NOVARA
CONSULTING GROUP

This document was translated from English by machine and may contain mistakes. The English original is at https://www.novaracg.com/2026/10/08/the-record-belongs-to-the-signer-a-reply-to-jeff-panasuik/.

记录属于签署者：对 Jeff Panasuik 的回应

Heather Grizzle · October 8, 2026

对*SL2T*上*Jeff Panasuik*的回复：为什么应以手语原文而非英语译文，作为聋人所言的记录。

目录

1. 摘要
2. Jeff说对的地方
3. 框架的交汇之处
4. 方向性,再探讨
5. 前提条件,以及核实应当归属何处
6. 回答这个问题
7. 这对SLAT意味着什么改变
8. 参考文献

简要说明

- Jeff Panasuik说得对,我对SL2T的分析确实没有记录任何聋人核实步骤,也没有引用任何手语来源的资料。对此我毫无保留地接受。
- 他说得也对:方向性是架构本身的属性,而不仅仅取决于系统被部署在哪里。
- 我与他的分歧在于:治理与架构并非两条独立的轴线。他所说的每一个会改变其判断的条件,实际上都是一个治理条件。
- 我对他的问题的回答是:机构无法重新解释的保障措施,正是那些属于聋人本人的保障措施。合格口译员的权利,以及以手语源文件作为记录。

摘要

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.

我毫无保留地接受他在方法论上的纠正,也接受方向性是架构本身的属性,而不仅仅是部署方式的属性。我与他分歧之处在于:治理与架构是否处于不同的轴线上。他所举出的每一个能改变其判断的条件,都是一个治理条件;而他所指出的架构问题,恰恰可以通过一个治理条件来回应:记录应当是手语源文件,而非英语输出,记录的应是聋人所说的内容;而获得合格人工口译员的权利应归属于聋人本人,而非购买该系统的机构。

Jeff说对的地方

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.

他在注译于整个故事中所处位置的问题上也是对的。去掉注译是这次发布中最有力的架构决定,而一篇无法认可这一点的批评文章,就根本没有真正触及这个系统。他将已披露的失败模式解读为一种模式,而非一份清单,这是他文章中最有价值的部分。非手动语法、深度与接触、环境指代、舌部动作以及跨轮次的话语,并非五个互不相关的工程缺陷。它们是一种手语中那些线性书面目标无处安放的部分,而手部动作——书面目标能够接收的部分——正是该系统捕捉得最好的内容。

披露

Jeff在提出任何论点之前,就先披露了他在S5动态学习计划中的公益角色,以及他作为SignaVision Solutions首席架构师的商业利益。我在这里也应同样如此:我曾为Jeff撰写过论文,读者在权衡这篇回应时应当考虑到这层关系。

框架的交汇之处

Jeff写道,治理质量与架构方向是两条独立的轴线,一份合规性审查可以衡量前者,而不必解决后者。作为对单份文件审查所能做到的事情的描述,这是公允的。但作为对架构实际如何改变的描述,我认为并不完整,而他自己文章结尾的部分恰恰说明了这一点。

看看他所说的那五个会实质性改变其判断的条件:聋人签署者在治理投票和预算决定权中占多数;对语料库拥有自由、事先和知情的集体同意权;公开涉及同意、补偿和非手动覆盖范围的数据声明和模型卡;依据具有约束力的法规进行评估;以及在采购合同中保证配备合格口译员。这五项条件没有一项是架构规格。每一项说的都是:谁来决定、谁来同意、必须披露什么、适用何种法律、合同必须保证什么。它们都是治理条件,而且是好的治理条件。



Panasuik所列举的五个条件,回答的都是治理问题,而非架构问题。
Novara Consulting Group

这种殊途同归正是关键所在。输出目标不是自然法则;它是拥有决定权的人所做出的设计选择,是接受它的买家出资支持的,也是被奖励它的评估体系维持下去的。Jeff说得对,一份合规清单看不到锁着的门后面的那条走廊。但历来唯一能真正改变走廊方向的工具,都是那些改变谁掌握钥匙的工具:决策权、可强制执行的合同条款,以及供应商无法仅凭描述意图就能满足的证据要求。治理不是与架构不同的一条轴线,它是用来撼动架构的杠杆。

方向性,再探讨

我最初的论点是:SL2T是一个理解系统——机器读取聋人,听人一方再读取机器。当一个虚拟形象向聋人观众打手语打得很糟糕时,造成的伤害是无障碍的失败。而当一个模型把聋人签署者的话转录得很糟糕时,造成的伤害是误归责,因为这个人从未说过的话可能会以其名义进入书面记录。我当时把这种风险当作一种只要“支持性使用范围”保持有效就能被遏制的东西。

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.

我想补充的是,这对评估者来说意味着什么。如果方向性是架构层面的问题,那么任何手语转文本系统的每一次部署,默认都继承了误归责的风险,而举证责任应落在那个想把机器输出纳入记录的人身上,而不是落在事后才不得不发现错误的聋人身上。这就重新定义了治理问题。问题不再是供应商是否围绕危险用途划出了边界,而是是否有任何机制能从一开始就阻止这种输出成为记录。

前提条件,以及核实应当归属何处

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.

读写能力假设,体现在一项研究中

408
样本中的成年聋人ASL使用者

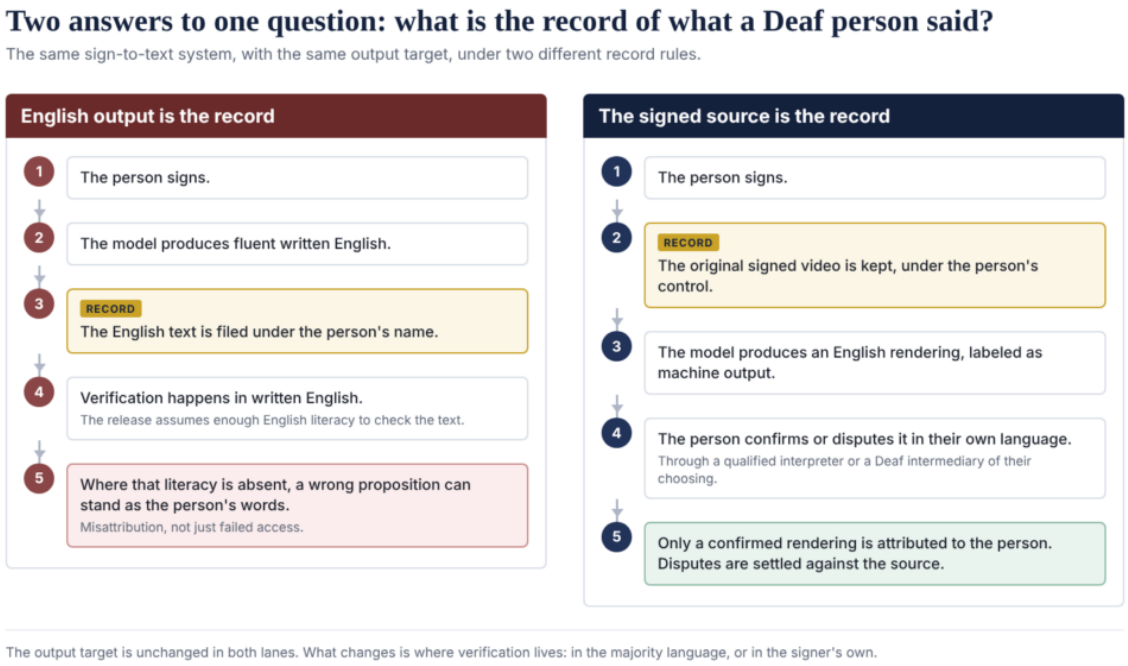
7.7
平均英语阅读年级当量

3.5
标准差,以年级水平计

来源:McKee等人,《患者教育与咨询》,2026年

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.



在两种路径中,输出目标是相同的。改变的是核实归属于何处。
Novara Consulting Group

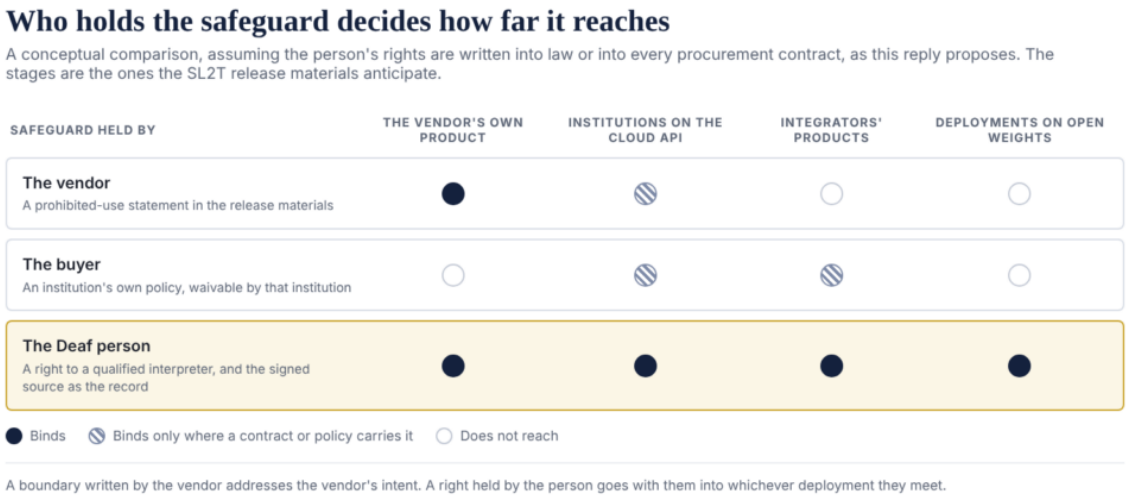
This is not a novel principle so much as a familiar one applied to a new artifact. A transcript is not testimony; the testimony is. An interpreted statement in a legal setting is checked against the speaker, not the reverse. Sign-to-text systems invert that relationship by default, and the inversion 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.

回答这个问题

Jeff问,哪项保障措施能给机构一些它无法重新解释、无法协商放弃、也无法悄悄忽视的东西。我的回答是:机构所持有的任何保障措施都无法通过这一检验,因为买家所掌握的东西,买家都可以放弃。能够在采购过程中幸存下来的保障措施,是那些属于聋人本人、并随其进入每一次部署的措施。我会列举两项。

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 [ncg_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 [ncg_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.

这两项保障措施都具备一个供应商划定的边界所不具备的特性。它们不是关于系统不应被用于何种用途的声明,而是附着于个人身上的权利,因此无论最终是哪个集成商、机构或雇主掌握了这项能力——包括那些在云端API和开放权重之后才出现的主体——这些权利都会对其产生约束力。



一个概念性的比较。谁掌握着保障措施,决定了它的效力能延伸多远。
Novara Consulting Group

这对SLAT意味着什么改变

Jeff把他的问题既提给了我的指数,也提给了他自己工作中的采购条款,因此公允地说,他的文章在我这一方带来了一些改变。有三点。

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.

> 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

第二,输出验证。迄今为止SLAT各项评估中最深层的发现[8]是:没有任何供应商发布过对其手语输出的独立评估。我已经起草了一份输出验证协议,将聋人评估者的参与视为定义性要素,而非可有可无的附加项;禁止将供应商提供的或演示用的内容纳入样本;并将误归责问题与准确性问题分开报告,以免Jeff和我都在谈论的那种失败被平均化而掩盖掉。这是一份草案,我在此公开它的存在,以便我对完成它负责。

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.

参考文献

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

TIPS

Know something about a sign language AI system that the public should know? Send it to tips@novaracg.com. Every tip is read by me and kept confidential. If you work in the UK, writing to us is not by itself a protected disclosure under the Public Interest Disclosure Act 1998, so before you share employer information, check your options with Protect (protect-advice.org.uk), the free whistleblowing advice line.

SUPPORT NCG

NCG does this work without vendor money, so our findings stay ours. If it has been useful to you, you can support NCG's independent work here. NCG is a for-profit company, not a charity. Contributions are not tax-deductible and buy no service or influence.

引用方式

Grizzle, H. M. (2026, October 8). 记录属于签署者：对 Jeff Panasuik 的回应. Novara Consulting Group. https://www.novaracg.com/zh/2026/10/08/the-record-belongs-to-the-signer-a-reply-to-jeff-panasuik/