

high precision. In 2024, after two decades of research, scientists finally succeeded in measuring the transition frequency with a 13-digit precision⁷.

The next step was to implement a system that periodically compares the frequency of the oscillator (the laser light wave) with the thorium-229 transition frequency to correct any drift in the frequency of the laser light. This major advance is now reported independently by Toscani De Col *et al.* and by Huang and colleagues.

Both groups used thorium-229 ions embedded in crystals of calcium fluoride, which enabled a large number of nuclei to be fixed into a small volume of space, and checked the transition frequency by analysing the absorption of laser light by the thorium ions (Fig. 1). However, the laser configurations used in the two studies were different. Toscani De Col and colleagues' laser configuration is simpler than that used by Huang *et al.*, but produces lower-power light and therefore a weaker correction signal from the thorium nuclei. The clock stability achieved in the two experiments is not yet comparable to that of state-of-the-art optical clocks, but the findings are nevertheless a milestone in the field.

One important aspect of a good clock is reproducibility – it must be possible for anyone to fabricate clocks that 'tick' at the same frequency. Huang *et al.* tested this by comparing the thorium-229 transition frequency with those of nuclei that had been embedded in crystals prepared by researchers in their collaboration who were based at various locations. They also compared their frequency measurements with those previously reported by a different research group⁸. These measurements were all in good agreement with each other.

The energy of the thorium-229 transition is so low because two key factors, arising from the strong and electromagnetic forces between particles in the nucleus, almost completely cancel each other out. Any changes to these two fundamental forces would therefore alter the transition frequency of the nuclear clock much more than they would alter the transition frequency of atomic clocks. This means that the nuclear clocks are extremely sensitive to variations of certain fundamental constants.

Such variations would be induced by the presence of a form of dark matter known as ultralight dark matter⁹, which scientists have previously searched for using atomic clocks. Toscani De Col *et al.* report that their thorium-229 nuclear clock is as sensitive to ultralight dark matter as are state-of-the-art optical clocks, even though its stability is currently orders of magnitude lower than that of atomic clocks. These findings demonstrate the future potential of the nuclear clock for probing fundamental physics.

Several technical aspects of the nuclear

clocks should now be worked on to improve performance. For example, the lasers used to probe the nuclei have low power (about 200 nanowatts at most). Increasing the power would enable the nuclei to be interrogated more quickly, which would make the clocks more robust against short-term drift of the oscillator frequency.

Another attractive avenue of research would be to find a way to manipulate the nuclear states using coherent control – a technique based on quantum mechanics that enables dynamic processes to be controlled by light. Not only would this make the clocks more robust, it would also enable the nuclei to be used for quantum information processing. Finally, an accurate estimate of what the nuclear transition frequency would be in the complete absence of external electric fields – including the field induced by the ions present

in calcium fluoride crystals – would open the way towards nuclear clocks becoming the new time standard.

Akio Kawasaki is at the National Metrology Institute of Japan, National Institute of Advanced Industrial Science and Technology, Tsukuba, Ibaraki 305-8563, Japan.
e-mail: akio.kawasaki@aist.go.jp

1. Toscani De Col, L. *et al.* *Nature* **658**, 897–902 (2026).
2. Huang, B. *et al.* *Nature* **658**, 903–908 (2026).
3. Aepli, A., Kim, K., Warfield, W., Safronova, M. S. & Ye, J. *Phys. Rev. Lett.* **133**, 023401 (2024).
4. Zhang, B. *et al.* *Phys. Rev. Lett.* **136**, 053202 (2026).
5. Peik, E. & Tamm, C. *Europhys. Lett.* **61**, 181 (2003).
6. von der Wense, L. *et al.* *Nature* **533**, 47–51 (2016).
7. Zhang, C. *et al.* *Nature* **633**, 63–70 (2024).
8. Ooi, T. *et al.* *Nature* **650**, 72–78 (2026).
9. Arvanitaki, A., Huang, J. & Van Tilburg, K. *Phys. Rev. D* **91**, 015015 (2015).

The author declares no competing interests.
This article was published online on 7 October 2026.

Neurodegeneration

Harmful tau spreads like prion proteins

Kayleigh Mason-Chalmers & Jose A. Rodriguez

Structural studies support the idea that fibrillized tau protein, a feature of neurodegenerative disorders, acts as a prion-like template to induce misfolding in healthy tau. **See p.972**

Prions are shape-shifting, self-assembling proteins that transform from a physiologically harmless version into one that is misfolded, prone to aggregation and infectious. What makes prions and prion-like proteins capable of causing disease is their ability to act as templates, forcing normal functional proteins into a toxic form. Several proteins that are associated with neurodegenerative diseases, such as amyloid- β , α -synuclein and tau, are thought to spread in a prion-like manner^{1,2}, but experimental support for some aspects of this hypothesis and its atomic basis have been lacking. In their paper on page 972, Lövestam *et al.*³ show that injecting human tau fibrils into mice causes mouse tau to take on the form of the injected structures, bolstering the theory that tau propagates similarly to prions.

The build-up of tau aggregates in neurons is linked to disorders such as Alzheimer's disease and corticobasal degeneration (CBD)^{4,5}. Electron microscopy has revealed that disease-associated tau aggregates adopt a structure known as an amyloid fold^{5,6}. They appear as rope-like fibrils built from many flattened tau monomers stacked tightly together in layers that are separated by about four atom widths⁷, with each layer repeating regularly along the

length of the fibril. The core of the fibril is ordered, but the flexible ends of each monomer create a disordered 'fuzzy' coat on the outside of the fibril⁸. The part of the monomer that forms the core is made of two polypeptide-chain motifs, β -strands and loops. These components are arranged into a collapsed, winding fold, which enables segments of amino-acid residues to interact tightly and exclude water⁷. This structure makes amyloid fibrils distinct from nearly all other protein folds⁷.

Remarkably, the structure of the tau fibril core might be disease-specific^{5,9} – for instance, the fold of an Alzheimer's-disease-associated fibril is distinct from that of one found in CBD. This has prompted the idea that the fibril structures might propagate a particular pathology through a prion-like mechanism, and that they could be used to distinguish between diseases.

Recreating those disease-specific folds in the laboratory has been a substantial challenge, however, because the structures of tau fibrils isolated from people with neurodegenerative conditions are difficult to recapitulate and propagate *in vitro* or in cell culture^{10–13}. Furthermore, although researchers have developed mouse models that replicate aspects of tau pathology⁴, these mice generally do not




Figure 1 | Propagation of human disease-related tau strains in mice. **a**, Misfolded tau proteins form fibrils that are associated with neurodegenerative diseases such as Alzheimer's disease (AD) and corticobasal dementia (CBD). Lövestam *et al.*³ show that human disease-specific strains of tau can induce the misfolding of normal tau in mice. **b**, The authors injected tau fibrils extracted from the brains of people with Alzheimer's

disease or CBD into the brains of genetically normal mice. After nine months, the core structures of the tau fibrils extracted from the mice matched those of the disease-specific human tau they were exposed to with atomic accuracy. The results suggest that human tau acted as a template for propagation in a mechanism that resembles the spread of self-replicating proteins called prions.

produce tau fibrils that fold in the same way as those involved in human disease^{13–15}.

During prion-like 'templating', the specific fold of a tau fibril acts as the structural pattern that guides the folding of other monomers, enabling them to adopt the same core structure and propagate as a distinct shape. This shape is a key characteristic of a prion 'strain'. Lövestam and colleagues' experiments bring the structures of two tau strains into focus. The authors injected fibril-containing brain extracts from people with Alzheimer's disease or CBD into the brains of healthy mice. One week after injection, there was no remaining trace of human tau fibrils. However, after a month, mice injected with the CBD extracts showed signs of abnormal clumps (known as inclusions) of tau. After three months, the mice injected with Alzheimer's-disease extracts also showed inclusions. In all cases, the inclusions contained mouse tau that was phosphorylated (modified with phosphate groups), which is a disease-associated form of the protein¹⁶.

After nine months, Lövestam and colleagues extracted the fibrils from mice and used a technique called cryogenic electron microscopy⁶ to visualize the atomic details of the fibril cores. After carefully comparing them with the structures of human Alzheimer's-disease and CBD fibrils, they found that the folds of the mouse and human tau fibril cores matched, down to about the size of an atom.

These experiments demonstrate that mice expressing tau normally can faithfully template disease-related human tau fibril folds (Fig. 1). The findings raise questions about

what the requirements are for the disease-related structures to be mimicked. Genetically normal mice seem to have the right factors and environment for faithful structural templating. Furthermore, despite the physiological and genetic differences between mice and humans, the fact that human disease structures can be propagated in mice could make this a useful experimental model for tau-related neurodegenerative diseases.

It is intriguing that both the Alzheimer's-disease and CBD tau strains propagated in mice, but at different rates, with diverse regional spreads in the brain and in distinct cell populations. Those differences underscore

"The authors' experiments bring the structures of two tau strains into focus."

the possible roles of varying cell types in the templating process. Perhaps each strain found its path of least resistance or greatest viability through the complex tissue of the mouse brain – although some of that variation might depend on where in the brain the human tau fibrils were injected.

More work is needed to understand the prion-like nature of tau fibrils. For example, it is not clear whether tau fibrils that are assembled in a test tube, which match the structures of disease folds, are also able to template *in vivo*. Unlike test-tube fibrils, the human-derived

fibrils used by Lövestam and colleagues were not fully purified away from other biological components, such as protein or polynucleotide complexes. Given the inability of other amyloid fibrils to template identical structures in mice¹⁷, the environment of the fibrils could have a role in their prion-like propagation.

Similarly, the role of post-translational modifications in the observed templating is not yet known. It is likely that the extracted Alzheimer's-disease and CBD tau fibrils had chemical modifications, especially the notable phosphorylation that is typical of disease-associated tau^{16,18}. It is also unknown whether the ordered core of injected fibrils is sufficient for faithful templating in this model, or whether their fuzzy outer coat plays a part.

Curiously, the mice in these experiments did not show behavioural symptoms of disease during the months that tau fibril templating took place. This could be due to many factors, including a species barrier that prevents the templated mouse tau fibrils from causing illness in mice¹⁹. However, given that passing prion strains from one mouse to the next, known as serial passage, can increase their lethality²⁰, prion-like tau strains could similarly become more lethal after adapting to their mouse hosts. If so, the authors' mouse model might prove uniquely valuable for testing modulators of tau aggregation that lead to therapies.

Kayleigh Mason-Chalmers and **Jose A. Rodriguez** are in the Department of Chemistry and Biochemistry, the UCLA-DOE Institute

for Genomics and Proteomics, and the STROBE NSF Science and Technology Center, University of California, Los Angeles (UCLA), Los Angeles, California 90095, USA. e-mail: jrodriguez@mbi.ucla.edu

1. Prusiner, S. B. *Science* **336**, 1511–1513 (2012).
2. Weingarten, M. D., Lockwood, A. H., Hwo, S. Y. & Kirschner, M. W. *Proc. Natl Acad. Sci. USA* **72**, 1858–1862 (1975).
3. Lövestam, S. et al. *Nature* **658**, 972–976 (2026).
4. Basheer, N. et al. *Acta Neuropathol. Commun.* **12**, 52 (2024).
5. Shi, Y. et al. *Nature* **598**, 359–363 (2021).
6. Fitzpatrick, A. W. P. et al. *Nature* **547**, 185–190 (2017).
7. Sawaya, M. R., Hughes, M. P., Rodriguez, J. A., Riek, R. & Eisenberg, D. S. *Cell* **184**, 4857–4873 (2021).
8. Wegmann, S., Medalsy, I. D., Mandelkow, E. & Müller, D. J.

9. Scheres, S. H. W. *Structure* **34**, 1061–1071 (2026).
10. Tarutani, A. et al. *FEBS Open Bio* **13**, 1394–1404 (2023).
11. Duan, P. et al. *J. Biol. Chem.* **300**, 107730 (2024).
12. Wang, J. et al. *Structure* **32**, 2251–2258 (2024).
13. Guo, J. L. et al. *J. Exp. Med.* **213**, 2635–2654 (2016).
14. Kano, M. et al. *Acta Neuropathol. Commun.* <https://doi.org/10.1186/s40478-026-02407-7> (2026).
15. Schweighauser, M. et al. *Acta Neuropathol. Commun.* **11**, 160 (2023).
16. Rawat, P. et al. *Int. J. Mol. Sci.* **23**, 12841 (2022).
17. Burger, D. et al. *Nature* **648**, 409–417 (2025).
18. Hanger, D. P., Betts, J. C., Loviny, T. L., Blackstock, W. P. & Anderton, B. H. *J. Neurochem.* **71**, 2465–2476 (1998).
19. Watts, J. C. & Prusiner, S. B. *J. Biol. Chem.* **289**, 19841–19849 (2014).
20. Ghaemmaghami, S. et al. *J. Mol. Biol.* **413**, 527–542 (2011).

The authors declare no competing interests.
This article was published online on 30 September 2026.

Machine learning

How LLMs can learn to find the ‘i’s in ‘intelligence’

Zhao Zhang & Yingfei Xiong

Most LLMs cannot reliably evaluate text on the level of individual letters. A technique called byteification retrofits existing models to enable it. **See p.915**

Strawberry contains the letter ‘r’ three times, but when asked, many large language models (LLMs) answer that the letter appears twice. This happens because most LLMs encode words as ‘tokens’ that represent sequences of letters. LLMs that operate in this way can achieve excellent performance, but they cannot access the individual characters in each word, which are encoded as binary sequences called bytes. On page 915, Minixhofer *et al.*¹ now report an approach called byteification that retrofits token-based LLMs to operate at the byte level. The authors show that byteified models can achieve competitive performance while retaining the ability to read individual characters.

When computers store text, each character is encoded by one or more bytes (a binary sequence of eight ones and zeros). Reading this sequence reveals each character one by one, but LLMs that process text as bytes have suboptimal performance, because reading these long numerical sequences is inefficient.

Instead, most LLMs break text down into a sequence of common subwords, known as tokens². LLMs process text by tokenizing it and then predict the next token in the sequence. For example, ‘strawberry’ could be represented as three tokens: ‘str’, ‘aw’ and ‘berry’. Each token would then be represented by a single number that identifies it in the model’s vocabulary. As a result, the sequence representing the word is shortened from ten bytes

to three tokens.

Models trained on token sequences have achieved better performance than those using byte-level approaches. Subword tokenization has therefore become the dominant approach for modern LLMs. The drawback is that once

‘berry’, for example, is represented by a single token, the model cannot access the individual letters in the subword³.

Researchers have tried to retain access to individual bytes without forcing the LLM to process each one separately. One such approach is called the byte latent transformer (BLT)⁴. In subword models, the possible tokens are predefined, but BLT dynamically divides the byte sequence into ‘patches’ by reading the bytes one by one and deciding whether to place a boundary to mark the end of one patch and the beginning of the next. In theory, it is possible to divide ‘strawberry’ into ten characters or three tokens, as in a subword model, based on the preceding text describing the task. Using this representation, the LLM predicts the next patch and then decodes the patches back into bytes.

However, a BLT model cannot inherit what a subword model has already learnt in a straightforward manner, because the patches it forms do not line up with the subword model’s vocabulary of tokens. In practice, BLT, as well as most other byte-level models, must be trained from scratch. This requirement is a serious obstacle, because a competitive LLM is an enormous investment, requiring substantial computational power, time and data, as well as the development of powerful algorithms.

Minixhofer and colleagues investigated whether an existing subword model could instead be retrofitted to work with bytes. To achieve this, the authors developed a technique called byteification. This uses a trained subword model but processes the model’s inputs and outputs using extra algorithmic




Figure 1 | Byteification for large language models. Most large language models (LLMs) process text as subwords, which means they often struggle to access information on the level of individual characters, which are generally encoded as bytes. Minixhofer *et al.*¹ report a technique called byteification, in which a conventional subword model can process text at the byte level. First, a local encoder (bottom) builds a mathematical representation of the bytes that represent a piece of text, and a boundary predictor divides these into ‘patches’. A conventional subword model then predicts the next patch and a local decoder (top) converts the patches back into bytes.