

The Record Belongs to the Signer: A Reply to Jeff Panasuik

Heather Grizzle · October 8, 2026

A reply to Jeff Panasuik on SL2T: why the signed source, not the English output, should be the record of what a Deaf person said.

Contents

1. In brief
2. What Jeff gets right
3. Where the frames meet
4. Direction, revisited
5. The precondition, and where verification should live
6. Answering the question
7. What this changes in SLAT
8. References

IN SHORT

- Jeff Panasuik is right that my analysis of SL2T documented no Deaf verification step and cited no signed-language source. I accept that without qualification.
- He is also right that direction is a property of the architecture, not only of where the system is deployed.
- Where I part from him: governance and architecture are not separate axes. Every condition he says would change his reading is a governance condition.
- My answer to his question: the safeguards an institution cannot reinterpret are the ones that belong to the Deaf person. A right to a qualified interpreter, and the signed source as the record.



In brief

Jeff Panasuik's "The Second Language Condition" ^[7] argues that Google DeepMind's SL2T removed gloss but kept the organizing logic that gloss used to expose: a written-English output target that decides what counts as success and what becomes the authoritative record. He reads my "The Coordinates of Trust" ^[5] as accurate within a governance frame and limited by it, and he closes by asking the SLAT Index which safeguard gives an institution something it cannot reinterpret, negotiate away or quietly ignore.

I accept his methodological correction without qualification, and I accept that direction is a property of the architecture rather than only of deployment. Where I part from him is the claim that governance and architecture sit on separate axes. Every condition he names as capable of changing his reading is a governance condition, and the architectural problem he identifies is answerable through one: the signed source, not the English output, must be the record of what a Deaf person said, and the right to a qualified human interpreter must belong to the Deaf person rather than to the institution buying the system.

What Jeff gets right

The methodological point lands, and I want to concede it plainly rather than manage it. "The Coordinates of Trust" analyzed a Deaf-access technology using two documents produced through a process the developer convened, facilitated and funded ^{[4][1]}. It documented no Deaf verification step, presented no independent assessment of the released model's output, and cited no signed-language source. Jeff is careful to say that he is not claiming those steps never happened, only that my published analysis does not establish that they did. That is the same distinction between absence and non-documentation that my article asked of DeepMind, and it binds my work exactly as it binds theirs. My finding that the strongest community-facing evaluation attached to a pre-release build identified a gap in the developer's record; Jeff is right that my article named the gap without closing it.

He is also right about where gloss sat in the story. Removing gloss was the strongest architectural decision in the release, and a critique that cannot credit it is not engaging with the system. His reading of the disclosed failure modes as a pattern rather than a list is the most useful thing in his article. Non-manual grammar, depth and



contact, environmental referents, the tongue, and discourse across turns are not five unrelated engineering gaps. They are the parts of a signed language that a linear written target has no place to put, and the hands, which a written target can receive, are what the system captures best.

DISCLOSURE

Jeff named his charitable role at S5 Dynamic Learning Initiative and his commercial interest as principal architect of SignaVision Solutions before making a single claim. I should do the same here: I have written papers for Jeff, and readers should weigh that relationship when they weigh this reply.

Where the frames meet


Jeff writes that governance quality and architectural direction are separate axes, and that a compliance reading can measure the first without settling the second. As a description of what a single document review can do, that is fair. As a description of how architecture actually changes, I think it is incomplete, and his own closing section shows why.

Look at the five conditions he says would materially change his reading: Deaf signers holding a majority of governance votes and budget discretion; free, prior and informed collective consent over the corpus; published data statements and model cards covering consent, compensation and non-manual coverage; assessment against binding regulation; and a qualified interpreter guaranteed in the procurement contract. Not one of them is an architectural specification. Each is a statement about who decides, who consents, what must be disclosed, which law applies and what a contract must guarantee. They are governance conditions, and they are good ones.



Where the frames meet

The five conditions Panasuik says would change his architectural reading, and the governance question each one answers.



None of the five is an architectural specification. Each changes who holds authority over the architecture, which is how an output target gets contested in practice.

Each of the five conditions Panasuik names answers a governance question, not an architectural one.
 Novara Consulting Group

That convergence is the point. An output target is not a law of nature; it is a design choice made by people with the authority to make it, funded by buyers who accept it, and kept in place by evaluation regimes that reward it. Jeff is right that a compliance checklist cannot see the corridor behind a locked door. But the only instruments that have ever redirected a corridor are the ones that change who holds the keys: decisional authority, enforceable contract terms, and evidence requirements that a vendor cannot satisfy by describing its intentions. Governance is not a different axis from architecture. It is the lever by which architecture is contested.

Direction, revisited

My original argument was that SL2T is a comprehension system: the machine reads the Deaf person and a hearing party reads the machine. When an avatar signs badly to a Deaf viewer, the harm is failed access. When a model transcribes a Deaf signer badly, the harm is misattribution, because words the person never produced can enter a written record under their name. I treated that risk as something the supported-use envelope contains for as long as the envelope holds.

Jeff's correction is that confinement does not change direction. A locked door may stop entry, but it does not change where the corridor leads, and the release materials



themselves anticipate cloud APIs, possible open weights, more languages and eventually generation. Once integrators hold the capability, the audience for the boundary is no longer the Deaf user but the employer, school, police service or benefits office deciding how to embed it. I accept this. Direction is a property of the system, and my article understated how temporary the envelope is likely to be.

What I would add is what follows for an evaluator. If direction is architectural, then every deployment of a sign-to-text system inherits the misattribution risk by default, and the burden should sit with whoever wants to put machine output into a record, not with the Deaf person who has to catch the error afterward. That reframes the governance question. It is no longer whether the vendor has drawn a boundary around dangerous uses. It is whether anything stops the output from becoming the record in the first place.

The precondition, and where verification should live

The strongest passage in Jeff's article concerns the release's own assumption that users have enough functional English literacy to verify generated text. I described that as a residual risk falling hardest on the users a sign-first tool might serve best. Jeff argues it is more than a distribution of risk: it is how the principal safeguard works, because the signer stays in the loop as verifier but must verify in written English. He cites a 2026 study of adult Deaf ASL users ^[6]. The study did not test AI translation, and Jeff says so, but it is enough to show that a safeguard requiring a signer to audit fluent English cannot be assumed equally available to everyone the system claims to serve.

THE LITERACY ASSUMPTION, IN ONE STUDY

408

adult Deaf ASL users in the sample



7.7

mean English reading grade-equivalent

3.5

standard deviation, in grade levels


Source: McKee et al., Patient Education and Counseling, 2026

I agree, and I think the remedy follows directly from his own analysis of what the system treats as the authoritative record. If the English output is the record, verification has to happen in English, and the precondition is unavoidable. If the signed source is the record, it does not. The original signed video, retained under the signer's control, becomes the evidence of what the person said, and the English text becomes what it actually is: a rendering, attributable to a machine, with no standing as the person's words until the person has confirmed it in their own language. Confirmation can then take forms that do not depend on English literacy at all, including review of the rendering by a qualified interpreter or a Deaf intermediary of the signer's choosing, and the person's right to dispute the rendering against the source.



Two answers to one question: what is the record of what a Deaf person said?

The same sign-to-text system, with the same output target, under two different record rules.



The output target is unchanged in both lanes. What changes is where verification lives: in the majority language, or in the signer's own.

The output target is the same in both lanes. What changes is where verification lives.

Novara Consulting Group

This is not a novel principle so much as a familiar one applied to a new artifact. A transcript is not testimony; the testimony is. An interpreted statement in a legal setting is checked against the speaker, not the reverse. Sign-to-text systems invert that relationship by default, and the inversion is exactly the continuity with Milan that Jeff identifies: competence in the majority language becomes the condition of being correctly represented in one's own words. Making the signed source the record removes that condition from the safeguard, even while written English remains the system's output target.

Answering the question

Jeff asks which safeguard gives an institution something it cannot reinterpret, negotiate away or quietly ignore. My answer is that no safeguard held by the institution can meet that test, because whatever the buyer holds, the buyer can waive. The safeguards that survive procurement are the ones that belong to the Deaf person and travel with them into every deployment. I would name two.



1. **A user-invocable right to a qualified human interpreter.** Not a recommendation in guidance and not a fallback the institution may offer at its discretion, but a contract term the Deaf person can invoke at any point in an interaction, with a stated remedy if it is refused and no power in the institution to waive it on the person's behalf. This answers the advisory committee's own highest-ranked concern, which Jeff notes the release materials do not guarantee. It also has a legal footing that predates AI: under Title II of the ADA, a public entity must give primary consideration to the auxiliary aid the individual requests [npg_citep key="cfr35160"], and Article 9(2)(e) of the CRPD names professional sign language interpreters among the forms of live assistance states parties must provide [npg_citep key="crpd2006"]. The choice already sits with the person in law. Procurement should stop quietly moving it to the buyer.
2. **The signed source is the record.** Machine-generated English never enters a file as a Deaf person's words unless that person has confirmed it in their own language, the signed source is retained under their control, and any dispute is resolved against the source rather than the text. This is the safeguard that answers Jeff's precondition, because it takes English literacy out of the verification path without pretending the output target has changed.

Both safeguards share a property that a vendor boundary lacks. They are not statements about what the system should not be used for. They are rights that attach to the person, so they bind whichever integrator, agency or employer ends up holding the capability, including the ones that arrive after the cloud API and the open weights.



Who holds the safeguard decides how far it reaches

A conceptual comparison, assuming the person's rights are written into law or into every procurement contract, as this reply proposes. The stages are the ones the SL2T release materials anticipate.

SAFEGUARD HELD BY	THE VENDOR'S OWN PRODUCT	INSTITUTIONS ON THE CLOUD API	INTEGRATORS' PRODUCTS	DEPLOYMENTS ON OPEN WEIGHTS
The vendor A prohibited-use statement in the release materials	●	◐	○	○
The buyer An institution's own policy, waivable by that institution	○	◐	◐	○
The Deaf person A right to a qualified interpreter, and the signed source as the record	●	●	●	●

● Binds ◐ Binds only where a contract or policy carries it ○ Does not reach

A boundary written by the vendor addresses the vendor's intent. A right held by the person goes with them into whichever deployment they meet.

A conceptual comparison. Who holds the safeguard decides how far it reaches.

Novara Consulting Group

What this changes in SLAT

Jeff put his question to my Index and to the procurement clauses in his own work, so it is fair to say what his article changes on my side. Three things.

First, the method gap. A governance reading of a Deaf-access technology that documents no Deaf verification step can describe a vendor's record but cannot independently test it. I am constituting a Deaf review body for the Index, and until it is seated I will say so on anything I publish rather than let the absence go undocumented. My further work on SL2T will carry that statement until the gap is closed.

Second, output validation. The deepest finding across the SLAT evaluations so far ^[8] is that no vendor has published an independent evaluation of its signed output. I have drafted an output validation protocol that treats Deaf assessor participation as definitional rather than additive, bars vendor-supplied or demonstration content from the sample, and reports misattribution separately from accuracy so that the failure Jeff

A boundary written in a PDF addresses the vendor's intent. A right held by the Deaf person addresses every deployment that follows.

Heather M. Grizzle



and I are both describing cannot be averaged away. It is a draft, and I am publishing its existence here so that I can be held to finishing it.

Third, the record rule. The Index already asks whether a Deaf user has a route to a qualified human. It does not yet ask whether machine output can enter a record as a person's words without that person's confirmation in their own language. Jeff's article persuades me it should, and I will propose that as a condition in the next revision of the Index. The question he closed on deserves a better answer than a sufficiency rating, and the best answer I have is a right the Deaf person holds rather than a promise the institution makes.

References

1. Google DeepMind Sign Language Team et al. (2026, August 12). AISLAC Joint Impact Report for SL2T 1.0. Google DeepMind.
2. U.S. Department of Justice. 28 CFR 35.160(b)(2), Communications: auxiliary aids and services. <https://www.ecfr.gov/current/title-28/chapter-I/part-35/subpart-E/section-35.160>
3. United Nations. (2006). Convention on the Rights of Persons with Disabilities, Article 9(2)(e).
4. Google DeepMind Sign Language Team. (2026, August 12). Putting sign language AI into users' hands. Google DeepMind.
5. Grizzle-Odland, H. M. (2026, August 22). The Coordinates of Trust: A governance reading of Google DeepMind's SL2T. Novara Consulting Group.
6. McKee, M. M., Plegue, M., Champlin, S., Hill, J., Panko, T., Buis, L. R., Sen, A., Paasche-Orlow, M. K., and Hauser, P. C. (2026). Predictors of health literacy among Deaf American Sign Language users. *Patient Education and Counseling*, 142, 109348. <https://doi.org/10.1016/j.pec.2025.109348>
7. Panasuik, J. (2026, August 26). The Second Language Condition: An Architectural Reading of SL2T. LinkedIn. <https://www.linkedin.com/pulse/second-language-condition-architectural-reading-sl2t-panasuik-m-sc--5ukic/>
8. Novara Consulting Group. (2026). SLAT Index Evaluation Standard, Version 2.0. <https://doi.org/10.5281/zenodo.21537271>



How to cite

Grizzle, H. M. (2026, October 8). The Record Belongs to the Signer: A Reply to Jeff Panasuik. Novara Consulting Group. <https://www.novaracg.com/2026/10/08/the-record-belongs-to-the-signer-a-reply-to-jeff-panasuik/>

