NOVARA
CONSULTING GROUP

Apple Could Have Bought the Avatar

Heather Grizzle · November 13, 2026

Apple put Deaf-directed human signing on Apple TV and passed on AI. Here is what that benchmark means for sign language AI procurement.

Contents

1. A film about a sit-in, and a feature that arrives with it
2. What Apple shipped
3. What Apple did not ship
4. Captions took forty years to become a button
5. What quality meant for captions, and what it should mean for signing
6. The objection a critic will raise, and should
7. The same week, the other model
8. What Apple has not told us
9. What a procurement office should take from this
10. The room where it gets decided
11. Tips

The world's most valuable consumer technology company put Deaf-directed human interpretation inside its streaming player and said plainly that AI-generated signing is not what the Deaf community asked for. Every institution buying sign language AI should read that as a benchmark, not a press release.

A film about a sit-in, and a feature that arrives with it

In April 1977, Judy Heumann and more than a hundred disabled activists walked into the federal building in San Francisco and refused to leave. They stayed for nearly four weeks. The demand was simple to state and enormous in consequence: sign the regulations that would make Section 504 of the Rehabilitation Act enforceable, so that any program taking federal money could no longer shut disabled people out. Deaf protesters were part of that occupation, and the communication inside the building ran

through interpreters and through people who refused to let anyone be left out of the room where decisions were being made. The lesson of 504 was never only that access should exist. It was that the people who need the access get to say what counts as access.

Today, November 13, 2026, Apple premieres Being Heumann, a film about that fight, and it uses the premiere to launch a new accessibility feature on the Apple TV app: built-in sign language interpretation. The timing is either a very good piece of corporate symbolism or a very good piece of corporate judgment. After reading what Apple shipped, and what it chose not to ship, I think it is mostly the second. And I think it matters far beyond streaming, because the decision Apple made this week is the same decision that transit agencies, state governments, universities, hospitals and courts are making right now, mostly in rooms where nobody at the table signs.

What Apple shipped

The feature is simple from the viewer's side. Turn on Sign Language in the Apple TV app settings, start a supported title, press the Sign Language button, and choose American Sign Language or British Sign Language. An interpreter appears in the corner of the frame, synced to playback. Titles that carry the option show a sign language icon on their page next to captions and audio description, and viewers can search for titles "with ASL" or "with BSL." Apple launched with two films, Being Heumann and F1, both in ASL and BSL, with Greyhound and Way of the Warrior Kid to follow later this year. It runs across iPhone, iPad, Mac, Apple TV and Vision Pro on the 27.2 software releases.

That is the product. The governance is underneath it, and it is the part worth studying. Apple describes the experience as designed and tested by Deaf and native signers. The interpreters are real people, working under Directors of Sign Language, the specialists who shape how a script becomes a signed performance that fits the film rather than a word-for-word relay. For Being Heumann, the sign language interpretation was directed by Daniel Durant, a Deaf actor who also appears in the film. Durant described the work as requiring care so that every translation carried the emotion and meaning of the story. That sentence is a quality standard. It says the measure of success is whether a Deaf viewer receives the film, not whether every spoken word has a corresponding sign on screen.

There is a craft tradition behind that sentence that most technology buyers have never heard of. In Deaf theatre and in film, a director of sign language, sometimes called a director of artistic sign language, does for signing what a dialogue coach and a translator do together for spoken performance. They decide how a line of English becomes a line of ASL that a Deaf audience will recognize as natural: where the emphasis goes, how a joke lands, how much space a character takes up, what the face is doing while the hands are moving. ASL carries grammar on the face and body, in eyebrow raises, mouth morphemes, head tilts and shifts of the shoulders. A signed performance that gets the handshapes right and the face wrong is not a slightly worse translation. It is often a different sentence. That is the layer a director of sign language exists to protect, and it is the layer that the newest commercial systems still struggle to reproduce reliably.

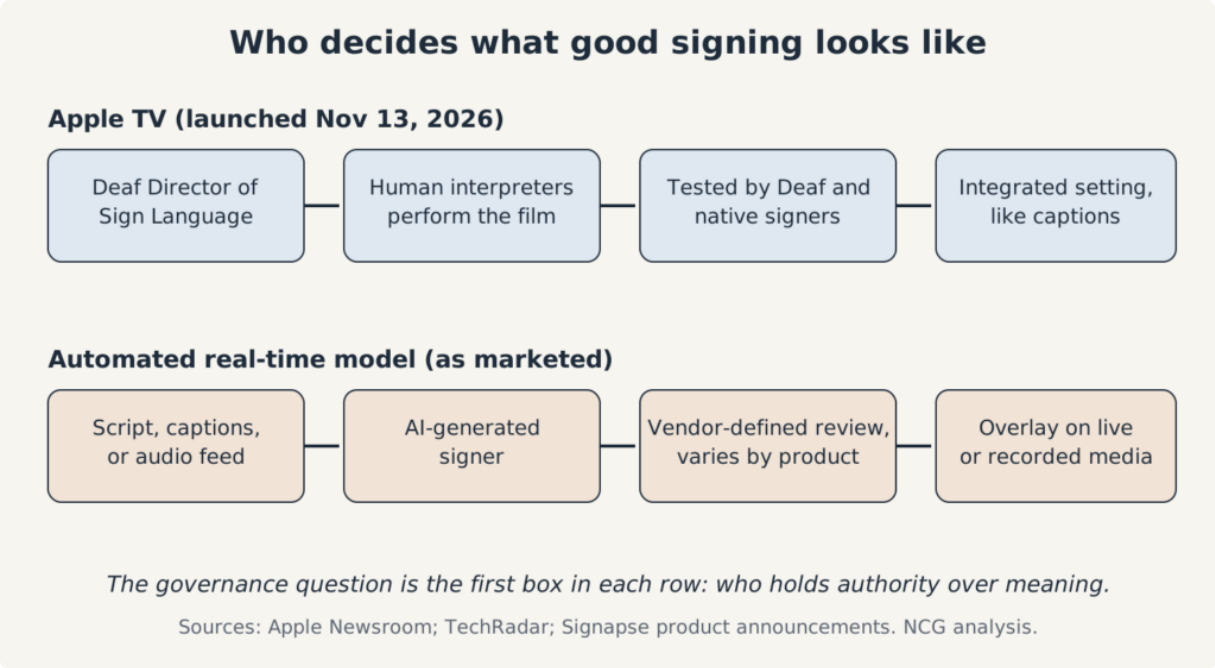


Figure 1. Two production chains for on-screen signing. The difference that matters for buyers is not human versus machine in the abstract. It is who holds authority over meaning at the first step, and who confirms comprehension before release.
Novara Consulting Group analysis

What Apple did not ship

Apple did not use AI-generated interpreters. Reporting on the launch is consistent on this point, and TechCrunch relayed Apple's reason in a form that should be printed on

the wall of every accessibility procurement office: AI-generated signing is not what the Deaf community has said it needs. TechRadar added that Apple is not, for now, using animation or AI to expand ASL coverage, and that there is no public word on whether it might later.

Consider who made that call. Apple is not a company that lacks machine learning talent, compute, video synthesis capability, or the money to acquire a startup that has all three. If any organization on earth could have built or bought a photorealistic signing avatar and attached it to its entire catalog by Friday, it was this one. It chose the slower, narrower, more expensive path, launched with two films instead of two thousand, and put a Deaf director in charge of the first one. That is a revealed preference. When the organization with the fewest technical constraints in the market picks human, Deaf-directed interpretation for prerecorded media, the argument that AI is the only practical option for prerecorded media gets much harder to make with a straight face.

It is also worth noticing what Apple did not say. It did not say AI signing is bad. It did not announce a principle that it will never use it. It said, in effect, that it listened to what Deaf people asked for and built that. That framing is more useful to a buyer than a moral position would be, because it is a method rather than a verdict. Ask the people who will use the access what they need, and let that answer set the specification. Public bodies have had that method available to them since 1977. Most of them are not using it for sign language AI.

Captions took forty years to become a button

To understand why a toggle in a settings menu matters, it helps to remember how long it took captions to get there, because sign language is now walking the same road, and the people who built captioning learned some hard lessons on the way.

Captions began as open captions, burned into the picture for everyone to see. WGBH in Boston captioned The French Chef in the early 1970s, and for years the only way a Deaf viewer could follow most programming was to find the rare broadcast that had been captioned that way, often on a delay. Closed captions arrived on American television in 1980, but receiving them meant buying a separate decoder box, so access depended on whether a Deaf household could afford and find the hardware. The turn came with the Television Decoder Circuitry Act of 1990, which required decoder chips to be built into televisions sold in the United States, so that from the mid-1990s captions were no longer an add-on a Deaf person had to go buy. They were inside the

product, a setting like volume. The 21st Century Communications and Video Accessibility Act of 2010 carried that logic into internet video, and in 2014 the Federal Communications Commission adopted caption quality standards built around accuracy, synchronicity, completeness and placement.

Read that history as a sequence and the pattern is plain. First access exists only in special versions. Then it exists as an extra the user must obtain. Then it is built into the product. Then, and only then, someone writes down what good looks like and holds the industry to it. Each step took years, and each one happened because Deaf people and their allies pushed for it.

Sign language on screen has mostly been stuck at the first stage. Signed versions of programs have tended to be separate productions, a special broadcast, a separate stream, a link to a different video, something a Deaf viewer has to go looking for. What Apple shipped this week is the jump to the third stage: signing as a setting inside the player, labeled and searchable, sitting next to captions. That is a bigger structural change than the launch slate suggests. And it raises the question captioning only answered at the fourth stage: who defines quality, and how is it measured?

What quality meant for captions, and what it should mean for signing

The FCC's four caption quality standards are a useful starting point because they were written after decades of evidence about how captions fail. Captions that are inaccurate change meaning. Captions that lag behind speech break comprehension. Captions that drop content leave gaps the viewer cannot see. Captions placed over important parts of the picture hide what they are supposed to support. Each of those failures has a direct counterpart in on-screen signing, and in most cases the signing version is harder to detect, because the person checking it often cannot read it.

Table 1. Caption quality standards and their sign language counterparts

Caption standard	What it guards against	Sign language counterpart	Who can verify it
Accuracy	Wrong words that change meaning	Meaning carried in the target sign language including facial grammar and not only handshapes	Fluent Deaf signers and qualified Deaf reviewers
Synchronicity	Text that lags behind speech	Signing timed to the scene so the viewer can watch both	Deaf viewers tested on real content
Completeness	Dropped lines and missing sound cues	Nothing summarized away; speaker changes and tone carried	Deaf reviewers comparing against the source
Placement	Text covering important visuals	Signer visible at usable size without hiding the picture	Deaf and low-vision viewers on real devices

Source: FCC caption quality standards (2014); Novara Consulting Group analysis

The last column is the point. For captions, a hearing quality checker can read the text and compare it against the audio. For signing, a hearing checker who does not sign fluently cannot tell whether the output is right. That is why the question of who directs and who tests is not a courtesy to the community. It is the only way the quality check can be performed at all. An institution that buys sign language output, human or automated, without fluent Deaf review of that output has not lowered its quality bar. It has removed the bar and kept the invoice.

Apple's launch answers the last column directly: Deaf and native signers designed and tested it, and a Deaf director shaped the first title. Whatever else one thinks of the product, that is a complete answer to the verification question, and it is more than

most public deployments of automated signing can show in their published materials today.

The objection a critic will raise, and should

There is an honest counterargument, and an article like this is only worth reading if it meets it head on. Apple is interpreting finished films. A film is fixed. It can be scheduled, rehearsed, directed, recorded, reviewed and corrected before a single viewer sees it. Live programming cannot. A city council meeting, a breaking emergency briefing, a sports broadcast, a train delay announcement composed thirty seconds before it is shown: these are the cases where interpreter supply genuinely runs out and where the vendors of automated signing make their strongest argument. Apple′s choice does not refute that argument. It does not even address it.

The second honest point is cost and pace. Human, directed interpretation of a feature film is expensive, and Apple′s launch slate shows the price: four titles named, two available on day one. A public agency with a fraction of Apple′s budget cannot simply copy the model across every video it publishes. Anyone who uses the Apple launch to argue that all sign language AI is illegitimate everywhere is overreaching, and a vendor would be right to say so.

But notice how much ground the counterargument concedes. The strongest case for automated signing is live, unbounded, time-critical content. The weakest case is prerecorded, finite, reviewable content. And a great deal of what institutions are actually buying sign language AI for sits in the second category: training videos, website explainers, recorded public service messages, onboarding modules, the same safety announcement played a thousand times. In those settings the scale excuse expires at the edit bay. There is time to bring in a Deaf director. There is time to test with Deaf viewers. The question becomes whether the buyer chose not to.

There is a subtler point inside the live category too. Much of what gets described as live is not really unbounded. A transit system announces delays, platform changes, closures and safety messages, but the set of things it says is large rather than infinite, and the most important of those messages can be translated, directed and checked by Deaf professionals in advance, then assembled when needed. An emergency management office knows today which twenty kinds of messages it will most need to send in a flood or a blizzard. Preparing those in sign language, with Deaf direction and Deaf review, before the emergency arrives, is a planning decision rather than a technology

problem. The truly unpredictable remainder is smaller than the sales pitch suggests, and that remainder is where the hard questions about automation belong.

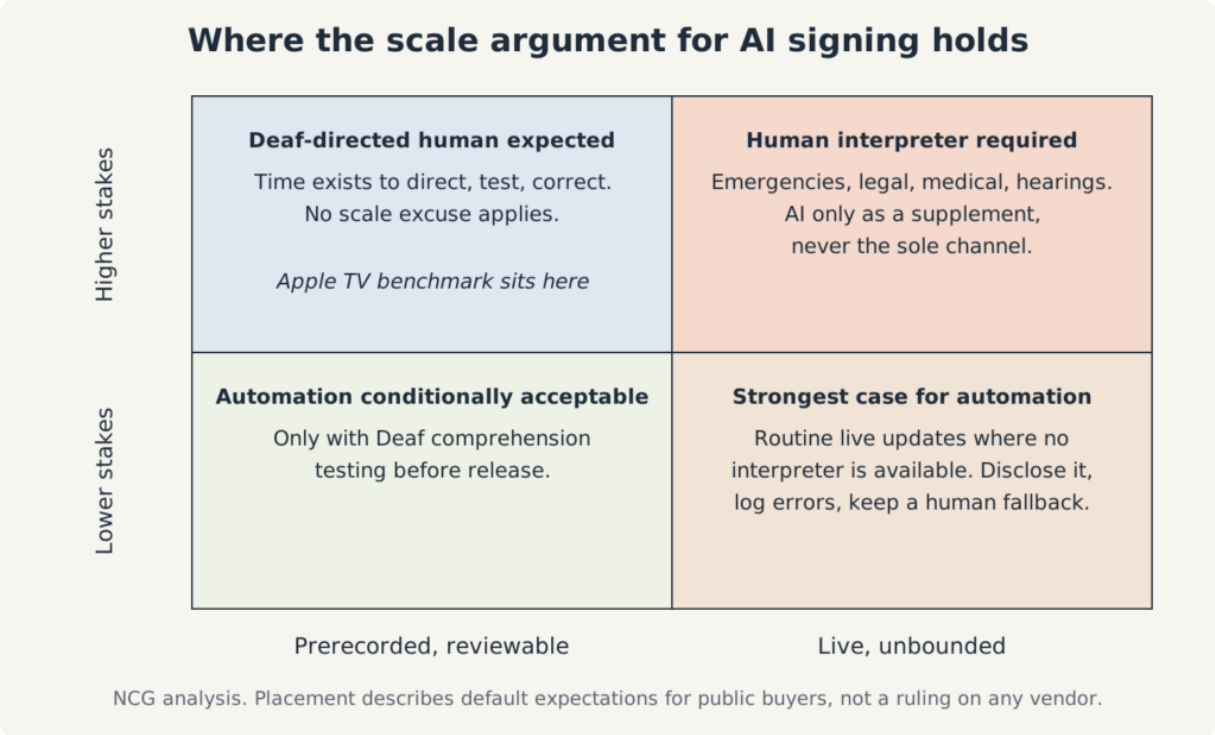


Figure 2. The scale argument is real, but it lives in one quadrant. Much of what institutions buy sign language AI for is prerecorded content, where Apple′s launch shows a Deaf-directed human default is achievable.
Novara Consulting Group analysis

The same week, the other model

The contrast is sharper because the automated model is moving fast in the same screen space Apple just entered. Signapse, the UK firm that is the most visible commercial developer of AI-generated signing, has spent the autumn pushing into live media. In late August it announced a partnership with SyncWords to bring real-time ASL translation to livestreams, alongside the subtitles and dubbing that platform already handles. On October 8 it launched SignStream Dev, a product for adding real-time sign language translation directly to live streams. Its signers are built from extensive recordings of professional interpreters, and its own materials say the digital signer is designed to complement, not replace, qualified interpreters.

Fairness matters here, and the record supports some credit. Signapse aims most of its new product at live content, which is the quadrant where the automation argument is

strongest. When NCG reviewed Signapse's public materials in September under the SLAT governance review, the company answered every question put to it and committed to a set of changes, which is more engagement than most of this market offers. The point of this article is not that one company is good and another is bad. It is that two very different answers to the same question are now on the market at the same moment, and buyers can no longer pretend there is only one.

The phrase to watch is "complement, not replace." It is the right phrase. It is also a claim that only a buyer can make true. A vendor can design a product to complement interpreters, and a procurement office can still buy it to replace them, line by line, budget by budget, because the avatar is cheaper and nobody wrote down where it was supposed to stop. A contract that says nothing about where automated signing may and may not be used will drift, over a few budget cycles, toward using it everywhere. The Apple launch gives buyers something they have been missing: a public, high-profile reference point for what the human, Deaf-directed version looks like when an organization decides it is worth doing, and therefore a reference point against which to write that boundary down.

What Apple has not told us

Credit where it is due should not become a free pass, and the launch leaves real questions open that Deaf viewers and anyone treating this as a benchmark should want answered.

The first is scale and commitment. Two films on day one and two more named for later this year is a pilot, not a catalog. Apple has not said how many titles it plans to interpret, on what timeline, or whether the feature will reach its back catalog, including the titles Deaf audiences might most expect, such as CODA, the Academy Award winning film that starred Deaf actors and was directed by Siân Heder, who also directed Being Heumann. A benchmark that covers four films is a statement of what is possible. It becomes a standard only when it is sustained.

The second is who did the work and on what terms. Apple's materials credit Deaf and native signers and name Durant for Being Heumann. They do not, in what has been published so far, describe whether Deaf interpreters performed on screen, how the interpreters and directors were credited and paid, or what process governs the titles that follow. These are not gotcha questions. They are the details that distinguish a

Deaf-led production model from a Deaf-consulted one, and the difference matters to anyone trying to copy the model.

The third is the footage itself. Apple now holds professionally directed, high-quality recordings of interpreters signing feature films in two sign languages, synchronized to scripts. In a field where the scarcity of good training data is the central technical bottleneck, that is an unusually valuable corpus. Nothing in Apple's announcement suggests it will be used for anything other than the feature it powers, and nothing here should be read as implying otherwise. But the question every institution should ask of any signed recording it commissions is the same one: who owns this footage, what were the signers told about its future use, and can it be used to train a system that replaces them? Apple's answer to that question would be as useful a benchmark as the feature itself.

What a procurement office should take from this

The first question for any sign language access purchase is now a question Apple answered in public: who directs the signing, and are they Deaf? Not who consulted, not who sat on an advisory panel, but who held authority over how meaning moved from English into a signed language for this specific content. If the answer is a model, the buyer should ask who sets and checks the model's standard of meaning, and should expect that answer in writing.

The second question is who confirmed comprehension before release. Apple says its experience was tested by Deaf and native signers. A buyer of any sign language product, human or automated, can ask for the same thing: who tested it, how many people, with what content, and what happened to the errors they found. Public evidence does not yet show that most deployments of automated signing in public settings carry that kind of documented, Deaf-led comprehension testing. Where it is missing, the buyer is accepting the vendor's internal standard as the public's standard.

The third question is which quadrant the content sits in. If it is prerecorded and finite, the buyer should be able to explain in writing why Deaf-directed human interpretation was not chosen, because the largest consumer technology company in the world just demonstrated that it can be chosen. If it is live and routine, automation may be defensible, provided the public is told what it is watching, errors are logged, and a human channel remains available. If it is live and high stakes, an AI signer should never be the only way a Deaf person receives the information.

The fourth question is whether the access is integrated or bolted on. Apple placed sign language in the player settings, next to captions and audio description, searchable and labeled. That design choice treats signing as a native part of the product rather than a separate deliverable a Deaf viewer has to go hunting for. Public bodies that publish video can apply the same test to their own sites at almost no cost.

The fifth question is about the footage. Any contract that commissions signed recordings, from a vendor or from interpreters directly, should say who owns them, whether they may be used to train automated systems, and whether the signers consented to that use. It is far easier to write that clause before the recordings exist than to argue about it afterward.

> **FIVE QUESTIONS FOR ANY SIGN LANGUAGE ACCESS PURCHASE**
>
> Who directs the signing, and are they Deaf? Who confirmed comprehension before release, and what happened to the errors? Is the content prerecorded, live and routine, or live and high stakes? Is the access built into the product or bolted on beside it? Who owns the signed footage, and can it be used to train systems?

The room where it gets decided

Judy Heumann's sit-in worked because disabled people refused to let access be defined for them by officials who would not have to live with the result. Nearly fifty years later, the same fight is happening in quieter rooms: procurement meetings, vendor demos, budget reviews where an avatar is presented as the efficient choice and nobody at the table signs. Captioning needed four decades and an act of Congress to go from a special broadcast to a button, and another two decades to get a written definition of quality. Sign language on screen does not have to take that long, because the map already exists.

Apple's decision this week does not settle the debate over sign language AI. Live content still needs better answers than anyone has yet shown, and automation will have a place in that answer. What the launch does settle is narrower and more useful. For content that can be directed, tested and corrected before anyone sees it, a Deaf-directed human standard is not a luxury that only idealists ask for. It is what the most re-

sourced company in the industry chose when it had every option on the table. Any institution that chooses less now owes the public a reason.

Tips

NCG welcomes confidential tips about sign language AI deployments, procurement decisions and vendor practices at tips@novaracg.com. For readers in the United Kingdom: writing to NCG is not, by itself, a protected disclosure under the Public Interest Disclosure Act 1998. If you are considering whistleblowing about your employer, independent advice is available from Protect at protect-advice.org.uk.

How to cite

Grizzle, H. M. (2026, November 13). Apple Could Have Bought the Avatar. Novara Consulting Group. https://www.novaracg.com/2026/11/13/apple-could-have-bought-the-avatar/