https://doi.org/10.1038/s44460-025-00011-1

AI and the rise of intelligent sensing

Check for updates

From wearable health monitors to autonomous robots, AI is reshaping how sensors collect, interpret, and act on data, with co-design and benchmarking key to their success, finds Andy Extance.

CREDIT: DRAFTER123

Huanyu 'Larry' Cheng from Pennsylvania State University went from being skeptical when his collaborators first proposed using artificial intelligence (AI) to analyse their sensor data to being "amazed". Their stretchable, wearable, sensor was recording vibration and electrical muscle signals in patients who'd had throat surgery[1]. Classical data analytics had failed in their aim of estimating the progress of the patients' rehabilitation from the chaotic-seeming data.

After taking on the proposal of Hongcheng Xu and Libo Gao at Xidian University in China, Cheng says that he was impressed by how AI "classified the information in a well-organized manner". "It helped us to tell the story of information we cannot really see through the traditional lenses of data analytics," says Cheng.

Having been converted by this first study, Cheng is excited by the potential AI brings to sensing in healthcare, through its key strength in data analysis. He wants to exploit AI to go from using sensor data to estimate people's current health to predict what will happen to their future health. Cheng hopes it might be possible to predict the varied responses different individuals have to COVID-19 infection, for example.

Today AI and sensing are together going through a double revolution. Sensors are proliferating, for example with modern cameras and imaging devices creating vast streams of data. And AI's expansion is even more explosive, slipping into and feeding upon these data streams.

These trends are now shaping sensor research. Therefore, *Nature Sensors* interviewed Cheng and six other leading researchers at the interface of AI and sensing to identify how the two disciplines can together enable future technologies. Their responses follow three main directions.

The first is in data analysis. AI can extract meaning from highly complex sensor data. When data is abundant, AI can power transformative applications such as robots that are intelligent, autonomous and able to adapt to changes in themselves and their surroundings. And where data is scarce, expensive and hard-to-obtain, AI can help understand it better.

The next important direction is simultaneously optimizing the sensor producing the data and the AI analysing it. This can help reduce the growing amounts of energy that both sensing and using AI consume. It can also reduce or eliminate the need to transmit data to be processed remotely, which leads to latency, delaying the process of sensing and responding.

The final important direction of AI and sensor research is to properly assess AI-enabled sensing by benchmarking system performance and privacy based on their form. Building on this approach will enable AI's successful and fair deployment in sensing.

By working in these three directions, sensing researchers can deliver critical advances in applications such as machine vision, robotics, healthcare and neuromorphic sensing. They can also work towards AI tools that can work on different types of sensing hardware, and sustainable approaches to implementing the new technology.

Model behaviour

In most cases in sensing, the type of AI used differs from the kinds most familiarly used in products like ChatGPT, Microsoft Copilot and Google Gemini. These large language models (LLMs) use self-attention, looking at nearby words to help process and understand the meaning of each word in long sequences of text in one go. As the data points produced by sensors are not related in the same way as words in a sentence, sensing AI is often based on recurrent neural networks (RNNs) that preceded LLMs technologically. In language, RNNs would process a sequence word-by-word. In sensing they usually process data point-by-point in a one-dimensional sequence.

RNNs enabled AI to process such sequential data by including a 'memory' or hidden state to connect each data point with the previous one. Early RNNs suffered from a problem known as vanishing gradients that fatally slowed down the systems, which long short-term memory (LSTM) RNNs resolved.

For two-dimensional image processing, convolutional neural networks (CNNs) are the foundation of modern computer vision. They use randomly generated two-dimensional filters that scan across and down an image, extracting local spatial features.

Today, Sami Haddadin from the Mohamed bin Zayed University of Artificial Intelligence in Abu Dhabi uses 'physics-informed' learning to teach robots manipulation[2]. Haddadin says that developing his team's own dedicated AI techniques is the best way to take full advantage of relatively small and costly experimental data sets.

Haddadin adds that robotics is driving large, complex AI systems known as world models. These typically represent three-dimensional environments, including their physical and spatial properties. They require large volumes of real-world data, mostly images and video currently. Haddadin calls using AI to extract meaning from images in this way and bringing together different modes of sensing, 'making sense out of the sensor'. He dreams of AI-powered understanding of sensor data enabling a fully autonomous robot that can perceive, learn, and act like humans. "We're obviously pretty far away from that," he says.

At the opposite end of the complexity scale, to help tackle cost and energy consumption challenges, sensor developers can use TinyML, where ML stands for machine learning. TinyML can run simple versions of existing neural network architecture on devices orders of magnitude cheaper and less power hungry than the chips that run LLMs. They can run directly on sensors and process data locally at the edge without having to send data to the cloud, although their capabilities are more limited.

Cheng and colleagues demonstrated how TinyML could help in data analysis by making predictions about diseases affecting motor neurons with sufficiently reliable statistical power in babies. The team used wearable inertial measurement unit (IMU) sensors to capture miniscule motions[3]. Although they only measured 23 infants, they could use a bespoke TinyML model for a linear regression analysis of the data. "Where we have limited data, AI can synthesize critical information," Cheng stresses.

Future vision

AI has similarly improved the usefulness of low-resolution visual data. For example, in 2017, Aydogan Ozcan and his team at University of California, Los Angeles, showed this by using CNNs to significantly improve the resolution of microscopes without otherwise adapting them[4]. They trained a CNN with low- and high-resolution versions of images of tissue samples. After training, they could input low-resolution images, and the CNN would output higher-resolution versions. Ozcan says that it should be possible to apply AI to improve 'any measurement device' in this way.

Optimizing sensors specifically to communicate with AI rather than humans can make them 'cheaper, faster and better', Ozcan says. As one example, he and his UCLA colleagues designed an inexpensive paper-based flow assay for Lyme disease[5]. While in some ways like lateral flow tests for COVID-19, the pattern that their test forms is more complex. The assay therefore captures the pattern with a digital camera and classifies it with AI. "AI builds a classifier to tell you, based on these spots, if they're positive or negative," Ozcan explains.

Unlike such applications that use few, low-resolution images, machine vision typically requires high data volumes, costly hardware and uses lots of energy, explains Jianshi Tang from Tsinghua University in Beijing, China. This application spans image, infrared, radar and ultrasonic sensors, and the amount of data recorded is growing exponentially. This also leads to latency arising from issues like translating the images from analogue to digital form. The solution that Tang and his colleagues propose would come in the form of optoelectronic memristors that directly compress light intensity, wavelength and polarization into machine interpretable patterns[6].

Inspired by the human sensory system, the Tsinghua University researchers intend to integrate these optoelectronic memristors directly with CNNs. Tang and his colleagues argue that such neuromorphic devices could minimize latency and energy consumption by making decisions at the network's edge, rather than centrally. "We want to bring the sensing and the computing together closer and closer so that we can minimize the data transfer cost," he says.

This would also enhance privacy, as sensitive data never leaves the device. Such edge AI could also control intelligent systems such as autonomous robots. Energy efficiency improvements could enable more sophisticated wearable health monitors and environmental sensors for tracking tiny creatures like bees. If such neuromorphic systems can deliver on these promises, they will be centrally important to sensing.

Opportunities and challenges in tactile sensing

In robotics, integrating AI and sensing delivers systems that perceive, interpret, and interact with the world. Aude Billard from the Swiss Federal Technology Institute of Lausanne (EPFL) is primarily interested in how sensors can help control such robotic systems. Her team investigates soft robotics, which are flexible and deformable and often use tactile sensors, rather than vision.

In the past, tactile sensors were rigid, measuring solely pressure along the sensors' surface, Billard explains. Today, they are flexible and stretchable sensors that can detect forces in other directions, as well as strain. Such sensors enable fine manipulation and can be placed all over robot bodies, including in joints that bend.

Tactile sensors produce data two orders of magnitude faster than vision. And while pressure sensors are highly precise, the values they produce can drift from their calibration, for example as temperature changes. There were two options to deal with this issue, Billard explains. "Either stop working with sensors and wait until someone designs a better one, or try to make sense of it, by modelling the drift over time," she says.

As AI techniques became able to model the complex nonlinear effects, people have opted for the latter option. This helped enhance the sensors' accuracy and reliability. Billard explains that when she started using LSTMs around the year 2000, they could help predict pressure values of tactile sensors and how they might drift over time.

Another issue with tactile sensors is that when the robot touches something, the contact also deforms the robot, affecting the measurement, says Cecilia Laschi from the National University of Singapore. Human brains predict sensory input to minimize the need for constant sensing, Laschi explains. Researchers emulate this process in robots to reduce computational load while enhancing performance. Using a predictive system based on conventional modelling rather than AI, Laschi and colleagues were able to use tactile sensing to enable soft robots to perceive contacts and make decisions in real time, autonomously steering around the walls of a maze[7]. This approach helps distinguish deformations caused by movement from the robot itself from contact with another surface, she explains, and could be enabled by AI.

Laschi's team also uses continual learning algorithms, which involve sequentially training models for new tasks while preserving previously learned tasks. They allow soft robots to adapt as they degrade or as their morphology changes, ensuring long-term autonomy. Soft actuators puncture or leak, but continual learning controllers adapt to damage and then readapt as self-healing materials recover. The result is resilient soft robots that learn, endure wear, and relearn tasks.

Using continuous learning and other approaches to enable legacy systems to adapt to new hardware will be essential as sensing and AI systems become ever more common. Billard explains that in trying to adapt an algorithm to different sensors it's necessary to separate the fundamental perception that they perform, from the specific format of information the sensor generates.

What an AI system has learned with an older, V1, sensor may not be applicable to a newer, V2, sensor with higher resolution, different sensitivity, or entirely new capabilities, Billard explains. "One must determine which aspects of the V1 data remain valid for V2 and develop methods to reconcile discrepancies – essentially filling the sensory gap between the two generations of hardware," she underlines. In planning for such issues, data analysis can lead into the second key direction for AI's use in sensing – device design.

Co-design and end-to-end optimization

Laschi notes that carefully arranging the shape of soft robots and their touch and strain sensors can help minimize the amount of computational processing they need to do. It's even possible to build completely electronics-free robots, with mechanical and sensing devices powered by fluidics. Here, AI can help in the form of evolutionary algorithms, which mutate and select the robots' forms. "After several generations of evolution, you find the best design," Laschi says[8].

Ozcan is also leveraging AI in the design process to push the boundaries of what sensors can do. He calls his philosophy end-to-end design, co-optimizing sensors and AI. Rather than evolutionary algorithms, his team uses a machine learning technique called decision tree learning. This approach can explore vast design spaces and identify optimal configurations, pruning parts of the design one-by-one to see which are the most important for performance. "Sensors to deploy at large scale need to be simple to use, simple to fabricate, and cost effective," Ozcan says. "You use decision tree learning to throw away things that you don't need. You cut out branches that don't change the result much."

Using this approach means that Ozcan's sensors are extremely compact. That brings three benefits. First, the sensors consume very little power. Second, they cost less per test. Third, they avoid the tendency of AI to over-fit to the examples they are trained on. "If you throw an extremely large network to a relatively simple problem, then the network memorizes the answers, and it doesn't generalize, it doesn't learn."

Haddadin's team combines AI and real-world testing in an approach he calls 'artificial physical evolution'. An AI synthesis system called CREATOR first produces robots and/or sensor designs and then manufactures those designs, then uses the data recorded to produce further designs. His team is currently using this approach to develop new fingertips for robots, together with an evolutionary aspect to select the fittest. They are exploring different fingertip surface geometries with soft silicone-based materials, and magnetic pressure and strain sensors. "You can only simulate what you can measure and understand," he explains. "Better sensors empower better simulations. AI empowers better simulations, and builds better sensors."

Benchmarking and transparency

The fitness concept provides an opportunity to tackle the challenge of quantifying the performances of different sensors against each other. Haddadin's team has recently published a 'tree of robots', mapping robot capabilities independently of their format[9]. It ranks industrial robots based on task performance assessed by measurable benchmarks including force sensing, enabling an objective, open leaderboard. The robots ranked are more fit in some applications than others.

For Haddadin, for AI-enabled sensing and robotics to grow "we need global efforts instead of individual labs or companies driving the field". "I do not believe that we can achieve such major advancements with isolated efforts," he says. He would like to see information about sensors shared equally openly, because they play such a central role in providing information to other systems.

Mona Sloane from the University of Virginia in Charlottesville, USA and colleagues have proposed a similar framework for considering AI sensing risks[10]. That's partly because privacy and security are growing concerns as vision systems become more pervasive. But AI can help by processing data locally, minimizing the risk of sensitive information being exposed. "Many of our current approaches to AI risk, including privacy, start from data and model, but they don't take into account the material conditions under which data is produced by sensing in the first place," she says.

Sloane and her colleagues argue that assessments should include factors like the properties of sensors that shape and constrain how they can be used, and the economic practice and social frameworks that make sensors commercially viable. "If we make this adjustment, we can actually surface privacy risks earlier, not just after privacy issues have showed up and been flagged by users," Sloane says.

Such a framework would be akin to the Tree of Robots in that it's an effort to more holistically assess a technology while it is being designed, Sloane notes. "If we can elevate privacy to the same level of concern as performance, that will be a big win for everyone," she adds.

Sensing sustainably

TinyML has many obvious advantages in data security and beyond, Sloane observes, including lowering energy consumption. "As we are seeing the potential negative effects of AI's scaling laws – the bigger the model, the better – the more relevant systems with smaller energy and processing footprints become," Sloane says. TinyML can shift who holds the power in machine learning. She says farming communities have been "long engaged in a tension-filled relationship with smart farming equipment vendors about smart system coercion and the right to repair". They could use TinyML for community-driven soil analytics.

Some researchers might consider these and other not-wholly-technical issues, including bias and interpretability, beyond their remit. Yet with sensing and AI integrating rapidly, it's ever more important to account for them. "I think designers are in a good position to begin incorporating questions about calibration, documentation, data profusion, privacy, and waste into standard design reviews," Sloane says. "I also think that they ought to collaborate with social scientists and community organizations so that sensing systems are co-designed with those who will live with them."

Policymakers clearly have a role to play in ensuring this, and Sloane argues that their task is even bigger. They should build on existing risk-based AI governance tools, like the US National Institute of Standards and Technology's AI risk management framework, she says. They also should consider intervening in the social and economic systems that drive sensor proliferation, Sloane suggests.

This could involve incentivizing data-minimal architectures and addressing environmental impacts through product standards.

Towards transformative sensing

The challenges facing AI-enabled sensing therefore span multiple scales. At the global level, the technology must serve our societies fairly and sustainably. Between research groups, scientists must find ways to achieve robust benchmarking and comparison of their systems. These high-level goals will provide a secure framework to address technical issues. In doing so, developers can overcome the many data challenges they face, including by simultaneously designing sensors and the algorithms connected to them.

As such, although researchers' focus is currently on optimizing existing technologies, transformative changes will follow as the AI and sensing fields further integrate, says Cheng. Going beyond improving existing technology may lead to "something magnificent", he adds.

Andy Extance

Freelance writer, Exeter, UK.

e-mail: aextance@virginmedia.com

Published online: 15 January 2026

References

1. Xu, H. et al. *Nat. Commun.* **14**, 7769 (2023).
2. Voigt, F. & Haddadin, S. Multi-level structure vs. end-to-end-learning in high-performance tactile robotic manipulation. In *Conference on Robot Learning 2020* https://corlconf.github.io/corl2020/paper_516/ (2020).
3. Bao, B. et al. *Adv. Sci.* **11**, 2306025 (2024).
4. Zhang, Y. et al. *Optica* **4**, 1437–1443 (2017).
5. Ghosh, R. et al. *Nat. Commun.* **15**, 7124 (2024).
6. Wang, X. et al. *npj Unconv. Comput.* **2**, 15 (2025).
7. Wang, P. et al. *Nat. Commun.* **15**, 9978 (2024).
8. Corucci, F., Cheney, N., Kriegman, S., Bongard, J. & Laschi, C. *Front. Robot. AI* **4**, 258163 (2017).
9. Kirschner, R. J., Karacan, K., Melone, A. & Haddadin, S. *Nat. Mach. Intell.* **7**, 459–470 (2025).
10. Sloane, M. et al. *Nat. Mach. Intell.* **7**, 334–345 (2025).