The HealthTech Herald

Your trusted source for the latest breakthroughs at the intersection of health and technology.

EDITION 5 · NOVEMBER 2026



IN THIS EDITION

The Heart Valve That Grows With a Child

Inside the First Expandable Pulmonary Valve

Teaching the Immune System to Grab Cancer

Tarlatamab and the Toughest Lung Cancer

When the App Becomes Part of the Medicine

The FDA's New Rules for Prescription Drug Software

Planning the Stent Before the Patient Arrives

How a CT Scan Matched Ultrasound Inside the Heart

An AI That Talks Back About Seizures

A First Test of Patients and Machines Managing Epilepsy Together

Dr Sukh Rashpal Sandhu | Subscribe free on LinkedIn

EDITION 5 | NOVEMBER 2026

Contents

The HealthTech Herald: November 2026 3

The Heart Valve That Grows With a Child: Inside the First Expandable Pulmonary Valve... 4

Teaching the Immune System to Grab Cancer: Tarlatamab and the Toughest Lung Cancer 11

When the App Becomes Part of the Medicine: The FDA's New Rules for Prescription Drug Software..... 19

Planning the Stent Before the Patient Arrives: How a CT Scan Matched Ultrasound Inside the Heart..... 27

An AI That Talks Back About Seizures: A First Test of Patients and Machines Managing Epilepsy Together 35

Follow us on LinkedIn..... 43

About The HealthTech Herald..... 45

INTRODUCTION

The HealthTech Herald: November 2026

By Dr Sukh Rashpal Sandhu

Edition 5 of The HealthTech Herald: 5 pieces for November 2026.

Your trusted source for the latest breakthroughs at the intersection of health and technology.

This edition brings together 5 new features written for this edition. New features are marked "Written for this edition".

In this edition

- The Heart Valve That Grows With a Child: Inside the First Expandable Pulmonary Valve
- Teaching the Immune System to Grab Cancer: Tarlatamab and the Toughest Lung Cancer
- When the App Becomes Part of the Medicine: The FDA's New Rules for Prescription Drug Software
- Planning the Stent Before the Patient Arrives: How a CT Scan Matched Ultrasound Inside the Heart
- An AI That Talks Back About Seizures: A First Test of Patients and Machines Managing Epilepsy Together

Thank you to every reader who subscribes, shares and responds.

Warm regards,
Dr Sukh Rashpal Sandhu
Author, The HealthTech Herald

THE HEALTHTECH HERALD

The Heart Valve That Grows With a Child: Inside the First Expandable Pulmonary Valve

By Dr Sukh Rashpal Sandhu | Written for this edition

On 1 October 2026 the US Food and Drug Administration approved the Autus Size-Adjustable Valve, the first heart valve designed to be widened after surgery as a child grows. It is also the first US-approved valve with synthetic polymer leaflets instead of animal tissue. For families facing a childhood of repeat open-heart operations, it offers a new possibility, with some honest caveats.

The problem every paediatric heart surgeon knows

Adult heart valve replacement has a straightforward logic. A surgeon removes a failing valve and implants a new one sized to fit. The patient's heart will not get any bigger, so the main question is how long the replacement valve will last.

Children break that logic. A two-year-old's heart is a fraction of the size it will be at eighteen. Any valve implanted in a toddler will eventually be too small, no matter how well it works. And the materials used in most replacement valves, tissue taken from cows or pigs and chemically treated, tend to wear out faster in young bodies. The FDA notes that animal tissue tends to harden and break down faster in children, often requiring earlier repeat surgery.

The result is a familiar and difficult pattern. A child born with a serious valve defect may have a first operation as a baby, another as a toddler, another in primary school, and more through adolescence. Each operation means opening the chest, stopping the heart, using a heart-lung bypass machine, and weeks of recovery. Each one also adds scar tissue, which makes the next operation harder.

The pulmonary valve is particularly affected. This valve sits between the right ventricle, the pumping chamber that sends blood to the lungs, and the pulmonary artery. A range of congenital heart conditions involve a narrowed, blocked or missing pulmonary valve.

Who is affected

Congenital heart defects are the most common type of birth defect. The FDA cites a figure of about one in every 100 births in the United States. According to the US Centers for Disease Control and Prevention, pulmonary valve atresia (where the valve has not formed an opening) and pulmonary valve stenosis (where it is narrowed) together affect around 4,200 US babies each year.

Many of these children undergo early surgery or catheter procedures to open up the path from the heart to the lungs. These repairs save lives, but they often leave the valve unable to close properly. Blood then leaks backwards into the right ventricle with each heartbeat, a condition called pulmonary regurgitation. Over years, that extra load can stretch and weaken the right ventricle. Eventually a new valve is needed, and because the child is still growing, that valve may itself need to be replaced as the child outgrows it.

What makes the Autus valve different

The Autus Size-Adjustable Valve was developed by Autus Valve Technologies and is being brought to market by Edwards Lifesciences, a major maker of structural heart devices. It has two features that set it apart.

1. It can be expanded after surgery

The valve is implanted by a surgeon at a diameter of about 13 millimetres, a size suitable for a toddler or preschool-aged child. Later, as the child grows, cardiologists can widen it to about 22 millimetres, close to the size of an adult pulmonary valve.

The widening does not require another open-heart operation. Instead, it is done using a balloon catheter, a thin tube threaded through a blood vessel, usually from the groin, up to the heart. The balloon is inflated inside the valve to stretch its frame to a larger size. This is a type of procedure that paediatric cardiologists already perform for other purposes, and it typically involves a much shorter recovery than surgery.

Edwards describes this as allowing the valve to be expanded "through a minimally invasive transcatheter procedure, potentially reducing the need for repeat open-heart surgeries."

2. Its leaflets are synthetic, not animal tissue

Every heart valve has leaflets, the thin flaps that open to let blood through and close to stop it flowing back. In most biological replacement valves, those leaflets are made from treated animal tissue. In the Autus valve, they are made from a polymer. TCTMD, the cardiology news service, reports the material as expanded polytetrafluoroethylene (ePTFE), a synthetic material with a long history in medical implants.

According to the FDA, this is the first time the agency has approved a heart valve using polymeric leaflet material for any indication. That is significant in itself. Researchers have tried for decades to develop durable synthetic leaflets that avoid the calcification and wear that affect animal tissue, especially in young patients. The polymer also has a practical advantage for an expandable design: it can be engineered to tolerate being stretched.

The evidence behind the approval

The approval was based on a clinical study, registered as NCT05006404, which enrolled 62 children at 12 sites in the United States. The FDA and Edwards reported the following results.

Procedural safety. All implant procedures were successful. There were no deaths, no blood clots in the valve requiring treatment, no strokes, and no other clot-related complications that caused symptoms.

Six-month performance. Among the first 60 patients, blood flow through the right side of the heart was within acceptable limits at six months, and leakage through the valve was no more than mild.

Early complications. Edwards reported that all patients were free from device-related complications through 30 days.

One-year results. According to Edwards, among patients who had completed one-year follow-up, all showed consistent outcomes: a pressure difference across the valve (mean right ventricular outflow tract gradient) of 40 mmHg or less, less than moderate leakage, and no repeat procedures on the valve. The FDA said one-year follow-up suggests the synthetic leaflets are holding up well over time.

Expansion in practice. Two children whose valves began to perform less well as they outgrew them had their valves successfully expanded using the catheter method, without surgery.

Issues identified. Three patients had a fracture in the valve's frame, and two had reduced movement of one leaflet. None of these findings caused symptoms. The FDA has flagged them as risks that longer-term studies need to understand.

Measure	Reported result
Children enrolled	62, at 12 US sites
Implant success	All procedures successful
Deaths, strokes, symptomatic clots	None reported
Six-month valve performance (first 60 patients)	Acceptable flow, no more than mild leakage
Valves expanded by catheter	2, both successful
Frame fractures	3, without symptoms
Reduced leaflet motion	2, without symptoms
Planned follow-up	10 years

Why an expandable valve is an engineering puzzle

It is worth pausing on why nobody had done this before. Making a valve bigger sounds simple. In practice it is one of the harder problems in device design.

A heart valve works only if its leaflets meet cleanly when they close. Engineers call this coaptation. If the leaflets are too short for the opening, a gap remains in the middle and blood leaks backwards. If they are too long, they can fold, flutter or wear unevenly. A conventional valve is designed with leaflets matched to one fixed opening size.

An expandable valve has to perform well at a range of sizes: when it is first implanted in a small child, and again after each widening. That means the frame must stretch in a controlled, predictable way, without cracking or distorting, and the leaflets must still meet properly across that range. It also means the materials must tolerate being stretched while continuing to open and close tens of millions of times a year, since a child's heart beats faster than an adult's.

This helps explain why the frame fractures and the reduced leaflet motion seen in a few patients matter so much to the FDA's ongoing scrutiny, even though none caused symptoms. They are exactly the kinds of issues engineers would expect to watch closely in a design of this kind.

Surgery compared with a catheter expansion

For families, the practical difference between an open-heart operation and a catheter procedure can be enormous. The comparison below is general, and every child's situation is different.

	Open-heart valve replacement	Catheter balloon expansion
Access	Chest opened through the breastbone	Thin tube through a blood vessel
Heart-lung bypass	Usually required	Not required
Typical recovery	Days in intensive care and a longer hospital stay	Often a short hospital stay
Effect on future surgery	Adds scar tissue in the chest	Avoids adding chest scar tissue

None of this means catheter procedures are risk-free. They carry their own risks, including bleeding, blood vessel injury and rhythm disturbances, and they are performed by specialist teams. But for a child who might otherwise face several sternotomies over the course of childhood, replacing even one of those operations with a catheter procedure is a meaningful change.

Questions families may want to ask

For parents of a child with congenital pulmonary valve disease, news of a new device naturally raises questions. These are general prompts for a conversation with the child's cardiology team, not medical advice:

1. Is my child likely to need a pulmonary valve replacement, and roughly when?
2. What valve options are available to us now, and what are the trade-offs between them?
3. Is an expandable valve relevant to my child's anatomy and age, and is it available here?
4. If it is not yet available, are there clinical trials or international programmes we should know about?
5. How will my child's heart be monitored over the coming years, whatever valve is chosen?

The honest limits

This is a first approval, and the evidence base reflects that. Several points deserve emphasis.

The numbers are small. Sixty-two children is a reasonable size for a paediatric device study, where patients are relatively few and every child enrolled is precious. But it is a small group from which to draw firm conclusions about rarer complications.

Experience with expansion is very limited. The central promise of the device, that it can be widened repeatedly as a child grows, has so far been tested in only two children in the pivotal study. The FDA is explicit that experience with expansion is limited and that longer-term studies are underway to show that expansion can keep the valve functioning and lead to fewer surgeries.

Long-term durability is not yet known. The follow-up reported so far runs to a year. Children who receive this valve at two or three years old will need it to work, with expansions, for a decade or more. Edwards says patients in the pivotal study will be followed for ten years.

Frame fractures need watching. Three frame fractures without symptoms is not alarming in itself, since fractures in some types of stented valves can be harmless. But in a valve designed to be stretched multiple times, how the frame behaves over years and repeated expansions is a key question.

It will not eliminate all surgery. Even in the best case, a 22-millimetre maximum means that some children will eventually outgrow the valve. The aim is fewer trips to the operating room, not none. FDA device chief Dr Michelle Tarver framed the promise exactly that way, saying the valve "gives doctors the option of far fewer trips to the operating room."

How the FDA got it to market

The regulatory story is instructive for anyone interested in how paediatric devices reach patients.

Children have long been called "therapeutic orphans" in medical technology. The market for paediatric devices is small, the engineering is harder because children grow, and clinical trials are more complex to run ethically. As a result, many devices used in children were originally designed for adults and used off-label or adapted.

The Autus valve came through several FDA pathways designed to address this:

- **Breakthrough Device designation.** This status is given to devices that may provide more effective treatment for life-threatening or irreversibly debilitating conditions. It provides closer interaction with FDA reviewers during development.
- **Total Product Life Cycle Advisory Program (TAP).** This is a voluntary FDA programme intended to speed up the development of safe and effective devices by providing early, frequent and strategic advice. The Autus valve is the first paediatric approval to come out of TAP.
- **Premarket approval (PMA).** This is the FDA's most stringent review pathway, generally used for high-risk, Class III devices such as heart valves.

The use of TAP is notable. It suggests that intensive early engagement between regulators and developers can help bring complex paediatric devices to market, a model that may be useful for other conditions where children have been underserved.

What clinicians are saying

Dr Emile Bacha, surgeon-in-chief at NewYork-Presbyterian/Columbia University Irving Medical Center and an investigator in the trial, said in Edwards' announcement that children with congenital pulmonary valve disease "have long needed prosthetic valves that can account for growth over time." He added that the valve gives congenital heart teams "a new way to think about long-term care planning." Edwards discloses that Dr Bacha has received reimbursement from the company for travel and lodging related to educational activities.

That phrase, long-term care planning, captures the real shift. For a child born with a pulmonary valve defect, care is not a single event but a lifelong journey of procedures, monitoring and decisions. A valve that can grow changes how that journey might be mapped. Instead of planning the next surgery, a team might plan the next catheter expansion.

Why it matters beyond one device

Growth-accommodating implants

The Autus valve is part of a broader push in medicine toward implants that adapt to the body rather than requiring the body to adapt to them. In paediatrics this has included devices to correct spinal curvature that can be lengthened without surgery and research into tissue-engineered heart valves that might, in theory, grow with the patient. The Autus approach is more mechanical: a frame designed to be stretched. But it reflects the same idea, that children's devices should be designed for children's bodies, not shrunk-down versions of adult equipment.

Synthetic valve leaflets

The approval of polymer leaflets may also have implications for adult cardiology. Animal-tissue valves are used widely in older adults because they do not require lifelong blood-thinning medication, but they wear out. Mechanical valves last longer but require lifelong anticoagulants. A durable synthetic leaflet that combines the advantages of both has long been a goal. This approval is for children's pulmonary valves, which experience lower pressures than valves on the left side of the heart. It does not mean polymer valves are ready for every use. But it is a regulatory milestone that researchers in the field will note.

Fewer operations, fewer risks

Every open-heart operation carries risks: bleeding, infection, complications from bypass, effects on the developing brain from anaesthesia and surgery, and the emotional toll on children and families. If expandable valves reduce the number of such operations a child needs, the benefits may extend well beyond the heart itself. That is the hope. The evidence to prove it will take years to collect.

What this means for Australian families

The approval is from the US FDA and applies to the United States. Australia's Therapeutic Goods Administration makes its own independent decisions on medical devices, and there

is no indication yet of when, or whether, the Autus valve will be submitted or approved here.

Australian children with congenital heart disease are treated in specialist paediatric cardiac centres, and these teams follow international developments closely. Families whose children may face future pulmonary valve replacement can reasonably ask their cardiology team about the device and whether it may become relevant to their child's care. But there is no reason to delay or change any planned treatment on the strength of this news. As with any new device, the first priority is long-term evidence that it is safe and effective.

A quiet revolution in a small field

Paediatric cardiac surgery is a field of small numbers and large stakes. Compared with adult heart care, relatively few children in any one country need the kind of valve replacement the Autus device addresses. That makes it a hard market for innovation, and progress often comes slowly.

The Autus valve does not solve every problem. It is based on a small study, its central feature has been used in only a handful of patients, and it will need a decade of follow-up to prove its worth. But it represents a genuinely new idea brought through a rigorous regulatory process: a valve designed from the start to grow with a child.

If the longer-term results confirm the early findings, the main change for families may be simple: fewer operations during childhood.

Sources

- US Food and Drug Administration, "FDA Approves First Heart Valve Designed to Grow with Children", press announcement, 1 October 2026. https://www.fda.gov/news-events/press-announcements/fda-approves-first-heart-valve-designed-grow-children
- Edwards Lifesciences, "Edwards Lifesciences Receives FDA Approval for AUTUS Valve, the First Surgical Pulmonary Valve for Pediatric Patients", news release, 1 October 2026. https://s27.q4cdn.com/788244549/files/doc_news/Edwards-Lifesciences-Receives-FDA-Approval-for-AUTUS-Valve-the-First-Surgical-Pulmonary-Valve-for-Pediatric-Patients-2026.pdf
- TCTMD, "FDA Approves First Expandable Option for Pediatric Patients With Pulmonary Valve Disease", October 2026. https://www.tctmd.com/news/fda-approves-first-expandable-option-pediatric-patients-pulmonary-valve-disease
- Cardiovascular Business, "FDA approves first heart valve that grows with children", October 2026. https://cardiovascularbusiness.com/topics/congenital-heart-disease/fda-approves-first-heart-valve-grows-children
- AJMC, "FDA Approves First Heart Valve to Keep Pace With Children's Growth", 5 October 2026. https://www.ajmc.com/view/fda-approves-first-heart-valve-to-keep-pace-with-children-s-growth
- ClinicalTrials.gov, study record NCT05006404. https://clinicaltrials.gov/study/NCT05006404

THE HEALTHTECH HERALD

Teaching the Immune System to Grab Cancer: Tarlatamab and the Toughest Lung Cancer

By Dr Sukh Rashpal Sandhu | Written for this edition

At the European Society for Medical Oncology Congress in Madrid in October 2026, researchers reported that adding tarlatamab, a "bispecific" immune therapy, to standard maintenance treatment extended median survival in extensive-stage small cell lung cancer from 12.7 months to 21.1 months. The results were published at the same time in the New England Journal of Medicine. Here is how the technology works, what the trial showed, and where the caution lies.

A cancer that has resisted progress

Small cell lung cancer is one of the most aggressive cancers known. It accounts for a minority of lung cancers, but it grows quickly, spreads early and is strongly linked to smoking. Most people are diagnosed once the cancer has already spread beyond one area of the chest, a stage doctors call extensive-stage disease. According to ecancer's report of the trial, about 80% of people with small cell lung cancer are diagnosed at this extensive stage.

For decades, treatment barely changed. Patients typically received platinum-based chemotherapy combined with a drug called etoposide. The cancer often shrank dramatically at first, which gave patients and families real hope. But it almost always came back, often within months, and when it did, it was usually resistant to further treatment.

The last significant step forward came in the late 2010s and early 2020s, when trials showed that adding an immune checkpoint inhibitor, a drug that releases a brake on the immune system, to first-line chemotherapy modestly improved survival. Durvalumab, sold as Imfinzi, is one of those drugs. Today the standard approach for many patients is chemotherapy plus a checkpoint inhibitor, followed by maintenance with the checkpoint inhibitor alone.

Even so, as ecancer reports, median survival with first-line therapy has remained around 12 to 13 months. Dr Jacob Sands of the Dana-Farber Cancer Institute, who presented the new results, said that given the aggressive nature of the disease, "many patients quickly relapse on current therapy."

The technology: a molecule with two hands

Tarlatamab, sold as Imdelltra by Amgen, belongs to a class of medicines called bispecific T-cell engagers. The name describes what it does.

Most antibody medicines have one target. They recognise one specific molecule and stick to it. A bispecific antibody is engineered to recognise two different targets at once, like a molecule with two hands.

In tarlatamab's case:

- **One hand grabs CD3**, a protein found on the surface of T cells, the immune system's specialist killer cells.
- **The other hand grabs DLL3**, a protein found on the surface of most small cell lung cancer cells but rarely on healthy adult cells.

By holding on to both at once, tarlatamab physically pulls a T cell up against a cancer cell. Once they are bridged together, the T cell is activated and attacks the cancer cell it has been brought into contact with. In effect, the drug bypasses the usual process by which the immune system has to find and recognise a tumour on its own, which many cancers are good at evading.

The choice of DLL3 is important. A good target for this kind of therapy should be common on cancer cells and rare on healthy tissue, so the immune attack is focused where it is needed. DLL3 fits that profile in small cell lung cancer, which is why it has become a major focus of drug development in this disease.

Bispecific T-cell engagers are not entirely new. The first, blinatumomab, was approved for a form of leukaemia in 2014, and several others have since been approved for blood cancers. Tarlatamab is notable because it extends the approach to a solid tumour, which has proved harder, and because it targets one of the most difficult solid tumours of all.

Tarlatamab was already approved by the US FDA in 2024 for people with extensive-stage small cell lung cancer whose disease had progressed after chemotherapy. The new trial asked a different question: what happens if it is used earlier, as part of first-line maintenance, before the cancer has a chance to come back?

The trial: DeLLphi-305

DeLLphi-305 was a phase 3, open-label, multicentre, randomised trial. It enrolled 563 adults with extensive-stage small cell lung cancer whose disease had not progressed after initial treatment with durvalumab plus platinum chemotherapy and etoposide.

Participants were randomly assigned, in equal numbers, to one of two maintenance treatments:

- **Combination:** tarlatamab 10 mg every two weeks plus durvalumab 1,500 mg every four weeks (282 patients)
- **Standard:** durvalumab 1,500 mg every four weeks alone (281 patients).

The main question, the primary endpoint, was overall survival: how long people lived. Amgen announced on 8 September 2026 that the trial had met this endpoint at a planned

interim analysis, without releasing the numbers. Full results were presented by Dr Sands at the ESMO Congress 2026 Presidential Symposium in Madrid on 24 October 2026, and published simultaneously in the New England Journal of Medicine.

The results

The median survival, hazard ratio and grade 3 or higher adverse event figures below match the NEJM paper by Maurice Pérol and colleagues (data cut-off 7 July 2026). The remaining figures are as reported by OncLive and ecancer from the congress presentation.

Survival

Measure	Tarlatamab plus durvalumab	Durvalumab alone
Median overall survival	21.1 months	12.7 months
Alive at 6 months	92%	75%
Alive at 12 months	79%	53%
Hazard ratio for death	0.42 (95% CI 0.32 to 0.54)	

A hazard ratio of 0.42 means that, at any given point during the follow-up period, people receiving the combination had a 58% lower risk of dying than those on durvalumab alone. Put more simply, median survival extended by more than eight months, and at the one-year mark, nearly four in five people on the combination were alive compared with about half on standard treatment.

The benefit was reported as consistent across subgroups, including patients whose cancer had spread to the brain or liver, two groups that typically do worse.

Cancer control

Measure	Tarlatamab plus durvalumab	Durvalumab alone
Median progression-free survival	7.4 months	1.9 months
Free of progression at 12 months	38%	10%
Objective response rate	26.4%	4.1%
Median duration of response	12.7 months	5.6 months

Progression-free survival measures how long the cancer was kept from growing or spreading. The difference here is large: a median of 1.9 months on durvalumab alone, which illustrates how quickly this cancer usually returns, compared with 7.4 months on the combination.

The side effects

Activating the immune system this powerfully has costs. The most characteristic side effect of T-cell engagers is cytokine release syndrome, or CRS. When large numbers of T cells are activated at once, they release chemical messengers called cytokines, which can cause fever, low blood pressure, low oxygen levels and, in severe cases, organ problems.

In DeLLphi-305, according to OncLive:

- **CRS** occurred in 49.6% of patients on the combination, compared with 1.4% on durvalumab alone. Most cases were mild (grade 1 or 2). Grade 2 CRS occurred in 7.4%, one patient (0.4%) had grade 3 CRS, and there were no grade 4 or fatal cases.
- **Management of CRS** included corticosteroids in 21.6% of combination patients, tocilizumab (an anti-inflammatory antibody) in 3.9%, and low-flow oxygen in 1.4%.
- **ICANS**, a neurological side effect associated with immune effector therapies that can cause confusion or difficulty speaking, occurred in 1.1% of combination patients, with none graded 3 or higher.
- **Serious side effects overall.** Grade 3 or higher treatment-emergent adverse events occurred in 46.8% of combination patients and 40.6% of those on durvalumab alone. Grade 3 or higher treatment-related adverse events occurred in 20.9% compared with 10.8%.
- **Stopping treatment.** Only 3.2% of patients stopped tarlatamab because of treatment-related side effects.
- **Treatment-related deaths.** None occurred with the combination, and two (0.7%) occurred with durvalumab alone.

Other common side effects included low white blood cell counts, decreased appetite, fatigue and weight loss.

It is worth noting that people in the combination group stayed on treatment much longer: a median of 46.9 weeks compared with 8.1 weeks for durvalumab alone. Some of the difference in side effect rates simply reflects longer exposure, because people who are doing well stay on treatment and have more time to experience side effects.

Reading the numbers: a quick guide

Cancer trial results come with their own vocabulary. For readers who do not spend their days with oncology journals, here is what the main terms mean.

- **Median overall survival.** The point at which half of the patients in a group have died and half are still alive. It does not mean every patient lives exactly that long. Some live much longer, some shorter.
- **Hazard ratio.** A comparison of the risk of an event, here death, between two groups over the follow-up period. A hazard ratio of 1 means no difference. A hazard ratio of 0.42 means the risk in the treatment group was 42% of the risk in the control group, that is, 58% lower.
- **95% confidence interval (CI).** The range within which the true effect is likely to lie. For overall survival here, the range of 0.32 to 0.54 sits well below 1, which is why the result is considered statistically robust.
- **Progression-free survival.** How long patients live without their cancer growing or spreading.
- **Objective response rate.** The proportion of patients whose tumours shrank by a defined amount on scans.
- **Open-label.** Both patients and doctors knew which treatment was being given.
- **Interim analysis.** A pre-planned look at the data before the trial is fully complete, usually to see whether a treatment is clearly working or clearly not.

Why cytokine release syndrome needs a prepared team

Cytokine release syndrome is not unique to tarlatamab. It is a known effect of several immune therapies that activate T cells strongly, including CAR-T cell therapies and other bispecific antibodies used in blood cancers. Over the past decade, cancer centres have built up considerable experience in recognising and treating it early.

In practical terms, that means:

1. **Monitoring after early doses.** CRS is most likely to occur after the first doses, when T cells are first activated. Patients are watched closely during this period.
2. **Early recognition.** Fever is often the first sign. Teams use standard grading systems to judge severity and decide on treatment.
3. **Prompt treatment.** Mild cases may need only supportive care such as fluids and fever medication. More significant cases are treated with corticosteroids or tocilizumab, which blocks a key inflammatory messenger.
4. **Patient education.** Patients and carers are taught what symptoms to watch for at home and when to seek help urgently.

The DeLLphi-305 data, in which most CRS was mild and there were no grade 4 or fatal cases, suggest that with this kind of preparation the side effect can be managed in most patients. But it underlines why the treatment needs to be delivered in centres with the right expertise.

Questions patients may want to ask

For people living with extensive-stage small cell lung cancer, and their families, these are general prompts for a discussion with an oncologist, not medical advice:

1. Am I in the group of patients this trial studied, that is, did my cancer respond to first-line chemotherapy and immunotherapy?
2. Is tarlatamab available to me now, in any setting, and on what terms?
3. Are there clinical trials of tarlatamab or similar DLL3-targeting treatments I could be eligible for?
4. What would treatment involve in practice, including travel, monitoring and time in hospital?
5. What side effects should I watch for, and who do I contact if they occur?

What the experts said

Dr Sands told the congress, as reported by ecancer, that "this trial is positive in every way." According to OncLive, he concluded that the outcomes support the combination "as a new standard of care first-line maintenance therapy" in extensive-stage disease.

Congress commentary was broadly enthusiastic, although some observers urged caution about comparing its results directly with other trials, since different studies enrol different patients and use different control arms.

The caveats

This is a strong result, but like every trial it has limits that should be understood.

It is an interim analysis. The survival results come from a planned interim look at the data, with median follow-up of about 16 months. Longer follow-up will show whether the survival gap holds, widens or narrows, and how many people remain alive at two and three years.

It was open-label. Patients and doctors knew which treatment was being given. For a hard endpoint like overall survival, this matters less, because death is not open to interpretation. But it can influence how side effects are reported and how other decisions are made.

It only included people who responded to first-line treatment. Participants had to have completed initial chemotherapy and immunotherapy without their cancer progressing. People whose cancer grew during first-line treatment, or who were too unwell to continue, were not included. The results apply to a selected group, not everyone diagnosed with the disease.

Delivery is demanding. Tarlatamab is given by infusion every two weeks, and because of the risk of CRS, the early doses require careful monitoring. This needs hospitals and clinics with experience in managing immune therapy side effects. Not every regional centre is equipped to do this, which may affect access.

It is not yet approved for this use. Tarlatamab is approved in the United States for later-line treatment, but not for first-line maintenance. The DeLLphi-305 results may support future applications to regulators, but each country will make its own decision. Australian patients should talk to their oncologist about what is available locally, including through clinical trials.

Cost will be a question. Novel immune therapies are expensive, and adding a second drug to maintenance over many months will be costly for health systems. In Australia, public subsidy decisions are made through the Pharmaceutical Benefits Advisory Committee, which weighs clinical benefit against cost.

Why this matters

For patients with small cell lung cancer

This disease has long been among the hardest to treat in oncology. Survival gains in first-line treatment over the past decade have been measured in weeks or a few months. An extension of median survival of more than eight months, with nearly 80% of patients alive at one year, is a substantial change in a field that has seen few of them. Ecancer reports that DeLLphi-305 is the first phase 3 study of a bispecific T-cell engager to show an overall survival benefit in first-line maintenance for this disease.

For cancer immunotherapy

The first wave of cancer immunotherapy, checkpoint inhibitors, works by releasing the brakes on the immune system. That helps when the immune system already recognises

the cancer. Many cancers, including small cell lung cancer, are good at hiding, so releasing the brakes is not enough.

T-cell engagers take a more direct approach. They do not wait for the immune system to find the cancer; they physically bring the immune cell to it. DeLLphi-305 suggests that combining the two approaches, releasing the brakes with durvalumab and building a bridge with tarlatamab, may work better than either alone in this setting.

For the design of future treatments

The DeLLphi-305 result is likely to accelerate interest in other DLL3-targeting therapies, including other bispecific molecules, antibody-drug conjugates and cell therapies. It also adds to growing evidence that solid tumours, long considered a harder target for T-cell engagers than blood cancers, can respond to this approach when the right target is chosen.

Not every new approach in this disease has succeeded. In late September 2026, according to OncLive and a Merck news release, Daiichi Sankyo and Merck voluntarily withdrew an accelerated approval application for a different drug, ifinatamab deruxtecan, in previously treated small cell lung cancer, after the FDA indicated that the phase 2 data did not meet the requirements for accelerated approval. The contrast is a reminder that large, randomised phase 3 trials with survival endpoints remain the most convincing evidence a new therapy can provide.

A word on prevention

No discussion of small cell lung cancer is complete without noting that the overwhelming majority of cases are linked to smoking. New treatments matter enormously for people already diagnosed. But the single most effective way to reduce deaths from this cancer remains preventing people from starting to smoke and helping those who do smoke to stop.

Australia has also begun a National Lung Cancer Screening Program, launched in July 2025, for eligible people at high risk because of their smoking history. Screening is aimed at finding lung cancers early, when more treatment options exist. Small cell lung cancer is often found late even with screening because of how fast it grows, but screening and smoking cessation together remain the foundation of reducing lung cancer deaths.

The verdict

DeLLphi-305 is one of the most significant results in small cell lung cancer in many years. Adding tarlatamab to durvalumab maintenance extended median survival from 12.7 to 21.1 months, reduced the risk of death by 58%, and kept the cancer in check for much longer, in a large, randomised phase 3 trial published in a leading journal.

It is not a cure. Most patients will still see their cancer return. The treatment brings real side effects, particularly cytokine release syndrome, and requires experienced teams to deliver safely. The results are interim, apply to people who responded to first-line treatment, and await regulatory and funding decisions in each country.

But for a disease in which progress has been painfully slow, the trial shows that an engineered antibody that bridges a T cell to a tumour cell can make a measurable difference to how long people live.

Sources

- ecancer, "ESMO 2026: Immunotherapy combo extends survival for patients with aggressive small cell lung cancer", 24 October 2026. https://ecancer.org/en/news/29090-esmo-2026-immunotherapy-combo-extends-survival-for-patients-with-aggressive-small-cell-lung-cancer
- OncLive, "DeLLphi-305 Data Support Frontline Maintenance Tarlatamab Plus Durvalumab as a New SOC in ES-SCLC", 24 October 2026. https://www.onclive.com/view/dellphi-305-data-support-frontline-maintenance-tarlatamab-plus-durvalumab-as-a-new-soc-in-es-sclc
- Pérol M, Yoshida T, Ahn MJ, et al., and Sands J. "Tarlatamab-Durvalumab Maintenance for Extensive-Stage Small-Cell Lung Cancer." New England Journal of Medicine, published online 24 October 2026. https://doi.org/10.1056/NEJMoa2609482
- OncLive, "Ifinatamab Deruxtecan BLA Voluntarily Withdrawn in Previously Treated Extensive-Stage SCLC", 25 September 2026. https://www.onclive.com/view/ifinatamab-deruxtecan-bla-voluntarily-withdrawn-in-previously-treated-extensive-stage-sclc
- CancerNetwork, "Tarlatamab Maintenance Combo Extends Survival in Extensive-Stage SCLC", 24 October 2026. https://www.cancernetwork.com/view/tarlatamab-maintenance-combo-extends-survival-in-extensive-stage-sclc
- CURE, "Phase 3 Trial Shows Adding Imdelltra to First-Line Treatment Helps Patients With SCLC Live Longer", 5 October 2026. https://www.curetoday.com/view/phase-3-trial-shows-adding-imdelltra-to-first-line-treatment-helps-patients-with-sclc-live-longer

THE HEALTHTECH HERALD

When the App Becomes Part of the Medicine: The FDA's New Rules for Prescription Drug Software

By Dr Sukh Rashpal Sandhu | Written for this edition

On 20 October 2026 the US Food and Drug Administration finalised guidance on how it regulates apps and software that drug companies provide alongside prescription medicines. Its central idea is simple and far-reaching: whatever a sponsor's app shows a patient or doctor about a medicine counts as part of that medicine's labelling. Here is what that means for the companion apps now common in modern treatment, and why it matters well beyond the United States.

The rise of the companion app

Not long ago, a prescription medicine came with two things: the medicine itself and a folded paper leaflet. Today, a growing number come with a third: software.

These apps take many forms. Some connect to a device that delivers the medicine, such as a smart inhaler that records each puff or an injection pen that logs each dose. Others stand alone on a phone: dose reminders, symptom diaries, coaching programmes, or tools that explain how to take the medicine and what side effects to watch for. Some send data to a patient's doctor. Some use that data to suggest actions.

For patients, these tools can be genuinely helpful. Remembering doses is hard, especially for people managing several medicines or long-term conditions. Recording symptoms can help doctors make better decisions. Clear, interactive instructions can make complex devices easier to use correctly.

But companion apps also raise new questions. If an app tells a patient something about their medicine, who checks that it is accurate? If it highlights benefits more than risks, is that information or advertising? If a feature is added after the medicine is approved, does a regulator need to know? Until recently, the rules in the United States were not entirely clear.

The guidance: what it says

The FDA's final guidance is titled *Regulatory Considerations for Prescription Drug Use-Related Software*. It finalises a draft first issued on 19 September 2023, and the FDA announced its availability in the Federal Register on 21 October 2026.

The guidance introduces a few key terms.

Prescription drug use-related software. In the FDA's words, this generally includes software that is disseminated by or on behalf of a drug sponsor and that produces an output that "supplements, explains, or is otherwise textually related to" one or more of the sponsor's drug products. Crucially, this applies whether or not the software is also regulated as a medical device.

Software function. A distinct purpose of the software. One app might have several functions, for example, logging doses, sending reminders and displaying educational content.

End-user output. Any material or content the software presents to the end user, which the FDA defines as a patient, caregiver or health care practitioner.

The core position is this: the FDA considers end-user output to be a type of prescription drug labelling.

That may sound like regulatory jargon, but it is a significant statement. In US law, labelling is not just the printed label on a bottle. It includes the official prescribing information, patient leaflets and promotional materials. Labelling must be accurate, balanced and not misleading, and it is subject to FDA oversight. By classing app content as labelling, the FDA is making clear that what an app says about a medicine is held to the same broad standards as what a company prints about it.

The guidance covers prescription drugs and drug-led combination products, those where a drug and a device are combined and the drug plays the primary role, assigned to the FDA's drug centre (CDER) or biologics centre (CBER).

Two types of software functions

The guidance draws a central distinction between two kinds of software functions.

Device-connected functions

These are software functions that rely on data from a device that is part of the drug product. The law firm Orrick, in its analysis of the guidance, gives the example of an app paired with a connected inhaler, where the inhaler is the device component of the product. Other examples might include apps connected to injection pens or other drug delivery devices.

Standalone functions

These are software functions that do not depend on a device constituent. Orrick's examples include symptom trackers, dose reminders and coaching tools.

When software belongs in the prescribing information

One of the most practical questions the guidance answers is when a software function should be described in a medicine's prescribing information, the official document that tells doctors how to use it.

The FDA's answer, as summarised by commentators including Orrick and Akin Gump:

- Simply distributing an app for use with a medicine does not, by itself, require that app to appear in the prescribing information.
- **Device-connected functions** generally belong in the prescribing information.
- **Standalone functions** generally should not be described in the prescribing information, unless there is an additional reason, such as evidence that the function is necessary for the safe and effective use of the medicine, or that it provides a meaningful improvement in a clinical outcome.
- Content from most standalone apps is expected to be treated as promotional material.

The final guidance also introduces a change from the draft that is worth noting. For device-connected software functions where there is no evidence of a meaningful improvement in a clinical outcome, the FDA now recommends adding a statement in the prescribing information saying so. In other words, if an app connected to an inhaler has not been shown to improve health outcomes, doctors reading the official prescribing information should be told that.

This is an important safeguard for clinical decision-making. It means the presence of an app should not be mistaken for evidence that the app improves care.

An illustration

To make the distinction concrete, consider a hypothetical example. It is not drawn from any real product, and it simplifies the guidance considerably.

Imagine a company that makes an asthma inhaler with a small sensor attached. The sensor records each time the inhaler is used and sends the data to a phone app. The app shows the patient a chart of their usage and reminds them when a dose is due. The same app also includes a separate section with general tips on avoiding asthma triggers.

Under the framework described in the guidance:

- The usage chart depends on data from the inhaler's sensor, which is part of the product. That is a **device-connected** function, and it would generally be described in the prescribing information. If the company has no evidence that using the chart improves asthma outcomes, the FDA's recommendation is that the prescribing information say so.
- The trigger tips do not depend on the sensor. That is a **standalone** function. Unless the company can show it is necessary for safe and effective use, or that it meaningfully improves outcomes, it would generally not appear in the prescribing information, and its content would likely be treated as promotional material, subject to the rules for promotion.
- In both cases, what the app shows the patient is labelling, and the company is responsible for making sure it is accurate and balanced.

The point of the example is that one app can contain functions in different categories, each with different consequences. That is why the guidance asks companies to work through their software function by function rather than app by app.

What changed from the 2023 draft

According to the FDA's Federal Register notice, the main changes from the draft to the final guidance are:

1. clarifying the considerations for standalone software functions
2. introducing a decision tree to guide sponsors through labelling determinations
3. providing illustrative examples of prescribing information language for both device-connected and standalone software
4. recommending a statement in the prescribing information for device-connected software functions for which there is no meaningful improvement in a clinical outcome.

The FDA also said it made editorial changes to improve clarity.

The decision tree is designed to help companies work out, function by function, which category a piece of software falls into and what labelling consequences follow. Commentators note that it addresses drug labelling only; separate device requirements may still apply if the software is also a medical device.

How this fits with medical device rules

Software that diagnoses, treats or monitors a condition may itself be regulated as a medical device, under the FDA's device centre (CDRH). The new guidance does not replace those rules. Instead, the FDA says the recommendations are intended to align with initiatives across all its product centres, including digital health initiatives at CDRH, and to support efficient, coordinated review when drug-related software is also reviewed as a device.

In practice, a single app could be subject to both frameworks: device rules covering whether the software is safe and effective at what it does, and drug labelling rules covering what it says about the medicine. The FDA's aim is to make those two frameworks work together rather than in conflict.

What it means for patients

For patients, the guidance is mostly invisible, but its effects could be significant.

More accountability for app content. Because app content about a medicine is now clearly treated as labelling, companies have a stronger obligation to make sure it is accurate, balanced and consistent with what regulators have approved. Information about risks should not be buried while benefits are highlighted.

Clearer signals about evidence. The recommendation that prescribing information state when a device-connected app has not been shown to improve outcomes gives doctors clearer information. That, in turn, helps them advise patients honestly about whether an app is likely to help.

Recognition that apps change. Software is updated far more often than printed leaflets. Commentators recommend that companies put change-control processes in place so that new features, content changes and algorithm updates are assessed for their

labelling implications before release. That should reduce the risk of an app drifting away from approved information over time.

Privacy remains a separate question. The guidance is about labelling, not data protection. Many companion apps collect sensitive health information. Orrick notes that companies should align privacy notices, consent and data-sharing practices with how the app is described in labelling and promotion, and comply with privacy and consumer protection laws. But the FDA guidance itself does not set privacy rules. Patients should still read privacy terms carefully and ask who will see their data.

What it means for clinicians

Doctors, pharmacists and nurses increasingly encounter companion apps in practice, sometimes recommended by drug companies, sometimes found by patients themselves. The guidance offers clinicians a useful mental model.

- If an app is described in the official prescribing information, there is usually a specific reason: it is connected to a device component, or there is evidence it matters for safe use or outcomes.
- If an app is not described in the prescribing information, it is most likely a standalone support tool, and its content is generally treated as promotional.
- If the prescribing information says a device-connected app has not been shown to improve clinical outcomes, that is useful context when discussing it with patients.

None of this means standalone apps are unhelpful. A good dose reminder may be valuable for a patient who struggles to remember medicines. But clinicians can distinguish between tools backed by evidence of clinical benefit and tools that are useful conveniences.

What it means for the industry

For pharmaceutical companies, the guidance creates both obligations and clarity. Law firm commentary has suggested several practical steps for sponsors, including:

1. **Take inventory.** Identify every app, web tool and software function linked to a prescription product, including those built or run by vendors and agencies, and map what each one shows users.
2. **Classify each function.** Decide whether it is device-connected or standalone, and document the reasons and supporting evidence.
3. **Review content properly.** Route app content through the same medical, legal and regulatory review processes used for other labelling and promotion.
4. **Control changes.** Assess updates for labelling and submission impact before they go live.
5. **Plan early.** For medicines in development, consider during clinical trial planning whether software functions will need evidence to support inclusion in the prescribing information.
6. **Update contracts.** Make sure agreements with technology partners cover who approves content and how regulatory obligations are shared, since content distributed on a sponsor's behalf is covered.

The guidance also notes that, like all FDA guidance, it represents the agency's current thinking and is not legally binding. Companies can use an alternative approach if it satisfies the relevant laws and regulations. In practice, however, guidance strongly shapes how companies behave, because it signals how the FDA will assess their materials.

Why this matters beyond the United States

The guidance applies only in the United States. But medicines and their companion apps are developed and marketed globally, and decisions made by the FDA often influence how companies design products everywhere.

In Australia, the Therapeutic Goods Administration regulates software that meets the definition of a medical device, under a framework that was updated in 2021 to address software-based devices more directly. Advertising of prescription medicines to the public is also tightly restricted in Australia, which differs from the US approach. The FDA guidance does not change Australian law, and Australian patients and clinicians should not assume that US labelling principles apply here.

But the questions the guidance tackles are universal:

- When a medicine comes with an app, who is responsible for what the app says?
- How should regulators distinguish between helpful information and promotion?
- How should evidence about an app's benefit be communicated to doctors?
- How should regulators keep up with software that changes constantly?

Every health system with digitally enabled medicines will need answers to these questions. The FDA's framework offers one model.

The bigger picture: medicines are becoming systems

Step back, and the guidance reflects a deeper shift in what a "medicine" is.

For most of the twentieth century, a medicine was a chemical in a tablet or vial. Its effect depended on the molecule and how it was taken. Today, many treatments are better understood as systems: a molecule, a delivery device, a digital tool that guides its use, and a stream of data that may flow back to clinicians. In some cases, the software may influence whether a medicine works as well in real life as it did in clinical trials, simply by helping people take it correctly.

Regulators built their frameworks around the molecule. They are now adapting them to cover the whole system. This guidance is one step in that adaptation. It does not answer every question. It does not, for example, set out how artificial intelligence features in companion apps should be assessed, and it does not deal with privacy. But it establishes a clear principle that will shape what comes next: if a company's software says something about its medicine, that statement carries the same responsibility as anything else the company says about it.

What to watch

Several developments will be worth watching as the guidance takes effect.

How companies respond. Existing companion apps may be reviewed, revised or, in some cases, withdrawn as companies apply the new framework.

How the FDA enforces it. Guidance becomes meaningful through how regulators apply it in reviews and enforcement. Early examples will show how strictly the FDA reads the distinction between device-connected and standalone functions.

Evidence of benefit. The guidance creates an incentive for companies to generate evidence that their apps improve outcomes, because that evidence can support inclusion in the prescribing information. Over time, that could lead to more rigorous testing of digital tools.

International alignment. Other regulators, including in Europe and Australia, may look at the FDA's approach as they develop their own policies on digitally enabled medicines.

A note for readers who use these apps

If you use an app that came with a prescription medicine, a few general points are worth keeping in mind. This is not medical advice, and your doctor or pharmacist is the best source of guidance for your situation.

- An app is a support tool, not a replacement for advice from your health professionals.
- If an app tells you something that conflicts with what your doctor or pharmacist said, ask them about it.
- Check what data the app collects, where it goes, and whether you can control it.
- Do not change your dose or stop a medicine based on an app alone.
- Report problems with an app, including confusing information, to your health professional and, where relevant, to the company or regulator.

The verdict

The FDA's final guidance on prescription drug use-related software is not a dramatic breakthrough of the kind that fills headlines. There is no new drug, device or cure. But it settles an important question for the digital age of medicine: the words and numbers that a drug company's app shows a patient are part of the medicine's labelling, and they carry the same responsibilities.

For patients, that should mean more accurate and balanced information. For clinicians, it means clearer signals about whether an app is backed by evidence. For companies, it means more work, but also more clarity. And for regulators around the world, it offers one model for managing the growing number of medicines that come with software.

Sources

- US Food and Drug Administration, "Regulatory Considerations for Prescription Drug Use-Related Software; Guidance for Industry; Availability", Federal Register notice, Docket No. FDA-2023-D-2482, scheduled for publication 21 October 2026. https://public-inspection.federalregister.gov/2026-21516.pdf

- US Food and Drug Administration, guidance documents for drugs. https://www.fda.gov/drugs/guidance-compliance-regulatory-information/guidances-drugs
- Orrick, "There's an App for That, and Now FDA Has a Framework for It", 23 October 2026. https://www.orrick.com/en/Insights/2026/10/Theres-an-App-for-That-and-Now-FDA-Has-a-Framework-for-It
- Akin Gump, "FDA Finalizes Guidance on Prescription Drug Use-Related Software", 27 October 2026. https://www.akingump.com/en/insights/alerts/fda-finalizes-guidance-on-prescription-drug-use-related-software
- Pharmaceutical Commerce, "FDA Finalizes Guidance on Prescription Drug Companion Apps", 20 October 2026. https://www.pharmaceuticalcommerce.com/view/fda-finalizes-guidance-on-prescription-drug-companion-apps
- NSF, "FDA final guidance on prescription drug apps and software", October 2026. https://www.nsf.org/life-science-regulatory-news/fda-final-guidance-on-prescription-drug-apps-and-software

THE HEALTHTECH HERALD

Planning the Stent Before the Patient Arrives: How a CT Scan Matched Ultrasound Inside the Heart

By Dr Sukh Rashpal Sandhu | Written for this edition

At the Transcatheter Cardiovascular Therapeutics conference in San Diego on 1 November 2026, researchers reported the P4 trial, published the same day in the New England Journal of Medicine. In 1,105 patients needing a coronary stent, planning the procedure from a CT scan done beforehand worked as well as the current gold standard, ultrasound from inside the artery, and procedures were shorter. Here is what that means, why it matters, and why the experts are saying "not yet" as well as "game-changing".

The procedure millions of people have

When the arteries that feed the heart muscle become narrowed by fatty plaque, the result can be chest pain, breathlessness, or in the worst case a heart attack. One of the most common treatments is percutaneous coronary intervention, usually shortened to PCI and known to most people as "having a stent".

In a PCI, a cardiologist threads a thin tube, called a catheter, through an artery in the wrist or groin up to the heart. A tiny balloon is inflated at the narrowed spot to push the plaque aside, and a small metal mesh tube, the stent, is left in place to hold the artery open. It is done in a specialised room called the cardiac catheterisation laboratory, or cath lab, usually under local anaesthetic and light sedation.

PCI is performed on a vast scale around the world, and it has become safer and more effective over decades. But a stent only works well if it is the right size, placed in the right spot, and fully expanded against the artery wall. A stent that is too short can miss part of the plaque. One that is too narrow or not properly expanded is more likely to fail later, either by narrowing again or by forming a clot.

How cardiologists currently "see" inside the artery

The traditional way to guide PCI is angiography: X-ray dye is injected into the coronary arteries and moving X-ray images show where blood is flowing. Angiography shows the shape of the channel, a bit like a silhouette. It does not show the plaque itself, how hard or soft it is, or the true size of the artery wall.

To fill that gap, many cardiologists use intravascular imaging. The most common form is intravascular ultrasound, or IVUS. A miniature ultrasound probe on the end of a catheter

is passed into the artery during the procedure, and it creates cross-sectional images from the inside. IVUS shows the vessel's true size, the amount and type of plaque, and, after the stent goes in, whether it has expanded properly.

Several large trials and meta-analyses have found that IVUS-guided PCI leads to better outcomes than angiography alone, with fewer repeat procedures and heart attacks. As a result, intravascular imaging has become a recommended standard for many PCI procedures in international guidelines. But IVUS has drawbacks. It adds time, cost and equipment. It requires an extra catheter inside the coronary artery. And all the planning happens in real time, in the cath lab, with the patient on the table.

The idea: do the planning first

Coronary CT angiography, or CCTA, is a non-invasive scan. The patient lies in a CT scanner, receives an injection of contrast dye into an arm vein, and the scanner builds a detailed three-dimensional image of the heart and its arteries. CCTA is increasingly used as a first test for people with chest pain, partly because it can rule out significant disease quickly.

Over the past decade, software has made CCTA far more powerful. Computer models can now estimate not only how narrow an artery is, but whether the narrowing actually restricts blood flow. This is called CT-derived fractional flow reserve, or FFR-CT. Other tools can characterise plaque: how much there is, whether it is calcified, and where it begins and ends.

The idea behind P4 was simple. If many patients already have a detailed CCTA before they come to the cath lab, why not use it to plan the whole procedure in advance? Cardiologists could decide beforehand how to prepare the lesion, what length and diameter of stent to use, where exactly the stent should land, and which X-ray angles would show the lesion best. Then, in the cath lab, they would carry out a plan rather than make it up as they go.

Lead investigator Dr Carlos Collet of the Cardiovascular Research Foundation in New York put the problem bluntly at TCT: "We don't plan the [PCI] procedures."

The trial: P4

The full name of P4 is the Precise Procedural PCI Plan trial. It was an international, randomised, open-label, non-inferiority trial with blinded adjudication of outcomes. Results were published in the New England Journal of Medicine by Collet and colleagues. The trial was investigator-initiated and funded by GE HealthCare and HeartFlow, two companies that make CT and CT-analysis technology.

According to the published abstract and reports from CardiologyNowNews.org and SOLACI:

- **Patients:** 1,105 adults at 23 international sites, with a median age of 65; about 78% were men.
- **Condition:** about 98% had stable coronary artery disease (what cardiologists call chronic coronary syndromes), and all had a CCTA before the procedure, most of them done as part of routine care before enrolling.

- **Eligibility:** at least one lesion narrowing the artery by more than 70%, with a CT-derived FFR of 0.80 or less, indicating the narrowing was likely to be restricting blood flow, and a clinical reason for the procedure.

Patients were randomly assigned to one of two strategies:

1. **CCTA-guided PCI (556 patients).** The procedure was planned from the CT scan, covering the artery's anatomy, plaque features, calcium, the amount of heart muscle at risk and blood flow physiology. The plan specified lesion preparation, stent length and diameter, landing zones and X-ray angles. IVUS was not allowed except as a bailout. Techniques to enhance the stent image on X-ray were mandatory.
2. **IVUS-guided PCI (549 patients).** IVUS was mandatory before and after stenting, in line with current best practice.

The main outcome was target-vessel failure at 12 months: a combination of death from heart causes, heart attack related to the treated artery, or a repeat procedure on the treated artery because of reduced blood flow.

The trial was designed as a non-inferiority study. The question was not whether CCTA planning was better than IVUS, but whether it was not meaningfully worse.

The results

Measure	CCTA-guided	IVUS-guided
Patients	556	549
Target-vessel failure at 12 months	4.1% (23 patients)	4.6% (25 patients)
Clinical procedural complications	6 patients	9 patients
Median procedure time	46 minutes	63 minutes
Contrast volume (as reported by SOLACI)	150 mL	170 mL

The difference in target-vessel failure was minus 0.42 percentage points, with a 95% confidence interval of minus 2.82 to 1.99. The prespecified non-inferiority margin was 4.5 percentage points, so the result met the test for non-inferiority (P below 0.001). It did not show superiority (P equal to 0.73), and there was no difference in the individual components of the main outcome.

In plain terms: in this group of patients, planning the stent from a CT scan done beforehand produced one-year results that were statistically no worse than the current gold standard of ultrasound from inside the artery. And procedures were, on average, about 17 minutes shorter.

The authors concluded that key PCI planning decisions "can be shifted from the cardiac catheterization laboratory to preprocedural CCTA."

Why shorter procedures matter

A 17-minute saving might not sound dramatic. But in a busy cath lab, time matters for several reasons.

- **For patients**, less time on the table means less time sedated, less time with catheters in the arteries and, potentially, less X-ray exposure during the procedure.
- **For hospitals**, the cath lab is an expensive resource with long waiting lists in many health systems. Collet noted that the time saving "means you can do a couple more [PCIs] a day."
- **For staff**, shorter, more predictable procedures can mean less fatigue and better scheduling.

There is also a less obvious benefit. When the plan is made in advance, the patient can be involved in discussions about the procedure before they are lying in the cath lab. The authors suggested this offers an opportunity to enhance shared decision-making.

The caveats, and why experts said "not yet"

The response at TCT was enthusiastic but measured. Several experts quoted by CardiologyNowNews.org praised the trial and in the same breath set out why it should not change practice overnight.

The patients were relatively straightforward. The median SYNTAX score, a measure of how complex a person's coronary disease is, was about 8 to 9, which is low. Collet himself agreed that anatomical complexity was relatively low. Dr Kevin Bainey of the Mazankowski Alberta Heart Institute in Canada, commenting to CardiologyNowNews.org, called P4 a "great first step" but said IVUS "really shines in complex PCI and left-main PCI," referring to the most challenging procedures.

Events were fewer than expected. The one-year rate of target-vessel failure was lower than the investigators anticipated, which reduces a trial's ability to detect differences. Commentators also noted that, with event rates around 4%, the non-inferiority margin of 4.5 percentage points can be questioned.

The CT scan's own contrast and radiation were not counted. CCTA requires contrast dye and X-ray exposure. The investigators did not account for the dye and radiation from the planning CT scan in the comparison. The modest savings during the procedure therefore do not tell the whole story.

After the stent goes in, CT cannot help. Collet acknowledged that "once you place the stent, the CT can't help you anymore because it was acquired before [the PCI]." IVUS can check stent expansion in real time; a pre-procedure CT cannot. The trial addressed this partly by mandating stent-enhancement imaging techniques in the CCTA arm.

The software is still maturing. Discussants noted that the trial initially needed six separate software programs to do the planning, although sponsors have since developed integrated software. Dr Margaret McEntegart of NewYork-Presbyterian/Columbia University Irving Medical Center, a TCT panellist, stressed that the information needs to be presented through an intuitive, simplified interface if it is to be used widely.

Costs have not been analysed. No formal cost-effectiveness analysis has yet been done, and affordability was raised as an open question.

Skills need to change. Dr Emmanouil Brilakis of the Minneapolis Heart Institute Foundation, a discussant, called P4 a "landmark" trial but said cardiologists "need to learn

how to actually use the CT, how to look at the CT images" before the approach becomes mainstream. The session moderator, Dr Ajay Kirtane of NewYork-Presbyterian/Columbia University Irving Medical Center, went further, warning that cardiologists who do not learn to read CT risk becoming "obsolete".

Dr B. Hadley Wilson, a past president of the American College of Cardiology who was not involved in the study, summed up the mood: it was "a game-changing trial," but "we're not there yet."

The technology behind the plan

It is worth understanding what sits behind a CCTA-based plan, because it is a good example of how software is changing medicine.

A CT scanner produces hundreds of cross-sectional images. On their own, these require expert interpretation. Increasingly, computer models do much of the heavy lifting:

- **Segmentation.** Software traces the outline of each coronary artery and its walls in three dimensions.
- **Plaque analysis.** Algorithms estimate the amount and type of plaque, including calcified and non-calcified components.
- **Physiology.** Computational fluid dynamics models simulate blood flow through the arteries to estimate FFR-CT, effectively predicting whether a narrowing limits blood flow without passing a pressure wire into the artery.
- **Planning.** Tools can then suggest stent lengths, diameters and landing zones and identify the X-ray angles that will show the lesion best.

Several of these tools use machine learning, and some are commercial products that have regulatory clearance in various countries. HeartFlow, a company that makes CT-based heart analysis software and one of the trial's two funders, issued a press release about P4 at the conference. Readers should note that commercial interests are involved in this field, and that the evidence for each specific tool needs to be judged on its own merits.

Two pathways, side by side

To make the difference concrete, here is a simplified picture of how the two strategies in P4 differ from a patient's point of view. Real care varies from hospital to hospital.

Step	IVUS-guided pathway	CCTA-guided pathway
Before the procedure	CT or other tests may show a narrowing; detailed planning waits for the cath lab	The existing CT scan is analysed in detail and a written procedural plan is prepared
In the cath lab	An ultrasound catheter is passed into the artery to measure it before stenting	The cardiologist follows the pre-prepared plan, without an ultrasound catheter
Checking the result	IVUS is repeated after stenting to confirm expansion	Stent-enhancement X-ray imaging is used to check the result
Typical time in P4	Median 63 minutes	Median 46 minutes

The key difference is not the stent itself but when and how the decisions about it are made. In one pathway, measurements are taken inside the artery during the procedure. In the other, they are taken from the CT scan beforehand.

The bigger trend: moving decisions out of the cath lab

P4 sits within a broader shift in cardiology. For decades, the cath lab was where both diagnosis and treatment happened. Patients with suspected coronary disease were often sent for invasive angiography to find out what was going on, and treatment decisions were made on the spot.

That model is changing. Non-invasive CT is increasingly used to diagnose coronary disease, decide who needs an invasive procedure and, now, plan the procedure itself. The aim is to reserve the invasive cath lab for treatment rather than exploration, with decisions made earlier and more deliberately.

The P4 investigators have already started a follow-up trial. Subgroup analyses suggested that the most complex patients may have gained the most from CT-based planning, and that has led to P5, a study in more complex disease that has begun enrolling patients. If CCTA planning holds up in complex cases, the argument for using it widely will be much stronger.

What this means for Australian patients

The P4 trial was international, and its results are relevant to cardiology everywhere, including Australia, where CT coronary angiography is widely available in major centres and PCI is a common procedure.

But a few points are worth stressing:

1. **Nothing changes immediately.** Practice guidelines change slowly and only after evidence accumulates. Cardiologists will want to see longer follow-up, results in complex disease, and cost analyses before making CT-based planning routine.
2. **IVUS remains an excellent tool.** P4 showed CT planning was not worse than IVUS in a selected group, not that IVUS is unnecessary.
3. **Not everyone needs or can have a CT.** Some patients cannot have contrast dye because of kidney problems or allergies, and others need urgent procedures, for example during a heart attack, where there is no time for a planning scan. P4 studied mainly stable patients.
4. **Ask questions.** If you are being considered for a stent, it is reasonable to ask your cardiologist what imaging will be used to plan and guide the procedure, and why.

This is general information, not medical advice. Decisions about heart procedures should always be made with your own cardiologist.

A quick glossary

- **PCI (percutaneous coronary intervention).** A procedure to open narrowed heart arteries using a balloon and usually a stent, done through a catheter.
- **Angiography.** X-ray imaging of blood vessels using contrast dye.

- **IVUS (intravascular ultrasound).** An ultrasound probe on a catheter that images arteries from the inside.
- **CCTA (coronary CT angiography).** A non-invasive CT scan that builds a three-dimensional image of the heart arteries.
- **FFR (fractional flow reserve).** A measure of whether a narrowing restricts blood flow. Values of 0.80 or less are generally considered significant.
- **Target-vessel failure.** A combined outcome of heart death, heart attack in the treated artery, or the need to treat the same artery again.
- **Non-inferiority trial.** A trial designed to show a new approach is not meaningfully worse than an established one, often because it has other advantages such as being simpler or cheaper.

The verdict

P4 is a carefully designed, randomised trial showing that, in relatively uncomplicated stable coronary disease, planning a stent from a CT scan done in advance gives one-year results statistically no worse than guiding it with ultrasound inside the artery, with shorter procedures. It points to a future in which the hard thinking happens before the patient reaches the cath lab, supported by software that turns a CT scan into a detailed plan.

It is not yet a reason to abandon intravascular imaging. The patients studied were not complex, event rates were low, costs have not been analysed, and cardiologists will need new skills and better software to use the approach well. The P5 trial, in more complex disease, will be the real test.

For now, P4 gives strong evidence that much of the planning for a routine stent procedure can be done from a CT scan before the patient reaches the cath lab.

Sources

- Collet C, Mizukami T, Bouisset F, et al. "Coronary CT Angiography versus Intravascular Ultrasound Guidance for PCI." New England Journal of Medicine, published 1 November 2026. https://www.nejm.org/doi/full/10.1056/NEJMoa2610771
- American College of Cardiology, "P4 and TRANSFORM II: CCTA-Guided vs. IVUS-Guided PCI; Sirolimus-Coated Balloon vs. Drug-Eluting Stents", 1 November 2026. https://www.acc.org/Latest-in-Cardiology/Articles/2026/11/01/00/11/sun-207pm-combo1-tct-2026
- CardiologyNowNews.org, "CCTA-Based PCI Planning Noninferior to IVUS Guidance: P4", November 2026. https://cardiologynownews.org/?p=139872
- SOLACI, "TCT 2026 | P4: Coronary CT Angiography vs. IVUS for PCI Guidance", 2 November 2026. https://solaci.org/en/2026/11/02/tct-2026-p4-coronary-ct-angiography-vs-ivus-for-pci-guidance/
- TCTMD, "CCTA Noninferior to IVUS for PCI Guidance: P4 Trial", 1 November 2026. https://www.tctmd.com/news/ccta-noninferior-ivus-pci-guidance-p4-trial
- HeartFlow, "TCT 2026: CT-guided PCI Matches IVUS Outcomes in Randomized P4 Trial", GlobeNewswire, 1 November 2026. https://www.globenewswire.com/news-release/2026/11/01/3392793/0/en/tct-2026-ct-guided-pci-matches-ivus-outcomes-in-

randomized-p4-trial-heartflow-pci-navigator-changes-pci-plans-in-approximately-70-of-patients.html

THE HEALTHTECH HERALD

An AI That Talks Back About Seizures: A First Test of Patients and Machines Managing Epilepsy Together

By Dr Sukh Rashpal Sandhu | Written for this edition

A study published in *Nature Medicine* on 13 October 2026 tested something new in epilepsy care: an artificial intelligence system that read patients' brainwaves in real time, alerted them when a seizure looked likely, answered their questions about their own data and learned from their feedback. It involved only 13 people in one hospital unit, and the authors are clear it is a proof of concept. But it offers an early look at a different relationship between patients and medical devices, one where the device talks with the patient.

A condition that is hard to predict

Epilepsy is one of the most common serious brain conditions. The authors of the new study cite an estimate of about 70 million people living with epilepsy worldwide, and note that more than a third continue to have seizures despite taking medication.

For people whose seizures are not controlled, one of the hardest parts of life is unpredictability. A seizure can strike at work, while driving, in the shower or in the middle of the night. Many people with epilepsy describe the uncertainty itself as a burden: not knowing whether today will be a good day, whether it is safe to go out alone, or whether a change in medication is helping.

Doctors face a related problem. Much of what they know about a patient's seizures comes from the patient's own diary. But people often do not notice, or cannot remember, their seizures, especially those that cause loss of awareness or happen during sleep. Seizure diaries are known to be incomplete.

Technology has started to help. Electroencephalography, or EEG, records the brain's electrical activity through sensors on the scalp or, in some patients, electrodes placed inside the skull. Computer algorithms can analyse EEG signals to detect seizures and, increasingly, to estimate the risk that one is coming. Some implanted devices already monitor brain activity continuously and respond automatically.

But as the study's authors point out, these devices typically do not share what they know with patients in real time. The data flows to clinicians, often reviewed weeks later. The authors describe this as a missed opportunity.

The idea: a device that converses

The research team, working at the Hospital of the University of Pennsylvania in the United States, set out to test whether an AI system could act as a kind of conversational partner in epilepsy care. The senior author is Brian Litt, a neurologist and engineer known for his work on brain devices.

Their system combined several elements:

- **Live brain monitoring.** EEG data was streamed in real time from the hospital to a secure, privacy-compliant cloud computing environment.
- **AI detectors.** A set of machine learning models analysed the brain signals to estimate the probability of a seizure, detect abnormal electrical spikes between seizures, classify sleep stages and track a pattern of brain-wide synchrony that the researchers proposed as a marker of how much anti-seizure medication was in the patient's system.
- **A smartphone app.** Patients used the app for two-way communication with the system.
- **A language model.** A commercial large language model, run in a privacy-compliant environment, handled the conversation. Its role was deliberately limited to delivering information and retrieving data, not making clinical judgements.

When a patient sent a message, the system worked out whether it was general conversation or a question about their data. If it was a data question, a separate component looked up the answer in the patient's own records and reported back. The system also sent patients alerts when it estimated a seizure was likely, notified them of detected seizures and spikes, and asked them to complete short surveys three times a day.

Crucially, the system was designed to learn from patients. When a patient confirmed or rejected an alert, that feedback could be used to fine-tune the seizure detector; in this study the fine-tuning was tested in a simulation rather than live. The researchers also simulated a clinician reviewing, once a day, ten brainwave segments the model was most unsure about, adding expert labels.

The study: 13 patients in an epilepsy monitoring unit

The study took place in a level 4 epilepsy monitoring unit, a specialised hospital ward where people with epilepsy stay for several days while their brain activity is recorded continuously. Patients are often admitted to these units to work out where their seizures start, for example when surgery is being considered. In the unit, medications are frequently reduced to allow seizures to be recorded.

The design was prospective, observational and run at a single centre. Thirteen adults took part: 11 had scalp EEG and two had electrodes placed inside the skull (intracranial EEG). Median participation was about 4.9 days. The researchers set a target of 10 to 15 patients to test feasibility, not to prove that the system improved outcomes.

Several safety measures were built in:

- Standard care was not changed. The AI system ran alongside usual monitoring.

- Patients were told the app was a research tool and were instructed to use the usual seizure alarm button for events.
- Only the validated baseline seizure detector was used for live alerts; the fine-tuned versions were tested offline.
- Every interaction between patients and the AI was reviewed daily by humans.
- Messages in both directions passed through moderation filters designed to catch inappropriate content.
- The language model provider was contractually barred from training on or retaining patients' data, and all computation stayed within privacy-compliant infrastructure.
- The system was tested by the researchers, including attempts to make it misbehave, before patients used it.

What happened

Patients used it

Over the study, 1,307 messages were exchanged. The system started 900 of them (68.9%), mostly alerts, notifications and surveys. Patients started 407 (31.1%). Of the patient-initiated messages, 88.5% were general conversation rather than data queries.

Patients responded to 67.2% of real-time seizure alerts. Survey response rates were highest in the morning (62.5%), lower in the afternoon (43.8%) and in between in the evening (54.2%).

It was fast and mostly accurate

General conversational replies took a median of about two seconds. Data queries took about nine seconds. The part of the system that decided whether a message was conversation or a data question was correct 89.7% of the time (70 of 78 cases checked). The data component answered correctly 92.1% of the time (35 of 38 cases).

Patients found it usable

Five patients completed a standard usability questionnaire, the System Usability Scale. Their average score was 83 out of 100, which falls in the range usually described as excellent. Five responses is a very small number, so this should be read as encouraging rather than conclusive.

Moderation flagged few messages

Of 407 patient messages, six (1.47%) were flagged by the moderation filter. None of the 900 AI messages was flagged. The authors report that all flags were benign.

Feedback reduced false alarms

The most interesting technical result concerns false alarms, a major problem in seizure detection. Alerts that fire when nothing is happening quickly erode trust and become a nuisance.

In an offline simulation of how the fine-tuned detector would have performed:

Learning approach	False alarms per hour, before	False alarms per hour, after	Change

Learning approach	False alarms per hour, before	False alarms per hour, after	Change
Patient feedback (using push-button alarms as labels)	0.469	0.097	about 77% fewer
Clinician review (5 patients)	0.210	0.044	about 83% fewer

Using push-button alarms as the reference, the reduction in false alarms came without losing sensitivity, the ability to catch real seizures. The clinician-review approach also improved a combined measure of accuracy, the F1 score, by about 272%.

There was an important exception. When the researchers used patients' own confirmations of seizures in the app as the reference, false alarms also fell by about 80% in the subset analysed, but those labels were sparse: patients' app confirmations covered only five of the 59 seizures that clinicians confirmed, and the authors state that sensitivity did not improve. Part of the explanation is that patients rarely reported seizures through the app, because they had been told to use the standard alarm button.

A possible medication marker

The researchers also found that their measure of brain-wide synchrony rose as estimated medication levels fell, with a modest correlation. If confirmed in larger studies, a marker like this could one day help track whether medication levels are adequate, but this finding is preliminary.

What the authors say are the limits

The authors are unusually candid about the study's limits, and they deserve to be heard.

It is tiny. Thirteen patients at one hospital is a feasibility study, not evidence that the approach improves care. Participants may also have been more interested in technology than typical patients.

The setting is artificial. In an epilepsy monitoring unit, medications are reduced and seizures are deliberately provoked. Patients are lying in bed, connected to equipment and watched by staff. Life at home is very different.

The surveys were not validated. The questionnaires were built by the researchers, so their results are treated as signs of engagement, not as measures of health outcomes.

Not all signals were used. Video, heart rhythm and other data that clinicians normally use were not included.

Alarms remain a burden. Even after fine-tuning, the authors estimate a false alarm rate of roughly two to four a day, which could still be burdensome for patients.

The detectors were not state of the art. The researchers chose lightweight models suited to real-time use, not the most accurate models available.

Wider risks need attention. The authors highlight concerns about over-reliance on AI, unclear lines of clinical responsibility, and the need for continuous evaluation as AI models change over time.

Why this matters

It would be easy to dismiss a 13-patient study. But the paper is significant less for its numbers than for the model of care it tests.

From silent devices to talking ones

Most medical devices today are silent. A pacemaker, an insulin pump or an implanted brain stimulator does its job and stores data for the clinician. The patient often knows little about what the device is recording.

The Penn team's argument is that devices should converse with patients. Imagine a person with an implanted seizure monitor being able to ask, "How has my seizure risk been this week?" or "Did I have any events last night?" and getting a clear answer. Imagine the device asking, "You seemed to have an unusual pattern at 3 am. Did you notice anything?" and using the answer to improve.

That is a different relationship from the one most patients have with medical technology. It treats the patient as an active participant whose observations make the device smarter.

Patients as teachers

The finding that patient feedback reduced false alarms is particularly intriguing. Machine learning models are usually trained on data labelled by experts, which is expensive and slow. If patients themselves can help train a detector, simply by confirming or rejecting alerts, devices could become personalised to each patient's brain in a way that is hard to achieve otherwise.

The study also shows the limits of this approach. Patient reports are imperfect, and the detector's sensitivity did not improve when patient app reports were used as the reference. Combining patient feedback with clinician review may be more reliable than either alone.

Guardrails by design

For a health technology newsletter, the most instructive part of the study may be its safety architecture. The language model was restricted to information delivery and data retrieval, not clinical reasoning. Live alerts came only from a validated detector. Patients were told to use standard alarm systems. Every interaction was reviewed by humans each day. Messages were moderated in both directions. Data was protected contractually and technically.

This is the kind of careful, layered design that responsible use of AI in medicine requires. It also shows how much work goes into making such a system safe, even in a supervised hospital setting. Moving to unsupervised use at home will require even more.

Where this fits among epilepsy technologies

Epilepsy has been a testing ground for medical technology for decades, and the Penn system builds on several existing strands.

- **Implanted stimulators.** Devices such as vagus nerve stimulators and responsive neurostimulation systems are already used in some people with drug-resistant

epilepsy. Responsive systems monitor brain activity through implanted electrodes and deliver small electrical pulses when they detect abnormal patterns. They act automatically and largely silently.
- **Wearable detectors.** Wrist-worn and other wearable devices that detect certain types of seizures, for example through movement and skin signals, have been cleared by regulators in some countries. They usually alert a carer rather than the patient.
- **Seizure diary apps.** Many smartphone apps let people log seizures, triggers and medications, but they depend entirely on what the person records.
- **Seizure forecasting research.** Researchers in several countries, including Australia, have studied long-term brain recordings to show that seizure risk often follows daily and multi-day cycles, raising the prospect of seizure "forecasts" similar to weather forecasts.

What is new in the Penn study is the combination: real-time brain signals, AI interpretation, and a two-way conversation in which the patient can ask questions and the device can learn from the answers. It is an attempt to join up pieces that have mostly existed separately.

A short glossary

- **EEG (electroencephalography).** Recording of the brain's electrical activity, usually through sensors on the scalp.
- **Intracranial EEG.** EEG recorded from electrodes placed inside the skull, closer to the brain, giving clearer signals. It requires surgery.
- **Interictal spikes.** Brief abnormal electrical discharges that occur between seizures and can indicate how irritable the brain is.
- **Sensitivity.** The proportion of real seizures a detector correctly identifies.
- **False alarm rate.** How often a detector signals a seizure when none is occurring.
- **Fine-tuning.** Adjusting an existing machine learning model using new, specific data, here from individual patients.
- **Large language model.** A type of AI system trained on large amounts of text that can hold conversations in everyday language.
- **Proof of concept.** An early study designed to show an idea can work in practice, not to prove that it improves outcomes.

The road ahead

The authors identify outpatient use, in people's everyday lives, as the true test of whether this approach adds value. They also mention a planned trial in Parkinson's disease using deep brain stimulation, suggesting the same conversational approach could extend beyond epilepsy.

Several big questions will need answers:

1. **Does it improve health?** The study measured engagement, accuracy and usability, not seizures, injuries, quality of life or medication management. Future trials will need to show real benefit.

2. **Who is responsible?** If an AI system tells a patient their seizure risk is low and they then have a seizure while driving, where does responsibility lie? Clear clinical and legal frameworks will be needed.
3. **How will it be regulated?** A system that interprets brain signals and communicates risk to patients is likely to be regulated as a medical device in most countries, including by Australia's Therapeutic Goods Administration. The authors note that regulatory pathways will likely be gradual.
4. **Will patients trust it, and should they?** Trust must be earned through evidence, and it can be lost quickly through false alarms or mistakes.
5. **Who gets access?** The system depended on high-quality EEG, secure cloud computing and continuous supervision. Making such technology available outside specialist centres, and to people without the latest smartphones or reliable internet, is a separate challenge.

What people with epilepsy should take from this

If you or someone you care for lives with epilepsy, this research is a glimpse of what may be possible, not something available today. A few general points, offered as information rather than medical advice:

- The system tested was a research tool used in a hospital unit for a few days. It is not a product you can buy or use.
- Keep using the seizure management plan and safety measures agreed with your neurologist.
- Existing seizure detection devices and apps vary widely in quality and evidence. Ask your epilepsy team which, if any, are suitable for you.
- If you are interested in research, ask your neurologist about clinical trials. Many advances in epilepsy care have come from patients willing to take part in studies.

The verdict

The Penn study is small, short and carried out in an artificial hospital setting. It does not show that a conversational AI improves the health of people with epilepsy. Its authors say as much.

What it does show is that such a system can work in practice: streaming brain signals in real time, alerting patients, answering their questions about their own data within seconds, learning from their feedback to cut false alarms, and doing all of this under careful human oversight. Patients engaged with it and, in a small sample, found it easy to use.

The study's most useful idea may be that patients do not have to be passive recipients of what their devices record. If later studies at home confirm these early results, a device might be able to tell patients what it is detecting and learn from their answers.

Sources

- Goldblum Z, Shi H, Xu Z, et al., Litt B. "Human-artificial intelligence collaboration for comanaging epilepsy in real time." Nature Medicine, published 13 October 2026. https://www.nature.com/articles/s41591-026-04680-4
- Nature Medicine, "Making AI work in healthcare", editorial, 14 October 2026. https://www.nature.com/articles/s41591-026-04726-7

STAY CONNECTED

Follow us on LinkedIn

Enjoyed this edition? Each of these newsletters publishes new pieces on LinkedIn. Subscribe, like and share to read them first, and join the conversation in the comments.



The HealthTech Herald You are reading this newsletter

The latest breakthroughs at the intersection of health and technology.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/the-healthtech-herald-7338913774701027328/



Letters of the Heart

Real stories, timeless wisdom, and soul-deep inspiration.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/letters-of-the-heart-7329273590376210432/



The Hope Dispatch

Uplifting news, inspiring stories, and positive change from around the world.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/the-hope-dispatch-7338912777081638914/



Connecting the Dots

Where compliance, risk management, digital innovation, and bottom-line growth intersect.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/connecting-the-dots-7397392757360648192/



Truth, Translated

Plain language for a complicated world: education, geopolitics, political accountability and Big Tech.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/truth-translated-7450724196525035520/

Follow the author: Dr Sukh Rashpal Sandhu on LinkedIn

BACK PAGE

About The HealthTech Herald

The HealthTech Herald is a LinkedIn newsletter written by Dr Sukh Rashpal Sandhu. Your trusted source for the latest breakthroughs at the intersection of health and technology.

Read and subscribe

New editions continue to be published on LinkedIn. Search for The HealthTech Herald on LinkedIn to subscribe.

Disclaimer

The views in this collection are the author's own. Pieces are reproduced as first published, and facts, figures and links are current at their original publication date. Nothing in this collection is legal, financial, medical, migration or regulatory advice.

Copyright

Copyright Dr Sukh Rashpal Sandhu 2026. All rights reserved. Images are reproduced as they appeared in the original posts.

The HealthTech Herald | Dr Sukh Rashpal Sandhu | sukh@sukhsandhu.com

Masthead

Author: Dr Sukh Rashpal Sandhu

Publisher: Dr Sukh Rashpal Sandhu

Edition: 5, November 2026

Contact: sukh@sukhsandhu.com

The HealthTech Herald is published on LinkedIn by Dr Sukh Rashpal Sandhu. This edition collects the pieces in the order they were first published, with their original images.