NOVARA
CONSULTING GROUP

Sign-Speak: Deaf-owned at the top, consumer terms underneath

Heather Grizzle · October 21, 2026

NCG's governance record on Sign-Speak: Deaf-owned and piloted by Syracuse, but its only public terms are consumer terms that forbid institutional use.

Contents

1. Why this record exists
2. What exists today
3. Questions put to Sign-Speak
4. Assessment
5. Sign-Speak's response
6. Frequently asked questions
7. Sources

Sign-Speak is the only company on NCG's roster that is Deaf-owned, United States-based, built for ASL and deployed by a United States university. The only terms it publishes were written for someone using the app on their own.

IN BRIEF

Sign-Speak holds three things most sign language AI providers do not: Deaf ownership, a route to human interpreters, and a published list of uses its product must not be put to. Its only public terms are a 2024 consumer licence that restricts use to personal, non-commercial purposes, and the company has now confirmed that its institutional deployments and enterprise products each run under separately negotiated agreements, which are under NDA. In written answers on 17 and 18 October 2026, Sign-Speak told NCG that signed video from the Syracuse pilot, its other institutional deployments and SignCaption has not been used to train its models, that its training data comes primarily from compensated Deaf contributors, and that every deployment of its automated interpreting includes a one-button route to a human interpreter at no cost to the user. It also said that a contributor who wants to withdraw consent is required to repay what they were paid, with exceptions decided case by case. None of these commitments yet appears in a published document a buyer can read, and the only accuracy evidence offered is a pilot study written by the company's founders. The company's answers are published in full below.

AT A GLANCE

9 Apr 2024

Date of Sign-Speak's only published terms and privacy policy

$0

Cost to the user of the human interpreter route, according to Sign-Speak

66

Paid participants across the company's only published accuracy study

0

Independent evaluations of Sign-Speak's output NCG could locate

Why this record exists

Sign-Speak is the fourth sign language AI provider NCG has reviewed under the method it first applied to Signapse in September 2026. The method does not change from company to company. NCG reads what the company publishes, puts specific questions to it in writing, sends the company its understanding of any answers before they are published, and publishes a record that ties every statement to a date and a document. The method is set out in the SLAT Index framework (https://www.slatindex.org/framework/).

Sign-Speak matters to United States public buyers more directly than any other company NCG has reviewed. It works in American Sign Language in both directions, recognising signing and turning English into ASL through a digital avatar. Syracuse University made it available to every student, faculty member, staff member and visitor between 1 November and 16 December 2025, and the National Susan B. Anthony Museum and House in Rochester began offering it to visitors in September 2025. Its enterprise products are sold for live events, meetings and broadcasts, and for integration into other companies' platforms through an API and SDK. A university accessibility office, a state agency or a city deciding whether to follow Syracuse needs to be able to read the governance behind the product before it signs. This record sets out what that governance looks like from the outside, document by document, with the date each was retrieved.

What exists today

Terms of Service. Sign-Speak publishes its customer terms at sign-speak.com/legal/customer, headed as an End User License Agreement and defined throughout as Terms of Service, last revised 9 April 2024 and retrieved by NCG on 7 October 2026. They are governed by New York law. Section 5.1 provides that a user "may access and use the Service only for your personal, non-commercial purposes," and it also bars

making the Service available to third parties or running it as a service bureau for others' benefit. A university that offers the service to its whole campus, a museum that offers it to visitors through a QR code and a company that builds it into its own product are each doing what that clause forbids. The published terms therefore cannot be the instrument that governs an institutional deployment. Those deployments must run under separate agreements, and none is public. On 18 October 2026 Sign-Speak confirmed that the Syracuse pilot, the museum deployment, SignLive, SignCaption Enterprise and the API each operate under a separately negotiated agreement rather than the published terms, and it said on 17 October that those agreements are under NDA and that its restrictions on use, and on the use of data, are typically higher in them. NCG also asked whether self-serve API developers accept the published terms. The company's one-line confirmation did not address that part. Asked again, Sign-Speak's co-founder and chief product officer, Nikolas Kelly, said on 20 October that for nearly all of its clients the company ends up negotiating a bespoke agreement that sets out restricted uses and protections for the end consumer. That describes the usual case rather than the exception, so which terms bind a developer who signs up without a negotiated agreement remains unclear.

The same clause excludes use for health diagnostics, for operating equipment whose failure could cause serious harm or in medical emergencies, for deciding eligibility for financial or lending services, and for any interaction or communication with police or other law enforcement. A list this specific is rare in this market, and a buyer can contract against it. It does not cover education accommodations, legal proceedings, employment decisions or disciplinary processes, in each of which a communication failure carries consequences, and nothing public says whether the exclusions are enforced in the product or exist only in the licence. Section 7.1 grants the company a licence over content a user uploads that is "perpetual" and "irrevocable" and extends expressly to using that content to "develop and train our AI Models." Section 15 requires individual arbitration, with hearings in New York by default, and waives class actions, a term many public institutions cannot accept without amendment. Section 4.1 sets a minimum age of 13, with a parent or guardian accepting the terms for any user under the age of majority.

Privacy policy. The privacy policy at sign-speak.com/legal/privacy carries the same date, 9 April 2024, and was retrieved on 7 October 2026. It confirms that uploaded content "may include voice recordings, video recordings, and written text" and that Sign-Speak "will use this content to assess the performance of and improve our own

products, technologies, and services, including by training our own AI models." It says uploaded content is shared with "the AI service providers that power our products" without naming them and, to the company's credit, that "we do not allow providers to train their AI models using your data." It states that "we keep personal information in perpetuity" for the purposes it describes. It describes no way for a person to ask for their information or their video to be deleted, offering only marketing opt-outs and cookie controls. It does not treat signed video as biometric or sensitive information, although a video of a person signing carries their face, their body and their movement, and it does not address children. In its written answers Sign-Speak described deletion routes the policy does not. SignCaption users can delete their account, which the company says deletes their data, and anyone else can email contact-us@sign-speak.com, after which the company says the data is removed by Sign-Speak and by its downstream processors, naming Amazon Web Services as the most notable. Because it does not try to re-identify users who have not supplied an identifying label, the company says it asks a person when they used the system so it can find their data, and that data from several deployments is disposed of automatically. Sign-Speak says its Terms of Service and Privacy Policy are being revised. No revised version had been published when this record was prepared, so the versions of 9 April 2024 are the ones cited.

Training data. In an Amazon Web Services profile published on 1 May 2025, Sign-Speak is described as having accumulated "a large corpus of data," most of it video, "gathered from product pilots and through gamification," and training within the Deaf community is described as crucial to the product. A January 2024 accelerator profile reported that the company had increased its datasets more than tenfold. Read alongside section 7.1 and the privacy policy's retention statement, those accounts raise a question any deploying institution has to be able to answer for its own users: whether signed video captured during a pilot like Syracuse's became training data, what the people signing agreed to, in what language that was explained to them, and whether they can take it back. No published document answers it. Sign-Speak has now answered in writing. On 17 October 2026 it told NCG that, to date, signed video from the Syracuse pilot, its institutional deployments and SignCaption has not been used to train its models, and that its primary training data comes from Deaf individuals who are compensated for their data and time. On 18 October it described how those contributors consent: the consent is presented first in English, and a Deaf member of the company then explains it in ASL on a video call. A contributor who later wants to withdraw consent is required to pay back the money they were paid for the data, with

exceptions handled case by case, and the company says it does not share or sell contributor data. NCG has no means of verifying these statements and records them as the company's. The repayment condition is discussed in the assessment below.

What is on sale now. Sign-Speak's homepage, retrieved on 7 October 2026, leads with SignCaption, which captions ASL video and synchronises the captions to the signs. It describes the real-time tool for ASL-to-voice, voice-to-ASL and captioning as available through a "Pioneers Program," a limited pilot for existing SignCaption users. The enterprise page offers SignCaption Enterprise; SignLive, in which software connects to a venue's sound system and "an automated avatar" interprets in ASL for live events, meetings and broadcasts; and the API and SDK. The current commercial status of the two-way real-time tool that Syracuse and the museum deployed is therefore unclear from the public record, and a buyer should establish which product, which version and which terms it is being offered before comparing it with what Syracuse tested.

Accuracy. The enterprise page states that "Our AI is Deaf-vetted for accuracy, protecting your brand," and offers a dashboard to "access and edit all AI generated captions" to "ensure 100% accuracy." That describes people correcting captions by hand. It is not a measure of the system. No figure, method, sample or denominator appears on the pages NCG retrieved on 7 October 2026. NCG's July 2026 review recorded a company-reported BLEU score of about 60, described as comparable to an average human interpreter. NCG could not find that figure on the company's current pages or in the AWS profile, and records it as previously published rather than current. BLEU counts word overlap between a machine translation and a reference text. It says nothing about ASL grammar, facial grammar, spatial reference, or whether a Deaf person understood what they were shown. No independent evaluation of Sign-Speak's output in either direction has been located, and who does the vetting, with what qualifications, on what sample and against what standard is not public. Asked what Deaf-vetted means, Sign-Speak said on 17 October that it runs case studies and quality assurance with Deaf people, designed the product with Deaf people, and that many of its staff are members of the Deaf community. It pointed NCG to a paper in the Extended Abstracts of CHI 2025, "Evaluating an AI Bi-Directional System for Communication Between Deaf and Hard of Hearing Individuals and Hearing Persons: A Pilot Case Study," whose authors are Zhiwei Yu and Sign-Speak's three co-founders, and on 18 October said that accuracy is measured and reported in both directions. Sign-Speak sent NCG the paper's full text on 20 October, and NCG has checked that statement against it. The paper does report a figure for each direction, but the two are different

kinds of measure. For ASL to English, it reports a SacreBLEU score of 55.1 ± 6.67 on held-out test sentences from signers who were not in the training data, set against a BLEU of 27.6 ± 5.26 for human interpreters whose source and method the paper does not give. For English to ASL, it reports an accuracy of 78% ± 1.7% for the company's avatar and 61% ± 2.4% for HandTalk's, as reported by the 48 Deaf and hard of hearing participants in a usability study of avatar preference, and sets that against 72.7% for expert human interpreters from a separate 2015 study. The first is an automated word-overlap score of the kind described above. The second is participants' own report of what they understood. Neither measures comprehension under a fixed protocol, and neither can be read against the human figure beside it as like for like. The paper's two usability tests drew 18 and 48 participants, each paid with a $25 gift card, and all four authors work for Sign-Speak. Kelly noted that while the numbers were self-reported, the paper passed review by CHI's reviewers, and that the avatar and recognition models have improved since it was published. NCG has no figures for the current versions. A peer-reviewed pilot case study written by the company is evidence, and publishing it is to the company's credit, but it is not an independent evaluation. The company also said it is happy to ask Syracuse for the rights to publish a case study of the pilot.

What the CHI 2025 paper reports, by direction

Direction	What the paper reports	How it was measured	Comparison the paper offers
ASL to English (recognition)	SacreBLEU 55.1 ± 6.67	Automated word-overlap score on held-out test sentences from signers not in the training data	Human interpreter BLEU of 27.6 ± 5.26; source and method not given
English to ASL (avatar)	78% ± 1.7%, against 61% ± 2.4% for HandTalk's avatar	Accuracy as reported by 48 Deaf and hard of hearing participants in an avatar-preference study	72.7% for expert human interpreters, from a separate 2015 study

Source: Yu, Wilkins, Payano and Kelly, CHI Extended Abstracts 2025; full text supplied by Sign-Speak, 20 October 2026

Human interpreters. The enterprise page describes the model as "Humans when you need them, automation when you don't want the wait." Reporting on the museum deployment in September 2025 describes video remote interpreting that connects visitors to human interpreters "for complex situations." A human route of this kind is rare in this market, and Sign-Speak deserves credit for it. Sign-Speak's written answers define the route more precisely than its pages do. It says that in every deployment of its automated interpreting service a user can roll over to a human interpreter through video remote interpreting with a single button at any time, that the interpreters come from several VRI partners, and that Sign-Speak pays the difference so the human route is always free to the end user. That is a stronger arrangement than most of this market offers. The human is offered rather than required: the answers describe no setting in which the system hands a conversation to a person without the user asking, including the settings the company's own licence rules out, and they do not say what certification the partner interpreters hold, how quickly one is reached, or which parts of SignLive output reach an audience with no human review.

Reporting a wrong rendering. NCG asked what route a Deaf user has, in ASL, to report a wrong rendering. Sign-Speak said developers can connect a /feedback endpoint

to raise rendering or recognition accuracy issues, that its automated interpreting service has its own feedback mechanisms, that anyone can write to contact-us@sign-speak.com, that staff have joined ASL video calls through its Intercom support channel to resolve problems, and that ASL support is available during customer service hours. It did not say whether it keeps an incident record, and it did not answer who is responsible when a wrong rendering reaches a user through a third party's integration, or what use restrictions it places on integrators.

Deaf ownership. Sign-Speak describes itself as "Deaf owned" on its team page and as "a Deaf led company" on its enterprise page, and its homepage says it is "Built by Deaf people, for Deaf people." Public reporting identifies co-founder Nikolas Kelly, an alumnus of the National Technical Institute for the Deaf at Rochester Institute of Technology, as the company's Deaf owner and chief product officer, alongside co-founders Yamillet Payano, chief executive, and Nicholas Wilkins, chief technology officer. Deaf ownership is the strongest form of Deaf authority the SLAT framework contemplates, and almost no other company in the field has it. Below the level of ownership, Sign-Speak says it deliberately has no separate Deaf role or advisory board, because Deaf staff hold roles throughout the company; that most staff, Deaf staff included, have "break-glass" mechanisms to stop a deployment; and that several Deaf people hold roles with explicit authority to halt a release or deployment. It says community members who contribute signing are compensated and that it has paid millions of dollars back into the Deaf community. None of this is written into a published document, and NCG has not seen the authority it describes.

The Syracuse pilot. Syracuse University's announcement of 12 November 2025 says the service was available to test until 16 December 2025 for "lectures, meetings, social events and everyday conversations," and that usage data would inform whether it "can be offered as a permanent campus resource." The university's accessibility page, retrieved on 7 October 2026, is still live, carries no dates, describes the tool as designed "to make it easier for Deaf, Hard of Hearing, and hearing students to communicate naturally in class," and states that "Sign-Speak is not intended to replace, but rather to supplement live ASL interpretation on campus." That is the right scope statement, and the university deserves credit for putting it in writing. Neither Syracuse nor Sign-Speak has published an outcome, an evaluation or a decision on permanent availability. Whether the pilot measured anything beyond usage, and whether it measured linguistic accuracy at all, has not been stated.

The company. Sign-Speak Inc. gives its address as 104 East Ave, Suite 205, Rochester, New York. Public profiles associate it with Techstars, Google for Startups, Launch NY, the Boston Impact Initiative and 2Gether-International, and with the Syracuse Surge and AWS Impact accelerators, and report a convertible note of undisclosed size in April 2026. NCG has not verified the company's state of incorporation or ownership structure from a primary register and makes no statement about either beyond the company's own description of itself as Deaf-owned.

Questions put to Sign-Speak

NCG sent the following questions to Sign-Speak on 7 October 2026, with a pre-publication copy of this record, and gave the company until 19 October 2026 to answer. Sign-Speak answered on 17 October and answered four follow-up questions on 18 October. Both responses are published in full below, and the sections above have been updated where the answers bear on them.

1. The only published terms restrict use to personal, non-commercial purposes and forbid making the service available to third parties. Which terms governed the Syracuse University pilot and the National Susan B. Anthony Museum deployment, and which govern SignLive, SignCaption Enterprise and the API and SDK? Do the section 5.1 exclusions carry into them, and will Sign-Speak vary the section 15 arbitration and class action waiver for public institutions?
2. Sign-Speak has described its training data as gathered in part from product pilots and through gamification, and section 7.1 licenses uploaded content perpetually and irrevocably for training. Was signed video captured during the Syracuse pilot, the museum deployment or any other institutional deployment used to train Sign-Speak's models? If so, what did the people signing consent to, how and in what language was that consent sought, and can a signer withdraw their video and have it removed from training data?
3. The privacy policy keeps personal information in perpetuity and describes no deletion process. How can a person have their signed video deleted? Which AI service providers receive uploaded signed video, where is it held, and for how long?
4. What does "Deaf-vetted for accuracy" mean in practice: who vets, with what qualifications, on what sample, against what criteria, and with what result? Is there

measured accuracy for each direction, ASL to English and English to ASL? Does the BLEU figure NCG recorded in July 2026 still stand, and for which version? Has anyone independent of Sign-Speak evaluated the output, and will the Syracuse pilot's findings be published?

5. "Humans when you need them": who decides when a human interpreter is brought in, what certification do those interpreters hold, how quickly is one reached, who pays, and is a human ever required rather than offered, including in the settings section 5.1 excludes and in education accommodations? Which parts of SignLive output reach an audience without human review?
6. Below the level of ownership, is there a Deaf review body with a charter, and can a Deaf role at the company require a change or halt a release or deployment over commercial objection? Is that written down? Are Deaf community members who contribute signing through pilots or gamified collection paid?
7. What route does a Deaf user have, in ASL, to report a wrong rendering, and does the company keep an incident record? When Sign-Speak output reaches a user through a third party's integration, who is responsible for a wrong rendering, and what use restrictions does Sign-Speak place on integrators?

Assessment

This assessment rests on the public record as of 7 October 2026 and on Sign-Speak's written answers of 17 and 18 October 2026.

Sign-Speak holds three things most of this market does not. Its ownership is Deaf, which puts Deaf judgment at the top of the company rather than in an advisory seat. It offers a route to human interpreters, which is the escalation path high-stakes access needs. And it publishes a list of excluded uses specific enough to contract against, while barring its own AI providers from training on users' data. A public buyer should give the company credit for all three, and should give Syracuse credit for scoping its pilot honestly in public.

Sign-Speak's answers close several of the questions the public record left open, and they close them in the company's favour. It says no signed video from Syracuse, its other institutional deployments or SignCaption has gone into training, that a person can have their video deleted on request, that the people whose signing does train its

models are paid and are taken through consent in ASL, and that the human interpreter route is free and one button away in every deployment. If those statements hold, Sign-Speak's practice is well ahead of its paperwork.

That gap is now the finding. Every one of those commitments exists in an email to NCG, not in a document a buyer, a student or a contributor can read. The published terms still take a perpetual, irrevocable right to train on uploaded video, the privacy policy still keeps personal information in perpetuity and describes no deletion route, and the agreements that actually govern institutional use are under NDA. A buyer cannot rely on what it cannot see, and a person signing into the camera at Syracuse had no way to know any of this. The company says its terms are being revised, and the revision is where these answers belong.

What Sign-Speak told NCG in writing	What a buyer or signer can read today
Signed video from Syracuse, other deployments and SignCaption has not been used for training	The published terms grant a perpetual, irrevocable licence to train on uploaded content
Video can be deleted on request, including from AWS	The privacy policy keeps personal information in perpetuity and describes no deletion route
A human interpreter is one button away and free in every automated deployment	No published document describes the human route
Institutional use runs under negotiated agreements with tighter restrictions	Those agreements are under NDA, and the only public terms are a 2024 consumer licence
Paid contributors are taken through consent in ASL on a video call	No contributor agreement is published
Accuracy is measured in both directions	The only published evidence is the company's own pilot study, with a different kind of measure in each direction

One answer cuts the other way. A contributor who wants to withdraw consent to the use of their signing is required to repay what they were paid for it, with exceptions decided case by case. Paying people for their signing is right, and Sign-Speak deserves credit for it. Conditioning withdrawal on repayment means that the practical ability to take one's face and language back out of a model depends on being able to afford to re-

turn the money. Sign-Speak confirmed this description on 20 October: a contributor can either repay the money or discuss an exception with the company. It did not say whether the contributor agreement itself states the condition. A buyer whose service will rest on models trained on that data should ask how often withdrawal has been requested, how often the exception has been granted, and to see the clause in the contributor agreement.

Accuracy remains reported by the company rather than independently measured. The evidence offered is a peer-reviewed pilot case study written by Sign-Speak's own staff, which pairs an automated score in one direction with participants' self-reported accuracy in the other, on models the company says have since changed. The Syracuse pilot's findings are still unpublished.

A buyer considering Sign-Speak, whether to follow Syracuse onto a campus or to build it into a service, should ask for four things before signing. The first is the agreement that will govern the deployment, with the company's answers here written into it: no use of its users' signed video for training without separate, informed consent sought in ASL, a deletion route with a stated time limit, and a named list of the processors that receive the video. The second is a written definition of the human interpreter route, stating the certification of the interpreters, how quickly one arrives and when a human is required rather than offered. The third is a statement of who answers for a wrong rendering, including one that reaches users through an integrator. The fourth is measured accuracy evidence in both directions, from the Syracuse pilot or from the buyer's own evaluation by Deaf raters, before the tool is used anywhere a student's access depends on it.

☐ The agreement that will govern the deployment, ruling out training on users' signed video without separate, informed consent sought in ASL

☐ A deletion route with a stated time limit, and a named list of the processors that receive the video

☐ A written definition of the human interpreter route: interpreter certification, response time, and when a human is required rather than offered

☐ A statement of who answers for a wrong rendering, including one that reaches users through an integrator

☐ Measured accuracy evidence in both directions, from the Syracuse pilot or the buyer's own evaluation by Deaf raters

This record makes no allegation of misconduct and no finding that any person's signed video was used for training. Statements attributed to Sign-Speak are the company's, given in writing, and NCG has not verified them independently. Terms negotiated between Sign-Speak and an individual customer are not visible to NCG and may differ from the published terms. Nothing here is legal advice.

Sign-Speak's response

Sign-Speak's chief executive, Yamillet Payano, answered NCG's seven questions in writing on 17 October 2026 and four follow-up questions on 18 October 2026, and its co-founder and chief product officer, Nikolas Kelly, answered three final questions on 20 October 2026. All three responses are reproduced below as sent, with only formatting, the email signature and the confidentiality notice removed.

7 Oct 2026
Questions sent
NCG sends Sign-Speak the pre-publication record and seven questions.

9 Oct 2026
Deadline extended
Sign-Speak asks for more time; NCG extends the deadline to 19 October and sets publication for 21 October.

17 Oct 2026
First answers
Yamillet Payano answers the seven questions.

18 Oct 2026
Follow-up answers
Payano answers four follow-up questions.

20 Oct 2026
Final answers and paper
NCG sends the sections as they will read; Nikolas Kelly answers three final questions and supplies the CHI paper.

21 Oct 2026

Published

The record publishes at 12:00 PM Central.

Response of 17 October 2026

Hello Heather,

Happy Saturday! Thank you for the answers above. Please find our answers below, and thank you for your patience:

1. Which terms governed the Syracuse pilot and the museum deployment, and which govern SignLive, SignCaption Enterprise and the API, given that the published terms restrict use to personal, non-commercial purposes?

Unfortunately terms with individual businesses are negotiated on a case-by-case basis and are behind NDA so we would be unable to share specifics with you. This is standard within commercial Agreements to have the terms sealed under NDA. However, if you have any specific questions, we would be happy to answer them as long as it does not breach any existing duties. I would additionally mention that our restrictions (Sign-Speak's usage restrictions and prohibitions on usage of data) are typically higher within these B2B deployments.

2. Was signed video from the Syracuse pilot or other institutional deployments used to train Sign-Speak's models, and if so, what did signers consent to, in what language, and can they withdraw it?

To date, signed videos from both Syracuse pilots, institutional deployments, and SignCaption have not been used to train our models. Our current language does provide mechanisms to withdraw consent if the users email us, and for several of our deployments data is automatically disposed. I should additionally mention that we do not make efforts to re-identify users unless they provide specific user labels (e.g. within SignCaption) so we tend to ask for details of when users interacted with the systems so we can wipe out data accordingly.

3. What deletion route exists for signed video, and which AI service providers receive it?

If it is through SignCaption, we have a dedicated account deletion mechanic which deletes data accordingly. Otherwise, users can always email us. Us and any down-

stream processors (most notably AWS, along with any other processors) have the data removed. Individuals can additionally always email us at contact-us@sign-speak.com if they have a bespoke data deletion request.

4. What does "Deaf-vetted for accuracy" mean in practice, is there measured accuracy in each direction, and will the Syracuse findings be published?

We are certainly happy to ask Syracuse for rights to publish a case study; in practice Deaf-vetted means that we ran case studies and QA with Deaf folks, and designed it with Deaf people; many team members are additionally members of the community itself. For more information, we recommend seeing our peer-reviewed paper in CHI https://doi.org/10.1145/3706599.3706678. This is only a small subset of the studies which we've done (as we tend to run similar studies to these whenever we do model improvements).

5. When is a human interpreter brought in, who are they, and is a human ever required rather than offered?

Human interpreters can be requested at any time via a button press. With a single button, it rolls over to VRI (human interpreter). This is true for every one of our AIS deployments (every deployment that we have had has required a human interpreter option); this is provided via various VRI partners, and Sign-Speak pays the difference, and importantly, it's always provided to the end-user free of charge.

6. Below the level of ownership, can a Deaf role at Sign-Speak halt a release or deployment, and are community members who contribute signing compensated?

Members who contribute signing are compensated, we have a vibrant community. Most members of our team (Deaf people included) have access to various "break-glass" mechanisms to stop deployments, especially at the product and engineering level. Several Deaf people are additionally in roles who have explicit authority to halt a release or deployment. We don't have a "Deaf role" in our company as we believe that folks should be integrated throughout the org (and Deaf people hold a variety of roles) rather than just relegated to an advisory board. I'm proud to mention that we have paid millions of dollars back into the Deaf community!

7. What route does a Deaf user have, in ASL, to report a wrong rendering, and who is responsible for a wrong rendering delivered through a third party's integration?

Great question! We have /feedback which developers can hook up to raise issues on rendering or SLR accuracy. For AIS we additionally have mechanisms to specifically provide feedback, and we are always happy to hear and respond to issues (contact-us@sign-speak.com) to ensure we can further our mission to provide the best service to the Community. Additionally, via intercom, we have previously hopped onto ASL video calls with different folks to remedy any issues which arise from using any of our platforms. We have ASL support available within any customer service hours.

General Notes:

I should mention that all commercial deployments of our automated interpreting systems have an interpreter option provided, as we believe in a collaborative approach.

Our primary training data source is sourced from Deaf individuals who are compensated for their data and time. As mentioned above, Data from these deployment's (I believe you mentioned pilot data for example) have not been used for training purposes.

We did publish a peer reviewed paper https://doi.org/10.1145/3706599.3706678 (we've only gotten better from here!) on our metrics and evaluation methodology.

I should additionally mention that our websites, ToS, and Privacy Policies are all currently undergoing a revamp, though we are trying our best not to rush as we want to update them properly. For feedback and data-deletion, our email is always open at contact-us@sign-speak.com

We are more than happy to hop onto a call to answer any additional questions you may have, and thank you for your patience again!

Best,

Yamillet Payano, Sign-Speak CEO and Co-Founder

Response of 18 October 2026

Good Evening Heather!

Here are some of the responses to your new set of questions:

1. Terms. I understand the contents of your customer agreements are under NDA. Without disclosing their contents, can you confirm that the Syracuse pilot, the museum deployment, SignLive, SignCaption Enterprise and the API each operate under a

separately negotiated agreement rather than the published Terms of Service of 9 April 2024? If self-serve API developers accept the published terms, please say so.

Yes, that is correct.

2. Training data consent. You note that your primary training data comes from compensated Deaf contributors. What do those contributors consent to, is that consent presented in ASL, English or both, and can they withdraw it?

The consent is first presented in English, then a Deaf member hops onto a video call and explains it in ASL. As we compensate individuals, in order for them to withdraw consent, we require that they "pay back" money given in consideration for that data (e.g. if I paid you $600 for you to sign something, and then you want to withdraw consent, we require that they remit the money paid for that data since). There are exceptions of course which are processed on a case by case basis. For clarity, we also do not share or sell this data due to privacy concerns raised by the contributors.

3. Accuracy. The record will cite your CHI paper (https://doi.org/10.1145/3706599.3706678). Please confirm whether it reports measured accuracy in both directions, ASL to English and English to ASL, and if only one, which.

Yes, accuracy is measured and reported in both directions.

4. Revised documents. The record will cite your Terms of Service and Privacy Policy of 9 April 2024 and note that a revision is underway. If a revised version is published before 21 October, please tell me and I will cite it with its retrieval date.

Unfortunately revised documents are still underway, but we will notify you once they are updated!

Let me know if you have any other additional questions.

Best,

Yamillet Payano, Sign-Speak CEO and Co-Founder

Response of 20 October 2026

Hello Heather,

Thank you for sending this, we appreciate you soliciting clarification during this process and your patience while we answered. Per your questions:

1. Self-serve API. Your 18 October answer confirms that the named deployments and products run under separately negotiated agreements. It does not say which terms bind a developer who signs up without one, so the record notes that as unclear.

For pretty much all of our clients, we end up negotiating a bespoke agreement to tailor and fit to their specific needs, as well as lay out various restricted use cases and protections for the end consumer.

2. The CHI paper. The record cites it, but NCG could not retrieve the full text before publication and says so. If you can send a copy by 9:00 AM tomorrow, I will check the both-directions statement against it.

The paper is attached to this email. I would like to clarify that while the numbers were self-reported, they were peer-reviewed by CHI reviewers and therefore passed scrutiny by researchers in the field. Please note that while both our Avatar and SLR have improved since publication, we prefer to only include peer-reviewed research in this record.

3. Consent withdrawal. The record reports, in your words, that a contributor who withdraws consent is required to repay what they were paid, with exceptions case by case. If the contributor agreement itself states that condition, or states it differently, please tell me.

Your reporting is correct, contributors can either repay the money paid or can discuss with us an exception.

Best,

Nikolas Kelly, Co-Founder and Chief Product Officer, Sign-Speak

Frequently asked questions

+ **What is Sign-Speak?**

  Sign-Speak Inc. is a Rochester, New York company that describes itself as Deaf-owned and Deaf-led. Its products recognise ASL and render it as spoken or written English, render English into ASL through a digital avatar, and caption ASL video. It

sells SignCaption, SignLive for live events and an API and SDK for integration into other platforms.

+ Where has Sign-Speak been used?

Syracuse University ran a campus-wide pilot from 1 November to 16 December 2025, and the National Susan B. Anthony Museum and House in Rochester began offering it to visitors in September 2025. No outcome from either has been published.

+ What does Sign-Speak say its product should not be used for?

Section 5.1 of its published terms excludes health diagnostics, medical emergencies and safety-critical equipment, decisions on eligibility for financial or lending services, and any interaction with police or other law enforcement.

+ Does Sign-Speak train its AI on users' video?

Its published terms and privacy policy allow uploaded content, including video, to be used to train its own AI models under a licence described as perpetual and irrevocable, and say personal information is kept in perpetuity. In written answers to NCG in October 2026, the company said signed video from the Syracuse pilot, its institutional deployments and SignCaption has not been used for training, and that its training data comes primarily from Deaf contributors who are paid. Its privacy policy says the third-party AI providers it uses are not allowed to train on users' data.

+ Can a contributor take their signing back out of Sign-Speak's training data?

Sign-Speak says a contributor can withdraw consent, but is required to repay what they were paid for the data, with exceptions decided case by case.

+ Does Sign-Speak use human interpreters?

Sign-Speak says every deployment of its automated interpreting includes a single button that rolls a user over to a human interpreter through video remote interpreting, provided by partner agencies at no cost to the user. The human is offered rather than required, and the interpreters' certification and response times are not stated.

+ What should an institution ask for before using Sign-Speak?

The agreement that will govern its deployment, with the company's written answers built into it: no training on its users' signed video without separate informed consent in ASL, a deletion route with a time limit, and a named list of processors. It should also ask for a written definition of the human interpreter route, a statement of who answers for a wrong rendering, and measured accuracy evidence in both directions before the tool is relied on where access is legally required.

+ What is the SLAT Index?

The SLAT (Sign Language Access Trust) Index is the framework Novara Consulting Group uses to review sign language AI providers. Its method is published at https://www.slatindex.org/framework/.

Sources

1. Sign-Speak Inc., Terms of Service (End User License Agreement), last revised 9 April 2024, retrieved 7 October 2026. https://sign-speak.com/legal/customer
2. Sign-Speak Inc., Privacy Policy, last updated 9 April 2024, retrieved 7 October 2026. https://sign-speak.com/legal/privacy
3. Sign-Speak Inc., homepage, For Enterprises, Team and Contact Us pages, retrieved 7 October 2026. https://www.sign-speak.com/

4. Amazon Web Services, "Sign-Speak builds with AI on AWS to create accessible experiences," 1 May 2025, retrieved 7 October 2026. https://aws.amazon.com/aws-startups/learn/sign-speak-builds-with-ai-on-aws-to-create-accessible-experiences/
5. Syracuse University News, "University Piloting AI-Powered ASL Interpreting Service Sign-Speak," 12 November 2025, retrieved 7 October 2026. https://news.syr.edu/2025/11/12/university-piloting-ai-powered-asl-interpreting-service-sign-speak/
6. Syracuse University Information Technology Services, "Sign-Speak Interpreting Service," undated, retrieved 7 October 2026. https://itsaccessibility.syr.edu/sign-speak-interpreting-service/
7. WXXI News, reporting on Sign-Speak at the National Susan B. Anthony Museum and House, 12 September 2025, as reproduced by CO/AI, retrieved 7 October 2026. https://getcoai.com/news/ai-sign-language-platform-debuts-at-susan-b-anthony-museum/
8. CenterState CEO, "Syracuse Surge Accelerator: Meet the Teams: Sign-Speak," 10 January 2024, retrieved 7 October 2026. https://www.centerstateceo.com/news/syracuse-surge-accelerator-meet-teams-sign-speak
9. VCBacked, company profile for Sign-Speak, retrieved 7 October 2026. https://www.vcbacked.co/company/sign-speak
10. Yu, Z., Wilkins, N., Payano, Y. and Kelly, N., "Evaluating an AI Bi-Directional System for Communication Between Deaf and Hard of Hearing Individuals and Hearing Persons: A Pilot Case Study," Extended Abstracts of the CHI Conference on Human Factors in Computing Systems, April 2025, full text supplied by Sign-Speak on 20 October 2026. https://doi.org/10.1145/3706599.3706678
11. Sign-Speak Inc., written responses to NCG from Yamillet Payano, chief executive, 17 October 2026 and 18 October 2026, and from Nikolas Kelly, co-founder and chief product officer, 20 October 2026, published in full above.
12. Novara Consulting Group, SLAT-SNAP-SIGNSPEAK v0.3, public-evidence review, July 2026 (internal).

How to cite

Grizzle, H. M. (2026, October 21). Sign-Speak: Deaf-owned at the top, consumer terms underneath. Novara Consulting Group. https://www.novaracg.com/2026/10/21/sign-speak-deaf-owned-at-the-top-consumer-terms-underneath/