

AISLAC Joint Impact Report for SL2T 1.0

Google DeepMind Sign Language Team
gdm-sign-language-team@google.com

Sean Forbes
Deaf Professional Arts Network (DPAN)

Nicholas Kiego
National Association of the Deaf (NAD)

Gary Behm, Joseph Riggio
Rochester Institute of Technology, National Technical Institute for the Deaf (RIT/NTID)

Yauheni Koraneu, John Moore
World Federation of the Deaf (WFD)

Regan Thibodeau
Independent subject matter expert

1. Introduction

Spoken language understanding technologies have been deployed for decades, enabling features such as voice input to mobile devices, autocaptioning for videos, and assistive applications such as Live Transcribe. A comparable input modality for sign language users has been unavailable. To address this gap, Google DeepMind has developed a sign-language-to-text model (SL2T 1.0) for translating from American Sign Language (ASL) into English, and is deploying it in two Android products: Gboard and Live Transcribe. This Joint Impact Report is co-authored by Google DeepMind team and representatives of the AI Sign Language Advisory Committee (AISLAC), an advisory body convened by Google that brings together leaders from Deaf advocacy organizations, academic and cultural institutions, and professional interpreters (including representatives from the National Association of the Deaf, RIT/NTID, DPAN, and the World Federation of the Deaf). Together, we detail the practical uses, technical limitations, and operational boundaries that govern this release. The use cases targeted with this release of Gboard are messaging, text dictation, search queries and assistant commands, and informal one-on-one conversations in low stakes situations like a coffee shop, restaurant, or retailer using Live Transcribe. As we discuss in depth below, high-stakes or multi-party conversations in medical, legal, and academic settings are explicitly out of scope, as the model has not been designed or evaluated for these. This release represents a step toward digital language equity and linguistic autonomy. The SL2T 1.0 features add to the Deaf community's portfolio of communication options, empowering signers to choose the modality that best fits their needs in a greater variety of contexts.

2. Model Overview

The SL2T 1.0 model is a transformer-based network that conditions on a sequence representing a signer's movements and generates a text translation autoregressively.

2.1. Input representation

The input representation used by the SL2T 1.0 model consists of 2D landmark coordinates for 130 key points covering the face, body, and hands extracted frame-by-frame from a video of the signer

using MediaPipe Holistic¹ (MPH). Relative to raw pixels, this landmark representation abstracts away from redundant visual detail, while discarding some contextually relevant features for sign language understanding. For instance, MPH’s face model does not track landmarks on the tongue, which is a critical articulator in ASL, and the representation also captures no information about contents of the signer’s environment, which can be used for referencing real-world entities. Further, the lack of depth ordering means the distinction between contact and hovering can be difficult to perceive. Despite these limits, this representation is efficient for low-bandwidth communication bottlenecks and provides strong privacy safeguards by replacing photo-realistic detail, physical appearance, and environmental context with abstract geometric coordinates. An illustration of what the model “sees” for a single frame taken from the FLEURS-ASL dataset is given in Figure 1.



(a) Frame from a video showing a signer.



(b) Landmark representation of the same frame.

Figure 1 | An illustration of what model “sees” for a single frame taken from the FLEURS-ASL dataset.

Although a great deal of detail is discarded, as we show below, in practice, the representation is rich enough that with adequate training, SL2T 1.0 performs very well on translation problems in a variety of domains.

Augmentation. Beyond efficiency and privacy, the 2D landmark representation is well suited for targeted data augmentation strategies during model training. Standard geometric transformations (e.g., introducing rotational noise or applying “keystone” shifts to keypoint coordinates) increase model robustness against varied camera angles and body orientations. Crucially, specific pose augmentations directly expand real-world usability for diverse signing habits: simulating one-handed signing by masking keypoints from the non-dominant arm accommodates users who hold their mobile device with one hand while signing with the other. Additionally, left-right mirroring transformations are applied during training to balance hand representation and improve accuracy for left-handed signers.

2.2. Model Deployment

For production model deployment, video processing is done on the client (e.g., the mobile device) and then the video is immediately discarded and only the landmark representation is sent to Google’s

¹<https://ai.googleblog.com/2020/12/mediapipe-holistic-simultaneous-face.html>

Table 1 | Performance of the SL2T 1.0 model on various ASL–English translation benchmarks.

Dataset	BLEU	BLEURT
FLEURS-ASL (sd-test, zero shot)	25	70
FLEURS-ASL (1-handed)	32	74
PRESTO-ASL	67	85
FSboard ⁴	64% accuracy (exact match)	

servers for processing. No logs of user inputs or system outputs are retained on Google’s side, unless explicitly authorized by the user in the context of a model evaluation study.

2.3. Benchmark Performance

The model development team measures quality across numerous benchmarks that assess various aspects of model performance. To give a sample of what the team tracks, we report metrics that capture performance in several critical domains:

- FLEURS-ASL: a translation of the FLEURS dataset into ASL by Certified Deaf Interpreters (CDIs) recorded in studio conditions. The contents are complicated, abstract language.²
- FLEURS-ASL (1 handed): a translation of part of the FLEURS dataset into 1-handed ASL by some of the same CDIs.
- PRESTO-ASL: a collection of “AI assistant” phrases performed by Deaf signers to a docked tablet device.
- FSboard: a dataset of fingerspelling from Deaf signers recorded on mobile devices.³

The results are shown in Table 1 in both BLEU, and BLEURT, and in exact match accuracy for FSboard.

2.4. Testing Methodology

SL2T 1.0 was put into the hands of real users in a series of progressive stages. Internally, the team worked with Deaf Googlers on pre-release candidates, uncovering model behaviors and user interface issues that existing benchmarks did not capture. Subsequently, an external multi-day diary study was conducted with Deaf participants to evaluate the model across a broad spectrum of everyday, low-stakes smartphone tasks in their own environments. Participants tested the input feature across five key daily scenarios:

- Search and navigation, including looking up addresses in Google Maps, searching for videos on YouTube, and browsing products on retail sites;
- Social media and messaging, including typing posts and comments on platforms like Instagram, TikTok, and Facebook, as well as chatting with friends and family in apps like WhatsApp and Google Messages;

²Garrett Tanzer, “[FLEURS-ASL: Including American Sign Language in Massively Multilingual Multitask Evaluation](#),” in Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL), 2025.

³Manfred Georg, Garrett Tanzer, Esha Uboweja, Saad Hassan, Maximus Shengelia, Sam Sepah, Sean Forbes, and Thad Starner, “[FSboard: Over 3 Million Characters of ASL Fingerspelling Collected via Smartphones](#),” in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2025.

⁴The version of FSboard we used differs slightly from the published version. We used Gemini to adapt the lowercase references to standard English orthographic conventions as well as to expand abbreviations that would be conventionally written out.

- Casual live conversations, using Live Transcribe for rapid, 1:1 back-and-forth exchanges with hearing individuals in routine, real-time interactions;
- Light productivity and notes, for example, drafting notes, brainstorming in Google Docs, and composing informal emails in Gmail; and
- Conversational AI use cases such as interacting with digital assistants and querying models like Google Gemini.

Finally, the North American AISLAC members were given hands-on access to evaluate the model in depth within their organizational and personal workflows, and their feedback directly informs the findings, mitigations, and roadmap detailed in this report.

3. Applications

The SL2T 1.0 model is designed to expand communication equity by offering Deaf signers an on-demand, sign-based text input modality. To maintain safety and translation integrity, the feature is explicitly bounded around low-stakes daily scenarios where the user maintains complete control over what is generated as they work with the tool and model.

3.1. Design Principle: Active Signer Involvement

A core design principle across both Gboard and Live Transcribe is that the signer remains directly in the loop as a verifier, and the model functions as an assistive draft-generation tool rather than an autonomous interpreter. This style of interaction is instantiated in two ways:

- **Text Preview and Editing.** In both Gboard and Live Transcribe, model outputs are presented to the user who can review the generated English text and make real-time corrections, additions, or edits using the standard touch keyboard before sending a message or showing it to their conversation partner.
- **Controlled Display in Live Transcribe.** In face-to-face conversations using Live Transcribe, the Deaf user has full elective control over when and how their translated text is displayed or voiced to the conversation partner. Output is not automatically broadcast; it is shown only when explicitly released by the user.

Note on User Prerequisites: This human-in-the-loop safeguard presumes that the user possesses functional English literacy to verify the generated text and sufficient device proficiency to operate on-screen editing controls. Where these conditions are not met, translation errors may pass unnoticed.

3.2. Primary Supported Scenarios

SL2T 1.0 is targeted toward informal, everyday digital tasks and casual interactions, including:

- Search and navigation, such as dictating locations, addresses, or queries into apps like Google Maps, YouTube, and web browsers;
- Conversational AI and assistant queries, such as Google Gemini;
- Message and document drafting, such as composing informal emails, taking personal notes, journaling, and chatting in messaging apps like Google Messages or WhatsApp;
- Elective face-to-face interactions (Live Transcribe), such as engaging in brief, two-way turn-taking during routine, low-stakes service interactions—such as ordering at a coffee shop or

restaurant, asking for assistance at a retail counter, or brief casual exchanges with neighbors and colleagues.

3.3. Prohibited and High-Risk Applications

While SL2T 1.0 offers communication autonomy for low-risk, everyday interactions, it has not been designed or validated for high-stakes or legally sensitive domains, and deploying SL2T 1.0 in these settings introduces unacceptable safety hazards and misinterpretation risks. Explicitly prohibited contexts include:

- Healthcare and medical settings, including clinical consultations, triage, emergency medical services, mental health counseling, or informed consent procedures;
- Legal and law enforcement, including police interrogations, courtroom proceedings, legal depositions, attorney-client consultations, or sworn statements;
- Academic evaluation and K-12/higher education classroom instruction, including formal educational assessments, individualized education program (IEP) meetings, or replacing classroom interpreters;
- Employment and workplace decisions, including job interviews, disciplinary hearings, performance reviews, or formal HR grievances; and
- Government and public services, including high-stakes administrative hearings or benefits determinations.

Protecting the Role of Human Interpreters

As we discuss below in more detail, a core concern raised by AISLAC members is the risk of institutional cost-cutting. Public agencies, healthcare providers, and private businesses may be tempted to opt for software solutions over live, qualified human interpreters due to convenience and lower cost. Beyond the fact that it fundamentally does not perform the complete interpretation task—it does not translate spoken language into sign language or have the ability to ask for clarifications—SL2T 1.0 lacks the dynamic pragmatic understanding, ethical judgment, cultural competence, and legal accountability of certified human interpreters. Google DeepMind and AISLAC explicitly state that this model must not be used by third parties to evade statutory obligations under accessibility laws. We do not believe that SL2T 1.0 satisfies legal obligations for reasonable accommodations under the Americans with Disabilities Act (ADA), Section 504/508, or international disability frameworks that are met by human interpreters.

4. AISLAC Evaluation & Community Feedback

4.1. Committee Governance & Evaluation Methodology

The AI Sign Language Advisory Committee (AISLAC) was established by Google DeepMind to anchor sign language AI research in the institutional expertise and lived experiences of the Deaf and Hard of Hearing community. To maintain a clear boundary between advisory oversight and core engineering, AISLAC external engagements are managed through two facilitators from Google's Accessibility organization who operate independently of GDM's and Android's research and engineering teams. For this initial release targeting ASL within the United States market, evaluation was conducted with North American AISLAC organizational representatives and independent subject matter experts (SMEs), and these likewise contributed to this report. Members were provided with Pixel test devices running pre-release builds of Gboard and Live Transcribe and tested against an early version of the

SL2T model. Advisors evaluated the technology through unstructured daily usage and structured task-based testing across personal, administrative, and professional workflows. Their collective feedback provides a consensus assessment of the model's real-world utility, linguistic accuracy, and operational limitations, and directly informs the issue triage and engineering roadmap detailed below.

4.2. Summary of Findings

Overall feedback from committee members was highly favorable. Advisors described the technology as intuitive, responsive, and potentially transformative for everyday digital accessibility.

Key Strengths

- **High Utility in Short-Form Tasks:** The system performed well for keyword search, navigation inputs (e.g., Google Maps), digital assistant queries (e.g., Google Gemini), and short back-and-forth messaging.
- **Productivity Gains:** Advisors reported strong performance when drafting notes and brainstorming within Google Docs, noting that sign-to-text significantly reduced typing fatigue.
- **Turn-Taking in Live Transcribe:** The dual-interface layout in Live Transcribe (displaying incoming speech captions on top and providing a dedicated sign-to-text drafting area on the bottom) was highlighted as a major advancement for casual, face-to-face interactions with hearing individuals.
- **Signing Modality Flexibility:** The model successfully translated variations between American Sign Language (ASL), Pidgin Signed English (PSE), and Signed Exact English (SEE), many regional sign variations, as well as both one-handed and two-handed signing.

Primary Limitations

- **Grammatical Nuance & Non-Manual Markers (NMMs):** The model recognizes individual lexical signs more reliably than full ASL sentence structures. It sometimes misses critical Non-Manual Markers—such as facial expressions, eyebrow positioning, and head nods or shakes—which convey negation, questioning, and core grammatical affect. This is particularly the case when the NMMs are not supported by hand signing or are the dominant feature of the signed phrase.
- **Polysemous Signs:** Signs with multiple contextual meanings (e.g., identical manual signs for “Sunday,” “great,” and “wonderful”) present recognition challenges, although this generally succeeds if sufficient disambiguating context is present.
- **Regional Sign Variants:** The model recognizes signs specific to particular geographic regions, but inconsistently.
- **Editing & Formatting Controls:** Users noted a lack of inline dictation commands for punctuation, line breaks, and text deletion.
- **System & Workflow Friction:** Evaluation revealed occasional application-state freezes when switching between apps, sensitivity in the floating trigger icon, and multi-step friction when alternating between speech and sign modes in Live Transcribe.

4.3. Priority Issues and their Mitigations

In accordance with the AISLAC Governance Charter, feedback from testing was classified by the AISLAC members into operational priority bands. Below are the full set of identified P0 (Critical) and P1 (High Priority) issues, along with their technical root causes and implemented or planned engineering mitigations. For each, we add a tag indicating whether the issue is addressed at the model layer ([model]) or the application layer ([application]) or elsewhere.

P0 Issues (Critical Hazards and Structural Safety Concerns)

- **P0-1: Risk of Institutional Misuse and Interpreter Substitution [other]**

- *Context:* National advocacy bodies and advisors expressed concern that public agencies, healthcare providers, or employers might inappropriately substitute automated sign-to-text tools for certified human interpreters in high-stakes settings due to cost or convenience.
- *Mitigation:* Google DeepMind and AISLAC have published this joint policy guidance stating that SL2T 1.0 is strictly an assistive input tool for low-stakes communication and does not satisfy statutory requirements for reasonable accommodations, particularly in legal, medical, or formal settings.

- **P0-2: Numerical Pattern Distortion [model]**

- *Context:* Language modeling biases caused the system to consistently distort specific multi-digit numeric sequences (e.g., systematically transcribing the Maine area code 207 as 2007).
- *Mitigation:* The modeling team is developing adversarial evaluations and training procedures to train and select models that do not exhibit this behavior. The release version of the model has improved somewhat in this regard.

P1 Issues (High Priority Linguistic & User Experience Gaps)

- **P1-1: Hallucinations and “Ghost Text” [model]**

- *Context:* Hallucinations (generation of random text not being signed) sometimes occur. Triggering contexts include when a second person enters the frame, and the signer not actively signing but pausing to organize their thoughts.
- *Mitigation:* The modeling team is working to measure hallucinations with increasingly adversarial evaluations and to address them with new training techniques. According to these metrics and internal testing, the release version of the model has already significantly reduced (but not completely eliminated) hallucinations of the kind experienced by AISLAC evaluators.

- **P1-2: Inconsistent sensitivity to Non-Manual Markers (NMMs) [model / application]**

- *Context:* Meanings expressed primarily with NMMs (with or without minimal hand signing) fail to translate. While the model demonstrates imperfect sensitivity to NMMs in conjunction with hand signing, lexemes conveyed purely with NMMs often fail to translate entirely.
- *Mitigation:* Targeted evaluations will be used to track sensitivity to NMM-mediated constructions, enabling subsequent model developments to address them, and improvements to the session management logic to prevent hands from going out of frame from preventing translation.

- **P1-3: Regional Sign & Polysemous Disambiguation [model]**

- *Context:* Localized vocabulary (e.g., regional signs for specific cities or concepts) and signs sharing identical manual forms across multiple meanings were frequently misrecognized, especially without context.
- *Mitigation:* Post-evaluation improvements to the model are addressing performance on polysemy, and regional variation will be addressed via new data acquisition efforts. Personalized word/sign lists will eventually be necessary to address some translation challenges, such as place names that share the same sign, since disambiguating context will not generally be present in the model's context.

- **P1-4: Absence of Formatting & Action Dictation Commands [model / application]**

- *Context:* Users could not dictate structural formatting commands (new line, paragraph), edit commands (delete, backspace), or action commands (send).
- *Mitigation:* A subsequent release will include dictation and formatting commands.

- **P1-5: Icon Discoverability [application]**

- *Context:* The floating trigger icon was difficult to locate for some users.
- *Mitigation:* The Gboard product team is working on discoverability enhancements.

- **P1-6: Live Transcribe Workflow [application]**

- *Context:* Alternating between incoming speech captions and sign-to-text drafting required multiple manual taps (pausing, signing, stopping, and clearing text), creating conversational friction.
- *Mitigation:* UX engineers are improving the workflow prior to the general launch.

- **P1-7: Application Lifecycle Trigger Freeze [application]**

- *Context:* Navigating away from an active app and returning caused the sign-to-text trigger to become unresponsive until the keyboard application was force-quit.
- *Mitigation:* This bug should be fixed prior to the general launch.

5. GDM Model Roadmap

In addition to the findings compiled through AISLAC advisory evaluations, Google DeepMind maintains an active, ongoing evaluation framework to map the technical boundaries of the SL2T 1.0 architecture.

5.1. Model Iteration & Evaluation Delta

It is important to note that the pre-release software builds tested by AISLAC advisors in July 2026 represented an early pre-release evaluation checkpoint. As part of our iterative development process, substantial model training and post-processing refinements have occurred between those evaluation builds and the launch candidate. As a result, several limitations identified during early testing—particularly regarding short-clip translation fidelity, one-handed sign recognition, streaming output jitter, and post-signing text generation—have been dramatically improved for the public release. Other complex features, such as explicit punctuation dictation and conversational context retention, are active priorities targeted for near-term post-launch updates.

5.2. Detailed Technical Boundaries

Since sign language understanding technology is being made generally available for the first time with this release, being transparent to users about its current technical limitations is of paramount importance. The following summarizes the GDM research team's list of limitations, which were informed by the testing processes detailed above. Most of these are subjects of active research.

5.2.1. *Vision & Perception Constraints*

Environmental & Lighting Sensitivity: Model accuracy degrades in low-light environments, under harsh backlighting or heavy directional shadows, or when signers wear clothing that reduces visual contrast against their hands and face.

Single-Subject Tracking Scope: The pose tracking pipeline currently tracks the largest subject detected within the camera frame. It does not support multi-person scene understanding, signer identification, or dynamic multi-signer turn-taking.

Anatomical & Pose Variations: Because the MediaPipe landmark model expects a standard hand structure, it may misestimate keypoints or hallucinate missing landmarks for signers with limb differences or non-standard hand geometries.

Atypical Camera Orientations: Model performance drops when the signer is partially outside the camera's field of view or positioned at extreme angles (e.g., phone resting flat on a surface or signer lying sideways).

5.2.2. *Core Translation Performance & Linguistic Nuance*

Grammatical Complexity: Linear signing or signing influenced by English structures (e.g., PSE or SEE) and signing accompanied by clear mouthings translate with higher fidelity. Conversely, complex sign language structures—such as spatial indexing, classifier predicates—are translated less robustly.

One-Handed vs. Two-Handed Signing: While one-handed data augmentation significantly improves flexibility, one-handed signs are sometimes more ambiguous, and two-handed signing typically yields higher translation accuracy, particularly on brief phrases.

Slang & Informal Sociolects: Highly informal, localized slang, or private conversational topics that are less prevalent in foundational training corpora exhibit higher error rates.

5.2.3. *Hallucinations & Streaming Prediction Artifacts*

Post-Signing Motion ("Ghost Text"): In earlier builds, keeping hands stationary in the frame after completing a phrase occasionally triggered erroneous text generation. Post-signing speech may still be erroneously transcribed via lipreading.

Autoregressive Predictive Bias: During real-time streaming translation, the model may occasionally over-predict or guess the completion of a phrase prior to the signer finishing the utterance.

Proper Noun & Acronym Expansion: The language decoder may occasionally guess expanded names or acronyms when fingerspelling or name signs cannot be definitively inferred from context.

5.2.4. Contextual Awareness & Integration Limits

Short Context Windows: Input clips are capped at 60 seconds (or when hands leave the frame). The model currently evaluates each clip independently, without retaining multi-turn conversation history or integrating the conversation partner’s messages in Live Transcribe. This limitation results in suboptimal translations in longer sessions.

Lack of Application & Environment Awareness: The model operates independently of application context—it does not adjust its output formality or formatting or change its priors on vocabulary whether operating inside a messaging app, a search engine, or a document editor.

User-Specific Personalization: The system does not support any mechanism for providing custom word lists or signs (e.g., name signs) or to set user-defined dialect preferences.

5.2.5. Dictation, Formatting & Editor Commands

Punctuation & Capitalization: Dictating explicit punctuation marks (e.g., comma, period, question mark) or capitalization rules is not yet consistently supported.

Editing & Navigation Commands: The feature does not currently support hands-free dictation of text editing actions (e.g., Delete, Backspace, Send) or document formatting commands (New Line, Tab).

Emojis: Dictating or inserting emojis through sign input is not currently supported.

5.2.6. Evaluative & Demographic Boundaries

Children’s Signing: Due to legal and privacy protections regarding children’s data, the model was neither trained nor formally evaluated on signers under 18 years of age.

Motor Diversity: Formal evaluation metrics across signers with motor disabilities or atypical movement patterns remain limited.

Demographic Representation: While the model was trained on diverse signers, benchmark disaggregation across sociolinguistic variables, such as gender, age, ethnicity, region, and across sociolects (such as Black ASL) remains an ongoing research priority.

5.3. Roadmap

To address these limitations, the GDM team prioritizes technical enhancements across three release horizons:

1. Integrated at Public Launch:

- Transition to an upgraded model that reduces translation error rates on short phrases, reduces hallucinations, and more effectively handles nonmanual signals, more complex grammatical constructions, and polysemous & regional signs.
- Stabilization of streaming translation outputs to reduce visual text jitter.

2. Targeted for Near-Term Updates:

- Rollout of dictation support for punctuation, newlines, emojis, and basic editing shortcuts.
- End-pointing fixes for more robustly handling clipping.

- Extending conversation memory across sequential clips in Live Transcribe and longer contexts in Gboard which will improve multi-turn and longer-form translation quality.

3. Long-Term Research Frontiers:

- Improved sensitivity to NMMs will appear even at launch, but addressing this comprehensively will occur over successive generations of the model.
- The current model is capable of translating a single signer and the supporting applications operate on a single signer. Subsequent releases will relax this constraint and provide support for dealing with multiple signers / multiple people more effectively.
- Expanding community-led data collection partnerships to evaluate and ensure model accuracy across regional ASL dialects, Black ASL (BASL), and other sign languages of the world.

5.3.1. Community Reciprocity and Responsible Open Research

A guiding principle of our work is that sign language AI must be developed jointly with the signing community. In creating our model evaluation benchmarks, we prioritized working directly with Deaf signers, Certified Deaf Interpreters, and CODAs. Furthermore, we recognize that true digital language equity requires supporting Deaf researchers, developers, and Deaf-owned businesses with foundational tools. Providing programmatic access to SL2T 1.0 through Google Cloud APIs, as well as exploring open-weights model releases (such as SignGemma), has tremendous potential to democratize technology and support the Deaf innovation ecosystem. At the same time, broader model availability introduces significant governance challenges. Rather than making unilateral decisions, Google DeepMind remains committed to exploring these questions in consultation with the AISLAC.

6. Conclusion

This Joint Impact Report describes the SL2T 1.0 model and two applications, Gboard and Live Transcribe, that are poised to make Sign Language AI a tangible daily reality for Deaf users, the way voice dictation for text input and conversational AI already are for hearing users. Although the model has limitations, the signer is kept in the loop whenever they choose to use this technology and thereby can ensure that their intents are communicated accurately. This opens up the potential for the experience of technology accessed for the first time in many people's native language, and for increased modalities for interaction between hearing and Deaf people in the world. The AISLAC feedback is being used by the GDM team to prioritize subsequent research directions, but also to give our users a sense for how community-led oversight shapes our work. By presenting the genuine reactions and critical perspectives of Deaf leaders across varied fields, we aim to build trust and empower users to make informed choices about how they incorporate these tools into their daily lives.