difficult to directly correlate nanoscale fibril mechanics in 3D with the deformation and organization of the individual fibres in which they are embedded.

Newham *et al.* have now united 3D pCT and SAXS with image correlation to study intact rat intervertebral discs under compression (Fig. 1). The authors used 3D pCT to map out the structures and positions of individual fibres in the tissue. This map was then used to estimate which fibres contributed the most to the observed scattering signals when the intervertebral discs were analysed using SAXS. This allowed the nanoscale and microscale properties of individual fibres to be determined. Use of the map also reduced the number of X-ray scans needed for the 3D SAXS analysis, because the SAXS data were no longer needed to determine the orientations of the fibres. Finally, image correlation was used to compare data before and after compression, so that changes in collagen fibrils could be related to the deformation and architecture of the fibres.

The experiments show that the discs sustain their load-bearing function by coordinating load-sharing between collagen fibres and fibrils, with strain in fibres mostly being transmitted directly to fibrils. Moreover, Newham *et al.* showed that molecular pre-strain — any stretch or compression of the collagen molecules that existed before the load was applied — correlates with local tissue organization and influences fibril deformation and strain distribution in the discs. The authors also demonstrated that TomoSAXS can be extended to study larger human intervertebral discs, although this experiment was performed under static conditions without mechanical loading.

Several technical challenges still need to be solved before TomoSAXS can reach its full potential. One issue is that the experiments currently require high doses of radiation at a synchrotron facility (a source of extremely bright X-rays), of which there are relatively few worldwide. Moreover, each rat experiment took approximately four hours and the intervertebral discs had to be moved back and forth between two experimental stations at the synchrotron because different stations were needed to run the pCT and SAXS analyses. This requirement limits the ability of TomoSAXS to study mechanical processes that occur rapidly in soft tissues or involve dynamic loading.

These limitations do not diminish the importance of the method, however — they just set the starting point for continued development. It will be interesting to see whether TomoSAXS can be adapted to capture tissue behaviour during movements that more closely replicate the dynamic loadings that occur *in vivo* than did those studied in the current experiments, and with lower radiation exposure. Using one experimental station at the synchrotron, rather than two, would speed up the process[10], for example. The ability to correlate nanoscale structure with 3D tissue deformation could make TomoSAXS a powerful method for studying how hierarchical load distribution in biological tissues changes with ageing, disease or during a medical treatment — not only in intervertebral discs, but also in other collagenous tissues and fibrous materials.

Hanna Isaksson is in the Department of Biomedical Engineering, Lund University, 22363 Lund, Sweden.
e-mail: hanna.isaksson@bme.lth.se

1. Newham, E. *et al. Nature* **659**, 91–99 (2026).
2. Momose, A., Takeda, T., Itai, Y. & Hirano, K. *Nature Med.* **2**, 473–475 (1996).
3. Bay, B. K., Smith, T. S., Fyhrie, D. P. & Saad, M. *Exp. Mech.* **39**, 217–226 (1999).
4. Disney, C. M. *et al. Acta Biomater.* **92**, 290–304 (2019).
5. Dejea, H. *et al. Adv. Sci.* **11**, 2308811 (2024).
6. Fratzl, P., Jakob, H. F., Rinnerthaler, S., Roschger, P. & Klaushofer, K. *J. Appl. Crystallogr.* **30**, 765–769 (1997).
7. Gupta, H. S. *et al. Proc. Natl Acad. Sci. USA* **103**, 17741–17746 (2006).
8. Liebi, M. *et al. Nature* **527**, 349–352 (2015).
9. Schaff, F. *et al. Nature* **527**, 353–356 (2015).
10. Nygård, K. *et al. J. Synchrotron Radiat.* **31**, 363–377 (2024).

The author declares no competing interests.

Medical research

Using AI to provide personalized health advice

Emma Pierson

The use of artificial-intelligence tools in medicine is increasing. What types of technology might be on the horizon and how can such health-advice tools be tested? **See p.197**

Although people report increasing trepidation about the effects of artificial intelligence across many domains, health care emerges as a rare area of some optimism[1]. AI tools, in the form of large language models (LLMs), such as the chatbot ChatGPT, are widely used for health advice. It is estimated that one-third of US adults turn to chatbots for answers to health questions[2], and globally it has been suggested that about 230 million people ask ChatGPT health and wellness questions every week[3]. On page 197, Heydari *et al.*[4] present their response to this enormous demand by creating and assessing an AI personal health 'agent', which is designed to analyse an individual's health data and offer advice.

The system that Heydari and colleagues present combines three specialized AI agents. One acts as a data scientist that analyses the user's sleep, physical activity and other data, such as information from smart watches (Fig. 1). Another acts as a health-domain expert that incorporates clinical knowledge and published scientific research. The third functions as a health coach that helps the user to set goals and make plans to achieve them. A central AI 'orchestrator' supervises the three agents, directing tasks to the appropriate specialist. The authors did a great deal of work to separately test each component of their system. For example, they show that both users and health specialists prefer the health coach AI agent to the original LLM on which it relies.

Heydari and colleagues' work is a promising first step towards developing AI systems that can offer personalized health advice. User and expert assessments and objective evaluations — such as how accurately the system answers health questions about heart health, fitness, sleep and other relevant factors — suggest that the three specialized agents all generally outperformed the original LLM.

But much work remains to be done. As the authors caution, their system is a research prototype and is not a replacement for professional health care. Their comprehensive evaluations show that the system is not perfect. For example, the data-scientist agent writes computer code to analyse health data that, more than one-fifth of the time, cannot pass the automated tests that the authors designed to test its correctness.

Real-world roll out of the system would also raise privacy and security challenges. Agents that access sensitive user health data must be carefully controlled. Moreover, deployment of another AI consumer health tool, ChatGPT Health, raised concerns when an independent evaluation[5] found that it did not treat all emergency situations with the urgency they required, highlighting the caution and careful evaluation that real-world use necessitates.

Indeed, one of the most important challenges to overcome for such medical AI tools is how to evaluate them[6]. Heydari and colleagues' work highlights two evaluation challenges that are noteworthy: first, evaluating whether AI health systems improve

long-term health outcomes, and second, conducting rigorous evaluations while keeping pace with the AI frontier.

First, the authors focused on short-term assessments of their system — whether it can have high-quality conversations and carry out the tasks it is designed to perform. A natural next step would be to measure whether it improves long-term health outcomes. The best way to assess this would be with a type of clinical trial called a randomized controlled trial (RCT), in which a treatment group is given access to an AI system and the outcomes of this group are compared with the those of a control group that was not given access to such a system. Ideally, trials should track both long-term health outcomes and interactions with the AI system to enable measurement of how the technology causally affects a person's health and to determine what types of interaction predict better results.

Some completed trials highlight the potential of AI tools to improve health outcomes. An RCT of an AI chatbot, called Therabot, that provides mental-health support showed that users experienced reduced symptoms of depression, anxiety and other conditions relative to those of people in a control group[7]. An RCT of an AI-based diabetes prevention coach showed that it yielded similar benefits to those achieved when a participant was coached by a human[8].

Another RCT of an AI system, which is designed to help individuals reduce high blood pressure through personalized education and coaching, showed that AI use resulted in a reduction in blood pressure among the treatment group compared with that of a control group[9]. A chatbot designed to help people improve their lifestyle behaviours, when tested in an RCT, showed that, compared with the control group, it produced improvements in diet and physical activity[10]. Collectively, this body of evidence shows that AI-based interventions can improve health outcomes, and these improvements can be measured using the standard, well-developed techniques for assessing other types of medical intervention.

"One of the most important challenges to overcome for such medical AI tools is how to evaluate them."

But RCTs take time, as did the meticulous work of Heydari and colleagues. Such work must precede any RCT. This raises the second challenge — careful evaluations are time-consuming but AI is advancing extremely quickly. Many of the evaluations presented in the authors' work are based on the LLM Gemini 2.0, which was released in 2024. However, such is the pace of AI developments that this model is no longer close to state-of-the-art. The AI diabetes RCT trial[8] that began in 2021, started a year before ChatGPT was released to the public, and this diabetes prevention tool does not make use of LLMs at all.

How do scientists rigorously assess safety and effectiveness of such systems while keeping pace with the AI frontier? The answer might be that it won't be possible to do so — and, given a choice between keeping pace or achieving a rigorous evaluation, the latter must win in high-stakes settings such as health care. To take an optimistic viewpoint, older evaluations and frameworks can often be repurposed to analyse systems that use new AI models, with performance tending to improve as the models improve.

For example, Heydari and colleagues' framework is designed to be compatible with other LLMs. Hopefully the performance of their framework will only improve when stronger core LLMs are used — and the evaluations the authors designed provide a reusable road map for assessing whether such an improvement occurs. Then, the assessment of new generations of AI models would not require starting work from scratch. Instead it would be a matter of reapplying the lessons that have already been learnt from carefully developed evaluations.

AI systems that provide health advice have great potential to improve the accessibility of health information, including for underserved populations who have historically lacked ready access to medical advice[11]. Rigorous assessments of short-term safety and



RICARDO RUBIO/EUROPA PRESS/GETTY

Figure 1 | Smart-watch use. These devices can record personal health data that can be used by an artificial-intelligence tool, such as the one created by Heydari and colleagues[4].

effectiveness, such as those carried out by Heydari and colleagues, followed by RCTs to assess effects on long-term health outcomes, will measure whether AI systems can truly achieve their potential.

Emma Pierson is at the University of California, Berkeley, Berkeley, California 94703, USA.
e-mail: emmapierson@berkeley.edu

1. https://www.pewresearch.org/wp-content/uploads/sites/20/2025/03/pi_2025.04.03_us-public-and-ai-experts_topline.pdf
2. https://www.kff.org/public-opinion/kff-tracking-poll-on-health-information-and-trust-use-of-ai-for-health-information-and-advice/
3. *Nature* **658**, 7–8 (2026).
4. Heydari, A. A. *et al. Nature* **659**, 197–207 (2026).
5. Ramaswamy, A. *et al. Nature Med.* **32**, 1671–1675 (2026).
6. Agrawal, M., Chen, I. Y., Gulamali, F. & Joshi, S. *npj Digit Med.* **8**, 600 (2025).
7. Heinz, M. V. *et al. N. Engl. J. Med. AI* https://doi.org/10.1056/AIoa2400802 (2025).
8. Mathioudakis, N. *et al. JAMA* **334**, 2079–2089 (2025).
9. Kario, K. *et al. Eur. Heart J.* **42**, 4111–4122 (2021).
10. Wong, C. H. Y. *et al. npj Digit. Med.* **9**, 657 (2026).
11. Pierson, E. *et al. N. Engl. J. Med. AI* https://doi.org/10.1056/AIp2400889 (2025).

The author declares no competing interests.
This article was published online on 21 October 2026.

Materials science

Large-scale single crystals printed layer by layer

Gregory S. Rohrer

A 3D-printing method has made single crystals of metals and alloys in complex shapes and up to one metre in length, as needed for advanced technological applications. **See p.100**

The dazzling gemstones used in jewellery are single crystals of minerals, prized for their beauty and uniformity. Single crystals of materials also have crucial functions in silicon microprocessors and jet-engine turbines. However, technological applications have been limited because such crystals are rare in nature and are difficult or expensive to make for most synthetic materials. The breakthrough reported by Liu *et al.*[1] (page 100) is therefore newsworthy: the authors have discovered a way to produce single crystals of metals in a wide range of shapes and sizes, including complex geometries that cannot be achieved by established crystal-growth methods. The method they used is a type of additive manufacturing – which means that it works like a 3D printer[2] – and has the advantage of being implemented using readily available equipment.

Liu and colleagues' method bears some resemblance to the Verneuil process (also known as the flame fusion process), which was developed in 1893 to produce gemstones for watches and jewellery[3]. In the Verneuil process, powder of a raw material is melted by dropping it through a flame, and the resulting liquid droplets land on a rotating solid 'seed' crystal. The droplets solidify so that the atoms in the material form a lattice with the same crystallographic orientation as the seed crystal, which grows as more powder is melted and deposited. But although the process is fast and simple, the quality of the resulting crystals is not as good as can be achieved by the methods used to grow crystals for semiconductor applications.

In Liu and co-workers' method, raw material in the form of a wire is melted by an electric arc, producing droplets that are continuously deposited onto a growing crystal (Fig. 1). However, the authors' method differs from the Verneuil process in an important way. Instead of using a constant heat source for melting, the current that produces the arc is pulsed so that a large amount of heat is injected quickly (in 0.5 seconds) – not only into the wire, but also into a thin layer on the surface of the growing crystal. This pulse re-melts some of the previously solidified material on the crystal, after which the current is lowered to enable re-solidification over a longer time interval (9.5 s). The cycle of re-melting and re-solidification provides a self-correction mechanism by eliminating any misoriented crystal grains that might have formed on the surface.

Most processes that use a melt for growing single crystals require a mould to hold the liquid material, as is the case with the method used to produce single-crystal nickel-based alloys for use in turbine blades[4]. Liu and colleagues' approach eliminates this constraint and the contamination that can come with it. Moreover, the lack of a mould makes it possible to grow crystals to indefinite lengths – the authors report that they have grown single crystals as long as 1 metre. The use of moulds is also challenging when melts are extremely hot. Liu *et al.* show that their mould-free method enables the growth of single crystals of tungsten, which melts at 3,422 °C.

Synthetic and natural materials are nearly always polycrystalline – made up of many microscopic crystals joined together at interfaces called grain boundaries[5]. These boundaries are often considered undesirable because they can make transparent materials opaque[6]; cause structural materials to become susceptible to

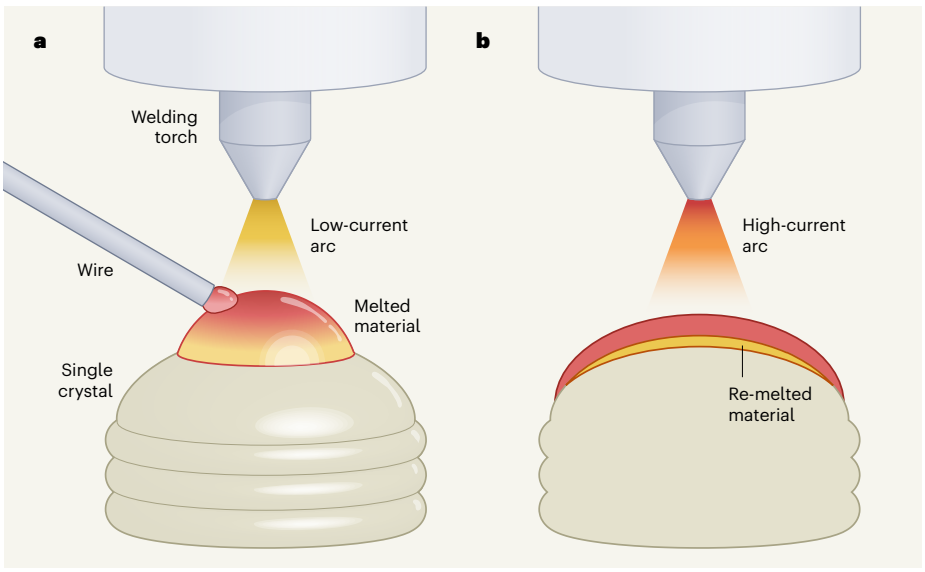


Figure 1 | A process for layer-by-layer printing of single crystals of metals and alloys. **a**, In Liu and colleagues' method[1], a wire is melted onto a previously printed layer of the same material using an electric arc, which is generated using a low current. The melt does not spread widely on the surface. **b**, The arc current is then increased to produce a short burst of intense heating, causing the deposited material (the red layer in this cross section of the printed object) to cover the whole surface. The atoms in the melt align with the crystal lattice of the previous layer on cooling (not shown), enlarging the single crystal. A thin volume of material in the previous layer also melts, including any tiny misaligned crystal grains that formed on that layer. The method enables single crystals up to 1 metre in length to be printed. (Adapted from Fig. 1k and Supp. Fig. 2h of ref. 1.)