S.C.G. IS COMING

SPATIAL COMPUTING GLASSES / DATALINK TECHNOLOGY: FUTURE SMART WEARABLE TECHNOLOGY WILL REPLACE THE CELLPHONE IN POPULARITY AND FUNCTION BY 2030.

coroflot.com/GeneKWalker G.K. Walker, Design Research

AI Assistant features in active development by the major cell phone and personal computer company leaders include:

Micro-LCD Dual Lens or wrap-around screen design will feature translucent menus, glossy icons, and streamlined menus that create a seamless experience as you move from one task to the next.

A customizable, bespoke AI assistant that understands what's on your screen; all you do is voice-command your customized AI agent questions about what you're seeing, or have your AI assistant search for a product online, or automatically add an event to your calendar.

An intuitive, heuristic AI assistant that executes actions together, like having it transcribe a lecture, compare it to your notes, and then summarize the key points you missed.

Live Translation brings real-time language translation to Messages, FaceTime, and calls, with processing done locally on-device to maintain privacy.

Multimodal capabilities: read, interpret, and convert planning files (including blurry maps and handwritten notes) into digital formats.

AI Call Screen and Hold Assist, including real-time direct-to-voice mail for spam calls. Also, when making a voice-command call, if you are placed on hold and that dreaded hold music kicks in, you can tune it out and be notified when a live agent is available.

Multitask Partition Window System: Live camera view of surroundings provides added security; the user can send voice-command emails and create a note to play a podcast.

Your Artificial Intelligence assistant can help generate responses.

For example, a student or professional can more easily build a report to compare class or meeting notes to the audio transcription to see what may have been missed.

Also, a dedicated function to summarize text and create AI images.

AI is the real star of the SCG glasses. Spatial Computing Glasses and contact lenses integrate AI to make the eyewear an “always present AI assistant” that can see and hear what you do. With built-in cameras, mics, and speakers, the glasses will allow the AI assistant to interpret context and offer live help, translating a street sign or identifying a landmark. By leaning into the strengths of AI, the Spatial Computing Glasses become an active personal tool.

The frames feature a built-in display of text and graphics onto the lenses the user sees floating in their visual field, with real-time language translation, telepresence, GPS, and interaction with a customized AI assistant.

PREDICTIONS

MORE PROFESSIONAL WOMEN THAN MEN WILL APPRECIATE THE ADDED PERSONAL SECURITY OF THE EVOLVING TECHNOLOGY.

POTENTIAL APPLICATIONS ARE AS LIMITLESS AS THE USER'S IMAGINATION...

AS SCG TECHNOLOGY EVOLVES AND BECOMES MORE REFINED, SOPHISTICATED, AND MINIATURIZED, SMART WEARABLES WILL BE SEEN AS A SIGN OF PRESTIGE AND AN ESSENTIAL TOOL FOR SUCCESSFUL DAILY LIVING.

SCG INTEGRATES AI INTO EYEWEAR AS A BESPOKE PERSONAL AI ASSISTANT, ENABLING USERS TO PROCESS RAW DATA, NEWS, AND INFORMATION EFFICIENTLY.

WITH BUILT-IN MICRO CAMERAS, MICS, AND REAR SKULLBONE CONDUCTION SILENT SPEAKERS, THE GLASSES ALLOW THE AI AGENT TO INTERPRET CONTEXT AND OFFER REAL-TIME, CONCIERGE-STYLE ASSISTANCE.

THE USER SEES IN THEIR VISUAL FIELD REAL-TIME TRANSLATION OF FOREIGN LANGUAGES, OPEN-SOURCE DATA, DOCUMENT DICTATION AND REVIEW, SECURE TELEPRESENCE, GPS POSITIONING AND NAVIGATION, AND 24/7 INTERACTION WITH YOUR AI AGENT.

How Sound Waves Are Conducted Through the Bones of the Skull

Sound typically travels through the air to our ears. While this pathway is well-known, sound also reaches our inner ear through the bones of the skull. This article explores this alternative route of sound transmission.

Sound's Usual Journey and Beyond

Sound usually reaches our inner ear through air conduction. Sound waves travel through the air into the outer ear, funneled through the ear canal to the eardrum. The eardrum, a thin membrane, vibrates in response to these incoming sound waves. These vibrations are then transferred to three tiny bones in the middle ear: the malleus (hammer), incus (anvil), and stapes (stirrup). These ossicles amplify and transmit the vibrations to the oval window, a membrane that separates the middle ear from the fluid-filled inner ear.

Bone conduction offers a distinct pathway for sound transmission. In bone conduction, sound vibrations bypass the outer and middle ear structures entirely. Instead, these vibrations are transmitted directly through the bones of the skull to the inner ear. This alternative mechanism allows sound to be perceived even if the ear canal is blocked or if there are issues with the middle ear.

Holographic-inspired lenses could unlock '3rd dimension of imaging' in future VR headsets and smart glasses

News By Andrea Saravia Pérez published November 7, 2024

Future VR headsets could use a new type of lens inspired by holographic devices. The bilayer bifocal lens relies on external voltage to change the intensities in the foci.

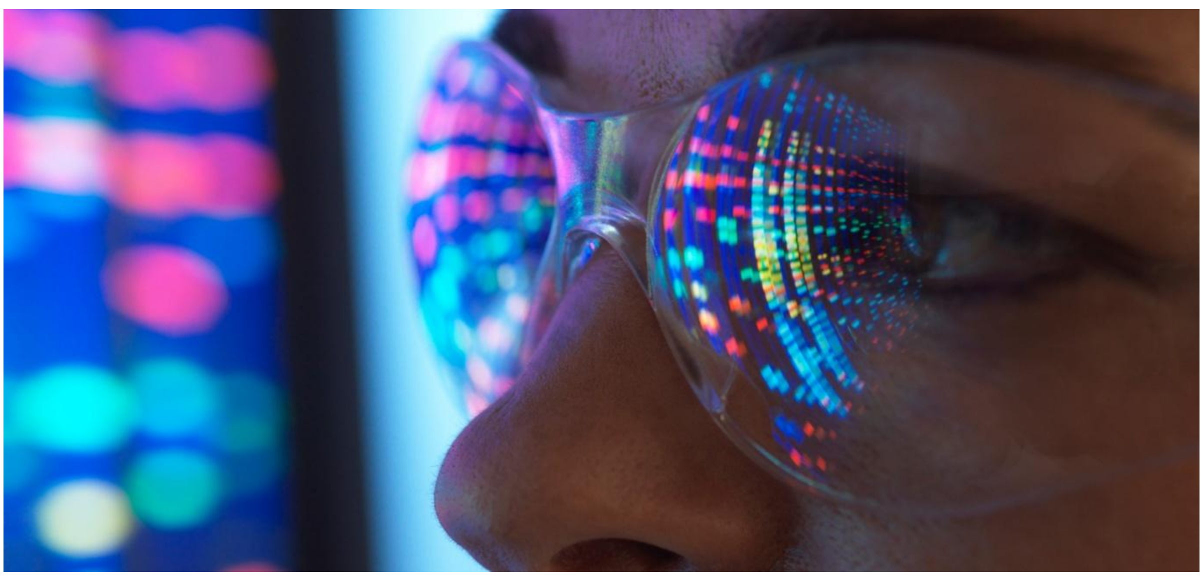

(Image credit: Andrew Brookes/Getty Images)

Future virtual reality (VR) headsets could use a new type of lens inspired by holographic devices, researchers in China say. This proposed new type of bifocal lens can switch between one focus and another at the flick of a switch, letting wearers witness intensities in the lenses change akin to a hologram.

These lenses would be made from two layers of liquid crystal structures that can switch between two foci with external voltage. The researchers described their findings in a new study published Oct. 1 in the journal Optics Letters.
The technology could be applied in imaging devices, optical computing and optical interconnectivity — in addition to future mixed-reality and VR headsets, the team said.

Specifically, they can be used for polarization imaging — often used to enhance image contrast or for edge imagining, which highlights the outline of objects or see finer details. Polarization can be referred to as the light's third dimension, with polarization cameras often detecting physical properties unseen by conventional imaging.

"We believe that the light control mechanism we created using the multilayer structure could also be used to design other optical devices, including holographic devices and beam generators, or for optical image processing," study lead author Fan Fan, professor of physics and electronics at Hunan University, said in a statement.

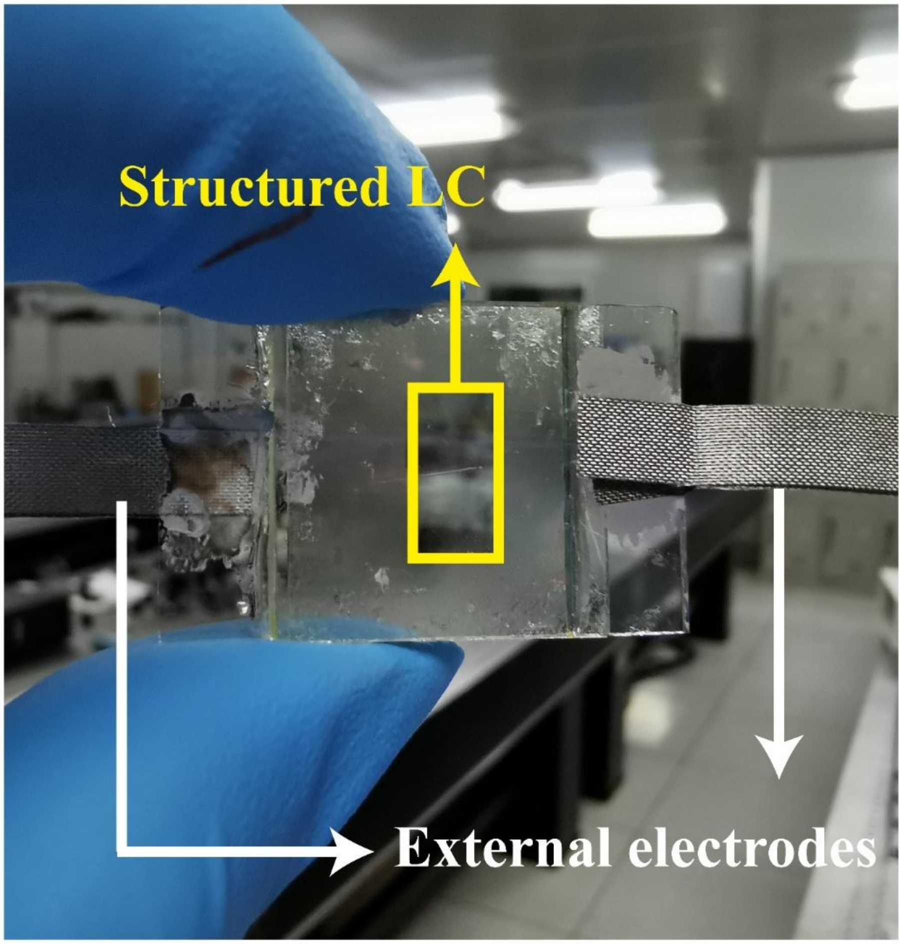

Researchers developed a bifocal lens based on two layers of liquid crystal (LC) structures. The intensities for the two focal lengths can be easily adjusted by applying external voltage. (Image credit: Fan Fan, Hunan University)

The team focused on creating bilayer structures, rather than single-layer structures, which most liquid-crystal devices are made from. The structures were made from a liquid crystal cell as well as a liquid crystal polymer — both of which are standard materials in the development of lenses intended for holographic imaging uses. This let researchers alter the intensity of the two foci. The development of multi-functional holographic devices was an inspiration but the technology could be used "beyond the field of holographic displays," Fan said.

Some bifocal lenses can create different focal points based on the incident light's polarization — but this new design enables the focus to switch on command, manipulating the polarization states of the output beams. The liquid crystal layer also allows the lenses to switch rapidly between focus points when voltage is applied.

The scientists are planning to use the new lenses in several kinds of multifunctional devices. For the optical components found in this technology to be more practical, however, they stressed the cost of the mass production of components would have to be lowered. If this happens, the team could design and incorporate fast and accurate layer-to-layer alignment technology, the researchers said in the statement.

What are neural processing units (NPUs) and why are they so important to modern computing?

Features By Tim Danton published May 12, 2025

Neural processing unts (NPUs) are the latest chips you might find in smartphones and laptops — but what are they ard why are they so important?

Ever since the dawn of computing, people have compared machines to brains. This includes two founding fathers of computing — John von Neumann wrote a book called "The Computer and the Brain" while Alan Turing was quoted in 1949 saying: "Eventually I do not see why [a computer] may not compete on equal terms with the human intellect in most fields."

The only problem with this comparison is that the traditional processor — the central processing unit (CPU) — doesn't mimic the brain at all. CPUs are far too mathematical and logical. The neural processing unit (NPU), on the other hand, takes an entirely different approach: simulating the structure of the human brain in its very circuitry.

Yet mimicking the workings of the human brain electronically is far from a new idea.

The birth of NPUs

Literal electronic brains date back to the birth of modern-day computing in the mid-1940s, specifically to a "neural network" of circuitry created by neurophysiologist Warren McCulloch and logician Walter Pitts. McCulloch's pioneering work spurred further research during the 1950s and 1960s, only for the idea to fall out of fashion — perhaps due to a lack of progress compared to the rising number-crunching power of classical computers. "There were a few isolated people in Japan and Germany [working on neural networks] but it was not a field," Yann LeCun, a French-American computer scientist widely considered one of 'godfathers' of AI, said of his time working with Geoffrey Hinton, another of the field's pioneers, on neural networks in the early 1980s. "It started being a field again in 1986."

Yet for neural networks to regain their foothold as a respected part of computer science, it took the success of speech recognition in the early 2000s., Even then, LeCun said: "We didn't want to use the word neuron nets because it had a bad reputation, so we changed the name to deep learning."

The term NPUs would come in the late 1990s, but it has taken the deep pockets of Apple, IBM and Google to move it from university labs and into the mainstream. These tech companies invested billions of dollars into the development of silicon, crystallizing all the past work into a product that fits inside our phones and laptops: a processor that takes inspiration from the human brain. LeCun's fortunes have also improved for the better: he is now chief AI scientist at Meta.

How do NPUs work?

In some ways, the NPUs of today aren't that different from those created by McCulloch and Pitts: their structure mimics the brain through a parallel architecture. This means that rather than tackling a problem in sequence, an NPU will simultaneously run millions, even trillions, of mini computations simultaneously. This is what the term "tera operations per second," or TOPS, refers to.

But here's where things get complicated. NPUs rely on deep learning instruction sets, which have already been trained on vast amounts of existing data. Take the example of edge detection in photos, which usually relies on convolutional neural networks (CNNs).
In a CNN, the convolution layer runs a filter (called a "kernel") over every area of the image, which will hunt for patterns that it suspects — thanks to its training — are edges. Each mathematical operation the NPU performs is called a convolution, which creates a feature map over the image. The software will repeat this process until it reaches the point where it is confident it has found edges.

NPUs are outstanding at performing convolutionary operations, being able to execute them at great speed and with low power demands. This is especially true when compared to CPUs. However, graphics processing units (GPUs), which also use parallel processing, are less optimized for this task and therefore less efficient. This drop in efficiency makes a big difference when it comes to the battery life of our devices.

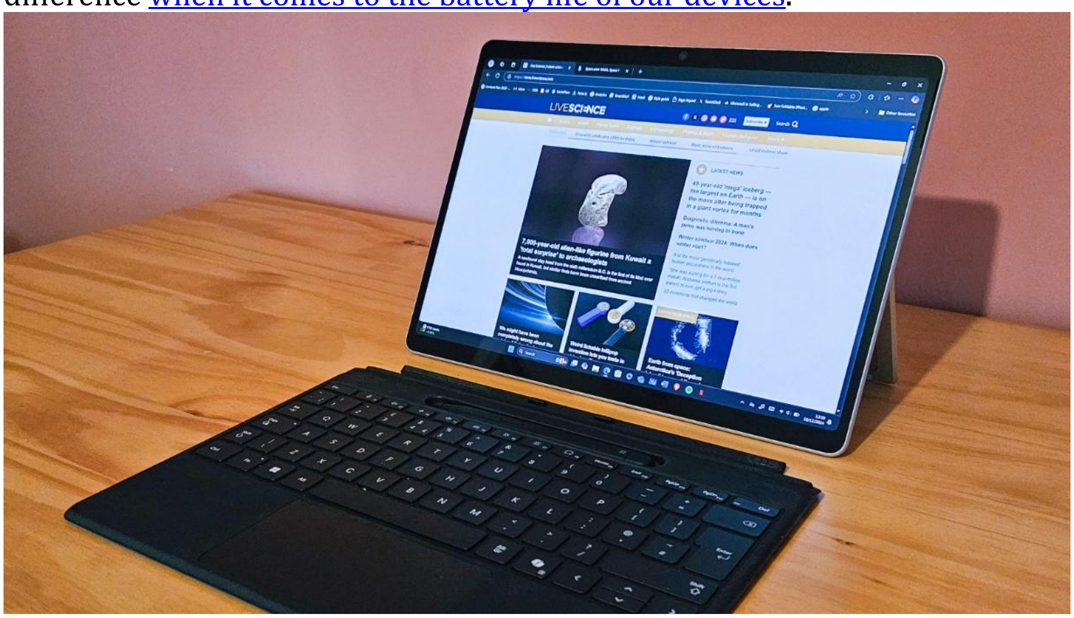

The Microsoft Surface Pro 11 makes use of NPUs. (Image credit: Keumars Afifi-Sabet/Future)

What are NPUs now used for?

Perhaps surprisingly, the first phones to include an NPU date back to 2017. That's when Huawei released the Mate 10 and Apple debuted its A11 Bionic chipset in the iPhone X. But neither of these NPUs was very powerful — each having less than 1 TOPS compared to the 45 TOPS NPUs in a modern-day Qualcomm Snapdragon X chipset, fitted into our best laptops. It has also taken several years for applications to appear that can take advantage of the chips' unique structure.

Yet just eight years later, AI applications are everywhere. For example, if you own a recent phone that includes the option to remove people from photos — that almost certainly uses an NPU. Likewise, Google's "Circle to Search" feature, or "Add Me" uses a NPU-powered form of augmented reality (AR) to place you in the photo after you've already taken the original shot.

NPUs have now spread to laptops too. Last year, Microsoft announced "a new category of Windows PCs designed for AI, Copilot+ PCs." These required NPUs with at least 40 TOPS, which unfortunately for AMD and Intel (whose early NPUs only ran at 15 TOPS), ruled them out of the race. But their loss was Qualcomm's gain, as all of its Snapdragon X processors exceeded that threshold with NPUs rated at 45 TOPS. Models that take advantage of these new chips include the Microsoft Surface Laptop and Snapdragon versions of the Acer Swift AI series.

Both AMD and Intel have now released chips that meet Microsoft's minimum requirements, meaning far more laptops are on the market with the "Copilot+ PC" branding. Yet there's a sting in this tail: more affordable laptops (less than $800) are now likely to still use lesser processors that don't meet the Copilot+ PC criteria.

What are the best Copilot+ PC features?

But why should you pay more for a Copilot+ PC? Microsoft hopes to tempt you with a number of exclusive features, and frankly, the most impressive one is also the most controversial. Called Recall, this promises a "photographic memory" that enables you to rediscover something you've previously seen in Windows 11.

Each snapshot taken by Recall is analysed by the NPU, using context, optical character recognition (OCR) and sentiment analysis to create an index that you can then search — at which point Recall will take you back in time through a visual timeline. After a shaky launch, where it was attacked for the lack of security or user control over what snapshots were stored, Microsoft said it spent more time reworking the feature to be more secure.

Other features build upon what has come before. Image Creator uses the NPU to turn text into images, an enhanced version of Windows Studio Effects adds creative filters to your video calls and Live Captions deploys the NPU to translate any video you're watching.

Companies like Acer, HP and Lenovo have released their own local AI tools that can analyse documents stored on your PC and provide summaries and sentiment analysis. Such tools are only likely to improve over time.

What's likely to happen next with NPUs?

For the next few years, some AI experts contend that NPUs will follow a similar path to CPUs in their early days — something close to Moore's Law, with a doubling of TOPS every year or two. With that power will come far greater abilities, to the point where you can create realistic AI artwork locally on your computer rather than resorting to programs such as Midjourney.

Over time, as software matures along with the hardware, and more developers take advantage of it, we expect to see the emergence of personal AI agents that understand us because they have been "living" inside our computers as we work. They won't just serve as memory joggers but perform actions on our behalf.

NPUs will also likely find a home in more devices than our phones and laptops. TVs will produce personalized news services using your favourite avatar presenter; your fitness tracker will recommend workouts based on your mood and the time until your next meeting. Who knows, your best friend may one day be a humanoid robot who understands you better than any human.

Research on 40 hertz brainwaves has shown that sensory stimulation at this frequency can **improve brain health**, bolster cognitive function, foster neuron growth, and combat neurodegenerative disease in both humans and animals[1].

Human oscillatory brain activity near 40 Hz coexists with cognitive temporal binding.

M Joliot [1], U Ribary [1], R Llinás [1]

PMCID: PMC45309 PMID: 7972135

Abstract

Spontaneous oscillatory electrical activity at a frequency near 40 Hz in the human brain and its reset by sensory stimulation have been proposed to be related to cognitive processing and to the temporal binding of sensory stimuli. These experiments were designed to test this hypothesis and to determine specifically whether the minimal interval required to identify separate auditory stimuli correlates with the reset of the 40-Hz magnetic signal. Auditory clicks were presented at varying times, while magnetic activity was recorded from awake human subjects. Experimental and modeling results indicate a stimulus-interval-dependent response with a critical interval of 12-15 ms. At shorter intervals only one 40-Hz response, to the first stimulus, was observed. With longer intervals, a second 40-Hz wave abruptly appeared, which coincided with the subject's perception of a second distinct auditory stimulus. These results indicate that

oscillatory activity near 40 Hz represents a neurophysiological correlate to the temporal processing of auditory stimuli. It also supports the view that 40-Hz activity not only relates to primary sensory processing, but also could reflect the temporal binding underlying cognition.

40 Hz Brain Waves: Implications for Brain Health

Explore how 40 Hz brain waves relate to cognitive function, neural processing, and auditory perception, with insights from current research. Published Apr 29, 2025

Brain activity is characterized by rhythmic electrical patterns, with different frequencies linked to various mental states and functions. Among these, 40 Hz brain waves have gained attention for their role in cognitive performance and neurological health. Research suggests they contribute to memory, perception, and neuroprotection against conditions like Alzheimer's disease.

Understanding their influence on cognition and sensory processing could inform therapeutic applications and cognitive enhancement strategies.

Relationship to Gamma Brain Waves

Gamma brain waves, oscillating between 30 and 100 Hz, are linked to higher cognitive functions, including attention, perception, and memory consolidation. Within this range, 40 Hz activity is particularly significant, synchronizing neural activity across brain regions to enhance communication between cortical and subcortical structures. This synchronization is essential for integrating sensory information and maintaining coherent cognitive processing.

Studies using electroencephalography (EEG) and magnetoencephalography (MEG) reveal that 40 Hz activity is prominent in tasks involving working memory and selective attention. A study in Nature Neuroscience found increased 40 Hz power in the prefrontal and parietal cortices during visual discrimination tasks, indicating a role in top-down cognitive control. Disruptions in gamma synchronization, particularly at 40 Hz, are observed in neurological disorders like schizophrenia and Alzheimer's disease, reinforcing their importance in brain function.

The generation of 40 Hz gamma waves results from interactions between excitatory pyramidal neurons and inhibitory interneurons, particularly parvalbumin-expressing interneurons. These fast-spiking interneurons regulate neuronal firing, creating rhythmic oscillations that enhance signal transmission and reduce background noise. Studies using optogenetics show that activating these interneurons at 40 Hz improves cognitive performance in animal models, highlighting their role in gamma synchronization and cognitive health.

Neural Mechanisms

The generation and propagation of 40 Hz brain waves rely on interactions between excitatory and inhibitory neuronal circuits, particularly within the neocortex and hippocampus. Parvalbumin-expressing interneurons synchronize neuronal firing by providing rhythmic inhibitory inputs, maintaining an excitatory-inhibitory balance crucial for working memory and sensory integration.

Large-scale network dynamics also contribute to 40 Hz oscillations. The thalamocortical system plays a role in maintaining gamma coherence across brain regions, with the thalamus modulating cortical rhythms to stabilize 40 Hz activity. This coordination strengthens sensory processing and perceptual accuracy, particularly during attention-demanding tasks. Functional imaging studies confirm that gamma synchronization between the thalamus and cortex correlates with improved cognitive performance.

On a molecular level, neurotransmitter systems influence 40 Hz rhythm stability. Gamma-aminobutyric acid (GABA) is central to this process, with reduced inhibitory signaling leading to desynchronized neural activity. In schizophrenia, disruptions in GABAergic interneuron function impair gamma oscillations, affecting attention and working memory. Cholinergic input from the basal forebrain further enhances gamma synchronization, underscoring the role of neuromodulatory systems in cognitive function.

Role in Auditory Processing

40 Hz brain waves play a key role in auditory processing, shaping how the brain interprets and organizes sound. These oscillations synchronize neural populations involved in detecting temporal patterns, essential for speech perception and music appreciation. When auditory stimuli align with this frequency, cortical responses become more synchronized, improving sound clarity and discrimination. This is particularly beneficial in noisy environments, where precise timing helps filter relevant auditory information.

Neural entrainment to 40 Hz auditory stimuli occurs in both cortical and subcortical structures. The auditory cortex, particularly primary and secondary regions, exhibits enhanced phase-locking to 40 Hz modulated tones, reinforcing temporal sound processing. Subcortical structures like the inferior colliculus and medial geniculate body contribute to early encoding, ensuring temporal integrity before higher-order processing. EEG studies show that individuals with stronger 40 Hz auditory steady-state responses perform better in speech-in-noise tasks, highlighting their functional significance.

Disruptions in 40 Hz auditory processing are linked to neurological conditions affecting speech and communication. In schizophrenia, deficits in gamma synchronization impair auditory perception, making phoneme distinction and speech tracking difficult. Age-related hearing decline is also associated with reduced 40 Hz phase-locking, contributing to difficulties in understanding spoken language in noisy settings. These findings suggest that interventions targeting gamma synchronization, such as auditory stimulation techniques, may help mitigate auditory processing deficits.

Binaural Beat Phenomena

Binaural beats occur when two slightly different frequencies are presented separately to each ear, causing the brain to perceive a third, phantom tone corresponding to the frequency difference. When this difference is 40 Hz, the brain may synchronize with this rhythm, potentially influencing cognitive function and sensory processing.

Neuroscientific studies using EEG suggest that binaural beats at gamma frequencies may alter cortical activity. Some research indicates increased gamma power in frontal and temporal regions, areas linked to attention and information processing. However, individual responsiveness varies based on baseline brainwave patterns, attentional state, and auditory sensitivity. While some studies report cognitive benefits from 40 Hz binaural beats, others find minimal effects, highlighting the need for further research.

Cognitive Correlates in Research

Research on 40 Hz brain waves links them to cognitive performance, particularly in memory retention, attention regulation, and perceptual integration. EEG and MEG studies show that tasks requiring sustained focus or rapid information processing elicit increased 40 Hz activity in the prefrontal cortex and hippocampus, suggesting a role in encoding and retrieving information. A study in The Journal of Neuroscience found that individuals with stronger 40 Hz power during memory tasks exhibited better recall accuracy, reinforcing their role in cognitive function

Emerging research explores 40 Hz stimulation as a potential intervention for cognitive decline. Clinical trials using non-invasive techniques like transcranial alternating current stimulation (tACS) and 40 Hz light flicker therapy report promising results in individuals with mild cognitive impairment. A study in Neuron found that daily exposure to 40 Hz visual and auditory stimulation increased gamma coherence and improved cognitive assessments over several weeks. These findings suggest that modulating 40 Hz activity could enhance neural plasticity and support cognitive resilience, particularly in aging populations at risk for neurodegenerative conditions.

Evidence that gamma rhythm stimulation can treat neurological disorders is emerging

Researchers survey a broadening landscape of studies showing what's known, and what remains to be found, about the therapeutic potential of noninvasive sensory, electrical, or magnetic stimulation of gamma brain rhythms.

David Orenstein | The Picower Institute for Learning and Memory
Publication Date: January 18, 2024

A surprising MIT study published in *Nature* at the end of 2016 helped to spur interest in the possibility that light flickering at the frequency of a particular gamma-band brain rhythm could produce meaningful therapeutic effects for people with Alzheimer's disease. In a new review paper in the *Journal of Internal Medicine*, the lab that led those studies takes stock of what a growing number of scientists worldwide have been finding out since then in dozens of clinical and lab benchtop studies.

Brain rhythms (also called brain "waves" or "oscillations") arise from the synchronized network activity of brain cells and circuits as they coordinate to enable brain functions such as perception or cognition. Lower-range gamma-frequency rhythms, those around 40 cycles a second, or hertz (Hz), are particularly important for memory processes, and MIT's research has shown that they are also associated with specific changes at the cellular and molecular level. The 2016 study and many others since then have produced evidence, initially in animals and more recently in humans, that various noninvasive means of enhancing the power and synchrony of 40Hz gamma rhythms helps to reduce Alzheimer's pathology and its consequences.

"What started in 2016 with optogenetic and visual stimulation in mice has expanded to a multitude of stimulation paradigms, a wide range of human clinical studies with promising results, and is narrowing in on the mechanisms underlying this phenomenon," write the authors including Li-Huei Tsai, Picower Professor in The Picower Institute for Learning and Memory and the Department of Brain and Cognitive Sciences at MIT.

Though the number of studies and methods has increased and the data have typically suggested beneficial clinical effects, the article's authors also clearly caution that the clinical evidence remains preliminary and that animal studies intended to discern how the approach works have been instructive, but not definitive.

"Research into the clinical potential of these interventions is still in its nascent stages," the researchers, led by MIT postdoc Cristina Blanco-Duque, write in introducing the review. "The precise mechanisms underpinning the beneficial effects of gamma stimulation in Alzheimer's disease are not yet fully elucidated, but preclinical studies have provided relevant insights."

Preliminarily promising

The authors list and summarize results from 16 clinical studies published over the last several years. These employ gamma-frequency sensory stimulation (e.g., exposure to light, sound, tactile vibration, or a combination); transcranial alternating current stimulation (tACS), in which a brain region is stimulated via scalp electrodes; or transcranial magnetic stimulation (TMS), in which electric currents are induced in a brain region using magnetic fields. The studies also vary in their sample size, design, duration, and in what effects they assessed. Some of the sensory studies using light have tested different colors and different exact frequencies. And while some studies show that sensory stimulation appears to affect multiple regions in the brain, tACS and TMS are more regionally focused (though those brain regions still connect and interact with others).

Given the variances, the clinical studies taken together offer a blend of uneven but encouraging evidence, the authors write. Across clinical studies involving patients with Alzheimer's disease, sensory stimulation has proven safe and well-tolerated. Multiple sensory studies have measured increases in gamma power and brain network connectivity. Sensory studies have also reported improvements in memory and/or cognition, as well as sleep. Some have yielded apparent physiological benefits such as reduction of brain atrophy, in one case, and changes in immune system activity in another. So far, sensory studies have not shown reductions in Alzheimer's hallmark proteins, amyloid or tau.

Clinical studies stimulating 40Hz rhythms using tACS, ranging in sample size from only one to as many as 60, are the most numerous so far, and many have shown similar benefits. Most report benefits to cognition, executive function, and/or memory (depending sometimes on the brain region stimulated), and some have assessed that benefits endure even after treatment concludes. Some have shown effects on measures of tau and amyloid, blood flow, neuromodulatory chemical activity, or immune activity. Finally, a 40Hz stimulation clinical study using TMS in 37 patients found improvements in cognition, prevention of brain atrophy, and increased brain connectivity.

"The most important test for gamma stimulation is without a doubt whether it is safe and beneficial for patients," the authors write. "So far, results from several small trials on sensory gamma stimulation suggest that it is safe, evokes rhythmic EEG brain responses, and there are promising signs for AD [Alzheimer's disease] symptoms and pathology. Similarly, studies on transcranial stimulation report the potential to benefit memory and global cognitive function even beyond the end of treatment."

Studying underlying mechanisms

In parallel, dozens more studies have shown significant benefits in mice including reductions in amyloid and tau, preservation of brain tissue, and improvements in memory. But animal studies also have offered researchers a window into the cellular and molecular mechanisms by which gamma stimulation might have these effects. Before MIT's original studies in 2016 and 2019, researchers had not attributed molecular changes in brain cells to changes in brain rhythms, but those and other studies have now shown that they affect not only the molecular state of neurons, but also the brain's microglia immune cells, astrocyte cells that play key roles in regulating circulation, and indeed the brain's vasculature system. A hypothesis of Tsai's lab right now is that sensory gamma stimulation might promote the clearance of amyloid and tau via increased circulatory activity of brain fluids.

A hotly debated aspect of gamma stimulation is how it affects the electrical activity of neurons, and how pervasively. Studies indicate that inhibitory "interneurons" are especially affected, though, offering a clue about how increased gamma activity, and its physiological effects, might propagate.

"The field has generated tantalizing leads on how gamma stimulation may translate into beneficial effects on the cellular and molecular level," the authors write.

As the authors make clear that more definitive clinical studies are needed, they note that at the moment, there are now 15 new clinical studies of gamma stimulation underway. Among these is a phase 3 clinical trial by the company

Cognito Therapeutics, which has licensed MIT's technology. That study plans to enroll hundreds of participants. Meanwhile, some recent or new clinical and preclinical studies have begun looking at whether gamma stimulation may be applicable to neurological disorders other than Alzheimer's, including stroke or Down syndrome. In experiments with mouse models, for example, an MIT team has been testing gamma stimulation's potential to help with cognitive effects of chemotherapy, or "chemobrain."

"Larger clinical studies are required to ascertain the long-term benefits of gamma stimulation," the authors conclude. "In animal models the focus should be on delineating the mechanism of gamma stimulation and providing further proof of principle studies on what other applications gamma stimulation may have."

In addition to Tsai and Blanco-Duque, the paper's other authors are Diane Chan, Martin Kahn, and Mitch Murdock.

JUNE 12, 2025

A mild spinal zap can cut brain-computer interface learning time in half

by University of Texas at Austin

edited by Sadie Harley, reviewed by Robert Egan

From left to right: Former students Satyam Kumar and Hussein Alawieh and professor José del R. Millán operate a robotic arm using a brain-computer interface. Credit: The University of Texas at Austin

Through a device called a brain-computer interface (BCI) it's possible to control a robotic arm or a wheelchair with thoughts alone. But for many users, learning to operate these systems is slow, difficult and, in some cases, unattainable.

Researchers at The University of Texas at Austin discovered a novel way to accelerate this learning process: a gentle electrical nudge to the spine before BCI training.

In a study published in the *Proceedings of the National Academy of Sciences*, researchers from the Cockrell School of Engineering and Dell Medical School found that noninvasive spinal stimulation can help the user focus on the task at hand, significantly speeding up the learning curve for brain-computer interfaces. This stimulation was shown to cut learning time in half.

"By using spinal stimulation to prime the brain, we're not just speeding up learning—we're also making it possible for people who previously struggled to use BCIs to gain control," said José del R. Millán, professor in the Cockrell School's Chandra Family Department of Electrical and Computer Engineering and the Department of Neurology at Dell Med.

"This opens up exciting possibilities for motor rehabilitation and assistive technology."

BCIs detect brain signals associated with movement intentions and translate them into commands for external devices. These systems rely on specific brain activity patterns called sensorimotor rhythms, which are generated when a person imagines moving a limb to control a device.

The researchers used a technique called transcutaneous electrical spinal stimulation, which involves delivering mild electrical pulses to the spinal cord through electrodes placed on the skin. The stimulation temporarily inhibits certain areas of the brain, allowing the neural activity associated with motor imagery to become more focused and stable.

This "preconditioning" effect helps users produce stronger and more consistent brain signals, making it easier for the BCI system to interpret their intentions.

"Think of it like tuning a radio to the right frequency," said Hussein Alawieh, a former graduate student in Millán's lab who was the first author of the study. "Spinal stimulation helps the brain filter out noise and focus on the signals that matter most for controlling the BCI. This makes the learning process faster and more effective."

The researchers conducted experiments involving 20 healthy participants and two individuals with spinal cord injuries. Participants were divided into two groups: one received spinal stimulation before each training session, while the other rested for the same amount of time. Here's what happened:

- Faster learning: Participants who received spinal stimulation showed significant improvements in BCI performance after just two training sessions, compared to five sessions for the control group.
- Improved accuracy: By the end of the training, the stimulation group achieved higher accuracy in controlling the BCI, with stronger and more focused brain activity patterns.
- Long-lasting effects: The benefits of the spinal stimulation persisted for at least a week after training, suggesting that the technique helps users retain their skills over time.

The researchers also tested their technique on individuals who had previously failed to learn BCI control using traditional methods. After undergoing the stimulation protocol, all participants in this "slow learner" group successfully gained control of the system, a sign that this technique could open up BCI to more potential users.

BCIs are already used to help individuals with paralysis regain some level of independence. In addition, this technique could be used as part of rehabilitation programs for stroke survivors and others with motor impairments.

BCIs have been shown to promote brain plasticity—the ability of the brain to reorganize itself and form new connections, which is critical for recovery. Faster and more reliable BCI control could enhance these therapeutic effects, potentially leading to better outcomes.

While this study focused on hand movements, the researchers believe their approach could be extended to more complex tasks, such as controlling robotic limbs with multiple degrees of freedom. They also plan to explore using their spinal stimulation technique in other populations, including individuals with severe neurological conditions.

"Our ultimate goal is to improve quality of life for people with motor impairments," said Millán. "Whether it's helping someone regain the ability to move their arm or enabling them to operate a wheelchair with their thoughts, this technology has the potential to make a real difference."

More information: Hussein Alawieh et al, Electrical spinal cord stimulation promotes focal sensorimotor activation that accelerates brain–computer interface skill learning, *Proceedings of the National Academy of Sciences* (2025). DOI: 10.1073/pnas.2418920122
Journal information: Proceedings of the National Academy of Sciences
Provided by University of Texas at Austin

Star-shaped brain cells may underpin the brain's massive memory storage

June 9, 2025

A new machine learning model shows that star-shaped brain cells may be responsible for the brain's memory capacity, and someday, it could inspire advances in AI and Alzheimer's research.

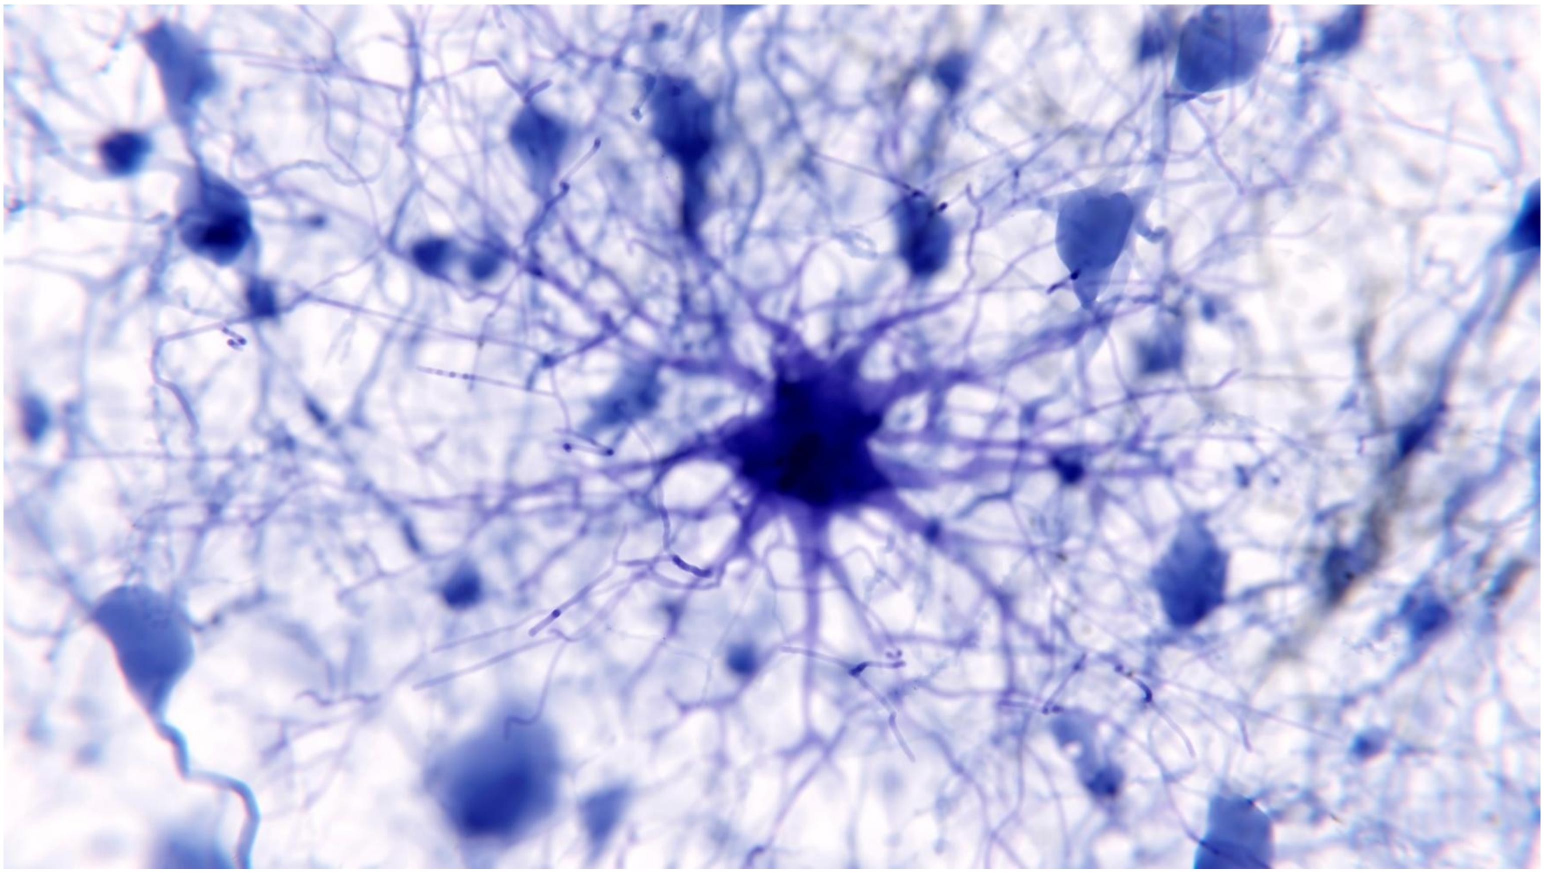

Astrocytes are star-shaped cells in the brain that may play and unsung role in memory. (Image credit: JUAN GAERTNER/SCIENCE PHOTO LIBRARY via Getty Images)

For decades, scientists believed neurons were the brain's sole architects of thought and memory — but now, new research suggests that another, often-overlooked type of brain cell may play a more central role in memory than previously thought.

The study, published in May in the journal PNAS, proposes that these other brain cells, called astrocytes, could be responsible for the brain's impressive memory-storage capacity through a newly discovered kind of network architecture.
Astrocytes are star-shaped cells that perform many maintenance tasks in the brain, including clearing cellular debris, supplying neurons with nutrients and regulating blood flow. They also sport thin branching structures, known as processes, that wrap around the

points where neurons exchange messages. This wrapping forms what is called a tripartite synapse, a kind of three-way handshake involving the two connected neurons and the astrocyte.

"You can imagine an astrocyte as an octopus with millions of tentacles," said lead author Leo Kozachkov, who was a PhD student at MIT at the time the study was conducted and is now a postdoctoral fellow at IBM Research in Yorktown Heights, New York. "The head of the octopus is the cell body, and the tentacles are 'processes' that wrap around nearby synapses," Kozachkov told Live Science in an email.
Astrocytes don't transmit electrical impulses like neurons do. Instead, they communicate via calcium signaling, sending waves of charged calcium particles within and between cells. Studies have shown that astrocytes respond to synaptic activity by altering their internal calcium levels. These changes can then trigger the release of chemical messengers from the astrocyte into the synapse.

"These processes act as tiny calcium computers, sensing when information is sent through the synapse, passing that information to other processes, and then receiving feedback in return," Kozachkov said. Ultimately, this chain email gets back to the neurons, which adjust their activity in turn. However, researchers don't yet fully understand the precise computational functions astrocytes perform with the information they receive from neurons.

To better understand this function, Kozachkov and his colleagues turned to machine learning architectures that are capable of representing complex interactions between many actors, rather than capturing only simple connections between pairs of units.

Traditional machine learning networks that link only pairs of neurons might encode limited information, said senior study author Dmitry Krotov, a research staff member at the MIT-IBM Watson AI Lab and IBM Research. Because a single astrocyte could connect to thousands of synapses, the team hypothesized that astrocytes might mediate communication across all of these connections. That could explain how the brain achieves its massive storage capabilities, they proposed.
"The unique anatomical structure of astrocytes provides a very natural and tempting way to design these large information storage systems in biological hardware," Kozachkov told Live Science in an email.

The researchers also hypothesized that astrocytes store memories through gradual changes in their internal calcium patterns and that these patterns are then translated back into signals that get sent to neurons in the form of chemical messengers. In this model, each astrocyte process, rather than the whole cell, functions as a distinct computational unit, the team proposed.

"Our model does not need a lot of neurons to store a lot of memories," Kozachkov said. "This is a significant advantage from an energy efficiency perspective, since neurons are metabolically 'expensive.'"

The model offers a "biologically grounded explanation" for how these memory storage systems might operate in the brain, said Maurizio de Pittà, an assistant professor at the Krembil Research Institute in Toronto, Canada, who was not involved in the work.
Past studies with high-resolution microscopes have supported this view, showing that astrocyte processes are interwoven throughout the brain and make contact with multiple synapses.
However, de Pittà told Live Science in an email that "models are powerful tools, but they remain approximations of the real world." He also cautioned that current technologies can not yet fully capture the dynamics unfolding in the human brain in real time, and that level of detail would be needed to validate the hypothesis.

Although scientists are starting to realize that astrocytes play a role in how we form memories, de Pittà said, we still don't have clear proof that the specific, calcium-based interactions between these cells and brain actually help create, store or recall memories, as suggested by the MIT team. If the team's model proves correct, though, the implications could offer a new way to think about brain storage, suggesting that memory capacity could scale with the number of astrocyte-synapse interactions present in the brain.

The model also offers potential therapeutic targets for neurodegenerative diseases, the study authors said.

"Astrocytes are known to be implicated in Alzheimer's and other memory disorders: our model provides a computational view of what might be going wrong," Kozachkov said. "Potentially, our mathematical model may inspire the search for new therapeutic targets: precise modulation of astrocyte process connectivity or signaling could restore or compensate for lost memory function."

However, much more research would be needed for this work to be translated into clinical treatments.

Beyond neuroscience, the model may point to applications in artificial intelligence. The model could help researchers create brain-like hardware systems, de Pittà said. Such systems could use dense memory architectures that enable them to store huge amounts of information and recall it efficiently, using very little energy, just like our brains do. This could be used for a wide array of applications, such as voice recognition; robotics and autonomous systems; AI assistants; or brain-machine interfaces and "neuroprosthetics."

Light and Sound at 40Hz Show Promise for Boosting Memory and Brain Growth in Down Syndrome

Muhammad Tuhin April 28, 2025

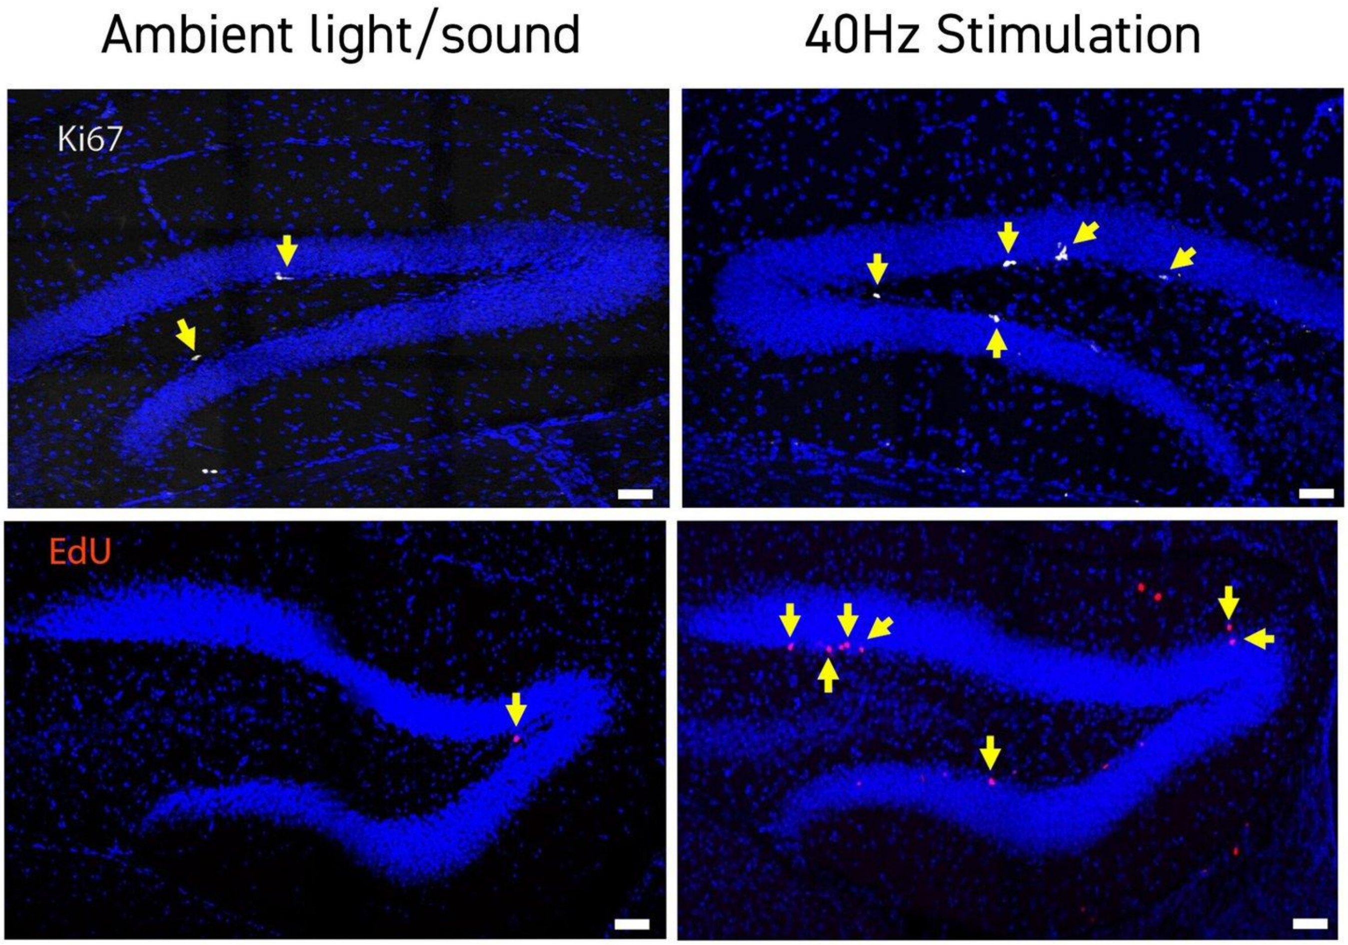

Images from the research paper show an increase in neurogenesis (as indicated by two markers: Ki67 and EdU) in mice exposed to 40Hz stimulation compared to those exposed only to ambient light and sound. Yellow arrows highlight instances of the markers. Credit: Alana Down Syndrome Center at MIT

Imagine a simple flicker of light or a gentle hum of sound—not invasive surgeries or potent pharmaceuticals—triggering the brain's ability to heal itself. This is not science fiction; it's the frontier of neuroscience. Recent studies are uncovering remarkable ways sensory stimulation at a frequency of 40 hertz, a rhythm matching the brain's natural gamma waves, can bolster cognitive function, foster neuron growth, and even combat neurodegenerative disease.

At the center of these discoveries is a technique called **GENUS**—**Gamma Entrainment Using Sensory Stimulation**. This method, using non-invasive light, sound, or touch at the 40Hz frequency, has already shown promise against Alzheimer's disease. Now, scientists at MIT's **Picower Institute for Learning and Memory** and the **Alana Down Syndrome Center** have extended its potential to another condition: **Down syndrome**.

Their findings, recently published in *PLOS ONE*, suggest that bathing the brain in a rhythm of light and sound could help correct key neurological deficits—even encouraging the birth of brand-new neurons.

Cracking the Code of Down Syndrome

Down syndrome arises when an individual carries an extra copy of chromosome 21, leading to lifelong challenges in cognition and memory. While the condition is complex and deeply rooted in genetics, researchers have long searched for ways to mitigate its impact on brain development and function.

In the latest study, a team led by Dr. **Li-Huei Tsai**, a distinguished neuroscientist and director of both the Picower Institute and the Alana Center, focused on whether GENUS could provide therapeutic effects in a mouse model of Down syndrome. These mice, called **Ts65Dn**, mirror many but not all characteristics of the human condition.

For three weeks, the mice were exposed daily to light and sound pulses oscillating precisely at 40Hz for one hour a day. The results were not just encouraging—they were groundbreaking.

Mice receiving the stimulation performed significantly better on tests of **short-term memory**. They remembered spatial environments more accurately, distinguished new objects from familiar ones with greater ease, and showed pronounced changes within a critical brain structure: the **hippocampus**, known as the memory center.

A "reel" describing a new study by MIT researchers finding benefits in Down syndrome mouse models from 40Hz sensory stimulation. Credit: Alana Down Syndrome Center at MIT

But why were these changes happening? To answer that, the scientists ventured deep inside the brain's cellular and molecular machinery.

Lighting Up the Brain's Blueprint

The hippocampus didn't just behave differently; it **looked** different at a cellular level. Using sophisticated techniques like **single-cell RNA sequencing**, the researchers analyzed nearly 16,000 individual brain cells, mapping how their genetic expression shifted after stimulation.

They discovered that mice exposed to GENUS had **revamped the way neurons organized their connections**. Key genes involved in forming **synapses**—the intricate communication points between neurons—were upregulated. Moreover, the architecture of these connections became denser and healthier, particularly in a vital hippocampal subregion called the **dentate gyrus**.

Under the microscope, it became clear: the brains of stimulated mice had **more synapses**, **stronger connectivity**, and potentially, **better communication pathways** than those of unstimulated controls.

And the changes didn't stop there.

Birthing New Neurons

Among the most stunning discoveries was the stimulation's effect on **neurogenesis**—the creation of new neurons from neural stem cells. For decades, scientists believed that adult brains were largely static, incapable of producing new brain cells. Today, we know that the adult brain, especially regions like the hippocampus, can indeed spawn new neurons under the right conditions.

The MIT team found that GENUS amplified the expression of a critical gene regulator called **TCF4**, known to govern the birth of neurons. Mice treated with 40Hz stimulation exhibited **significantly more new neurons** in their dentate gyrus compared to unstimulated mice.

This marked the first time any study had linked GENUS directly with increased neurogenesis.

Dr. **Md Rezaul Islam**, a leading author of the paper, explained: “The increase in functional synapses we observed is likely related to the increase in adult neurogenesis induced by GENUS. It’s an exciting finding because it suggests that we are not just preserving brain cells—we’re actually helping the brain regenerate.”

A Shield Against Aging and Alzheimer’s

Down syndrome and Alzheimer’s disease share a troubling link: nearly 90% of individuals with Down syndrome develop Alzheimer’s-like symptoms by the time they reach middle age. The Ts65Dn mice also exhibit signs of Alzheimer’s pathology as they age.

Here again, 40Hz stimulation offered hope.

The researchers found that sensory-stimulated mice retained **higher levels of Reelin**, a crucial protein for maintaining brain plasticity and resilience against Alzheimer’s. Reelin-expressing neurons are particularly vulnerable to degeneration, but in the stimulated mice, these neurons thrived.

In addition, clusters of genes usually associated with healthy brain aging stayed active longer in stimulated mice. In unstimulated mice, these genes declined—as they do in natural aging and Alzheimer’s progression.

In other words, GENUS not only improved memory and neuron growth but also appeared to slow the aging process in critical brain circuits.

A Symphony of Repair

Taken together, the findings paint a vivid picture of how GENUS might work. By synchronizing brain waves at 40Hz, the stimulation seems to set off a **homeostatic repair response**, rebalancing genetic activity, promoting synaptic health, encouraging the birth of new neurons, and enhancing resilience to degeneration.

It’s as if the brain, when bathed in this rhythmic stimulation, remembers how to heal itself.

Dr. Tsai emphasized the broader implications: “We are increasingly seeing that GENUS doesn’t just address one type of brain pathology. Whether it’s Alzheimer’s, chemo brain, stroke, or Down syndrome, this form of sensory stimulation seems to awaken the brain's innate capacity for restoration."

Limitations and The Road Ahead

Despite the enthusiasm, Tsai and her team urge caution. The study was conducted in mice, and even though the Ts65Dn model captures many features of Down syndrome, it does not fully replicate the human condition. Furthermore, all test subjects were male mice, leaving open questions about gender-specific effects.

Also, the study primarily assessed short-term memory improvements. Longer-term cognitive outcomes and the effects on other brain regions—like the prefrontal cortex, critical for decision-making—remain unexplored.

Nevertheless, small human trials of GENUS are now underway at MIT. Early observations suggest the method is safe and well-tolerated, but whether it can provide measurable benefits for people with Down syndrome, Alzheimer's, or other conditions remains to be seen.

The Dream of Non-Invasive Neurological Therapy

The beauty of GENUS lies in its **simplicity** and **non-invasiveness**. No surgeries. No drugs. Just carefully calibrated sensory inputs tuned to a brain's natural rhythm.

If future human studies confirm these findings, GENUS could open a revolutionary new chapter in brain medicine. Imagine treatment rooms where patients undergo relaxing sessions of flickering lights and harmonious sounds—recharging their brain's healing mechanisms, building new neurons, and staving off cognitive decline.

In a field dominated by molecular therapies and surgical interventions, GENUS offers a stunningly elegant alternative: healing the mind through its own music.

Reference: Md Rezaul Islam et al, Multisensory gamma stimulation enhances adult neurogenesis and improves cognitive function in male mice with Down Syndrome, *PLOS One* (2025). DOI: 10.1371/journal.pone.0317428

Love this? Share it and help us spark curiosity about science!

MIT News | Massachusetts Institute of Technology

Study reveals ways in which 40Hz sensory stimulation may preserve brain's "white matter"

Gamma frequency light and sound stimulation preserves myelination in mouse models and reveals molecular mechanisms that may underlie the benefit.

David Orenstein | The Picower Institute for Learning and Memory

Publication Date: August 13, 2024

Early-stage trials in Alzheimer's disease patients and studies in mouse models of the disease have suggested positive impacts on pathology and symptoms from exposure to light and sound presented at the "gamma" band frequency of 40 hertz (Hz). A new study zeroes in on how 40Hz sensory stimulation helps to sustain an essential process in which the signal-sending branches of neurons, called axons, are wrapped in a fatty insulation called myelin. Often called the brain's "white matter," myelin protects axons and insures better electrical signal transmission in brain circuits.

"Previous publications from our lab have mainly focused on neuronal protection," says Li-Huei Tsai, Picower Professor in The Picower Institute for Learning and Memory and the Department of Brain and Cognitive Sciences at MIT and senior author of the new open-access study in *Nature Communications*. Tsai also leads MIT's Aging Brain Initiative. "But this study shows that it's not just the gray matter, but also the white matter that's protected by this method."

This year Cognito Therapeutics, the spinoff company that licensed MIT's sensory stimulation technology, published phase II human trial results in the *Journal of Alzheimer's Disease* indicating that 40Hz light and sound stimulation significantly slowed the loss of myelin in volunteers with Alzheimer's. Also this year, Tsai's lab published a study showing that gamma sensory stimulation helped mice withstand neurological effects of chemotherapy medicines, including by preserving myelin. In the new study, members of Tsai's lab led by former postdoc Daniela Rodrigues Amorim used a common mouse model of myelin loss — a diet with the chemical cuprizone — to explore how sensory stimulation preserves myelination.

Amorim and Tsai's team found that 40Hz light and sound not only preserved myelination in the brains of cuprizone-exposed mice, it also appeared to protect oligodendrocytes (the cells that myelinate neural axons), sustain the electrical performance of neurons, and preserve a key marker of axon structural integrity. When the team looked into the molecular underpinnings of these benefits, they found clear signs of specific mechanisms including preservation of neural circuit connections called synapses; a reduction in a cause of oligodendrocyte death called "ferroptosis;" reduced inflammation; and an increase in the ability of microglia brain cells to clean up myelin damage so that new myelin could be restored.

"Gamma stimulation promotes a healthy environment," says Amorim, who is now a Marie Curie Fellow at the University of Galway in Ireland. "There are several ways we are seeing different effects."

The findings suggest that gamma sensory stimulation may help not only Alzheimer's disease patients but also people battling other diseases involving myelin loss, such as multiple sclerosis, the authors wrote in the study.

Maintaining myelin

To conduct the study, Tsai and Amorim's team fed some male mice a diet with cuprizone and gave other male mice a normal diet for six weeks. Halfway into that period, when cuprizone is known to begin causing its most acute effects on myelination, they exposed some mice from each group to gamma sensory stimulation for the remaining three weeks. In this way they had four groups: completely unaffected mice, mice that received no cuprizone but did get gamma stimulation, mice that received cuprizone and constant (but not 40Hz) light and sound as a control, and mice that received cuprizone and also gamma stimulation.

After the six weeks elapsed, the scientists measured signs of myelination throughout the brains of the mice in each group. Mice that weren't fed cuprizone maintained healthy levels, as expected. Mice that were fed cuprizone and didn't receive 40Hz gamma sensory stimulation showed drastic levels of myelin loss. Cuprizone-fed mice that received 40Hz stimulation retained significantly more myelin, rivaling the health of mice never fed cuprizone by some, but not all, measures.

The researchers also looked at numbers of oligodendrocytes to see if they survived better with sensory stimulation. Several measures revealed that in mice fed cuprizone, oligodendrocytes in the corpus callosum region of the brain (a key point for the transit of neural signals because it connects the brain's hemispheres) were markedly reduced. But in mice fed cuprizone and also treated with gamma stimulation, the number of cells were much closer to healthy levels.

Electrophysiological tests among neural axons in the corpus callosum showed that gamma sensory stimulation was associated with improved electrical performance in cuprizone-fed mice who received gamma stimulation compared to cuprizone-fed mice left untreated by 40Hz stimulation. And when researchers looked in the anterior cingulate cortex region of the brain, they saw that MAP2, a protein that signals the structural integrity of axons, was much better preserved in mice that received cuprizone and gamma stimulation compared to cuprizone-fed mice who did not.

A key goal of the study was to identify possible ways in which 40Hz sensory stimulation may protect myelin.

To find out, the researchers conducted a sweeping assessment of protein expression in each mouse group and identified which proteins were differentially expressed based on cuprizone diet and exposure to gamma frequency stimulation. The analysis revealed distinct sets of effects between the cuprizone mice exposed to control stimulation and cuprizone-plus-gamma mice.

A highlight of one set of effects was the increase in MAP2 in gamma-treated cuprizone-fed mice. A highlight of another set was that cuprizone mice who received control stimulation showed a substantial deficit in expression of proteins associated with synapses. The gamma-treated cuprizone-fed mice did not show any significant loss, mirroring results in a 2019 Alzheimer's 40Hz study that showed synaptic preservation. This result is important, the researchers wrote, because neural circuit activity, which depends on maintaining synapses, is associated with preserving myelin. They confirmed the protein expression results by looking directly at brain tissues.

Another set of protein expression results hinted at another important mechanism: ferroptosis. This phenomenon, in which errant metabolism of iron leads to a lethal buildup of reactive oxygen species in cells, is a known problem for oligodendrocytes in the cuprizone mouse model. Among the signs was an increase in cuprizone-fed, control stimulation mice in expression of the protein HMGB1, which is a marker of ferroptosis-associated damage that triggers an inflammatory response. Gamma stimulation, however, reduced levels of HMGB1.

Looking more deeply at the cellular and molecular response to cuprizone demyelination and the effects of gamma stimulation, the team assessed gene expression using single-cell RNA sequencing technology. They found that astrocytes and microglia became very inflammatory in cuprizone-control mice but gamma stimulation calmed that response. Fewer cells became inflammatory and direct observations of tissue showed that microglia became more proficient at clearing away myelin debris, a key step in effecting repairs.

The team also learned more about how oligodendrocytes in cuprizone-fed mice exposed to 40Hz sensory stimulation managed to survive better. Expression of protective proteins such as HSP70 increased and as did expression of GPX4, a master regulator of processes that constrain ferroptosis.

In addition to Amorim and Tsai, the paper's other authors are Lorenzo Bozzelli, TaeHyun Kim, Liwang Liu, Oliver Gibson, Cheng-Yi Yang, Mitch Murdock, Fabiola Galiana-Meléndez, Brooke Schatz, Alexis Davison, Md Rezaul Islam, Dong Shin Park, Ravikiran M. Raju, Fatema Abdurrob, Alissa J. Nelson, Jian Min Ren, Vicky Yang and Matthew P. Stokes.

Fundacion Bancaria la Caixa, The JPB Foundation, The Picower Institute for Learning and Memory, the Carol and Gene Ludwig Family Foundation, Lester A. Gimpelson, Eduardo Eurnekian, The Dolby Family, Kathy and Miguel Octavio, the Marc Haas Foundation, Ben Lenail and Laurie Yoler, and the U.S. National Institutes of Health provided funding for the study.

Tuning the Brain: The Science of 40Hz Stimulation and Mental Fitness

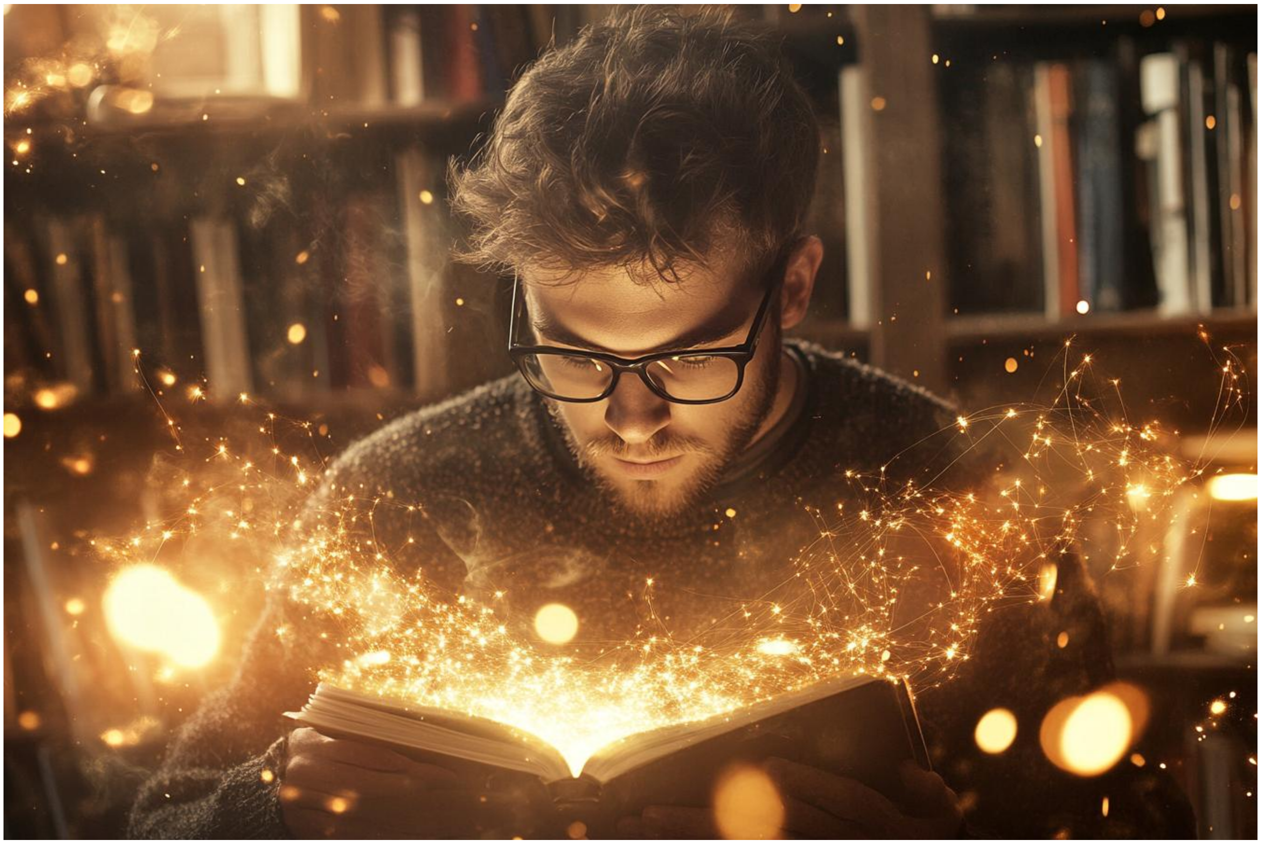

The Brain's Rhythm: Why 40Hz Matters

Over the past decade, scientists and cognitive performance researchers have been increasingly drawn to the 40Hz frequency—a specific gamma wave oscillation that appears to play a crucial role in brain function. Gamma waves, which range from 30 to 80Hz, are essential for processing information, memory formation, and neural communication. However, recent studies suggest that 40Hz stimulation, in particular, may enhance cognitive performance, improve focus, and even slow neurodegenerative processes.

Disruptions in neural coordination—whether due to chronic stress, sleep deprivation, aging, or neurological conditions—can weaken cognitive function, making it harder to focus, recall information, and sustain mental clarity. Researchers at MIT's Picower Institute have found that exposure to 40Hz sensory stimulation helps restore neural synchrony, potentially improving brain health and cognitive resilience. This discovery has sparked growing interest in how 40Hz can be leveraged for mental fitness and overall brain optimization.

While gamma waves span from 30 to 80Hz, the 40Hz frequency has attracted special attention. Studies show that sensory stimulation at this frequency can enhance cognitive function, improve focus, and potentially slow down the progression of Alzheimer's disease. This raises an important question: how can we leverage this frequency to support mental fitness in practical ways?

To answer that, let's explore the science behind 40Hz brain stimulation, its impact on neurotransmitters, and how sound-based neurostimulation can be harnessed to optimize cognitive performance.

The Neuroscience of 40Hz: How It Affects the Brain

Neurons communicate through electrical oscillations, forming brainwaves that correspond to different mental states. Delta waves (0.5–4Hz) dominate deep sleep, alpha waves (8–12Hz) emerge during relaxation, and beta waves (12–30Hz) accompany active thinking. Gamma waves (30–80Hz), particularly at 40Hz, are associated with:

- Memory formation – Essential for processing and retaining information.
- Attention and focus – Enhances cognitive performance in high-demand tasks.
- Sensory perception – Integrates multiple senses to create a unified experience.

- Neural synchronization – Facilitates efficient communication between different brain regions.

Recent breakthroughs from MIT's Picower Institute for Learning and Memory have revealed that 40Hz sensory stimulation (light and sound) enhances brain function by preserving white matter—the brain's communication superhighway. Motivated by the growing need for non-invasive interventions in neurodegenerative conditions, researchers aimed to investigate how 40Hz stimulation might counteract cognitive decline. Their study used a combination of light and auditory stimulation on mice, exposing them to flickering lights and synchronized sound pulses at 40Hz. This dual-sensory approach significantly reduced amyloid plaques, a hallmark of Alzheimer's disease, while also strengthening synaptic connections. This research is particularly noteworthy because it suggests that external rhythmic stimulation can restore neural function and potentially mitigate the effects of neurodegenerative diseases—a breakthrough with implications for both clinical and consumer neurotechnology applications.

How 40Hz Influences Neurotransmitters

Gamma oscillations at 40Hz influence critical neurotransmitters, including:

- Dopamine – Enhances motivation, reward processing, and learning.
- Acetylcholine – Improves memory retention and attention.
- Glutamate – Boosts cognitive processing speed.
- Serotonin – Supports mood regulation and emotional stability.

A study published in Nature demonstrated that 40Hz auditory stimulation increases dopamine release in brain regions responsible for motivation and goal-directed behavior. The researchers were motivated by previous findings linking gamma oscillations to cognitive function and wanted to determine whether externally applied 40Hz sound waves could actively modulate brain chemistry. In their study, they exposed mice to 40Hz auditory tones for prolonged periods and observed significant increases in dopamine activity in the prefrontal cortex and striatum, regions associated with motivation, learning, and executive function. This suggests that 40Hz stimulation may reinforce cognitive states linked to sustained attention and problem-solving. This means that listening to sound waves at this frequency may enhance focus and mental clarity—an exciting prospect for high-performers seeking a cognitive edge.

Building upon these findings in animal models, researchers sought to determine whether similar effects could be observed in humans. A notable study published in The Journal of Neuroscience investigated how 40Hz auditory stimulation influences neural activity in the human auditory cortex. Researchers combined magnetoencephalography (MEG) recordings with pharmacological interventions in healthy participants to assess the role of specific neurotransmitter systems. They discovered that enhancing GABAergic (gamma-aminobutyric acid) transmission increased the amplitude of the 40Hz auditory steady-state response (ASSR), suggesting that 40Hz stimulation engages inhibitory neural circuits in the human brain.

This finding is particularly relevant because inhibitory circuits, especially those involving GABAergic neurons, help regulate and balance brain activity. When these circuits function optimally, they prevent excessive neural excitability, which is crucial for maintaining focus, cognitive stability, and preventing mental fatigue. Strengthening these circuits through 40Hz stimulation may enhance information processing efficiency and cognitive performance, particularly in tasks requiring sustained attention and working memory.

Another study explored the potential therapeutic applications of 40Hz sensory stimulation in humans. Researchers found that 40Hz visual and auditory stimulation increased functional neuronal connectivity and improved disease biomarkers in individuals with early Alzheimer's disease. This suggests that 40Hz stimulation may have beneficial effects on cognitive function and neurodegenerative conditions.

These human studies provide promising evidence that 40Hz stimulation can modulate brain activity and potentially enhance cognitive functions. While more research is needed to fully understand the mechanisms and long-term effects, these findings support the potential of 40Hz auditory stimulation as a non-invasive tool for mental fitness and cognitive enhancement.

Practical Applications: Using 40Hz to Improve Mental Fitness

So, how can you apply 40Hz stimulation to enhance focus, creativity, and overall brain function? Here are three science-backed ways to integrate it into your routine:

1. Enhancing Focus and Productivity

If you struggle with distractions, 40Hz auditory stimulation may help sharpen concentration. Research suggests that listening to 40Hz binaural beats or isochronic tones before and during cognitive tasks can improve sustained attention and reaction time.

Try This: Before tackling deep work, play a 40Hz binaural beat track through high-fidelity headphones, such as enophones, which are designed to deliver optimized sound stimulation. Studies suggest that as little as five minutes of exposure can prime your brain for focus.

2. Boosting Memory and Learning

Memory formation relies on neural synchrony, which is strengthened by gamma wave activity. Research has shown that students exposed to 40Hz sound while studying retain more information and recall details faster than those without stimulation. A recent study conducted at a leading cognitive neuroscience institute investigated this phenomenon by exposing participants to controlled 40Hz auditory stimulation while engaging in memory-based tasks. The results indicated a statistically significant improvement in working memory, recall speed, and information retention compared to control groups. This suggests that 40Hz stimulation may enhance the brain's ability to encode and retrieve information, making it a promising tool for learners, students, and professionals looking to optimize cognitive performance.

Try This: When learning new material, experiment with 40Hz sound therapy to see if it improves your retention.

3. Supporting Relaxation and Sleep

Interestingly, 40Hz stimulation doesn't just enhance focus—it may also improve sleep quality. Researchers at Stanford University, led by experts in sleep and cognitive neuroscience, found that 40Hz brainwave entrainment before bed led to deeper, more restorative sleep cycles by reducing overactivity in the prefrontal cortex. Their research was motivated by growing evidence that disrupted gamma oscillations contribute to sleep disturbances and cognitive decline.

Using EEG recordings, they observed that participants exposed to 40Hz auditory stimulation before sleep experienced increased slow-wave activity, which is critical for memory consolidation and emotional regulation. These findings suggest that targeted 40Hz stimulation could be an effective, non-invasive strategy for enhancing sleep quality and cognitive resilience.

Try This: If you experience restless nights, try low-volume 40Hz sound exposure before bed to encourage relaxation and deeper sleep.

Where the Science Is Unclear

While 40Hz stimulation has shown promise, some studies yield mixed results. Here are a few areas where research is ongoing:

- Individual Variability – Some individuals respond strongly to 40Hz stimulation, while others show minimal effects.
- Long-Term Impact – More human trials are needed to determine if chronic 40Hz exposure has sustained benefits.
- Optimal Delivery Methods – Auditory, visual, and tactile stimulation may have different effects on the brain.

Despite these uncertainties, researchers remain optimistic about the therapeutic potential of 40Hz gamma waves.

Experimenting with 40Hz: How enophones Can Help

If you're interested in experimenting with 40Hz auditory neurostimulation, enophones provide a seamless way to do so. The eno app features a growing selection of soundscapes that incorporate 40Hz neurostimulation on its own and in combination with other frequencies, designed to support different target mental states such as focus, relaxation, and creativity.

The stimulation is applied through precisely tuned auditory signals, delivered in real-time based on the user's brain activity, ensuring a personalized and optimized neurostimulation experience.

How It Works: When enophones detect that your focus is slipping, they can subtly adjust the audio environment to reintroduce 40Hz stimulation, reinforcing cognitive engagement.

Since every brain is unique, the best way to determine if 40Hz stimulation works for you is to experiment with real-time EEG feedback. Try listening to gamma-enhancing soundscapes during work, study, or relaxation, and track how your mind responds.

Curious to explore more? Check out eno's brain-responsive sound technology and see how it aligns with your mental fitness goals.

Electrons spiral with a purpose: A new platform decodes their selective spin

Researchers have engineered an artificial, controllable system that can mimic the conditions under which chiral-induced spin selectivity effect occurs, possibly transforming how we study quantum transport.

Jun 14, 2025

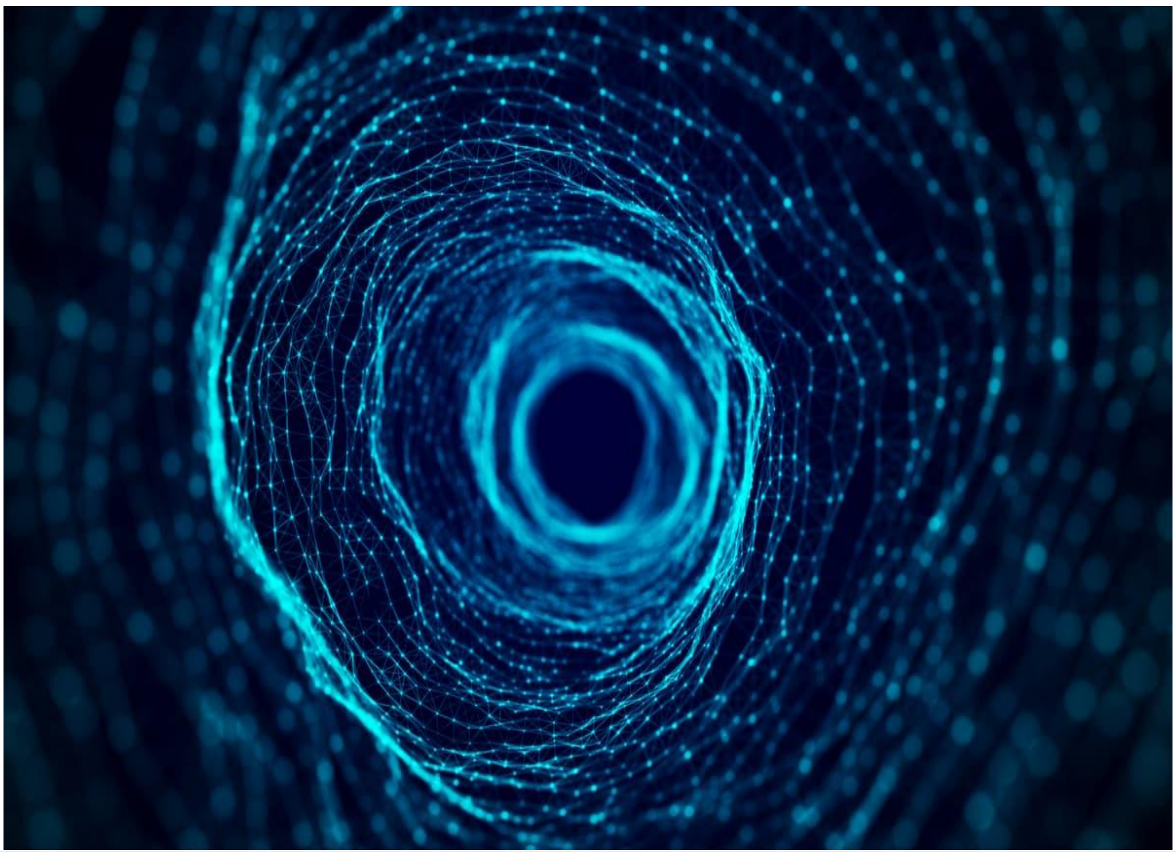

Concept image of a quantum spiral.

By combining the principles of physics, chemistry, and biology, scientists have crafted a special programmable platform to explore one of the most puzzling quantum mysteries of our time: why electrons seem to choose sides when passing through certain twisted molecules.

This behavior, known as the chiral-induced spin selectivity (CISS) effect, has baffled researchers for over two decades. It shows up in biological processes like photosynthesis and cellular respiration, yet no one fully understands how or why it happens.

Now, researchers from the University of Pittsburgh have engineered an artificial, controllable system that can mimic the conditions under which this strange effect occurs. Their approach could reshape how we study quantum transport and might also help us design new materials for electronics, energy, and even medicine.

"The beauty of our approach is not that it mimics chemistry or biology exactly, but that it allows us to isolate and study individual processes that are relevant in chiral quantum transport," said François Damanet, a physicist and one of the members of the research team.

Drawing quantum spirals to unlock electron secrets

Back in the late 1990s, scientists Ron Naaman and David Waldeck made a surprising discovery. When electrons pass through films of chiral (twisted) molecules, how easily they can move is decided by their spin, which is a quantum property. Instead of a small noticeable effect, they saw spin-dependent changes as high as 20 percent, a result that stunned the scientific community.

Since then, the CISS effect has popped up in various biological systems, yet researchers haven't been able to pin down the exact mechanism behind it. This is because real biological molecules are complex. They're soft, flexible, constantly moving, and surrounded by water, all of which makes it nearly impossible to isolate the role of chirality alone.

That's where the new platform comes in. The researchers did not try to recreate biology. Instead, they built a clean, programmable playground for electrons. Using a technique developed in 2008, they worked with a special material made from layers of lanthanum aluminate ($LaAlO_3$) and strontium titanate ($SrTiO_3$).

By using a fine-tipped microscopic pen, they could draw paths where electrons can travel. To make those paths chiral, they introduced a clever twist: the probe not only moved in a wavy, serpentine pattern across the surface, but its voltage was also modulated up and down in sync. This combination created spiral-like channels that broke mirror symmetry, the key ingredient of chirality.

These artificial chiral waveguides weren't just pretty shapes. When electrons flowed through them, surprising quantum effects emerged. The team saw unusual conductance patterns and even observed electrons pairing up in ways that shouldn't be possible under strong magnetic fields.

Theoretical models suggested that the spiral geometry created a kind of engineered spin-orbit coupling, which locked the electrons' spin to their direction of motion, just like some theories had proposed for the CISS effect in molecules.

What makes this platform so powerful is that it's fully programmable. Researchers can change the shape, size, and strength of the chiral patterns, erase them, and write new ones, all on the same device. "We can systematically vary parameters like the pitch, amplitude, and coupling strength of chiral modulations—something impossible with fixed structures," Damanet said.

The programmable platform is unlike anything else

This new platform doesn't try to copy molecules atom-for-atom. Instead, it gives scientists something they've never had before: precise control. In biological systems, everything is messy—molecules wiggle, environments shift, and vibrations interfere with measurements.

However, on this programmable platform, each variable can be changed independently, allowing researchers to test exactly how chirality affects quantum transport. This could help settle long-standing debates about whether spin-orbit interactions, molecular vibrations, or other mechanisms drive the CISS effect.

While the system operates at ultra-cold temperatures and uses inorganic materials, it sets the stage for future hybrid setups that could combine these solid-state tools with real molecules.

The team is already exploring ways to pair their platform with organic materials or carbon nanotubes, and even to run experiments at higher temperatures. The goal isn't to replace biological studies, but to work alongside them, much like how wind tunnels help engineers test aircraft designs before real-world flights.

If successful, this approach could help scientists not only solve the CISS puzzle but also understand other complex quantum systems. It could inspire new materials for spintronics, where electron spin is used in computing, or guide the design of efficient catalysts and bio-inspired energy devices.

The study is published in the journal *Science Advances*.

What is Moore’s Law and does this decades-old computing prophecy still hold true?

Features By Tim Danton published May 16, 2025

Moore's Law was an off-hand prediction that came to be one of the prevailing laws of modern computing — but what did it predict, and can we still rely on it?

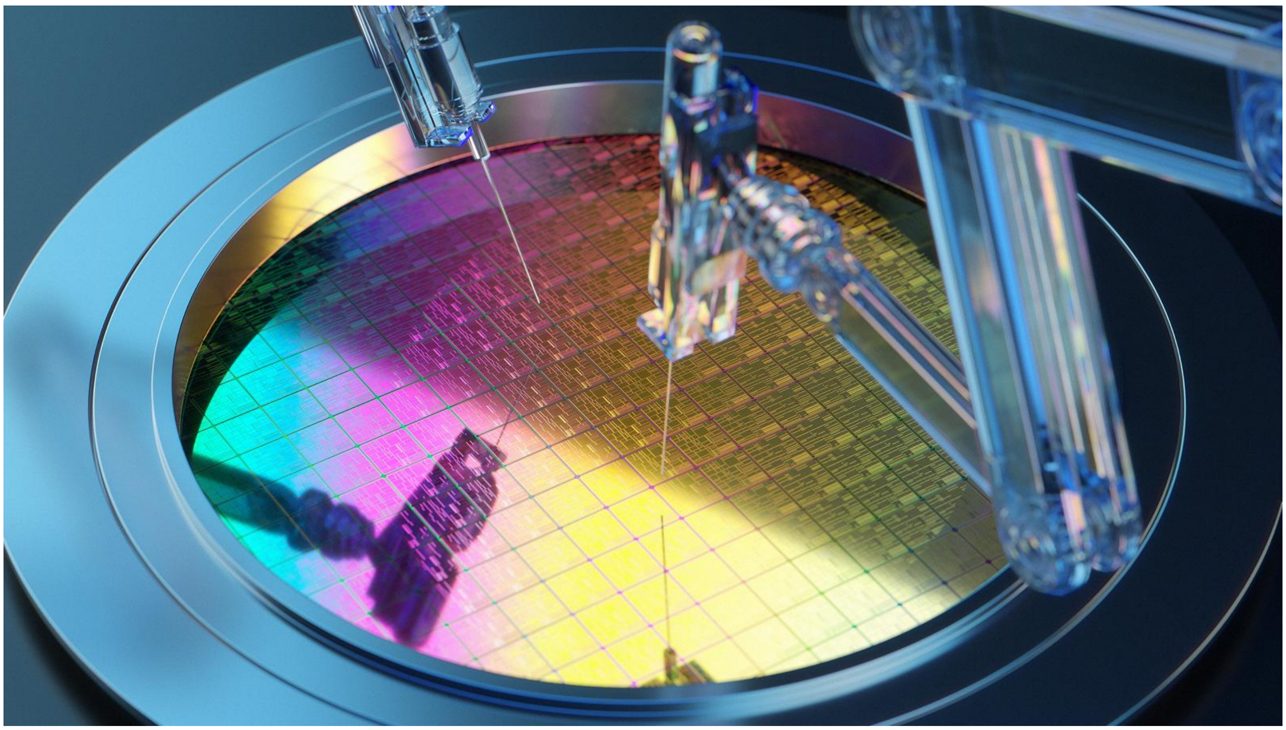

(Image credit: Yuichiro Chino/Getty Images)

Gordon Moore wasn’t overly enthused when he was asked to write an article to celebrate the 35th anniversary of Electronic Magazine in 1965. "I was given the chore of predicting what would happen in silicon components in the next ten years," he recalled 40 years later. But as Director of R&D at Fairfield Semiconductor, which had developed the breakthrough planar transistor in 1959, he was perfectly placed to assess the progress that had been made in six short years.

In particular, Moore noticed that Fairfield had doubled the number of transistors that could be placed on a chip each year — being able to squeeze 60 where there had once been two. He then "blindly extrapolated for about ten years and said, okay, in 1975 we’ll have about 60 thousand components on a chip." In other words, every year the number had doubled and Moore thought it would continue to double. His prediction was neat and easy to understand — but most of all, it worked.

The idea was quickly dubbed Moore's Law, and it mostly held true until 1975. (To be strictly accurate, the number doubled nine times over ten years rather than ten times over ten years). Seeing complications further down the line, Moore revised his prediction to a doubling every two years, and remarkably, his prediction once again proved to be (roughly) accurate for the next 40 years.

Moore's self-fulfilling prophecy

One reason for the success of Moore's prediction is that it became a guide — almost a target — for chip designers. This was especially the case for Intel, the company that Gordon Moore co-founded with Robert Noyce in 1968. Moore and Noyce, one of the engineers behind the planar process, saw a potential in integrated circuits that the recession-hit and cautious Fairfield did not.

In 1971, Intel would have its first big hit: the 4004 microprocessor. It included 2,300 transistors measuring 10 microns thick — five times slimmer than a strand of human hair. A little over ten years later, Intel introduced the 80286 processor, with 134,000 transistors each measuring 1.5 microns (because of this, it's referred to as a "1.5-micron process"). These developments emerged very much in line with the revised Moore's Law.

When looking back over the years that followed — the 1980s, 1990s and early 2000s — it may seem like the path of progress was smooth. Moore's Law kept holding, after all. But that was only possible due to a series of major breakthroughs, each solving a problem that at one time seemed impossible.

Some were based on material science, such as improving the "doping" methods that insert impurities into a semiconductor to better control its conductivity. Or the creation of complementary metal oxide superconductor (CMOS) technology in the mid-1980s, which brought lower power consumption and thus less heat. Other breakthroughs came in the manufacturing process, such as the development of extreme ultraviolet lithography (EUV) to etch patterns onto ever smaller wafers.

"I was given the chore of predicting what would happen in silicon components in the next ten years," Moore said 40 years after he made his first prediction. (Image credit: © Roger Ressmeyer/CORBIS/VCG via Getty Images)

And the innovations didn't stop. We mentioned the planar processor, where transistors sit on a level plane, right at the start of this article. It took years of research and development — (four Japanese researchers at Delta created the first vertical design for a processor in 1989) — but when it arrived, the vertical FinFET processor gave Moore's Law fresh life in 2012 in the form of Intel's third-generation Core i3, i5 and i7 processors. These used a 22nm processor and packed up to 1.4 billion processors.

These are just a handful of the innovations that Gordon Moore could never have foreseen yet enabled his law to hold true. But there was one seemingly impossible problem looming on the horizon — physics.

Why smaller isn't always better

A strand of hair is around 50 microns thick. A mote of dust around five microns. A bacterial cell, such as Mycoplasma, measures 0.5 microns. Now, consider that modern transistors are often 0.005 microns thick, or 5 nanometers (5nm) and you'll realise we're approaching atomic levels. We mean that literally: the space between the centre of two adjacent silicon atoms is around 0.235 nanometers, so you can squeeze around 21 into a 5nm space.

Then consider that the latest CPU manufacturing processes have reduced yet further, from 5nm to 2nm, meaning space for eight silicon atoms. At this point, we begin to reach the point where quantum mechanical effects prevail, such as quantum tunneling, which causes electrons leak. This is not a property you want in a transistor.

All of this means that the straightforward approach of "make things smaller" no longer works on its own. That's why we have seen a shift away from miniaturization and instead towards more sophisticated processors, with every chip in every device you own now including many different cores so that tasks can be split between them.

Is Moore's Law still relevant?

In the simplest sense, no. The days when we could double the number of transistors on a chip every two years are far behind us. However, Moore's Law has acted as a pacesetter in a decades-long race to create chips that perform more complicated tasks quicker, especially as our expectations for continual progress continue.

To measure its success, consider that if Moore's Law had suggested a doubling every 10 years instead of every two, then we would be stuck with 1980s-era computers. Steve Jobs would never have been able to announce the iPhone in 2007, establishing the smartphone era.

This pace-setting is something that's now demanded not only by consumers, but the boards of technology companies. It's one of the drivers behind the neural processing units (NPUs) inside recent processors, capable of running local artificial intelligence (AI) tasks that are beyond the reach of conventional CPUs. For now, this technology can perform simple tasks such as removing unwanted people from the background of our photos, but this is just the beginning.

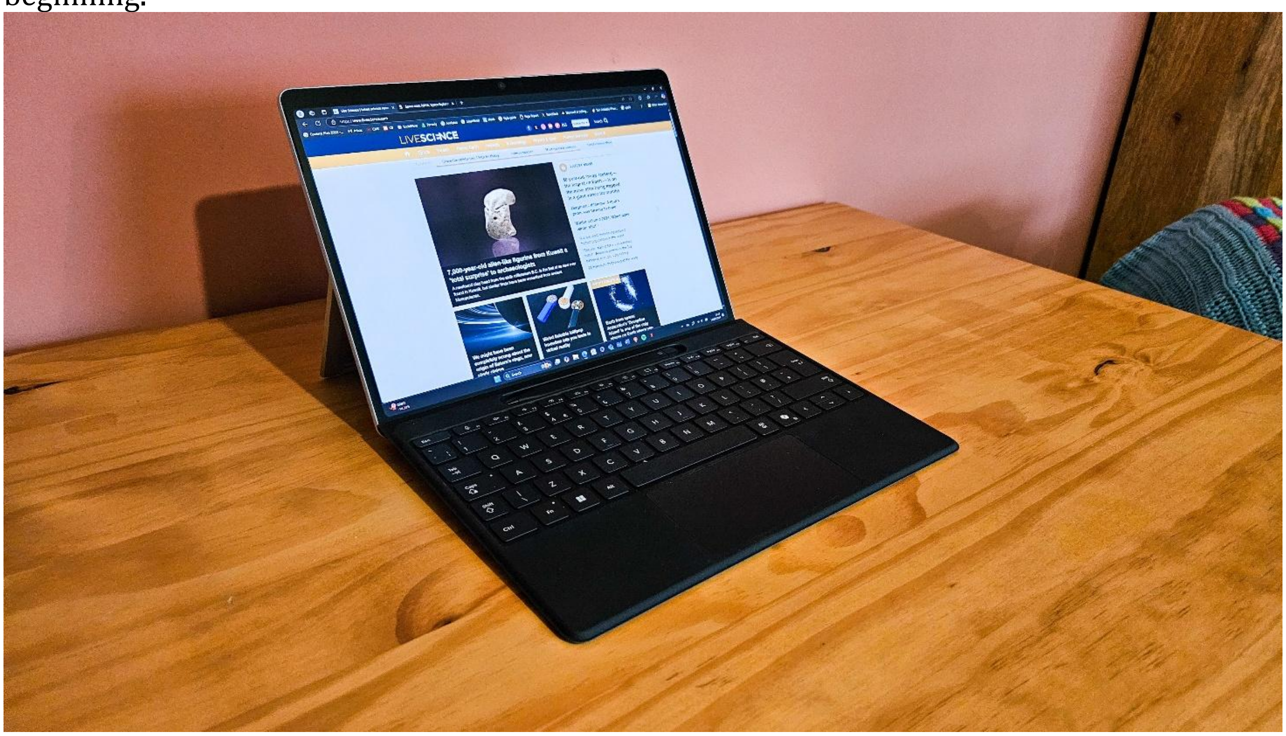

The newest laptops, including the Microsoft Surface Pro 11 (pictured) are fitted with NPUs, which allow for specialized AI workloads. (Image credit: Keumars Afifi-Sabet/Future)

NPUs are big news today, as are the incredible Nvidia-powered equivalents in data centers that drive ChatGPT, Midjourney and the other AI services we are gradually coming to rely on. Meanwhile, it seems that personal AI assistants are a heartbeat away, with even bigger leaps likely to come in the next decade.

We can't yet be sure what those leaps will entail. What we can say is that developments are currently happening in university research labs and R&D divisions in megacorporations such as Intel. One of those labs might yet work out a way to cram yet more transistors into even smaller areas — or perhaps move away from transistors altogether — but that seems unlikely.

Instead, Moore's Law lives on as an expectation of the pace of progress. An expectation that every tech company from DeepSeek to Meta to OpenAI will continue to use as their guide.

Scientists clear major roadblocks in mission to build powerful AI photonic chips

News By Matthew Spink, Demosthenes Koutsogeorgis published May 11, 2025

Two studies show major progress in the field of photonic microchips.

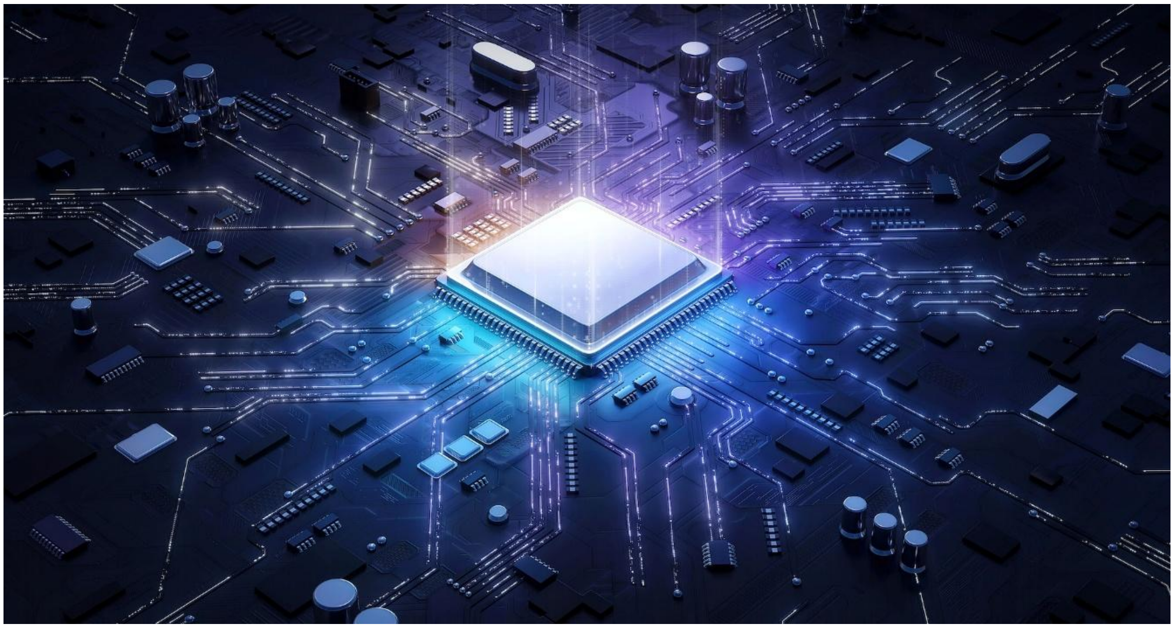

Two studies show major progress in the field of photonic microchips. (Image credit: 3dartists via Shutterstock)

Electronic microchips are at the heart of the modern world. They're found in our laptops, our smartphones, our cars and our household appliances. For years, manufacturers have been making them more powerful and efficient, which increases the performance of our electronic devices.

But that trend is now faltering because of the increased cost and complexity of manufacturing chips, as well as performance limits set by the laws of physics. This is happening just as there's a need for increased computing power because of the boom in artificial intelligence (AI).

An alternative to the electronic microchips we currently use are photonic chips. These use light instead of electricity to achieve higher performance. However, photonic chips have not yet taken off due to a number of hurdles. Now, two papers published in Nature address some of these roadblocks, offering essential stepping stones to achieving the computing power required by complex artificial intelligence systems.
By using light (photons) instead of electricity (electrons) for the transport and processing of information, photonic computing promises higher speeds and greater bandwidths with greater efficiency. This is because it does not suffer from the loss of electrical current due to a phenomenon known as resistance, as well as unwanted heat loss from electrical components.
Photonic computing is also particularly suited for performing what are known as matrix multiplications — mathematical operations that are fundamental to AI.
Those are some of the benefits. The challenges, however, are not trivial. In the past, the performance of photonic chips has generally been studied in isolation. But because of the dominance of electronics in modern technology, photonic hardware will need to be

However, converting photons into electrical signals can slow down processing times since light operates at higher speeds. Photonic computing is also based around analogue operations rather than digital ones. This can reduce precision and limit the type of computing tasks that can be carried out.

It's also difficult to scale them up from small prototypes because large-scale photonic circuits cannot currently be fabricated with sufficient accuracy. Photonic computing will require its own software and algorithms, compounding the challenges of integration and compatibility with other technology.

The two new papers in Nature address many of these hurdles. Bo Peng, from Singapore-based company Lightelligence, and colleagues demonstrate a new type of processor for photonic computing called a Photonic Arithmetic Computing Engine (Pace). This processor has a low latency, which means that there is a minimal delay between an input or command and the corresponding response or action by the computer.

The large-scale Pace processor, which has more than 16,000 photonic components, can solve difficult computing tasks, demonstrating the feasibility of the system for real world applications. The processor shows how integration of photonic and electronic hardware, accuracy, and the need for different software and algorithms can be resolved. It also demonstrates that the technology can be scaled up.

This marks a significant development, despite some speed limitations of the current hardware.

In a separate paper, Nicholas Harris, from California-based company Lightmatter, and colleagues describe a photonic processor that was able to run two AI systems with accuracy similar to those of conventional electronic processors. The authors demonstrated the effectiveness of their photonic processor through generating Shakespeare-like text, accurately classifying movie reviews and playing classic Atari computer games such as Pac-Man.

The platform is also potentially scalable, though in this case limitations of the materials and engineering used curtailed one measure of the processor's speed and its overall computational capabilities.

Both teams suggest that their photonic systems can be part of scalable next generation hardware that can support the use of AI. This would finally make photonics viable, though further refinements will be needed. These will involve the use of more effective materials or designs.

World's first light-powered neural processing units (NPUs) could massively reduce energy consumption in AI data centers

News By Owen Hughes published April 3, 2025

Q.ANT's new chip uses photon power in a bid to solve AI's big energy issue. It's also 50 times faster than silicon-based equivalents, the company says.

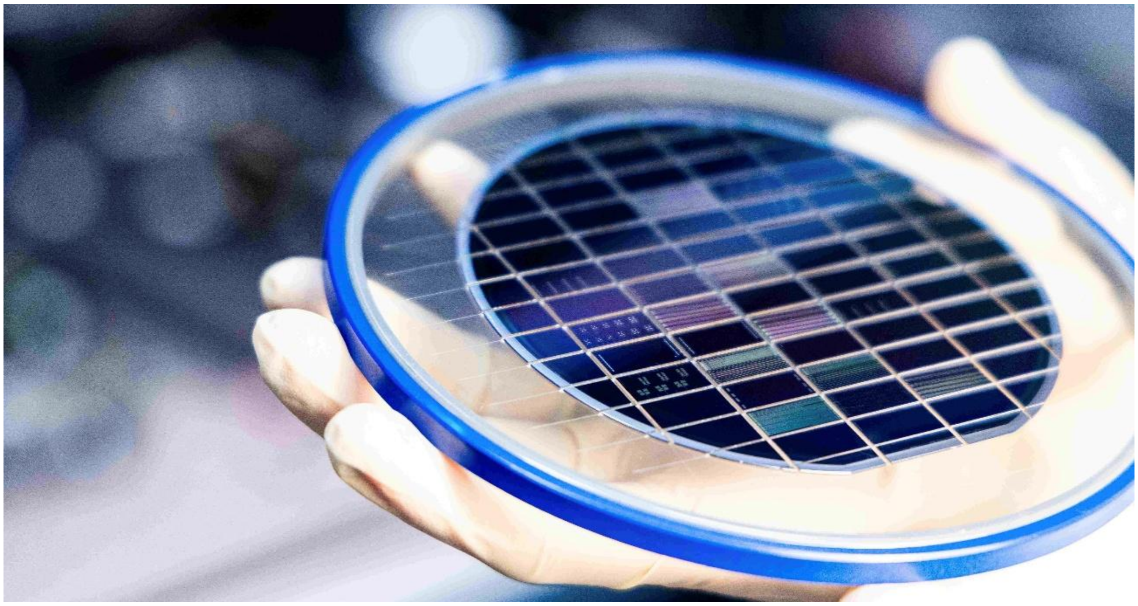

The Q.ANT wafer based on Thin Film Lithium Niobate enables photonic integrated circuits with high-precision, high-speed optical modulation, low noise and reduced thermal dissipation. (Image credit: © Q.ANT)

A light-powered computer chip designed to drive artificial intelligence (AI) data centers and make high-performance computing (HPC) more sustainable has entered production.

In a statement published Feb. 24, representatives from analog photonic chip company Q.ANT said its photonic AI chip could deliver a 30-fold increase in energy efficiency and a 50-fold boost in computing speed compared with conventional, silicon-based computer chips.

Pilot production of the new chip is now underway at IMS Chips in Stuttgart, Germany, where Q.ANT has invested 14 million euros ($15.1 million) to repurpose an existing semiconductor factory to fabricate its new, light-powered chip.

Because the chip is being produced on a repurposed facility instead of a specialist production line, the company believes it can bring the technology to market much more quickly. The chip can also integrate with the existing HPC servers, potentially accelerating adoption, Q.ANT representatives said.

"By 2030, we aim to make our photonic processors a scalable, energy-efficient cornerstone of AI infrastructure," Michael Förtsch, chief executive of Q.ANT, said in the statement.

Photonic computing

Photonic chips could solve a massive challenge faced by existing processor technology, particularly as AI and other data- and resource-intensive computing applications grow.

Traditional silicon chips control electrical signals using tiny switches called transistors. Photonic chips, by contrast, process data using light particles (photons), which are massless and can travel much faster than electrons do in conventional computer chips. Photons don't emit heat in the same way electrons carrying an electrical charge do. As such, using photonic chips in applications involving complex, energy-intensive computations like AI could overcome the limitations of classic silicon chip architecture and thus vastly accelerate the computers' processing speed and reduce their energy consumption.

"This comes at a critical time for the computing industry, as the exponential growth of AI and data-intensive applications will soon overwhelm the current data center infrastructure," Jens Anders, a professor at the University of Stuttgart and director and chief executive of IMS Chips, said in the statement. Anders added that the two companies aimed to establish "a scalable model for energy-efficient computing."

Q.ANT's chip is built using thin-film lithium niobate (TFLN), a crystalline compound applied to a wafer that forms the basis of the company's photonic chip. TFLN is increasingly catching the attention of photonics researchers and quantum scientists for its potential in next-generation computing. When an electric field is applied to the material, it can be used to control the speed and phase of light waves, thereby enabling it to modulate optical signals with extreme precision.

The pilot production line has been set up specifically to manufacture chips that incorporate TFLN, with Q.ANT aiming to fabricate 1,000 wafers per year.

"As AI and data-intensive applications push conventional semiconductor technology to its limits, we need to rethink the way we approach computing at the core," Förtsch said. "With this pilot line, we are accelerating time to market and laying the foundation for photonic processors to become standard coprocessors in high-performance computing."

AI-designed chips are so weird that 'humans cannot really understand them' — but they perform better than anything we've created

News By Tim Danton published February 20, 2025

AI models have, within hours, created more efficient wireless chips through deep learning, but it is unclear how their 'randomly shaped' designs were produced.

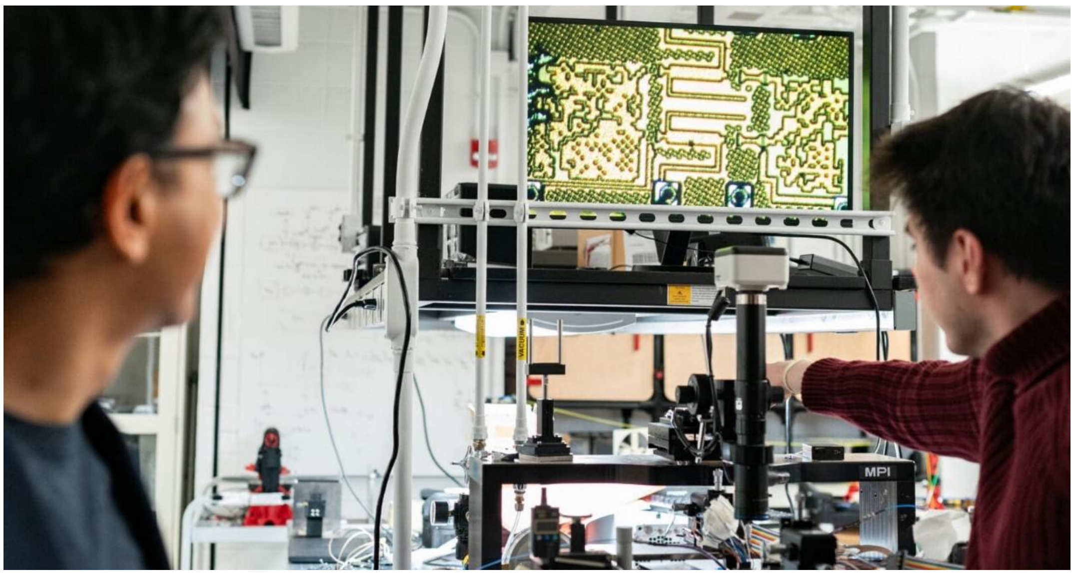

The AI treated the chip design as one complete system rather than a collection of parts. (Image credit: Princeton University)

Engineering researchers have demonstrated that artificial intelligence (AI) can design complex wireless chips in hours, a feat that would have taken humans weeks to complete. Not only did the chip designs prove more efficient, the AI took a radically different approach — one that a human circuit designer would have been highly unlikely to devise. The researchers outlined their findings in a study published Dec. 30 2024 in the journal Nature Communications.

The research focused on millimeter-wave (mm-Wave) wireless chips, which present some of the biggest challenges facing manufacturers due to their complexity and need for miniaturization. These chips are used in 5G modems, now commonly found in phones.

Manufacturers currently rely on a mix of human expertise, bespoke circuit designs and established templates. Each new design then goes through a slow process of optimization, based on trial and error because it is often so complex that a human cannot fully understand what is happening inside the chip. This leads to a cautious, iterative approach based on what has worked before.

In this case, however, researchers at Princeton Engineering and the Indian Institute of Technology posited that deep-learning-based AI models could use an inverse design method — one that specifies the desired output and leaves the algorithm to determine the inputs and parameters.

The AI also considers each chip as a single artifact, rather than a collection of existing elements that need to be combined. This means that established chip design templates, the ones that no one understands but probably hide inefficiencies, are cast aside.

The future of chip design?

In this experiment, the resulting structures "look randomly shaped," said lead author Kaushik Sengupta, a professor of electrical and computer engineering at Princeton. "Humans cannot really understand them."
And when Sengupta'steam manufactured the chips, they found the AI creations hit performance levels beyond those of existing designs.

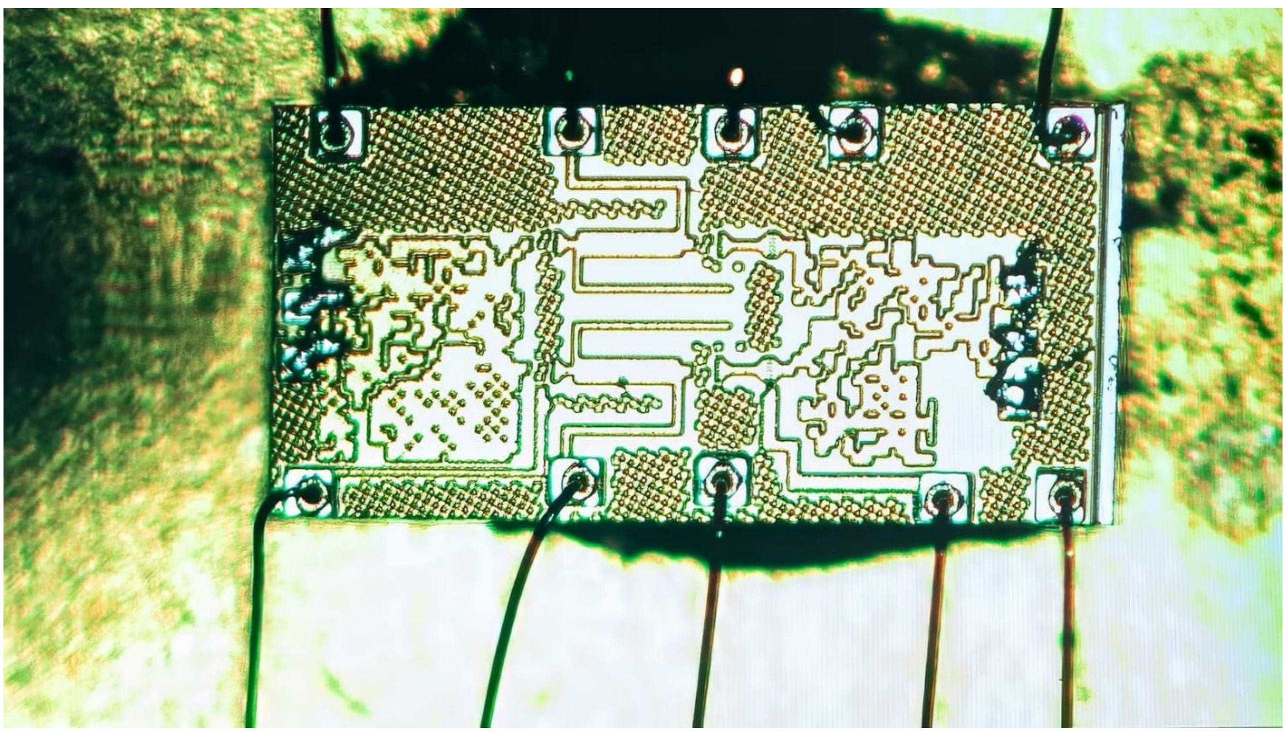

(Image credit: Princeton University)
Although the findings suggest that the design of such complex chips could be handed over to AI, Sengputa was keen to point out that pitfalls remain "that still require human designers to correct." In particular, many of the designs produced by the algorithm did not work– equivalent to the "hallucinations" produced by current generative AI tools.

"The point is not to replace human designers with tools," said Sengputa. "The point is to enhance productivity with new tools."

The speed with which iterative designs can be developed opens up new possibilities, too. Some chip designs can be geared towards energy efficiency, others to outright performance or to extending the frequency range.

Wireless chips are of growing importance, with an ever-growing demand for miniaturization, so this research is a valuable step forward. But Sengupta said that if his team's method can be extended to other parts of a circuit's design, it could change the way we design electronics in the future. "This is just the tip of the iceberg in terms of what the future holds for the field."

AI could crack unsolvable problems — and humans won't be able to understand the results

Opinion By Ehsan Nabavi published January 5, 2025

AI promises to accelerate scientific discovery, but if scientists aren't careful public trust may be left behind.

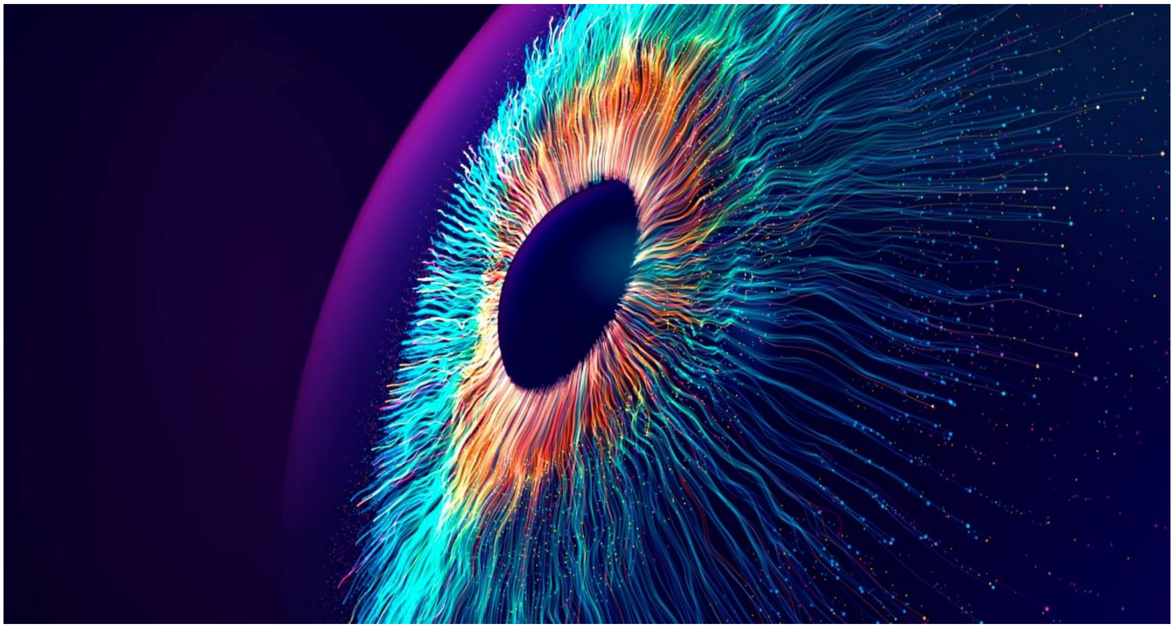

(Image credit: Andriy Onufriyenko via Getty Images)

Artificial intelligence (AI) has taken centre stage in basic science. The five winners of the 2024 Nobel Prizes in Chemistry and Physics shared a common thread: AI. Indeed, many scientists — including the Nobel committees — are celebrating AI as a force for transforming science.

As one of the laureates put it, AI's potential for accelerating scientific discovery makes it "one of the most transformative technologies in human history". But what will this transformation really mean for science?

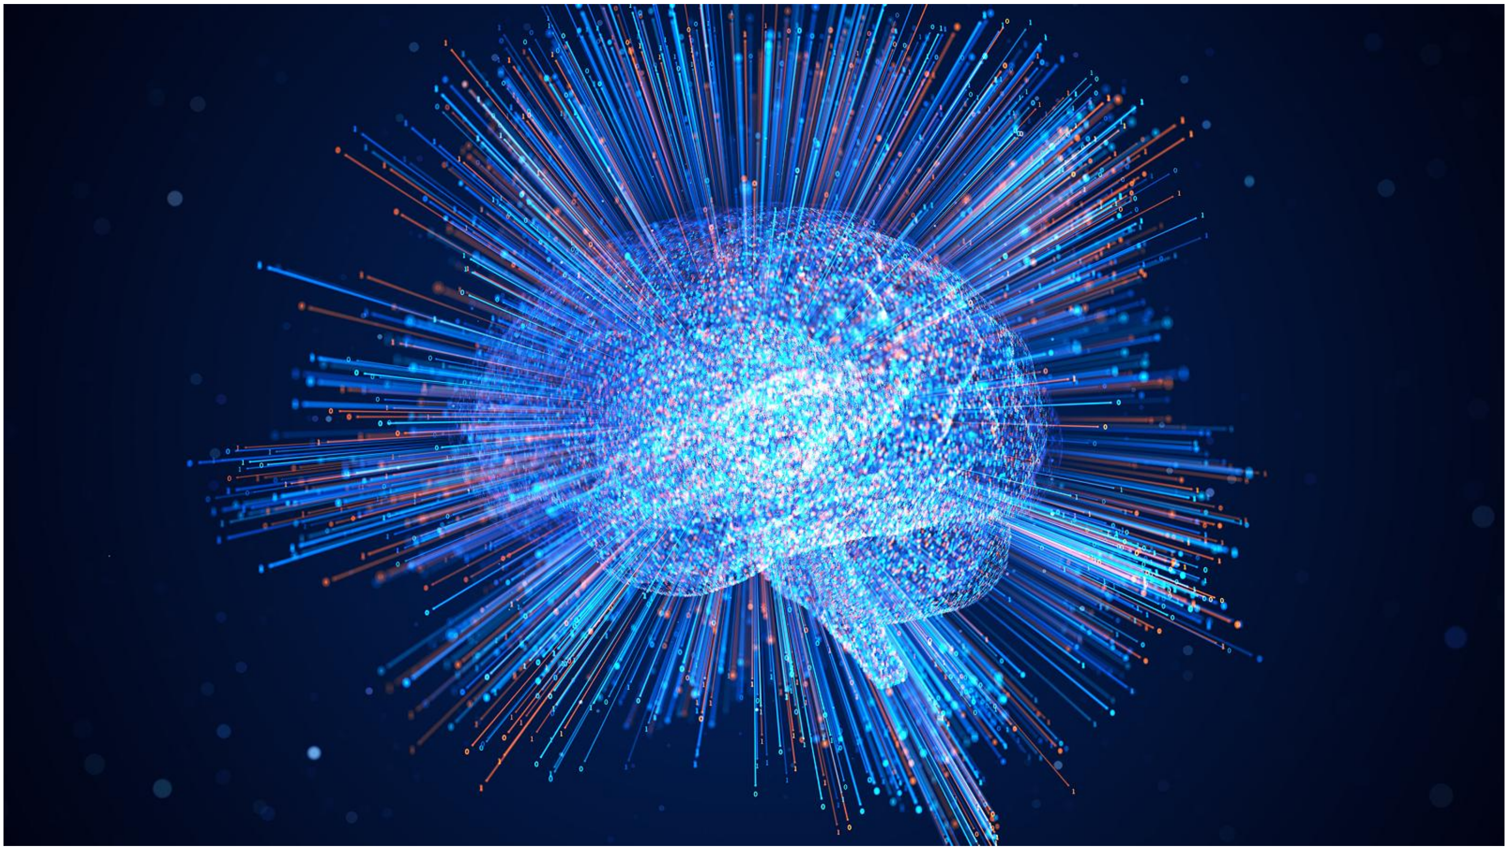

AI promises to help scientists do more, faster, with less money. But it brings a host of new concerns, too — and if scientists rush ahead with AI adoption they risk transforming science into something that escapes public understanding and trust, and fails to meet the needs of society.

The illusions of understanding

Experts have already identified at least three illusions that can ensnare researchers using AI.

The first is the "illusion of explanatory depth". Just because an AI model excels at predicting a phenomenon — like AlphaFold, which won the Nobel Prize in Chemistry for its predictions of protein structures — that doesn't mean it can accurately explain it. Research in neuroscience has already shown that AI models designed for optimised prediction can lead to misleading conclusions about the underlying neurobiological mechanisms.

Second is the "illusion of exploratory breadth". Scientists might think they are investigating all testable hypotheses in their exploratory research, when in fact they are only looking at a limited set of hypotheses that can be tested using AI.

Finally, the "illusion of objectivity". Scientists may believe AI models are free from bias, or that they can account for all possible human biases. In reality, however, all AI models inevitably reflect the biases present in their training data and the intentions of their developers.

Cheaper and faster science

One of the main reasons for AI's increasing appeal in science is its potential to produce more results, faster, and at a much lower cost.

An extreme example of this push is the "AI Scientist" machine recently developed by Sakana AI Labs. The company's vision is to develop a "fully AI-driven system for automated scientific discovery", where each idea can be turned into a full research paper for just US$15 — though critics said the system produced "endless scientific slop".

Do we really want a future where research papers can be produced with just a few clicks, simply to "accelerate" the production of science? This risks inundating the scientific ecosystem with papers with no meaning and value, further straining an already overburdened peer-review system.

We might find ourselves in a world where science, as we once knew it, is buried under the noise of AI-generated content.

A lack of context

The rise of AI in science comes at a time when public trust in science and scientists is still fairly high , but we can't take it for granted. Trust is complex and fragile.

As we learned during the COVID pandemic, calls to "trust the science" can fall short because scientific evidence and computational models are often contested, incomplete, or open to various interpretations.

However, the world faces any number of problems, such as climate change, biodiversity loss, and social inequality, that require public policies crafted with expert judgement. This judgement must also be sensitive to specific situations, gathering input from various disciplines and lived experiences that must be interpreted through the lens of local culture and values.

As an International Science Council report published last year argued, science must recognise nuance and context to rebuild public trust. Letting AI shape the future of science may undermine hard-won progress in this area.
If we allow AI to take the lead in scientific inquiry, we risk creating a monoculture of knowledge that prioritizes the kinds of questions, methods, perspectives and experts best suited for AI.
This can move us away from the transdisciplinary approach essential for responsible AI, as well as the nuanced public reasoning and dialogue needed to tackle our social and environmental challenges.

A new social contract for science

As the 21st century began, some argued scientists had a renewed social contract in which scientists focus their talents on the most pressing issues of our time in exchange for public funding. The goal is to help society move toward a more sustainable biosphere — one that is ecologically sound, economically viable and socially just.
The rise of AI presents scientists with an opportunity not just to fulfil their responsibilities but to revitalize the contract itself. However, scientific communities will need to address some important questions about the use of AI first.
For example, is using AI in science a kind of "outsourcing" that could compromise the integrity of publicly funded work? How should this be handled?

What about the growing environmental footprint of AI? And how can researchers remain aligned with society's expectations while integrating AI into the research pipeline?
The idea of transforming science with AI without first establishing this social contract risks putting the cart before the horse.

Letting AI shape our research priorities without input from diverse voices and disciplines can lead to a mismatch with what society actually needs and result in poorly allocated resources.

Science should benefit society as a whole. Scientists need to engage in real conversations about the future of AI within their community of practice and with research stakeholders. These discussions should address the dimensions of this renewed social contract, reflecting shared goals and values.

It's time to actively explore the various futures that AI for science enables or blocks — and establish the necessary standards and guidelines to harness its potential responsibly.

AI 'hallucinations' can lead to catastrophic mistakes, but a new approach makes automated decisions more reliable

News By Nicholas Fearn published August 30, 2024

Researchers have developed a new method to improve the accuracy and transparency of automated anomaly detection systems deployed in critical infrastructure.

(Image credit: SEAN GLADWELL/Getty Images)

Scientists have developed a new, multi-stage method to ensure artificial intelligence (AI) systems that are designed to identify anomalies make fewer mistakes and produce explainable and easy-to-understand recommendations.
Recent advances have made AI a valuable tool to help human operators detect and address issues affecting critical infrastructure such as power stations, gas pipelines and dams. But despite showing plenty of potential, models may generate inaccurate or vague results — known as "hallucinations."

Hallucinations are common in large language models (LLMs) like ChatGPT and Google Gemini. They stem from low-quality or biased training data and user prompts that lack additional context, according to Google Cloud.

Some algorithms also exclude humans from the decision-making process — the user enters a prompt, and the AI does the rest, without explaining how it made a prediction. When applying this technology to a serious area like critical infrastructure, a major concern is whether AI's lack of accountability and trust could result in human operators making the wrong decisions.

Some anomaly detection systems have previously been constrained by so-called "black box" AI algorithms, for example. These are characterized by opaque decision-making processes that generate recommendations difficult for humans to understand. This makes it hard for plant operators to determine, for example, the algorithm's rationale for identifying an anomaly.

A multi-stage approach

To increase AI's reliability and minimize problems such as hallucinations, researchers have proposed four measures, outlining their proposals in a paper published July 1 at the CPSS '24 conference. In the study, they focused on AI used for critical national infrastructure (CNI), such as water treatment.
First, the scientists deploy two anomaly detection systems, known as Empirical Cumulative Distribution-based Outlier Detection (ECOD) and Deep Support Vector Data Description (DeepSVDD), to identify a range of attack scenarios in datasets taken from the Secure Water Treatment (SWaT). This system is used for water treatment system research and training.

The researchers said both systems had short training times, provided fast anomaly detection and were efficient — enabling them to detect myriad attack scenarios. But, as noted by Rajvardhan Oak, an applied scientist at Microsoft and computer science researcher at UC Davis, ECOD had a "slightly higher recall and F1 score" than DeepSVDD. He explained that F1 scores account for the precision of anomaly data and the number of anomalies identified, allowing users to determine the "optimal operating point."
Secondly, the researchers combined these anomaly detectors with eXplainable AI (XAI) — tools that help humans better understand and assess the results generated by AI systems — to make them more trustworthy and transparent.

They found that XAI models like Shapley Additive Explanations (SHAP), which allow users to understand the role different features of a machine learning model play in making predictions, can provide highly accurate insights into AI-based recommendations and improve human decision-making.

The third component revolved around human oversight and accountability. The researchers said humans can question AI algorithms' validity when provided with clear explanations of AI-based recommendations. They could also use these to make more informed decisions regarding CNI.

The final part of this method is a scoring system that measures the accuracy of AI explanations. These scores give human operators more confidence in the AI-based insights they are reading. Sarad Venugopalan, co-author of the study, said this scoring system — which is still in development — depends on the "AI/ML model, the setup of the application use-case, and the correctness of the values input to the scoring algorithm."

Improving AI transparency

Speaking to Live Science, Venugopalan went on to explain that this method aims to provide plant operators with the ability to check whether AI recommendations are correct or not.

"This is done via message notifications to the operator and includes the reasons why it was sent," he said. "It allows the operator to verify its correctness using the information provided by the AI, and resources available to them."

Encouraged by this research and how it presents a solution for the AI black box problem, Rajvardhan Oak said: "With explanations attached to AI model findings, it is easier for subject matter experts to understand the anomaly, and for senior leadership to confidently make critical decisions. For example, knowing exactly why certain web traffic is anomalous makes it easier to justify blocking or penalizing it."

Eerke Boiten, a cybersecurity professor at De Montfort University, also sees the benefits of using explainable AI systems for anomaly detection in CNI. He said it will ensure humans are always kept in the loop when making crucial decisions based on AI recommendations. "This research is not about reducing hallucinations, but about responsibly using other AI approaches that do not cause them," he added.

AI 'hallucinates' constantly, but there's a solution

Opinion By Artur Garcez published June 6, 2025

Neurosymbolic AI combines the learning of LLMs with teaching the machine formal rules that should make them more reliable and energy efficient.

The main problem with big tech's experiment with artificial intelligence (AI) is not that it could take over humanity. It's that large language models (LLMs) like Open AI's ChatGPT, Google's Gemini and Meta's Llama continue to get things wrong, and the problem is intractable.

Known as hallucinations, the most prominent example was perhaps the case of US law professor Jonathan Turley, who was falsely accused of sexual harassment by ChatGPT in 2023.

OpenAI's solution seems to have been to basically "disappear" Turley by programming ChatGPT to say it can't respond to questions about him, which is clearly not a fair or satisfactory solution. Trying to solve hallucinations after the event and case by case is clearly not the way to go.

The same can be said of LLMs amplifying stereotypes or giving western-centric answers. There's also a total lack of accountability in the face of this widespread misinformation, since it's difficult to ascertain how the LLM reached this conclusion in the first place.

We saw a fierce debate about these problems after the 2023 release of GPT-4, the most recent major paradigm in OpenAI's LLM development. Arguably the debate has cooled since then, though without justification.

The EU passed its AI Act in record time in 2024, for instance, in a bid to be world leader in overseeing this field. But the act relies heavily on AI companies to regulate themselves without really addressing the issues in question. It hasn't stopped tech companies from releasing LLMs worldwide to hundreds of millions of users and collecting their data without proper scrutiny.

Meanwhile, the latest tests indicate that even the most sophisticated LLMs remain unreliable. Despite this, the leading AI companies still resist taking responsibility for errors.

Unfortunately LLMs' tendencies to misinform and reproduce bias can't be solved with gradual improvements over time. And with the advent of agentic AI, where users will soon

be able to assign projects to an LLM such as, say, booking their holiday or optimising the payment of all their bills each month, the potential for trouble is set to multiply.
The emerging field of neurosymbolic AI could solve these issues, while also reducing the enormous amounts of data required for training LLMs. So what is neurosymbolic AI and how does it work?

The LLM problem

LLMs work using a technique called deep learning, where they are given vast amounts of text data and use advanced statistics to infer patterns that determine what the next word or phrase in any given response should be. Each model — along with all the patterns it has learned — is stored in arrays of powerful computers in large data centers known as neural networks.

LLMs can appear to reason using a process called chain-of-thought, where they generate multi-step responses that mimic how humans might logically arrive at a conclusion, based on patterns seen in the training data.

Undoubtedly, LLMs are a great engineering achievement. They are impressive at summarizing text and translating, and may improve the productivity of those diligent and knowledgeable enough to spot their mistakes. Nevertheless they have great potential to mislead because their conclusions are always based on probabilities — not understanding.

A popular workaround is called "human-in-the-loop": making sure that humans using AIs still make the final decisions. However, apportioning blame to humans does not solve the problem. They'll still often be misled by misinformation.

LLMs now need so much training data to advance that we're now having to feed them synthetic data, meaning data created by LLMs. This data can copy and amplify existing errors from its own source data, such that new models inherit the weaknesses of old ones. As a result, the cost of programming AIs to be more accurate after their training — known as "post-hoc model alignment" — is skyrocketing.
It also becomes increasingly difficult for programmers to see what's going wrong because the number of steps in the model's thought process become ever larger, making it harder and harder to correct for errors.

Neurosymbolic AI combines the predictive learning of neural networks with teaching the AI a series of formal rules that humans learn to be able to deliberate more reliably. These include logic rules, like "if a then b", such as "if it's raining then everything outside is normally wet"; mathematical rules, like "if a = b and b = c then a = c"; and the agreed upon meanings of things like words, diagrams and symbols. Some of these will be inputted directly into the AI system, while it will deduce others itself by analyzing its training data and doing "knowledge extraction".

This should create an AI that will never hallucinate and will learn faster and smarter by organising its knowledge into clear, reusable parts. For example if the AI has a rule about things being wet outside when it rains, there's no need for it to retain every example of the things that might be wet outside — the rule can be applied to any new object, even one it has never seen before.

During model development, neurosymbolic AI also integrates learning and formal reasoning using a process known as the "neurosymbolic cycle". This involves a partially trained AI extracting rules from its training data then instilling this consolidated knowledge back into the network before further training with data.

This is more energy efficient because the AI needn't store as much data, while the AI is more accountable because it's easier for a user to control how it reaches particular conclusions and improves over time. It's also fairer because it can be made to follow pre-existing rules, such as: "For any decision made by the AI, the outcome must not depend on a person's race or gender".

The third wave

The first wave of AI in the 1980s, known as symbolic AI, was actually based on teaching computers formal rules that they could then apply to new information. Deep learning followed as the second wave in the 2010s, and many see neurosymbolic AI as the third.
It's easiest to apply neurosymbolic principles to AI in niche areas, because the rules can be clearly defined. So it's no surprise that we've seen it first emerge in Google's AlphaFold, which predicts protein structures to help with drug discovery; and AlphaGeometry, which solves complex geometry problems.
For more broad-based AIs, China's DeepSeek uses a learning technique called "distillation" which is a step in the same direction. But to make neurosymbolic AI fully feasible for general models, there still needs to be more research to refine their ability to discern general rules and perform knowledge extraction.

It's unclear to what extent LLM makers are working on this already. They certainly sound like they're heading in the direction of trying to teach their models to think more cleverly, but they also seem wedded to the need to scale up with ever larger amounts of data.

The reality is that if AI is going to keep advancing, we will need systems that adapt to novelty from only a few examples, that check their understanding, that can multitask and reuse knowledge to improve data efficiency and that can reason reliably in sophisticated ways.

This way, well designed digital technology could potentially even offer an alternative to regulation, because the checks and balances would be built into the architecture and perhaps standardized across the industry. There's a long way to go, but at least there's a path ahead.

AI hallucinates more frequently as it gets more advanced — is there any way to stop it from happening, and should we even try?

OpenAI's most advanced reasoning model is smarter than ever — but it hallucinates more than previous models, too.

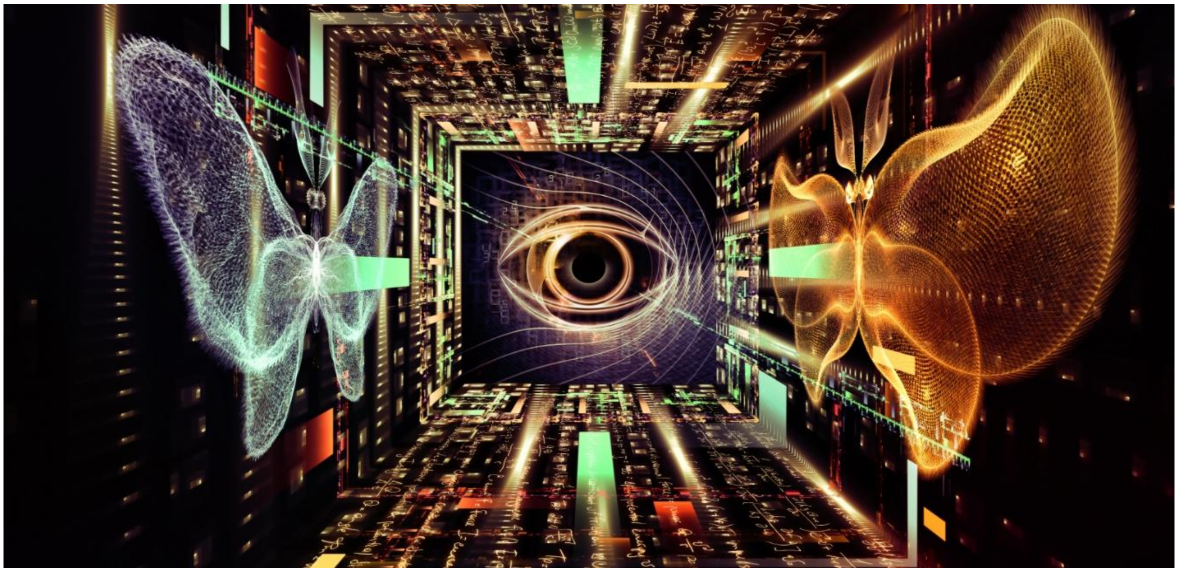

(Image credit: agsandrew/ Shutterstock)

The more advanced artificial intelligence (AI) gets, the more it "hallucinates" and provides incorrect and inaccurate information.
Research conducted by OpenAI found that its latest and most powerful reasoning models, o3 and o4-mini, hallucinated 33% and 48% of the time, respectively, when tested by OpenAI's PersonQA benchmark. That's more than double the rate of the older o1 model. While o3 delivers more accurate information than its predecessor, it appears to come at the cost of more inaccurate hallucinations.
This raises a concern over the accuracy and reliability of large language models (LLMs) such as AI chatbots, said Eleanor Watson, an Institute of Electrical and Electronics Engineers (IEEE) member and AI ethics engineer at Singularity University.
"When a system outputs fabricated information — such as invented facts, citations or events — with the same fluency and coherence it uses for accurate content, it risks misleading users in subtle and consequential ways," Watson told Live Science.

The issue of hallucination highlights the need to carefully assess and supervise the information AI systems produce when using LLMs and reasoning models, experts say.

Do AIs dream of electric sheep?

The crux of a reasoning model is that it can handle complex tasks by essentially breaking them down into individual components and coming up with solutions to tackle them. Rather than seeking to kick out answers based on statistical probability, reasoning models come up with strategies to solve a problem, much like how humans think.

In order to develop creative, and potentially novel, solutions to problems, AI needs to hallucinate —otherwise it's limited by rigid data its LLM ingests.

"It's important to note that hallucination is a feature, not a bug, of AI," Sohrob Kazerounian, an AI researcher at Vectra AI, told Live Science. "To paraphrase a colleague of mine, 'Everything an LLM outputs is a hallucination. It's just that some of those hallucinations are true.' If an AI only generated verbatim outputs that it had seen during training, all of AI would reduce to a massive search problem."
"You would only be able to generate computer code that had been written before, find proteins and molecules whose properties had already been studied and described, and answer homework questions that had already previously been asked before. You would not, however, be able to ask the LLM to write the lyrics for a concept album focused on the AI singularity, blending the lyrical stylings of Snoop Dogg and Bob Dylan."

In effect, LLMs and the AI systems they power need to hallucinate in order to create, rather than simply serve up existing information. It is similar, conceptually, to the way that humans dream or imagine scenarios when conjuring new ideas.

Thinking too much outside the box

However, AI hallucinations present a problem when it comes to delivering accurate and correct information, especially if users take the information at face value without any checks or oversight.
"This is especially problematic in domains where decisions depend on factual precision, like medicine, law or finance," Watson said. "While more advanced models may reduce the frequency of obvious factual mistakes, the issue persists in more subtle forms. Over time, confabulation erodes the perception of AI systems as trustworthy instruments and can produce material harms when unverified content is acted upon."

And this problem looks to be exacerbated as AI advances. "As model capabilities improve, errors often become less overt but more difficult to detect," Watson noted. "Fabricated content is increasingly embedded within plausible narratives and coherent reasoning chains. This introduces a particular risk: users may be unaware that errors are present and may treat outputs as definitive when they are not. The problem shifts from filtering out crude errors to identifying subtle distortions that may only reveal themselves under close scrutiny."

Kazerounian backed this viewpoint up. "Despite the general belief that the problem of AI hallucination can and will get better over time, it appears that the most recent generation of advanced reasoning models may have actually begun to hallucinate more than their simpler counterparts — and there are no agreed-upon explanations for why this is," he said.

The situation is further complicated because it can be very difficult to ascertain how LLMs come up with their answers; a parallel could be drawn here with how we still don't really know, comprehensively, how a human brain works.

In a recent essay, Dario Amodei, the CEO of AI company Anthropic, highlighted a lack of understanding in how AIs come up with answers and information. "When a generative AI system does something, like summarize a financial document, we have no idea, at a specific or precise level, why it makes the choices it does — why it chooses certain words over others, or why it occasionally makes a mistake despite usually being accurate," he wrote. The problems caused by AI hallucinating inaccurate information are already very real, Kazerounian noted. "There is no universal, verifiable, way to get an LLM to correctly answer questions being asked about some corpus of data it has access to," he said. "The examples of non-existent hallucinated references, customer-facing chatbots making up company policy, and so on, are now all too common."

Crushing dreams

Both Kazerounian and Watson told Live Science that, ultimately, AI hallucinations may be difficult to eliminate. But there could be ways to mitigate the issue.

Watson suggested that "retrieval-augmented generation," which grounds a model's outputs in curated external knowledge sources, could help ensure that AI-produced information is anchored by verifiable data.

"Another approach involves introducing structure into the model's reasoning. By prompting it to check its own outputs, compare different perspectives, or follow logical steps, scaffolded reasoning frameworks reduce the risk of unconstrained speculation and improve consistency," Watson, noting this could be aided by training to shape a model to prioritize accuracy, and reinforcement training from human or AI evaluators to encourage an LLM to deliver more disciplined, grounded responses.

"Finally, systems can be designed to recognise their own uncertainty. Rather than defaulting to confident answers, models can be taught to flag when they're unsure or to defer to human judgement when appropriate," Watson added. "While these strategies don't eliminate the risk of confabulation entirely, they offer a practical path forward to make AI outputs more reliable."

Given that AI hallucination may be nearly impossible to eliminate, especially in advanced models, Kazerounian concluded that ultimately the information that LLMs produce will need to be treated with the "same skepticism we reserve for human counterparts."

Scientists discover major differences in how humans and AI 'think' — and the implications could be significant

News By Drew Turney published April 1, 2025

Study finds that AI fundamentally lacks the human capability to make creative mental connections, raising warning signs for how we deploy AI tools.

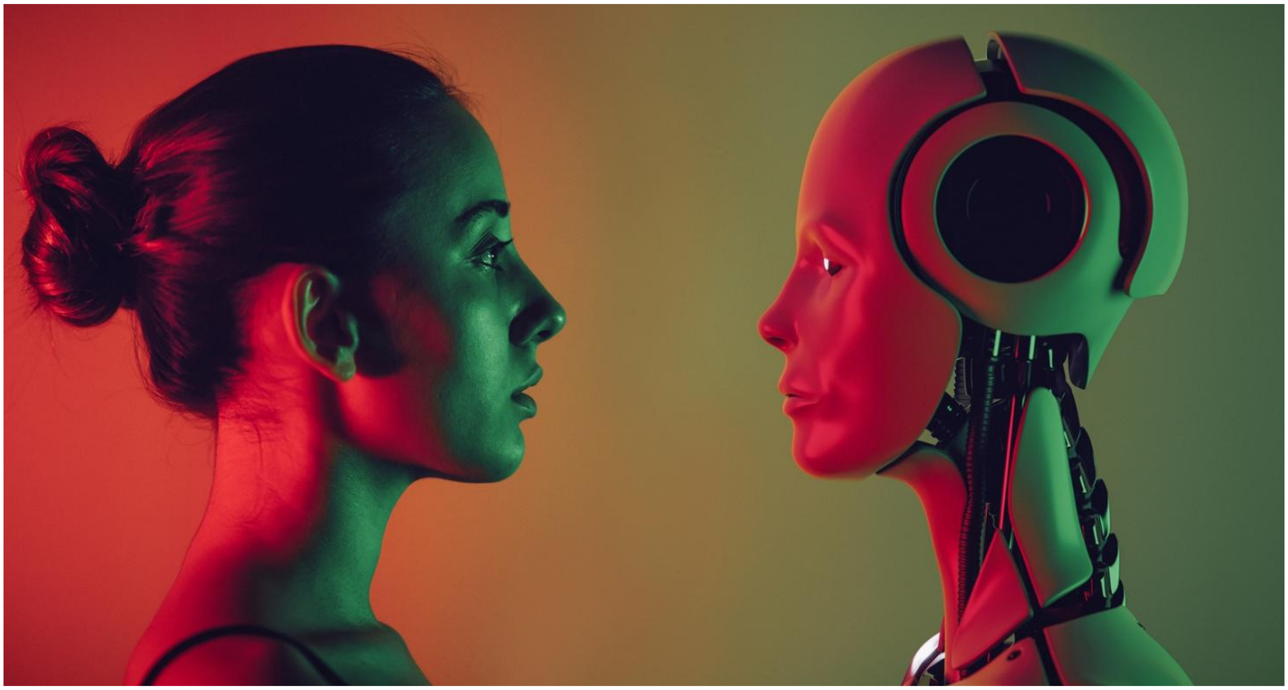

AI models struggle to form analogies when considering complex subjects, like humans can, meaning their use in real-world decision making could be risky. (Image credit: imaginima/Getty Images)

We know that artificial intelligence (AI) can't think the same way as a person, but new research has revealed how this difference might affect AI's decision-making, leading to real-world ramifications humans might be unprepared for.

The study, published Feb. 2025 in the journal Transactions on Machine Learning Research, examined how well large language models (LLMs) can form analogies.
They found that in both simple letter-string analogies and digital matrix problems — where the task was to complete a matrix by identifying the missing digit — humans performed well but AI performance declined sharply.

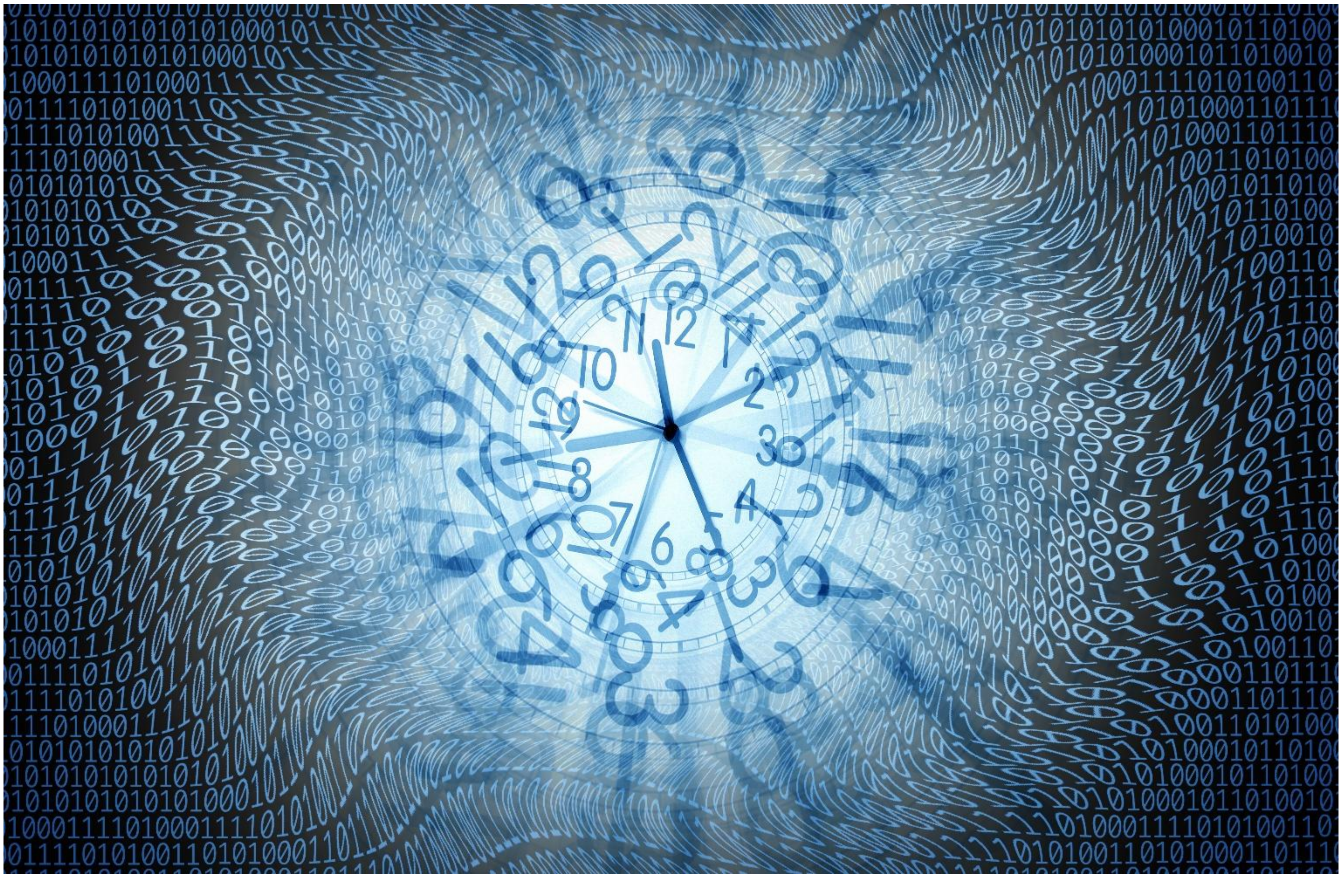

While testing the robustness of humans and AI models on story-based analogy problems, the study found the models were susceptible to answer-order effects — differences in responses due to the order of treatments in an experiment — and may have also been more likely to paraphrase.

Altogether, the study concluded that AI models lack “zero-shot” learning abilities, where a learner observes samples from classes that weren't present during training and makes predictions about the class they belong to according to the question.

"Letter string analogies have the form of 'if abcd goes to abce, what does ijkl go to?' Most humans will answer 'ijkm', and [AI] tends to give this response too," Lewis told Live Science. "But another problem might be 'if abbcd goes to abcd, what does ijkkl go to? Humans will tend to answer 'ijkl' – the pattern is to remove the repeated element. But GPT-4 tends to get problems [like these] wrong."

Why it matters that AI can't think like humans

Lewis said that while we can abstract from specific patterns to more general rules, LLMs don't have that capability. "They're good at identifying and matching patterns, but not at generalizing from those patterns."

Most AI applications rely to some extent on volume — the more training data is available, the more patterns are identified. But Lewis stressed pattern-matching and abstraction aren't the same thing. "It's less about what's in the data, and more about how data is used," she added.

To give a sense of the implications, AI is increasingly used in the legal sphere for research, case law analysis and sentencing recommendations. But with a lower ability to make analogies, it may fail to recognize how legal precedents apply to slightly different cases when they arise.

Given this lack of robustness might affect real-world outcomes, the study pointed out that this served as evidence that we need to carefully evaluate AI systems not just for accuracy but also for robustness in their cognitive capabilities.

MIT invents new way for QPUs to communicate — paving the way for a scalable 'quantum supercomputer'

News By Alan Bradley published April 6, 2025

A new device enables remote entanglement, allowing distant quantum processors to communicate with one another with reduced error rates.

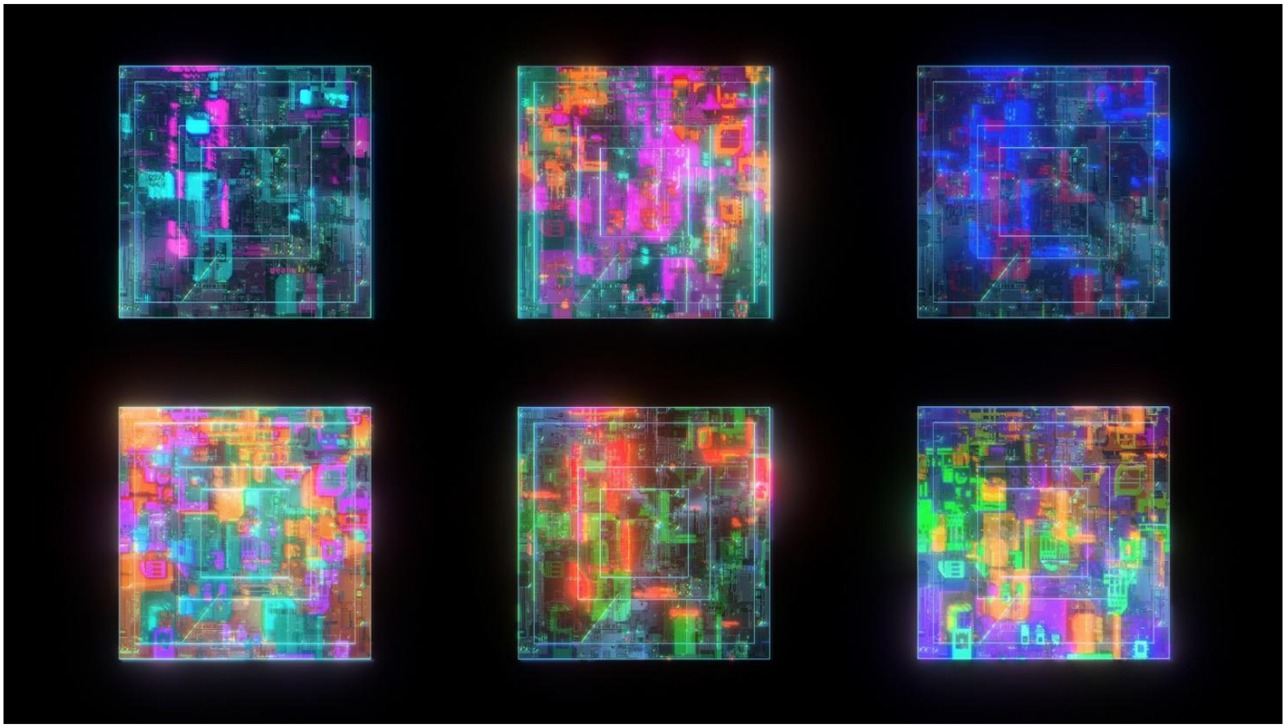

(Image credit: da-kuk/Getty Images)

Researchers have created a device that allows quantum processors to communicate with each other directly — an important step in developing practical quantum computers. It could mean both faster and less error-prone communication between processors.

Existing quantum architecture offers only limited communication between separate quantum processing units (QPUs). Such communication is "point-to-point," meaning that information has to be transferred in a chain across several nodes before reaching its destination. This increases the possibility of exposing the quantum information to noise and makes it more likely for errors to occur.

However, the new device developed by MIT scientists allows for "all-to-all" communication, so that all processors in a single network can communicate directly with any other processor. The researchers outlined their "remote entanglement" approach in a new study published March 21 in the journal Nature Physics.

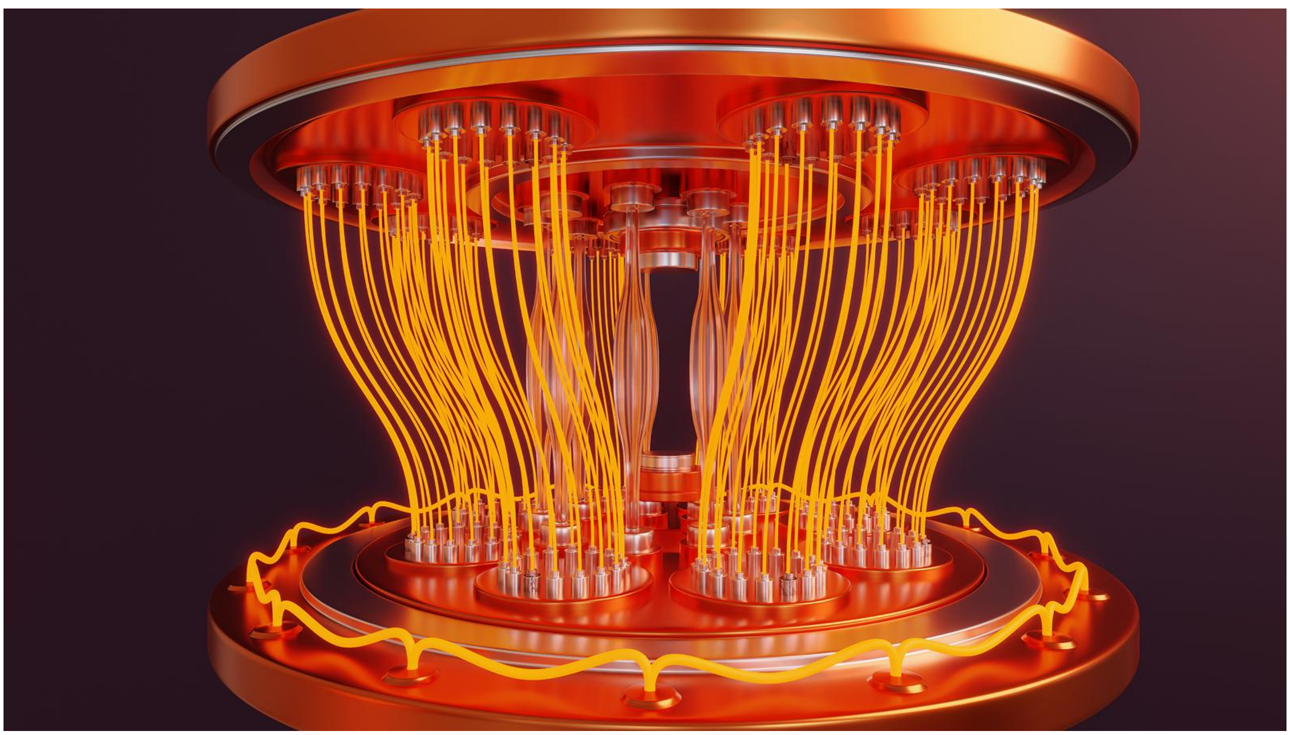

Remote entanglement is a state where two particles share the same state, and changes to one automatically affect the other. The distance between the two can be vast, with no currently known limit.

In testing, the researchers connected two quantum processors by way of modules, each comprising four qubits. Some of the qubits in each module were tasked with sending photons, light particles that can be used to transmit quantum data, while others were assigned to storing data.
The modules were linked together with a superconducting wire called a waveguide, with the modules serving as an interface between the larger quantum processors and the waveguide. The scientists said that any number of processors could be connected in this way, creating a highly scalable network.

The researchers then used microwave pulses to spark an individual qubit into emitting photons in either direction across the waveguide.

"Pitching and catching photons enables us to create a 'quantum interconnect' between nonlocal quantum processors, and with quantum interconnects comes remote entanglement," said senior author of the study William D. Oliver, Associate Director of the Research Laboratory of Electronics at MIT, in a statement.

Photonic distortion

Entanglement is a state where two particles become connected and share information, even at vast distances. A change in one entangled particle will immediately affect its partner. It's a critical phenomenon for quantum computing because it allows qubits to be correlated and act as a single system. This, in turn, lets us create algorithms that are impossible with classical computers.
However, just moving photons back and forth between modules doesn't automatically create entanglement. To achieve that, the team had to specially prepare both the qubits and the photon, so that after being transferred, the modules shared a single photon.

To force the two modules to share the same photon, they had to interrupt photon emission pulses at the halfway point. This essentially meant that half of the photon was absorbed on the receiving end while half was retained by the emitting module.

The problem with this method is that the photon becomes distorted while traveling across the waveguide, which can impact absorption and interrupt entanglement. To overcome this flaw in the architecture, the team had to distort the photons to encourage maximum absorption. By distorting photons prior to transmission, they were able to raise absorption levels to 60%, enough to ensure entanglement.

The work is broadly applicable to practical quantum computing applications, according to lead author of the study Aziza Almanakly, an electrical engineering and computer science graduate student.
"In principle, our remote entanglement generation protocol can also be expanded to other kinds of quantum computers and bigger quantum internet systems," Almanakly said.

Flat, razor-thin telescope lens could change the game in deep space imaging — and production could start soon

News By Owen Hughes published March 28, 2025

Scientists have developed an impossibly thin telescope lens that addresses a key astronomical challenge in a new study funded by NASA and DARPA.

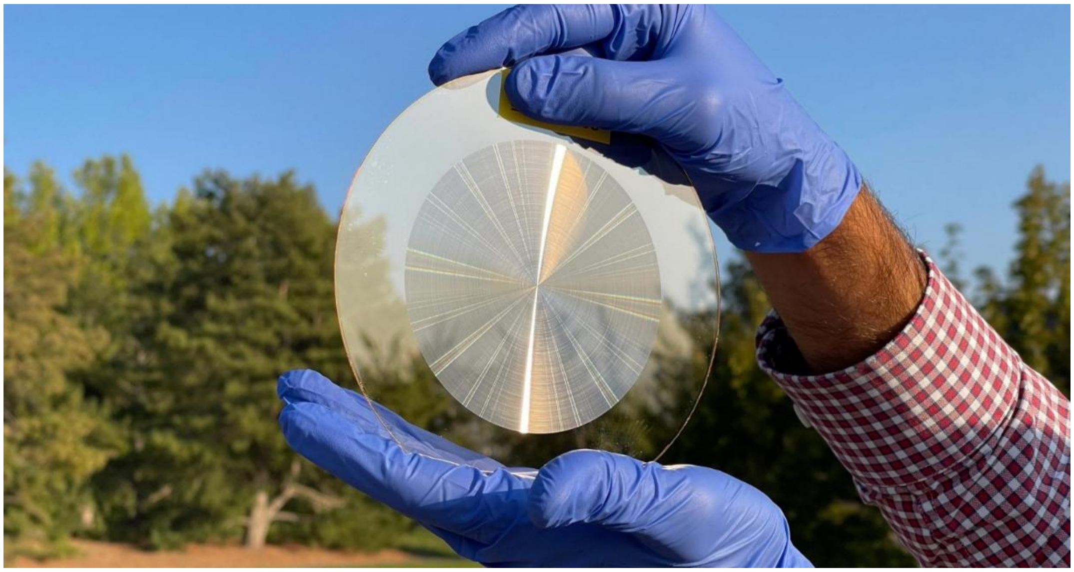

(Image credit: The Menon Lab/The University of Utah)

A new type of flat, razor-thin telescope lens could transform deep-space stargazing by making it possible to mount lightweight but powerful telescopes onto aircraft and satellites, scientists say.

Refractor telescopes normally use curved lenses to magnify distant objects through a process called refraction. Similar to a magnifying glass, the curved lens of a telescope bends light and directs it to a focal point, making objects appear larger.
However, traditional lenses quickly become impractical for space telescopes studying stars or galaxies millions of light years away. This is because the further away an object is, the more magnification is required to bring it into focus, and therefore the thicker and heavier the lens needs to be.

That's why scientists have explored flat lenses, which should in theory be lighter and less bulky. The challenge with them, however, is that light interacts with them differently than with curved lenses.

Visible light is a type of electromagnetic radiation, which is transmitted in waves or particles at different wavelengths and frequencies. When light passes through a flat lens, it diffracts, scattering wavelengths in multiple directions and resulting in a blurry, unfocused image.
But a new "multilevel diffractive lens" (MDL) developed by scientists features a multi-level structure consisting of "microscopically small concentric rings." These effectively channel different wavelengths of light towards the same focal point to create a sharp, color-accurate image.

The new 100-millimeter (3.9-inch) diameter lens, which has a 200 mm (7.8 in) focal length, is just 2.4 micrometers thick. Optimized for the 400 to 800 nm wavelength range for visible light, this lens is much lighter than a conventional curved lens and eliminates color distortions.

The scientists published their findings Feb. 3 in the journal Applied Physics Letters. The study was funded by the Defense Advanced Research Projects Agency (DARPA), NASA and the Office of Naval Research.
"Our demonstration is a stepping stone towards creating very large aperture lightweight flat lenses with the capability of capturing full-color images for use in air- and space-based telescopes," lead study author Apratim Majumder, assistant professor in electrical and computer engineering at the University of Utah, said in a statement.

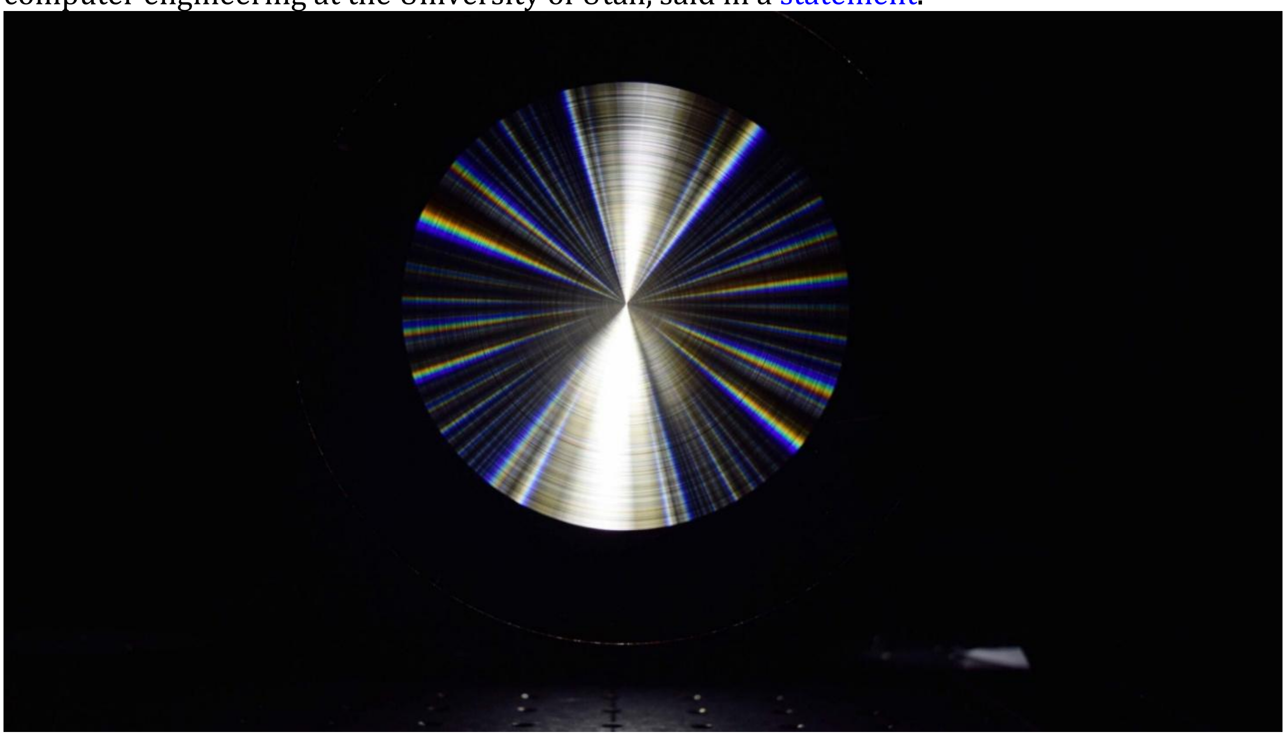

(Image credit: The Menon Lab/The University of Utah)

Ahead of the curve

Scientists have designed flat lenses in the past, most notably the fresnel zone plate (FZP), which features concentric ridges etched across the surface. However, the ridges of FZPs break light into separate wavelengths and diffract them at different angles, resulting in color distortions.

The MDL is unique in that its concentric rings exist at varying depths within the lens itself. As light passes through, the microscopic indentations adjust how different wavelengths diffract, preventing them from spreading apart as they normally would. This controlled diffraction brings all wavelengths of light into focus at the same time, resulting in a sharper, color-accurate image.

As well as avoiding the color distortions of FZPs, the researchers said the new flat lens offered the same light-bending power as traditional curved lenses. In the study, they used the MDL to capture images of the sun and moon. Lunar images they took revealed key geological features, while they also used it in solar imaging to capture visible sunspots.

"Simulating the performance of these lenses over a very large bandwidth, from visible to near-infrared, involved solving complex computational problems involving very large datasets," Majumder said in the statement. "Once we optimized the design of the lens' microstructures, the manufacturing process required very stringent process control and environmental stability."

The researchers said the technology had applications in astronomy, astrophotography and other "long-range imaging tasks" including "airborne and space-based imaging applications." What's more, production may not be far off.

"Our computational techniques suggested we could design multi-level diffractive flat lenses with large apertures that could focus light across the visible spectrum and we have the resources in the Utah Nanofab to actually make them," study co-author Rajesh Menon, professor of electrical and computer engineering at University of Utah, said in the statement.

Groundbreaking amplifier could lead to 'super lasers' that make the internet 10 times faster

News By Peter Ray Allison last updated June 2, 2025

Scientists have designed an amplifier that can transmit 10 times more information per second than current fiber-optic systems can, which could be helpful for medical treatment and diagnosis.

(Image credit: Baac3nes via Getty Images)

Scientists have developed a new type of laser amplifier that can transmit information 10 times faster than current technology.
Laser amplifiers boost the intensity of light beams. This particular amplifier achieves a tenfold increase in transmission speed by expanding the bandwidth, or wavelengths of light, at which the lasers can transmit information.

The amount of information we generate and transmit is growing every day. Due to the proliferation of streaming services, smart devices and generative AI, Nokia Bell Labs predicted in their Global Network Traffic Report that the amount of data traffic will double by 2030.
Current optical-based telecommunication systems transmit information by sending pulses of laser light through fiber-optic cables, which are thin strands of glass. The capacity — the amount of information that can be transmitted — is determined by the amplifier's bandwidth (the wavelengths of light that it can amplify). As data traffic increases, bandwidth therefore becomes crucial.

Most lasers used for modern telecommunications, such as internet communications, require an amplifier. These work by a process called stimulated emission, which uses an incoming photon to stimulate the release of another photon with the same energy and direction.

Scientists have now designed a new type of laser technology that can transmit information using a technology called high-efficiency optical amplification. The researchers published their findings April 9 in the journal Nature.
"The amplifiers currently used in optical communication systems have a bandwidth of approximately 30 nanometers," lead author Peter Andrekson, a professor of photonics at Chalmers University of Technology in Sweden, said in a statement. "Our amplifier, however, boasts a bandwidth of 300 nanometers, enabling it to transmit ten times more data per second than those of existing systems."
The new amplifier is made of silicon nitride, a hardened ceramic material that is resistant to high temperatures. The amplifier uses spiral-shaped waveguides to efficiently direct the laser pulses to remove anomalies from the signal. The technology has also been miniaturized so that multiple amplifiers can fit onto a small chip.

The researchers chose spiral waveguides over other waveguide types because they enable longer optical paths to be created within a small area. This enhances useful effects such as four-wave mixing, which occurs when two or more optical frequencies are combined together to amplify the output with minimal noise (external interference that can disrupt the quality of the signal).

Because the speed of light is constant, the laser light itself does not travel any faster than that from conventional lasers. However, the larger bandwidth enables the new amplifier to transmit 10 times more data than conventional lasers can.

The amplifier currently functions in a wavelength range of light 1,400 to 1,700 nanometers, which is within the short-wave infrared range. The next stage in the research will be to see how it operates over other wavelengths, such as those for visible light (400 to 700 nanometers) and a broader range of infrared light (2,000 to 4,000 nanometers).

The new amplifier has multiple potential applications, including medical imaging, holography, spectroscopy and microscopy, according to the statement. The miniaturization of the technology could also make lasers for light-based applications smaller and more affordable.

"Minor adjustments to the design would enable the amplification of visible and infrared light as well," Andrekson said. "This means the amplifier could be utilised in laser systems for medical diagnostics, analysis, and treatment. A large bandwidth allows for more precise analyses and imaging of tissues and organs, facilitating earlier detection of diseases."

'Crazy idea' memory device could slash AI energy consumption by up to 2,500 times

News By Owen Hughes published August 6, 2024

By performing computations directly inside memory cells, CRAM will dramatically reduce power demands for AI workloads. Scientists claim it's a solution to AI's huge energy consumption.

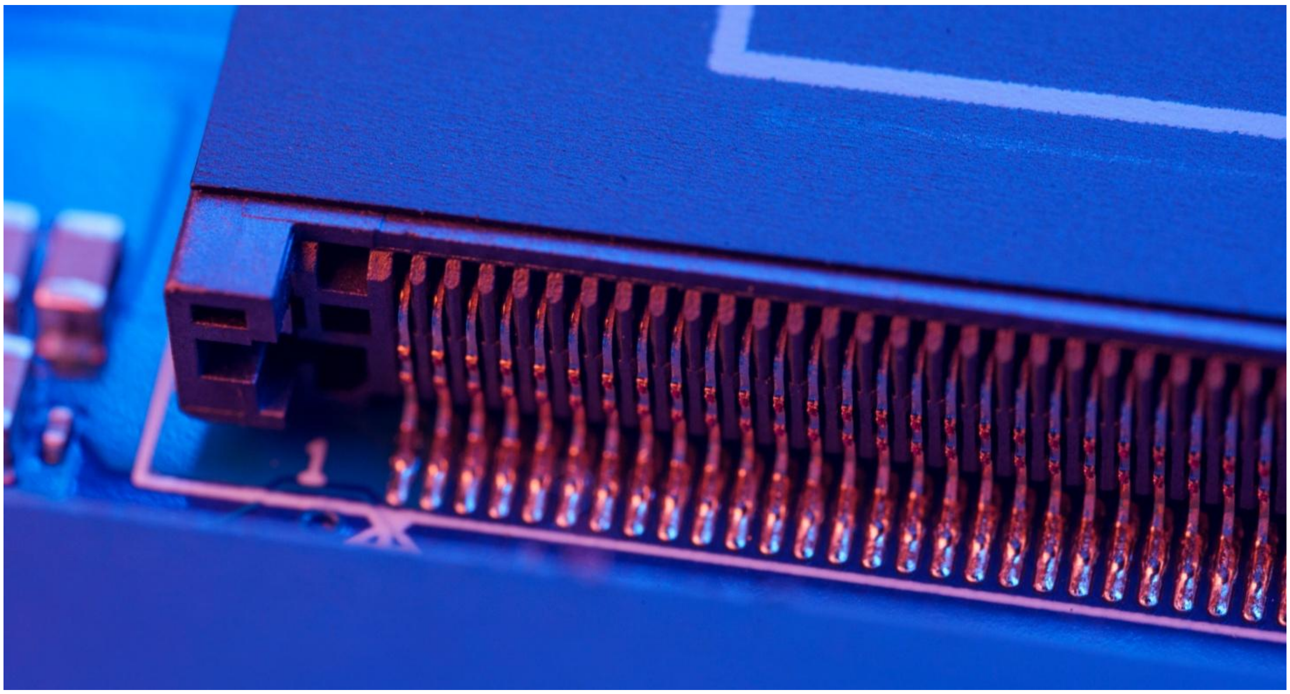

(Image credit: Serhii Prystupa/Getty Images)

Researchers have developed a new type of memory device that they say could reduce the energy consumption of artificial intelligence (AI) by at least 1,000.
Called computational random-access memory (CRAM), the new device performs computations directly within its memory cells, eliminating the need to transfer data across different parts of a computer.

In traditional computing, data constantly moves between the processor (where data is processed) and the memory (where data is stored) — in most computers this is the RAM module. This process is particularly energy-intensive in AI applications, which typically involve complex computations and massive amounts of data.

-

According to figures from the International Energy Agency, global energy consumption for AI could double from 460 terawatt-hours (TWh) in 2022 to 1,000 TWh in 2026 — equivalent to Japan's total electricity consumption.

Related: Intel unveils largest-ever AI 'neuromorphic computer' that mimics the human brain

In a peer-reviewed study published July 25 in the journal npj Unconventional Computing, researchers demonstrated that CRAM could perform key AI tasks like scalar addition and matrix multiplication in 434 nanoseconds, using just 0.47 microjoules of energy. This is some 2,500 times less energy compared to conventional memory systems that have separate logic and memory components, the researchers said.

The research, which has been 20 years in the making, received financial backing from the U.S. Defense Advanced Research Projects Agency (DARPA), as well as the National Institute of Standards and Technology, the National Science Foundation and the tech company Cisco.

Jian-Ping Wang, a senior author of the paper and a professor in the University of Minnesota's department of electrical and computer engineering, said the researchers' proposal to use memory cells for computing was initially deemed "crazy."

"With an evolving group of students since 2003 and a true interdisciplinary faculty team built at the University of Minnesota — from physics, materials science and engineering, computer science and engineering, to modeling and benchmarking, and hardware creation — [we] now have demonstrated that this kind of technology is feasible and is ready to be incorporated into technology," Wang said in a statement.

The most efficient RAM devices typically use four or five transistors to store a single bit of data (either 1 or 0).

CRAM gets its efficiency from something called "magnetic tunnel junctions" (MTJs). An MTJ is a small device that uses the spin of electrons to store data instead of relying on electrical charges, like traditional memory. This makes it faster, more energy-efficient and able to withstand wear and tear better than conventional memory chips like RAM.

CRAM is also adaptable to different AI algorithms, the researchers said, making it a flexible and energy-efficient solution for AI computing.

The focus will now turn to industry, where the research team hopes to demonstrate CRAM on a wider scale and work with semiconductor companies to scale the technology.

'It might pave the way for novel forms of artistic expression': Generative AI isn't a threat to artists — it's an opportunity to redefine art itself

News By Remo Pareschi published November 23, 2024

Rather than fearing the rise of generative AI, new technologies may allow creatives to define and express themselves in completely new ways.

One of the key areas that artificial intelligence (AI) threatens to disrupt is human creativity — and the rise of generative AI has certainly thrown art into the spotlight. While fears remain that AI may replace human input and agency across society, a different approach suggests humans will meld with AI in some capacity — with the new technologies augmenting us rather than undermining us.

In "Centaur Art: The Future of Art in the Age of Generative AI" (Springer, 2024), computer scientist Remo Pareschi explores the notion of "centauric intelligence" — an integration of human and computing intelligence — and its impact on the future of art. In this excerpt, Pareschi explains how our primal fears are misguided and argues that the rise of AI may, in fact, help human creative endeavors transcend to new heights.

As we revisit our initial, most pressing concern — will artificial intelligence surpass human creativity in artistic endeavors? — it's crucial to recognize that the question, as commonly posed, is somewhat misdirected. Unlike games such as chess or Go, art cannot be measured by a simple win-lose metric. The fear, however, is similar: could humans be supplanted in artistic production as they have been in high-level chess?

Our analysis, grounded in cognitive considerations and carried out in the pages to come, suggests a nuanced answer. At the pinnacle of creativity — where inspiration, conception and originality reign — AI is an enhancer, aiding in realizing powerful and original works. In such contexts, the human artist's role remains paramount, with AI serving as a tool to augment their creative vision. Conversely, AI's role can become more structured or repetitive. It can automate the production of routine outputs such as certain types of commercial illustrations, brochures, or video game characters. This automation significantly impacts professionals in these fields, who may rely on such work for their livelihood. Consequently, these individuals may view the advent of AI technologies with apprehension and concern.

Economic and legal implications

The economic threat of generative platforms is closely intertwined with concerns about copyright infringement. A notable legal action in this arena occurred in January 2023, when a group of artists filed a class-action lawsuit against Stability AI, Midjourney, and DeviantArt. The lawsuit centered on the alleged unauthorized use of artists' works to train AI tools, sparking a debate over the legal and ethical boundaries of AI in art.

The highlighted lawsuit underscores the ongoing tension between AI's innovative potential in art and individual artists' rights. Plaintiffs in the case argued that AI tools were creating derivative works based on their styles without proper authorization or compensation. In contrast, the AI art companies defended their actions, asserting that the AI-generated images were transformative and original, thus not violating any laws. In a significant development in October 2023, a US judge dismissed most of these claims due to a lack of direct infringement evidence.

On the other hand, to complicate matters further, the US Copyright Office in September 2023 rejected copyright protection for *Theatre D'opera Spatial*, an artwork predominantly created by AI and crafted by artist Jason Allen. This artwork had previously won an art contest, but the Copyright Office's decision was based on the lack of significant human intervention in its creation. This stance contrasts with the position of Lawrence Lessig, a Harvard Law School professor and a renowned expert on Internet and law issues. Lessig advocates for the recognition of copyright in prompt-generated artworks, arguing that they are original and creative, embodying the human input of the prompter. He believes these works should be treated on par with other technologically aided art forms and that they have the potential to invigorate the art world by inspiring current artists and drawing in new audiences interested in AI's artistic capabilities.
Interestingly, international perspectives on this issue vary. In November 2023, a Chinese court ruled that AI-generated content could be protected under copyright law, which starkly contrasts with the human authorship requirement under U.S. copyright law.

Thus, despite these legal battles, the issue of copyright infringement in AI art remains unresolved, with the potential for future challenges and changes in various jurisdictions. This uncertainty underscores the evolving nature of art in the digital age and the need for a balanced approach that respects both innovation and artists' rights.

A striking example of the social impact of generative platforms is the publication of *Sunyata* by Eris Edizioni, an Italian graphic novel authored by philosopher and digital artist Francesco D'Isla that combines AI-generated images with traditional text. The novel ignited controversy among artists and comic book creators, who criticized its use of AI as undermining artistic integrity and economic fairness. The author and publisher defended

their work as a legitimate artistic endeavor, emphasizing their careful use of prompts and adherence to a Creative Commons license.
This case exemplifies the broader debate surrounding AI in art: balancing technological innovation with ethical considerations and the economic interests of human artists. As generative platforms continue to evolve, these discussions will likely intensify, shaping the future of art in the digital era. And yet, we might say that this is nothing new, with history repeating itself — the intersection of technology and art has always been a crucible of innovation and controversy.

Exploring new artistic avenues

Indeed, throughout history, technological advancements, from ink and paper to the development of cameras and computers, have continually opened new avenues for artistic expression. Yet, each technological leap has also brought ethical and social challenges, echoing today's tensions between generativity and control, authenticity and originality, and the impact of digital platforms on the art market.

Historically, resistance to new art forms and technologies is not a novel phenomenon but a recurring pattern. For instance, the advent of photography in the 19th century was initially met with skepticism by some critics and artists, who viewed it as a mechanical, uncreative process threatening traditional art forms. Similarly, introducing sound and color in cinema faced opposition, with concerns about diminishing the artistic value of silent and black-and-white films. Just as photography and color cinema once disrupted artistic norms, today's generative platforms like DALL-E, Stable Diffusion, and MidJourney are provoking similar debates in the art world.

To navigate this complex landscape, the insights of Walter Benjamin and John Maynard Keynes offer valuable perspectives. Benjamin was concerned about losing aura due to the mechanical reproduction of the work of art, but also recognized its democratizing potential. Generative platforms like DALL-E, Stable Diffusion, and MidJourney take this democratization one step further, making artistic creation more accessible and fostering a sense of community and creativity. Allowing users to generate and share art potentially fulfills Benjamin's social role in art.

While Benjamin's insights shed light on the cultural implications of technological advancements in art, John Maynard Keynes's concept of technological unemployment offers a crucial economic perspective. Among his many contributions to economic theory, Keynes, a prominent economist of the 20th century, explored the job loss caused by technological advancements, a form of structural unemployment that he viewed as a 'temporary phase of maladjustment'. This perspective is particularly relevant when considering the fears among artists about AI-powered generative platforms. While some artists worry about the potential for AI to diminish the demand for human-made art, thus impacting their livelihoods, it's crucial to recognize that technological changes can also create new job opportunities and artistic avenues.

Keynes' insights remind us that, despite initial disruptions, technological advancements often lead to the emergence of new roles and industries. In the context of art, generative AI might challenge traditional practices and pave the way for novel forms of artistic expression and collaboration. This evolution necessitates a redefinition of art and its creators, a task we aim to address in this book.