Eleven Signers Gave It 2.67 Out of 5. The Announcement Says They Understood 89 Percent.

Heather Grizzle · November 21, 2026

IIIT Hyderabad's SignPlug signs web text in Indian Sign Language. Its Deaf testers rated understandability 2.67 of 5. Coverage reported 89 percent.

Contents

1. A demo, a jubilee and a number
2. What the students got right
3. Eighty-nine percent of a four-option quiz
4. What the signers said when asked directly
5. How one number left the others behind
6. The face is not a feature request
7. Three descriptions of the step where meaning is decided
8. Whose hands are in the widget
9. A plugin for any website is a governance decision
10. What a web team should get in writing
11. Questions for the SignPlug team
12. The table and the headline
13. Sources

Two undergraduates in Hyderabad built a browser plugin that signs web text in Indian Sign Language, tested it with eleven Deaf and hard-of-hearing signers, and posted the results, including the ones that do not flatter the work. Then their institute's press office and the newspapers kept one number. The students did the honest part. What happened to

their numbers on the way to the headline is the part buyers need to study.

IN BRIEF

- SignPlug, from the Centre for Visual Information Technology at IIIT Hyderabad, lets a reader select text on a web page and watch a 3D avatar sign it in Indian Sign Language (ISL). The paper has been accepted at ICVGIP 2026 and the code is public.
- The institute's announcement on 17 November 2026 says ISL signers understood about 89 percent of the avatar's sentences. The project page shows what that number is: eleven signers picking the right sentence out of four options, 332 times out of 374.
- Asked directly, the same eleven signers rated understandability 2.67 out of 5, hand shape 2.61, body posture 2.50 and naturalness 2.48. None of those ratings appears in the institute's announcement or in the newspaper coverage I read.
- The avatar has no face, and the team says so plainly. In a signed language the face carries grammar, so this is a missing part of the language, not a missing finish.
- The plugin is described as something that can be added to any website, and the work is funded through India's national language translation mission. Any web team thinking of installing it should ask for the full study first, and should not call a site accessible in ISL because a widget is on it.

SIGNPLUG'S USER STUDY, AS POSTED ON ITS PROJECT PAGE

11

Deaf and hard-of-hearing ISL signers in the study

34
avatar sentences each signer watched

88.8%
correct picks in a four-option task, where guessing scores 25 per cent

2.67 / 5
understandability, rated by the same signers

2.48 / 5
naturalness, rated by the same signers

Source: SignPlug project page, checked 18 November 2026

A demo, a jubilee and a number

The Centre for Visual Information Technology at IIIT Hyderabad recently marked its silver jubilee, and one of the demonstrations on show came from two third-year BTech students, Masthan Sharif Shaik and Mohammed Raza. According to the institute's own account, Andrew Zisserman of Oxford and Google DeepMind was among the people who watched. What they showed is easy to describe. A reader highlights a sentence on a web page, clicks a button marked Click To Sign, and a three-dimensional figure in the corner of the screen signs the sentence in Indian Sign Language.

On 17 November 2026 the institute published a feature about it under the headline IIIT-H Develops Accessibility Tool For Hearing Impaired. The next day the Deccan Chronicle ran its own report, headlined with the phrase hearing-challenged people. Both carried the same result: when the plugin was tested with deaf and hard-of-hearing signers, they "understood about 89% of sentences" signed by the avatar.

I am DeafBlind and my language is American Sign Language. I do not sign ISL, I cannot judge this avatar's signing, and nothing below is a verdict on it. What I can do is read. So I read the project page the students put up, the code they released, the institute's feature and the newspaper report, and I laid them next to each other. The paper itself is not public yet, and it may answer some of what follows. But the gap I found is not in the students' work. It is between what they reported and what was repeated.

What the students got right

I want to spend real time here, because this project does several things that companies with far more money have not done. The motion comes from real signers. Instead of animating signs from written notation, the team took the videos in the official ISL dictionary published by the Indian Sign Language Research and Training Centre, used computer vision to recover the body, hand and finger positions, and stored them as reusable 3D poses. They then noticed that the automatic step kept getting fingers wrong, and that in a signed language a wrong finger can be a different word. So they built a correction tool, which they call Gizmo, and fixed the key poses of each sign by hand. That is the right instinct. It treats the hands as language and not as animation.

They also tested it with the people it is for. Eleven Deaf and hard-of-hearing ISL signers watched thirty-four sentences signed by the avatar and were asked about them. The team then published the scores on its project page, the low ones along with the high one, and wrote down what the participants told them: that facial expression matters and is missing, and that some signs change form depending on the sentence around them. The page ends by naming both as the next steps. They released their code. And the whole thing runs on an ordinary laptop, which matters in a country where the lab's own materials cite an estimated 63 million deaf people and about 500 qualified sign language instructors.

Most sign language AI material I read this year would not survive that list. No published test with Deaf users, no ratings, no stated limits. These two students and their co-authors did all three before their conference talk. Keep that in mind through everything that follows, because the problem is what was done with their candor.

Eighty-nine percent of a four-option quiz

Here is what the project page says about the study, in full: eleven signers, thirty-four sentences, and 88.8 percent of sentences identified correctly, which it spells out as 332 out of 374 in a four-option task. Eleven times thirty-four is 374, so each signer appears to have judged every sentence.

A four-option task means the signer watches the avatar and then chooses, from four written or signed candidates, which sentence it was. That is a legitimate research design, and for a first paper it is a reasonable one. But it measures something narrower than understanding. A person who has no idea what the avatar signed will still be right one time in four by guessing. A person who catches one clear sign, the one for home or for eat, can often rule out three of the four options without following the grammar at all. Picking a sentence out of a lineup is recognition with help. Understanding is when someone hands you nothing and asks what the avatar said.

The page does not yet say how the three wrong options were written, whether they differed from the right one by a single sign or by the whole sentence, or whether the thirty-four sentences were chosen because the dictionary covered them well. The paper may. Until it does, 88.8 percent tells us the avatar's signing carried enough signal for eleven fluent signers to beat a multiple-choice test most of the time. It does not tell us that a Deaf reader who meets an unfamiliar paragraph on a government website will understand what the avatar makes of it.

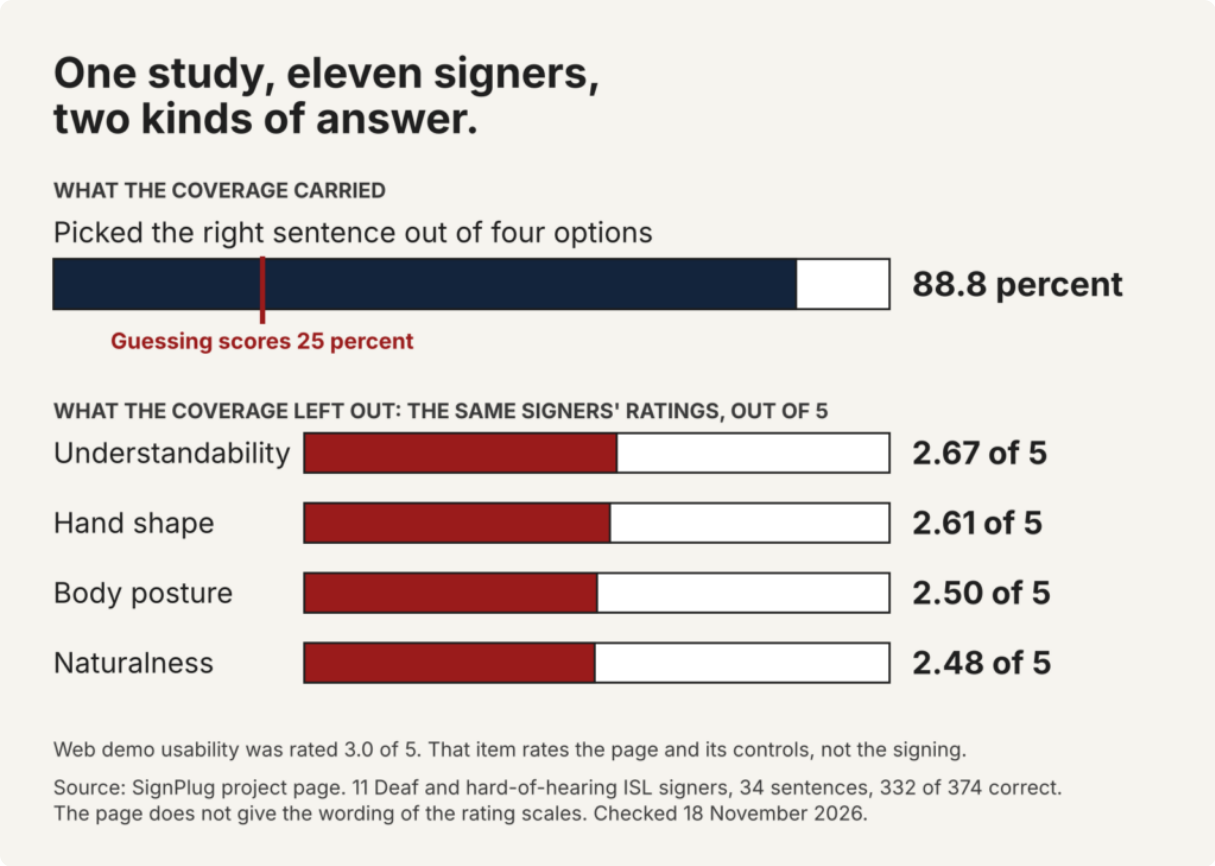


Figure 1. One study, the same eleven signers, two kinds of answer. The multiple-choice score reached the headlines. The signers' own ratings did not.
Novara Consulting Group analysis of the SignPlug project page

What the signers said when asked directly

The same eleven people were also asked to rate what they saw, on scales out of five. Understandability came in at 2.67. Hand shape was 2.61. Body posture was 2.50 and naturalness was 2.48. The one item that reached 3.0 was the usability of the web demo, which is a rating of the page and its buttons, not of the signing.

The project page does not print the wording of the scales, so I will not put labels on those numbers that the team did not. But every rating of the signing itself sits in the lower half of a five-point range, and the people giving those ratings are the only ones in this story qualified to give them. Set that beside 88.8 percent and the two results describe the same avatar honestly: fluent signers can usually work out which of four sentences it means, and they do not find it easy or natural to watch.

There is a third number, from the team's automated check. To show that hand correction helps, they fed videos of the avatar's individual signs to a sign recognition model called Uni-Sign and asked it to find the real dictionary video of the same sign. Before correction, the model's first choice was right 10.17 percent of the time. After correction, 35.17 percent. That is a large improvement, and it supports the students' decision to fix fingers by hand. It also means that after correction a recognition model still picked the wrong sign first about two times in three. A model is not a Deaf viewer, and I would not lean on that figure for more than it can hold. It is consistent with the ratings, though, and not with the word understood.

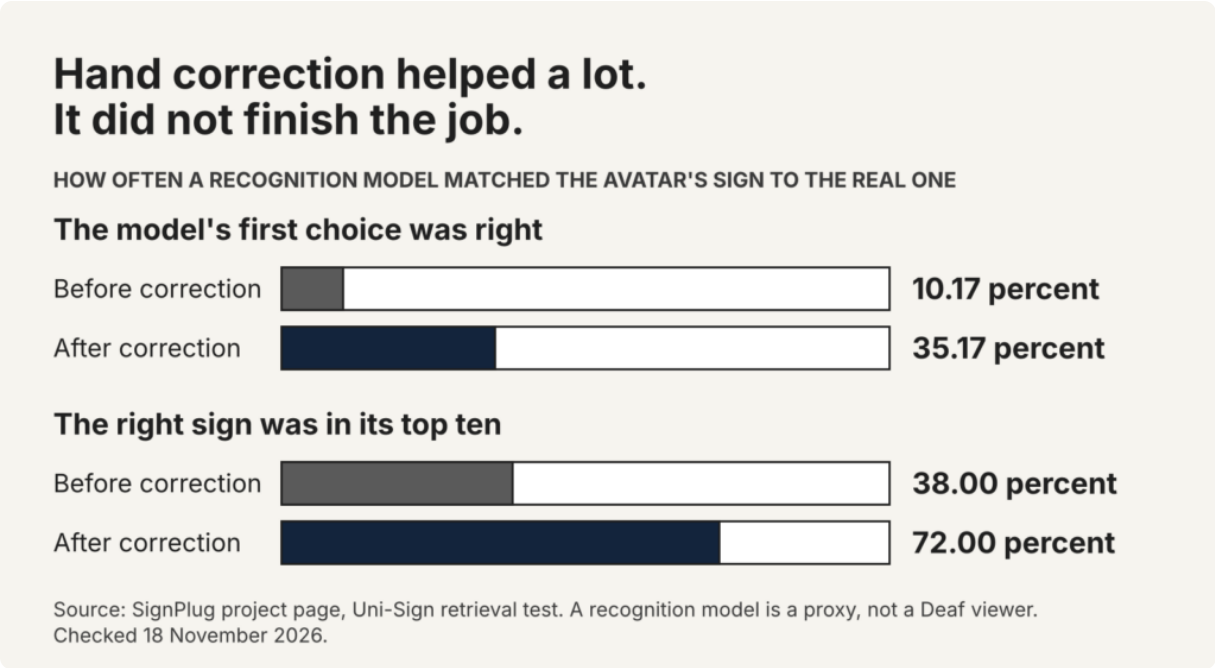


Figure 2. The team's own automated check. Correcting key poses by hand more than tripled the first-choice match rate, and still left the first choice wrong about two times in three.
Novara Consulting Group chart of figures on the SignPlug project page

How one number left the others behind

None of this was hidden. It is all on a public page that the institute's feature links to. But follow each fact from the project page to the press office to the newspaper and watch what survives the trip.

What each account of the SignPlug study told its readers

Fact about the study	Project page	Institute feature	Newspaper report
Number of signers	11	Not given	Not given
Kind of task	Four-option identification	Not given	Not given
Score on that task	88.8 percent	About 89 percent understood	About 89 percent understood
Understandability rating	2.67 out of 5	Not given	Not given
Naturalness rating	2.48 out of 5	Not given	Not given
Missing facial grammar	Stated	Stated	Stated
A Deaf person quoted	Participants' feedback summarized	No	No
Words used for the readers it serves	Deaf and hard-of-hearing	Hearing impaired	Hearing-challenged

Source: SignPlug project page; IIIT Hyderabad blog (17 November 2026); Deccan Chronicle (18 November 2026). Checked 18 November 2026.

The verb changed first. Identified, with a note that there were four options, became understood, with no note. Then the denominator went: neither account says eleven. Then the ratings went. What was left is a sentence that sounds like a finding about comprehension by a community, and it is now the sentence that will be quoted. I have watched this happen to vendor numbers all year. It is more striking here because nobody was selling anything. A careful result was simply rounded up at each handoff by people who meant well.

The last row of the table belongs to the same pattern. The students' page says Deaf and hard-of-hearing. The institute's headline says hearing impaired, and its text describes sign language as gestures. The newspaper says hearing-challenged. No Deaf person is quoted in either piece, including the eleven who did the testing. A tool for signers was announced to the public without a signer in the story.

The face is not a feature request

The team is direct about the avatar's largest limit. It has hands and a body and no working face. Their page says the system covers manual articulators only, and both the institute feature and the newspaper report repeat that facial expressions will come in a later version.

I would ask every reader who does not sign to slow down on that sentence, because in English it sounds like a cosmetic gap, the way a cartoon character might lack eyebrows. It is not. In signed languages the face and the upper body do grammatical work. Whether a sentence is a question or a statement, whether something did or did not happen, how big, how soon, how certain: much of that can live off the hands. The journal I founded on sign language AI is called The Fifth Parameter for this reason. A signing figure without a face is closer to English with the punctuation and many of the small grammatical words removed than to English spoken in a flat voice.

The participants raised a second problem that the team also wrote down. Signs in a dictionary are recorded one at a time, in their plain form. Signs in a sentence change. The page gives the example of the verb go inside a sentence about a man going home. A system that looks up each sign in a dictionary and blends from one to the next produces a chain of dictionary entries in a new order. That is useful for vocabulary. It is a different thing from a signed sentence, and the 2.48 for naturalness is the participants saying so with a number.

Then there is fingerspelling. When the dictionary has no sign for a word, the plugin spells it letter by letter. Consider who the plugin is for. The institute's own framing is that many deaf readers have sign language as a first language and find written English, Hindi or Telugu hard going. A reader like that selects a sentence because the English is the obstacle. If the hard word is the one with no dictionary entry, the avatar hands the same English word back, one letter at a time.

Three descriptions of the step where meaning is decided

Before the avatar moves, the English sentence has to become a sequence of ISL signs in ISL order. Everything after that is lookup and animation. This is the step where the meaning is either carried or lost, and the public material describes it three different ways.

How the English-to-ISL step is described

Source	How English becomes ISL sign order	Accuracy reported
Institute feature	An offline preprocessing step	None
Project page	A large language model at the time of use	None
Released code	A rule-based demo translator run from a small web server	None

Source: IIIT Hyderabad blog; SignPlug project page; SignPlug code repository. Checked 18 November 2026.

These may all be true of different versions. A research demo often has one path for the paper and another for the public code. But they matter differently to a buyer. A fixed set of rules will make the same mistake every time and can be audited. A language model can produce a different sign order for the same sentence on different days and cannot easily be audited at all. And the claim on the project page, echoed in the institute feature, that the plugin runs with "no GPU and no server-side model" sits awkwardly beside a translation step that, in one description, is a language model, and in the released code is a program running on a server.

What none of the three provides is a measure of how often that step gets the ISL right. The user study, as described so far, tests whether signers can identify what the avatar signed. It does not appear to test whether what the avatar was told to sign was a faithful ISL rendering of the English on the page. A reader could understand the avatar perfectly and still be misinformed, because the mistake was made before the avatar started.

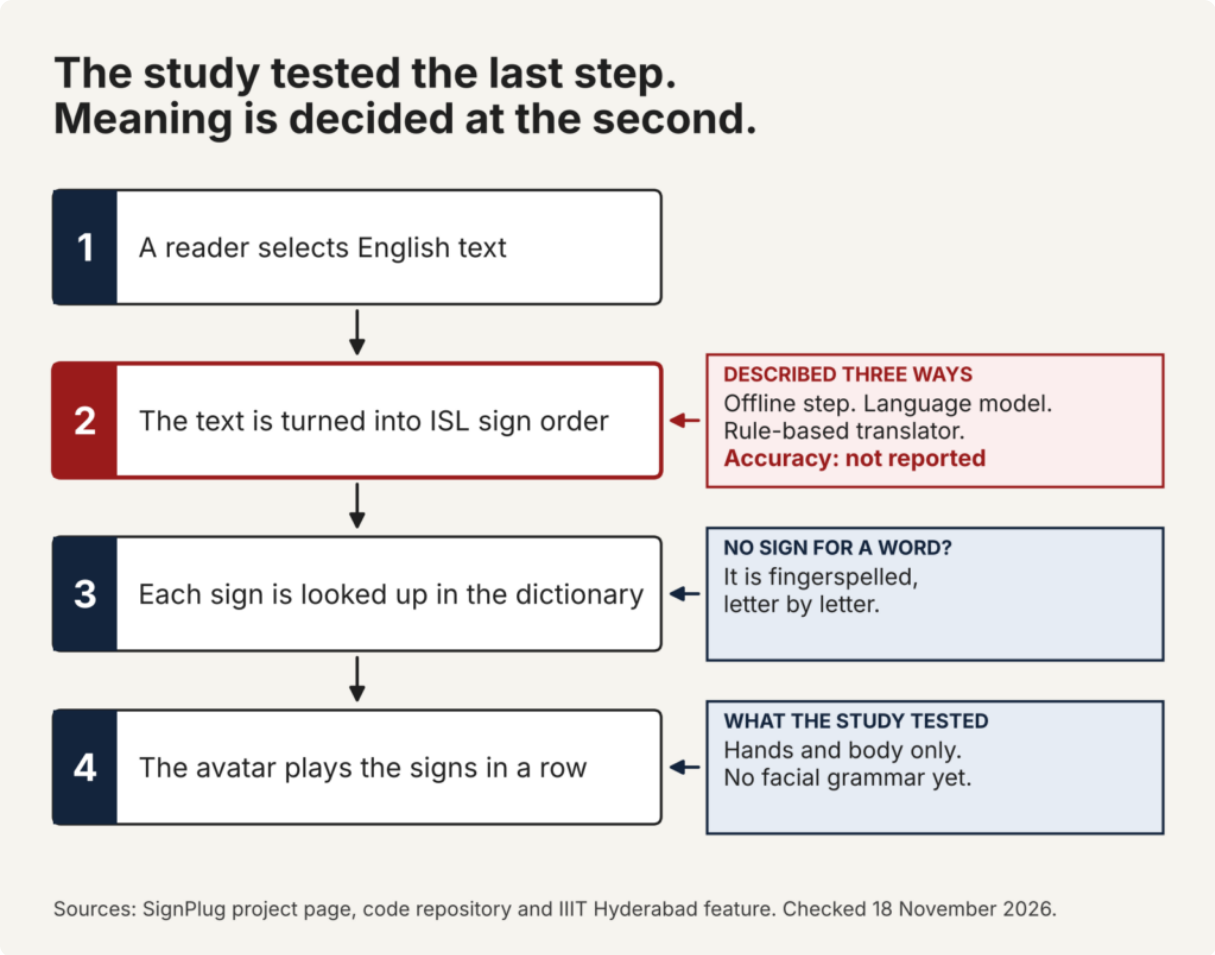


Figure 3. The pipeline in four steps. The published study looks at whether signers can identify the avatar's output. The step that decides what the avatar should say has no reported accuracy.
Novara Consulting Group analysis of the SignPlug project page and code repository

Whose hands are in the widget

The raw material for every sign is a video of a person. ISLRTC, a government institute under India's Department of Empowerment of Persons with Disabilities, has published a video dictionary of 10,000 terms, and it is listed on the national open data portal with a stated purpose: to help people learn ISL. SignPlug reconstructs the signers in those videos into 3D poses, corrects them, and moves the result onto a single neutral avatar so that, the project page says, the signs are independent of the original signer.

The dictionary is a public government resource, and its presence on the open data portal suggests reuse may well be permitted. I am not suggesting the team did anything improper. But a license and consent are two different things, and this field has a

habit of treating the first as if it settled the second. The people who signed those 10,000 entries were recorded to teach a language. Whether they were told their movements might be extracted, edited and replayed through an avatar on any website is a question the public material does not address. The code repository handles the videos carefully, noting that they belong to the ISLRTC dictionary and that "the dictionary is not released with this code." That caution is to the team's credit, and it is also a sign that the question is live. I have written about it at length in The Corpus Question.

One more practical point for anyone who pictures this going into production. The avatar is built on a body model called SMPL-X, which the repository notes is licensed separately. Its standard license covers non-commercial scientific research and prohibits commercial use without a separate agreement. A university demo is squarely inside that. A plugin sold or bundled into a paid service may not be.

A plugin for any website is a governance decision

The feature says the plugin can be added to any website. That phrase is the reason I am writing about a student project at all.

Web accessibility has been through this once already. For years companies sold a line of code that a site owner could paste in to make a website accessible, and disabled users spent those years explaining that the widget did not do what the sales page said. In April 2025 the US Federal Trade Commission finalized an order requiring one of the best-known of those companies, accessiBe, to pay 1 million dollars over claims that its AI plugin could make any website compliant with accessibility guidelines. The lesson was not that the code was worthless. It was that a tool installed by the site owner, judged by the site owner, and advertised by the site owner as proof of access, puts the disabled user last in line.

A signing widget has the same shape. The web team that installs it almost never signs. They cannot see whether the avatar is right. What they can see is that their page now has a sign language button, and that an announcement from a respected institute says signers understood about 89 percent of it. And this is not a hobby project that will stay in a lab. The repository credits its support to India's Ministry of Electronics and Information Technology through the Bhashini national language translation mission, whose purpose is to let citizens reach digital services in their own languages. Where that funding points is my inference, not an announcement: toward public websites, where the content is health advice, exam notices, benefit rules and legal forms.

None of that is a reason to stop the research. ISL is badly under-resourced, human ISL translation of the whole web is not coming soon, and a vocabulary-level signing aid on a learning site is a different risk from an avatar on a hospital page. The demo the team chose, an ISL course page, is about the right level. The World Federation of the Deaf and the World Association of Sign Language Interpreters said in 2018 that avatars should not stand in for human signers, with special concern for information that is live, complicated or carries real weight in Deaf people's lives. A plugin that works on any text on any site has no way to know which kind of page it is on. Somebody has to decide that, in writing, before it is installed.

What a web team should get in writing

AI signing is not inherently unacceptable. It has to be evidenced. If a ministry, university, bank or hospital is offered this plugin, or any select-and-sign avatar, these are the things to settle first.

1. The provider MUST supply the full study: how many Deaf signers took part, how they were recruited, what the task was, how the wrong options were written, the wording of every rating scale, and the results for each sentence.
2. Comprehension MUST be tested with open responses on the kind of content the site carries, not only with multiple-choice picks from prepared sentences, and errors on critical content MUST be reported separately.
3. The provider MUST state which method turns the page's text into ISL sign order on the live system, and MUST report how often Deaf ISL signers judge that rendering correct.
4. The widget MUST NOT be used on health, legal, financial, examination or emergency content until facial grammar is implemented and tested with Deaf signers.
5. The widget MUST tell the user on screen that the signing is machine-generated and has no facial grammar, MUST offer a route to human-made ISL video or a human interpreter where one exists, and MUST let a user report a wrong rendering in ISL.
6. The site MUST NOT describe itself as accessible in ISL on the strength of the widget.

7. The provider MUST state the terms under which the dictionary videos were used, what the recorded signers consented to, and the license covering every third-party model in the product.
8. Deaf ISL signers from Deaf-led organizations MUST hold authority to approve or block a release, not only a seat in the test group.
9. For fixed pages that matter, the buyer SHOULD pay Deaf translators for human-signed ISL video first and treat the avatar as a supplement.

Questions for the SignPlug team

The paper is not yet public, and several of these may be answered in it. I am putting them on the record now, and the team and the institute are welcome to answer. Anything they send will be added to this piece in full.

1. In the four-option task, how were the three incorrect options constructed, and how were the thirty-four sentences selected?
2. What were the anchors on the five-point scales, and what was the spread of ratings across the eleven signers?
3. Which method produced the ISL sign order for the study sentences: the language model, the rule-based translator, or hand-written glosses?
4. How many signs are in the Widget Dictionary today, and how often did test sentences fall back on fingerspelling?
5. Were Deaf ISL signers involved in designing the system or the study, and are any among the authors?
6. Under what terms were the ISLRTC dictionary videos used, and was ISLRTC or any Deaf organization consulted?
7. Will the institute correct its announcement to say that the 89 percent figure comes from a four-option identification task with eleven signers, and add the ratings?

The table and the headline

I would hand the SignPlug project page to any procurement officer as an example of what a first report on a signing avatar should contain. It names its participants as

Deaf and hard-of-hearing signers. It gives a count. It says what the task was. It prints ratings that a marketing department would have buried, and it lists what the participants said was missing. If the vendors I have reviewed this year published a page like it, my work would be shorter.

The headline is what will travel, though. Somewhere a web team will read that deaf signers understood 89 percent of what an avatar signed, paste a script into a site, and tick a box. The eleven people who sat through thirty-four sentences and wrote down 2.67 will not be in the room. The students put their answer on the page. The least the rest of us can do is repeat all of it.

> **METHOD AND DISCLOSURE**
>
> Novara Consulting Group is independent of IIIT Hyderabad, the SignPlug authors, ISLRTC and every other organization named here, and holds no financial interest in any sign language AI vendor. This piece compares the SignPlug project page, its public code repository, the institute's announcement and one newspaper report, all checked on 18 November 2026. The paper had not been released at that date. I sign ASL, not ISL, and NCG has not evaluated the avatar's signing. This piece makes no allegation of misconduct and no finding about the accuracy of SignPlug's output. The students are named because they are the named authors of public work, and the criticism here is directed at how their results were reported and at how such a tool could be deployed. Nothing here is legal advice.

Sources

All sources were checked on 18 November 2026.

Shaik, M. S., Rahman, M. R. U., Rathore, A. S., Mazumder, S., Jawahar, C. V., and Namboodiri, V. P. (2026). SignPlug: ISL plugin for the web. Project page. Paper accepted at ICVGIP 2026.

IIITH-Research. (2026). SignPlug code repository. GitHub.

Chebbi, S. (2026, November 17). IIIT-H develops accessibility tool for hearing impaired. IIIT Hyderabad Blog.

Dammala, R. (2026, November 18). IIIT-H students make online reading easier for hearing-challenged people. Deccan Chronicle.

Centre for Visual Information Technology. (2026). Projects from CVIT that clinched ANRF ARG awards. CVIT, IIIT Hyderabad.

Government of India. (n.d.). Indian Sign Language Dictionary. Open Government Data Platform India.

Indian Sign Language Research and Training Centre. (n.d.). Frequently asked questions. ISLRTC.

Max Planck Institute for Intelligent Systems. (n.d.). SMPL-X model license. MPI.

Federal Trade Commission. (2025, April 22). FTC approves final order requiring accessiBe to pay $1 million. FTC.

Office of the Principal Scientific Adviser to the Government of India. (n.d.). Natural Language Translation mission. PSA.

World Federation of the Deaf and World Association of Sign Language Interpreters. (2018, March 14). WFD and WASLI statement on use of signing avatars. WFD.

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

> **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.

How to cite

Grizzle, H. M. (2026, November 21). Eleven Signers Gave It 2.67 Out of 5. The Announcement Says They Understood 89 Percent.. Novara Consulting Group. https://www.novaracg.com/2026/11/21/eleven-signers-gave-it-2-67-out-of-5-the-announcement-says-they-understood-89-percent/