

Beyond ASL: AI for Low-Resource Sign Languages

A Research Agenda for Indian and African Contexts

Authors

Lead Researcher: Pranay M Mahendrakar

Mail: pranaymahendrakar2001@gmail.com

Contact: +91 6361723454

Company: SONYTECH

Date: April 2026

Abstract

There are over 300 sign languages in active use worldwide, serving an estimated 70 million Deaf signers. Despite this, the overwhelming majority of AI research on sign languages targets American Sign Language (ASL), with smaller but growing efforts on German, British, Chinese, and Indian Sign Language. Regional Indian variants and African sign languages remain almost entirely unaddressed. This paper argues that closing this gap is not, as commonly framed, a question of building more datasets in the same paradigm. We make three claims. First, the term "low-resource" means something structurally different for sign languages than for spoken languages — sign languages are not signed versions of spoken languages, they have no standardised writing system, they are multimodal, and within-country dialectal variation is unusually high — and naive transfer of low-resource spoken-NLP techniques produces misleading evaluations and brittle systems. Second, four bottlenecks (data scarcity, inappropriate benchmarks, the limits of cross-language transfer, and structural exclusion of the Deaf community from research design) compound rather than substitute for each other; addressing only one will not produce usable systems. Third, a Deaf-led research agenda focused on Indian and African sign languages is both ethically necessary and scientifically productive — these contexts surface problems that ASL-centric research has been able to ignore. We propose specific methodological commitments, six concrete research questions, and a strategic case for why Indian Sign Language should serve as a reference setting for the next phase of work.

Keywords

Sign language processing, low-resource languages, Indian Sign Language, African sign languages, Deaf community, participatory design, multilingual AI, accessibility

1. Introduction

The World Federation of the Deaf recognises over 300 distinct sign languages in active use worldwide, with an estimated 70 million Deaf signers (Bragg et al., 2019). Despite this scale and despite a decade of methodological progress in sign language recognition, translation, and generation, AI research has been concentrated on a small handful of these languages — overwhelmingly American Sign Language (ASL), with substantial but secondary attention to German Sign Language (DGS), British Sign Language (BSL), Chinese Sign Language (CSL), and, more recently, Indian Sign Language (ISL). The remaining hundreds of sign languages, including most regional Indian variants and almost all African sign languages, are either entirely absent from the research record or represented by isolated proof-of-concept datasets.

It is tempting to frame this gap as a familiar low-resource NLP problem and to recommend the standard remedies: more data collection, multilingual pretraining, transfer learning from high-resource sign languages, perhaps Bible-corpus-based parallel data of the kind that bootstrapped low-resource spoken-language translation. The AfriSign dataset (Gueuwou et al., 2023) explicitly takes this route. So does much of the recent ISL work.

This paper argues that the framing is wrong. "Low-resource" means something structurally different in sign languages than in spoken languages, and the techniques that worked for spoken low-resource languages either do not apply or produce systems that look good on benchmarks while failing in deployment. The path to useful sign language AI for under-served communities is not faster execution of the existing playbook; it is a different playbook.

We make three claims. First, four properties of sign languages — linguistic independence from co-located spoken languages, the absence of a standardised writing system, intrinsic multimodality, and unusually high within-country dialectal variation — distinguish them from spoken low-resource languages in ways that demand different methods. Second, the four bottlenecks holding back sign language AI (data scarcity, inappropriate benchmarks, limited cross-language transfer, and the structural exclusion of the Deaf community from research design) compound rather than substitute for one another. Third, a Deaf-led research agenda focused on Indian and African sign languages is both ethically necessary — there is a substantial recent literature making this case (De Meulder et al., 2024) — and scientifically productive, because these contexts surface problems that ASL-centric research has been able to ignore.

The contribution of this paper is a structured agenda, not new empirical results. Section 2 reviews the current state of the field. Section 3 develops the structural argument about why low-resource sign language is not analogous to low-resource spoken language. Section 4 analyses the four compounding bottlenecks. Section 5 proposes a research agenda with specific commitments and questions. Section 6 makes the case for Indian Sign Language as a strategic reference setting.

2. The Current Landscape

2.1 ASL Hegemony and Its Recent Loosening

Most foundational work in modern sign language processing has used ASL. The MS-ASL dataset (Joze and Koller, 2018) and WLASL (Li et al., 2020) became standard benchmarks for isolated sign recognition. Continuous sign language recognition and translation have been advanced primarily on the German Sign Language RWTH-PHOENIX-Weather corpus (Camgöz et al., 2018; Camgöz et al., 2020), which despite its name is German rather than American. The First WMT Shared Task on Sign Language Translation (Müller et al., 2022) was a watershed in standardising evaluation, but again focused on German Sign Language.

In the past three years this concentration has loosened. The INCLUDE dataset (Sridhar et al., 2020) provided 4,287 videos covering 263 word signs across 15 categories for Indian Sign Language. A 2025 IET survey (Damdoo et al., 2025) catalogues a substantial recent literature on ISL recognition and translation, predominantly using transfer learning from ImageNet-pretrained vision backbones combined with sequence models. AfriSign (Gueuwou et al., 2023) introduced the first multilingual dataset for African sign

languages, covering Ghanaian, Nigerian, Kenyan, Zambian, Zimbabwean, and South African sign languages, all derived from Jehovah's Witness Bible translation videos. The Sign Language Translation and Avatar Technologies workshop reached its ninth edition in 2025 (Nunnari et al., 2025), reflecting a growing institutional infrastructure.

2.2 What Is Still Missing

The expansion has not, however, reached the languages that needed it most. Within India, ISL itself is increasingly studied at the national level, but regional variants — which differ substantially in lexicon and to a lesser extent in grammar across linguistic regions of India — receive virtually no attention. Within Africa, AfriSign is the only multilingual dataset of which we are aware, and its Bible-translation provenance constrains it to a religiously narrow vocabulary that does not transfer well to everyday communicative settings. Most national African sign languages are entirely undocumented in the AI research record.

The Bragg et al. (2019) interdisciplinary perspective on sign language processing remains, six years on, an accurate diagnosis of the field's blind spots. Yin et al. (2021) made a similar argument in the NLP context with their call to include signed languages in NLP. Both papers identified the same problems we revisit here. That those problems persist is itself diagnostic.

3. Why Low-Resource Means Something Different in Sign Languages

Standard low-resource NLP techniques rest on assumptions that hold for most spoken languages. They do not all hold for sign languages.

3.1 Sign Languages Are Not Signed Spoken Languages

This point is widely acknowledged in the linguistics literature and routinely forgotten in the engineering literature. American Sign Language is not signed English; Indian Sign Language is not signed Hindi or signed English. Each sign language has its own grammar, its own lexicon, and its own historical development, often unrelated to the dominant spoken language of the same region. Translation systems that treat sign-to-text as a transliteration problem, or that assume a sign language's syntax tracks the local spoken language, produce structurally wrong outputs that may nevertheless score acceptably on token-overlap metrics. This is a core source of the gap between benchmark performance and deployment usefulness in sign language translation.

3.2 No Standardised Writing System

Spoken low-resource NLP can fall back on whatever orthography the language uses. Sign languages have no comparable default. Linguistic gloss notations exist (and are used in datasets like RWTH-PHOENIX), but glossing is contested, lossy, and not used by Deaf communities outside of academic and pedagogical contexts. SignWriting is a phonetic system but not widely adopted. The practical consequence is that the standard NLP pattern of "collect text, train a language model, then ground it" is not directly available for sign languages — there is no equivalent text corpus to begin with. Approaches that nevertheless lean on glossed corpora are training models on a representation that the target community does not use.

3.3 Intrinsic Multimodality

Sign languages convey grammatical and lexical information through manual signs (handshape, location, movement, palm orientation) and non-manual markers (facial expressions, head tilts, mouth shapes, body posture). Negation, question type, conditional mood, topic marking, and emphasis are typically expressed through non-manual channels. A system that captures only the hands is not a low-resolution sign language system; it is one that has discarded the grammar. This makes data collection genuinely harder than for spoken low-resource languages — single-microphone audio has no comparable problem — and it makes evaluation harder because systems can fail in linguistically meaningful ways that token-level metrics do not detect.

3.4 High Within-Country Dialectal Variation

Indian Sign Language exhibits substantial regional variation in lexicon — signers from Kerala, Maharashtra, West Bengal, and Tamil Nadu may use distinct signs for the same concept — even though the underlying grammar is more uniform. Similar patterns appear across African sign languages, where individual towns and schools have historically developed local conventions. This is a different distribution problem than spoken-language dialectal variation, because there is no orthographic regularisation pulling toward a standard form. A model trained on signs collected in one Indian city may generalise poorly across the country, and benchmarks that do not stratify by region overstate national-scale performance.

4. Four Compounding Bottlenecks

4.1 Data Scarcity

The numerical gap is stark. The MS-ASL dataset contains roughly one thousand signs across approximately twenty-five thousand video samples; WLASL covers two thousand

glosses. INCLUDE for ISL contains 263 word signs. AfriSign covers six African sign languages but is constrained to Bible vocabulary. For most regional Indian variants and most African sign languages, no public dataset of meaningful scale exists. Building one is expensive — high-quality sign data requires native signers, multiple camera angles, and ideally annotation by linguistically trained Deaf reviewers — and the funding model that has supported ASL data collection has not been replicated elsewhere.

4.2 Inappropriate Benchmarks

The benchmarks the field uses inherit assumptions from spoken-language NLP that do not transfer well. BLEU and chrF treat sign-to-text translation as a word-overlap problem, which understates failures of the kind described in §3.1 (correctly worded outputs with structurally wrong syntax) and overstates progress when models simply learn surface lexical mappings. Isolated sign recognition benchmarks abstract away the continuous, co-articulated nature of natural signing. Recognition-then-translation pipelines often score well on glossed corpora and poorly on naturalistic video. None of these issues are unique to low-resource settings, but they bind tighter when the underlying datasets are smaller and the room for spurious correlations is larger.

4.3 The Limits of Cross-Language Transfer

In low-resource spoken NLP, multilingual pretraining on related high-resource languages reliably helps. The analogous move in sign language — pretrain on ASL or DGS, fine-tune on ISL or a Nigerian sign language — works less well, for two reasons. First, sign languages are typically less related across regions than the spoken languages of the same regions; transfer from ASL to BSL, despite both being Anglophone-region sign languages, is harder than the spoken-language analogy suggests. Second, the visual feature representations that pretrain well on ASL video are biased toward the handshape and motion distributions of ASL signers; they need substantial adaptation for signers using different visual conventions. Cross-language transfer in sign languages is a research question, not a default tool.

4.4 Structural Exclusion of the Deaf Community

This is the bottleneck the engineering literature has been slowest to confront. The dominant pattern in sign language AI has been hearing researchers building tools intended for Deaf users, with Deaf consultants brought in late, if at all. The recent literature on this is increasingly direct. De Meulder et al. (2024) document the structural barriers to genuine co-creation in sign language technology projects and offer five lessons, including "redistribute power through deaf leadership" and "recognise and resource deaf partners"

invisible labour." Recent ASSETS-venue work has empirically documented the gap between ML practitioners' priorities and Deaf signers' priorities in sign language AI.

This is not only an ethics issue, though it is that. It is also a methodological issue: systems built without Deaf leadership routinely solve problems Deaf users do not have, ignore problems they do have, and fail evaluations that Deaf evaluators would have predicted. The engineering and the ethics are not separable here.

5. A Research Agenda

5.1 Methodological Commitments

We propose four commitments any serious project in this area should adopt. First, Deaf leadership in research design, not consultation after the fact — meaning Deaf researchers as principal investigators or co-PIs, not advisory-board members. Second, evaluation by Deaf evaluators using tasks that reflect actual communicative use, not only token-overlap metrics on glossed corpora. Third, explicit reporting of dialectal coverage in datasets and benchmarks, with stratified evaluation. Fourth, public release of data collection protocols and consent procedures, so that the methodological and ethical work is itself a research contribution.

5.2 Six Specific Research Questions

1. How well does cross-sign-language transfer actually work? A systematic study of pretrain-on-ASL/fine-tune-on-ISL versus train-from-scratch-on-ISL, controlling for total compute, would settle a question that is currently answered by intuition rather than evidence.
2. What does within-country dialectal degradation look like? Train ISL recognition models on signers from one Indian region and evaluate on signers from another. Quantify the drop. Repeat across multiple region pairs.
3. Can self-supervised video pretraining substitute for annotated sign data? Pretraining on large quantities of unlabelled signing video — feasible to collect from broadcast media — followed by limited supervised fine-tuning is plausible but mostly untested in sign languages.
4. What benchmarks would a Deaf-led evaluation framework actually look like? This is not a rhetorical question. Co-designed benchmarks for sign language production and translation that emphasise communicative success over surface fidelity are a tractable methodological contribution.

5. How much non-manual marker information do current systems actually use? Ablation studies on facial-expression and body-pose channels would reveal whether models that nominally use full multimodal input are functionally hand-only.
6. Can community-collected data, gathered with appropriate consent and remuneration, scale to the levels needed? The data-collection-as-research-question framing is itself underdeveloped in sign language AI.

5.3 Why Indian Sign Language Is the Strategic Reference Setting

Several factors make ISL particularly well suited to anchor the next phase of low-resource sign language work. The signing population is large (estimates above five million). India has substantial existing AI research infrastructure and a maturing tradition of low-resource NLP work on Indian spoken languages — the AI4Bharat initiative being the most visible example — that could plausibly be extended. The high regional dialectal variation makes ISL a natural test bed for the within-country generalisation question (§5.2.2) that any sign language AI system will eventually face. And the existence of the INCLUDE dataset, whatever its limitations, provides a non-zero starting point.

None of this argues that African sign languages should be deprioritised. They should not. But ISL provides a setting where methodological progress can be made on infrastructure and benchmarks that will be reusable across the broader low-resource sign language landscape.

6. Conclusion

The dominant framing of low-resource sign language as a scaled-down version of low-resource spoken language has produced a research culture that systematically underestimates the difficulty of the problem. It treats the absence of data as the central obstacle when in fact the absence of standardised writing systems, the multimodal grammar, the regional dialectal variation, and the historical exclusion of Deaf communities from research design are at least as important. Each of these is a research question in its own right; none of them is solved by another dataset.

The next decade of useful work in this space will, we suggest, be defined less by which architectures researchers apply and more by who is in the room when problems are scoped, and by how seriously the field engages with the linguistic and social structure of the languages it is trying to serve. Indian and African sign languages are not edge cases. They are where the field's standard assumptions break first, and

consequently where the most informative research can be done.

References

- Bragg, D., Koller, O., Bellard, M., Berke, L., Boudreault, P., Braffort, A., Caselli, N., Huenerfauth, M., Kacorri, H., Verhoef, T., Vogler, C., & Morris, M. R. (2019). Sign language recognition, generation, and translation: An interdisciplinary perspective. *ASSETS 2019*, 16–31. arXiv:1908.08597.
- Camgöz, N. C., Hadfield, S., Koller, O., Ney, H., & Bowden, R. (2018). Neural sign language translation. *CVPR 2018*, 7784–7793.
- Camgöz, N. C., Koller, O., Hadfield, S., & Bowden, R. (2020). Sign language transformers: Joint end-to-end sign language recognition and translation. *CVPR 2020*, 10023–10033.
- Damdoo, R., et al. (2025). An integrative survey on Indian sign language recognition and translation. *IET Image Processing*.
- De Meulder, M., et al. (2024). Lessons in co-creation: The inconvenient truths of inclusive sign language technology development. arXiv:2408.13171.
- Gueuwou, S., Takyi, K., Müller, M., Nyarko, M. S., Adade, R., & Gyening, R.-M. O. M. (2023). AfriSign: Machine translation for African sign languages. 4th Workshop on African NLP, ICLR 2023.
- Joze, H. R. V., & Koller, O. (2019). MS-ASL: A large-scale data set and benchmark for understanding American Sign Language. *BMVC 2019*.
- Li, D., Rodriguez, C., Yu, X., & Li, H. (2020). Word-level deep sign language recognition from video: A new large-scale dataset and methods comparison (WLASL). *WACV 2020*.
- Müller, M., Ebling, S., Avramidis, E., Battisti, A., Berger, M., Bowden, R., et al. (2022). Findings of the First WMT Shared Task on Sign Language Translation (WMT-SLT 2022). *Proceedings of the Seventh Conference on Machine Translation (WMT)*, 744–772.
- Nunnari, F., Luna Jiménez, C., Wolfe, R., McDonald, J. C., Filhol, M., Efthimiou, E., Fotinea, E., & Hanke, T. (2025). 9th Workshop on Sign Language Translation and Avatar Technologies (SLTAT 2025). arXiv:2508.08050.
- Sridhar, A., Ganesan, R. G., Kumar, P., & Khapra, M. (2020). INCLUDE: A large-scale dataset for Indian Sign Language recognition. *Proceedings of the 28th ACM International Conference on Multimedia*, 1366–1375.
- Yin, K., Moryossef, A., Forster, J., Ye, X., Tan, X., Forster, M., et al. (2021). Including signed languages in natural language processing. *ACL-IJCNLP 2021*.
- World Federation of the Deaf. (2021). World Federation of the Deaf statistics on Deaf populations and sign languages. WFD.