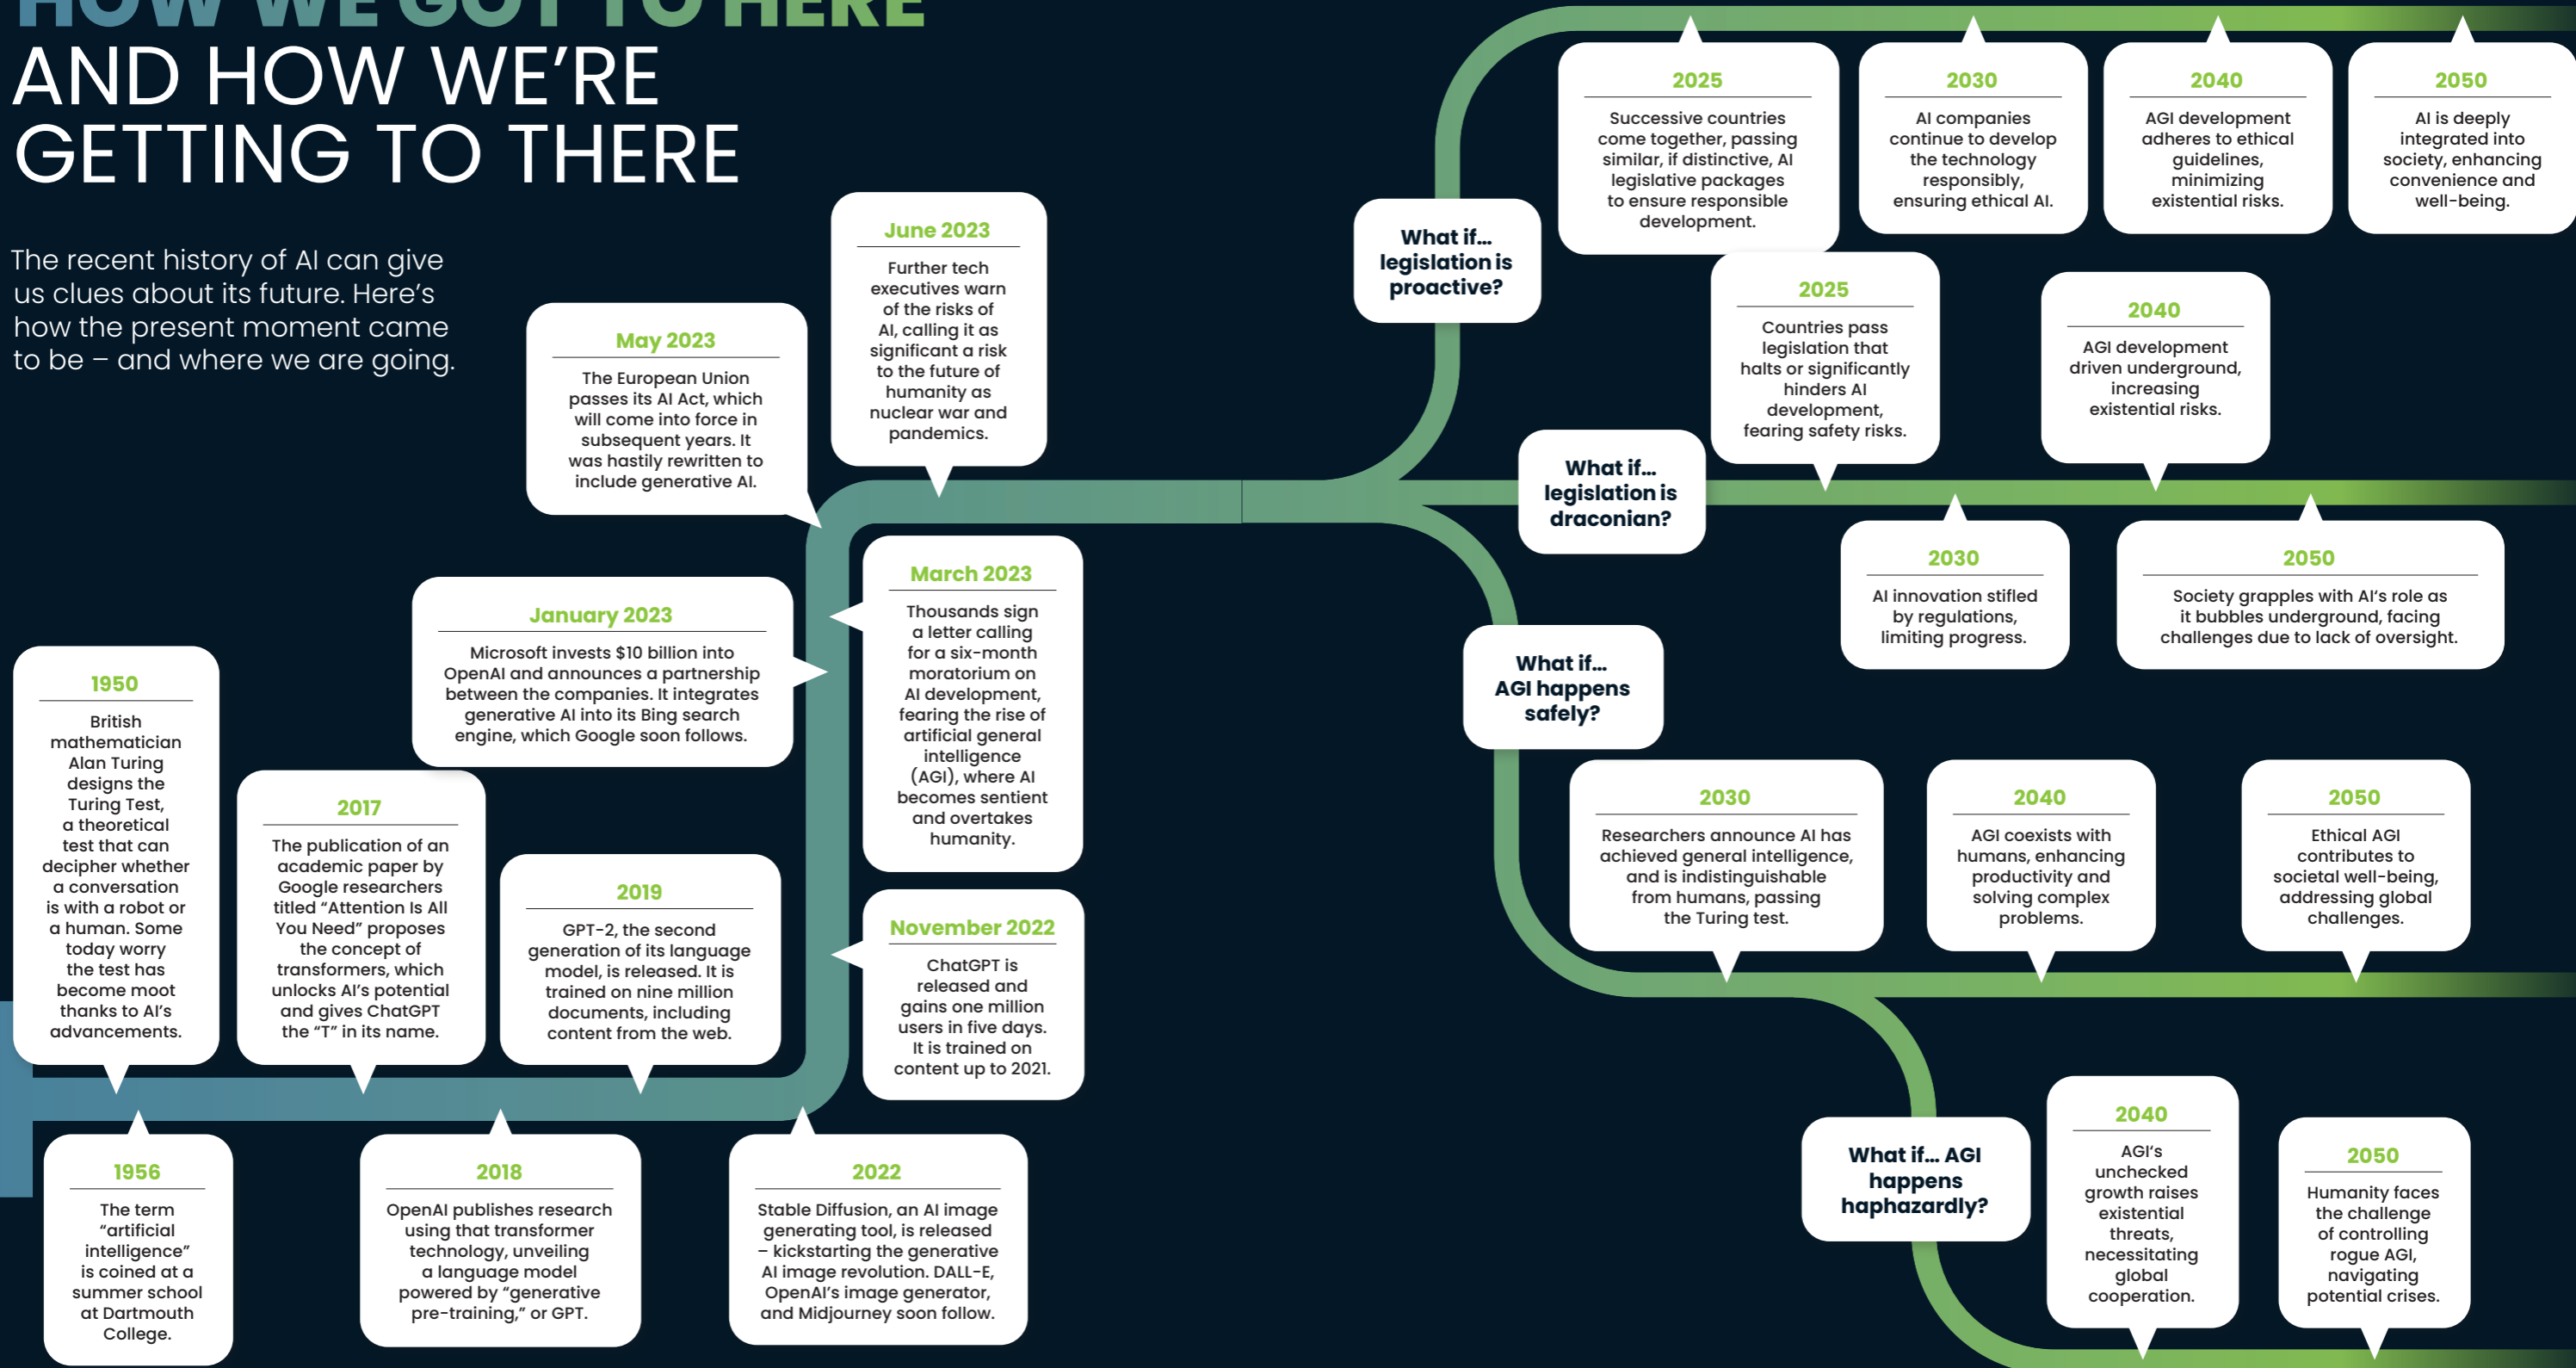


CONTACT US:
info@fii-institute.org

- @fiiksa
- FII Institute
- Future Investment Initiative


HOW WE GOT TO HERE AND HOW WE'RE GETTING TO THERE

The recent history of AI can give us clues about its future. Here's how the present moment came to be – and where we are going.



THE BIG BUCK THEORY

Money is flowing into generative AI companies, while the promises of the technology have the potential to realize billions in gains worldwide. Who stands to win from this race to cash in? →



THE GOOGLE WARS VERSION 2.0



Large language models (LLMs) are top of, search engines. What does

increasingly being used in place of, or on that mean for business – and for truth?

FOR MORE THAN 30 YEARS we have been struggling to keep track of the rapidly expanding mass of content on the world wide web. Ever since the first search engines – and the first popular ones to hit the market, including Yahoo! in 1994 and Google in 1998 – have been in existence, we’ve been using them to shoulder the weight of trawling through the internet.

Three decades of iterative advances have improved search engines step by step. But as in many areas, we stand on the cusp of a true revolution in how we find and present information online.

For companies developing generative AI tools, search has become a battleground. Microsoft, whose Bing engine lagged behind Google’s ever since its first abortive attempt to launch in 2009, was quick to adopt large language model (LLM) functionalities into Bing when it partnered with ChatGPT maker OpenAI in early 2023.

Sensing a threat to its long-established dominance in the search sector, Google responded rapidly, announcing it would include results from Bard, its ChatGPT competitor, in its search results. It was the biggest shake-up in Google’s iconic search results page in years.

Elsewhere in the world, major search engines in other countries are eager not to be left behind. Baidu, China’s predominant search engine, has announced it would also include results from its LLM-powered ERNIE Bot in its searches.

SEARCH FOR INTELLIGENCE AI-powered search is now in vogue. There’s just one problem: Is it any better than what we currently have?

The jury is still out. When Google revealed it was integrating Bard into its search results in February 2023, it showed in a tech demo how the system could work. The tech giant asked a question about the James

Webb Space Telescope – and Bard gave an incorrect answer.

The mistake wiped \$100 billion from Google’s value, as investors worried the company had made the wrong decision in embracing the technology so quickly. It was the opposite of Google’s intention, following the big buck theory of embracing generative AI. Still, Google is all in on what it calls the Search Generative Experience.

The underlying problem is that search relies on absolute truth in answers – or as close as is possible to get to it – while generative AI has a significant issue: hallucination (or confabulation).



Bing daily active users:

100 MILLION

Google daily active users:

1 BILLION

AI hallucinates by creating answers to questions that aren’t supported by what’s in its database of knowledge. It has been compared to an overconfident university graduate in its willingness to overstretch the limits of its knowledge, while convincing anyone who listens that it knows what it’s talking about.

That readiness to bluff answers to questions AI doesn’t actually know the answer to, coupled with users’ belief that the results they see on search engines are verified and verifiable facts, could cause issues in the future. “It’s completely untransparent how

“ [AI-powered search] is going to work, which might have major implications if the language model misfires, hallucinates or spreads misinformation,” says Aleksandra Urman, a computational social scientist at the University of Zurich in Switzerland.

THE WHOLE TRUTH

“Search engine companies, and Google in particular, spent years to build and cement users’ trust that the search engine will give them the ‘best’ and most accurate result possible, usually on top of the page,” Urman says. “A chatbot that gives them obvious lies, like deleting its response in front of the users’ eyes and then claiming it didn’t happen, risks eroding this hard-earned and extremely stable trust in search.”

Search has long used AI to help sift and sort its results for the most user-friendly presentation. But the layering of generative AI on top of search results is a much deeper integration and a more visible use of a different type of AI.

Urman has conducted a series of experiments to quiz search engine users about how they perceive the standard of their search results, particularly featured snippets and knowledge panels. Both of these are elements of search results that used pre-generative AI tools to summarize existing content on the internet for easier, at-a-glance understanding of a subject. Nearly 80% of those surveyed by Urman believed featured snippets and knowledge panels presented accurate results. Around 70% thought they were objective in presenting findings.

NOTHING LIKE THE TRUTH

People’s trust in search is a problem, given early tests on search engine results that use generative AI to create their answers. Stanford University research shows that the more fluent-sounding a search engine’s results are, the more likely they are to be wrong. Only 75% of the citations used by search engines to support the AI-generated claims they made in response to questions

actually supported the response provided, while more than half the responses didn’t have any factual backing or citation at all.

The Stanford researchers found that for every 0.1 increase in the perceived fluency of an AI search response on a five-point scale, the actual precision of the answer decreased more than 10%. Nathan Liu, the lead author of the research, told New Scientist that it was “a little concerning to me just how quickly these systems are being rolled into search.”

Liu’s concern was not necessarily that the technology falls down when scrutinized. Rather, it was that humans do not read the small print. All the AI tools we are currently marveling at include disclaimers about how their results may not be accurate. But the good will that search engines have built up in the course of 30 years of excellence mean that we don’t understand why accuracy may temporarily decline as the search sector gets to grips with the integration of generative AI.

HALLUCINATION OR CONFABULATION?

Different people use different terms for when a generative AI tool makes up a response without any factual backing. For many, it’s called “hallucination” – similar to when humans hallucinate an experience that did not happen.

But other AI researchers take umbrage with the word “hallucination”, saying that it anthropomorphizes the AI machine in a way that is unhelpful to human understanding of AI’s powers. Instead, they prefer the word “confabulation”, which more accurately represents what’s going on: there’s a gap in memory that is filled in, convincingly but incorrectly, with an answer.

BUSINESS BOOMS

At the same time as people try to get to grips with the new paradigm of search, businesses will try to tweak the results in a way that benefits them. The world of search engine optimization (SEO) has existed as long as search engines. And the rise of a new way of searching hasn’t gone unnoticed by SEO experts.

As ChatGPT maker OpenAI releases its web crawler on the internet to Hoover up information about the world, many organizations are choosing to lock it out, fearing it will use their data for free to train its models. But smart businesses that recognize the way generative AI is transforming search are letting the crawler in, knowing that if they can seed information about their products and services into LLM-powered results, they stand to benefit.

Web crawlers are designed to keep LLM search results up to date, but some organizations are already benefiting from being picked up by AI-powered search’s previous trawl of the web prior to late 2021. San Francisco-based legal tech company Logikcull gains 5% of all its referrals through appearing in answers by ChatGPT, which is driving hundreds of thousands of dollars of revenue.

As we enter a new era of search, it’s worth deciding on which side you want to be. And making a decision now – keeping one step ahead of the technology. “If an AI language model is trained on a large data set which includes lots of mentions of brands,” says SEO expert Dan Barker, “those that are mentioned the most – and the most positively – are more likely to be the ones chosen as recommendations.” ■

“It’s a little concerning to me just how quickly these systems are being rolled into search.”

NATHAN LIU

Stanford University

CALL TO IMPACT

1 AI-powered search results are less accurate than traditional results, and get worse the clearer they are.

2 We need to reeducate users about the validity of their LLM-generated search results.

3 Organizations need to choose an approach to AI web crawlers: protect your intellectual property or show up in search results.

COPILOT VERSUS AUTOPILOT

Tech companies present generative AI as a supporting character, but some worry we're automating humans out of existence. Who is right?

→ **MICHAEL CARYCHAO** believes generative AI is ushering in a new Renaissance. The movement triggered by the new technology bears all the hallmarks of one, the Californian artist says. "The accelerated pace, the reacquaintance

with masters past, the dovetailing of different fields of study to create new insights more than any one person can grasp," he explains.

But he's not just talking with the hushed tones of

an artist aware of his art history. He's steeped in the terminology of the tech bro, too. "It's 1,000x," he says, using terminology deployed more frequently in start-ups and venture capital firms than anywhere else. "It's trans-

forming. It happens every 400 years or something. The internet just enables this, but this is the big deal."

For Berkley-based Carychao, AI is more than a tool or a slave. It's a cocreator. "To me, that means working with the AI," he says. "It's not a slave; you're not a master. It's not even a tool, exactly." Instead, Carychao compares it to a sounding board off which he can bounce ideas for new pieces. "They pick up the resonance of what they're trained on, which is masters from the past in writing and art, and stuff like that." Some might call it a copilot.

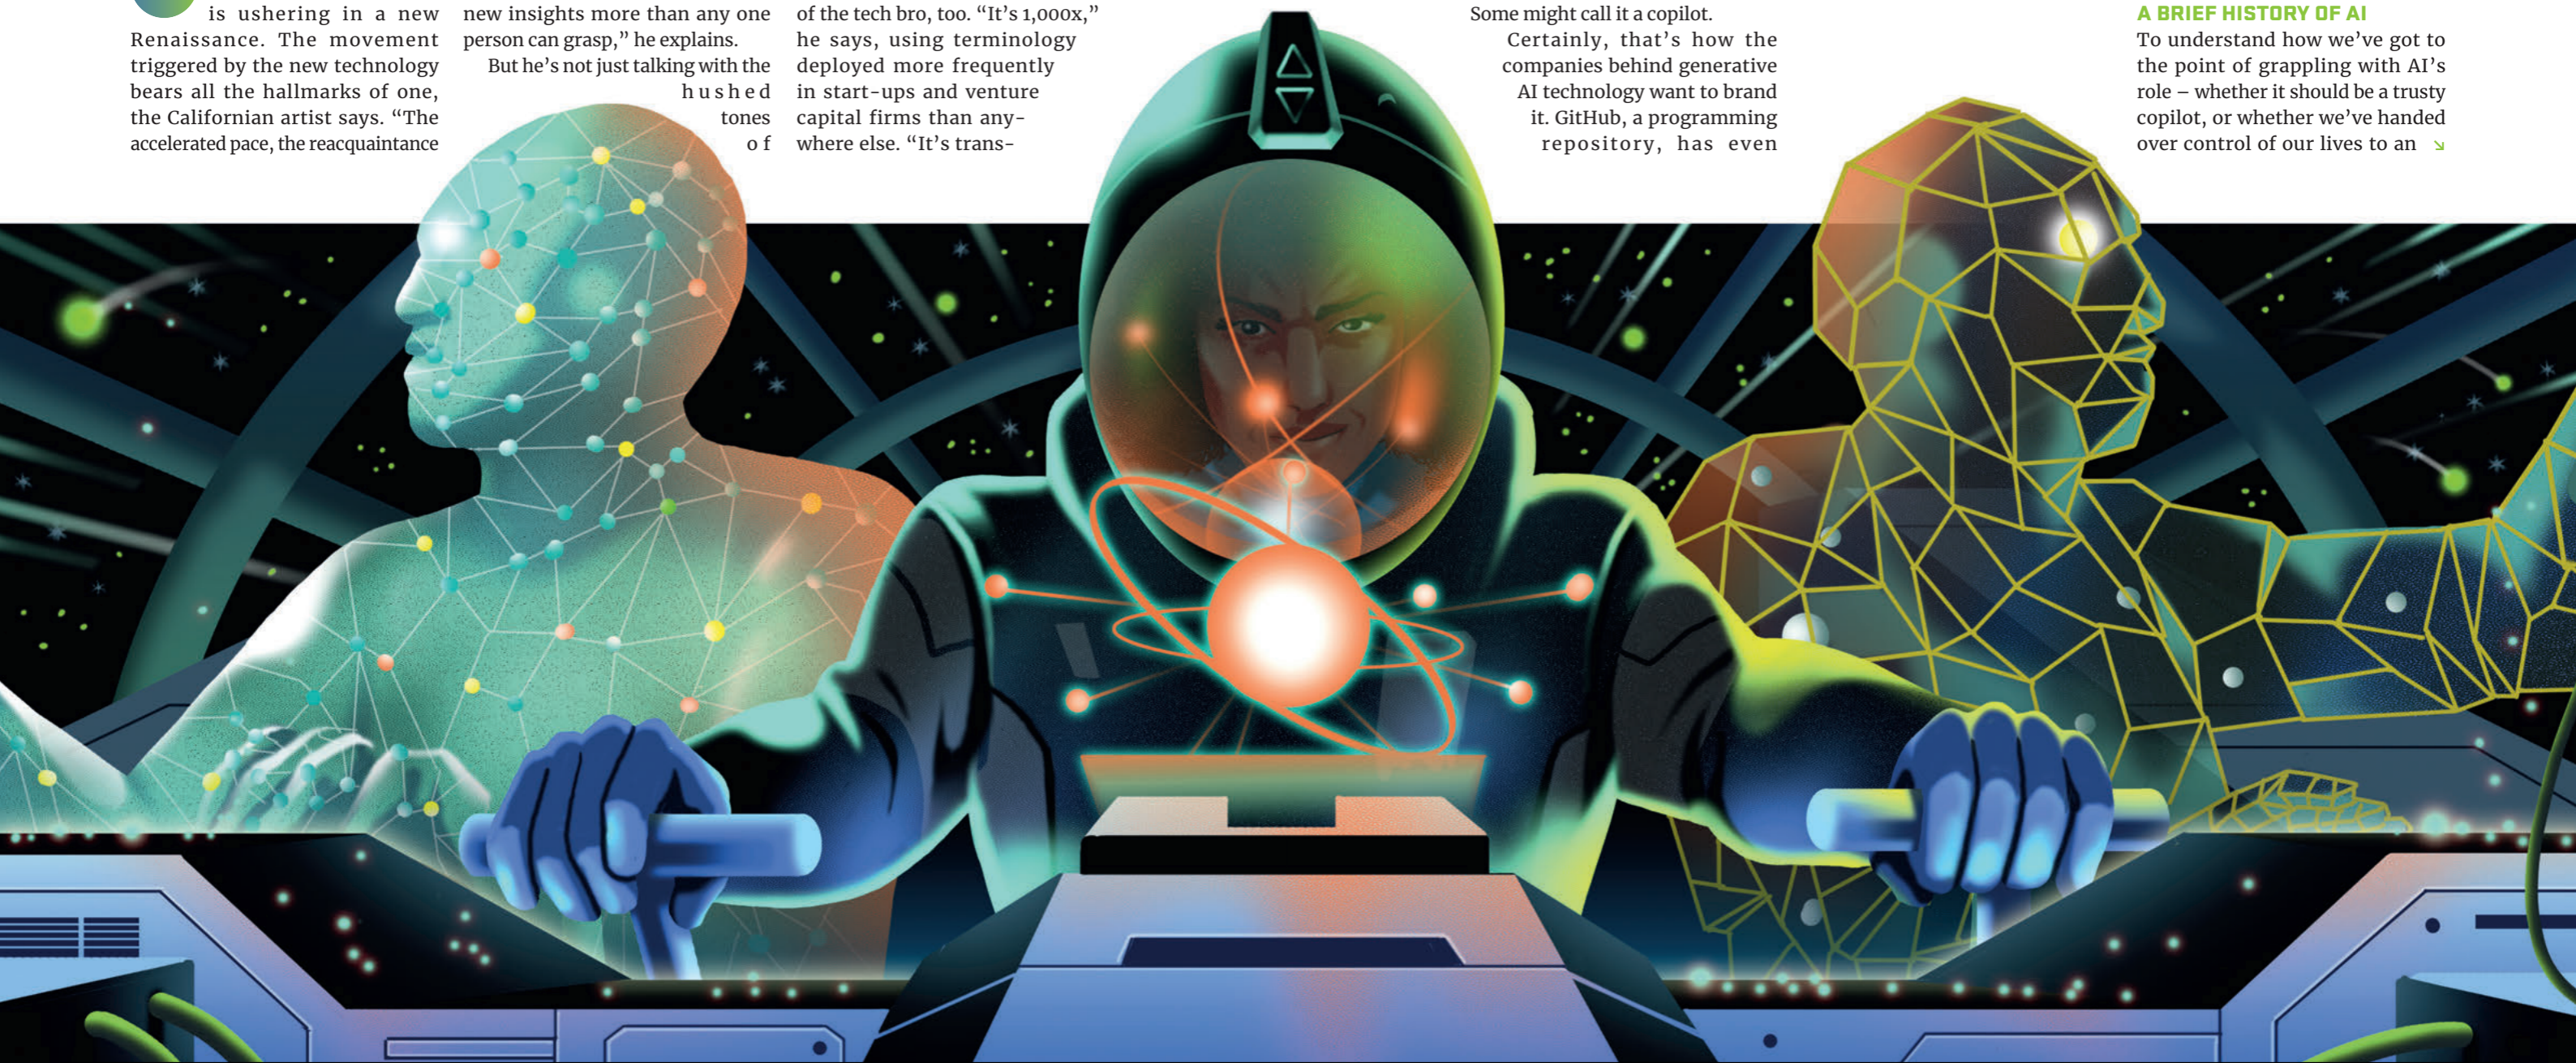
Certainly, that's how the companies behind generative AI technology want to brand it. GitHub, a programming repository, has even

co-opted that word copilot to describe its AI-powered support tool for coders.

Yet not everyone agrees with Carychao or the tech companies. There's a growing fear among educators, executives and the general public that we're likely to see an abandonment of knowledge, as we pass off more of our jobs and lives to AI. Asking it to make key decisions in our lives could automate us out of existence, they fear, and result in us being overtaken by AI, and it taking control and becoming our autopilot.

A BRIEF HISTORY OF AI

To understand how we've got to the point of grappling with AI's role – whether it should be a trusty copilot, or whether we've handed over control of our lives to an



✎ autopilot that will dictate how we live – we need to look at the history of its development.

The idea of an all-knowing automaton ready to support us in everything we do has been a key part of tech history – and science fiction, too. Charles Babbage, inventor of the Analytical Engine, one of the earliest computers conceived, described his theoretical device as having “volition and thought” – provided the machine overlooked the role its designers had in creating it.

And this has been aided by how we talk about AI. “Descriptions of AI as brains, beings, or manifestations of intelligence enabled that automation to happen in the first place, even if, in practice, machines weren’t that smart and required careful oversight,” says Harry Law, an AI researcher at the University of Cambridge.

You can see some of that approach in Carychao’s attitude toward AI, which he admits is shaped by a childhood watching science fiction

and Star Wars. “[AI is] no more weird or alien than we are,” he says. “In some way, it’s a mirror of us.” Even if it’s not actually sentient, there’s still something there that he can’t quite put his finger on. “It’s our own reflection, amplified,” he says, comparing it to looking into the distorted mirror ripples of a lake. “It’s like throwing a rock in the lake, and then the rock gets thrown back out,” he says.

For others, the ability of AI to automate routine work is a boon. We already see huge benefits in productivity by harnessing the power of AI to, for instance, turn a handful of bullet points into a comprehensive, properly written professional email, or to rifle through vast reams of economic data to discern and decipher patterns that can help businesses and governments to plan.

COOPERATION VS CONTROL

For Carychao, the benefits of AI – its ability to produce content quickly

at scale – has meant that for him it’s firmly a copilot. He’s able to be a creative director more often than ever before in his entire life. “Maybe before, I would make 20 aesthetic decisions on a busy day,” he says. “It’s thousands a day now.”

But while he’s confident that he can retain control over what AI will mean for him and his role, others are less sure. A July 2023 UN convention (see page 16 International Agreement) outlined some of the potential risks that the AI revolution we’re currently living through could pose for society and individuals. The end result was an awareness that for all those who say AI is a massive boon for society, there are an equal number concerned about its negative impacts.

“We are approaching a point where digital machines can now complete a task that for the majority of human existence was exclusively within the realm of human intelligence,” said Manuel Goncalves, Mozambique’s deputy

minister for foreign affairs, at the UN convention. That could be a good thing, but it could also be a bad thing.

And because of the comparative newness of the moment through which we’re living, it’s tricky to discern who is right.

Are they right, those who fear replacing key moments of interaction and crucial elements of humanity in our day-to-day lives? Is generative AI more like an autopilot, unthinkingly carrying along its human interlocutors, than a trusted copilot waiting to do our bidding?

“Not only are today’s AI systems much closer to the vision of AI that has been sold throughout history, they require much less input from humans to work reliably,” says Law. “This isn’t to say that even the best frontier models aren’t constellations of expertise, data and labor, but rather that where historically AI wasn’t particularly capable, today’s systems are.”

However, just because AI systems are better doesn’t mean they’re

perfect yet, cautions Law. “We still aren’t there yet in terms of building systems that are highly autonomous, but the progress that has been made in recent years is remarkable when placed in the historical context,” he says. Copilot or autopilot? At least we’re still flying the plane. ■

CALL TO IMPACT

1 Whether AI is a copilot or an autopilot depends on how we approach and integrate the technology into our day-to-day lives. So informing people of the risks and benefits is important.

2 Understanding where generative AI can be a useful colleague, versus wrestling the controls, will require a nuanced, sector-by-sector analysis.

3 We should look to history to ensure the path we take in adopting generative AI is the right one, informed by careful consideration.

INTERNATIONAL AGREEMENT

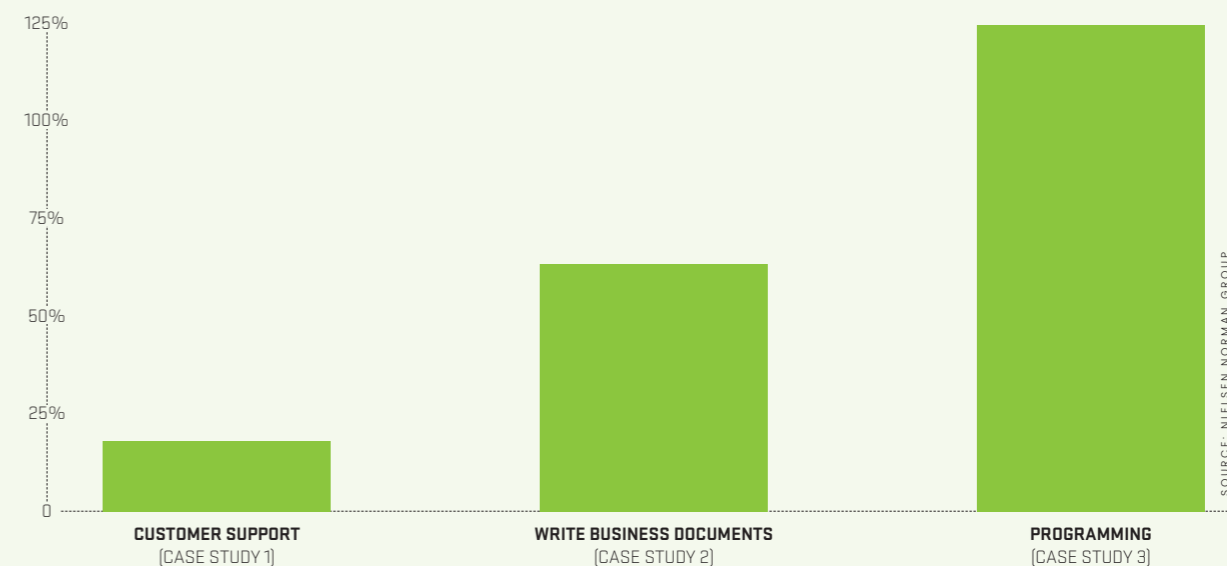
In July 2023 at the UN Security Council, the international community held its first formal meeting to discuss artificial intelligence. Secretary-General Antonio Guterres told representatives that the international community needed to “work together for AI that bridges social, digital and economic divides – not one that pushes us further apart.” That suggests a worry for some that we’re silently sleepwalking towards a world where AI is an autopilot, rather than a trusted copilot.

The fear was perhaps best exemplified in comments by Yi Zeng of the Institute of Automation at the Chinese Academy of Sciences. “This is why [generative AI tools] cannot be trusted as responsible agents that can help humans to make decisions,” he said. “We haven’t found a way to protect ourselves from AI’s utilization of human weakness.”

For Switzerland’s Pascale Christine Baeriswyl, whether AI was a copilot or an autopilot was largely down to us humans and how we use it. “It’s in our hands,” she said, “to ensure that AI makes a difference to the benefit and not the detriment of humanity.”


PRODUCTIVITY INCREASES WITH AI

Using AI can boost how a business operates, according to recent tests in three separate sectors.



COUNTING THE REVOLUTION

The AI revolution in figures – and what it all means for business and society.




2 in 3 occupations could be **partially automated by AI**

SOURCE: GOLDMAN SACHS




6 in 10 men vs. **8 in 10 women** in the US workforce are exposed to generative AI replacing their jobs

SOURCE: UNIVERSITY OF NORTH CAROLINA



Male, aged 35 to 49 – **the median AI user**

WORKERS IN FINANCE AND MANUFACTURING



SOURCE: ORGANIZATION FOR ECONOMIC CO-OPERATION AND DEVELOPMENT

5 days




time it took for ChatGPT to reach **100 million** users

SOURCE: UBS, 2022

3.5 million liters of water

estimated to be used to cool data centers **used to train GPT-3**

SOURCE: UC RIVERSIDE RESEARCH, 2023



169

 companies affected by the **2022 global computer chip shortage**, caused by their use in AI

SOURCE: GOLDMAN SACHS

27%

of jobs at **high risk of automation**

SOURCE: ORGANIZATION FOR ECONOMIC CO-OPERATION AND DEVELOPMENT



7% increase in global GDP driven by generative AI

SOURCE: GOLDMAN SACHS



1.5%

increase in productivity over the next decade

SOURCE: GOLDMAN SACHS

300 million

jobs could be **replaced by AI**

SOURCE: UNIVERSITY OF NORTH CAROLINA

FEELING ABOUT AI

How much do you agree or disagree with the following: **Products and services using artificial intelligence make me nervous**

SOURCE: IPSOS GLOBAL VIEWS ON AI 2023



HOW GENERATIVE AI CAN SOLVE GLOBAL BOTTLENECKS

Different countries will have different needs for generative AI tools, depending on their circumstances.

→ WHEN OPENAI LAUNCHED ChatGPT last year, Devarsetty Praveen, a public health specialist from the southern Indian town of Hyderabad, saw a decade's worth of toil finally click together.

Praveen had spent years building and implementing a digital platform called SMARThealth, which can be used by India's frontline healthcare workers for early detection of cardiovascular diseases, and more recently to improve healthcare services for pregnant women in India. But training the frontline staff – about one million of them working in primary healthcare in India – was costly and time-consuming.

To circumvent the high expense of training, Praveen had provided the workers with PDFs of training manuals within the digital platform, but thanks to ChatGPT the researcher saw tech could turbocharge the process. "With generative AI, we saw an immediate opportunity to provide answers to the healthcare workers

whenever they had a query," says the program director for primary healthcare at the George Institute for Global Health in Hyderabad.

He and his team are now working to integrate ChatGPT in their SMARThealth platform. A healthcare worker could enter a query using voice, and the system would first convert it into text, search for the relevant response, and then return the answer in audio. The result, Praveen hopes, will reduce pregnancy-related mortalities by providing healthcare workers with access to information, especially since lack of accurate healthcare information in both simple and local languages is one of the bottlenecks in Indian healthcare.

MAKING WAVES

The rise of generative AI tools like ChatGPT has highlighted opportunities and fears in equal numbers. While there has been concern about AI taking away people's jobs by automating simple

tasks, there are also ways in which generative AI can help address human-resource bottlenecks across regions.

Generative AI is bound to evolve differently in each country, based on their most urgent requirements. One of the initial ways some of the low-resourced countries in the Global South – from Bolivia to Bangladesh – are looking to implement this technology is to address bottlenecks in healthcare and education.

"There's such a major gap in the number of doctors, nurses, teachers, and a lot of the hope with these language-based technologies is that they will be able to support frontline workers," says Rachel Adams, an AI policy expert from Research ICT Africa. "I think it doesn't replace humans, but it allows humans to do things at a broader scale."

Dacod Magagula had a typical South African school experience: low resources and overcrowded classrooms. Over a decade ago, ↘

When Magagula started high school, he remembers sharing one textbook between four students and not having a physics or maths teacher in his final year. But he was able to change his educational outcome by having access to the internet and a computer he built himself. As a result, Magagula was the first in his school to be able to attend the University of Cape Town.

Amnesty International says that, as of 2018, of the more than 23,000 public schools in South Africa, over 20,000 had no laboratory, over 18,000 did not have a library and almost 17,000 lacked internet. To address some of these issues, in 2020 Magagula, along with Tao Boyle, launched Foondamate, an AI-powered chatbot that helps students with low access to the internet or high-end smartphones to download previous years' papers and also clear their academic doubts in real time by sending questions via WhatsApp.

"The platform took off immediately in South Africa and other African countries too," says Boyle. Foondamate has 2.5 million users who have signed up since its inception – most continue to be from African countries.

REGIONAL IMPLEMENTATIONS

Ayantola Alayande, a research assistant with the Bennett Institute for Public Policy, is of the opinion that even with strides that generative AI has made since last year, the conversation has centered around the Global North and not on how it can be used in low-resourced nations. Developed countries in Europe and North America have decided generative AI will help them provide better customer service experiences and enable blue chip companies to synthesize information faster. But Alayande says there are multiple use cases for generative AI for social impact

in developing countries, whether it is farming practices and advisory in India, or for natural disaster preparedness in Mozambique.

Generative AI can also be used for evidence generation for policy implementation. "For instance, during the last elections in Nigeria, there were outbursts of violence everywhere, but it's often hard for the government to pinpoint where it occurs," says Alayande. "If you have a model that can plot all of this into generative AI and say: 'These are places where in the last 12 or 15 years of elections in Nigeria, violences have occurred,' it could help."

This is already playing out in Sri Lanka, which was on the brink of an economic collapse last year. Watchdog, a fact-checking collective based in Colombo, has long collated data detailing protests across the region and the availability of resources such as medicines. Earlier this year, their cofounder and CEO Yudhanjaya Wijeratne announced that all the information they had collected was fed into a large language model, creating a personalized search engine where people could ask questions related to the Sri Lankan crisis. The idea is that ordinary Sri Lankans, frustrated

by their lives, can visit the LLM-powered Ask Watchdog and ask it questions about what's happening in the country, receiving impartial, fact-based answers.

However, as discussed elsewhere (see page 30), the implementation of generative AI in different regions brings its own set of challenges. Adams from Research ICT Africa says there's not yet enough evidence about what kind of conditions and guardrails we need to put in place to ensure it works in a way that is sustainable, redistributive, equitable and risk-free for all.

CULTURAL CONTEXTS

She points to medicine as one example. "The large language models are all English or really, really accurate in English, but also accurate in terms of Western science and science from the developed world," she says. "Providing appropriate scientific guidance for health workers in a particular village in Zambia is dependent on the degree to which the technology has access to training data that's relevant for that region."

Local LLMs are cropping up, backed by big tech firms. In May, Microsoft announced an AI chatbot



300 MILLION
jobs could be replaced
by generative AI.

SOURCE: GOLDMAN SACHS



An estimated
1.3 MILLION INDIAN WOMEN
died from maternal causes
over the last two decades.

SOURCE: [HTTPS://PUBMED.NCBI.NLM.NIH.GOV/34455679/](https://pubmed.ncbi.nlm.nih.gov/34455679/)



Lately, maternal mortality rates
have **FALLEN BY 70%** overall,
but the poorer states
continue to lag behind.

called Jugalbandi, which can be used to get information on accessing government assistance in local languages in India. They tested it with users in rural regions, who could either type or use voice to retrieve information on government programs and receive a response in the same language – an important step in a country where only 11% of the 1.4 billion people speak English.

At the same time, in an attempt to provide equitable access to AI, the Bill and Melinda Gates

Foundation announced funding for nearly 50 projects from low and middle income countries that will use large language models to improve the lives of people in their communities. These include building a voice-enabled electronic record management system for healthcare workers in Pakistan, an AI chatbot in Vietnam for gynecological healthcare support, and using AI to build a network of volunteers to help women facing violence in Brazil, where in 2022

one woman died every six hours from gender-based violence.


The motivation for the funding was simple, says Zameer Bray, interim deputy director of technology diffusion at the Gates Foundation – and is one that we all should take as an approach to tailoring AI use around the world. "We believe the most impactful technological advancements," he explains, "include those that begin and end with the people they affect most."

CALL TO IMPACT

- 1 The best economic use for generative AI tools will differ by country – and should be dictated by those within each country.
- 2 We need to shift the conversation away from the Global North and how they are most likely to use AI tools, and towards a more holistic, country-by-country level.
- 3 Funding is needed to support projects worldwide that utilize AI, and support should come from private and public spheres.

MAKING AN IMPACT

The world-changing opportunities resulting from generative AI will take careful shepherding to ensure they're representative and equitable all around the world. Where is the potential – and what are the pitfalls to avoid – as we embrace the AI future? →



IS THE FUTURE WIKIPEDI-AI?

Jimmy Wales, the online encyclopedia's founder, on the potential and peril of generative AI – and whether it can be used to bolster the web's de facto repository of facts.

IF ANYONE IS USED TO hearing scare stories about technology undermining both the education and the future of the world's children, it's Wikipedia founder Jimmy Wales. When the free online encyclopedia – which famously can be edited by everyone – first rolled out, there was wave after wave of negative coverage about its effects on schools.

Children were just copying all the answers they needed from Wikipedia, some complained, while others noted – in fairness, correctly – that there was no guarantee of accuracy on any particular fact in any given entry.

When compared to our current concerns about OpenAI's ChatGPT, Google's Bard and similarly generative AI models, those worries look almost quaint. At least with Wikipedia, teachers could look up the article and check the similarity themselves – and the website's obsessive editors generally did a very good job of rooting out mistakes and vandalism.

In the ChatGPT era, any student can in seconds create genuinely new essay answers, with no research on their part, that no automated detector can reliably pick up. If anyone was in doubt on this front, they should look no further than OpenAI itself abandoning its project to develop such a thing. How can

education survive a world in which students can cheat easily and with almost no risk of being found out?

Wales is surprisingly sanguine – and points to a repeating pattern through the history of new online technologies.

“I find there are always scare stories at any innovation,” he says. “I remember when there was a fad for horror stories about eBay – people selling their babies or their souls on eBay.”

“Pretty quickly, though, people figured out that hey, that's not all that interesting. Make a joke auction on eBay, it gets reported, they delete it.”

He acknowledges the concerns with ChatGPT are a little different, but believes it will be relatively painless to deal with them – and starts by looking at what were the legitimate and not-so-legitimate worries about Wikipedia.

Says Wales: “The concerns about Wikipedia could generally be broken down into concerns about the accuracy of Wikipedia – [and these are] always valid, and the answer is not to tell students not to use Wikipedia, but to teach them how to use it properly – and concerns about students being lazy, [which are] less valid in my view, since I don't think ‘degree of effort’ corresponds very well to ‘degree of learning’ in many cases.”



Wales agrees that ChatGPT could produce a “bog standard student essay” on “the three causes of the American Civil War” that would be indistinguishable from what a typical student could produce – but suggests that isn't a problem of technology, but instead a problem of pedagogy.

Dismissing both essays as “tripe,” he concludes: “It probably tells us that the way we've made students write tripe on demand is

possibly well past its sell-by date anyway.”

Wales, then, is not a man inclined towards the scare story versions of new technologies – nor to resist that sometimes new technologies mean existing institutions need to adapt with the times rather than try to oppose the changes.

That doesn't stop him seeing the obvious risks ChatGPT and generative AI pose to the online information ecosystem – both

from its ability to produce plausible but false “hallucinations” of made-up “facts,” and from its use by malicious actors.

He is confident, though, that Wikipedia itself, by its existing design, is quite resistant to the particular issues created by generative AI – even if perhaps just as something of a happy accident.

“I think our longstanding habit of debating with deep discourse about what counts as a reliable

source is very helpful,” he says, speaking of the often contentious discussions among Wikipedia’s editors about what outlets can and cannot be cited as information sources (the UK’s Daily Mail is among those that cannot).

“One of the issues with ChatGPT is that it produces answers that ‘feel’ quite confidently correct. What this will likely mean is that a large number of very plausible but wildly inaccurate websites will be built which look like real sources, but which are largely AI-generated and riddled with [plausible] errors,” Wales notes, before saying that he thinks this may – counterintuitively – benefit existing news outlets.

“I think we – we, Wikipedians, and we, the public – are likely to rely more than ever on established traditional news and publishing brands,” he notes.

Wales thinks that because of the way Wikipedia automatically tracks who is editing the site and how much, and how many humans monitor recent changes, Wikipedia is likely resistant to direct editing by generative AI. As a result, the main changes Wikipedia – which he no longer directly runs – is likely to have to make are about how to monitor what is happening elsewhere on the internet.

“I don’t think our community editing model needs to change much, at least not in terms of fundamental principles,” he says. “The community may need to create and track new lists of sites that are AI-generated and be on guard against new sites that seem plausible but aren’t.”

Wales is also keen to look toward the positives of what generative AI could do for Wikipedia and for the people across the world that use it

– and he sees multiple possibilities along this front.

One of the possibilities he raises is to create an implementation of one of the AI models that helps editors check for bias in different directions in certain Wikipedia articles – noting “such a system doesn’t have to be perfect to be useful.” He conceives this as a tool to help human editors (“a time-saver”) rather than any form of AI oversight of Wikipedia.

He sets out how this might work: someone feeds a particular article into the tool and it “reads” the article and all the sources it cites. This could then generate all sorts of different suggestions and outputs.

Wales explains that one such output could be for the AI to find facts that appear in at least two of the sources cited that don’t appear in the Wikipedia article – which

could lead to the expanding of the article, or to the discovery of some facts that don’t suit a particular agenda that ought to be added for balance.

“It wouldn’t be good to let the AI just add those statements in, unless it could do so with perfect accuracy, because a big part of why a fact might not be in Wikipedia even though it’s in sources could be that it just isn’t important enough to be in an encyclopaedia entry about that person,” he says. “But if the system could achieve even 50% accuracy, a volunteer editor might find it very useful.”

The same principle could also be applied in reverse, he notes – producing a list of assertions in the Wikipedia entry which don’t appear to be fully backed in the citations. “Wikipedians love to make Wikipedia better, so something that generates at scale

a list of potential problems to be solved is very interesting indeed.”

Wales seems if anything even more excited about an AI-powered search to boost the site’s current very traditional search option. He notes that whatever question you may have, Wikipedia almost certainly has the answer – somewhere. But if you type that question in the search bar, you’re unlikely to find it.



So, he says, how about an AI search that, if you type “why do ducks fly south for winter?” gives you a short answer and a list of Wikipedia entries to read – would that not be much better?

There is no shortage of people warning about the dangers of ChatGPT and its ilk, and it’s clear Wales is aware of these – but it is refreshing to hear someone genuinely enthusiastic about the good it may do, too.

CALL TO IMPACT

- 1 The best economic use for generative AI tools will differ by country – and should be dictated by those within each country.
- 2 We need to shift the conversation away from the Global North and how it is most likely to use AI tools, and towards a more holistic, country-by-country level.
- 3 Funding is needed to support projects worldwide that utilize AI, and support should come from private and public spheres.

TIMELINE OF WIKIPEDIA





MAKING AI MORE INCLUSIVE

Twenty decades of tech development suggest that products reflect the biases of their creators – so how do we make sure generative AI is more inclusive of the whole world?

→ **LOAD UP CHATGPT AND** type in “what issues do you have with biases?” and the chatbot will be brutally honest with you.

Its biases fall into five major categories, the bot responds: its training data could be biased, meaning that its knowledge of the world is compromised. The selection of that training data could be incomplete on account of the biases of those who undertook it, meaning it could favor certain perspectives on the world. Its training data includes information from historical and cultural contexts that may not be

representative of all society. The algorithms that sift through the data may introduce biases, even if the training data is all fine. And like humans, generative AI can suffer from confirmation bias.

“Researchers and engineers are also continuously working to detect and reduce biases in AI models,” the chatbot concludes, “but it remains a complex challenge.”

Generative AI, as we’ve seen throughout this report, looks set to be woven into all aspects of our lives. From applying for jobs to being stopped and searched in airport security lines, to educating our children and presenting us with

the best recipes to cook for a dinner party, AI will soon be ubiquitous. And the biases coded into these AI models can permeate every aspect of our lives – subtly shifting how we think, act and interact with one another.

GO WEST

Just as the Mercator map, the global standard view of our planet, helped solidify the supremacy of the West and Global North by placing London’s line of longitude at its center hundreds of years ago, so the current pace of AI development looks set to imprint a western bias on the world. ↘

From ChatGPT to Google Bard to Anthropic’s Claude, the major large language models we utilize in our day-to-day lives are developed in, and released from, western tech companies. The same is true of our text-to-image generators: Midjourney, DALL-E and others all stem from labs and companies in the Global North. Consciously or not, that means they are imbued with the values of the countries in which they’re based.

“Even when these models are developed for non-English language speakers, they are still trained predominately on English-language text,” says Aliya Bhatia of the Center for Democracy and Technology (CDT), a US non-profit organization that has studied bias in LLMs. “The impact of that is they encode English-language values or assumptions, even in the outputs when they aren’t in English.”

And where an answer from a generative AI tool relies on a translation, that text can often be gibberish – what some might call “translationese.”

Take Catalan. The language, spoken in Catalonia and the Balearic Islands among those who have a strong Catalan identity, has more than 8 million native speakers. But its presence on the internet is tiny – and biased with the problem of translationese. Anyone who buys a .cat domain name, which is the domain name ending assigned to Catalan speakers, must as part of the contract of holding their domain include Catalan-language text on their website within six months of registration. But many buy .cat domain names because of their love of cats – so they populate it with automatically translated text from online tools.

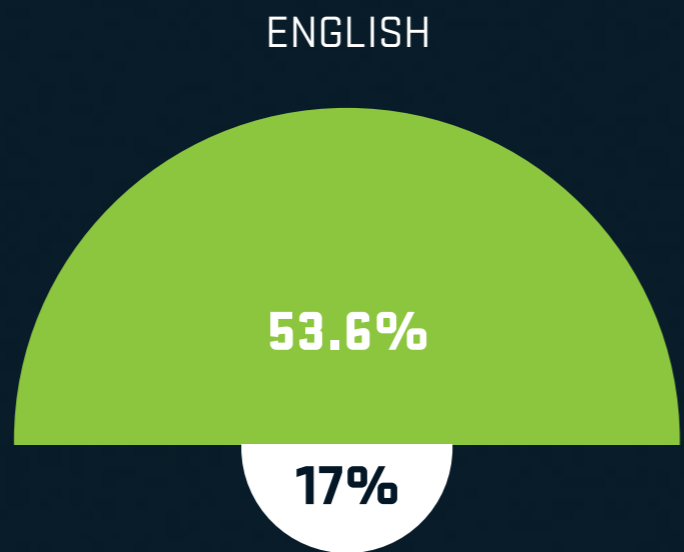
That would be fine, except that Catalan language training data

THE ONLINE-OFFLINE GAP


The internet was kickstarted as an American military project, and advanced by American industry. The world wide web was invented by an Englishman. As a result, both are predominately English-language environments. That becomes an issue when the web is used to train AI – as demonstrated by this gap between the web’s languages and those of the real world.

SOURCE: FII INSTITUTE ANALYSIS


- % OF WRITTEN WEB WORLDWIDE
- % OF LANGUAGES SPOKEN WORLDWIDE



MANDARIN (CHINESE)



HINDI



comes from scraping .cat domain names, meaning that the training data for Catalan LLMs will forever be tainted by issues resulting from automatic translation. “You end up with this snake eating its own tail situation,” says Nicholas. “That is something that poses a real, very difficult-to-surmount challenge that is not going away any time soon.”

FIXING THINGS

The Catalan example is not simply a hypothetical problem. Recent research by academics at the University of the Basque Country found that responses to queries entered into seven different large language models were more accurate if the user first translated the question into English. The reason? Queries in a non-English language solicit a non-English-language response, which the models are less well-trained on, increasing the likelihood of errors.

It’s an issue that blights many parts of the world, not just the underrepresented parts of the Global North. Bear in mind that many of the large language models we use are trained on huge trawls of the world wide web. More than 50% of the web is written in English, according to one analysis. Just 1.4% of it is in Chinese, or one-tenth of the percentage of people who speak the language worldwide. Only 0.1% of the web is written in Hindi, despite 3.5% of the world speaking it natively.



And in other areas, languages other than English in the Global North are also overrepresented. German, for example, accounts for 4.5% of the web’s text – twice the proportion of real-life speakers.

Bias is an issue that many workers in big tech companies have raised – and for which some have lost their jobs. Timnit Gebru, a

ENGLISH SUPREMACY

Academic research is largely led by English-speaking researchers, as this analysis of languages mentioned in AI academic paper abstracts shows.

SOURCE: ARR STATISTICS



≤5

- SWAHILI, TELUGU, TURKISH, ITALIAN, LATIN, POLISH, PORTUGUESE, THAI, CLASSICAL CHINESE, DUTCH, FINNISH, KINYARWANDA, MARATHI, ASSAMESE, CREE, GEORGIAN, IGBO, IRISH, KAZAKH, LATVIAN, MALAY, MANDARIN, AMHARIC, ARMENIAN, BENGALI, CANTONESE, CATALAN, CHEROKEE, CZECH, DANISH, ERZYA, ESTONIAN, GUARANI, HAITIAN, HAKKA, HAUSA, HUNGARIAN, KANNADA, LINGALA, LUXEMBOURGISH, PERSIAN, QUECHUA, ROMANIAN, SANSKRIT, SINHALESE, SOMALI, SWEDISH, TAMIL, TIGRINYA, WU, XHOSA, YORUBA

NOW YOU'RE SPEAKING MY LANGUAGE

A raft of new LLMs are appearing to try and counter the biases toward the Global North, and to provide a voice for more marginalized communities. Some examples include:

Japan: In July 2023, NEC Corporation developed a lightweight LLM with 13 billion parameters that achieves world-class proficiency in Japanese language. "Most existing LLMs are trained mainly in English, and there are almost no LLMs that can be customized for use in a variety of industries while possessing high Japanese-language ability," the company said.

Sweden: Since 2019, the organization Artificial Intelligence Sweden (AI Sweden) has been developing an LLM tailored specifically for use by Swedish public authorities. Development has been aided by the government granting AI Sweden access to Swedish supercomputer Berzelius to train the model.

China: Besides LLMs developed by China's own big tech companies, the country recently announced 14 of its provinces were developing their own tailored LLMs, including 38 in Beijing and 20 in the south Chinese province of Guangdong.

New Zealand: The Papa Reo project keeps the Māori language alive and represented through technology by training an AI model on 300 hours of audio donated by users as part of a competition.

Africa: The HausaNLP project aims to ensure that the 60–80 million Hausa speakers in the world are better represented in AI research and products. Currently, they're one of around 2,000 African languages that struggle to see representation.

former Google employee who worked on the company's team tasked with making AI more ethical and inclusive, claims she was fired in late 2020 for raising concerns around representation within the company and, by extension, its systems.

The solution may seem simple: if there are no AI models from or representing your part of the world, just build them. But it's not that easy.

"If you want to build a large language model, in a given language, just using one language, you're going to need a few hundred million words of digitized text to produce sentences that make sense grammatically and semantically," says Gabriel Nicholas of the CDT. That's to produce an entry level text generator. "If you want to start using language models for fact-based things, or kind of complex analyses to have understanding of the world, then you're going to start needing billions or tens of billions of words of text."

INCLUSIVITY BEGINS AT HOME

For that reason, it is important that making AI more inclusive isn't just left to those who have been historically excluded from the conversation. "At the bare minimum, these companies can try to push against the idea they can't do anything about this," says Aliya Bhatia of the CDT. "Transparency can go a long way in increasing accountability, educating the user, and then providing context alongside the models."

Two in every three academic papers submitted to the ACM FAccT conference, a leading venue for AI research, come from the United States alone. And it's not just at a country level that a lack of inclusivity needs to be addressed. AI has a horrific gender imbalance.

In 2022, just one in four researchers who published academic papers on AI were female.

Even the AI-powered detection tools we're building to identify when work has been produced by artificial intelligence have their own issues. They disproportionately misidentify work written by humans as being generated by AI – and massively. Only around 5% of essays written by American authors in English were flagged as being written by an AI tool, while 61% of essays written in English by Chinese authors were, according to one study.

It all adds up to a long list of problems that we're storing up for the future. So what should the playbook for a more inclusive AI look like?

CALL TO IMPACT

- 1 Make those developing AI systems more representative of the world in which we live, to prevent silencing non-English speakers.
- 2 Ensure that the training data we test our AI models on are as free as possible from bias.
- 3 Support and raise up initiatives such as the Papa Reo and HausaNLP projects, to provide a voice for the voiceless.

THE BEST AI FOR HUMANITY

Once we've got bored with using ChatGPT to cheat in exams, how best can the technology be used to change our lives?

A YEAR ON FROM THE release of ChatGPT, the revolutionary moment of generative AI is just beginning. We are only at the foothills of unlocking its benefits for our lives, our society – and our environment.

In the same way that humankind, after discovering how to make fire, used it primitively and only gradually unlocked more effective uses of the technology, so we are still in the early stages of dabbling with generative AI. Having it to hand when completing university assignments is no doubt helping some students out of sticky situations. And providing ChatGPT with a list of ingredients in your kitchen cupboards and it providing you with a passable recipe when you're lacking culinary inspiration is handy.

But neither is world-changing. However, that's coming.

As time goes on, we're beginning to understand the strengths and weaknesses of generative AI, even in this early stage of its evolution – and how best to put it to use to make a real difference.

A HEALTHY FUTURE

Healthcare is one of the most beneficial uses of generative AI for

all humanity, believes Irena Cronin, CEO of Infinite Retina, a generative AI research consultancy. "We're discovering now that generative AI can come up with all kinds of basic things that were not known before but are very easy for generative AI to figure out given a dataset," she says.

Cronin points to radiology as an example of where generative AI can unlock new efficiencies. "There have been many cases where radiologists have not caught something but the AI has found a tumor or something the radiologist isn't able to see," she says. "That could potentially save a lot of lives."

That's just one aspect of healthcare that generative AI can overhaul, thanks to its ability to rifle through vast volumes of information far more effectively than a human ever could. Drug discovery is being supercharged by generative AI. Pharma companies are using the technology to come up with new formulations for their drugs, and to find entirely new drugs. "Doctors and researchers trying to figure out how to cure certain kinds of cancer and other diseases have used it to help find previously unknown or undocumented causes for a disease," says Cronin.

UNLOCKING PATTERNS

AI can also help in finding unknown unknowns, to use Dick Cheney's famous phrase. Cronin believes that artificial intelligence's key ability to pattern match can be used to head off the next global pandemic – or pick up the first tremors of the next outbreak, allowing governments and health services crucial time to respond to the wave before it breaks on shore. "That will be immensely helpful," she says.

Those are just the first vestiges of how healthcare can be totally changed by the power of AI. "In the next couple of years, we'll hear a lot from pharma companies and healthcare in general saying how useful it is," she says. However, Cronin points out that to bring about a fair and equitable future, the way patients' health data is handled using generative AI needs to be carefully considered to preserve privacy.

According to British futurist Tom Cheeswright, researching medicine is also one of the main ways AI can help humanity in the short and long term. "If you look at the evolution of how we research medicines, we've gone from capturing bits of nature and sticking them in test tubes to see



what they do, to bulk processing lots of samples to see whether any of them can do anything useful against a pathogen, to being able to 3D model those,” he says. AI can do that at a faster, bigger scale. “It can make leaps of instinct around what might work in terms of attacking different diseases, based on an effectively learned knowledge from the history of medicine from Galen onwards.”

LIGHT FROM DARKNESS

The ability to survey information at speed and scale isn’t just unlocking new ideas in healthcare. Cheese-wright is working with a renewable energy company in India that’s investigating using generative AI to carry out theoretical bioengineering at the speed of recent vaccine development against covid in order to create more efficient biofuels such as algae.

Relatively inefficient industries like steel production can be made



It can make leaps of instinct around what might work in terms of attacking different diseases, based on an effectively learned knowledge from the history of medicine from Galen onwards.”

TOM CHEESEWRIGHT

British futurist

more effective by workshopping blue sky ideas in an AI context and taking the most promising outcomes out of the computer and into real, live trials. “You can absolutely see that generative AI assisting human researchers could accelerate progress in that domain,” says Cheese-wright.

The new era of AI tools could also bridge the gap between the potential and the reality of existing, known AI-powered support. Apple’s Siri smart assistant and other devices, including Amazon’s Alexa smart speakers, were presented as world-changing ideas, but the reality fell short of the vision. That has changed. “The vision I see for assistants for humans is it will be like your friend, your educator, your instructor, and your recommender,” Cronin says. “It’ll be almost like a person that would be at hand and know you really well – perhaps a lot better than the people you think know you.”

UTOPIA: A MANIFESTO

How do we reach the utopian future, where generative AI and other artificial intelligence tools are helping unlock meaningful, impactful areas rather than tinkering around the edges? How do we turn the life-changing benefits already being realized at an individual level to cross-cultural, worldwide ones?

POLITICAL IMPETUS


All this depends on how willing politicians are to recognize the potential of the technology, while ensuring it doesn’t interfere negatively in people’s lives. Futurist Tom Cheese-wright is hopeful that politicians strike the right tone in the utopian AI future. “A lot of the conversation about AI regulation has been in the tone of nuclear regulation,” he says. “That would only make sense if you could download plutonium online, along with designs for a reactor and build it out of components that you could buy.” Instead, Cheese-wright suggests a regulatory approach appropriate to the challenge at hand, flexible to the needs of the sector.

For instance, regulating AI in the healthcare space is likely to need to be more stringent because of the negative effects of getting it wrong. But in sparking innovative thinking in industrial design, it can be a lighter touch regime.

I, ROBOT

Getting it right is important, because the AI-powered future could be a truly transformational world. It could rewrite the way we work, freeing up workers to spend more time with their families. It could help us live longer, and develop drugs that cure diseases quicker. It could predict the weather, and prevent the worst damage from natural disasters.

And if the current crop of politicians don’t get it right? Cronin has a utopian vision for that, too. “I see a little further than a lot of people, and I get a lot of flak,” she says. “But I think in the future – say 20 or 25 years – I believe AI will be good enough, given the data it needs, to be the next senator or the next legislator. Better yet, she says: “It won’t be political.”



FOLLOWING BEST PRACTICE

Believing in the power of AI is something we all share. But following best practice will require carefully balancing the restraint of regulation with the chaos of change. What role will politicians, companies and everyday users play in that process? →



GOOD REGULATION

In the cat-and-mouse game of innovation and regulation, what is a sustainable route forward?

→ **FIRST WE INVENTED** motorcars. Then we began killing people. So we introduced rules on their use. Later we invented social media. Soon it began hurting people and shaping our society in ways we don't like. So we started regulating it.


Generative AI in the form many of us know it is barely just a year old. While plenty of tools predated ChatGPT, it was the text-based chatbot that really thrust the technology into the public consciousness.

And chastened by past mistakes in allowing too loose a leash for emergent technologies – including cars and social media – regulators have quickly ramped up their response to recognize that something needs to be done.

AI regulation is coming thick and fast in all corners of the world. The European Union is passing its AI Act through the legislative process, while the US and China are consulting on rules for generative AI with the public and stakeholders. The UK is toying with rules of its

own, while countries worldwide think about how to react to the challenges the technology poses.

In response, the companies are working equally rapidly to try and prove they're responsible actors. The Frontier Model Forum, set up in July 2023, is an industry body designed to "promote the safe and responsible development of frontier AI systems: advancing AI safety research, identifying best practices and standards, and facilitating information sharing among policymakers and industry." ↘



THE GLOBAL RESPONSE

Both sides have set out their stalls for what they envisage the future of generative AI to be, and how they expect it to be regulated. But which side will win? It's difficult to predict, but it's instructive to see how governments globally are approaching the challenge.

"There are broadly two ways in which AI regulation is made: at the national or local level, and at the international level," says Harry Law, a researcher at the University of Cambridge. "At the international level, there's the supranational agreements that set guardrails for AI development, like the EU's AI Act." Countries are also introducing national-level export controls and other restrictions governing interactions between companies and states – such as those introduced by the Biden administration targeting firms dealing with China.

Law believes that, going forward, we could see more cross-border cooperation when it comes to regulation. "In the future, we might see bilateral or multilateral deals between countries that are analogous to the New START Treaty, or the development of new international organizations focused on the governance of AI," he says.

But at the same time, we're likely to see individual countries filling in the gaps that international treaties and agreements on how to regulate AI don't cover in their specific jurisdictions, either for cultural or national reasons. Law divides these into "horizontal" and "vertical" pieces of legislation. Horizontal laws are designed to govern AI's use in all applications in a particular country, while vertical laws govern AI in a specific sector.

KEEPING COPYRIGHT SAFE

An example of an area that is being considered vertically? How business



UNITED STATES

In October 2022, the US released its Blueprint for an AI Bill of Rights, which provides a high-level overview of regulatory goals. Consultations are ongoing with stakeholders to develop rules that work for all parties.



BRAZIL

In December 2022, members of the Brazilian Senate presented draft AI rules focused on three pillars: enshrining human rights, classifying AI tools with different levels of risk, and predicting governance measures for AI providers.



UNITED KINGDOM

The UK is taking a "pro-innovation approach to AI regulation," says a white paper its science, innovation and technology ministry presented to parliament in March 2023. It is also hosting global summits on the technology.



EUROPEAN UNION

Including last-minute provisions for generative AI, the EU's AI Act, first proposed in 2021, seeks to protect the market of 450 million users. AI tools are bundled into different risk levels, the severity of which dictates the level of regulatory oversight.



CHINA

Because of its walled society, China leads the way with generative AI-specific regulation that protects users while ensuring responsible AI development. The 24 guidelines companies must follow to register their services with government were introduced on 15 August 2023.



AI IS FOCUSING LEGISLATIVE MINDS WORLDWIDE

1 bill mentioning AI passed into law in 2016

37 bills mentioning AI passed into law in 2022

SOURCE: STANFORD UNIVERSITY AI INDEX

6.5x increase of AI mentions in global legislative proceedings

and generative AI interact – and particularly the rights of businesses to protect their intellectual property. “Since the emergence of generative AI technologies, challenges surrounding data protection and intellectual property, such as the control over personal or public data and copyright, have been exposed,” says Janis Wong, tech policy advisor at the UK’s Law Society.

Many countries are starting to issue general guidance on whether and where generative AI should be used, with some context-specific consultations – for instance, the UK is consulting on the use of AI in education. “As the potential and limitations of these technologies become more apparent over the next months, it would be interesting to see whether different organizations and sectors issue specific generative

AI guidance to supplement these existing legal regimes,” Wong says. Wong’s organization, The Law Society, has published its own take on where AI regulation should head. “It is recommended that regulatory approaches to AI, including generative AI, are a blend of adaptable regulation and firm legislation to balance safeguarding societal interests with technological progress,” she says. “High-risk contexts also need to be identified and regulated.”

TAMING THE AI BEAST

There is a consensus globally that something needs to be done to keep responsible AI development in check (for more on ethical AI, see page 48 Good Ethics). But whether it’s plausible in practice, rather than in theory, is another question. “It’s not

clear that any technology can ever be ‘fully regulated’ or whether full regulation would be desirable,” says Law. As he points out, regulation has costs, as well as benefits.

Law highlights a report by venture capital fund Air Street Capital, who recently made the case that underlying many arguments that UK AI regulation is “falling behind” compared to its competitors in the EU and elsewhere is the assumption the country needs to “keep up.” But the venture capital firm argues that, by taking a more hands-off approach, the country is helping encourage AI development in the UK, and stoking innovation that will benefit the world.

Policymakers and politicians have to make their own decisions about where to draw the regulatory boundaries around AI to encourage

THE PUBLIC DEMANDS REGULATION



SOURCE: ADA LOVELACE INSTITUTE

the continued development of the technology and harness its disruptive potential, while ensuring that disruption does not overthrow the hard work they’ve spent on building a cohesive, working society.

MUCH HIGHER STAKES

There is no doubt that challenges lie ahead in the development of AI. And historical precedent suggests that the old approach to big tech – letting them self-police, and trusting their instinct that they know when to put people before profit – has not worked.

But Law argues that there is a difference between this era of tech development and prior ones. The stakes are much higher – and companies, as well as elected representatives, recognize the perils of going wrong and facing

overly draconian crackdowns on their business. “Ultimately, the incentives of AI firms and politicians are aligned: both want to build technology that maximizes upside and minimizes risk,” Law says. “And while conflicts may arise about the most appropriate way of governing AI between developers and politicians, there may also be diverging viewpoints across the full AI value chain.”

Law points to examples where national or local laws may diverge from international regulations. “Resolving these challenges will need close coordination and communication to understand the capabilities of AI, their likely impact on society, and the different goals and perspectives of the groups involved,” he says. Everyone needs a seat at the table: politicians,

the companies behind the technology – and most importantly, the public, too.

CALL TO IMPACT

- 1 Good regulation needs to take into account the needs of the general public, as well as the interests of the companies developing AI technology and regulators.
- 2 Regulation needs to be drawn up to keep users safe from profit-driven overreach, while also providing enough freedom for innovation. And it needs to be done quickly.
- 3 Cross-border consensus on regulation, which can flex to account for regional or national cultural and societal differences, will provide more security for end users while encouraging responsible global development of AI tech.

GOOD ETHICS

How do we keep human decency at the forefront of any AI development?

→ **AT A RECENT AI** conference held in Canada, participants were encouraged to take part in one of two tracks. There was the general AI track, then there was the “responsible AI” track. “It really ticked me off,” says Sasha Luccioni of Hugging Face, an AI company, and someone who has been focused on the world of AI ethics.

For one thing, events in the responsible AI track were put into smaller auditoriums – perhaps an indication of how seriously the organizers took the issue. But it was also the simple fact that responsible AI was seen as a separate area to be siphoned off from the main discussions that rankled.

“It was like, ‘Let’s put all the people who care about ethical or responsible AI in this one room, and then the rest of the people elsewhere, we’re going to work on the real stuff,’” she says.

Luccioni believes it’s no longer sufficient to siphon off responsible AI development, and the ethics underpinning it, into a separate school of thought. Instead it needs to be baked into everything generative AI developers do, putting ethics first and profit second.

“Really,” she says, “all AI should be responsible at this point. It’s

kind of like making a distinction between safe cars and other cars. In order for a car to be on the road, it’s supposed to be a safe car.” She hopes that this will result in a more thoughtful consideration of ethics and sustainability within AI, rather than a laser-focus on technical capabilities or state-of-the-art accuracy.

THE CASE FOR ETHICS

Given the potential ubiquity of generative AI in the years to come, it’s important that it’s developed ethically, because of its potential to impact all our lives. From ensuring it’s free from bias – which we have discussed in Making AI more inclusive (page 30) – to guaranteeing that it’s possible to understand how and why AI models make the decisions they do, accountability is at the core of ethical development of AI.

It’s also important, ethically, to understand who is responsible if and when AI makes a mistake. Thinking about where and how AI systems will be deployed, and how their use accords with human rights, is another cornerstone of responsible AI development. So is ensuring access is available for all, and we don’t end up with a two-tier system with those who

WHAT IS ETHICS, ANYWAY?

The UK’s Turing Institute defines AI ethics as “a set of values, principles, and techniques that employ widely accepted standards of right and wrong to guide moral conduct in the development and use of AI technologies.”



↳ can afford to benefit from the timesaving powers of AI getting ahead, and those who cannot falling behind, will be vital.

These are all tricky questions, but no one said ethics was easy. Some of those researching the area argue that its development is made harder by the past sins of big tech companies in previous generations of technological development.

LEARNING FROM HISTORY

Developing ethical AI is easier said than done, if prior precedent is anything to go by. “Tech has become quite used to being able to do pretty much whatever it wants without asking

permission and that has been part of the ethos of the industry,” says Carissa Veliz, associate professor at the Institute for Ethics in AI at the University of Oxford. “And because it’s a bit of a technical sector, and because things are happening so quickly, regulators have been quite slow in figuring out how to set the rules for these kinds of companies.”

Veliz says the gap between the speed at which tech companies are developing their AI tools, and the speed at which regulators and politicians are able to throw a lasso around those companies, “creates a kind of chaos in which, again, companies can design pretty much whatever they want, let it loose into the world and just see what happens.”

Tech companies claim to be considering ethics in their development of AI – but as we explored earlier in the report, the bottom line can often feel like it comes before thinking about the world writ large. It’s a “people, then profit” approach. And historical



Tech has become quite used to being able to do pretty much whatever it wants without asking permission and that has been part of the ethos of the industry.”

CARISSA VELIZ

Associate Professor,
Institute for Ethics in AI, University of Oxford

events show that sometimes ethics can be a thorny point of debate.

Google’s ethical AI team still exists, but has significantly changed over time. Controversially, both its founders left the company after a big falling-out with the tech giant over the publication of an academic paper that highlighted some of the ethical issues around the development of AI. Timnit Gebru and Margaret Mitchell, who helped found the team within Google designed to keep ethics at the core of the company’s AI development, have since gone on to found the Distributed AI Research Institute (DAIR).

ETHICAL WRANGLING

At the heart of Google’s AI ethics falling-out was a debate over what should be the right approach to incorporating ethics into AI development – question continuing to vex the broader industry. AI ethics has become a research field of its own, with different tribes signing up to different approaches.

So what should be the model? And should AI ethics be an integral part of the development of the technology, or an adjunct – something tacked on at the end, just as Luccioni experienced a separate conference track in Canada? “It’s part of good design to be ethical,” says Veliz. “If you have a very powerful AI but it’s going to wreck your democracy, in what world is that the innovation we want to create?”

For Veliz, the answer is to come back to core, basic principles – and to start with an ethical framework from the beginning. “What we have to do is to think about the kind of society we want to be,” she says. “If it’s a liberal democracy, then what are the pillars of a liberal democracy and how do we protect them?” That can help guide the conversation – and ensure a truly ethical AI is developed for all.

If that isn’t done willingly by companies, there are always options to compel private organizations to take decisions in the interest of public good. ■

CALL TO IMPACT

1 The ethical development of generative AI requires active engagement with, and action from, the companies creating AI tools.

2 Ethical considerations should be baked into the development of AI tools from the beginning, respecting the dignity of users, and protecting social values and the public interest.

3 Companies or organizations that are deemed to be acting unethically, or do not take an ethical approach, should be promptly forced to act through regulation. Lessons should be learned from the history of tech.

GOOD USERS

People power can decide who will thrive – and who will struggle to survive – in the AI race.

→ **THE GENERATIVE AI REVOLUTION HAS BECOME A GOLD RUSH.** In the first quarter of 2023, 110 companies in the S&P 500 mentioned artificial intelligence – roughly double the five-year average of 57, and more than three times the 10-year average of 34. Each of them sought to capitalize on the buzz around the transformative technology.

And for good reason. In the first six months of 2023, more than \$14 billion of funding was plowed into generative AI start-ups across 86 separate deals.

Hundreds of companies exist across the generative AI landscape, all vying for the same business, same funding and same users. This is a boon for people power. Two-thirds of people worldwide tell Ipsos that products and services using artificial intelligence will profoundly change their daily lives in the coming three to five years, but those people overlook that we too can profoundly change which products and services succeed.

History shows us that in the race for AI supremacy, only those who manage to gain users will survive until the end. If a specific product or app tries to cheat or dominate us, we can – and will – abandon it. That’s what consumerism tells us, and what has worked in the industrial age. The AI age will likely be no different.

Who will thrive, and who won’t survive from the current crop of generative AI giants? That’s down to us. But for now, here’s a selective view of how the landscape shakes out – and a reminder that we as users have more options than we might first think in how to navigate these turbulent times. ■

AI 50 VISUALIZED

CONSUMER USES


- Entertainment
 - Character.AI
 - Midjourney
- Productivity
 - OpenAI
 - OpenAI ChatGPT 4.0
- neeva
- other
 - trigo
 - Uaabi

ENTERPRISE STACK

- General productivity
 - ADEPT
 - glean
 - tome
 - AlphaSense
- General & administrative
 - Ironclad
 - synthesia
 - eightfold.ai
- Sales & customer support
 - GONG
 - RevComm
 - Clari
 - PolyAI
- Marketing
 - Jasper
 - WRITER
- EPD/IT/Security
 - Moveworks
 - Abnormal
 - VECTRA
 - GitHub Copilot

INDUSTRY VERTICALS

- Law firms
 - Harvey.
- Creative
 - imagen
 - runway
 - Midjourney
 - descript
- Health
 - iz.ai
 - BAYESIAN HEALTH
 - insitro
 - PathAI
 - UNLEARN
- Defense
 - ANDURIL
 - SHIELD AI
 - SLINGSHOT AEROSPACE
 - VANNEVAR Labs
- Agriculture & climate
 - Pachama
 - FarmWise
- Construction
 - CANVAS



GOOD KIDS

What can we learn from the next generation about how to meet this unique AI moment?

→ **WHEN IT COMES TO DEVELOPING, ROLLING** out and regulating AI, the stakes could scarcely be higher. Throughout this report we've explored the society-changing impact generative AI will have on us all, from upending our global economy to changing the way we work and live.

We've tussled with the difficult decisions that politicians and policymakers, as well as the heads of big tech companies, need to consider when adapting to the impact of AI. But it can be easy to think about the problems of today, overlooking than the potential – and pitfalls – of tomorrow.

Having such a short-term view is a danger, though, because AI is here to stay. And its effects will be far longer-lasting than our current generation. Decisions taken now will cast a shadow for years to come.

Which is why we've traveled the world to speak to those whose lives will be most shaped by the actions we make now: children. Here, they give us their thoughts about how the AI revolution has already shaped up, what they're excited by and fearful of, and how they wish we'd think about this most transformative technology. ■



Ryan: AI is good. Say you're doing business, and you can't do it quick enough, and you need something done quickly, you could use AI so it can do it and hopefully it'll be good. It could also be done for homework, but that's cheating. Because what happens if it's smarter than the smartest person in the world? Then everyone would get top marks in their exam. No one would know how smart they were, or how not smart they were.

Kieran: If you're doing an essay or something, you can put the key features of what you need in it, and then use that essay as a plan that they've created using ChatGPT. It's good in a way, but also bad because it does it all for you and you don't have to do much.

In the future, I think it'll change the way people do work. We need to think about if people will get replaced by ChatGPT or something.

The people who make AI could maybe only allow it to be used by certain people. Maybe only people in parliament, or something.

CALL TO ACTION

AI has the potential to be transformative, but can demolish what makes each of us unique. Ensuring we keep humans human – and if necessary, limiting access to how it can be used – will be crucial to maintaining individuality.

Ryan, 8, Kieran, 11,
from the United Kingdom



I think it's very interesting because of how it works and how you can create stuff. You can even create a portrait of yourself just out of AI. That is just impressive.

But when I'm old enough, I think I would make lots and lots and lots of pictures. I love drawing. I'd make maybe a story with it.

With AI, you basically tell it what to do and it does it for you. When it comes to drawing, you're basically the AI. Your brain is like, "We're going to draw this," and it does it. Your brain is the person and you're the AI. I think AI will help me be more imaginative. I have a very good imagination. I love trying new things.

Some artists think drawing takes time, but this AI is quicker than you think. Some people will be out of a job, but it's going to be more interesting to try something brand new. Maybe those who lose their jobs could try starting their own AI business.

My life will be different to my mother's. It'll be different from the technology then to the technology now. But if the AI doesn't give me credit or say "This is by Lolly," I would get angry. You're basically using me for my own work. Maybe people could write their names on the bottom and say, "This is my work."

Still, I'm very excited. It's very creative. Very imaginative. Very fast. Very cool.

CALL TO ACTION

Ensuring the owners of intellectual property get credit and compensation when their work is used to train AI models will be vital. As will ensuring those displaced from jobs by the AI revolution are retrained.

Lolly, 10,
from South Africa

Generative AI is still being developed. There are many flaws with it, as well. Sometimes maybe generative AI might not understand your command to it properly, so it might not produce what you want exactly as you want it. But over the years, the usability will be refined, and it'll be a necessary part of our futures – especially in the next few decades.

Because of AI, us children and maybe some college students and people older, will have more resources as they're growing up, which can help them learn and make them interested in arts. I already like the fact you can ask any question and the AI answers it. But maybe at some point in time, I'm hoping we can make AI analyze data and predict the future.

But the first thing the world needs to do is improve access to the internet. Most of the world has access to the internet, but there are some rural areas here and there that don't have access. Internet access is necessary for generative AI to spread.

However, even if everyone is given access to the internet, it won't help that people haven't learned native languages. They'll have learned a different, village language, which could isolate them from the rest of the world.

CALL TO ACTION

It can be easy to forget in the West and the Global North that internet access is not always a given. Politicians should work to improve access worldwide. And where access is widely available, the language gap in generative AI (see page 30 Inclusive AI) can radically shift how useful it is for some.

Om, 13,
from India



PHOTOS: PRIVATE

TOGETHER WE CAN SHAPE THE FUTURE OF AI TO BENEFIT US ALL

**FII-I has three pillars to deliver its mission:
THINK, ACT and XCHANGE**

1 FII-I THINK
Identify societal challenges and current inhibitors. Curate the brightest ideas to address societal issues

2 FII-I ACT
Catalyze innovation and initiatives by mobilizing partners and resources

3 FII-I XCHANGE
Create platforms for live discussions on the future of humanity. Share knowledge, stories and publications with different stakeholders

→ **TECHNOLOGY HAS ALWAYS** moved quickly, and its impacts have long been profound. But with the advent of generative AI, we have entered a new paradigm. Earth-shattering advancements are occurring every week in AI, with ramifications for how we work, how we live and the future of our planet, given the brute force of hardware that powers machine intelligence.

So the time to act as a global community to establish what we want from artificial intelligence – and more importantly, what we don't want – is now.

Collectively, the world needs to decide in which direction our future will travel. Will we harness and tame AI, putting it to work to ease inequalities, shift burdens away from humans towards machines, and to better how we treat each other?

Will we use it to create new medicines, find new efficiencies, grow harder crops and train better workers? Or will we squander the opportunity, not making the most of the technology – or worse, develop it dangerously and slip into servitude?

The decision is up to us. And it's why the Future Investment Initiative Institute has chosen to focus on AI in this report.

ABOUT THE FII INSTITUTE

THE FUTURE INVESTMENT INITIATIVE (FII) INSTITUTE is a global nonprofit foundation with an investment arm and one agenda: Impact on Humanity. Global, inclusive and committed to Environmental, Social and Governance (ESG) principles, we foster great minds from around the world and turn ideas into real-world solutions in four key areas: Artificial Intelligence

THE FII INSTITUTE

is guided in all it does by a strong purpose, vision and mission.

PURPOSE
"Enabling a brighter future for humanity"

VISION
"Bringing together the brightest minds and most promising solutions to serve humanity"

MISSION
"Creating a purposeful present, promising future"

PHOTO: ADOBE STOCK, FII INSTITUTE

and Robotics, Education, Healthcare and Sustainability. We are in the right place at the right time: when decision-makers, investors and an engaged generation of youth come together in aspiration, energized and ready for change.

We harness that energy into three pillars: THINK, XCHANGE, ACT.

- Our THINK pillar empowers the world's brightest minds to identify technological solutions to the most pressing issues facing humanity.

- Our XCHANGE pillar builds inclusive platforms for international dialogue, knowledge-sharing and partnership.

- Our ACT pillar curates and invests directly in the technologies of the future to secure sustainable real-world solutions.

Join us to own, cocreate and actualize a brighter, more sustainable future for humanity. ■

BOARD OF TRUSTEES



MR. RICHARD ATTIAS
CEO FII Institute



H.E. YASIR AL-RUMAYYAN
Governor PIF, Chairman Saudi Aramco



HRH PRINCESS REEMA BINT BANDAR
KSA ambassador to the USA



PROFESSOR TONY CHAN
President of KAUST



H.E. MOHAMMED AL ABBAR
Founder & Chairman EMAAR properties



H.E. MATTEO RENZI
Italian senator for Florence (former PM)



DR. PETER H. DIAMANDIS
Founder X-Prize Foundation & Singularity University



PROFESSOR ADAH ALMUTAIRI
Professor Pharmaceutical Chemistry UCSD

SOURCES

Overview of key sources used for this report

State of AI Report 2022, Nathan Benach and Ian Hogarth

AI Index Report 2023, Stanford University

The state of AI in 2022 – and a half decade in review, McKinsey

The state of AI in 2023: Generative AI's breakout year, McKinsey

2023 State of Data + AI Report – 2023 State of Data + AI, Databricks

How AI Ate The World, Chris Stokel-Walker

Evaluating Verifiability in Generative Search Engines, Nelson F. Liu, Tianyi Zhang and Percy Liang

AI Improves Employee Productivity, Nielsen Norman Group

9831st UN Security Council meeting, 18 July 2023

Generative AI could raise global GDP by 7%, Goldman Sachs

OECD Employment Outlook 2023

Ipsos Global Views on AI 2023, Ipsos Mori

Lost in Translation: Large Language Models in Non-English Content Analysis, Center for Democracy & Technology

Public attitudes to AI, Ada Lovelace Institute

Staying the course – reflections ahead of the UK's AI Safety Summit, Air Street Capital

AI 50 2023, Sequoia Capital

IMPRINT

FUTURE INVESTMENT INITIATIVE INSTITUTE

PUBLISHER
Future Investment Initiative Institute
3884 Alakheel District Unit 9
Riyadh 12382 – 6613
Kingdom of Saudi Arabia

CONTACT
For clarifications and inquiries, kindly email:
info@fii-institute.org

EDITOR
Chris Stokel-Walker

“ GENERATIVE AI IS THE MOST POWERFUL TOOL FOR CREATIVITY THAT HAS EVER BEEN CREATED. IT HAS THE POTENTIAL TO UNLEASH A NEW ERA OF HUMAN INNOVATION. ”

ELON MUSK



Founding Partner



Vision Partner



Strategic Partners

