

This document was translated from English by machine and may contain mistakes. The English original is at <https://www.novaracg.com/2026/10/28/human-review-optional-deaf-ai-and-the-airport-sign-language-pitch/>.

Human Review, Optional: Deaf AI and the Airport Sign Language Pitch

Heather Grizzle · October 28, 2026

Deaf AI tells students its airport signing is interpreter-reviewed and tells airports review is optional. What buyers should require first.

Inhalt

1. The announcement you cannot hear
2. Version one: what the campus press heard
3. Version two: what airports are told
4. Version three: what funders and the university are told
5. Three audiences, three products
6. The review step is the product
7. Whose face is that?
8. Involvement is not authority
9. An airport has already pulled the plug once
10. What an airport should require before a pilot becomes a contract
11. Questions put to Deaf AI
12. The same information, at the same time
13. Quellen

Deaf AI, a Toronto startup selling sign language airport announcements, describes its product differently depending on the audience. The student paper heard about recorded interpreters and human review. Airport executives were offered "automatically translated," "real-time," and human



review as an optional extra. The one safeguard that makes the first story believable is the thing the second story lets a buyer switch off.

KURZ GEFASST

- Deaf AI is piloting signed announcements at Region of Waterloo International Airport in Ontario and is pitching the product to US airports through a one-pager hosted by the American Association of Airport Executives (AAAE).
- On 27 October 2026 The Eyeopener described a cautious design: an interpreter records each announcement, face-swap software alters the face, and interpreters review every clip before it is used.
- The AAAE one-pager describes something else: announcements "automatically translated into ASL" in real time, "optional human-reviewed outputs," testing "in controlled environments," and a claim that the product supports ADA compliance.
- Across six public descriptions there are no published accuracy results, no named customers, and no named Deaf person with authority to approve or reject what goes on the screen.
- On 26 October, fourteen days after an ADA complaint, CVG took a different vendor's AI signer off its train screens. Any airport considering Deaf AI should ask its questions before the pilot, not after.

DEAF AI'S PUBLIC RECORD, BY THE NUMBERS

6

public descriptions of
the product compared

1

0



describes human review
as part of the process

published accuracy or
user testing results

0

customers named

7

written questions put to
Deaf AI

Source: The Eyeopener; Deaf AI one-pager on the AAAE site; MaRS; TMU; Accelerator Centre; Mitacs; deaf-ai.com. Checked 28 October 2026.

The announcement you cannot hear

I am DeafBlind, and I know the airport moment by heart. The gate area goes quiet in that particular way. Heads tilt toward the ceiling speakers. Half the room starts grabbing bags. Something was just said, and it could be a gate change, a delay, a final boarding call, or my own name being paged to the desk. The departure board might catch up in a few minutes, or it might not. I have retinitis pigmentosa, and I cannot see anything above eye level, which is exactly where airports mount their screens. So I read the room, a skill I have been sharpening since I was a kid at the Atlanta Area School for the Deaf. Reading the room is a skill. It is not access.

That gap is real, and it is the gap Deaf AI says it exists to close. On 27 October 2026, The Eyeopener, the student paper at Toronto Metropolitan University, profiled a fifth-year computer engineering student working as a research assistant at the company. It is a warm profile. The student is hard of hearing, wears cochlear devices, and says he wants to help people with disabilities. I take that at face value, and nothing in this piece is about him.

This is about the product, because the product described to a student reporter is not the product described to the people who buy it. I put the company's public materials



side by side and found three different systems: one for the campus press, one for airport executives, and one for research funders. When the whole value of a technology rests on whether a Deaf person can trust what shows up on the screen, three inconsistent accounts of how it works is the finding.

Version one: what the campus press heard

The Eyeopener's account of the airport pilot is the most concrete description of the system anywhere in the public record, and it is a cautious one. According to the article, a sign language interpreter records the standard set of announcements an airport makes. Face-swap software then alters the clips so the interpreter is not identifiable. Human interpreters review the altered videos for accuracy before they are used, because, as the student explains, small changes in movement can change meaning in a signed language. The pilot is described as running at smaller airports in southern Ontario, and a companion navigation app is planned that would guide travelers to check-in, security and gates with sign language video in ASL and in Langue des signes québécoise (LSQ).

Read plainly, that is not machine translation. It is a library of human-produced signed clips with a synthetic face laid over them and a human check at the end. It is a modest design, and in a domain as repetitive as airport announcements a modest design can be a sensible one. A finite set of phrases, recorded by a fluent signer and checked by people who know the language, is the kind of system that could actually be validated, because the whole output can be reviewed before any passenger sees it.

The article also reports what the system does not yet have. It gives no accuracy figures, no results from the pilot, no description of who the reviewing interpreters are, and no statement on whether Deaf people review the output. The one outside voice in the story, Deborah Fels, who directs TMU's Inclusive Media and Design Centre, does not assess Deaf AI at all. Speaking generally, she says she sees more promise in better AI captioning than in AI sign language interpreting.

Version two: what airports are told

The document an airport executive would see is different. Deaf AI has a one-page pitch hosted on the website of the American Association of Airport Executives, the main US trade association for airport operators, presented in connection with the Air-



port Consortium on Transformation and the association's innovation program. The page is undated. Its only named person is the company's chief executive, Mehdi Ma-soumi.

The pitch describes a platform that takes gate changes, delays and final boarding calls and has them "automatically translated into ASL" in real time, shown on screens the airport already owns or on passengers' phones. Quality assurance is covered in two phrases: the system has been "tested in controlled environments," and the offer includes "optional human-reviewed outputs." The page says the solution supports compliance with regulations, giving "ADA accessibility requirements" as its example. It offers a "pilot-friendly, low-cost model" that needs no major infrastructure investment. It names no customers, gives no prices, and publishes no accuracy data.

The company's other sales-facing profiles use the same language. Its MaRS Discovery District listing describes "AI-based real-time voice-to-sign language interpretation" and claims "reliable concept-based interpretations" in controlled communication settings, again with no figures behind the word reliable. That listing also names a partnership with Omnivex, a digital signage software company. The company's own website says it trains AI systems to produce sign language and describes a goal of "digital humans as sign language interpreters."

None of this mentions a recorded interpreter. None of it mentions face swapping. And the interpreter review that the student paper presented as part of how the system works appears here as an add-on the buyer may or may not choose.

Version three: what funders and the university are told

The research-facing record describes a third system. A Mitacs research project between Deaf AI and Toronto Metropolitan University, supervised by Naimul Khan, is titled "An intelligent framework for scalable sign language synthesis." It describes generating signing through a virtual avatar, using deep neural networks to deal with hand self-occlusion and motion blur, trained on a large dataset captured with head-mounted displays such as the Apple Vision Pro. That is generative signing, the hardest and least proven category in the field.

TMU's own news story from 24 September 2026 sits between the two. It identifies the pilot site as Region of Waterloo International Airport and describes the student's work as including ASL video processing, API testing, cloud deployment, and impro-



ving the quality and resolution of "AI-generated video." A June 2026 profile from the Accelerator Centre, which runs the aerospace accelerator the company attended, says the Waterloo pilot began in May 2026 and quotes the airport's operations manager saying the technology "has the potential" to transform airport communication. Neither piece describes how the signing is produced, and neither reports a result.

The kindest reading is that these are stages of one roadmap: recorded clips now, generative avatars later, with the sales copy written for the destination rather than the current location. That may well be true. But a buyer signing a contract buys what the contract describes, and the document written for buyers describes the destination.

Three audiences, three products

Set next to each other, the company's public descriptions disagree on the two things a buyer most needs to know: how the signing is made, and whether a human checks it.



How Deaf AI's airport product is described, by audience

Au- dience	Quelle	How the signing is made	Human review	Evidence of accuracy
Campus press	The Eyeope- ner (27 Oct 2026)	Interpreter records clips; face-swap soft- ware alters the face	Interpre- ters review each clip	None repor- ted
Airport buyers	AAAE one- pager (unda- ted)	Automatically trans- lated into ASL in real time	Optional	Tested in con- trolled envi- ronments
Accele- rator lis- ting	MaRS profi- le (undated)	AI-based real-time voice-to-sign inter- pretation	Not menti- oned	Called reliable with no figu- res
Univer- sity news	TMU News (24 Sep 2026)	AI-generated video	Not menti- oned	Pilot running with no re- sults
Resea- rch fun- ders	Mitacs pro- ject (unda- ted)	Neural avatar syn- thesis	Not menti- oned	No evaluation plan descri- bed
Compa- ny web- site	deaf-ai.com	AI trained to produ- ce sign language; goal of digital hu- mans	Not menti- oned	None publis- hed

Source: Sources listed at the end of this piece; all checked 28 October 2026

Only the student paper, the source with the least commercial weight, describes a re-view step as part of the process. The document an airport would actually rely on in a purchasing decision is the one that makes it optional.




The review step is the product


In sign language AI, a human review step is not a nice extra on top of the technology. For a system like the one the student paper described, it is the only thing standing between a fluent interpreter's original signing and whatever the face-swap software did to it. Face swapping changes the face, and in a signed language the face is grammar. Eyebrow position marks questions. Mouth morphemes change the meaning of a sign. Head movement and eye gaze carry negation, conditionals and who is doing what to whom. A process that regenerates the face of a signed clip is editing the grammatical layer of the message, and the only evidence that it did so without damage is a fluent person watching the result.

Review is part of the process in one account. Optional in the other.

WHAT THE STUDENT PAPER WAS TOLD



WHAT AIRPORTS ARE TOLD



If the buyer declines review, no one checks the clip before passengers see it

The same company, two accounts of the same product. Where the review step is optional, nothing stands between the translation and the passenger.



So when the airport pitch lists human review as optional, it is offering the buyer a version of the product with no evidence of accuracy at all. A procurement officer comparing a low-cost pilot with and without review will reasonably assume the reviewed version is the premium tier and the unreviewed version is the standard one. For a Deaf traveler standing at gate B7, the unreviewed version is a screen showing a signer whose face was produced by software, delivering a gate change that nobody fluent has checked.

The same logic applies with more force to the "automatically translated" and "real-time" language. A finite library of pre-recorded, pre-reviewed announcements can be validated once and reused. A system that translates whatever the gate agent says, live, cannot be reviewed before it is seen, because the review would have to happen faster than the announcement. Real-time and human-reviewed pull in opposite directions. A buyer told both should ask which one the airport is actually getting, announcement by announcement.

PROCEDURAL ASSURANCE VERSUS EVIDENTIARY ASSURANCE

Procedural assurance is a sentence in a brochure that says outputs can be reviewed. Evidentiary assurance is a published record showing who reviewed which outputs, against what standard, with what error rate, and what happened when they found an error. The AAAE one-pager offers the first. The public record contains none of the second.

Whose face is that?

The face swap raises a question none of the public materials answer. If an interpreter's signing is recorded once and then replayed at airports under a different face, the interpreter's hands, timing, body and language are still theirs. Only the face changed, and the face is the part of a signed language that carries grammar. So who agreed to what? Did the interpreter consent to their signing being reused indefinitely, at sites they will never know about, under a face that is not theirs? Were they paid once or per use? Can they pull their work back out? And whose face is on the screen: a synthetic one, or another real person's?

I am not assuming the answers are bad. The student paper frames the swap as anonymization, which suggests someone was trying to protect the interpreter. But protecti-



on is a claim, and claims need records. Two days before the Eyeopener story ran, New Zealand's NZSL Board said it plainly in a statement on AI and New Zealand Sign Language: public visibility does not mean people have consented to their language, image, identity, movement or biometric data being used for AI. That principle does not stop at the border. An interpreter whose signing becomes an airport's permanent announcement library is exactly the person it protects.

I made the general version of this argument in "The Corpus Question": whose signing builds these systems, what they agreed to, what they were paid, and whether they can withdraw. Deaf AI is a live case of the same question, with a smaller corpus and a face swap on top.

Involvement is not authority

Every profile of Deaf AI says the company works with the Deaf community. TMU's story mentions Deaf community members and a Deaf mentor. The Accelerator Centre profile mentions Deaf advisors, mentors and leaders, and Children of Deaf Adults who help bridge the two cultures. The company's TMU startup page describes a "community-informed approach." None of the public materials describes a Deaf person as a founder or executive, and none names a Deaf person who has authority to approve or reject what goes on the screen.

That distinction matters more than it sounds. Advice can be taken or ignored. Sign-off cannot. In a system whose only quality control is human review, the question that decides whether a Deaf traveler can trust the screen is who does the reviewing and whether their "no" stops the clip from shipping. The Eyeopener says interpreters review the output. Interpreters are skilled professionals, and many of them are excellent. But a hearing interpreter checking whether a clip matches the English script is doing a different job from a Deaf native signer checking whether the clip is clear, natural and correct for the people who will actually rely on it. Both checks have value, and a buyer should know which one, if either, is happening.

The NZSL Board drew the same line on 25 October: technology can improve access, but Deaf communities have to be involved in development, and human expertise has to stay central where the stakes are high. Airport operations are not the highest-stakes setting there is, but a missed gate change or a misread security instruction is not trivial either, and the person who carries that cost is never the vendor.



An airport has already pulled the plug once

This is not hypothetical, and US airports already have a precedent to learn from. On 12 October 2026 I filed an ADA Title II and Section 504 complaint against Cincinnati/Northern Kentucky International Airport over AI-generated ASL announcements from a different vendor, Signapse, playing on the terminal train platform screens. The complaint rested on two recordings in which the signer and the screen gave a Deaf traveler two different trains. Fourteen days later CVG Airport Authority wrote back to say it had taken the signed announcements off the train screens, that they will not return until they pass a review that includes Deaf ASL users, and that the complaint had been forwarded to the Federal Aviation Administration.

I count that as a win, and not only for me. It shows that an airport will act when it is shown evidence, and it shows how cheaply that evidence can be gathered: the second recording came from a Deaf man who rides that train as an ordinary passenger, with a phone. It also shows what it costs an airport to find out after deployment rather than before. CVG now has a review, a federal referral and a public record to manage, all over a feature that was supposed to make it look more accessible.

Any US airport considering Deaf AI through the AAAE channel should read the CVG story first. The vendor is different and the technology may be different, but the buying pattern is the same: a company claims ADA support, publishes no accuracy evidence, and lets the airport discover the gap through its Deaf passengers. Under the Title II regulations, a qualified interpreter is one who can interpret effectively, accurately and impartially. A claim of ADA support for a product whose accuracy has never been published is a claim the vendor makes and the airport ends up answering for.

What an airport should require before a pilot becomes a contract

None of this means an airport should refuse to talk to Deaf AI. A recorded, reviewed announcement library for a predictable set of messages could be a reasonable product, and it would be a real improvement over silence. AI signing is not inherently unacceptable. It has to be evidenced. Before any pilot turns into a purchase order, an airport should get the following in writing.



1. The provider **MUST** state which method produces each announcement: a recorded human signer, a face-swapped recording, a generated avatar, or live machine translation. Where the answer differs by announcement type, the provider **MUST** list them.
2. Human review **MUST** be included for every announcement in the contract, not offered as an option, and the provider **MUST** say who performs it: Deaf native signers, hearing interpreters, or both, and whether a reviewer can block a clip from going live.
3. The provider **MUST** publish accuracy evidence from testing with Deaf users in airport conditions, with the error rate on critical messages such as gate changes, security instructions and emergency information reported separately from the average.
4. The system **MUST NOT** improvise machine signing for a message that has no approved clip. Silence, text only, or a handoff to a human are all defensible answers. An unreviewed signed message is not.
5. The deployment **MUST** give a Deaf passenger a way to report a wrong announcement from the gate, with a committed response time from the airport and the vendor.
6. The provider **MUST** disclose the consent, compensation and withdrawal terms for every interpreter whose signing appears on screen, and state plainly whose face the passenger is looking at.
7. The provider **SHOULD** state which sign languages each deployment covers, since ASL is not the only signed language Deaf travelers use and the company's own Canadian materials already promise LSQ.

An airport that cannot get those answers is not buying an accessibility service. It is buying risk, and the people who absorb that risk will be its Deaf passengers.

Questions put to Deaf AI

I have put the following questions to Deaf AI in writing. This piece is published ahead of any response, because every source it relies on is already public. If Deaf AI re-



sponds, its answers will be added here in full, and the relevant sections will be revised where the answers bear on them.

1. Which method produces the announcements currently shown at Region of Waterloo International Airport and any other pilot site: recorded interpreter video, face-swapped recordings, generated avatars, live machine translation, or a combination?
2. The AAAE one-pager describes announcements automatically translated into ASL in real time with optional human-reviewed outputs. The Eyeopener describes interpreter review as part of the process. Is human review included for every announcement a customer receives, or is it an option a customer can decline?
3. Who reviews output before it is shown, are any reviewers Deaf, and can a reviewer stop a clip from being used?
4. What testing has been done with Deaf users, in what conditions, and will the results be published?
5. What consent, compensation and withdrawal terms apply to the interpreters whose signing is recorded, and is the face shown on screen synthetic or another real person's?
6. Which airports, in Canada or the United States, are currently using or contracted to use the system?
7. What supports the statement that the product helps airports comply with ADA accessibility requirements?

The same information, at the same time

What Deaf travelers want from an airport is not complicated. We want the same information everyone else gets, at the same moment they get it, and we want it to be right. A signed announcement that might be wrong, with no way to know, is not the same information. It is a second channel that a Deaf passenger has to second-guess while the hearing passengers are already walking to the new gate.

Deaf AI may have a good product. The design the student paper described could be one. But the company is selling it to airports with the safeguard marked optional, and



until it publishes how the signing is made, who checks it, and how often it is wrong, the only people who will find out are the Deaf travelers standing under the screen.

METHOD AND DISCLOSURE

Novara Consulting Group is independent of Deaf AI, of Signapse and of every airport named here, and holds no financial interest in any sign language AI vendor. This piece compares statements Deaf AI and its partners have made in public sources, all checked on 28 October 2026, and identifies what a buyer cannot currently verify. It makes no allegation of misconduct and no finding about the accuracy of any announcement Deaf AI has produced. NCG has not seen Deaf AI's signed output. The CVG account draws on NCG's own complaint and the airport's written correspondence with NCG. Nothing here is legal advice.

Quellen

All sources were checked on 28 October 2026.

Elsaesser, M. (2026, October 27). Student helps develop tech for sign language interpretation. The Eyeopener.

Deaf AI. (n.d.). Inclusive airport one-pager. Hosted by the American Association of Airport Executives.

MaRS Discovery District. (n.d.). Deaf AI company profile. MaRS.

MaRS Discovery District. (n.d.). Mehdi Masoumi, Deaf AI. MaRS.

Toronto Metropolitan University. (2026, September 24). This TMU student is helping bring sign language AI to Waterloo airport. TMU News.

Toronto Metropolitan University, Social Ventures Zone. (n.d.). Deaf AI startup profile. TMU.

Accelerator Centre. (2026, June 3). Deaf AI is closing the information gap for Deaf travelers. Accelerator Centre.

Mitacs. (n.d.). An intelligent framework for scalable sign language synthesis. Mitacs.

Deaf AI. (n.d.). Company website. deaf-ai.com.



Te Ao Māori News. (2026, October 25). Tāngata Turi raise alarms over AI risks to New Zealand Sign Language. Te Ao Māori News.

U.S. Department of Justice. 28 CFR 35.104, definitions (qualified interpreter). Electronic Code of Federal Regulations.

Novara Consulting Group. (2026, October 26). I Filed One Complaint. An Airport Took Its AI Signer Down.. NCG Insights.

Novara Consulting Group. (2026, October 25). Never Signed, or Signed Wrong: What Two Airports and a Rail FOI Teach Buyers of AI Signing. NCG Insights.

Novara Consulting Group. (2026, September 27). The Corpus Question. NCG Insights.

Zitierweise

Grizzle, H. M. (2026, October 28). Human Review, Optional: Deaf AI and the Airport Sign Language Pitch. Novara Consulting Group. <https://www.novaracg.com/de/2026/10/28/human-review-optional-deaf-ai-and-the-airport-sign-language-pitch/>

