

This document was translated from English by machine and may contain mistakes. The English original is at <https://www.novaracg.com/2026/10/07/umis-sign-the-app-is-launching-the-evidence-is-not-yet/>.

/sign da UMI: O Aplicativo Está Sendo Lançado, as Evidências Ainda Não

Heather Grizzle · October 7, 2026

UMI's /sign opens to testers October 15 and is offered to hospitals with no human in the loop. Its own code says the translation models are not live.

Conteúdo

1. A different kind of company
2. The code tells the truth, and that is the problem
3. Whose hands, and on what terms
4. What the scores can see
5. The emergency room question
6. Privacy without a policy
7. Authority still forming
8. O que um comprador deve perguntar
9. The questions put to UMI
10. Fontes

A sign language app opens to testers next week, and its own code says the translation models are not live yet



EM RESUMO

- UMI's /sign, a two-way app between Deaf and hearing people, opens a TestFlight beta on October 15 and launches on the App Store on November 15. UMI markets it to hospitals, with "no human in the loop."
- UMI's own public repository says the translation models are not yet live on the network meant to produce them, and lists consented data supply and evaluation metrics as unfinished.
- The only measures named are character and word error rates on English text. No Deaf reviewer, data contributor or advisor is named, and no privacy policy is linked.
- UMI deserves credit for publishing its code and its open gates. NCG has put seven questions to the company and will publish its answers as an addendum.

"A Deaf patient arrives at the ER." The sentence is not NCG's. It is the opening of a scenario on UMI's own website, on the page for /sign, an app that promises "real-time two-way communication between Deaf and hearing people" through an ordinary phone camera. On the same page, under a heading for hospitals, clinics and practices, UMI invites providers to "Bring /sign to your patients as an early clinical site." Asked whether a person is involved anywhere in the exchange, the page answers plainly: "No. The conversation is fully automated, both directions, with no human in the loop."

Picture what that means in the room. A Deaf patient, frightened and in pain, signs about her symptoms to a phone propped on a counter. The phone speaks English to a nurse who does not sign. The nurse answers aloud, and the phone renders the answer back as captions and signing. If the app gets the dosage, the allergy or the timing of the chest pain wrong, nobody in that loop is fluent enough to notice. That scene is an illustration, not a reported event, and NCG knows of no case in which it has happened. It is built from UMI's own words.

UMI announced /sign on October 2, 2026. Its TestFlight beta opens on October 15, and the App Store launch is set for November 15. NCG is publishing this analysis now, before the beta opens, rather than after, because the decisions it bears on are being



made this month. NCG has put seven questions to UMI in writing, with a deadline of 5:00 p.m. Pacific time on October 14. UMI's answers will be published in full as an addendum to this piece. Nothing here alleges misconduct, offers any view on investment, or constitutes legal advice.

A different kind of company

UMI, which styles itself umi inc., is unlike the other providers NCG has reviewed. Signapse and Kara Technologies build their systems in-house and sell them to institutions. UMI builds its translation models through a decentralized network, Bittensor Subnet 78, in which outside teams, called miners, submit models and automated validators score them. Co-founder and Head of AI Koyuki Nakamori, formerly head of AI at the OpenTensor Foundation behind Bittensor, describes it in the [launch announcement](#) as "a King of the Hill dynamic: the best models rise, and every score is public." The company's [home page](#) also links to a Coinbase price page for a UMI token. NCG offers no view on the token. It notes the structure because a network that pays for models by automated score, and a company whose public materials point to a traded asset, both create incentives a buyer will want to see balanced by someone with the authority to say a model is not good enough to release.

The code tells the truth, and that is the problem

The most useful document UMI has published is not its marketing. It is the README of its [public code repository](#), which is unusually honest. "SN78 is active on mainnet," it says, "but UMI translation weights are not." In plain terms, the network meant to produce and rank /sign's translation models is running, but no translation models are live on it yet. The README lists what is still open, including "consented challenge supply," "metric and canary studies" and "validator economics." The [/sign page](#) agrees: "The gesture-recognition layer that completes the loop is the layer we are training now and preparing to launch," and "Verified benchmarks will be published in the technology whitepaper as the pipeline is built." The page's speed figure, "under three seconds end-to-end," is described as a "design target."

That candor deserves credit. Few companies in this market publish a list of what they have not finished, and most would not have let a sentence like "translation weights are not" active stand in public a week before launch. But read alongside the calendar, the candor is the story. On the company's own account, the translation layer is still being



trained, the consent arrangements for its data are not in place, and the way it will measure accuracy has not been settled, while the app opens to testers on October 15 and is offered to hospitals now.

FIGURE 1

The launch clock, and what UMI's own code says is unfinished



Figure 1. The launch clock, beside what UMI's own repository lists as unfinished.

Whose hands, and on what terms

Every sign language model is built from people signing. UMI's public materials say little about who those people are. The home page refers to datasets "generated through its products and research" and to "permissioned, validated and appropriately governed data." The repository refers to "public and revealed batch material" and to the unfinished "consented challenge supply." The only dataset NCG found named anywhere is on UMI's [SignLab](#) page, a nine-clip, 101.5-second inspection set that the page itself says is "not a measure of training scale," where one clip is credited to "FLEURS-ASL / Google."

The [/sign](#) page says the product is "Built with the Deaf community, not around it," and refers to "paid contribution work, a genuine voice and a real partnership." Those are the right words. What sits behind them is not public: who the signers in the training and evaluation data are, what they agreed to, what "paid contribution work" pays and on what terms, and whether a signer who changes her mind can take her recordings back. On a decentralized network, where models are trained and scored by te-



ams outside the company, the last question is harder than usual, and it deserves a written answer before anyone signs into an app built on those hands.


What the scores can see

UMI's promise that "every score is public" sounds like accountability, and in one sense it is. The question is what the score measures. The only measures the repository names are "exact CER/WER," character and word error rates. Both compare the English text a model produces against a reference sentence. They can tell you whether a model wrote "chest" where the reference said "chest." They cannot tell you whether a Deaf person understood the signing the app produced in reply, whether the meaning of a signed sentence survived when the words did, or whether a model that scores well on clean reference video holds up on a frightened patient signing fast in bad light.

Those are the things that matter in an emergency room, and they can only be judged by people fluent in the language. No public UMI material mentions Deaf reviewers, native signers or human evaluation of any kind, and the "defined benchmarks" the launch announcement says validators score against are not named. A public leaderboard of error rates, however transparent, is a measure of the text. It is not yet a measure of communication.

FIGURE 2

What an error-rate score can see, and what it cannot



Source: README, github.com/Umi-BitSign/umi ("exact CER/WER"), retrieved October 7, 2026. The bottom row is NCG's interpretation.

Figure 2. What an error-rate score can see, and what it cannot. The bottom row is NCG's reading.




The emergency room question

UMI says, in its launch announcement, that /sign is "not as a replacement for qualified interpreters where one is required or appropriate," and on the /sign page that it is "built for the everyday and on-demand moments where no interpreter is available" and "expands access rather than removing the certified interpreter where one is required." That is the right line to draw. The difficulty is that the same page draws the emergency room inside it, and describes the product as "built to support effective-communication obligations under the ADA and Section 1557."

An emergency room is exactly the place where an interpreter is most likely to be required and least likely to be immediately available, which is why a fully automated tool will be tempting there. UMI's pages do not say who decides, in that room, that an interpreter "is required." They do not say what happens when the translation is wrong, who is responsible for the error, or how a Deaf patient reports it. Whether a fully automated system can meet a hospital's legal duty of effective communication is a question for lawyers and regulators, and nothing here answers it. But a hospital considering /sign should not let the software, or the vendor's marketing, make that call.

FIGURE 3

UMI's emergency-room scenario, and the question it leaves open



Source: umi.vision/sign, retrieved October 7, 2026. Nothing here is legal advice.

Figure 3. UMI's emergency-room scenario, and the question it leaves open.

Privacy without a policy

The /sign page describes the product as "Architected to support HIPAA and GDPR-grade privacy," running "over end-to-end encrypted streams." NCG could not find a privacy policy, terms of service or data retention statement linked anywhere on



umi.vision, and nothing says whether UMI will sign the business associate agreements HIPAA requires of vendors handling patient information. The repository says its public observer feeds never expose "raw video or private consent data." What it does not say is whether video of the people using /sign, Deaf patients included, is kept, for how long, where, whether it trains anything, and whether any of it reaches the miners and validators on a network that is open by design. "Architected to support" is a statement about intent. A hospital needs a contract.

Authority still forming

The /sign page carries the heading "Advisory board forming" and says: "We are assembling clinical, accessibility and community advisors, and we'll name them here as they join." The [leadership page](#) names three people: Michael Parker, co-founder and chief executive; Ms. Nakamori; and a pseudonymous lead development partner. None is identified as Deaf, as a fluent signer or as an interpreter. That is an absence of information, not a finding about any person, and it may change as the board forms. For now, nothing public describes a Deaf role at UMI with the authority to require a change to a model, or to stop one from shipping.

That is the question NCG has asked of every provider it has reviewed. Signapse [told NCG in September](#) that Deaf native signers sign off on every translation in its reviewed product before release, though it has not said the same of its real-time output. UMI is at an earlier stage: it is launching to the public, and to hospitals, before it has named a single Deaf person with a say over what the network releases.



FIGURE 4

What UMI has put on the record, and what it has not

On the record

- ✓ Its code, published under an open license
- ✓ A candid list of what is not finished
- ✓ States that /sign does not replace qualified interpreters where one is required
- ⚠ "Built with the Deaf community"; "paid contribution work" (terms not stated)
- ⚠ Advisory board "forming"; no members named

Not found

- ⚠ Named training and evaluation datasets, beyond one sample clip
- ⚠ Who the signers are, what they agreed to, how they are paid
- ⚠ Any accuracy figure, or any Deaf evaluator
- ⚠ A privacy policy, terms of service or retention rule
- ⚠ Any Deaf role with authority over release

Sources: umi.vision (home, /sign, SignLab, Leadership) and github.com/Umi-BitSign/umi, retrieved October 7, 2026.

Figure 4. What UMI has put on the record, and what it has not.

O que um comprador deve perguntar

A buyer considering /sign, and above all a hospital or clinic, should ask for four things before any use:

- Measured accuracy for the exact model it will receive, from an evaluation that includes Deaf native signers judging both directions, published or supplied before deployment.
- The data and consent record behind the training and evaluation sets: what was used, under what license, who signed, on what terms, and how a signer withdraws.
- A privacy policy, terms of service, a retention rule for video, a statement of what reaches miners or validators, and a signed business associate agreement.
- A written rule for when a qualified interpreter must be called, with that decision resting with the clinician and the Deaf patient, not with the software.

The questions put to UMI

NCG sent UMI seven questions with a pre-publication record. In short: what model testers will receive on October 15 and how accurate it is; which datasets train and evaluate it, who the signers are, and what they were paid and agreed to; what the validators' benchmarks are and whether Deaf native signers will judge comprehension before November 15; under what conditions /sign will be used in clinical settings and who decides an interpreter is required; whether video is retained, used for training or exposed to the network, and whether UMI will sign business associate agreements; who



will sit on the advisory board and whether any Deaf role can halt a release; and how users report errors and who is responsible for them. The deadline is 5:00 p.m. Pacific time on October 14, 2026. UMI's answers will be published in full as an addendum to this piece, and any correction of fact will be made here.

Fontes

- UMI, "UMI Launches Bittensor AI Network for Motion-to-Meaning Intelligence, Starting With Sign Language," GlobeNewswire, October 2, 2026.
 - umi inc., home page, retrieved October 7, 2026.
 - umi inc., /sign, retrieved October 7, 2026.
 - umi inc., SignLab, retrieved October 7, 2026.
 - umi inc., Leadership, retrieved October 7, 2026.
 - Umi-BitSign, "umi: Trust-minimized ASL-to-English translation subnet," README, retrieved October 7, 2026.
 - NCG, "Signapse: A Stop Right, Stated for One Product," October 7, 2026.
-

Como citar

Grizzle, H. M. (2026, October 7). UMI's /sign: The App Is Launching, the Evidence Is Not Yet. Novara Consulting Group. <https://www.novaracg.com/pt/2026/10/07/umis-sign-the-app-is-launching-the-evidence-is-not-yet/>

