



António Vieira Santos is a sociologist, technology strategist and accessibility practitioner based in Cork, Ireland. With over 30 years in the technology sector, he brings a social sciences perspective to questions of digital inclusion, artificial intelligence governance and the future of work. Until August 2026, he was Senior Innovation Strategist at Atos, where he co-founded the company's accessibility practice. He co-hosts AXSChat, one of the longest-running podcasts on accessibility and inclusion, now in its 12th year. His team was listed on the Forbes Accessibility 200, writes a column on artificial intelligence governance for AI.Mag in Germany, and co-authored a book on accessibility published by Universidade Católica Portuguesa. He has delivered keynotes at Hannover Messe and joined international events including GITEX, CES, VivaTech, Mobile World Congress and Web Summit. He served as a Next Generation Internet Ambassador for the European Commission.

When artificial intelligence forgets who it's for

Antonio Vieira Santos

For the past 12 years, I've co-hosted AXSChat alongside Debra Ruh and Neil Milliken. It's one of the longest-running podcasts on accessibility and inclusion. We've spoken to hundreds of guests from across the world, the majority of them people who live with disability every day. A guest from India, a guest from Brazil, a guest from Kenya. People building screen readers, running accessibility consultancies, teaching at universities, fighting for their rights in court.

Those conversations taught me something no amount of corporate strategy work ever could: the distance between what the technology industry promises disabled people and what it actually delivers.

In a recent episode, we spoke to the founder of Lumin, a voice-based email assistant designed with and for blind and low-vision users. He'd co-designed the product with 50 users at the Lighthouse for the Blind in San Francisco. That's the right way to build. It's also rare.

That distance is growing. And artificial intelligence is making it worse.

I want to explain how, because the evidence is now strong enough that the industry has run out of excuses. The failures are documented. The numbers are in. And the people affected have been saying this for years.

The machines that can't see their own users

Artificial intelligence is sold as the next step in making digital life easier for everyone. Voice assistants that understand you. Chatbots that answer your questions. Tools that describe images, read documents aloud, write in simpler language. For disabled people, these could be life-changing. Some of them already are.

Be My Eyes, a free app connecting blind users with sighted volunteers, now uses artificial intelligence to describe what a phone camera sees. Microsoft's Seeing AI does something similar, no account or payment needed. Google's Project Euphonia is building speech recognition that works for people whose speech patterns differ from the norm. These are real, working tools that make real differences in people's daily lives.

The problem is what happens everywhere else.

A blind user testing commercial artificial intelligence chatbots described what happens when they send a message: the screen reader jumps back to the top of the page. Text from the underlying website bleeds into the chatbot's reply. The words "Shop By Category" and "Browse Our Deals" get read aloud in the middle of an answer about something completely unrelated. The chatbot passes every automated accessibility test. A sighted person would never notice. A screen reader user can't get through a single conversation.

In July 2026, the blind user community on AppleVis documented that ChatGPT's newest features, including Deep Research and Agent Mode, were announced to screen readers as static text. Pressing

Enter or Space to select them opened the microphone instead. The features worked fine with a mouse. They were effectively unavailable to anyone using a keyboard or screen reader.

Microsoft 365 Copilot's own accessibility statement, republished by University College London in March 2025, rated the product "partially compliant" with accessibility standards. Documented failures included keyboard navigation that performed no action, a keyboard trap where focus disappeared, and a screen reader that stayed silent when navigating between results.

An academic study at the International Web for All Conference tested how well artificial intelligence tools generate accessible user interfaces: 21,880 assessments across 6 tools and 54 designers. The result: 29 percent compliance. When designers specifically asked the artificial intelligence to make the interface accessible, compliance went down.

The bias built in

The failures go beyond broken interfaces. Artificial intelligence carries the biases of the data it was trained on. And disabled people are barely in that data. A 2024 study from the University of Washington, presented at the Association for Computing Machinery Conference on Fairness, Accountability, and Transparency, tested what happens when you add disability-related credentials to a job application. A scholarship. An award. A panel seat. These additions made up less than 7 percent of the total document. The language

model, GPT-4, ranked the applications with disability credentials lower every time. Training the model on disability-justice principles reduced the bias. It didn't remove it.

This matters because artificial intelligence screening of job applications is now common. A 2021 Harvard Business School study found roughly 75 percent of large enterprises in the United States automate applicant screening. A class action lawsuit alleges that Workday's screening tools discriminate on the basis of disability. A federal judge ruled the tools can be considered an agent of the employer.

Speech recognition still performs badly for people with atypical speech. Research using the TORGO database, which includes speakers with cerebral palsy and motor neurone disease, found word error rates between 60 and 90 percent for people with low-intelligibility speech. For people without speech differences, the same systems scored between 4 and 7 percent. The gap has narrowed slightly. It remains enormous.

Google Research used text-to-image generators to create pictures of disabled people. The same result kept appearing: a person in a wheelchair. When asked for a story about a disabled person, language models produced familiar patterns. The disabled character needs to be rescued. The disabled character is inspirational for doing ordinary things. The disabled character wants to be fixed. A blind person supposedly can't play basketball. The models had never heard of wheelchair basketball.

Researchers at Google ran 19 focus groups with 56 adults with disabilities. When the participants saw the artificial intelligence's

stories and images, they used a phrase that sums up the whole problem: "The bot has a non-disabled mind."

A separate study by Ben Hutchinson and colleagues, published at the Association for Computational Linguistics in 2020, found that text-processing models used for content moderation rated language mentioning disability as more toxic and more negative. The models had learned from the way disability appears in public text, next to words about violence, addiction, and suffering. The effect: content moderation systems can suppress disabled people's own words about their own lives.

A 2026 study by Aktion Mensch, the German disability organisation, with the polling firm Ipsos, surveyed 574 people with disabilities. The findings confirmed both sides of the coin. 73 percent of respondents already used artificial intelligence tools, 42 percent regularly. But 28 percent of those who don't use artificial intelligence said inaccessible design was the reason. The demand exists. The access doesn't.

The price tag

There's a dimension of this problem that rarely gets the attention it deserves: cost.

The most capable artificial intelligence tools now sit behind a monthly payment. ChatGPT Plus costs 20 dollars per month. Microsoft Copilot Pro costs the same. Google's artificial intelligence Pro tier costs 19.99 dollars. Premium tiers run to 200 or 250 dollars a month.

Free versions exist. They run older, less capable models, have shorter time limits, and in some cases now show advertisements. ChatGPT's voice mode, which is the primary way many disabled users interact with the tool, is free but capped at a few minutes per day. When you hit the limit, the conversation ends. For someone who uses voice as their main way of interacting with a computer, that's a functional wall, not a sample.

Specialised assistive technology is even more expensive. The JAWS screen reader, one of the most widely used, costs about 90 dollars a year for a basic home licence and up to roughly 1,295 dollars for a permanent one.

Now consider who needs these tools most.

Across the European Union in 2024, the employment gap between disabled and non-disabled people was 24 percentage points. In Ireland, where I've lived and worked for 20 years, the European Disability Forum measured the gap at 38.6 percentage points, one of the worst in Europe. The National Disability Authority's own figures put the employment rate for working-age disabled people at 49.3 percent in 2022, compared to 70.8 percent for non-disabled people.

The International Labour Organization found a global disability pay gap of 12 percent. Three quarters of that gap can't be explained by differences in education, age, or type of work. In lower-income countries, the gap reaches 26 percent.

Then there's the extra cost of having a disability. In Ireland, government-commissioned research puts this at between 10,000

and 15,000 euros per year. In the United Kingdom, the disability charity Scope found extra costs of about 1,095 pounds per month, equal to 63 percent of household income after housing costs.

A 240-euro annual subscription for an artificial intelligence assistant competes directly with heating, food, and existing equipment. The people who would benefit most from artificial intelligence as an assistive tool are the people least able to pay for it.

We are building a two-tier artificial intelligence system. The capable version costs money. The limited version is free. The people on the wrong side of that divide are the ones the technology was supposed to help.

This pattern repeats a familiar story. Every time a new technology arrives that could reduce inequality, access to its best version gets priced in a way that widens the gap. Television, the internet, smartphones. Each time, we told ourselves the market would eventually fix it. With artificial intelligence and disability, that wait is too expensive and the stakes are too high.

The 299-dollar question

No product captures this tension better than Meta's Ray-Ban smart glasses.

At 299 dollars for the basic Meta Glasses line and 379 dollars for the Ray-Ban branded version, they are the cheapest artificial intelligence glasses on the market. Envision Glasses, the dedicated product built specifically for blind and low-vision users, costs between 1,899 and

2,499 dollars. Apple's Vision Pro costs 3,699 dollars. Samsung's Galaxy XR costs 1,799 dollars. Meta's glasses deliver much of what these expensive devices do, for a fraction of the price.

Blind users are embracing them. A blind writer for ABILITY Magazine described the glasses reading paragraphs aloud as she types, identifying cat-food flavours, and correctly naming her location on a street corner in New York. The American Foundation for the Blind reviewed them and called the experience "magical" for reading menus and mail. Through the free Be My Eyes app, a blind user can say "Hey Meta, Be My Eyes" and connect hands-free to a sighted volunteer through the glasses' camera. Be My Eyes now has over 10 million volunteers supporting roughly 1 million blind and low-vision users in over 185 languages.

The glasses look like ordinary Ray-Bans. They're sold at Target, Best Buy, LensCrafters, and Boots. They don't look like assistive technology. For many users, that matters.

The catch is privacy. And it's serious.

Because the glasses look like ordinary sunglasses, bystanders often don't know they're being recorded. In October 2024, two Harvard students combined the glasses' camera feed with facial recognition software and within seconds pulled up strangers' names, phone numbers, home addresses, and relatives' names. They built it to show what was possible.

In early 2026, Swedish newspapers revealed that footage from Meta glasses, including people undressing and using the bathroom, was

being reviewed by workers at a contractor in Kenya to train Meta's artificial intelligence. A class action lawsuit followed. Meta has also been developing a facial recognition feature called "Name Tag" that would let the glasses identify people by face. After journalists found dormant code for the feature already installed on devices, and after United States senators demanded answers, Meta removed the code. The company had abandoned facial recognition across its platforms in 2021 over exactly these concerns. It appears to be bringing it back through the glasses.

Venues are responding. Pub chains, theatres, and restaurants in the United Kingdom have banned the glasses. New York state courts banned recording-capable eyewear across all 1,240 courthouses. In Germany, the digital rights organisation HateAid filed a criminal complaint seeking a sales ban under the same telecommunications law used against a children's doll with a hidden microphone. The Federal Network Agency says it is monitoring the situation.

The American Foundation for the Blind describes the glasses as "an accessibility tool only by coincidence." They were designed for social media users. The accessibility benefit is real but accidental, and the privacy risk is built into the product's design.

Can you get the accessibility benefits without the surveillance? Researchers say yes. A 2025 study showed visual assistance for blind users can run entirely on the device, without sending footage to a company's servers. A peer-reviewed study at the 2026 Conference on Human Factors in Computing Systems proposed context-sensitive protections: automatic face blurring, location-based camera restrictions, and recording indicators that work

through sound and vibration, not just a small light that only 41 percent of bystanders considered sufficient.

The Centre for Democracy and Technology argues that Meta is using accessibility as cover for a tool that creates real privacy and security risks. The same organisation insists that disabled people should not lose access to useful technology because of how others misuse it. The answer is better design: on-device processing, data deleted after use, no reuse of footage for training, and recording indicators that everyone, including blind bystanders, can detect.

Starting at 299 dollars, Meta's glasses are the most affordable artificial intelligence visual aid available. That price matters enormously when the disability employment gap is 24 percentage points and the extra cost of disability is over 10,000 euros a year. The product is a case study in why accessibility and ethics can't be separated. Getting the price right while getting the privacy wrong helps no one in the long run.

Open source and the question of access

There is another path. Open models, released with their code and technical details available to anyone, have closed much of the quality gap with commercial products. Meta's Llama, Mistral, Alibaba's Qwen, Google's Gemma. These can be downloaded and used without paying a subscription.

Mozilla's Common Voice project has collected over 33,000 hours of speech data in 133 languages, all freely available. Google's Project Euphonia has open tools for building speech recognition that works

with non-standard speech. OpenAI's Whisper, an open speech recognition model, performs well on standard speech and, when adjusted for specific voices, shows real improvement for atypical speech too. The free screen reader NVDA, built by the charity NV Access, is now the most-used screen reader according to the WebAIM 2024 survey. It costs nothing and has community add-ons that connect it to artificial intelligence for image descriptions.

The catch is practical. Running an open model requires technical knowledge and hardware that most non-technical users don't have. The barrier moves from money to skill. For many disabled users, that's a different wall.

The opportunity is real. Open models can be adjusted for specific needs: dysarthric speech, sign language recognition, simplified text. What's missing is money and commitment. Public and charitable funding could turn these open tools into finished products that non-technical disabled users can install and use without help. The raw capability exists. The last step, making it usable by the people who need it, keeps getting skipped.

The rules that were supposed to protect people

Europe has built the strongest legal structure in the world for artificial intelligence and accessibility. The question is whether it's strong enough.

The European Union Artificial Intelligence Act requires high-risk artificial intelligence systems to meet accessibility standards and bans artificial intelligence that exploits people's vulnerabilities,

including disability. The European Accessibility Act, in effect since June 2025, requires private businesses to make digital products and services accessible. Germany put this into law as the Barrierefreiheitsstärkungsgesetz, with fines up to 100,000 euros.

The European Disability Forum called the result a "partial win." The accessibility requirements in the Artificial Intelligence Act apply only to high-risk systems. Most consumer artificial intelligence products, the chatbots and voice assistants and content generators that disabled people actually use daily, fall outside that definition.

Then came the rollback. In June 2026, the European Parliament voted through an "Artificial Intelligence Omnibus" package, 423 in favour, 57 against, that delayed and weakened several protections before they had even taken effect. Broader processing of sensitive health and disability data without individual consent. Equality bodies losing direct access to artificial intelligence documentation. Most high-risk obligations pushed to December 2027.

The German blind and visually impaired association, the Deutscher Blinden- und Sehbehindertenverband, responded with a resolution in June 2026 demanding that artificial intelligence be made inclusive. Their statement said: "Everyone is talking about artificial intelligence, but far too rarely is inclusion and accessibility considered." They called accessibility a "Grundbedingung," a fundamental precondition, not something to add later.

The United Nations Special Rapporteur on disability rights, Gerard Quinn, published a report in 2021 concluding that artificial

intelligence "can be a force for good" if tied to the goals of the Convention on the Rights of Persons with Disabilities. He urged governments to prohibit discriminatory artificial intelligence and to require reasonable accommodation. The Convention has been ratified by 190 countries. India ratified it in 2007. The framework exists. The implementation doesn't.

The Web Content Accessibility Guidelines, the main technical standard for digital accessibility, were written for web pages that sit still. They tell you how to make a button accessible, how to structure a heading, how to provide text alternatives for images. They give no guidance on the accessibility of content that artificial intelligence creates on the fly. They don't address bias in artificial intelligence outputs. They don't cover a chatbot that generates a different interface every time you use it. WebAIM's 2026 report found that 95.9 percent of the top one million web pages still had detectable accessibility failures, at an average of 56 errors per page. That's the baseline, and artificial intelligence is adding complexity on top of it.

What "designed for everyone" should mean.

I keep coming back to a phrase from disability rights: nothing about us without us.

The academic research, the product failures, the biased hiring tools, the speech recognition that doesn't work, the images that reduce disability to a wheelchair, the regulations that arrive too late: all of them share one root cause. Disabled people are not in the room when these systems are designed, built, tested, and launched.

They're not in the training data. Research on disability data in the United States found that out of 30 nationally representative federal surveys, none had a process to make data collection accessible to disabled people. If a survey website doesn't work with a screen reader, a blind person can't complete it. If disabled people can't participate, the data won't represent them. If the data doesn't represent them, the artificial intelligence won't work for them.

Artificial intelligence is built to find patterns in data and to perform best for the most common patterns. Disabled people sit further from the statistical average. An artificial intelligence trained on data that barely includes them will treat their needs as noise. That's the structural problem underneath every failure I've described.

They're not in the design process. A 2025 study of 25 artificial intelligence practitioners at a major technology company found that accessibility work was treated as separate, underfunded, and disconnected from the rest of responsible artificial intelligence work. One practitioner's comment became the study's title: "Accessibility people, you go work on that thing of yours over there."

They're not in the testing. Automated accessibility testing catches some failures. It misses the ones that matter most. A chatbot can score perfectly on every automated test and still trap a screen reader user in a loop they can't escape.

The exceptions prove what's possible. Voiceitt built speech recognition for non-standard speech by working with disability organisations from the start. Google's Project Euphonia collected over 1.5 million speech samples from about 3,000 speakers with speech differences. The Glazko study at the University of

Washington showed that training a language model on disability-justice principles measurably reduced bias. In each case, disabled people were in the data and in the room from the beginning.

What I'd change

After 30 years in technology, 20 of them working on accessibility inside a large European company, and 12 years of conversations on AXSChat with people who know this problem from the inside, my position is plain.

Accessibility is a precondition, not an add-on. If an artificial intelligence feature doesn't work with a keyboard and a screen reader before it launches, it shouldn't launch. If a hiring tool scores disabled candidates lower, it shouldn't be used for consequential decisions. If a speech recognition system has a 60 percent error rate for people with atypical speech, it isn't ready.

Cost must be treated as an accessibility barrier. Assistive technology funding programmes in Ireland, Germany, and the United Kingdom should explicitly cover artificial intelligence subscriptions. The United Nations Convention on the Rights of Persons with Disabilities already establishes assistive technology as a human right. The World Health Organization's 2022 report found that nearly one billion people worldwide are denied access to assistive products they need. Artificial intelligence tools now belong in that category.

Open source is the equaliser, but it needs investment. The free tools that work today, Be My Eyes, Seeing AI, NVDA, depend on corporate partnerships or charity. Community-driven, open fine-tuning of

models for atypical speech, sign language, and easy-to-read text is the most promising direction. It needs public funding.

Standards need to catch up. The Web Content Accessibility Guidelines and the European Accessibility Act were written before generative artificial intelligence existed. Disability organisations need to be writing the new standards, not reviewing them after the fact.

And companies need to stop treating accessibility as a separate concern handled by a separate team. The study of artificial intelligence practitioners I mentioned earlier captured this in its title: "Accessibility people, you go work on that thing of yours over there." Accessibility belongs in the same room as safety, security, and performance, with the same authority to delay a launch.

For India, this matters directly. India has one of the world's largest populations of people with disabilities, over 26 million by official count, and the real number is likely much higher. India's technology industry builds artificial intelligence products used by hundreds of millions of people. The choices made in Bangalore, Hyderabad, and Pune about how to train models, who to test with, and what to charge will shape whether artificial intelligence includes or excludes disabled people across South Asia and beyond.

The German blind association's word was right: Grundbedingung. A fundamental precondition. Accessibility doesn't come after the product is built. It's the condition under which building should begin.