This document was translated from English by machine and may contain mistakes. The English original is at https://www.novaracg.com/2026/11/27/the-fire-was-an-advertisement-what-a-false-alarm-teaches-anyone-buying-sign-language-ai/.

The Fire Was an Advertisement: What a False Alarm Teaches Anyone Buying Sign Language AI

Heather Grizzle · November 27, 2026

An AI fire system mistook an ad for a fire, and everyone saw. When sign language AI gets it wrong, only the Deaf viewer knows.

Spis treści

1. A picture of fire is not a fire
2. Fire alarms fail loudly. Sign language AI fails quietly.
3. Where the context lives in a sign
4. The demo is the happy path
5. Who checks the alarm
6. Five questions for anyone buying sign language AI
7. The difference between funny and dangerous
8. Bibliografia

W SKRÓCIE

- A story making the rounds on LinkedIn this week describes an AI fire detection system in China that saw flames in a restaurant advertisement and set off the building's firefighting system. It is a clean picture of the gap between recognizing a pattern and understanding what it means.
- Fire systems already live with this gap at enormous scale. One Chinese fire alarm cloud platform logged 265,749 suspected fires from August 2018 onward, and 409 of them were real.
- A fire alarm fails loudly: water, noise and a building full of witnesses. Sign language AI fails quietly. When an avatar signs the wrong meaning, the screen still looks fine, and the only person who can tell is the Deaf viewer.
- Buyers are mostly testing sign language AI on the happy path, the clean demo. The questions that matter are what the system does when it is confidently wrong, who is able to see that, and how the Deaf user can say so in the moment.

Somewhere in China, according to a story credited to Al Arabiya English and shared this week by the UX designer Niranjan Kumar M, a restaurant put up an advertisement with flames on it. Big, orange, appetizing flames, the kind that are supposed to make you think of a grill. An AI fire detection system pointed at that space saw the flames, decided there was a fire, and triggered the firefighting system [1].

I have not been able to trace the original report, so I am treating the details as a story rather than a record. But I do not need the details to be exact to recognize the shape of it, because China has been here before. In November 2018, traffic cameras in Ningbo that photograph jaywalkers and put their faces on a public shaming screen flagged Dong Mingzhu, the chair of Gree Electric Appliances, for crossing against the light. She was not there. Her face was printed on an advertisement on the side of a bus that drove through the intersection. Police deleted the photo and said they would upgrade the system [2].

Both stories get passed around because they are funny. A camera that cannot tell a woman from a picture of a woman. A fire system that cannot tell a fire from a menu. Niranjan's post makes the serious point in one line: "Recognising a pattern is not the

same as understanding its context." He was writing for product designers. I want to take that line somewhere he did not, into the part of AI I spend my working life on, which is sign language AI sold to governments, hospitals, schools and transit systems as an accessibility tool.

A picture of fire is not a fire

Start with what the fire detector actually did. A vision system was trained to find the visual signature of fire: color, shape, flicker, brightness against the background. It found that signature. By its own standard it worked. What it could not do was ask the next questions. Is this fire in the room, or on a poster in the room? Is anything burning? Is anyone in danger? Those questions are not about pixels. They are about context and consequence, and nobody had built them into the decision before the system was allowed to turn on the sprinklers.

This is not a rare edge case in fire detection. It is the main problem the field is working on. A paper published this March in the journal Fire describes a Chinese fire alarm cloud platform that logged 265,749 suspected fire alarms starting in August 2018. Of those, 409 were real fires. That is a false alarm rate of 99.85 percent, and the authors' own improved model, which they offer as a step forward, raised the share of alarms that turned out to be real from about 0.15 percent to about 29 percent on one test batch, while still catching more than 98 percent of the real fires [3].



One platform's alarm log from August 2018 onward. The gold segment is drawn wider than its true share so it is visible at all.
Data: Shu et al., Fire, 2026

I want to be fair to the engineers here. Fire detection is supposed to be sensitive. A system that misses a real fire can kill people, so designers accept a lot of false alarms to avoid one missed fire. That is a defensible trade. The point is that it is a trade, someone made it on purpose, and the people paying for it, building owners, fire departments, the neighbors who stand outside in the cold, can see the cost every time the alarm goes off for nothing.

That last part is the thing I want you to hold onto.

Fire alarms fail loudly. Sign language AI fails quietly.

When the fire system misfired on the restaurant ad, everyone knew. There was water. There were alarms. Staff and customers saw it, the fire department may have rolled, somebody filmed it, and the clip ended up on LinkedIn with a laughing emoji. A loud failure is embarrassing, but it is also self-reporting. It produces its own evidence, it lands on the desk of whoever owns the system, and it gets fixed, the same way the Ningbo police fixed their cameras once the whole internet had seen the bus.

Now picture a sign language avatar on a screen at a train platform, in an emergency room waiting area, or on a county website explaining how to apply for benefits. The system takes English text, generates a signing figure, and plays it. Suppose it gets the meaning wrong. Not garbled, not frozen, just wrong: a question signed as a statement, a warning signed as routine, a time or a gate or a dosage that comes out as something else.

What happens next? Nothing. The screen looks fine. The figure is moving its hands. To the hearing staff walking past, and to the procurement officer who signed the contract, the accessibility box has been checked and the job is done. The one person in the building who can see that the message is wrong is the Deaf person watching it, and in most deployments I have looked at, nobody has given that person a way to say so: no button, no reporting line, no log, no audit, no one whose job it is to read the complaints.

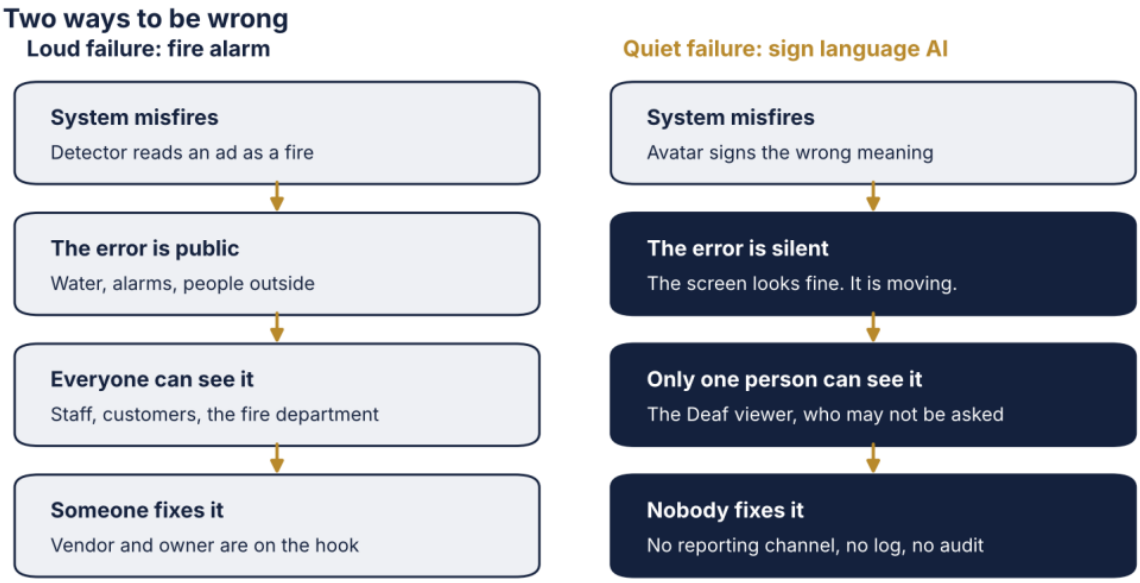

The fire alarm's mistake produces its own evidence. The avatar's mistake produces none unless someone designs a way for the Deaf user to report it.

I am DeafBlind. I have spent my whole life on the receiving end of communication that somebody else decided was good enough, and I can tell you how that usually goes. When the message is wrong, the burden of noticing, explaining and proving it falls on the Deaf person, who is often also the person with the least power in the room. A quiet failure does not stay quiet because it is small. It stays quiet because the only witness is the person nobody built a microphone for.

Where the context lives in a sign

The fire detector's problem was that it saw the right shape in the wrong context. Sign language AI has that same problem, with one difference: in a signed language, a lot of the meaning is the context.

Take two signs, YOU and GO. In American Sign Language, sign them with a neutral face and you have made a statement, roughly "you go" or "you're going." Sign the same two signs with your eyebrows raised and your head tilted slightly forward, and you have asked a yes or no question: "Are you going?" The hands did not change. The grammar moved to the face. Furrow the brows instead, and add a question sign, and you are into a different kind of question again. A system that gets the hands right and the face wrong has not made a small error. It has changed what was said.

That is before you get to everything else a fluent signer reads without thinking. Who is being quoted, which ASL shows by shifting the body and the eyes to take on another person's role. Which name sign belongs to which person, since name signs are given within a community and do not come from spelling. Which regional variant is in use, and whether the signer is using Black ASL, which has its own history and features. How the signing space has been set up, so that a point to the left means the person established on the left two sentences ago. And for DeafBlind people who receive language through touch, whether any of this can be received at all.



Both systems are built and tested mostly at steps 1 and 2. The harm, and the safety, sit at steps 3 and 4.

Most sign language AI is built and measured at the first two steps in that picture. Recognition systems are scored on whether they produced the right label for a clip. Generation systems are demonstrated on whether the hands look smooth and the avatar looks friendly. Those are real engineering achievements. They are also the equivalent of the fire detector correctly spotting orange flicker. They tell you nothing about whether anybody, at the end of the chain, got the right message.

The demo is the happy path

Niranjan ends his post with a list of questions for designers, and one of them is whether we have tested unexpected scenarios, not just ideal ones. In sign language AI procurement, the honest answer is mostly no.

A vendor demo is the happy path by design. One signer or one avatar, good lighting, a plain background, a short sentence chosen because the system handles it well, shown to a room of hearing decision makers who cannot judge the signing. The accuracy number on the slide, when there is one, usually comes from a test set that looks a lot like the training set: similar signers, similar cameras, similar vocabulary. None of that is dishonest. It is just not the restaurant ad.

The restaurant ad, for sign language AI, is the real world. It is the regional sign the training data never saw. It is the announcement written in bureaucratic English that does not map cleanly to any sign. It is the older Deaf man who signs the way people signed at his state school in 1975. It is the DeafBlind woman standing close to the screen because that is the only way she can see it at all. It is the emergency, the moment when the message matters most and the system has the least margin for error, which is the one moment nobody demos.

A fire detection designer at least knows the false alarm rate, because the false alarms are loud and they get counted. For most sign language AI deployments, nobody knows the error rate in the field, because nobody is counting, and nobody is counting because the errors make no noise.

Who checks the alarm

The fire world has a pattern worth borrowing. California's ALERTCalifornia program runs AI over a network of more than a thousand wildfire cameras, and when the AI flags what might be smoke, it does not dispatch anybody on its own. It flags the image to human operators, who check whether the machine is right and then alert firefighters [4]. The AI is fast and tireless. The human brings the context. The decision to act sits with the human.

That is the right shape, and it is the shape Niranjan is pointing at when he asks which decisions should be automated and which need human validation. Sign language AI needs it too, with one condition that hearing institutions keep skipping: the human

who validates has to be a Deaf person fluent in the language being produced. A hearing staff member who took a semester of ASL cannot catch a wrong eyebrow. A vendor's internal reviewer who is paid by the vendor is not independent. If the validation step does not include fluent Deaf reviewers, the system has a human in the loop who cannot see the fire.

Five questions for anyone buying sign language AI

Niranjan's questions were written for designers. Here they are rewritten for the person who signs the contract, which is where most of the harm gets locked in or locked out.

1. What does the system do when it is confidently wrong? Ask the vendor to show you a failure, not a success, and to explain how the system behaves when it meets a sign, a sentence or a signer it was not trained on.
2. Who can see the mistake? Name the person, by role, who would notice a wrong output in your deployment. If the only answer is "the Deaf user," you have a quiet failure by design.
3. How does a Deaf user report it, right then? A reporting channel that requires a phone call, an English email or a form buried three pages deep is not a channel. It should be usable in the user's language, at the moment of the problem, and someone should be accountable for reading it.
4. Who validated the output, and were they fluent Deaf signers independent of the vendor? Ask for names, qualifications and what they reviewed. An accuracy percentage without this is a brochure, not evidence.
5. Which decisions are automated, and which are not allowed to be? Emergency instructions, medical information and legal rights are the restaurant ad scenario every time. Decide in the contract which of those a machine may produce on its own, if any.

None of these questions requires the buyer to understand machine learning. They require the buyer to stop treating a moving avatar as proof that communication happened.

The difference between funny and dangerous

The fire detector in the story will probably be fine. Someone will retrain it, or move the camera, or tell the restaurant to change its sign. The failure was loud, so it will get fixed, and next year it will be a slide in someone's conference talk about edge cases.

The failures I worry about will not make anyone's slide. They happen on screens in places where Deaf people already have the least room to push back, and they look, to everyone but the Deaf person, like success. That is why NCG tests these systems before governments and institutions buy them, with Deaf evaluators fluent in the languages being produced, and why our standard treats a vendor's accuracy claim as a claim to be checked rather than as evidence.

Niranjan asked whether we are testing the limitations of AI as seriously as we are developing its capabilities. For sign language AI in public procurement, the answer today is no. The good news is that it does not take new technology to change that. It takes buyers who ask what happens when the fire is only a picture, and who make sure the person who can tell the difference is actually in the room.

Heather Grizzle-Odland is the founder of Novara Consulting Group, which evaluates sign language AI and other AI accessibility tools before institutions buy them. She is DeafBlind and lives in Minnesota.

> **TIPS**
>
> Know something about a sign language AI system that the public should know? Send it to tips@novaracg.com. Every tip is read by me and kept confidential. If you work in the UK, writing to us is not by itself a protected disclosure under the Public Interest Disclosure Act 1998, so before you share employer information, check your options with Protect (protect-advice.org.uk), the free whistleblowing advice line.

SUPPORT NCG

NCG does this work without vendor money, so our findings stay ours. If it has been useful to you, you can support NCG's independent work here. NCG is a for-profit company, not a charity. Contributions are not tax-deductible and buy no service or influence.

Bibliografia

1. Kumar M, N. (2026, November 25). AI detected a fire. It just didn't realise it was looking at a restaurant advertisement [LinkedIn post, citing Al Arabiya English]. https://www.linkedin.com/in/niruthedesigner/
2. Li, T. (2018, November 22). Facial recognition snares China's air con queen Dong Mingzhu for jaywalking, but it's not what it seems. South China Morning Post. https://www.scmp.com/tech/innovation/article/2174564/facial-recognition-catches-chinas-air-con-queen-dong-mingzhu
3. Shu, X., Li, C., Xu, Y., Wang, J., Huo, Y., & He, J. (2026). Research on highly suspected true alarm model for fire alarm data based on deep learning method. Fire, 9(3), 124. https://doi.org/10.3390/fire9030124
4. Agence France-Presse. (2023, September 12). California firefighters use AI to battle wildfires. Asharq Al-Awsat. https://english.aawsat.com/technology/4540976-california-firefighters-use-ai-battle-wildfires

---

Jak cytować

Grizzle, H. M. (2026, November 27). The Fire Was an Advertisement: What a False Alarm Teaches Anyone Buying Sign Language AI. Novara Consulting Group. https://www.novaracg.com/pl/2026/11/27/the-fire-was-an-advertisement-what-a-false-alarm-teaches-anyone-buying-sign-language-ai/