The SignBeach Dataset of Dutch Sign Language (NGT) signs

Annika Schiefner[1,2], **Gomèr Otterspeer**[1], **Beyza Sümer**[2], **Floris Roelofsen**[1]
[1]Institute for Logic, Language, and Computation, University of Amsterdam
[2]Amsterdam Center for Language and Communication, University of Amsterdam
{a.schiefner, g.otterspeer, b.sumer, f.roelofsen}@uva.nl

Abstract

This paper presents the SignBeach dataset, including 1401 lexical signs from Dutch Sign Language (NGT). The items in this dataset represent everyday vocabulary appropriate for primary school children and are part of a larger research project, investigating sign learning in a digital environment. Each sign is presented by four deaf signers in a controlled studio environment. For each item, high quality video recordings are available from five synchronised cameras, providing rich multi-view visual input suitable for linguistic analysis and the development of computer vision pipelines. In addition, we provide three types of computational derivatives: keypoint estimates using MediaPipe, handshape estimates using HaMeR, and 3D body reconstructions using SAM 3D Body. Signs are aligned with lexical entries in the NGT Signbank to provide interoperability of the database with other NGT resources. We outline the construction of the dataset and provide information on opportunities for reuse, for example in the context of psycholinguistic studies or in the context of sign language technology. All materials are available for non-commercial reuse under a CC BY-NC 4.0 license.

Keywords: Dutch Sign Language, NGT, sign language, lexical database, 3D-video recordings

1. Introduction

The primary contribution of this paper is a large-scale, multi-view, child-oriented lexical video database for NGT, enriched with aligned lexical identifies and multiple layers of computational derivatives, the *SignBeach* dataset. We outline the development of the dataset, describe its structure, and provide information on potential future scientific reuse[1]. The aim of the SignBeach project is to build a series of online learning games in which primary school-aged children throughout the Netherlands learn Dutch Sign Language (*Nederlandse Gebarentaal*, NGT), through simple translation games between NGT signs and Dutch words. The games are suitable for both deaf and hearing children, regardless of whether their preferred language is NGT or Dutch.

In order to select items that are suitable for this age group and purpose, we started from a book with signs that are relevant to children learning NGT, published by the *Nederlands Gebarencentrum* (engl. *Dutch Sign Language Centre*) (Nederlands Gebarencentrum, 2012). From there, we compared items to the NGT Signbank (Crasborn et al., 2020; Klomp et al., 2024) and created engaging and child-friendly recordings, linked to Signbank through alignment of ID glosses and item IDs. This process resulted in a dataset of 1401 individual lexical signs, each produced by four signers in a multi-camera setup. In addition, we provide derivative data and simple ways to link the dataset to other lexical resources of NGT to facilitate scientific reuse of the dataset.

[1]The dataset is available on Figshare at https://doi.org/10.21942/uva.c.8233876 under a CC BY-NC 4.0 license for non-commercial reuse.

2. Related work

Sign language databases differ in their purposes and thus in the types of information presented in each individual database. To clarify how the SignBeach dataset integrates with and complements existing resources, the following subsections present similar resources for NGT and other sign languages and highlight in particular how they address different research needs. We limit our review to lexical databases developed as resources for academic research that are similar to our projects in focus. Online dictionaries and materials developed in the context of language assessment tools are not covered here, as their purpose differs quite substantively from our target uses with important differences in design choices resulting from these different purposes.

2.1. NGT Signbank and related resources

Research on NGT has a long tradition within the context of sign language research as a scientific field of enquiry. NGT thus has a relatively well developed infrastructure of resources, building on the original development of the NGT Corpus (Crasborn and Zwitserlood, 2008), a collection of video materials of conversations in NGT including approximately 72 hours of recordings of which around 21% are annotated (Kopf et al., 2022). The NGT Signbank contributes a lexical database to that infrastructure,

Proceedings of the 12th Workshop on the Representation and Processing of Sign Languages, pages 450–458
16 May 2026.

providing standardised gloss-IDs for videos, alongside phonological and other linguistic information (Crasborn et al., 2020; Klomp et al., 2024). At the point of writing, the NGT Signbank includes 6863 individual public signs. By providing a set of standard gloss-IDs for individual signs with reference videos and phonological descriptions, the NGT Signbank serves as a useful connector, providing single identifiers for specific variants of signs. It is mainly designed as a resource for video annotation and can be imported as a controlled vocabulary in ELAN (2025; see also Crasborn et al., 2016). For this purpose, glosses are defined based primarily on manual forms, disregarding mouthings and other non-manual information. This, in turn, makes it a database that is suitable for looking up manual forms but currently provides very limited information on how these signs may look if produced in context with appropriate mouthings and other non-manuals. It does provide 'senses', alternative translations to Dutch and English, but these remain limited and additional information required for disambiguation is not available.

Besides typical academic licensing constraints around reuse, the videos are therefore not designed to be particularly engaging. By contrast, our project required signs targeted at children, presented in an engaging manner. In addition, in our dataset, signs are produced as translations for specific lexical items in Dutch, providing a dataset for a translation game between NGT signs and Dutch words. This requires the presence of mouthings, mouth gestures, and other non-manuals that contribute both semantic and pragmatic meaning appropriate to the game context. However, as mentioned above, the gloss-IDs of the NGT Signbank are a useful mechanism for aligning sign variants across datasets. In our dataset, we therefore identified all sign forms by their gloss-IDs, creating broad alignment with the NGT Signbank. In this way, we make sure that the variants recorded in our dataset can be easily identified and matched with signs used in other studies, opening up possibilities to bring together disparate pieces of information.

Another component of the existing infrastructure for the development of lexical datasets for NGT and other sign languages is the signCollect platform (Otterspeer et al., 2024). SignCollect provides a 'touchless' recording pipeline in which a signer can control the system through simple gestures recognized via computer vision (MediaPipe), enabling recordings of 60 to 120 signs per hour without the need for technical assistance staff. The platform integrates three key components: (i) signCollect Studio for multi-camera video recording with foot pedal or gesture-based control, automatic synchronization (using Tentacle Sync), and real-time review and approval of recordings; (ii) signCollect Dashboard for collaborative metadata management among team members; and (iii) signCollect Hub which connects the recording system to Global Signbank through automated synchronization of lexical entries. The present project makes use of the signCollect platform.

2.2. Lexical databases in other sign languages

Lexical databases are also available for a number of other sign languages. Some languages have Signbank lexicons that are similar to the NGT version and published under the database structure provided by the Global Signbank Project (Crasborn et al., 2018), for example the BSL Signbank (Fenlon et al., 2015), which streamlines annotations for the BSL Corpus (Schembri et al., 2017). Signbank databases thus primarily fulfil a lookup function for annotating video materials. In a similar vein, the German Sign Language (DGS) Corpus is linked with a number of lexical resources that serve similar purposes (Kopf et al., 2022; Otte et al., 2022; Schulder et al., 2024), as well as more public facing presentation modes that primarily serve as a language archive and resource for the deaf community (Konrad et al., 2020; Jahn et al., 2018).

Meanwhile, other lexical databases have been developed with the purpose of collecting and presenting detailed linguistic information and norms of a variety of psycholinguistic variables. These databases serve as a basis for developing psycholinguistic experimental stimuli and materials for standardised, educational or clinical tools. An early example is the lexical database for Spanish Sign Language (LSE) by Gutierrez-Sigut et al. (2016). This database includes highly detailed phonetic coding of signs and video materials that are visually optimised for neuroimaging studies and similar research that requires a high level of experimental control. Shortly after, ASL-LEX was created as a database that combines phonological information with different psycholinguistic variables, such as iconicity, familiarity and age of acquisition, measured through rating studies (Caselli et al., 2017; Sehyr et al., 2021). Initially, this resource was primarily created to allow for studying the effects of phonological neighbourhood density and creating stimulus sets for controlled psycholinguistic studies. In further iterations of the project, additional variables have been added and the database architecture has been made adaptable, creating a foundation for creating parallel databases in other sign languages (for examples, see Morgan et al., 2022; Trettenbrein et al., 2025). The database presented in this paper most closely resembles these types of databases in terms of structure and projected use. It is designed to present video materials that are recorded in a highly controlled setting alongside useful derivate data, to maximise opportunities for scientific reuse. As such, it directly addresses the need for larger, public video databases articulated by Bragg et al. (2019).

3. Dataset design

The dataset for this project combines video data, metadata on the individual signs and information for linking the materials to NGT Signbank (Crasborn et al., 2020; Klomp et al., 2024), and derivative data (see Table 1).

3.1. Item selection

Items were initially selected from the book "Mijn eerste 1500 gebaren" (engl. *My first 1500 signs,* Nederlands Gebarencentrum, 2012) and referenced against the NGT-Dutch online dictionary (Nederlands Gebarencentrum, 2024). These items were then compared against the signs available on the NGT Signbank (Crasborn et al., 2020; Klomp et al., 2024) and if the variant in the book was not available on Signbank, a different variant was chosen from Signbank. If no corresponding item was available on Signbank, the sign as produced in the NGT-Dutch online dictionary (Nederlands Gebarencentrum, 2024) was recorded by the first author as a reference for the recording sessions. This strategy was also employed for NGT compounds, for which there are typically no unique items on Signbank. During the recording sessions, signers were able to provide additional input to ensure widely used variants were favoured in the recording process. Where signers disagreed with the proposed variant, we recorded both the initial variant and the proposed variant and checked with the other signers to select the final variant. Wherever this process resulted in an alternative variant being chosen, the corresponding item on Signbank was identified in post-processing. The items are thus aligned with Signbank glosses and item IDs, which can be resolved into the URL leading to each sign.[2] After collecting the dataset input list, we transferred the items into lists in signCollect (Otterspeer et al., 2024) and recorded them in individual recording sessions with the signers.

3.2. Linking the dataset to other lexical resources

As mentioned above, our dataset is aligned with the NGT Signbank through ID glosses and sign IDs. The sign IDs can be used to derive the URL leading to the relevant Signbank entry. In our dataset, we have therefore opted for directly including the URL, to facilitate access to Signbank. This alignment with Signbank allows for easy retrieval of linguistic information about the individual items, such as phonological annotations, as well as alignment with any other resource that is linked to Signbank IDs.

It should be noted that Signbank items are presented without non-manual information, meaning that the same manual form may be associated with multiple meanings. Multiple items in our dataset can thus be based on the same Signbank item, resulting in the same gloss and Signbank ID (1401 items:1297 distinct glosses, including compounds). We are therefore using the Dutch translations to identify signs in our dataset, as glosses and Signbank IDs are not unique to individual items.

4. Recording video dataset

4.1. Signers

Each sign is presented by four distinct, deaf signers, two men and two women. The signers were professionals with ample experience in sign language video recordings, who reported using NGT regularly in their everyday lives, including in professional contexts. Each of the initial four signers recorded all 1401 items for the dataset. However, due to technical issues in the items from one recording session that could not be re-recorded with the same signer, two additional signers joined the project to create new recordings for these items for the final dataset. In the final recordings, each sign is recorded by two male and two female signers. The signers were reimbursed for their participation and contributed to the project as professional contractors, rather than as research participants.

All signers were chosen for their experience recording materials for children. They are deaf professionals, well-versed in working in a recording context and produced engaging videos of the lexical items. They were asked to sign in a child-friendly manner and to produce mouthings and mouth gestures as they felt was appropriate for the concepts. Overall, the signs are recorded in a friendly and expressive manner, making them feel more natural than typical recordings in the NGT Signbank. However, it should be noted that affect is influenced by the semantic content of the signs recorded (e.g. the sign for BE CAREFUL is signed with a stern face across signers) but probably also by signers' fatigue over long recording sessions. We varied the order of items across signers and recorded over multiple sessions to alleviate issues of fatigue and the resulting recordings feel natural.

[2]Note that some items, including most compounds, are not available on Signbank. The Signbank gloss and ID are marked as NA for those items.

4.2. Recording setup

The recordings were made at the SignLab of the University of Amsterdam using a five-camera setup that extends the signCollect infrastructure described in Otterspeer et al. (2024). All cameras are Sony FX30 units with 35mm lenses, recording at 4K (3840×2160) and 60fps. The middle camera (M) faces the signer directly; cameras L and R are placed at 45° to the left and right; cameras A and B are positioned at approximately 80° to capture the depth of the hand that the other angles miss (see Figure 1). To sync frames across cameras, a screen in the recording space displays QR codes encoding the gloss ID, date, and timestamp. These codes are decoded in post-processing to align the five video streams. The embedded metadata also ensures the recordings remain identifiable if they outlive the signCollect system.

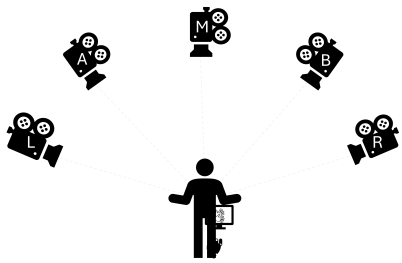

Figure 1: Schematic illustration of the recording setup.

For each individual sign, two signers produced the sign oriented straight towards the middle camera, while the other two signers were oriented at an angle, presenting the hands from the side. Which of the signers presented at which angle was rotated across signs. For the items in the game, that meant that each sign is available from both perspectives to somewhat mitigate the disadvantages of a 2-dimensional presentation of 3-dimensional materials (Watkins et al., 2024). In the present dataset, videos are available from all five cameras for both presentation angles.

Individual signs, timing information on the recording, and the review-screen for each sign were presented to the signer on a screen in front of them. The recording pipeline was set up for a default duration of four seconds, though signers were able to add time in one-second increments, if necessary. This means that videos may differ in length, particularly if they contain compounds. Navigating through the items, starting the recording, and adding additional time, was controlled by the signer through tapping a set of buttons on the foot. The first author of this paper served as a recording assistant, checking the quality of the recordings and alignment with the target variants along with the signer. In this way, we were able to efficiently record large numbers of items in a relatively short time frame.

4.3. Video processing pipeline

Raw footage is backed up in the cloud (UvA Research Drive), then videos of camera Left, Middle, and Right are automatically processed in DaVinci Resolve Headless mode to key out the green screen and replace it with a blue background. After that, they are cropped around the signer based on MediaPipe keypoints (Lugaresi et al., 2019), using the maximum extent of the detected hand and head landmarks while centered at waist height. The processed files are uploaded to the signCollect Hub for viewing and annotation. Then, derivative data are computed: MediaPipe pose estimates (`.pose`), HaMeR 3D hand meshes (Pavlakos et al., 2024a), and SAM 3D Body full-body reconstructions (Yang et al., 2025). Finally, all materials are archived on Figshare.

5. Computed derivatives

Our dataset provides both video files and computational derivatives from those videos for academic reuse (see Section 6 for information on licensing). Filenames are kept consistent between videos and computational derivatives for ease of alignment.

5.1. Keypoint estimates

For each video, we extract 2D and 3D pose estimates using MediaPipe Holistic (Lugaresi et al., 2019), which estimates body and hand keypoints. The output is stored in JSON format with the extension `.pose`, containing per-frame landmark coordinates for all detected keypoints.

Keypoint estimates can be used for pose estimation, to reconstruct three-dimensional information about the signs. These are particularly useful for extracting kinematic information about movement characteristics, such as the size of the signing space and velocity of movements throughout the sign. Similar representations have been used in sign language technology to drive sign detection and recognition and to analyse phonological information, such as hand location and movement (see, for example, Moryossef et al., 2020; Bragg et al., 2019).

5.2. Handshape estimates

3D hand mesh keypoints are extracted using HaMeR (Pavlakos et al., 2024a), a transformer-based model for hand mesh recovery from monocular video. For each frame, the model outputs 3D keypoint coordinates for both the left and right hand. Results are stored in JSON format with the extension `.hamer`, with keypoints ordered by frame.

HaMeR's 3D hand mesh keypoints create a reconstruction of the handshape through keypoints on the hand's surface. In this way, they allow reconstructing handshape estimations in occlusion situations (Pavlakos et al., 2024b). For research purposes, they can be used as a basis for categorising and comparing handshapes across signs and signers and may support the identification of the corresponding handshapes in fluent signing (Asasi et al., 2025; He et al., 2025).

5.3. 3D body reconstructions

Full-body 3D reconstructions are generated using SAM 3D Body (Yang et al., 2025), which estimates body, hand, and foot pose from a single image. The output is stored in NPZ format with the extension `.s3d` and can be converted to SMPL-X mesh representations for downstream use.

These 3D reconstructions can serve as a full-body representation of the signs. Providing a detailed representation of not only the hands but the full body, they allow the detailed tracking of bodily motions and poses over time. They can also serve as a basis for reconstruction of a sign in virtual avatars (Forte et al., 2023; Baltatzis et al., 2024).

6. The SignBeach dataset

6.1. Dataset structure

The SignBeach dataset contains recordings of 1401 lexical signs, each produced by four deaf signers, resulting in 5604 items. For each sign, video recordings are available from five camera angles (L, A, M, B, R) and in both orientations (front and side), resulting in 20 video files per lexical item (4 signers (2 front, 2 side) $\times$ 5 cameras).

The dataset is organized around Dutch target words (`wordNL`), which serve as the primary item identifiers. Each item is linked to the NGT Signbank through a gloss identifier (`glossID`) and a Signbank ID, from which the full Signbank URL can be reconstructed.[3] Table 1 summarizes the variables in the dataset.

Video files and their derivatives follow a consistent naming convention: `wordNL_signerID_orientation_cameraID_format.suffix` For example, `hond_theo_front_M_video.mp4` refers to the video recording (video) of the sign for 'hond' (dog), recorded by signer Theo, signing frontally (front) at the middle camera (M). This naming scheme allows immediate identification of the content without consulting external metadata, and enables straightforward parsing into dataset fields.

[3] Signs that are not available on Signbank have these fields marked as NA.

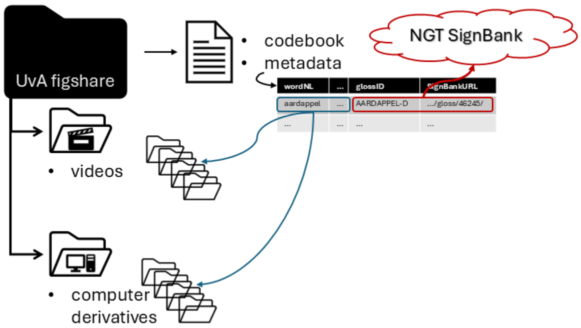

Figure 2: Illustration of the project architecture and links to the NGT Signbank as an external resource.

The derivative files follow the same base name with different extensions: `.pose` for MediaPipe estimates, `.hamer` for hand mesh data, and `.s3d` for full-body reconstructions.

6.2. Availability and license

The dataset is archived on Figshare at https://doi.org/10.21942/uva.c.8233876, organized into 50 sub-datasets grouped alphabetically by the first letter of the Dutch target word. This partitioning accommodates Figshare's limit of 500 files per dataset. Letters with many entries (e.g., 'S' with five parts, 'K' and 'V' with four parts each) are split across multiple datasets (see Figure 2). A Figshare Collection groups all datasets under a single citable DOI, with a total size of approximately 27 GB. All materials are released under a CC BY-NC 4.0 license for non-commercial reuse.

7. Conclusion and perspectives

This article presents a dataset of high quality videos for approx. 1400 NGT signs with derivational data suitable for scientific reuse, for example in the context of psycholinguistic studies or machine learning developments. High quality video recordings are available from five angles, providing 3D-information in the raw video data. In addition, we provide mediapipe, hammer, and SAM 3D reconstruction outputs for computational processing and analyses.

7.1. Links to other projects

The SignBeach dataset presented in this paper is the product of a larger project on sign learning in primary school-aged children. All resources and publications from this project will be be linked to the original video materials and derivational data provided in this dataset, permitting full academic reuse of these materials for future research.

In the SignBeach project, we are interested in investigating factors that contribute to sign learning. One factor that has been shown to impact sign

Variable	Description	Possible values
wordNL	Dutch target word	*Dutch words or expressions*
wordEN	English translation	*English words or expressions*
signerID	Signer identifier	[theo, tobias, manouk, eva]
orientation	Orientation of the signer relative to the central camera	[front, side]
cameraID	Camera identifier	[L, A, M, B, R] (left to right from signer's perspective)
glossID	ID gloss from the NGT Signbank	*e.g.*, HOND, SCHOOL; compounds: GLOSS+GLOSS
SignBankURL	URL to the sign entry on the NGT Signbank	signbank.cls.ru.nl/dictionary/gloss/{id}
video	Video file	filenames*_video.mp4
mediapipe	MediaPipe Holistic keypoint estimates	filenames*_mediapipe.pose
hamer	HaMeR hand mesh keypoint estimates	filenames*_hamer.hamer
s3d	SAM 3D Body full-body reconstruction	filenames*_s3d.s3d
*filenames	Naming system	*wordNL_signerID_orientation_cameraID*

Table 1: Variables in the SignBeach dataset. File extensions indicate format: `.pose` for MediaPipe JSON, `.hamer` for HaMeR JSON, and `.s3d` for SAM 3D Body NPZ.

learning is iconicity (for a range of different learning contexts and effects, see for example Ortega, 2017; Caselli and Pyers, 2017, 2020; Karadöller et al., 2024). In another study within this project, we therefore collected iconicity ratings from Dutch and German non-signers and Dutch signers and computed transparency estimates from translation attempts by Dutch non-signers for all items in the dataset (Schiefner et al., 2026). Since the obtained iconicity and transparency data is aligned with the items provided in the present dataset by corresponding filenames, the two resources can easily be combined to answer additional research questions.

Data about childrens' learning trajectories, in turn, will also be made available, allowing us to link the visual information derived from video materials with children's learning trajectories over time or at specific time points, such as at first exposure. This creates an integrated set of resources in which lexical, visual, and behavioural data is brought together. This will facilitate work on NGT lexical processing and acquisition, as well as comparative studies that situate NGT within a broader cross-linguistic context, e.g., by comparing the NGT dataset with data from ASL-LEX (Caselli et al., 2017; Sehyr et al., 2021) and similar resources.

7.2. Limitations

While the SignBeach dataset is designed to maximise opportunities for scientific reuse, a number of limitations need to be highlighted that will restrict interpretation and opportunities for analysis. First, the lexical coverage of this dataset is constrained by the goals of the SignBeach project. Items were selected to be child-friendly and broadly known. As a result, the database focuses on everyday vocabulary and more specialised as well as later-learned vocabulary is likely to be underrepresented. Within the concepts represented, variant choice was driven by existing dictionaries and signers' input during the recording process, likely favouring widely used variants and falling short of reflecting the lexical variation present in NGT.

Second, signs are recorded in isolation in a highly controlled studio environment using the signCollect pipeline (Otterspeer et al., 2024). This setting ensures high recording quality but does not represent naturalistic, conversational signing styles. However, the instructions to envision a child audience for the signs should result in relatively lively renditions, that are likely to be more natural than what is often used in lexical references.

Finally, the dataset also only includes productions from a limited and fairly homogenous set of four signers, an issue highlighted for a number of similar databases by Bragg et al. (2019). This design supports the analysis of articulatory variation across signers but does not reflect regional, generational, or sociolinguistic diversity in the NGT signing community, nor does our sample of signers reflect the demographic diversity of the deaf community in the Netherlands.

7.3. Perspectives for future use

We anticipate that the dataset presented in this paper provides extensive opportunities for scientific reuse. The video materials and computational derivatives provide relevant information for phonological and kinematic analysis of the lexical signs provided, with the parallel recordings from four distinct signers providing opportunities to study articulatory variation. In parallel with sensor-based motion capture data, high-quality recordings and video-derived pose and motion estimation, such as presented here, can significantly advance sign processing and analysis.

The video materials can also be reused as stimulus materials for future psycholinguistic studies, as is or in combination with the computational derivatives. For example, analysis of the computational derivatives may provide timing information for creating onset-controlled stimulus materials directly from the videos, or form the basis for creating synthetic representations of the signs by avatars that allow for targeted modification of the signs for even more controlled stimulus presentation.

Meanwhile, for computational research, the dataset provides a rich testbed for computer vision and sign language technology. This could include benchmarking pose estimation models, training and evaluating sign recognition systems, and exploring methods for realistic avatar animation grounded in naturalistic sign data.

By combining this rich, multi-angle video data with information from other studies in the SignBeach project, such as iconicity ratings, transparency scores, and information about the learnability of signs, analyses of visual information embedded in these lexical signs that contributes to these different factors will also be possible. Beyond the SignBeach project itself, the tight alignment with NGT Signbank gloss-IDs means that the SignBeach dataset easily integrates with the existing research infrastructure for NGT.

8. Acknowledgements

We thank Tobias de Ronde, Debby Musch, Machiel Ouwerkerk, Eva Prevaes, Theo Stoop and Manouk van Asperen who served as models for the video materials presented in this paper and gave valuable input in the sign selection process.

9. Bibliographical References

Sobhan Asasi, Mohamed Ilyas Lakhal, Ozge Mercanoglu Sincan, and Richard Bowden. 2025. Beyond gloss: A hand-centric framework for gloss-free sign language translation. ArXiv:2507.23575 [cs].

Vasileios Baltatzis, Rolandos Alexandros Potamias, Evangelos Ververas, Guanxiong Sun, Jiankang Deng, and Stefanos Zafeiriou. 2024. Neural sign actors: A diffusion model for 3D sign language production from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1985–1995.

Danielle Bragg, Oscar Koller, Mary Bellard, Larwan Berke, Patrick Boudreault, Annelies Braffort, Naomi Caselli, Matt Huenerfauth, Hernisa Kacorri, Tessa Verhoef, Christian Vogler, and Meredith Ringel Morris. 2019. Sign language recognition, generation, and translation: An interdisciplinary perspective. In *Proceedings of the 21st International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '19, pages 16–31, New York, NY, USA. Association for Computing Machinery.

Naomi K. Caselli and Jennie E. Pyers. 2017. The road to language learning is not entirely iconic: Iconicity, neighborhood density, and frequency facilitate acquisition of sign language. *Psychological Science*, 28(7):979–987.

Naomi K. Caselli and Jennie E. Pyers. 2020. Degree and not type of iconicity affects sign language vocabulary acquisition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 46(1):127–139.

Naomi K. Caselli, Zed Sevcikova Sehyr, Ariel M. Cohen-Goldberg, and Karen Emmorey. 2017. ASL-LEX: A lexical database of American Sign Language. *Behavior Research Methods*, 49(2):784–801.

Onno Crasborn, Richard Bank, Inge Zwiterlood, Els Van Der Kooij, Ellen Ormel, Johan Ros, Anique Schüller, Anne De Meijer, Merel Van Zuilen, Ellen Nauta, Frouke Van Winsum, and Max Vonk. 2020. NGT Signbank. Nijmegen, NL. Radboud University, Centre for Language Studies. Version: 3.0.

Onno Crasborn, Richard Bank, Inge Zwitserlood, Els van der Kooij, Anique Schüller, Ellen Ormel, Ellen Yassine Nauta, Merel van Zuilen, Frouke van Winsum, and Johan Ros. 2016. Linking Lexical and Corpus Data for Sign Languages: NGT Signbank and the Corpus NGT. In *Proceedings of the LREC2016 7th Workshop on the Representation and Processing of Sign Languages: Corpus Mining*, pages 41–46, Portorož, Slovenia. European Language Resources Association (ELRA).

Onno Crasborn and Inge Zwitserlood. 2008. The Corpus NGT: an online corpus for professionals and laymen. In *Construction and Exploitation of Sign Language Corpora*, pages 44–49, Paris. ELDA.

Onno Crasborn, Inge Zwitserlood, Els Van Der Kooij, and Anique Schüller. 2018. Global SignBank manual.

J. Fenlon, K. Cormier, and A. Schembri. 2015. Building BSL SignBank: The lemma dilemma revisited. *International Journal of Lexicography*, 28(2):169–206.

Maria-Paola Forte, Peter Kulits, Chun-Hao Paul Huang, Vasileios Choutas, Dimitrios Tzionas, Katherine J. Kuchenbecker, and Michael J. Black. 2023. Reconstructing signing avatars from video using linguistic priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12791–12801.

Eva Gutierrez-Sigut, Brendan Costello, Cristina Baus, and Manuel Carreiras. 2016. LSE-Sign: A lexical database for Spanish Sign Language. *Behavior Research Methods*, 48(1):123–137.

Low Jian He, Harry Walsh, Ozge Mercanoglu Sincan, and Richard Bowden. 2025. Hands-on: Segmenting individual signs from continuous sequences. In *2025 IEEE 19th International Conference on Automatic Face and Gesture Recognition (FG)*, pages 1–5, Tampa/Clearwater, FL, USA. IEEE.

Elena Jahn, Reiner Konrad, Gabriele Langer, Sven Wagner, and Thomas Hanke. 2018. Publishing DGS corpus data: Different formats for different needs. In *Workshop Proceedings. 8th Workshop on the Representation and Processing of Sign Languages: Involving the Language Community*, volume 12 May 2018. ELRA, pages 83–90, Miyazaki, Japan.

Dilay Z. Karadöller, David Peeters, Francie Manhardt, Aslı Özyürek, and Gerardo Ortega. 2024. Iconicity and gesture jointly facilitate learning of second language signs at first exposure in hearing nonsigners. *Language Learning*, 74(4):781–813.

Ulrika Klomp, Lisa Gierman, Pieter Manders, Ellen Nauta, Gomèr Otterspeer, Ray Pelupessy, Galya Stern, Dalene Venter, Casper Wubbolts, Marloes Oomen, and Floris Roelofsen. 2024. An extension of the NGT dataset in Global Signbank. In *Proceedings of the LREC-COLING 2024*, pages 178–183.

Reiner Konrad, Thomas Hanke, Gabriele Langer, Dolly Blanck, Julian Bleicken, Ilona Hofmann, Olga Jeziorski, Lutz König, Susanne König, Rie Nishio, Anja Regen, Uta Salden, Sven Wagner, Satu Worseck, Oliver Böse, Elena Jahn, and Marc Schulder. 2020. *MY DGS – annotated. Public corpus of German Sign Language*, 3rd edition. Universität Hamburg.

Maria Kopf, Marc Schulder, and Thomas Hanke. 2022. The Sign Language Dataset Compendium: Creating an Overview of Digital Linguistic Resources. In *Proceedings of the LREC2022 10th Workshop on the Representation and Processing of Sign Languages: Multilingual Sign Language Resources*, pages 102–109, Marseille, France. European Language Resources Association (ELRA).

Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Yong, Juhyun Lee, and others. 2019. Mediapipe: A framework for perceiving and processing reality. In *Third workshop on computer vision for AR/VR at IEEE computer vision and pattern recognition (CVPR)*.

Max Planck Institute for Psycholinguistics, The Language Archive. 2025. ELAN (Version 7.0).

Hope Morgan, Wendy Sandler, Rose Stamp, and Rama Novogrodsky. 2022. ISL-LEX v.1: An online lexical resource of Israeli Sign Language. In *Proceedings of the 10th Workshop on the Representation and Processing of Sign Languages*, pages 148–153, Marseille, FR.

Amit Moryossef, Ioannis Tsochantaridis, Roee Aharoni, Sarah Ebling, and Srini Narayanan. 2020. Real-time sign language detection using human pose estimation. In *Computer Vision – ECCV 2020 Workshops*, pages 237–248, Cham. Springer International Publishing.

Nederlands Gebarencentrum. 2012. *Mijn eerste 1500 gebaren*, 1 edition. Nederlands Gebarencentrum, Bunnik, NL.

Nederlands Gebarencentrum. 2024. *Online Woordenboek NGT*. Nederlands Gebarencentrum.

Gerardo Ortega. 2017. Iconicity and sign lexical acquisition: A review. *Frontiers in Psychology*, 8:1280.

Felicitas Otte, Elena Jahn, Cornelia Loos, Julian Bleicken, and Annika Herrmann. 2022. The DGS corpus as a linguistic resource: student research and beyond. *Linguistische Berichte*, (272):447.

Gomer Otterspeer, Ulrika Klomp, and Floris Roelofsen. 2024. SignCollect: A 'touchless' pipeline for constructing large-scale sign language repositories. In *Proceedings of the LREC2024 11th Workshop on the representation and processing of sign languages: Evaluation of sign language resources*, pages 269–275, Turin. European Language Resources Association.

Georgios Pavlakos, Dandan Shan, Ilija Radosavovic, Angjoo Kanazawa, David Fouhey, and Jitendra Malik. 2024a. HaMeR: Hand mesh recovery for the EgoExo4D hand pose challenge.

Georgios Pavlakos, Dandan Shan, Ilija Radosavovic, Angjoo Kanazawa, David Fouhey, and Jitendra Malik. 2024b. Reconstructing hands in 3D with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9826–9836.

Adam Schembri, Jordan Fenlon, Ramas Rentelis, and Kearsy Cormier. 2017. *British Sign Language Corpus Project: A corpus of digital video data and annotations of British Sign Language 2008-2017*, 3 edition. University College London, London.

Annika Schiefner, Floris Roelofsen, and Beyza Sümer. 2026. Iconicity and transparency in Dutch Sign Language. *Behavior Research Methods*, 58(87).

Marc Schulder, Dolly Blanck, Thomas Hanke, Ilona Hofmann, Sung-Eun Hong, Olga Jeziorski, Lutz König, Susanne König, Reiner Konrad, Gabriele Langer, Rie Nishio, and Christian Rathmann. 2024. Data statement for the public DGS corpus. *DGS-Korpus Project Note AP06-2020-01*. Version: 3.

Zed Sevcikova Sehyr, Naomi Caselli, Ariel M Cohen-Goldberg, and Karen Emmorey. 2021. The ASL-LEX 2.0 project: A database of lexical and phonological properties for 2,723 signs in American Sign Language. *The Journal of Deaf Studies and Deaf Education*, 26(2):263–277.

Patrick C. Trettenbrein, Jan Bümmerstede, Thomas A. Finkbeiner, Pia Gehlbach, Nina-Kristin Meister, Annika Schiefner, Petra B. Schumacher, Door Spruijt, Markus Steinbach, and Pamela Perniss. 2025. Towards a "DGS-LEX": A roadmap for the collaborative creation of a psycholinguistic database for German Sign Language (DGS). *Linguistische Berichte*, 281:91–100.

Freya Watkins, Diar Abdlkarim, Bodo Winter, and Robin L. Thompson. 2024. Viewing angle matters in British Sign Language processing. *Scientific Reports*, 14(1):1043.

Xitong Yang, Devansh Kukreja, Don Pinkus, Anushka Sagar, Taosha Fan, Jinhyung Park, Soyong Shin, Jinkun Cao, Jiawei Liu, Nicolas Ugrinovic, Matt Feiszli, Jitendra Malik, Piotr Dollar, and Kris Kitani. 2025. SAM 3D Body: Robust full-body human mesh recovery. *arXiv preprint*.